IEEE COMMUNICATIONS MAGAZINE

September 2017, Vol. 55, No. 9

- Advances in Network Services Chain
- Network Testing and Analytics
- Radio Communications
- Internet of Things

IEEE ComSoc IEEE Communications Society

A Publication of the IEEE Communications Society

www.comsoc.org

IEEE COMMUNICATIONS MAGAZINE

September 2017, vol. 55, no. 9

www.comsoc.org/commag

IEEE COMMUNICATIONS MAGAZINE (ISSN 0163-6804) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. Headquarters address: IEEE, 3 Park Avenue, 17th Floor, New York, NY 10016-5997, USA; tel: +1 (212) 705-8900; http://www.comsoc.org/commag. Responsibility for the contents rests upon authors of signed articles and not the IEEE or its members. Unless otherwise specified, the IEEE neither endorses nor sanctions any positions or actions espoused in *IEEE Communications Magazine.*

ANNUAL SUBSCRIPTION: $71: print, digital, and electronic. $33: digital and electronic. $1001: non-member print.

EDITORIAL CORRESPONDENCE: Address to: Editor-in-Chief, Osman S. Gebizlioglu, Huawei Technologies, 400 Crossing Blvd., 2nd Floor, Bridgewater, NJ 08807, USA; tel: +1 (908) 541-3591, e-mail: Osman.Gebizlioglu@huawei.com.

POSTMASTER: Send address changes to *IEEE Communications Magazine*, IEEE, 445 Hoes Lane, Piscataway, NJ 08855-1331. GST Registration No. 125634188. Printed in USA. Periodicals postage paid at New York, NY and at additional mailing offices. Canadian Post International Publications Mail (Canadian Distribution) Sales Agreement No. 40030962. Return undeliverable Canadian addresses to: Frontier, PO Box 1051, 1031 Helena Street, Fort Eire, ON L2A 6C7.

SUBSCRIPTIONS: Orders, address changes — IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08855-1331, USA; tel: +1 (732) 981-0060; e-mail: address.change@ieee.org.

ADVERTISING: Advertising is accepted at the discretion of the publisher. Address correspondence to: Advertising Manager, *IEEE Communications Magazine*, IEEE, 445 Hoes Lane, Piscataway, NJ 08855-1331.

SUBMISSIONS: The magazine welcomes tutorial or survey articles that span the breadth of communications. Submissions will normally be approximately 4500 words, with few mathematical formulas, accompanied by up to six figures and/or tables, with up to 10 carefully selected references. Electronic submissions are preferred, and should be submitted through Manuscript Central: http://mc.manuscriptcentral.com/commag-ieee. Submission instructions can be found at the following: http://www.comsoc.org/commag/paper-submission-guidelines. All submissions will be peer reviewed. For further information contact Tarek El-Bawab, Associate Editor-in-Chief (telbawab@ieee.org).

BPA WORLDWIDE

Network Testing and Analytics

Series Editors: Ying-Dar Lin, Irena Atov, and Erica Johnson

Radio Communications: Components, Systems, and Networks

Series Editors: Amitabh Mishra and Tom Alexander

Accepted from Open Call

IEEE COMSOC live WEBINAR

IEEE

IEEE ComSoc
IEEE Communications Society

mmWave 5G Phased Array and Beamforming System Design

5G

Thursday, 28 September 2017
1:00 PM ET \ 10:00 AM PT

Live webinar on mmWave 5G Phased Array and Beamforming System Design

- Introduces advanced modeling and simulation approaches for various beamforming techniques.
- Demonstrates phased array system design while using practical phased arrays in working, system-level beamforming scenarios.
- Features a live software demo with typical configurations, weightings and impairments that affect beam-level performance.
- Shows how to apply the arrays in communications link applications.

Featured Speaker

Sangkyo Shin
5G Product Planning and Application Developer
Keysight Technologies

Register at http://bit.ly/Beamforming

Sponsored Content by KEYSIGHT TECHNOLOGIES

IEEE GREEN ICT

Harvey Freeman

The IEEE sets up IEEE-wide initiatives via its Future Directions Committee (FDC) to ensure that the organization is prepared to address future needs, is ahead of and capitalizes on technology trends to maximize its impact and fulfil its mission of advancing technology for the benefit of humanity. Recent IEEE initiatives have focused on cloud computing, smart cities, transportation electrification and Green Information and Communications Technology (Green ICT).

The IEEE Green ICT initiative is currently in its third year, the final year funded by the IEEE New Initiatives Committee (NIC). NIC has provided approximately $1m over the three years to support the activities of the initiative. Upon graduating from the NIC, the IEEE Green ICT initiative will continue to be led by ComSoc with the active involvement of the other IEEE Societies.

Led by Co-Chairs Jaafar Elmirghani and Charles Despins, the IEEE Green ICT initiative has achieved outstanding results against its objectives and desired outcomes over its first 2.5 years.

Jaafar M. H. Elmirghani

Charles Despins

Jaafar M. H. Elmirghani is a Full Professor and Chair in Communication Networks and Systems, and the Director of the Institute of Communication and Power Networks within the School of Electronic and Electrical Engineering, University of Leeds, UK. He has co-authored *Photonic Switching Technology: Systems and Networks*, (Wiley) and has published over 450 papers. He has research interests in communication systems and networks. Prof. Elmirghani is Fellow of the IET, and a Chartered Engineer and Fellow of the Institute of Physics. He was Chairman of the IEEE ComSoc Transmission Access and Optical Systems (TAOS) technical committee and was Chairman of IEEE ComSoc Signal Processing and Communications Electronics (SPCE) technical committee, and an editor of *IEEE Communications Magazine*. He was awarded the IEEE ComSoc 2005 Hal Sobol Award, two IEEE ComSoc outstanding service awards (SPCE 2009 and TAOS 2015), the 2015 GreenTouch 1000x award, the IET 2016 Premium Award for work on Green Communications, and shared the 2016 Edison Award in the collective disruption category with a team of six from GreenTouch for their joint work on the GreenMeter.

Charles Despins' career has spanned more than 30 years in both the academic and industry segments of the information and communications technologies (ICT) sector. In addition to his academic work, he has held various posts in the private sector, namely at Bell Nordiq Group as vice-president and chief technology officer, as a consultant for wireless network deployments in India and China, and for 13 years as President and CEO of Prompt Inc., an ICT research and development consortium. He is currently Dean of Research at École de Technologie Supérieure (Université du Québec) in Montreal, Canada. Dr. Despins is a Fellow of the Engineering Institute of Canada and a recipient (2006) of the Outstanding Engineer award from IEEE Canada. He is a former recipient of the "Best Paper of the Year" award in *IEEE Transactions on Vehicular Technology*. He is currently a frequent advocate on issues regarding the opportunities ICT offers to achieve sustainability in the 21st century.

THE CHALLENGE

ICT has a carbon footprint comparable to the global aviation industry, each responsible for about 2 percent of the world's carbon emissions. A study by the Global e-Sustainability initiative (GeSi) shows that of 25 percent of ICT emissions are attributed to telecoms, 18 percent to data centers, and 57 percent to peripheral devices. Telecom companies are now the single largest users of energy in their countries, well ahead of the manufacturing industry. For example, Telecom Italia uses about 1 percent of the energy in Italy.

To understand the driving force behind the growth in ICT power consumption, we have to understand the traffic trends. Internet traffic used to grow at about 300 percent per year just before the turn of the millennium; it was expected to grow at even higher rates, which did not happen as a result of the dot-com bubble burst in 2000. Internet traffic, nonetheless, is now still growing at 30 percent to 40 percent per year! The aviation industry's growth is almost flat in comparison.

At 40 percent growth per year, traffic doubles every two years, multiplies by 30x in 10 years, and multiplies by 1000x in 20 years. So if we were to improve the energy efficiency of ICT hardware, software and networks by a factor of 1000, then in 20 years, we will consume the same amount of power as today, but accommodate 1000x more traffic. Hence the GreenTouch (one of the IEEE Green ICT Initiative partners) vision to improve energy efficiency by a factor of 1000. This is essential as traffic grows faster at present than improvement in technology (energy efficiency); as such, business as usual is not sufficient and new initiatives are needed to tackle the challenge.

To see which strands to tackle in telecom, we need to understand the traffic trends. Voice is the slowest growing, wireless data is the fastest growing traffic strand and Internet video is the largest traffic strand by volume followed by peer to peer (P2P) traffic. Therefore, to Green the Internet, we have to green video, P2P and wireless data mainly.

In addition to greening ICT in the manner outlined, a greater opportunity exists in the use of ICT to Green other sectors, including transport (by reducing journeys, for example), manufacturing and agriculture (both by introducing more efficient processes through sensing, processing and actuation, for example), and a host of other verticals such as healthcare and energy distribution via smart grid. The use of ICT to green other sectors (greening by ICT) has the potential to reduce the global carbon footprint by an amount equal to 10 times ICT's own carbon footprint, as shown by the SMARTer 2030 report, i.e. a 20 percent reduction in the global carbon footprint can be achieved through the use of ICT to green verticals.

THE INITIATIVE

The IEEE Green ICT initiative, through its many IEEE member Societies, addresses two main challenges: Greening ICT and Greening by ICT. The "sustainability through technology" challenge will remain relevant for a long time, is essential for the future of the planet, and is not simply a temporary "hot" technology trend. The initiative is currently hosted by the IEEE Communications Society, which is the lead society, with participation by many IEEE societies including the Computer Society, Vehicular Technology Society, Power and Energy Society, Industrial Electronics Society, Microwave Theory and Techniques Society, and the Society on Social Implications of Technology.

The challenge spans multiple disciplines where greening one part of ICT may lead to adverse effects elsewhere. For example, greening telecom by shifting processing to data centers may have adverse effects on data centers. Therefore, seamless collaboration is needed across IEEE societies.

The Green ICT Initiative's mission, therefore, is to "build a holistic approach to sustainability by incorporating green metrics throughout IEEE technical domains." Its specific objectives are to: (a) foster inter-society collaboration to build such a holistic approach to sustainability; (b) diversify IEEE membership through outreach to non-traditional IEEE communities; and (c) grow IEEE's influence and visibility in international forums. Viewed through the triple bottom line of sustainability (economic, environmental, social), the IEEE Green ICT Initiative, therefore, offers a compelling opportunity for IEEE to demonstrate the full impact of the technology innovation it supports.

The IEEE Green ICT Initiative desired outcomes from the outset included: (i) establishing and enhancing conferences and workshops to share knowledge about greening of/by ICT; (ii) establishing new publications to promote multidisciplinary approaches to incorporating Green ICT metrics in IEEE fields; (iii) developing IEEE standards to formalize IEEE expertise and leadership in tangible ways; (iv) developing training and awareness material to improve IEEE member knowledge and engagement; (v) engaging a broader Green ICT community through rich portal content.

CONFERENCES AND WORKSHOPS

In the IEEE Communications Society, the Transmission Access and Optical Systems (TAOS) technical committee started the first ICC/GLOBECOM track on green communications and networking through the efforts of Prof. Elmirghani, starting in 2009, which led to the first green track in ICC/GLOBECOM at GLOBECOM 2011. This effort has grown since with the Green Communications and Networking track elevated to a full symposium at ICC 2016 through the efforts of TAOS and the IEEE Green ICT initiative.

Over the past thee years the initiative has co-organized many tracks, sessions, panels and tutorials at every ICC and GLOBECOM. It has also organized similar activities at the leading conferences of other IEEE societies to foster inter-society collaboration including at ISTAS 2015-2017, VTC 2016 and VTC 2017, Intelect 2015-2017, GreenComm 2015-2017, ICUWB 2015, ICUNC 2015, ICTON2015-2017, Future Internet and Smart Cities, 2015, EuCNC 2015, GreenTouch showcase 2015, CCNC 2015-2017 and OFC 2015-2017.

The IEEE Green ICT initiative has started a new Summit on Green ICT for which a first roadmapping event took place at ICC 2017. The inaugural Summit on Greening through Information & Communications Technology, "The GtICT Summit," will take place in Paris on October 3, 2017. Its aims are to identify key technological, commercial, and public policy challenges and solutions to achieve sustainability through ICT. Prof. Charles Despins delivered a number of webinars and talks in English and French to promote the summit. The summit will include a World Café approach focused on stimulating dialogue across disciplines, a plenary on how to ensure that smart is also sustainable, together with views from government, from ICT verticals, and an IEEE Young Professional Green ICT Idea Competition on greening ICT and greening by ICT (http://greenict.ieee.org/summit/gtict-summit-2017).

PUBLICATIONS

Building on the successful ICC/GLOBECOM Green Communications and Networking track and then symposium, the initiative, with colleagues in ComSoc, was instrumental in introducing an *IEEE Journal on Selected Area in Communications* Series on Green Communications and Networking. The first issue, in December 2015, received 143 submissions; 39 papers were accepted with 552 total pages (by far the largest JSAC special issue in recent history; a typical issue includes 15 to 17 papers); 22 papers were transferred to the second issue. Issue 2, in June 2016, received 125 submissions; 23 papers were accepted for Issue 2 (in addition to 22 transferred from submissions to Issue 1), and as such, about 50 papers were published in Issue 2 with about 720-760 total pages. Issue 3, in December 2016, received 130 submissions. The effort led to the approval of the new *IEEE Transactions on Green Communications and Networking*. The first issue of this new journal was published in January 2017.

At ICC 2017, the IEEE Green ICT initiative received approval from the IEEE Communications Society Publications Council and Board of Governors for a phase I submission of a new Sustainable ICT Magazine where ComSoc is the lead financial and technical sponsor.

The *IEEE Institute* dedicated its March 2016 issue to the IEEE Green ICT initiative. *IEEE Spectrum* has published feature articles on the initiative, for example in February 2017. Members of the initiative contributed to the GreenTouch Final Results: the Green Meter Research Study White Paper 2015, together with numerous articles in industrial magazines and venues (http://greenict.ieee.org/blog).

STANDARDS

To maximize industrial engagement and impact, the initiative set up a very successful standardization program. It organized two rapid reaction workshops to mobilize the community, first in November 2015 in London and then in July 2016 in Chicago, with participation by several key players including Nokia, Ericsson, GSMA, Huawei and Qualcomm. The nine project authorization requests were approved by the IEEE Communications Society Standards Development Board in October 2016 and by the IEEE Standards Association in December 2016. Work on the standards, including final sponsor ballots, is due to conclude by September 2018. The nine standards are organized under three working groups:

GICT – GICT Emissions Working Group

IEEE P1922.1 is a standard for a method for calculating anticipated emissions caused by virtual machine migration and placement.

IEEE P1922.2 is a standard for a method to calculate near real-time emissions of information and communication technology infrastructure.

EECH – Energy Efficient Comm Hardware Working Group

IEEE P1923.1 is a standard for computation of energy efficiency upper bound for apparatus processing communication signal waveforms.

IEEE P1924.1 is a recommended practice for developing energy efficient power-proportional digital architectures.

EEICT – COM/SDB/Energy Efficient ICT Working Group

IEEE P1925.1 is a standard for Energy Efficient Dynamic Line Rate Transmission System.

IEEE P1926.1 is a standard for a Functional Architecture of Distributed Energy Efficient Big Data Processing.

IEEE P1927.1 is a standard for Services Provided by the Energy-efficient Orchestration and Management of Virtualized Distributed Data Centers Interconnected by a Virtualized Network.

IEEE P1928.1 is a standard for a Mechanism for Energy Efficient Virtual Machine Placement.

IEEE P1929.1 is an Architectural Framework for Energy Efficient Content Distribution which specifies a framework for designing energy efficient content distribution services, such as migration, placement, and replication, over networks.

TRAINING

The initiative organized a number of short courses at ICC and GLOBECOM delivered by Prof. Elmirghani which include: (i) "Greening the Internet," ICC'13: (ii) "Energy Efficient Core and Content Distribution Networks," SoftCOM'2013; (iii) "Greening Core, Data Center and Content Distribution Networks," ICC 2014; (iv) "Greening Cloud Networks," ICC'15; (v) "Greening Big Data Networks," ICC'16; (vi) "Greening Cloud and Virtualised Communication Networks," GLOBECOM'17. Prof. Fabrizio Granelli delivered the online training course "Designing the Green Internet" offered by ComSoc. The initiative co-organized several distinguished lecturer tours focusing on Green ICT. As part of the standardization effort, the initiative organized and delivered four webinars, one by each of the standards working groups and a broader webinar to promote and provide training in the areas of the standards.

We encourage the engagement of a broader Green ICT community through rich portal content and the initiative technical community. The IEEE Green ICT community already has over 4000 members. This is a call to participate in all the IEEE Green ICT initiative activities. Please see http://greenict.ieee.org, and contact the initiative Co-Chairs. Sustainability concerns all of IEEE!

UNDERSTANDING TELECOMMUNICATIONS BUSINESS

By Andy Valdar and Ian Morfett, The Institution of Engineering and Technology, 2016, ISBN 978-1-84919-745-8, softcover, 412 pages

Reviewer: Piotr Cholda

When an engineer thinks about the IT industry, a tendency to consider only the technical side of it is likely to appear. That is quite natural as successes in this area are obviously related to the innovation and exciting inventions that are presented throughout the pages of our magazine. However, as in the case of each industry, there are also other very important aspects, i.e., economic, legal, etc. Therefore, to be effective in practical IT, a network specialist needs not only to be expert in one's own field; they also need to be aware of the general business context, since in many cases this is the source of inspiration or, on the other hand, problems influencing the every-day operation. This book authored by two distinguished practitioners will help with getting the required knowledge.

The content is divided into 10 chapters, each of which is some 30-40 pages long. Hence, it is possible to get acquainted with a single piece of the presented material during an afternoon. While almost all of the presented notions are described in the context of IT issues (with some, more emphasis is given to telecommunications-related problems, although the general IT context is also present), some of the chapters cover very general economic problems related to business operation: organization around value chain analysis (chapter 1), business strategy concepts (chapter 3), corporate finance and governance (chapter 4), dealing with customers and marketing (chapter 7), product management (chapter 8), and people and organizational development (chapter 10). These general topics are intertwined with the chapters detailing specific aspects of the field interesting for us: regulation in telecommunications along with the recent emphasis on opening markets (chapter 2), economics of network design (chapter 5), planning of various types of networks (chapter 6), and management of networks and services (chapter 9). These chapters were most interesting for me as well as for my students with whom I discussed the book during a course on techno-economic problems of contemporary networking. An especially attractive aspect of these chapters was related to the fact that except for robustly covering the most important topics, the authors present many real-life illustrations of the presented ideas, thus facilitating the process of memorization due to coupling impacting stories with general concepts. Although throughout the book the British perspective of the authors is seen, there are many examples covering European, American and (rarely) Asian markets. Anyway, globalization, so important for today's big IT businesses, is tangible for the reader.

The book can be recommended to students as a handbook, being a perfect supplement for the lectures on entrepreneurship or practical project management courses that are typically a part of the curriculum for a future IT engineers. This is also a good summary of various aspects of networking from the organizational and global viewpoint, a perspective that might sometimes be lost when a student encounters various topics covered during different courses separately. On the other hand, the book will be useful for network specialists who would like to understand what are the basic categories of thinking for management or marketing staff.

UPDATED ON THE COMMUNICATIONS SOCIETY'S WEB SITE

www.comsoc.org/conferences

2017

SEPTEMBER

ITC29 2017 — International Teletraffic Congress, 4–8 Sept.
Genoa, Italy
https://itc29.org/

IEEE CSCN 2017 — IEEE Conference on Standards for Communications & Networking, 5–7 Sept.
Helsinki, Finland
http://cscn2017.ieee-cscn.org/

ICACCI 2017 — Int'l. Conference on Advances in Computing, Communications and Informatics, 13–16 Sept.
Udupi, India
http://icacci-conference.org/2017/

IEEE Sarnoff Symposium 2017, 18–20 Sept.
Newark, NJ
https://ewh.ieee.org/conf/sarnoff/2017/

SOFTCOM 2017 — Int'l. Conference on Software, Telecommunications and Computer Networks, 21–23 Sept.
Split, Croatia
http://softcom2017.fesb.unist.hr/

IEEE CLOUDNET 2017 — IEEE Int'l. Conference on Cloud Networking, 25–27 Sept.
Prague, Czech Republic
http://cloudnet2017.ieee-cloudnet.org/

OCTOBER

I3C 2017 — IoT Int'l, Innovation Conference, 5–7 Oct.
Saodoa. Morocco
http://i3c2017.emena.org/index.html

IEEE PIMRC 2017 — IEEE Int'l. Symposium on Personal, Indoor & Mobile Radio Communications, 8–13 Oct.
Montreal, Canada
http://pimrc2017.ieee-pimrc.org/2015/08/21/sample-news-post/

IEEE CNS 2017 — IEEE Conference on Communications and Network Security, 9–11 Oct.
Las Vegas, NV
http://cns2017.ieee-cns.org/

HONET-ICT 2017 — Int'l. Conference on Smart Cities: Improving Quality of Life Using ICT & IoT, 9–11 Oct.
Irbid, Jordan
http://honet-ict.org/

WCSP 2017 — Int'l. Conference on Wireless Communications and Signal Processing, 11–13 Oct.
Nanjing, China
http://www.ic-wcsp.org/

CyberC 2017 — Int'l. Conference on Cyber-Enabled Distributed Computing and Knowledge, 12–14 Oct.
Nanjing, China

IEEE HEALTHCOM 2017 — IEEE Int'l. Conference on e-Health Networking, Application & Services, 12–15 Oct.
Dalian, China
http://healthcom2017.ieee-healthcom.org/

IEEE SmartGridComm 2017 — IEEE International Conference on Smart Grid Communications, 16–19 Oct.
Dresden, Germany
http://sgc2017.ieee-smartgridcomm.org/

CSNet 2017 — Cyber Security in Networking Conference, 18–20 Oct.
Rio de Janeiro, Brazil
http://csnet2017.dnac.org/

ICTC 2017 — Int'l. Conference on Information and Communication Technology Convergence, 18–20 Oct.
Jeju Island, Korea
http://ictc2017.org/

ATC 2017 — Int'l. Conference on Advanced Technologies for Communications, 18–20 Oct.
Quynhon, Vietnam
http://atc-conf.org/

INTEC 2017 — Int'l. Conference on Internet of Things, Embedded Systems and Communications, 20–22 Oct.
Gafsa, Tunisia
http://www.iintec.org/

IEEE/CIC ICCC 2017 — IEEE/CIC Int'l. Conference on Communications in China, 22–24 Oct.
Qingdao, China
http://iccc2017.ieee-iccc.org/

MILCOM 2017 — Military Communications Conference, 23–25 Oct.
Baltimore, MD
http://events.afcea.org/milcom17/public/enter.aspx

Fog World Congress 2017, 30 Oct.–1 Nov.
Santa Clara, CA
http://www.fogworldcongress.com/

NOVEMBER

WINCOM 2017 — Int'l. Conference on Wireless Networks and Mobile Communications, 1–4 Nov.
Rabat, Morocco
http://www.wincom-conf.org/?p=welcome

IEEE NFV-SDN 2017 — IEEE Conference on Network Function Virtualization and Software Def ned Networks, 6–8 Nov.
Berlin, Germany
http://nfvsdn2017.ieee-nfvsdn.org/

FRUCT21 2017 — Conference of Open Innovations Association (FRUCT) 2017, 6–10 Nov.
Helsinki, Finland
http://fruct.org/conference21

IEEE LATINCOM 2017 — 9th Latin-American Conference on Communications, 8–10 Nov.
Guatemala City, Guatemala
http://latincom2017.ieee-comsoc-latincom.org/

–Communications Society portfolio events appear in bold colored print.
–Communications Society technically co-sponsored conferences appear in black italic print.

GCN GLOBAL COMMUNICATIONS NEWSLETTER

September 2017
ISSN 2374-1082

REGIONAL ACTIVITIES

Making San Diego Better Off: From Antenna Workshop to Science Fair

By Liangping Ma, InterDigital, Inc., San Diego, CA, USA (ComSoc Chapter Chair)

Like any other IEEE ComSoc chapter, we, the San Diego ComSoc chapter, are part of a much larger community, professionals in electrical engineering and computer science who may or may not be IEEE members, as well as young students who will become the next generation of electrical engineers and computer scientists. We strongly believe that the better off the larger community, the better off the IEEE. While serving IEEE members, we reach out to offer training opportunities to non-member professionals and to inspire the growth of young students (high school students and below).

In 2017, as a key part of our tradition, we continued to bring world-class researchers and professionals to our members. On June 12, in collaboration with the Vehicular Technology Society (VTS) chapter, chaired by Dr. Byung K. Yi, CTO of InterDigital, we invited Prof. Shiwen Mao of Auburn University, an IEEE VTS Distinguished Lecturer, to present his recent research on contact-free vital sign measurement for the healthcare Internet of Things. Besides theory and implementation, Prof. Mao gave a live demonstration of a smartphone app, which made the event very engaging to the audience.

As part of our year 2017 outreach effort, we worked closely with Jason Geurts of ANSYS Inc., a major computer-aided engineering software development company, and organized a hands-on interactive workshop on antenna design on April 26. Arien Sligar, lead application engineer of ANSYS, showed the ease, speed and accuracy of combining a high frequency structural simulator (HFSS) and the shooting and bouncing rays methodology (SBR+) in the new ANSYS Electronics Desktop platform to study electrically large scenes. The workshop was enjoyable and rewarding to many in the audience.

Prof. Shiwen Mao, an IEEE VTS Distinguished Lecturer, shown presenting to IEEE members.

We further reached out to young students, mainly high school students and middle school students in the greater San Diego area. We leveraged the annual Greater San Diego Science and Engineering Fair, which had 578 participating projects in 2017. The IEEE San Diego Section formed an award committee, led by Lawrence Hamerman, Member-at-Large of the IEEE San Diego Section. I was also on the committee, as the ComSoc chapter chair. The committee reviewed all projects, interviewed many participants, and selected six projects for an award of $200 per awardee. The award committee also provided valuable suggestions on how to improve the quality and impact of the projects.

The accomplishments were the result of the hard work of the IEEE San Diego ComSoc chapter and the section, with the generous support of corporate sponsors. Together, we are making the greater San Diego community better off.

Lawrence Hamerman (fourth from the right), Member-at-Large of IEEE San Diego Section, shown with the Science Fair awardees..

Arien Sligar, lead application engineer of ANSYS, shown teaching at the antenna hands-on workshop.

CONFERENCE PREVIEW

IEEE Green ICT Envisions a Smart and Sustainable Future

By Charles Despins, ÉTS Université du Québec, Montreal, Canada, Jaafar Elmirghani, University of Leeds, UK, and Thierry Klein, Nokia, NJ, USA, IEEE Green ICT Initiative Co-Chairs

Currently transforming all spheres of human activity, at a level greater than the industrial revolution, the Information and Communication Technology sector (ICT) is emerging in the 21st century as the dominant driver of sustainability, with the potential to reconcile economic growth, environmental protection and societal benefits. It is a key tool in the fight against climate change as it can enable a 20 percent reduction in global greenhouse gas emissions by 2030. However, achieving such outcomes will require a holistic approach to the proper design, broad application, widespread adoption and social acceptance of ICT products and solutions. As such, this "green" potential of ICT requires a complete rethinking of how not only we design but also how we use ICT in a sustainable fashion. It is a huge challenge that can only be satisfactorily addressed by bringing together the research community, ICT equipment and solution providers, practitioners in various ICT vertical markets, the standards community, and public policy and regulatory influencers and decision makers.

As IEEE is the world's largest technical professional organization dedicated to advancing technology for the benefit of humanity, the Green ICT initiative is committed to providing leadership on these issues. In this context, the mission of the IEEE Green ICT Initiative, launched in January 2015 by the IEEE Future Directions Committee and the IEEE Communications Society, has been defined as "build a holistic approach to sustainability by incorporating green metrics in various IEEE technical domains". Viewed through the triple bottom line of sustainability (economic, environmental and social), the Green ICT Initiative offers a compelling opportunity for IEEE to demonstrate the full impact of the technology innovation it supports.

Huge efforts have been deployed over the past two and a half years by a dedicated group of initiative volunteers from academia and industry. Nine new IEEE Green ICT standards are currently in development within three working groups and are the result of inter-Society and multidisciplinary collaboration. New publication opportunities have been created and raising awareness of the Green ICT opportunity has been an initiative focus through events of various IEEE Societies. The Initiative is now extending this collaborative approach by reaching out to non-traditional IEEE communities within industry, governments and society at large.

The Greening through ICT Summit: Smart Must Be Sustainable

A strategic planning workshop was held last May in Paris between 25 invited stakeholders and Initiative members in order to prepare our October 3rd Greening through ICT Summit. Diversity was a key component of this core group, with representatives from academia, industry and the public policy arena from the Americas, Asia, Europe and Africa.

The May 24th preparatory workshop for the October Summit.

One of its recommendations was that the Green ICT initiative must emphasize even more the crucial role of ICT in building a sustainable future in both an environment-friendly and sustainable way. This ambitious goal can only be achieved through a cross-disciplinary and systematic approach requiring broad participation of stakeholders from different backgrounds.

The inaugural Greening through ICT Summit will be held in Paris, France, on October 3, 2017. The unique event, held under the theme "Sustainability in a Connected World", will draw a diverse audience of ICT researchers, practitioners, technology providers, standards experts, and policy- and decision-makers, to collectively identify the technical, commercial and public policy challenges that must be addressed to achieve sustainability through ICT. Through unparalleled open discussions, networking and education opportunities, GtICT delegates will explore solutions for overcoming challenges to achieving sustainability in today's increasingly connected world.

Smart technologies hold untold potential for enhancing the lives of millions around the world. However, as emphasized by initiative co-chairs Charles Despins, Jaafar Elmirghani and Thierry Klein, these innovations must drive disruptive change with minimal environmental consequence, and with an eye toward maximum sustainability. Moreover, recent major developments in the Paris Climate Change agreement have moved prominent world leaders to reiterate their commitments, and this further heightens the importance of IEEE's role and actions in the sustainable development sphere.

As ICT must be part of the search for climate change solutions and the adoption of related public policies, the Summit will stimulate a dialogue between those who design technology, those who use it and those who regulate it. These specialists will share their experiences and challenges that impact energy efficiency, greenhouse gas emissions, technology life cycle management, public policy and social issues. This conversation is key to humanity's well-being in the 21st century and beyond.

Designed to stimulate interaction in an impartial environment, the event will consider an array of topics and viewpoints, including:

- ICT as a driver of sustainability through digital economy strategies.
- Convergence of power grids with the Internet.
- Impacts of regulatory and adoption policies on climate change initiatives.
- Improved management of technology life cycles.
- Vertical market reactions to the ICT-driven paradigm shift.

Graphic compilation of May 24th discussions (to be also used at October Summit).

(Continued on Newsletter page 4)

CHAPTER REPORT

IEEE ComSoc Kolkata Chapter: Winner of the 2016 Chapter Achievement Award

By Prof. Iti Saha Misra, ComSoc Kolkata Chapter Chair

To begin, ComSoc's Chairperson of the Kolkata Section, Prof. Iti Saha Misra, would like to convey her sincere thanks to the leaders of IEEE ComSoc Global Activities for considering the Kolkata Chapter as one of the winners for CAA-2016. The year 2016 was a very eventful year for the IEEE ComSoc Kolkata Chapter.

The Chapter organized two DLTs. On 20 June Prof. Luiz Da Silva of Trinity College Dublin delivered a lecture on "Sharing Infrastructure and Spectrum: A Vision for the Future of Mobile Networks" with 65 attendees from academia and industry; and on 29 September, Prof. Anura P. Jayasumana of Colorado State University, USA, delivered a lecture on "Internet of Things: A Pervasive Technology for Innovation", with 70 attendees. In addition to organizing these DLTs, the Chapter actively participated and provided technical sponsorship for International Conferences, RAIT-2016 at ISM Dhanbad and ICACNI-2016 at NIT Rourkela. The Chapter also organized six IEEE lecture meetings, three out-reach programs, a one-day Industry seminar, and a four-day workshop with the participation of the new ComSoc Student Branch Chapter at MSIT, Kolkata. The Chapter also convened six executive committee meetings, and participated in two IEEE 5G Summits held in Delhi and Patna.

Chapter activities started at the beginning of the year by hosting an IEEE Technical Lecture meeting on 8 January 2016 delivered by Dr. Samik Sengupta, Dept. of Computer Science and Engineering at the University of Nevada, Reno (UNR) with 60 participants, followed by almost monthly programs through the year. In all programs 10 to 20 IEEE members participated. To partially cover expenses the chapter has taken a small amount of registration fees for one DLT and collaborated with other institutes. In 2016 the one-day industry participated program was financed by Moulana Abul Kalam Azad University of Technology, Salt Lake Kolkata, attended by 120 undergraduate students and research scholars.

Organizing the DLT program on 29 September at Meghnath Saha Institute of Technology (MSIT) had special significance as the newly formed ComSoc Students Branch Chapter (30 students) was the co-organizer of the DLT. Moreover, it was the first DLT program in the MSIT where the young faculty members and students learned about the IEEE Distinguished Tour quality lecture. The DLT lecture topic was on the very relevant upcoming technology of "Internet of Things: A Pervasive Technology for Innovation", delivered by Prof. Anura Jayasumana.

The concept of the IEEE 5G Summit in Kolkata held on 17-18 March 2017 has been formalized through the ComSoc representation of IEEE 5G Summit at IIT Patna.

The IEEE Patna 5G Summit was the fourth in the series and the first one in India, and has held at IIT Patna on 29 March 2016 that provided a rich platform for industry leaders, innovators, and researchers from the industry and academic community to speak about 5G technology. The then Hon. Minister of IT (India), Sri Ravi Sankar Prasad graced the Inaugural program along with distinguished Professor Ramjee Prasad, Fellow IEEE and Founder President CTIF Global Capsule. Chairperson of the IEEE ComSoc Chapter, Prof. Iti Saha Misra, highlighted ComSoc activities for 5G initiatives around the world. She is one of the members of 5G Summit steering Committee.

Sri Ravi Sankar Prasad, Prof. Ramjee Prasad , Prof. Iti Saha Misra at the Lighting of the Lamp for 5G Summit.

MSIT Students Branch Chapter members and ComSoc organizers with DLT speaker, Prof. Anura Jaya Sumana at the centre (sitting).

DLT speaker Prof. Luiz Da Silva delivering his lecture at the K.P. Basu Memorial, Jadavpur University lecture.

RAIT (Recent Advancement in Information Technology) has become the signature conference of ComSoc, technically sponsored by the IEEE ComSoc Kolkata Chapter for the fourth time, (RAIT-2016) organized by The Department of Computer Science & Engineering of Indian School of Mines, Dhanbad 3-5 March 2016, which was a very technically rich and well organized conference, with the conference proceedings published in IEEE Explore.

ICACNI-2016, the 4th International Conference on Advanced Computing, Networking and Informatics, was hosted at NIT Rourkela, Odisha, India on 22-24 September 2016, technically co-sponsored by IEEE ComSoc Kolkata Chapter. There was a special event featuring a ComSoc student paper contest. The first and second prize winners based on quality of work and presentation were awarded Rs. 2500 and Rs. 1500, respectively, along with a certificate by ComSoc Kolkata Chapter.

The ComSoc Kolkata Chapter organized three outreach programs in 2016. In each program the ComSoc Chairperson delivered a motivational speech for the faculty, researchers and students to join IEEE followed, by a technical lecture. One such collaborative seminar program was held on 23 April 2016 jointly with IETE at Tripura State Centre along with 50 students studying engineering.

CONFERENCE REPORT

Security of Things Conference & Workshop at CommunicAsia 2017, Singapore

By Leo Hwa Chiang and Ewell Tan, IEEE APO Office, Singapore

The CommunicAsia 2017 Summit once again dominated the arena, being the most highlighted information and communications technology (ICT) event on 23-25 May 2017, at the Singapore Marina Bay Sands Convention Centre. This summit gathered ICT stakeholders, regulators, authorities, telecommunication firms, enterprises, technology firms, analysts and industry experts. A total of 40,000 trade professionals attended the exhibitions and the conference tracks.

Hwa (Leo) Chiang, the Director of IEEE Asia Business Development, orchestrated and curated a "Security of Things" Conference & Workshop on behalf of the IEEE Communications Society. This three-day conference attracted more than 50 participants and was one of the most successful and popular tracks among all the seven conference sessions. We invited a total of 18 speakers and seven workshop presenters. They hailed from the Singapore Cyber Security Agency, Nomura Research Institute, Edith Cowan University (ECU), Singapore University of Technology and Design (SUTD), KPMG, etc.

Security of Things Conference & Workshop.

Exhibit booth and trade show at CommunicAsia 2017.

In addition, also accepting an invitation was His Excellency, Mr. Mohamed Abulkheir, Egypt Ambassador to Singapore, who represents the United Nations Group of Government Experts (UNGGE) in Cyber Security to share the United Nations' policy with regards to cyber threats.

In collaboration with IEEE Marketing, Sales and Design (MSD) and the IEEE authorized dealer, InfoHost, Ewell Tang, the Project Manager of IEEE Asia-Pacific Limited, provided support and assistance in managing the IEEE Communications Society exhibit booth, to promote IEEE Xplore Digital Library, IEEE and IEEE Communications Society membership.

It was a productive and successful event in which all colleagues from IEEE Asia-Pacific, IEEE MSD and InfoHost worked together to make happen.

IEEE Green ICT/*Continued from page 2*

The Summit will be conducted under a stimulating "World Café" approach in order to foster a productive, open dialogue in a relaxed and casual atmosphere. The event will culminate with a declaration of a "Green ICT Vision for the Future" driven by the day's discussions.

The IEEE Green ICT Initiative has just held its first IEEE Young Professional Green ICT Idea Competition. Aimed at harnessing the input of younger membership with IEEE, notably those who will live with the consequences of climate change, the competition encouraged young professionals to submit creative Green ICT ideas and solutions. Competition winners will have the opportunity to present their ideas at the Summit.

The IEEE GtICT's theme of smart, sustainable technologies will also be highlighted during Le Salon de la Ville et des Territoires Intelligents, Durables et Connectés, a smart city and smart grid exhibition taking place immediately after the GtICT Summit, October 4-5, 2017 at the Paris Porte de Versailles.

We are looking forward to seeing many readers at the October 3 Summit in the 9th arrondissement of Paris (Opera district). Registration is now open through the Initiative Web portal at http://greenict.ieee.org.

GCN GLOBAL COMMUNICATIONS NEWSLETTER

Stefano Bregni
Editor-in-Chief
Politecnico di Milano — Dept. of Electronics and Information
Piazza Leonardo da Vinci 32, 20133 MILANO MI, Italy
Tel: +39-02-2399.3503 — Fax: +39-02-2399.3413
Email: bregni@elet.polimi.it, s.bregni@ieee.org

Fabrizio Granelli
Associate Editor-in-Chief
University of Trento
Email: fabrizio.granelli@unitn.it

IEEE Communications Society

Stefano Bregni, Vice-President for Member and Global Activities
Carlos Andres Lozano Garzon, Director of LA Region
Scott Atkinson, Director of NA Region
Andrzej Jajszczyk, Director of EMEA Region
Takaya Yamazato, Director of AP Region
Curtis Siller, Director of Sister and Related Societies

Regional Correspondents who Contributed to this Issue

Ewell Tan, Singapore (ewell.tan@ieee.org)

IEEE ComSoc IEEE Communications Society

www.comsoc.org/gcn
ISSN 2374-1082

Fog World Congress™

October 30 - November 1, 2017 | Santa Clara, California

Jointly produced by IEEE ComSoc IEEE Communications Society · OpenFog

4 TRACKS TO CHOOSE FROM:

Interactive Track

- Participate in the fog hackathon
- Gain insight into the latest research through the presentation of peer-reviewed papers in our conference program

Business Track

- Experience hands-on learning in our interactive workshop
- Visit the exhibits and demo area to see offerings from companies and universities who are leading the way in fog

Technology Track

Research Track

Don't miss this opportunity to be a part of the first global conference on fog computing.

www.fogworldcongress.com

For more information, please contact **info@fogworldcongress.com**.

Potential sponsors or exhibitors, please contact **greg@fogworldcongress.com**.

INTERNET OF THINGS: PART 4


Christos Verikoukis


Roberto Minerva


Mohsen Guizani


Soumya Kanti Datta


Yen-Kuang Chen


Hausi A. Muller

The Internet of Things (IoT) is seen as a set of vertical application domains that share a limited number of common basic functionalities. In this view, consumer-centric solutions, platforms, data management, and business models have to be developed and consolidated in order to deploy effective solutions in the specific fields. The availability of low-cost general-purpose processing and storage systems with sensing/actuation capabilities coupled with communication capabilities are broadening the possibilities of IoT, leading to open systems that will be highly programmable and virtualized, and will support large numbers of application programming interfaces (APIs). IoT emerges as a set of integrated technologies and new exciting solutions and services that are set to change the way people live and produce goods. IoT is considered by many as a fruitful technological sector in order to generate revenues. IoT covers a large wealth of consumer-centric technologies and it is applicable to an even larger set of application domains. Innovation will be nurtured and driven by the possibilities offered by the combination of increased technological capabilities, new business models and the rise of new ecosystems.

This Feature Topic (FT) issue addresses several promising approaches to sensors, actuators, and new consumer devices. New communication capabilities (from short-range to LPWAN to 4G and 5G networks, with NB-IoT). In addition, new communication protocols and the exploitation of NFV/SDN for better communications are considered:

- New solutions for large distributed systems (e.g., combination of cloud, grid, and edge/fog computing)
- New business models and ecosystems
- Consumer-centric aspects including IoT application development, utilization of semantics, and security, privacy, and trust

This timely Feature Topic (FT) issue has gathered articles from a wide range of perspectives from different industrial and research communities in IoT.

In response to the Call for Papers, 103 high-quality manuscripts were received, and after a very careful review process four outstanding papers have been selected for Part 4 of this Feature Topic, giving an overview of recent developments in IoT platforms for smart cities, interoperability, business models, security, privacy, LoRaWAN, energy efficiency, and channel access methods.

In the first article, Mehmood *et al.* devise a taxonomy to best bring forth a generic overview of the IoT paradigm for smart cities with major requirements such as integrated information, information and communication technologies (ICT), and network types. The authors give an overview of existing open source IoT platforms for realizing smart city applications followed by several exemplary case studies. Then they summarize the latest synergies and initiatives to promote smart cities' IoT with a list of research challenges.

The world is expected to face a food shortage in the very near future. IoT technologies have great potential for applications in food and agriculture. In the second article, Brewster *et al.* outline the challenges and constraints of such technologies in terms of the execution of IoT-based large-scale pilots (LSPs). The importance of addressing the interoperability challenges, new business models, and security, privacy, and data governance are reviewed. Finally, they highlight the fact that for IoT to be successful in this domain, a significant change of culture is required.

In the third article, "Understanding the Limits of LoRaWAN," F. Adelantado *et al.* discuss low-power wide area networking (LPWAN) technology, which offers long-range communication. LoRaWAN is one of the most used services. It promises ubiquitous connectivity in outdoor IoT applications, while keeping network structures, and management. But the technology has limitations that need to be addressed before its widespread use. The authors give an overview of the capabilities and limitations of LoRaWAN and list open research challenges.

Energy efficiency is an important topic in IoT deployment. H. Elhammouti *et al.*, in the fourth article, propose that better energy efficiency can be achieved by targeting satisfactory quality of service levels only, and they propose a game theoretical solution, the satisfaction equilibrium. Moreover, the authors propose fully distributed schemes in order to reach efficient satisfaction equilibria in both slow- and fast-fading channels.

In the fifth article, D. Zuchetto *et al.* provide a comparative overview of the uncoordinated channel access methods for IoT technologies, that is, ALOHA-based and Listen Before

Talk (LBT) schemes, in relation to the European Telecommunications Standards Institute (ETSI) and U.S. Federal Communications Commission (FCC) regulatory frameworks. In addition, they provide a performance comparison of these access schemes, in terms of successful transmissions and energy efficiency, in a typical IoT deployment.

The sixth article deals with Third Generation Partnership Project (3GPP) solutions for enabling massive cellular IoT and investigates the random access strategies for machine-to-machine communications. M.Shirvanimoghaddam *et al.* propose a massive non-orthogonal multiple access (NOMA) technique, which is presented as a promising solution to support a massive number of IoT devices in cellular networks.

T. Damla *et al.*, in the seventh article, debate whether the hype on IoT will one day be a reality. They provide failed testbed examples and state the reasons behind those failures. They introduce a new concept called value of information (VoI) that they apply to such IoT technologies. Then they discuss a formal model for the VoI and cost of privacy. Finally, they give some directions on future research in this fundamental area.

BIOGRAPHIES

CHRISTOS VERIKOUKIS [S'95, M'04, SM'07] (cveri@cttc.es) got his Ph.D. from Universidad Politècnica de Catalunya in 2000. He is currently a fellow researcher at CTTC, head of the SMARTECH Department, and an adjunct associate professor at the University of Barcelona. He has published 113 journal papers and over 180 conference papers. He is also a co-author of three books, 14 chapters in other books, and two patents. He is currently Chair of the IEEE ComSoc CSIM Technical Committee.

ROBERTO MINERVA holds a Ph.D. in computer science and telecommunications from Telecom Sud Paris, France, and a Master's degree in computer science from Bari University, Italy. He is the Chairman of the IEEE IoT Initiative, an effort to nurture a technical community and foster research in IoT. He is at TIMLab, involved in activities on SDN/NFV, 5G, big data, and architectures for IoT. He is the author of papers published in international conferences, books, and magazines.

MOHSEN GUIZANI [S'85, M'89, SM'99, F'09] received his B.S. (with distinction) and M.S. degrees in electrical engineering, and M.S. and Ph.D. degrees in computer engineering from Syracuse University in 1984, 1986, 1987, and 1990, respectively. He is currently a professor and the Electrical and Computer Engineering Department Chair at the University of Idaho. He currently serves on the Editorial Boards of several international technical journals. He is the author of nine books and more than 450 publications in refereed journals and conferences.

SOUMYA KANTI DATTA is a research engineer at EURECOM and a co-founder of an IoT startup, Future Tech Lab. His research focuses on innovation, standardization, and development of next-generation technologies in mobile computing, IoT, M2M communication, and security. He is an active member of the IEEE Consumer Electronics Society and W3C. He has published more than 40 papers in top IEEE conferences and journals. Currently he is involved in oneM2M and the W3C Web of Things Group.

YEN-KUANG CHEN [F'12] received his Ph.D. degree from Princeton University. He is a principal engineer at Intel Corporation, Santa Clara, California. His research areas span emerging applications that can utilize the true potential of IoT and computer architecture that can embrace emerging applications. He has 50+ U.S. patents, 20+ pending patent applications, and 90+ publications. He is the Editor-in-Chief of the *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*. He is a Distinguished Lecturer of the IEEE Circuits and Systems Society, 2016–2017.

HAUSI A. MULLER is a professor in the Department of Computer Science and associate dean of research in the Faculty of Engineering at the University of Victoria. He is a member of the IEEE Computer Society Board of Governors and the 2016–2017 Vice-President of the IEEE CS Technical and Conference Activities Board. His research interests include software engineering, software evolution, IoT, smart cyber physical systems, and self-adaptive systems. He is a Fellow of the Canadian Academy of Engineering.

Internet-of-Things-Based Smart Cities: Recent Advances and Challenges

Yasir Mehmood, Farhan Ahmad, Ibrar Yaqoob, Asma Adnane, Muhammad Imran, and Sghaier Guizani

The authors devise a taxonomy to best bring forth a generic overview of the IoT paradigm for smart cities, integrated ICT, network types, possible opportunities, and major requirements. Moreover, an overview of the up-to-date efforts from standard bodies is presented.

ABSTRACT

The Internet of Things is a novel cutting edge technology that proffers to connect a plethora of digital devices endowed with several sensing, actuation, and computing capabilities with the Internet, thus offering manifold new services in the context of a smart city. The appealing IoT services and big data analytics are enabling smart city initiatives all over the world. These services are transforming cities by improving infrastructure and transportation systems, reducing traffic congestion, providing waste management, and improving the quality of human life. In this article, we devise a taxonomy to best bring forth a generic overview of the IoT paradigm for smart cities, integrated ICT, network types, possible opportunities and major requirements. Moreover, an overview of the up-to-date efforts from standard bodies is presented. Later, we give an overview of existing open source IoT platforms for realizing smart city applications followed by several exemplary case studies. In addition, we summarize the latest synergies and initiatives worldwide taken to promote IoT in the context of smart cities. Finally, we highlight several challenges in order to give future research directions.

INTRODUCTION

The Internet of Things (IoT) is a revolutionary communication paradigm that aims to bring forth an invisible and innovative framework to connect a plethora of digital devices with the Internet. Thus, it intends to make the Internet more immersive and pervasive [1]. The emerging IoT market is continously gaining momentum as operators, vendors, manufacturers, and enterprises begin to recognize the opportunities it offers. According to the latest IDC forecast,[1] the worldwide IoT market will reach US$1.7 trillion in 2020 up from US$655.8 billion in 2014 with a compound annual growth rate of 16.9 percent. The devices alone are expected to represent 31.8 percent of the total worldwide IoT market in 2020. This greater percentage of the revenue in 2020 is expected through building IoT platforms, application softwares, and service-related offerings.

A smart city is a complex ecosystem characterized by the intensive use of information and communications technologies (ICT), aiming to make cities more attractive and more sustainable, and unique places for innovation and entrepreneurship [2]. The major stakeholders include application developers, service providers, citizens, government and public service providers, the research community, and platform developers. Furthermore, the smart city cycle consists of numerous ICT technologies, development platforms, maintenance, and sustainability, apps for evolving citizens, and technical, social, as well as economic key performance indicators (KPIs). Consequently, IoT systems will play a fundamental role in the deployment of large-scale heterogeneous infrastructures. A high-level illustration of an IoT-based smart city is given in Fig. 1.

IoT-based smart city applications can be categorized on the basis of network type, scalability, coverage, flexibility, heterogeneity, repeatability, and end-user involvements [3]. In general, these applications can be grouped into personal and home, utilities, mobile, and enterprises. For instance, *personal and home* applications include ubiquitous e-healthcare services to live independently via body area networks (BANs), which help doctors monitor patients remotely. *Utilities* applications include smart grid, smart metering/monitoring, water network monitoring, and video-based surveillance. Similarly, *mobile* applications include intelligent transportation system (ITS) and logistics, traffic management, congestion control, and waste management. Additionally, IoT-based enterprise applications usually consist of a network of things within a work environment.

Several research efforts have been made to integrate IoT with smart city environments. For instance, Zanella *et al.* [1] presented a comprehensive survey of the architectures, protocols, and enabling technologies for a web-service-based IoT framework in the Padova smart city project. The proof of concept implementation with numerous technical solutions aims to monitor street lighting, the quality of air, and identification of the most critical issues. A survey on the fundamental IoT elements in realizing smart cities was conducted in [4], which also described a case study on noise monitoring. Nathalie *et al.* [5] proposed a different perspective of smart cities in which IoT devices were considered service providers mimicking cloud-based services. The proposal offered a higher level of abstraction to deploy innovative ubiquitous applications by eliminating the barriers between physical IoT devices and the logical (cloud service providers) world. A generic top-down smart city architecture was proposed in [6]

[1] https://www.telecompaper.com/news/global-iot-market-to-reach-usd-17-tln-in-2020-idc–1085269, accessed October 20, 2016.

Digital Object Identifier: 10.1109/MCOM.2017.1600514

Yasir Mehmood (corresponding author) is with the University of Bremen; Farhan Ahmad and Asma Adnane are with the University of Derby; Ibrar Yaqoob is with the University of Malaya; Muhammad Imran is with King Saud University; Sghaier Guizani is with Alfaisal University.

in which service providers play a role of central information unit that is connected to a set of IoT-based services. It also offers IoT convergence and acceptance of numerous ICT technologies for realizing smart cities.

Although several studies exist on IoT and smart cities, convergence of these two areas grants further academic efforts for the flourishing of IoT-based smart cities. Thus, unlike other studies, this article best bring forths an IoT-based smart cities taxonomy, prime open source platforms, and case studies of recent deployments, as well as unearthing several open research challenges. The contributions of this study are as follows:

- First, we devise a taxonomy of the IoT-based smart city environment.
- We present an overview of major open platforms for smart cities.
- Further, we present recent synergies and a number of case studies on various smart city deployments reported by various enterprises.
- Finally, we unearth several IoT-related open research challenges to give future directions.

IoT-Based Smart City Taxonomy

This section presents a taxonomy of IoT-based smart cities that categorizes the literature on the basis of existing communication protocols, major service providers, network types, standardization efforts, offered services, and crucial requirements. An overview of the devised smart city taxonomy is depicted in Fig. 2.

Communication Protocols

IoT-based smart city realization significantly relies on numerous short- and wide-range communication protocols to transport data between devices and back-end servers. The most prominent short-range wireless technologies include ZigBee, Bluetooth, Wi-Fi, WiMAX, and IEEE 802.11p, which are primarily used in smart metering, e-healthcare, and vehicular communication. Wide-range technologies such as Global System for Mobile Communication (GSM) and general packet radio service (GPRS), Long Term Evolution (LTE), and LTE-Advanced are commonly utilized in ITS such as vehicle-to-infrastructure (V2I), mobile e-healthcare, smart grid, and infotainment services. Additionally, LTE-M is considered as an evolution for cellular IoT (C-IoT). In Release 13, the Third Generation Partnership Project (3GPP) plans to further improve coverage, battery lifetime, and device complexity [7]. Besides well-known existing protocols, the LoRa Alliance is standardizing the LoRaWAN protocol to support smart city applications, primarily ensuring interoperability between several operators. Moreover, SIGFOX is an ultra narrowband radio technology with full star-based infrastructure that offers a highly scalable global network for realizing smart city applications with extremely low power consumption. A comparative summary[2] of the major communication protocols is presented in Table 1.

Service Providers

Pike Research estimated that the smart cities market will grow to hundreds of billions of dollars by 2020, with an annual growth of nearly US$16 billion. IoT is recognized as a potential source to increase the revenue of service providers. Thus, well-known worldwide service providers have already started exploring this novel cutting-edge communication paradigm. Major service providers include Telefonica, SK Telecom, Nokia, Ericsson, Vodafone, NTT DOCOMO, Orange, Telenor, and AT&T, which offer a variety of services and platforms for smart city applications such as ITS and logistics, smart metering, home automation, and e-healthcare.

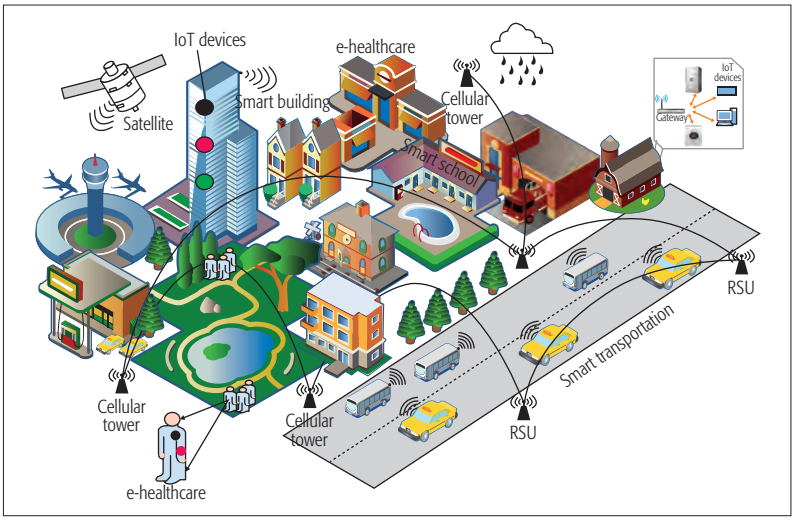

Figure 1. An illustration of an IoT-based smart city.

Network Types

IoT-based smart city applications rely on numerous network topologies to accomplish a fully autonomous environment. The capillary IoT networks offer services over a short range. Examples include wireless local area networks (WLANs), BANs, and wireless personal area networks (WPANs). The application areas include indoor e-healthcare services, home automation, and street lighting. On the other hand, applications such as ITS, mobile e-healthcare, and waste management use wide area networks (WANs), metropolitan area networks (MANs), and mobile communication networks. The above networks pose distinct features in terms of data, size, coverage, latency requirements, and capacity.

Activities of Standard Bodies

The vast smart city applications not only demand large scale deployment of numerous kinds of IoT devices, but also require device interoperability. Therefore, most prominent governing bodies such as the Internet Engineering Task Force (IETF), Third Generation Partnership Project (3GPP) European Telecommunications Standards Institute (ETSI), oneM2M, IEEE, and Open Mobile Alliance (OMA) are actively involved in developing standards to support smart city applications on a large scale. This section discusses the major contributions and ongoing activities of the prime standard bodies for enabling smart city applications.

IETF: The first IETF working group (WG), 6LoWPAN, standardized techniques for handling IoT small packets using header compression and neighbor discovery optimization. Moreover, the Routing Over Low-power and Lossy networks (ROLL) WG standardized Routing Protocol for Low Power and Lossy Networks (RPL) for smart

[2] https://www.global-logic.com/wp-content/uploads/2015/12/The-role-of-telecommunications-in-smart-cities.pdf, accessed December 10, 2016.

IoT offers diverse applications in a smart city, thus demands numerous requirements. For instance, IoT-based solutions are expected to be low cost, low energy consumption, high quality-of-service (QoS), wider coverage, increased flexibility, high security and privacy, ultra-dense deployments, and multivendor interoperability.

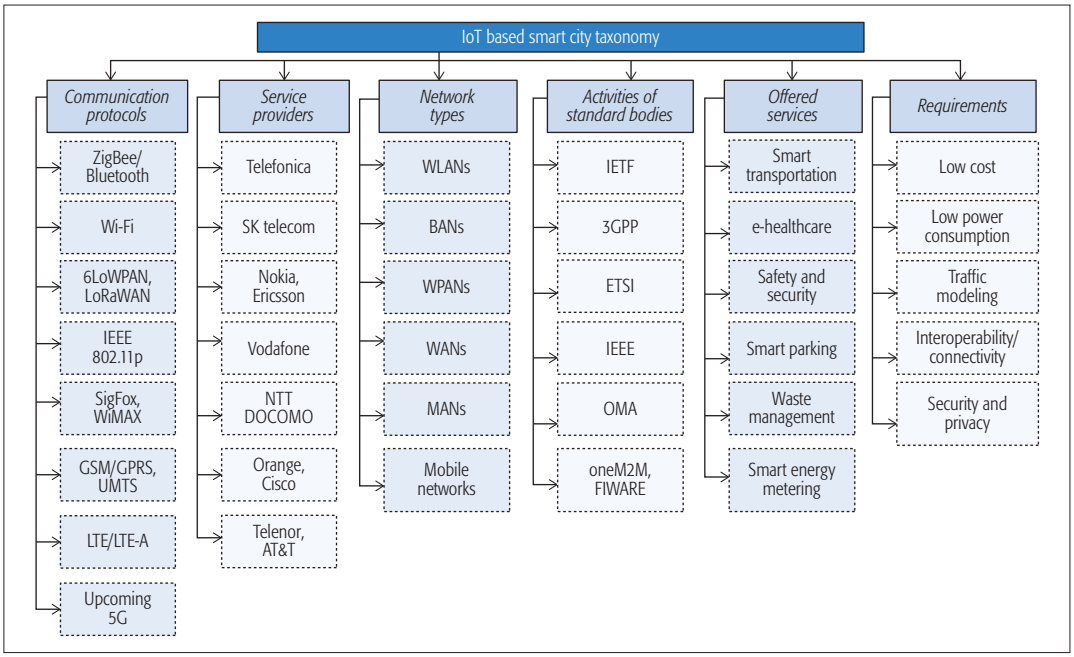

Figure 2. A representation of IoT-based smart city taxonomy.

city applications. In addition, several IETF WGs such as DICE are active in standardizing security profiles, such as Transport Layer Security (TLS) and Datagram-TLS (DTLS) for constrained IoT devices.

3GPP: in its latest Release 13, 3GPP standardized narrowband IoT (NB-IoT) to provide better network coverage for smart city applications by further reducing the bandwidth to 200 kHz (uplink/downlink), reducing throughput on a physical resource block (PRB) level, supporting massive IoT devices and low power consumption, and enhancing coverage extension by 20 dB [10]. As a result, NB-IoT meets the application requirements in the industrial, public, personal, and home domains. Additionally, 3GPP is introducing extended discontinuous transmission/reception (eDTX/eDRX) techniques in Release 13 to further reduce power consumption, and thus increase the device operating time.

ETSI: ETSI aims to deliver interoperable and cost-effective solutions to support smart city applications. Particularly, oneM2M is the global initiative by ETSI in cooperation with the member research institutes such as the Broadband Forum, OMA, and Continua to support IoT connectivity on a large scale. oneM2M aims to develop a single horizontal platform for enabling interoperability among all applications through a distributed software layer.[3] Additionally, it delivers architecture, requirements, application programming interface (API) specifications, and privacy and security solutions, as well as an interoperability framework for smart city applications. Consequently, the standardized APIs and open interfaces can be used within several systems for enabling a plethora of IoT devices to connect worldwide with the backend servers.

IEEE: In the context of smart cities, IEEE mainly focuses on optimization of the air interface for ultra-IoT deployments. Furthermore, IEEE focuses on the use of sub-6 GHz spectrum for IoT connectivity to support numerous smart city applications.

OMA: *OMA* standardized the OMA Lightweight M2M (OMALWM2M) protocol for resource constrained IoT device management for both sensor and cellular networks. OMALWM2M is located at the device end, and offers a communication path between an LWM2M client and an LWM2M server. Therefore, OMALWM2M is a light and compact protocol that is frequently used with the Constrained Application Protocol (CoAP), and offers an efficient resource data model for the resource constrained IoT devices. Additionally, it provides a choice for service providers to deploy IoT systems for supporting corresponding smart city applications.

Offered Services

IoT offers numerous services that are of great interest in the context of smart cities to not only improve the quality of human lives, but also leverage the city administration by reducing the operational costs [1]. Major offerings include smart lighting, waste and water management. For instance, smart IoT modules can be deployed within grid stations, homes, and workplaces for distributing and consuming energy efficiently. In *e-healthcare*, IoT devices can be positioned on the bodies of patients for monitoring health parameters such as temperature, pulse rate, and sugar level, and provide opportunities for doctors to regularly monitor their patients. Besides, *urban IoTs* can provide solutions to control traffic congestion through monitoring of traffic intensity using either GPS services in modern vehicles or WANs. In *waste management*, the truck route can be optimized based on the load level indication by smart waste containers. Consequently, it enhances the quality of recycling by reducing the cost of waste collection.

Requirements

IoT offers diverse applications in a smart city, and thus demands numerous requirements. For instance, IoT-based solutions are expected to have

[3] oneM2M-TS-0001-V-2014-08, accessed June 5, 2016.

Technology	Operating frequency	Data rate	Coverage	Latency	Power usage	Use cases
ZigBee	2.4 GHz 868 MHz, 915 MHz	250 kb/s	50–100 m	16 ms	Low	Smart metering, indoor e-healthcare
Bluetooth	2.4 GHz	25 Mb/s	10 m	100 ms	Low	Indoor e-healthcare
Wi-Fi	2.4 GHz5 GHz, 802.11n	54 Mb/s, 6.75 Gb/s	140 m 100 m	46 ms	Medium	Metering, waste management automation, energy management, infotainment, automation
IEEE 802.11p	5.85–5.925 GHz	6 Mb/s	1000 m		Low	Vehicular communication, V2V/V2I, infotainment
DSRC/WAVE	5.8, 5.9 GHz	6 Mb/s	1000 m	200 μs	Low	ITS (V2V/V2I)
DASH7	433, 868, 915 MHz	55.5 kb/s, 200 kb/s	1000 m	15 ms	Low	ITS, automation
6LoWPAN	2.4 GHz, 868, 915 MHz	250 kb/s	100 m		Low	ITS, smart metering, logistics
LoRaWAN	433, 868, 780, 915 MHz	50 kb/s	2–5 km		Low	ITS, smart metering, waste management
GSM/GPRS	850, 900, 1800, 1900 MHz	80–384 kb/s	5–30 km	1.5–3 s	High	ITS, smart metering, m-health, energy management, logistics, infotainment
3G	850 MHz	3 Mb/s	5–30 km	100 ms	High	ITS, smart metering, energy management, m-health, logistics, infotainment
LTE/LTE-Advanced	700, 750, 800, 1900, 2500 MHz	1 Gb/s, 500 Mb/s	5–30 km	5 ms	High	ITS, smart metering, mobile health, logistics, infotainment

Table 1. A summary of major communication protocols for realizing IoT-based smart cities [8].

low cost, low energy consumption, high quality of service (QoS), wider coverage, increased flexibility, high security and privacy, ultra-dense deployments, and multivendor interoperability. To fulfill the above requirements, several new techniques must be adopted. For instance, traffic modeling can play an essential role in handling massive IoT traffic. Therefore, instead of using the traditional source traffic modeling approach, where each IoT device accesses the network individually for sending and receiving important messages, an aggregated traffic modeling approach must be commonly used, as illustrated in Fig. 3. In this way, several IoT devices can share scarce resources for sending and receiving small-sized data. To achieve this, a gateway can be deployed, which may operate using any of the existing communication technologies. For instance, the data from IoT devices can be transported to a gateway using short-range communication standards such as ZigBee, Bluetooth, Wi-Fi, and other dedicated short-range communications (DSRC) protocols. In addition, LTE and LTE-A relays and femtocells can be deployed to perform data aggregation.

Aggregated traffic modeling can improve the performance of IoT networks by supporting enormous numbers of devices. Since tiny IoT systems also demand low power consumption in order to increase device lifetime (up to 10 to 15 years), data aggregation can be used to ensure low energy consumption through improved coverage along with the network capacity. This is achieved by improving the channel conditions using intermediate gateways and relays. Furthermore, provisioning of QoS is also essential in critical IoT applications such as e-healthcare and emergency alert. Therefore, incorporating data differentiation schemes along with the aggregation approach can potentially fulfill QoS requirements and satisfy the delay requirements of critical smart city applications. High-priority traffic (e.g., e-healthcare data) should be served immediately followed by low-priority traffic (e.g., regular monitoring). Besides traffic modeling, ultra-dense IoT realization demands strong collaboration among modem manufacturers and vendors in order to increase interoperability. Consequently, this will further reduce the associated cost factors. Besides, data security and integrity is of significant importance to ensure safe and secure IoT communications.

Besides traffic modeling, ultra-dense IoT realization demands strong collaboration among modem manufacturers and vendors in order to increase interoperability. Consequently, this will further reduce the associated cost factors. Besides, data security and integrity is of significant importance to ensure safe and secure IoT communications.

IoT Open Source Platforms

Open source implementations always play a vital role in sharing information in order to achieve multi-vendor interoperability. Worldwide, the following prime open communities offer easy and fast development platforms for smart cities.

FIWARE[4]

FIWARE is a standard open platform for realizing smart city applications. It was launched by the European Commission, and aims to develop the core future technologies in the IoT paradigm. It is based on software components named generic enablers. These components provide common functionalities to the multiple vertical sectors with

[4] http://www.fiware.org/about-us/, accessed October 10, 2016..

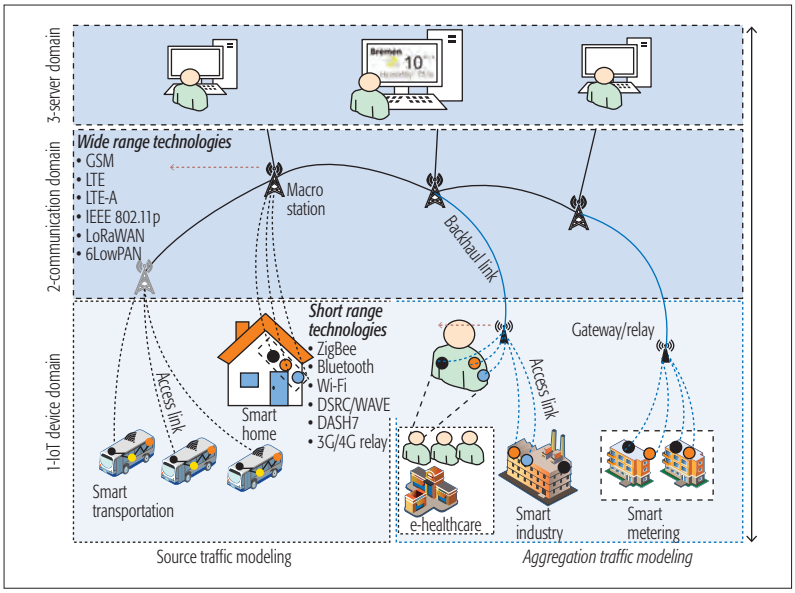

Figure 3. An illustration of traffic modeling for an IoT based smart city.

the objective of enabling interoperability among them. FIWARE enablers are classified into seven wide technical categories: cloud hosting, applications, services and data delivery, security, interface to networks and devices, advanced web, IoT services enablement, and data/context management. Moreover, FIWARE provides a simple and powerful set of APIs that ease the development of smart applications. Despite many advantages of FIWARE, the lack of a complete set of functionalities is one of the key issues that remains to be addressed.

OCEAN[5]

The Open Alliance for IoT Standard (OCEAN) was initiated in January 2015 by the Korea Electronics Technology Institute (KETI) and the Korean government. It is a global alliance that aims to bring forth open source implementations for smart cities based on IoT standards. Additionally, the initiative focuses on promoting the development and commercialization of platforms, products, and services by widespread adoption of IoT standards-compliant open source code. OCEAN is responsible for releasing the source code for IoT standards as open source, and helping vendors and developers collaborate with each other to create new innovative products and services under a global partnership, finally establishing a global IoT ecosystem for smart cities. Currently, OCEAN provides several oneM2M-based platforms for devices, gateways, and servers.

OM2M[6]

The OM2M project was initiated by the ECLIPSE Foundation to deliver an open source implementation of the oneM2M and SmartM2M standards. The primary goal of this initiative is to support the deployment of vertical applications and heterogeneous devices by providing a horizontal machine-to-machine (M2M) service platform for developing services independent of the underlying network. Thus, it provides a horizontal service common entity (SCE), which can be deployed in an M2M device, a gateway, or a server. The major SCE functionalities include application enablement, triggering, notification, security, persistence, interworking, and device management. Additionally, it provides RESTful interfaces for authentication, registration, resources discovery, containers management, synchronous/asynchronous communications, access rights authorization, groups organization, and re-targeting.

Open DayLight IoT Data Management[7]

The IoT Data Management (IoTDM) from the Open DayLight (ODL) project is about developing a data-centric middleware that will represent an IoT broker compliant with oneM2M. It will also authorize applications to recover data uploaded by any IoT device used in a smart city. The ODL platform is used to implement the oneM2M data store, which models a hierarchical containment tree, where each node in the tree represents a oneM2M resource. Typically, IoT devices and applications interact with the resource tree over standard protocols such as CoAP, Message Queue Telemetry Transport (MQTT), and HTTP.

Contiki[8]

Contiki is also an open platform that offers fast and easy development of numerous IoT-based smart city applications. It offers powerful Internet communication to tiny microcontrollers and operates at extremely reduced cost and low power. Additionally, it fully supports standard IPv4 and IPv6 protocols such as UDP, TCP, and HTTP. Besides, it offers support of the latest low-power wireless and mobile networks such as 6LowPAN, CoAP, and multihop RPL. Furthermore, it provides highly efficient memory allocation procedures for various smart city applications.

IoT Synergies and Case Studies

This section presents a number of case studies provided by different enterprises along with the modern IoT synergies for the realization of smart cities. The aim is to provide a summary of the current deployments and recent initiatives of IoT-based solutions to tackle various city-related issues. A detailed summary of the case studies and projects is given in Table 2.

Busan Green u-City[9]

Busan Green u-City (ubiquitous city) is the first IoT-enabled smart city in South Korea. It is one of the modern practical examples of IoT-based smart cities, using a cloud-based infrastructure to improve the efficacy of city management and local business opportunities, improving the quality of human lives. It is a public-private cooperation among the Busan city government, major technology suppliers, Cisco, and South Koreas largest telco, KT, with an approximate investment of US$452 million. The primary objective of this cooperation is to deliver an improved transportation system, e-healthcare services, increased jobs and business opportunities, and improved information accessibility through various devices and communication sources.

Smart Santander[10]

Santander, the Spanish city, is widely recognized as an IoT-based smart city. This *Future Internet Award* winning project is a cooperation

[5] http://www.iotocean.org/main/, accessed June 5, 2016.

[6] http://www.eclipse.org/om2m/, accessed June 5, 2016.

[7] https://www.opendaylight.org/,accessed June 5, 2016.

[8] http://www.contiki-os.org/, accessed October 10, 2016.

[9] http://www.gsma.com/connectedliving/busan-green-u-city-a-successful-example-of-a-smart-city-in-south-korea/, accessed April 15, 2016.

[10] http://www.smartsantander.eu/, accessed April 1, 2016.

City	Country	Population	Solutions	Major partners
Busan	South Korea	3.4 million	Safety service for childrens/elderly, drone based smart marine, smart parking, crosswalk and energy usage	Busan government, Cisco, ETRI, KETI, SK telecom, KT
Santander	Spain	0.1 million	Smart metering of temperature, traffic intensity, humidity, transportation plans, water needs, etc..	Ericsson, Telefonica, Telefonica I+D
Chicago	United States	2.7 million	Smart grid, smart living, emergency alert, reduced crime	Cisco, IBM, Chicago government
Milton Keynes	United Kingdom	0.2 million	Smart transportation, reduced corbon emission, smart energy, water management	Milton Keynes Council, Samsung, Huawei, CATAPULT, Cambridge University

Table 2. A comparison of case studies.

of 15 big companies and institutions including Ericsson, Telefonica, and several universities and research groups in Spain, Greece, Germany, Denmark, the United Kingdom, and Australia. The city is equipped with approximately 20,000 smart IoT devices that perform several intelligent tasks such as measurement of temperature, humidity, speed and position of vehicles, traffic intensity, public transportation conditions and timetables, air quality, and water networks. The acquired sensor data is transmitted to Munoz's laboratory and compiled into a big picture by a central computer. Thus, everything is recorded in this digital city.

Chicago[11]

The project focuses on infrastructure management, economic development, and community engagement to tackle major issues of education, economic development, crime, and transport in Chicago, Illinois. In cooperation with IBM, Chicago deployed around 300,000 smart IoT devices to support smart grid operations. The primary objective of this project is to reduce energy waste to save customers US$170 million. The project developed an analytics platform on Cisco technology that has helped to minimize crime rates in the city. Also, a model was created that has more than 31 variables to predict and prevent rodent infestations. Analytics is also incorporated to identify buildings that are anticipated to become vacant. Numerous apps have been built using 600 datasets of an open city portal to notify citizens about several unwanted situations expected within a territory.

Milton Keynes[12]

The Milton Keynes Smart project is coordinated by the Open University and aims to develop a data hub within the city that collects and manages data received from several smart devices. The project emphasizes control of corbon emissions and support of sustainable growth without deploying additional infrastructure. The project is a collaboration between Samsung, Huawei, CATAPULT, and Cambridge University that aims to bring forth innovative solutions involving the aforementioned issues. Also, the project aims to deliver an efficient transportation system, water usage, and smart energy solutions as well as focusing on business, education, and community engagement activities.

Open Research Challenges

Besides the aforementioned advances, there are several open research issues and challenges in adopting IoT for smart cities. The purpose of discussing these challenges is to give research directions to new researchers in this domain. Table 3 summarizes the future research directions along with their advantages and requirements.

Security, Privacy, and Trust

Security in general is required for every IoT device. As smart cities provide Internet connectivity to an ample variety of devices, security becomes a very critical challenge. According to HP, about 70 percent of IoT devices in a smart city were at risk of attack due to sufficient vulnerabilities such as insufficient authorization, inadequate software protections, and weak encrypted communication protocols [11]. These vulnerabilities instigate various threats and attacks, leading to several issues in terms of security and privacy. In order to design a successful IoT-based smart city, the following issues must be addressed aforementioned:

- Privacy-aware communication for user data should be provided.
- Simple, lightweight, and efficient security solutions should be designed to ensure data authenticity and integrity, and to provide secure communication between IoT devices and a cloud-based application center [13].
- Detailed risk assessment must be performed to identify present and newly emerging attacks based on vulnerabilities and threats. One such risk assessment framework is proposed by ENISA, which identifies possible emerging attacks in ITS [13].
- An active and adequate decentralized trust management system must be designed.
- Users' trust and consent must be ensured by providing strong privacy measures.

Interoperability

Interoperability is the capability of two different devices and networks to communicate with each other for the exchange of important information. Smart cities include IoT devices from a diverse range of domains (e.g., smart metering, e-healthcare, logistics, monitoring, and intelligent transport). In a smart city, interoperability plays a vital role in providing connectivity among devices operating with different communication technol-

According to HP, about 70 percent of IoT devices in a smart city were at risk of attack due to sufficient vulnerabilities such as insufficient authorization, inadequate software protections, and weak encrypted communication protocols. These vulnerabilities instigate various threats and attacks, leading to several issues in terms of security and privacy.

[11] http://www.smartchicagocollaborative.org/category/city-of-chicago/, accessed June 4, 2016.

[12] http://www.mksmart.org/, accessed: May 1, 2016.

ogies. For example, smart metering uses WLAN technologies as the underlaying communication protocols, while ITS mainly utilize DSRC and mobile technologies for communication. According to the World Economic Forum, interoperability between devices from different domains is a major barrier to IoT success due to lack of universal standards [14]. To overcome this barrier, interoperability issues should be identified at different levels (e.g., device, network, communication, application, and platform). To address these issues, an intelligent and holistic approach is required to provide connectivity to billions of IoT devices. For instance, standardization of oneM2M and FIWARE is a major step in overcoming the interoperability issues with the collaboration of world's largest standardization bodies such as ETSI, 3GPP, and OMA.

Low-Power and Low-Cost Communication

Usually, IoT devices are small in nature and equipped with a group of sensors. In order to operate these devices, a continuous source of energy is constantly required, which poses a significant challenge in terms of battery life and cost. To address these issues in IoT-based smart cities, the devices must feature low power consumption at very low cost. This can be achieved through advancements in the domain of wireless communication and micro-electronics.

Feature	Applications	Advantages	Research challenges	Major requirements
Security	ITS, e-healthcare, smart schools, logistics	Secure attack-free execution environment to deploy services	1) Lack of standardized security solutions without hindering data integrity 2) Secure deployment and integration of cloud-based services at the device and network levels 3) Efficient early identification of both insider and outsider threats	Identification of vulnerabilities in the network that serve as weak entry points for various attacks
Privacy	ITS, e-healthcare	Provides data protection and user privacy in the network	Ensuring users' anonymity in the IoT network for using particular services	Pervasive network model with strong encryption and cryptographic tools
Trust	ITS, e-healthcare	Ensures users' belief and trust that the desired services are free from vulnerabilities	1) Efficient decentralized trust management system 2) Intelligent trust evaluation during service unavailability and compromised IoT network	Decentralized trust model avoiding a single point of failure in the network
Risk management	ITS, indoor e-healthcare	Ensures security by identifying uncertain events and threats in the IoT network	1) Low-cost and efficient risk management systems to identify newly emerged attacks effectively 2) Fast and ultra-efficient risk decision mechanisms to counter identified threats	1) Detailed threat modeling to identify various threats in the network 2) Identify various risks areas through threat actors and asset-based threat modeling
Interoperability	ITS, smart home, personal e-health ecosystem	Provides a platform for two IoT devices from different domains to communicate	Integrating devices for vendor locked-in services	Generic, centralized, flexible, and open reference models for devices to integrate and communicate (e.g., IP, CoAP)
Low-power and low-cost communication	ITS, smart meters, e-healthcare	Provides a wide range of applications in IoT-based smart cities if communication is low-cost	How to prolong the battery life of the IoT devices?	Advancements in micro-electronics and wireless communication to provide low-cost communication and increased battery life
Big data	Smart meters, ITS, e-healthcare	Increases IoT network performance by processing useful information identified by authenticated sources (e.g, analyzing traffic data can reduce traffic congestion processing)	1) Lack of appropriate tools to handle the massively generated information 2) Protection of users' privacy and security 3) Efficient centralized data acquisition and information	1) Centralized big data processing centers 2) Public awareness to utilize resources in the IoT network safely
Connectivity	ITS, waste management, e-healthcare, smart industry	Ensures that IoT devices can communicate from various domains	How to ensure connectivity in a wide range of IoT devices during no communication network and high mobility?	1) Efficient usage of spectrum for IoT devices to communicate 2) Intelligent usage of every possible communication medium (e.g., WLAN, 3G, LTE, WiMAX) 3) Development of gossip-based algorithms to provide connectivity to IoT devices in the absence of a communication network

Table 3. Future research challenges.

Big Data Analytics

Big data analytics is one of the major research directions in the IoT-based smart cities paradigm [15]. Smart cities connecting billions of devices will provide a massive amount of information and data for analysis. This data can include information from surrounding environments (ITS) and user private data (smart hospitals). To analyze this data, intelligent techniques and algorithms are required. For instance, deep learning algorithms can be adopted to efficiently analyze immense information produced by locally connected devices. The major issues that must be addressed are:

- To respect user privacy during data analysis
- To provide data anonymity for sensitive data
- To provide infrastructure to collect, store, and analyze the huge amount of data
- To provide the required computation power to extract new knowledge from the data

Connectivity in IoT

An IoT-based smart city includes billions of devices in the network. The concept of smart cities can succeed only if it has the ability to provide connectivity to every available IoT device with sensing capabilities that produce significant information. In smart cities, IoT devices can use any available communication networks such as public Wi-Fi, Bluetooth, cellular networks (LTE/LTE-Advanced), and satellites to communicate with the cloud-based application center. However, ensuring connectivity in smart cities poses several challenges such as:

- Providing connectivity to devices with high mobility (e.g., high-speed trains and vehicles)
- Connectivity transition from device to network level and vice versa
- Ensuring connectivity to massively deployed devices in the absence of communication networks

Conclusions

This article has presented recent trends and advancements in the IoT-enabled smart cities paradigm. We have devised a taxonomy for IoT-based smart cities based on communication protocols, major service providers, network types, standard bodies, and major service requirements for the understanding of the reader. Based on the conducted study, we have concluded that smart city applications rely on several wireless technologies such as IEEE 802.11p, WAVE, SIGFOX, 6LoWPAN, and LTE/LTE-A. Furthermore, we have studied major open IoT platforms for the ease of researchers. In addition, a number of reported case studies of several of the newest IoT deployments and research projects have been presented to reveal an increasing trend of IoT deployments. In the end, we have unearthed several open research issues such as multi-vendor interoperability, low cost, low power consumption, and security, which demand considerable attention from our research community.

Acknowledgements

We especially thank the late Prof. Dr. Carmelita Görg, former head of ComNets, University of Bremen, Germany, for all her support and guidance. Furthermore, we thank the International Graduate School for Dynamics in Logistics, the doctoral training group of LogDynamics, University of Bremen, for financial support of this work. M. Imran's work is supported by the Deanship of Scientific Research at King Saud University through Research Group No. (RG # 1435-051).

References

[1] A. Zanella *et al.*, "Internet of Things for Smart Cities," *IEEE Internet of Things J.*, vol. 1, no. 1, 2014, pp. 22–32.
[2] J. Gubbi *et al.*, "Internet of Things (IoT): A Vision, Architectural Elements, and Future Directions," *Future Generation Computer Systems*, vol. 29, no. 7, 2013, pp. 1645–60.
[3] A. Gluhak *et al.*, "A Survey on Facilities for Experimental Internet of Things Research," *IEEE Commun. Mag.*, vol. 49, no. 11, Nov. 2011, pp. 58–67.
[4] J. Jin *et al.*, "An Information Framework for Creating a Smart City Through Internet of Things," *IEEE Internet of Things J.*, vol. 1, no. 2, 2014, pp. 112–21.
[5] N. Mitton *et al.*, "Combining Cloud and Sensors in a Smart City Environment," *EURASIP J. Wireless Commun. and Net.*, vol. 2012, no. 1, 2012, pp. 1–10.
[6] I. Ganchev, Z. Ji, and M. O'Droma, "A Generic IoT Architecture for Smart Cities," *25th IET Irish Signals & Systems Conf. 2014 and 2014 China-Ireland Int'l. Conf. Info. and Commun. Technologies*, 2013, pp. 196–99.
[7] R. Ratasuk *et al.*, "Narrowband LTE-M System for M2M Communication," *IEEE VTC-Fall*, 2014, pp. 1–5.
[8] P. Papadimitratos *et al.*, "Vehicular Communication Systems: Enabling Technologies, Applications, and Future Outlook on Intelligent Transportation," *IEEE Commun. Mag.*, vol. 47, no. 11, Nov. 2009, pp. 84–95.
[9] J.-S. Lee, Y.-W. Su, and C.-C. Shen, "A Comparative Study of Wireless Protocols: Bluetooth, Uwb, Zigbee, and Wi-fi," *33rd Annual Conf. IEEE Industrial Electronics Society*, 2007, pp. 46–51.
[10] X. Lin, A. Adhikary, and Y.-P. E. Wang, "Random Access Preamble Design and Detection for 3GPP Narrowband IoT Systems," arXiv preprint arXiv:1605.05384, 2016.
[11] A. Botta *et al.*, "Integration of Cloud Computing and Internet of Things: A Survey," *Future Generation Computer Systems*, vol. 56, 2016, pp. 684–700.
[12] S. Sicaria *et al.*, "Security, Privacy and Trust in Internet of Things: The Road Ahead," *Computer Networks*, vol. 76, 2015, pp. 146–64.
[13] C. Levy-Bencheton and E. Darra, "Cyber Security and Resilience of Intelligent Public Transport," tech. rep., ENISA, Dec. 2015.
[14] World Economic Forum Industrial Internet Survey, "Industrial Internet of Things: Unleashing the Potential of Connected Products and Services," tech. rep., Jan. 2015.
[15] I. Vilajosana *et al.*, "Bootstrapping Smart Cities through a Self-Sustainable Model Based on Big Data Flows," *IEEE Commun. Mag.*, vol. 51, no. 6, June 2013, pp. 128–34.

Biographies

Yasir Mehmood (ym@comnets.uni-bremen.de) received his Master's in electrical (telecommunications) engineering from the Military College of Signals (MCS), National University of Science and Technology (NUST) Islamabad, Pakistan. He is currently a doctoral researcher at the Sustainable Communication Networks (ComNets) research group, University of Bremen, Germany, in the framework of the International Graduate School (IGS) for Dynamic in Logistics (a doctoral training group at the University of Bremen). His major research area includes cellular communications, mobile M2M communications, and the cellular Internet of Things.

Farhan Ahmad (f.ahmad@derby.ac.uk) received his B.Sc. degree in electronics engineering from COMSATS Institute of Information Technology, Abbottabad, Pakistan, and his M.Sc. degree in communication and information technology from the University of Bremen in 2009 and 2014, respectively. He is currently a final year Ph.D. student in computer science and a postgraduate teaching assistant at the College of Engineering and Technology, University of Derby, United Kingdom. His research focuses on cyber security, risk-assessment, and trust in vehicular networks, M2M communications, and the Internet of Things.

Ibrar Yaqoob (ibraryaqoob@siswa.um.edu.my) received his B.S. (Hons.) degree in information technology from the University of the Punjab, Gujranwala campus, Pakistan, in 2012. He has been pursuing his Ph.D. degree in computer science at the University of Malaya, Malaysia, since November 2013. He won a scholarship for his Ph.D. and is also a Bright Spark Program research assistant. He has published a number of research

In smart cities, the IoT devices can use any available communication networks such as public Wi-Fi, Bluetooth, cellular networks (LTE/LTE-Advanced) and satellites to communicate with the cloud-based application center. However, ensuring connectivity in smart cities poses several challenges.

articles in refereed international journals and magazines . His numerous research articles are very famous and most downloaded in top journals. His research interests include big data, mobile cloud, the Internet of Things, cloud computing, and wireless networks.

ASMA ADNANE (a.adnane@derby.ac.uk) joined the University of Derby as a full-time senior lecturer in Networks and Security from the University of Leicester, United Kingdom, where she was a Knowledge Transfer Partnership (KTP) associate with CrowdLab as their database and security expert. She received her Ph.D. in computer science from the University of Rennes, France. She has published several papers in renowned conferences and journals focusing on ad hoc network security and trust management. She also worked as a research associate/lecturer in France at the University of Rennes, University of Nantes, and ENSI-Bourges. Her research interests include trust management in intelligent transport systems, smart cities, and network security.

MUHAMMAD IMRAN (dr.m.imran@ieee.org) has been an assistant professor in the College of Computer and Information Sciences, King Saud University (KSU), since 2011. He worked as a postdoctoral associate on joint research projects between KSU and the University of Sydney, Australia. He is a visiting scientist at Iowa State University. His research interests include mobile ad hoc and sensor networks, WBANs, M2M, IoT, SDN, fault-tolerant computing, and security and privacy. He has published a number of research papers in refereed international conferences and journals. His research is financially supported by several grants. Recently, the European Alliance for Innovation (EAI) appointed him Co-Editor-in-Chief of *EAI Transactions on Pervasive Health and Technology*. He also serves as an Associate Editor for *IEEE Access*, *IEEE Communications Magazine*, the *Wireless Communication and Mobile Computing Journal* (SCIE, Wiley), the *Ad Hoc & Sensor Wireless Networks Journal* (SCIE), *IET Wireless Sensor Systems*, the *International Journal of Autonomous and Adaptive Communication Systems* (Inderscience), and the *International Journal of Information Technology and Electrical Engineering*. He has served/serves as a Guest Editor for *IEEE Communications Magazine*, *Computer Networks* (SCIE, Elsevier), *MDPI Sensors* (SCIE), the *International Journal of Distributed Sensor Networks* (SCIE, Hindawi), the *Journal of Internet Technology* (SCIE), and the *International Journal of Autonomous and Adaptive Communications Systems*. He has been involved in more than 50 conferences and workshops in various capacities such as Chair, Co-Chair, and Technical Program Committee member. These include IEEE ICC, GLOBECOM, AINA, LCN, IWCMC, IFIP WWIC, and BWCCA. He has received a number of awards such as Asia Pacific Advanced Network fellowship.

SGHAIER GUIZANI (sguizani@alfaisal.edu) received his Ph.D. degree from the University of Quebec, Canada, in 2007. He is currently an assistant professor in the Electrical Engineering Department at Alfaisal University, Riyadh, Saudi Arabia. His research interests include communication networks and security (particularly wireless ad hoc, sensor networks, QoS, wireless sensor network security, and RFID/NFC application and security) and the Internet of Things. He has published a number of research papers in refereed international conferences and journals. He has served/is serving as an Associate Editor for *Security and Communication Networks* (Wiley), the *International Journal of Sensor Networks* (Inderscience), and the *Journal of Computer Systems, Networking, and Communications*. He has been involved in a number of conferences and workshops in various capacities.

IEEE COMSOC
TRAINING

Wednesday, 20 September 2017
9:00 am to 4:30 pm EDT

Online via WebEx!

Nanoscale Communication Networks:

How to Use IEEE 1906.1 for Improved Interoperability

In this course, instructor and Chair of IEEE 1906.1-2015, Stephen Bush, will provide a precise description of what a nanoscale communication network is and the minimum requirements to define it.

Highlights covered in this course include:

- A framework for nanoscale communication networks using universal building blocks
- The 20 standard metrics for nanoscale communication networks
- A reference model in the network simulation 3 (ns-3)
- Use-cases of the standard providing practical examples of applications of the standard.

Learn more and register at
http://www.comsoc.org/training

IEEE

IEEE ComSoc™
IEEE Communications Society

IoT in Agriculture: Designing a Europe-Wide Large-Scale Pilot

Christopher Brewster, Ioanna Roussaki, Nikos Kalatzis, Kevin Doolin, and Keith Ellis

The technologies associated with the Internet of Things have great potential for application in the domain of food and agriculture, especially in view of the societal and environmental challenges faced by this sector. From farm to fork, IoT technologies could transform the sector, contributing to food safety, and the reduction of agricultural inputs and food waste.

ABSTRACT

The technologies associated with the Internet of Things have great potential for application in the domain of food and agriculture, especially in view of the societal and environmental challenges faced by this sector. From farm to fork, IoT technologies could transform the sector, contributing to food safety, and the reduction of agricultural inputs and food waste. A major step toward greater uptake of these technologies will be the execution of IoT-based large-scale pilots (LSPs) in the entire supply chain. This article outlines the challenges and constraints that an LSP deployment of IoT in this domain must consider. Sectoral and technological challenges are described in order to identify a set of technological and agrifood requirements. An architecture based on a system of systems approach is briefly presented, the importance of addressing the interoperability challenges faced by this sector is highlighted, and we elaborate on requirements for new business models, security, privacy, and data governance. A description of the technologies and solutions involved in designing pilots for four agrifood domains (dairy, fruit, arable, meat and vegetable supply chain) is eventually provided. In conclusion, it is noted that for IoT to be successful in this domain, a significant change of culture is needed.

INTRODUCTION

The Internet of Things (IoT) provides a unique opportunity for technology to transform many industries [1, 2], including the food and agriculture sector. The agrifood sector has a rather low level of uptake of information and communications technology (ICT) and a relatively high cost of data capture [3]. The stack of technologies in IoT [4] includes sensors, actuators, drones, navigation systems, cloud-based data services, and analytics delivering a variety of decision support tools, and could significantly change this sector. In Europe, large-scale deployments, or pilots (LSPs), of IoT in the context of H2020 [5] will be funded by the European Commission. This article provides an overview of the potential role that IoT can play in the agrifood sector from the perspective of designing and specifying such an IoT-based LSP, thus enabling the readers to understand the associated opportunities, constraints, and requirements of the sector that IoT can address.

In the agrifood sector, IoT technologies appear under labels such as "precision agriculture" or "smart farming." A typical example is the use of GPS to control tractors (auto-guidance of machinery) to ensure precise coverage of a field, whether ploughing, planting, or engaging in some other activity. The gradual instrumentation of all stages of the agrifood sector leads to a wealth of new data-driven services. These can provide the farmer with advice as to where to spray or apply fertilizers, when to inseminate a dairy herd, and when to capture data required by regulatory or certification bodies. Data-driven services could help the logistics and supply chain by enabling optimal route planning, facilitating recalls in food crisis scenarios, or more generally improving stock taking and ordering processes [6]. Supermarket checkout counter data and loyalty cards can integrate well with various IoT eco-system elements. While the majority of these technologies and services exist, they are only deployed in a few instances.

The European Research and Innovation agenda includes deployment of IoT through integration of these technologies across the value chain and their operation on a large scale to respond to real needs of public authorities, citizens, and business [7]. Designing and executing such LSPs will reveal potential shortcomings in the technologies and help promote IoT in agriculture. This article elaborates on the design of an LSP that aims to address several major challenges of the agrifood sector via the exploitation of IoT technologies and their adoption by all stakeholders in the food supply chain.

IOT CHALLENGES AND CONSTRAINTS FOR THE AGRIFOOD SECTOR

The uptake of IoT in agriculture faces considerable challenges, but there are also specific drivers. ICT and corresponding data driven services have penetrated in large-scale industrial farming, especially in North America, and supermarkets in most developed countries. There are many areas, even in highly developed countries, where there has been little or no uptake of IoT. The main issues affecting the uptake of IoT in the agrifood sector in Europe are elaborated on hereafter.

SECTORAL ISSUES

Heterogeneity of the sector: There are a great variety of different types of actors in the food system ranging from very large (supermarkets, seed

Digital Object Identifier:
10.1109/MCOM.2017.1600528

Christopher Brewster is with Nederlandse Organisatie voor Toegepast Natuurwetenschappelijk Onderzoek; Ioanna Roussaki and Nikos Kalatzis are with the National Technical University of Athens; Kevin Doolin is with the Telecommunications Software & Systems Group; Keith Ellis is with INTEL Labs Europe.

and inputs suppliers, commodity traders) to very small (artisanal cheese makers, microbreweries, roadside fruit and vegetable sellers). Consequently, no single solution, whether technological, business model, or regulatory, will fit or accommodate the needs of all. Vineyards in Hungary need quite different solutions from arable farmers in North America. In the EU, for example, precision agriculture practices in arable farming have been widely adopted by large farmers in Central and Northern Europe, in order to increase production and enhance quality. However, in Southern Europe, the latest economic pressure in agriculture, the high farm segmentation and dispersion, as well as the increasing water scarcity requires the exploitation of precision irrigation techniques mainly for minimizing the usage of resources [3].

Farm sizes and capital investment costs: Larger more capital-intensive farms are much more receptive to the uptake of IoT technology, and also are recipients of such technology as part of the continuous investment in new equipment (e.g., tractors and farm equipment). Existing leading smart farming industrial solutions have either been designed for large farms, for example, MyJohnDeere (JohnDeere™), or operate only in limited geographical spaces, for example, FieldView (The Climate Corporation™) and Encirca (DuPont™), which offer services mainly in the United States and Canada. 365FarmNet is adapting the cost and type of offered services to the size of the holding, but its market penetration is still limited to Central Europe. The challenge is making IoT offerings sufficiently attractive to small-scale farmers with limited investment available for new technology and significant fears of data misuse.

Business models and business confidentiality: Appropriate business models are needed with the requisite level of confidentiality and control over data for which farmers are asking, but allowing farms and other agrifood actors to monetize the data they are producing. This is an area of contention, with large players like John Deere seeking to exploit the data captured by the machines they provide, and farmers resisting this as yet another loss of control and loss of value. The American Farm Bureau Federation has been leading a fight for farmers there to retain control and ownership of their data and recently set up the Agricultural Data Coalition.

User and societal acceptance: Education and training aspects are necessary to help end users understand the use and applicability of these new technologies. According to [8], 71 percent of EU farm managers were still operating on the basis of practical experience until recently, believing that they do not need such enhancements for their daily jobs and do not have time to learn. The adoption of smart technologies will undoubtedly be challenging for non-technologically literate persons. However, there are already education and training initiatives running across Europe aiming to disseminate IoT culture among youngsters and all stakeholders in the food chain.

Technological Issues

Lack of interoperability: Common building blocks, data protocols, and standards are needed for billions of devices to interoperate, and there are various appropriate standards in the agrifood domain in an attempt to reach an overall consensus in this area. Such standards exist for semantics and data modeling (e.g., AgroRDF, AgroVOC, agroXML), agri-machinery (e.g., ISOBUS), weather data (e.g., SWEET), for the supply chain (EPCIS from GS1), e-commerce retail stores (e.g., Good Relations and Schema.org), and numerous initiatives; for example, the IEEE Standards Association's IoT Related Standards; the International Telecommunication Union's (ITU's) IoT Global Standards Initiative; Onem2m, Open Interconnect Consortium; the AllSeen Alliance; and the IPSO Alliance). However, standards such as ISOBUS have not adapted to the pace of change, and most new machinery has proprietary connectivity with machinery of the same manufacturer. This leads to "vendor lock-in" and further resistance from farmers. Major initiatives are underway from both the Agricultural Electronics Foundation and AgGateway to overcome interoperability barriers. The challenge here is not the lack of standards, but the emergence of too many standards.

Lack of connectivity: A key challenge in many locations for the further development of IoT in agriculture is the lack of connectivity, that is, poor third/fourth generation (3G/4G) coverage (in spite of the much trumpeted wish to move to 5G). Low power wide area (LPWA) technologies like LoRa and SIGFOX provide a real opportunity to overcome such limitations [9], but they do not handle large datasets (e.g., originating from satellite imagery).

Data processing power: This may appear surprising, but the ability to access large-scale processing power at reasonable cost to solve complex calculations (e.g., traveling salesman type planning of field traversal) remains a challenge for small to medium farmers. The absence of data processing services significantly hinders IoT.

Lack of clear data governance: Regulations and legal frameworks are only slowly catching up with the current technological realities. Control and ownership of farm data is still contentious (as noted above). Large companies may want to conceive of themselves as "data companies" and fight initiatives to leave control of data in the hands of farmers and other primary actors.

Data security and privacy: Distinct from governance issues are the issues of security and data privacy. In a European Commission/International Data Corporation (EC/IDC) analysis of the EU demand for cloud computing services and barriers to uptake, the top five concerns among respondents directly or indirectly relate to security or privacy. This is indicative of the wider importance of such matters for IoT adoption in smartagri [7].

Despite these issues, there is a growing community of either technically literate young farmers, or hi-tech professionals and tech-enthusiasts with a strong interest in the agrifood domain. This has led to a proliferation of startups, hackathons, and many different initiatives which are gradually making the application of data science, sensors, and technology in general to agrifood an attractive and exciting prospect. The U.S.-based Food + tech Connect website is an excellent source of examples, but there are equally many such initiatives in Europe and East Asia.

Appropriate business models are needed with the requisite level of confidentiality and control over data that farmers are asking for, but allowing farms and other agrifood actors to monetise the data they are producing. This is an area of contention with large players like John Deere.

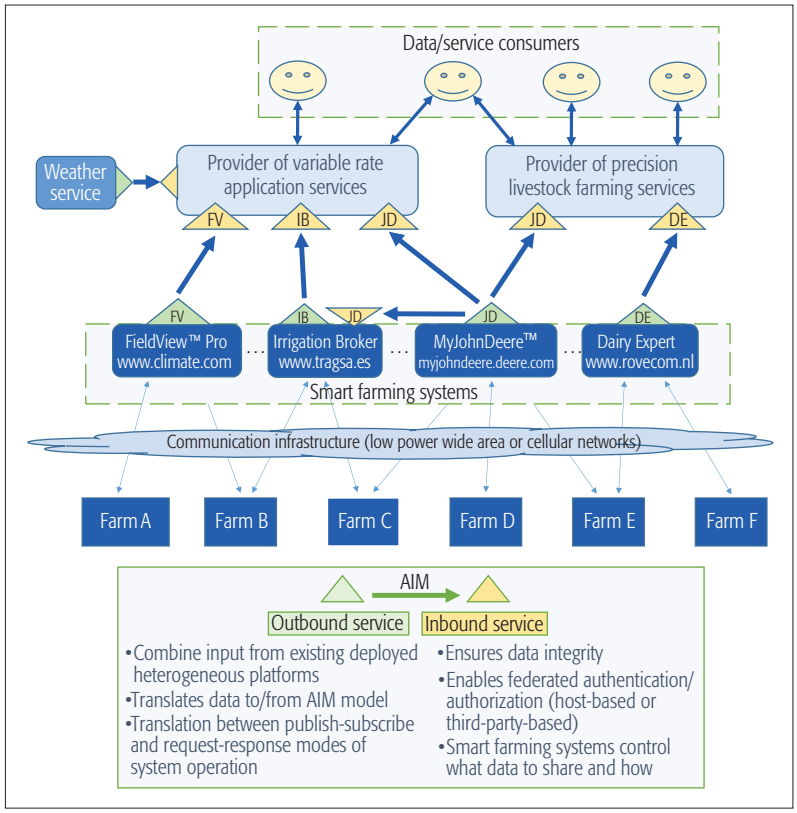

Figure 1. A system-of-systems architectural approach to an IoT-based large-scale pilot in agriculture.

In addition to tacking the sectoral and technological issues above, the LSPs need to address several objectives specific to the agrifood sector in order to convince users of the usefulness of IoT technologies in agriculture. These include the following:

- Lower production costs and increase yield quality/quantity.
- Increase productivity and animal health/welfare.
- Enable monitoring and control of plant/animal products during the entire life cycle for food traceability and increased food awareness for consumers.
- Reduce use of water and other natural resources, and improve soil quality.
- Minimize the ecological footprint and environmental impact of agricultural practice, and adapt crop management to requirements of climate change.
- Facilitate and enhance food safety/security.
- Ensure that certification schemes (e.g., organic) are effective and fraud-free across the entire food supply chain.
- Develop or enable business models adapted for an IoT ecosystem, and create new business and cooperation opportunities.

Overall Architecture

There is a variety of "precision agriculture" systems and platforms already deployed, employing many different communication, sensing, and data processing technologies. Analysis of the challenges indicate that the approach of building a new master system to incorporate others may not be feasible for an LSP due to potential scalability (e.g., maintaining state in a pub/sub approach) and governance (e.g., access to agricultural data) issues. Therefore, a system-of-systems (SoS) approach is proposed. This enables existing agriculture knowledge information systems (AKISs) to continue operating, but also allows those systems to make available and consume data from cooperating systems within the SoS. Additionally, the SoS can expose newer technologies and services that may be of interest to cooperating AKISs. This is more realistic and viable in terms of usability, market adoption, and sustainability. In order to realize this approach, the following two core functional requirements need to be fulfilled by the proposed solution:

1. Allow existing AKISs to offer their data to and consume data from their counterparts.
2. Extensive use of virtualization containers for services should be made to ensure rapid deployment once required

The proposed architecture (Fig. 1) consists of an inbound and an outbound service that support AKISs to expose and consume data, respectively. Rapid deployment is highly beneficial for survey services that might not require a continuous feed from a particular AKIS. Such a service would deploy and start an inbound service for that particular AKIS, gather necessary information, and then stop the service. The service will be packaged into a lightweight container along with all the software necessary to support self-contained deployment of the service (runtime environment, libraries for supported communication protocols, encryption techniques, etc.).

As data interoperability is of critical importance, the proposed solution provides the necessary data translation mechanisms combining the use of a semantic data model (Agriculture Information Model — AIM) along with the respective data translation/management/inference mechanisms adopting OMA Next Generation Sensors Initiative (NGSI) functional network application programming interfaces (APIs). In order to enable interoperability of heterogeneous data handling approaches, the inbound-outbound services, deployed on various AKISs, translate and exchange data based on the AIM common data format with the utilization of lightweight data wrappers/translators. For this conversion to be feasible, each AKIS needs to provide the specifications of the utilized data model-semantics, or it should parse returning content in AIM format. The AIM is not built *ab initio*, but incorporates and extends existing ontologies and vocabularies already available for this domain (e.g., agroRDF, GACS, EPCIS).

Inbound-outbound services maintain the necessary mechanisms for satisfying data security and privacy concerns (cf. below). They need to be trusted to be deployed and hosted by the AKIS on their own cyber-premises (i.e., hosting environments). This is an inherent data privacy protection feature as the owner of the data maintains the control/decision of which data are allowed to be shared with other entities. The services need to provide privacy and security functionalities, including user authentication and access authorization.

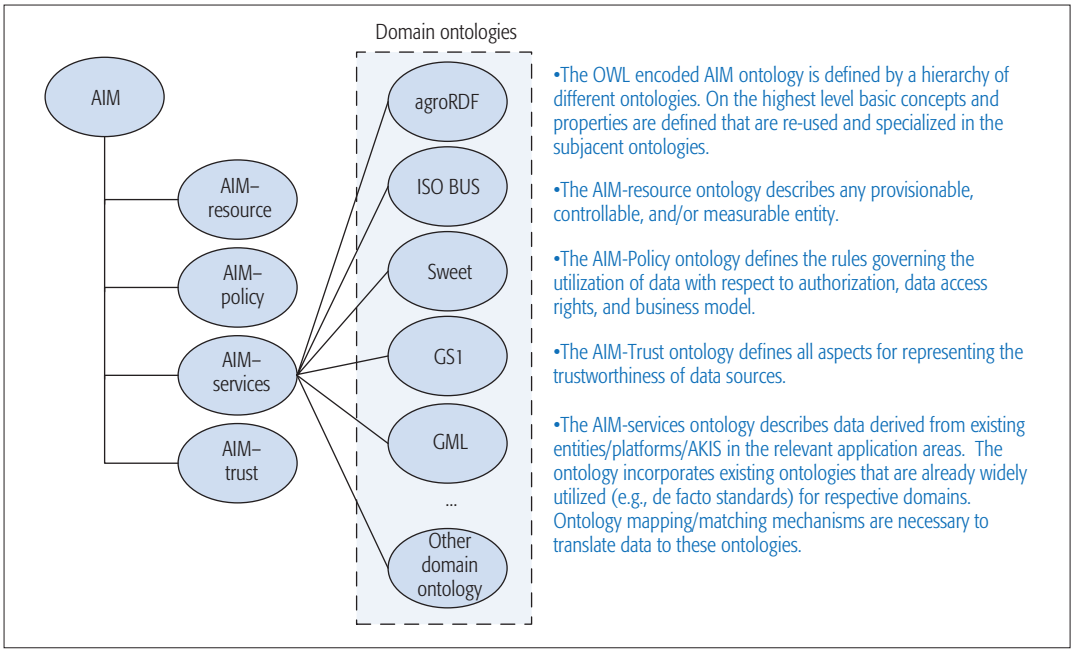

Figure 2. The Agricultural Information Model (AIM) structure.

As the dominating stakeholders are now creating powerful positions for themselves, there may be no incentive to enable radical, transformative changes and that recent developments will not truly make use of the affordances offered by smart technologies, or explore how to reconnect parties in the supply chain in entirely new ways.

Addressing Interoperability Aspects

Interoperability is a core issue for IoT in the agricultural domain [10]. For a viable ecosystem to develop, a hardware provider from country A should be able to offer services to farmers in country B using software from country C. Probably the most coordinated domain-specific effort in this direction is AgGateway, an association of agriculture and agtech companies started in the United States but expanding globally. Technical interoperability is not discussed below, as the protocols and infrastructure required are not specific to the agrifood sector, but are rather addressed by domain agnostic groups like the Open Channel Foundation (OCF) with initiatives such as the Alliance for the Internet of Things Innovation (AIOTI), set up to support dialogue among the various actors in Europe.

Syntactic Interoperability

In spite of the widespread adoption of XML and more recently JSON as standard syntaxes for data sharing, because of the wide variety of legacy systems it is often expedient to export in simple formats such as comma separated variables (CSV). Electronic data interchange (EDI), specifically EDIFACT, despite its expense and complexity, still plays a significant role in some agrifood sectors, mostly for invoices and similar types of messaging. Thus, a key requirement here is that all systems provide export facilities or API access that return standard formats, typically XML or JSON, and where possible legacy systems are provided with appropriate interchange gateways.

Semantic Interoperability

There are two types of semantic standards for data sharing relevant here. There are standards for on-farm operations (e.g., ISOBUS, AgroXML) and standards for data exchange across the supply chain (e.g., GS1 EPCIS, EDIFACT). For precision farming and IoT in the supply chain, there is no shortage of standards, but rather a lack of universal uptake. One of the most successful examples is the creation of unique identifiers for bovine animals in the EU, which has been gradually extended to other dairy animals. The United Nations and Global Standards One (GS1) have an initiative to offer global location numbers (GLNs) to small-scale farmers around the world, so-called blue numbers, for participation in supply chains. In the supply chain, the GS1 EPCIS standard (for barcodes and RFID) has established itself as the dominant standard, but while the core standard is agreed, slow progress has been made to incorporate additional information concerning production methods or other characteristics. There are many different standards in existence for the description of agricultural products including AGROVOC from the Food and Agriculture Organization (FAO), CBV from GS1, the NALT thesaurus, or the CAB Thesaurus, among others. While the GACS initiative is a step in the right direction, this largely is an outcome of "research data" rather than precision farming or supply chain requirements.

The proposed approach is to use semantic technology building on existing standards, extending where appropriate, and ensuring appropriate mappings so as to produce an integrated AIM (Fig. 2), as mentioned in the previous section. This fits with the ambitions of AgGateway in their SPADE and ADAPT projects for integration of data from agricultural machinery for farm management information systems. In Fig. 3, a proposal for SoS-based integrated information management from farm to fork is shown that supports syntactic and semantic interoperability, handling data collected from different sensors/devices/platforms and modeled with various local ontologies/schemas. The system considers the data governance/ownership/privacy policies, and if the necessary rights are in place, it supports agri-data fusion and exchange of data across formats.

Most IoT architectures assume all data is written to one blackboard, and all services have access to all data. This is not realistic from a business perspective, as most actors will refuse to participate. Architectures are needed that ensure each farmer controls the data from their own farm and can determine who has access and for which service.

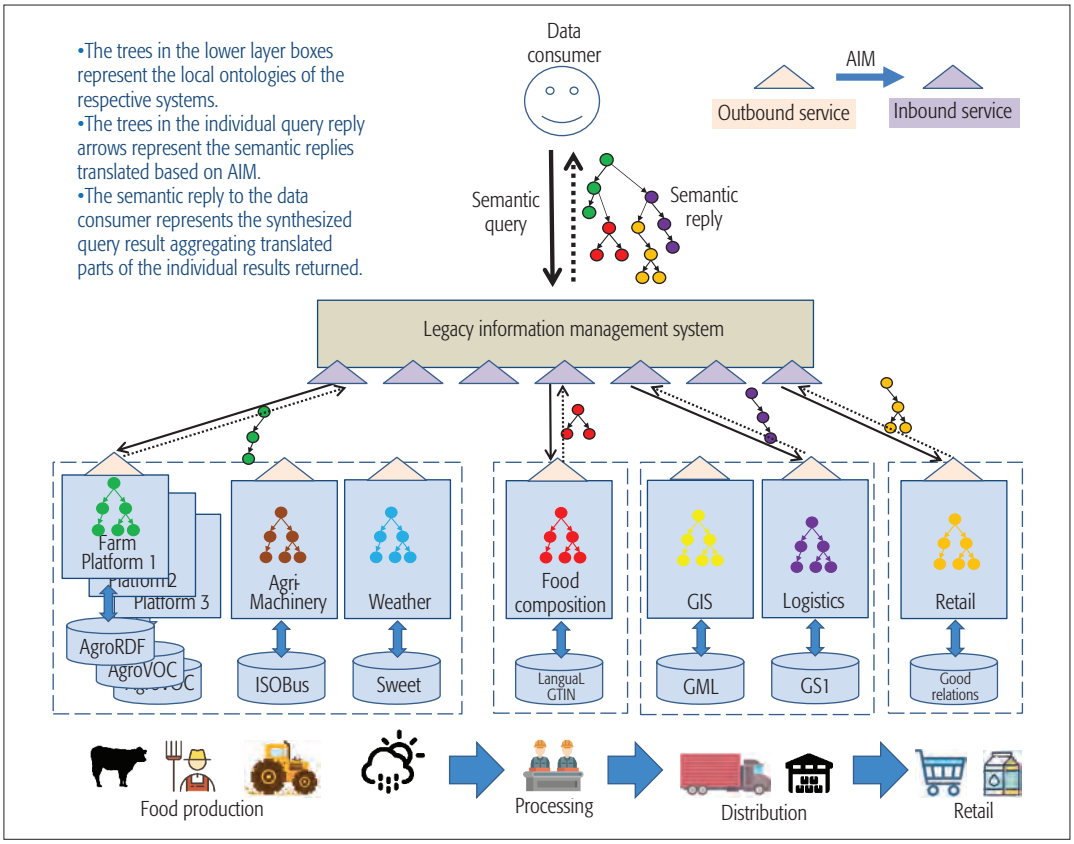

Figure 3. A farm-to-fork management information system ensuring data interoperability.

Other Critical Aspects for Pilot Specification

New Business Models

The food chain is quickly turning into a food-and-data chain [11]. Although there have been a great many data-driven startups in the agrifood domain in the last few years, the sector is being reshaped by large corporations. DuPont, Monsanto, and John Deere have acquired various data focused startups. As the dominating stakeholders are now creating powerful positions for themselves, there may be no incentive to enable radical, transformative changes, and recent developments will not truly make use of the affordances offered by smart technologies, or explore how to reconnect parties in the supply chain in entirely new ways. The future of farming involves engaged farmers becoming active prosumers of agri-data, rather than passive consumers of data analyzed by other parties. Thus, a new connected collaborative agriculture" (CCA) farming model is needed, focused on the farmer as the primary producer of agricultural data. We envision the emergence of a data cooperative where all data collected along the supply chain will be made available via a community platform, mirroring the open data movement. This is the basis for enabling the creation of new business models that will be of financial, environmental, and social significance for all entities in the food production value chain [8]. One way is by connecting to other industries such as pharmaceutical companies, insurance companies, and consumer groups.

The core business model will be built around the components applicable for IoT general business models regarding digital products [12]. These components are: physical freemium, digital add-on, digital lock-in, product as point of sales, object self-service, remote usage, and condition monitoring, while the new "sensor as a service" business model pattern emerges. This core business model will be adapted to the various use cases and will be vertically complemented by data brokerage, as well as linked data components.

Security and Privacy

In spite of the emergence of different cross-world initiatives in recent years — the International Electronics Recycling Congress (IERC), ITU Telecommunication Standardization Sector (ITU-T) Study Group 20 (SG20), the IEEE IoT Initiative, and the Internet Protocol Security Option (IPSO) Alliance are just some of them — there is a lack of a unified vision on security and privacy considerations in the IoT paradigm. Requirements for the agrifood domain include data confidentiality and authentication, access control within the IoT network, privacy and trust among users and things, and the enforcement of security and privacy policies [13]. To be able to carry out IoT-based LSPs in the agrifood domain, an efficient lightweight authentication, authorization, and access control solution for smart agriculture needs to be employed, building on best practices [14], but controlling data at the data model level. The success of IoT services could be threatened if privacy by design or data minimization principles are not supported.

Data Governance and Ownership

The proposed approach is to ensure that the architecture facilitates complete control of data by the primary data generator (i.e., the farmer,

the transportation company, the aggregator, etc.). This would provide the participant in an IoT ecosystem with a sense of control, and thus also both conform to the privacy requirements of the European Data Protection Regulation (EDPR) and enable farmers to treat data as a potential source of income. As farmers have generally lost out in recent years due to developments in the agrifood supply chain, control of data and the ability to see this as an income stream are important incentives for participation in the IoT ecosystem.

Most IoT architectures assume all data is written to one blackboard, and all services have access to all data. This is not realistic from a business perspective, as most actors will refuse to participate. Architectures are needed to ensure that each farmer controls the data from their own farm and can determine who has access and for which service.

Large-Scale Pilots

An LSP aims to evaluate the usability and usefulness of IoT technologies in agriculture, and four pilot domains are proposed here. We describe the main focus, the respective technical challenges addressed, the IoT technologies used, as well as the agrifood applications provided. The last subsection provides a set of metrics for the evaluation of the proposed pilots in a quantifiable manner. In Fig. 4, various representative applications of the four selected pilot domains are illustrated.

Dairy Pilot

Main focus: Demonstrating end-to-end *integration of heterogeneous data sources* throughout the value chain and *advanced decision making* for farm operations.

IoT HW/SW include: Devices/sensors for *position/location and activity monitoring of individual animals* (e.g., GPS trackers, proximity tags, neck/leg transmitters with accelerometer, rumen pH sensing) and environmental sensors (temperature, humidity, sound, gas sensors). Communication technologies include *IoT communication protocols*, such as Constrained Application Protocol (CoAP), message queueing telemetry transport (MQTT), and Advanced Message Queuing Protocol (AMQP); AND LPWA (LoRA, SigFox) or cellular networks (EDGE, HSPA, LTE, 5G). Platforms/systems and components include IoT-enabling (e.g., FIWARE Generic Enablers, components from FInish, FRACTALS, IoT6, iCore) and agri domain specific platforms (e.g., FIspace, 365FarmNet, AgroIT, Ermes, Agrocycle, Rovecom Dairy Expert). Analytics can provide advice to dairy farmers regarding animal/herd behavior (e.g., collar-based analytics, disease and pregnancy detection), and dairy quality (e.g., milk composition, quality, and quantity).

Fruit Pilot

Main focus: Demonstrating *interoperability of IoT systems* and support for advanced learning/reasoning/prediction over several farm-related elements.

IoT HW/SW include: devices/sensors for *air/soil/water monitoring* (e.g., evapotranspiration, water content, stem water potential, stem psychrometer sensors, CO_2 gas, IR and VIS absorbance and temperature, water nitrates and conductivity, humidity, radiation, nutrient levels, pH, cation-exchange capacity, and salinity), sun intensity, wind (direction/speed), and meteorological stations networks. *Agri-machinery monitoring* (e.g., sensors monitoring via the ISOBUS, temperature, pressure, electronic noses, product barcodes readers, RFID sensors, irrigation systems controlling solenoid valves and hydrometers, fertigation actuators, autonomous solar units). *Plant/fruit monitoring* (e.g., berry growth rate; sap-flow meter, dendrometer, drone and satellite imagery data for, e.g., fruit maturation, vegetation decay, quality forecast), crop and post-harvest monitoring (irrigation, pest, and quality alarms); biophysical (LAI, fPAR, etc.), and biochemical indicators (chlorophyll estimations allow recommendation of corrective actions, e.g., fertilization, directing and commanding variable rate production technologies). Communication technologies are similar to those in the Dairy pilot. Platforms/systems and components are also similar to those in the Dairy pilot. Forecast models are: disease/pest prevention, soil condition and quality analysis/estimation, crop management, and harvest period estimation. Also included are recommendations and control of variable rate production technologies based on indicators and crops current status, irrigation related recommendations, and machinery control, as well as vehicle tracking systems (e.g., V-Track AutoID middleware).

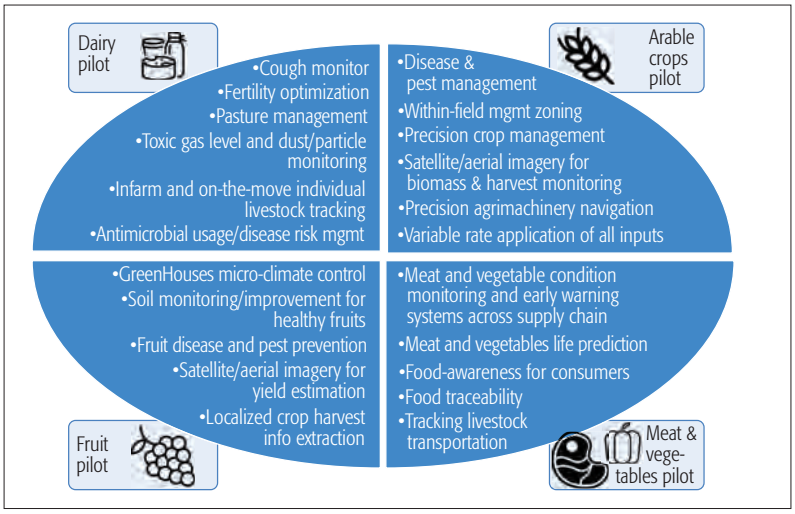

Figure 4. Applications in the various pilot domains.

Arable Crops Pilot

Main focus: Demonstrating *integration of fixed IoT systems with machinery IoT infrastructure, collective decision making, and cross-country interoperability.*

IoT HW/SW used: devices/sensors include *air/soil/water monitoring* (similar to Fruit), including conventional micro-meteorological stations; soil sensors, and proximal and remote crop sensors (NDVI, near-infra-red spectroscopy, hyperspectral images); sensors for water supply, soil water; leaf wetness and nutrients sensors; delivery points, hydrant remote controls, pumping stations, raft sensors, and cameras; and irrigation systems with hydrometers). *Agri-machinery monitoring* (e.g., RTK-DGPS for precise vehicle guid-

Domain	Key performance indicators
Business & sustainability	• Jobs# created by the pilot involved parties • Innovative business models# identified by the pilots • Companies# implementing new industrial/business processes after the pilot execution • Farmers/producers# profiting by granting access to the data collected on their farms after the pilot execution • Partnership projects# created by pilot involved parties based on learning from pilot activities • Local/rural businesses# involved in pilots with cross-border companies • Local businesses# willing to expand their business beyond borders • Pilot-involved parties# willing to continue exploiting the pilot deployment • Pilot sites# that will continue using the piloted deployments • Amount of additional investments committed pilot end-users during the project oriented to reinforce their IoT-based capabilities • Agri-Industry adoption of pilot approach over 5 years
End-user side	• Adherence to Privacy-by-design across pilot components and applications • Achievement of agreed, credible security format • Rating of user acceptance per pilot through qualitative and quantitative means • Qualitative analysis of user perception of privacy and security, vulnerability issues throughout the project.
Standardization	• Standards# used across pilot deployments • Companies# adopting open platforms / standards within the pilots • Partners# contributions to (pre)-standardization activities • Standardisation committees# contributed to • Waste reduction, traceability best practice adoption • EPCIS standards# uptake
Scale and more	• Pilot sites# • Services/applications# deployed at pilot sites • Applications# deployed on top of open platforms • Demand side actors# participating in the requirement definition phase • Individual users# of applications deployed on top of open platforms • Proprietary solutions integrated against open platforms/solutions across pilots • Inclusive and with participation levels recorded • Datasets# aggregated by the pilots

Table 1. Key performance indicators across the domains of the large scale pilots.

ance, onboard sensors, VIS-NIR spectrometer), historic satellite soil and crop variability data, sound-based pest detection, pheromone traps). Communication technologies are similar to those in the Dairy pilot. Platforms/systems and components are similar to those in the Dairy pilot. Crop factor and NDVI-based irrigation systems (both need formula and equations specifically adapted to farms, crops, and parcels using remote sensing, for example, MegaBroker (Tragsa, bb-smartworx). IoT technologies on farm machinery, farm management system (FMS) (e.g., MyJohnDeere by John Deere). Decision support systems to improve crops yield forecasting and to optimize the harvesting strategy. Yield forecasting systems (soil, crop cover, and grain quality data acquired with spectrometers, on-line grain selection system, other sensors and meteorological stations) (e.g., 365FarmNet™, OnFarm™, MyCrop, Yield Prophet).

Meat and Vegetables Pilot

Main focus: Demonstrating *interlinking of IoT systems of various stakeholders* across the entire food supply chain, *element monitoring and tracing* across all supply chain phases, and *security and privacy* of information collected.

IoT HW/SW used: For Meat, very similar to the Dairy Pilot, for Vegetables very similar to Fruit. Additional devices/sensors include RFID readers (FEIG/DTE), RFID tags (HID Global); slaughterhouse recordings (monitoring information regarding the kill time/origin of animal and safety compliance), sensors for recording the movement of vehicles. Environmental sensors for farms and trucks (temperature, humidity, luminosity, CO_2, noise, and ammonia), weight/load cells, cooling actuators (transport and shop shelve). Communication technologies similar to the Dairy pilot, plus vehicle-to-infrastructure, vehicle-to-vehicle, IoT/machine-to-machine communication protocols (technologies used for V2X communications include DSRC, cellular, RFID, IEEE 802.11). Platforms/systems and components similar to the Dairy pilot, plus livestock logistics, decision support systems for complying with legal regulations of animal welfare (e.g., Connecterra, QIfresh, VirtualVet), and fusion systems (i.e., fusion of sensor information for measurement of several parameters that are later integrated into products, e.g., farmsoft). Track & trace platforms for meat from farm to fork (e.g., fTrace and other EPCIS-based systems). Vehicle tracking systems (as above).

Pilot Performance Evaluation

The pilots proposed above need to be measured against well defined, quantifiable key performance indicators. Table 1 introduces an indicative set of such indicators.

Conclusions

This article aims to guide industry stakeholders and researchers who have undertaken the task to build large-scale pilots in agriculture that are heavily based on IoT technologies. The IoT-related challenges and constraints for the agrifood sector are described together with the core objectives of IoT-based LSPs. A system-of-systems architectural approach is proposed, with an emphasis on the interoperability aspects which are critical for the uptake of IoT technologies in the agrifood sector. The Agricultural Information Model approach is proposed to address semantic interoperability, and a farm-to-fork management information system solution ensuring data interoperability is outlined. There remain many challenges including the need for new business models, security and privacy, and data governance and ownership solutions, as they are critical for executing IoT-based LSPs in agrifood. Finally, a detailed account is presented of the most appropriate IoT technologies and agrifood applications to be used, as well as the main key performance indicators to which one can refer to evaluate the performance of the proposed LSPs in a quantifiable manner. The execution of such LSPs will undoubtedly promote the usage of IoT in agriculture, thus optimizing various operations in the entire food supply chain resulting in reduced effort and costs for the producers, and higher food quality and safety, as well as extended food awareness for the consumer. Nonetheless, the main barrier that needs to be overcome before IoT is extensively exploited by the stakeholders across the food supply chain is the change of culture that is needed to appreciate the advantages and opportunities provided by IoT technologies.

References

[1] L.D. Xu *et al.*, "Internet of Things in Industries: A Survey," *IEEE Trans. Industrial Informatics*, vol. 10, no. 4, 2014, pp. 2233–43.

[2] O. Vermesan and P. Friess, *Building the Hyperconnected Society: IoT Research and Innovation Value Chains, Ecosystems and Markets*, River Publishers, 2015.

[3] EIP-AGRI Focus Group, "Mainstreaming Precision Farming," Nov. 2015.

[4] A. Al-Fuqaha *et al.*, "Internet of Things: A Survey on Enabling Technologies, Protocols, and Applications," *IEEE Commun. Surveys & Tutorials*, vol. 17, no. 4, 2015, pp. 2347–76.

[5] EC, "Call — Internet of Things, in HORIZON 2020-Work Programme 2016-2017 Cross-Cutting Activities (Focus Areas)"; http://ec.europa.eu/research/participants/data/ref/h2020/wp/2016_2017/main/h2020-wp1617-focus_en.pdf, July 2016, accessed Nov. 25, 2016.

[6] C.N. Verdouw *et al.*,"Virtualization of Food Supply Chains with the Internet of Things," *J. Food Eng.*, vol. 176, 2015, pp. 128–36.

[7] Joint Research Centre of the EC, "Precision Agriculture: An Opportunity for EU Farmers — Potential Support with the CAP2014-2020"; http://www.europarl.europa.eu/RegData/etudes/note/join/2014/529049/IPOL-AGRI_NT(2014)529049_EN.pdf, June 2014, accessed Nov. 25, 2016.

[8] AIOTI WG06 — Smart Farming and Food Safety, "Smart Farming and Food Safety Internet of Things Applications — Challenges for Large Scale Implementations"; http://www.aioti.org/wp-content/uploads/2016/10/AIOTIWG06Report2015.pdf, Nov. 2015, accessed Nov. 25, 2016.

[9] L. Vangelista *et al.*, "Long-Range IoT Technologies: The Dawn of LoRaTM," *Proc. 1st Int'l. Conf. Future Access Enablers for Ubiquitous and Intelligent Infrastructures*, Ohrid, Republic of Macedonia, Sept. 2015.

[10] European Research Cluster on the Internet of Things, IERC, "Internet of Things — IoT Semantic Interoperability: Research Challenges, Best Practices, Recommendations and Next Steps"; http://www.internet-of-things-research.eu/pdf/IERC_Position_Paper_IoT_Semantic_Interoperability_Final.pdf, Mar. 2015, accessed Nov. 25, 2016.

[11] C. Brewster *et al.*, "Identifying the ICT Challenges of the Agri-Food Sector to Define the Architectural Requirements for a Future Internet Core Platform," *Proc. 2012 eChallenges Conf.*, Lisbon, Portugal, Oct. 2012.

[12] E. Fleisch *et al.*, "Business Models and the Internet of Things," *Proc. Interoperability and Open-Source Solutions for the Internet of Things Int'l. Wksp.*, Split, Croatia, Sept. 2014.

[13] J.H. Ziegeldorf *et al.*, "Privacy in the Internet of Things: Threats and Challenges," *Security and Commun. Networks*, vol. 7, no. 12, 2014, pp. 2728–42.

[14] S. Sicari *et al.*, "Security, Privacy and Trust in Internet of Things: The Road Ahead," *Computer Networks*, vol. 76, 2015, pp. 146–64.

Biographies

Christopher Brewster [M] is a senior scientist at TNO, and until recently a senior lecturer in Information Technology at Aston Business School, Aston University. He received a Ph.D. in computer science from the University of Sheffield, specializing in NLP and semantic technologies. His main research interests lie in the application of ICT to the food and agriculture system, including the use of semantic technologies, the Internet of Things, blockchains, and social implications of technology.

Ioanna Roussaki [M] received her Diploma in electrical and computer engineering in 1997 from the National Technical University of Athens (NTUA), Greece. In 2003, she received her Ph.D. in the area of telecommunications and computer networks. She has participated in many national and international research and development projects. Since 2015, she is an assistant Professor in the NTUA School of Electrical and Computer Engineering. Her research interests include the Internet of Things, context awareness, social-computing, and more.

Nikos Kalatzis received his Diploma in physics in 2000 from the University of Ioannina, Greece, and an M.Sc. degree in information security from the University of London, United Kingdom, in 2002. Since 2005, he has been a research associate at the School of Electrical and Computer Engineering, NTUA. His research interests include the Internet of Things, collaborative inference/learning algorithms, information security, user behavior modeling, and social media. He has participated in several international research projects.

Kevin Doolin is the director of EU Programmes in the Telecommunications Software & Systems Group (TSSG) in Ireland. Before joining TSSG, he worked as a project executive with Ireland's Investment and Development Agency in Waterford for 18 months. Prior to this, he worked for eight years with Ericsson Systems Expertise ending as a strategic product manager. He has worked in numerous European research projects, some of which he personally coordinated (e.g., SOCIETIES).

Keith A. Ellis [M] is a senior research scientist in the Internet of Things Systems Research Lab, Intel Labs. He investigates the feasibility, performance, and adoption of interoperable networks and intelligent edge platforms within various domains. His research interests also include embedded systems, data manageability, and decentralized systems. He holds an M.Sc. in innovation and technology management and a B.Sc. in technology, and has 19 years of industrial experience in manufacturing, ICT systems engineering, and research.

Understanding the Limits of LoRaWAN

Ferran Adelantado, Xavier Vilajosana, Pere Tuset-Peiro, Borja Martinez, Joan Melià-Seguí, and Thomas Watteyne

Low-power wide area networking technology offers long-range communication, which enables new types of services. Several solutions exist; LoRaWAN is arguably the most adopted. The authors provide an impartial and fair overview of the capabilities and limitations of LoRaWAN. They discuss those in the context of use cases, and list open research and development questions.

ABSTRACT

Low-power wide area networking technology offers long-range communication, which enables new types of services. Several solutions exist; LoRaWAN is arguably the most adopted. It promises ubiquitous connectivity in outdoor IoT applications, while keeping network structures and management simple. This technology has received a lot of attention in recent months from network operators and solution providers. However, the technology has limitations that need to be clearly understood to avoid inflated expectations and disillusionment. This article provides an impartial and fair overview of the capabilities and limitations of LoRaWAN. We discuss those in the context of use cases, and list open research and development questions.

INTRODUCTION

Network operators are starting to deploy horizontal machine-to-machine (M2M) solutions to cover a wide set of large-scale verticals, using low power wide area networking (LPWAN) technologies [1, 2]. Application domains include smart city, metering, on-street lighting control, and precision agriculture. LPWAN technologies combine low data rate and robust modulation to achieve multi-kilometer communication range. This enables simple star network topologies that simplify network deployment and maintenance [3]. While the benefits of these technologies are known and are often considered as the key enablers for some applications, their limitations are still not well understood [4, 5].

In this article we aim to provide an impartial overview of the limitations of long-range wide area networking (LoRaWAN) [6], one of the most successful technologies in the LPWAN space. LoRaWAN is a network stack rooted in the LoRa physical layer. LoRaWAN features a raw maximum data rate of 27 kb/s (50 kb/s when using frequency shift keying, FSK, instead of LoRa), and claims that a single gateway can collect data from thousands of nodes deployed kilometers away. These capabilities have really resonated with some solution providers and network operators, who have created a large momentum behind LoRaWAN to the point that it is sometimes touted as the connectivity enabler for any Internet of Things (IoT) use case [7].

The goal of this article is to bring some sanity to these statements by providing a comprehensive, fair, and independent analysis of what the capabilities and limitations of LoRaWAN are. We adopt a pragmatic approach, and identify in which use cases the technology works, and in which use cases it does not work. We provide an overview of LPWAN technologies, including cellular. We describe LoRaWAN technology in detail. We analyze the network capacity and scale limitations of the technology. We discuss the use cases where LoRaWAN works/does not work. We list open research and development challenges for the technology. We then conclude the article.

OVERVIEW OF LPWAN AND CELLULAR TECHNOLOGIES FOR IOT

LOW-POWER WIDE AREA ALTERNATIVES

Although LoRaWAN is one of the most adopted technologies for IoT, there is a wide range of LPWAN technologies in the market, such as Ingenu, Weightless W, N, and P, and SigFox [8].

Ingenu developed a proprietary LPWAN technology in the 2.4 GHz band, based on random phase multiple access (RPMA) to provide M2M industry solutions and private networks. The main asset of Ingenu in comparison to alternative solutions is high data rate up to 624 kb/s in the uplink and 156 kb/s in the downlink. On the contrary, the energy consumption is higher and the range is shorter (a range around 5–6 km) due to the high spectrum band used.

The Weightless Special Interest Group developed a set of three open standards for LPWAN: Weightless-W, Weightless-N, and Weightless-P. Weightless-W was developed as a bidirectional (uplink/downlink) solution to operate in TV white spaces (470–790 MHz). It is based on narrowband frequency-division multiple access (FDMA) channels with time-division duplex between uplink and downlink; data rate ranges from 1 kb/s to 1 Mb/s, and battery lifetime is around 3–5 years. Weightless-N was designed to expand the range of Weightless-W and reduce the power consumption (a battery lifetime up to 10 years) at the expense of data rate decrease (from up to 1 Mb/s in Weightless-W to 100 kb/s in Weightless-N). Unlike Weightless-W, Weightless-N is based on the ultra narrowband (UNB) technology and operates in the UHF 800–900 MHz band; it provides only uplink communication. Finally, Weightless-P is proposed as a high-performance two-way communication solution that can operate over 169, 433, 470, 780, 868, 915, and 923 MHz bands. However, cost of the terminals and power consumption are higher than in Weightless-N, with a battery lifetime of 3–8 years.

Digital Object Identifier:
10.1109/MCOM.2017.1600613

Ferran Adelantado, Pere Tuset-Peiro, Borja Martinez, and Joan Melià-Seguí are with IN3 at the Universitat Oberta de Catalunya; Xavier Vilajosana is with IN3 at Universitat Oberta de Catalunya and Worldsensing; Thomas Watteyne is with Inria.

Together with LoRaWAN, SigFox is one of the most adopted LPWAN solutions. It is a proprietary UNB solution that operates in the 869 MHz (Europe) and 915 MHz (North America) bands. Its signal is extremely narrowband (100 Hz bandwidth). It is based on random frequency and time-division multiple access (RFTDMA) and achieves a data rate around 100 b/s in the uplink, with a maximum packet payload of 12 bytes, and a number of packets per device that cannot exceed 14 packets/day. These tough restrictions, together with a business model where SigFox owns the network, have somewhat shifted the interest to LoRaWAN, which is considered more flexible and open.

Cellular Solutions for IoT

The Third Generation Partnership Project (3GPP) standardized a set of low-cost and low-complexity devices targeting machine-type communications (MTC) in Release 13. In particular, 3GPP addresses the IoT market from a three-fold approach by standardizing the enhanced MTC (eMTC), narrowband IoT (NB-IoT), and extended coverage GSM for IoT (EC-GSM-IoT) [9].

eMTC is an evolution of the work developed in Release 12 that can reach up to 1 Mb/s in the uplink and downlink, and operates in LTE bands with a 1.08 MHz bandwidth. NB-IoT is an alternative that, thanks to reduced complexity, has a lower cost at the expense of decreasing data rate (up to 250 kb/s in both directions). Finally, EC-GSM-IoT is an evolution of evolved general packet radio service (EGPRS) toward IoT, with data rate between 70 and 240 kb/s.

Although the approaches proposed by 3GPP reduce the energy consumption and cost of the devices, they have not yet caught up to their non-3GPP counterparts. For instance, module cost for LoRaWAN and SigFox is around $2–5 and for eMTC is still around $8–12. Despite the expected broad adoption of cellular IoT solutions supported by 3GPP, LoRaWAN presents some assets that prevail against these technologies in specific market niches. Current assets are:

- The number of LoRaWAN network deployments is increasing continuously, while on the other hand, few initial NB-IoT deployments have already been deployed.
- LoRaWAN operates in the industrial, scientific, and medical (ISM) band, whereas cellular IoT operates in licensed bands; this fact favors the deployment of private LoRaWAN networks without the involvement of mobile operators.
- LoRaWAN has backing from industry, including CISCO, IBM, and HP, among others. In the future, both technologies will probably coexist when 3GPP solutions are backed up by large volumes.

Overview of LoRaWAN

LoRa is the physical layer used in LoRaWAN. It features low-power operation (around 10 years of battery lifetime), low data rate (27 kb/s with spreading factor 7 and 500 kHz channel or 50 kb/s with FSK) and long communication range (2–5 km in urban areas and 15 km in suburban areas). It was developed by Cycleo, a French company acquired by Semtech. LoRaWAN networks are organized in a star-of-stars topology, in which gateway nodes relay messages between end devices and a central network server. End devices send data to gateways over a single wireless hop, and gateways are connected to the network server through a non-LoRaWAN network (e.g., IP over cellular or Ethernet). Communication is bidirectional, although uplink communication from end devices to the network server is strongly favored, as explained in the following [6].

LoRaWAN defines three types of devices (*Classes A*, *B*, and *C*) with different capabilities [6]. Class A devices use pure ALOHA access for the uplink. After sending a frame, a Class A device listens for a response during two downlink receive windows. Each receive window is defined by the duration, an offset time, and a data rate. Although offset time can be configured, the recommended value for each receive window is 1 s and 2 s, respectively. Downlink transmission is only allowed after a successful uplink transmission. The data rate used in the first downlink window is calculated as a function of the uplink data rate and the receive window offset. In the second window the data rate is fixed to the minimum, 0.3 kb/s. Therefore, downlink traffic cannot be transmitted until a successful uplink transmission is decoded by the gateway. The second receive window is disabled when downlink traffic is received by the end device in the first window. Class A is the class of LoRaWAN devices with the lowest power consumption. Class B devices are designed for applications with additional downlink traffic needs. These devices are synchronized using periodic beacons sent by the gateway to allow the scheduling of additional receive windows for downlink traffic without prior successful uplink transmissions. Obviously, a trade-off between downlink traffic and power consumption arises. Finally, Class C devices are always listening to the channel except when they are transmitting. Only Class A must be implemented in all end devices, and the other classes must remain compatible with Class A. In turn, Class C devices cannot implement Class B. The three classes can coexist in the same network, and devices can switch from one class to another. However, there is not a specific message defined by LoRaWAN to inform the gateway about the class of a device; this is up to the application.

The underlying PHY of the three classes is the same. Communication between end devices and gateways start with a *Join procedure* that can occur on multiple frequency channels (e.g., in EU863-870 ISM Band there are 3 channels of 125 kHz that must be supported by all end devices and 3 additional 125 kHz channels) by implementing pseudo-random channel hopping. Each frame is transmitted with a specific spreading factor (SF), defined as $SF = \log_2 (R_c/R_s)$, where R_s is the symbol rate and R_c is the chip rate. Accordingly, there is a trade-off between SF and communication range. The higher the SF (i.e., the slower the transmission), the longer the communication range. The codes used in the different SFs are orthogonal. This means that multiple frames can be exchanged in the network at the same time, as long as each one is sent with one of the six different SFs (from SF = 7 to SF = 12). Depending on the SF in use, LoRaWAN data rate ranges from 0.3 kb/s to 27 kb/s.

The three classes can coexist in the same network and devices can switch from one class to another. However, there is not a specific message defined by LoRaWAN to inform the gateway about the class of a device and this is up to the application.

The maximum duty cycle, defined as the maximum percentage of time during which an end device can occupy a channel, is a key constraint for networks operating in unlicensed bands. Therefore, the selection of the channel must implement pseudo-random channel hopping at each transmission and be compliant with the maximum duty cycle. For instance, the duty cycle is 1 percent in EU 868 for end devices.

The LoRa physical layer uses chirp spread spectrum (CSS) modulation, a spread spectrum technique where the signal is modulated by chirp pulses (frequency varying sinusoidal pulses), hence improving resilience and robustness against interference, Doppler effect, and multipath. Packets contain a preamble (typically with 8 symbols), a header (mandatory in explicit mode), the payload (with a maximum size between 51 and 222 bytes, depending on the SF), and a cyclic redundancy check (CRC) field (with configurations that provide a coding rate from 4/5 to 4/8). Typical bandwidth (BW) values are 125, 250, and 500 kHz in the HF ISM 868 and 915 MHz band, while they are 7.8, 10.4, 15.6, 20.8, 31.2, 41.7, and 62.5 kHz in the LF 160 and 480 MHz bands. The raw data rate varies according to the SF and the BW, and ranges between 22 b/s (BW = 7.8 kHz and SF = 12) to 27 kb/s (BW = 500 kHz and SF = 7) [2]. Frequency hopping is exploited at each transmission in order to mitigate external interference [10].

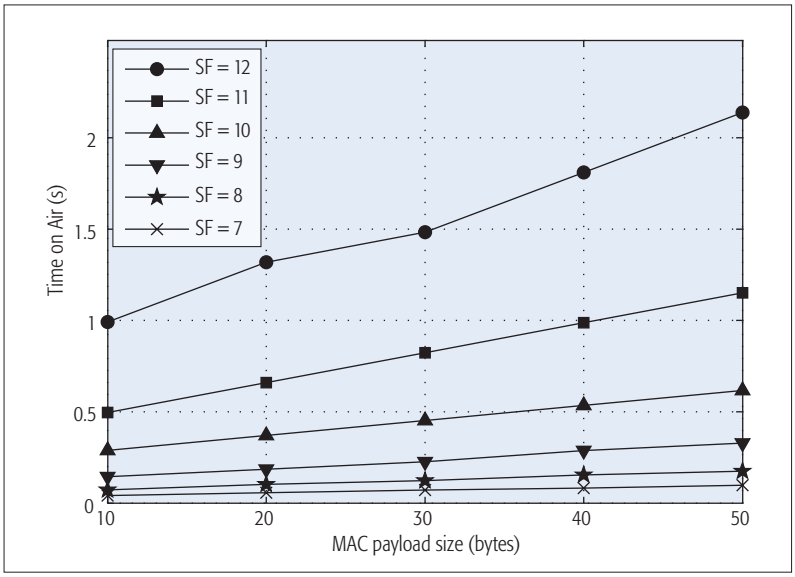

Figure 1. Time on Air of LoRaWAN with code rate 4/5 and a 125 kHz bandwidth.

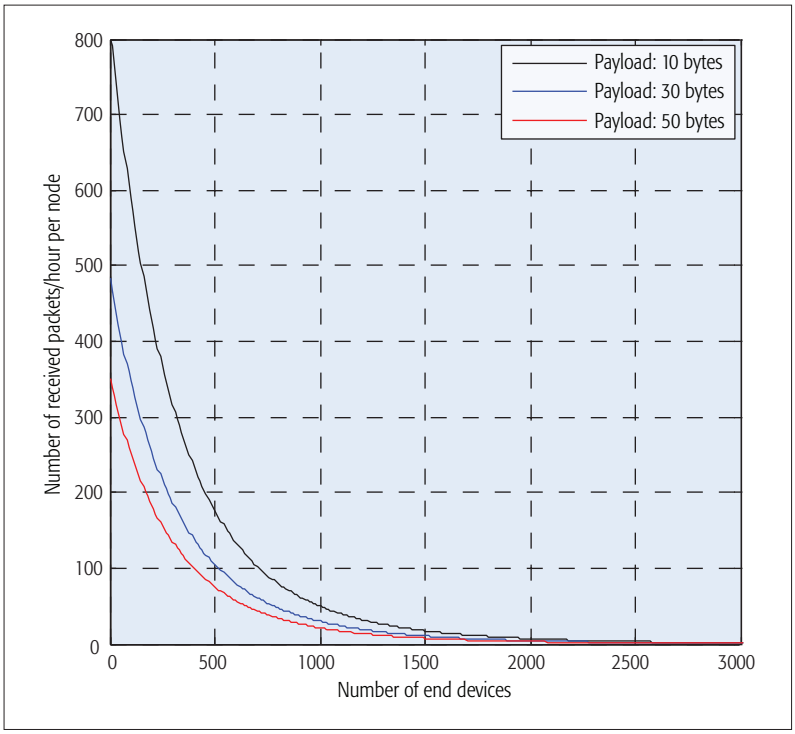

Figure 2. Number of packets received per hour when end devices attempt transmission at nd/T_{a_i} packets/s with coding rate 4/5 and n = 3 channels with 125 kHz bandwidth.

Capacity and Network Size Limitations

In this section we study the LoRaWAN network scale with respect to data rate, duty cycle regulations, and so on.

Network Size Limited by Duty Cycle

Although the performance of LoRaWAN is determined by PHY/medium access control (MAC) overviewed in the previous section, the duty cycle regulations in the ISM bands [11, 12] arise as a key limiting factor. If the maximum duty cycle in a sub-band is denoted by d and the packet transmission time, known as time on air, is denoted by T_a, each device must be silent in the sub-band for a minimum off-period $T_s = T_a(1/d - 1)$. For instance, the maximum duty cycle of the EU 868 ISM band is 1 percent, and it results in a maximum transmission time of 36 s/h in each sub-band for each end device. Figure 1 shows the time on air of a packet transmission with coding rate 4/5 over a 125 kHz BW channel. It is known that large SFs allow longer communication range. However, as observed in Fig. 1, large SFs also increase the time on air and, consequently, the off-period duration. This problem is exacerbated by the fact that large SFs are used more often than small SFs. For instance, considering a simple scenario with end devices distributed uniformly within a round-shaped area centered at the gateway, and a path loss calculated with the Okumura-Hata model for urban cells [13], the probability that an end device uses an SF i, p_i, would be p_{12} = 0.28, p_{11} = 0.20, p_{10} =0.14, p_9 = 0.10, p_8 = 0.08, and p_7 = 0.19.

Although Listen Before Talk is not precluded in LoRaWAN, only ALOHA access is mandatory. Accordingly, the LoRaWAN capacity can be calculated roughly as the superposition of independent ALOHA-based networks (one independent network for each channel and for each SF, since simultaneous transmissions only cause a collision if they both select the same SF and channel; no capture effect is considered). However, and in contrast to pure ALOHA, a LoRaWAN device using SF i cannot exceed a transmitted packet rate given by nd/T_{a_i}, where n is the number of channels, d is the duty cycle, and T_{a_i} is the time on air with SF i.

In the simple scenario described above, if all end devices transmit packets at the maximum packet rate nd/T_{a_i}, the number of packets successfully received by the gateway decreases, as shown in Fig. 2, where a network with n = 3 chan-

nels is analyzed. The number of received packets drops due to the effect of collisions.

In Fig. 3 the number of packets received successfully per hour and end device is shown for deployments with {250, 500, 1000, 5000} end devices and n = 3 channels. For low transmission rate values (in packets per hour), throughput is limited by collisions; for high values, the maximum duty cycle prevents end devices from increasing the packet transmission rate and stabilizes the throughput. For deployments with a "small" number of end devices, the duty cycle constraint limits the maximum throughput.

Table 1 summarizes the maximum throughput per end device and the probability of successful reception for a set of different deployments. The maximum throughput falls as the number of end devices grows.

Reliability and Densification Drain Network Capacity

In LoRaWAN, reliability is achieved through the acknowledgment of frames in the downlink. For Class A end devices, the acknowledgment can be transmitted in one of the two available receive windows; for Class B end devices, it is transmitted in one of the two receive windows or in an additional time-synchronized window; for Class C end devices, it can be transmitted at any time.

In LoRaWAN the capacity of the network is reduced not only due to transmissions in the downlink, but also due to the off-period time following those transmissions (gateways must be compliant with duty cycle regulation). Therefore, the design of the network and the applications that run on it must minimize the number of acknowledged frames to avoid the capacity drain. This side-effect calls into question the feasibility of deploying ultra-reliable services over large-scale LoRaWAN networks.

At this point of development of the technology, LoRaWAN faces deployment trends that can result in future inefficiencies. Specifically, LoRaWAN networks are being deployed following the cellular network model, that is, network operators provide connectivity as a service. This model is making gateways become base stations covering large areas. The increase in the number of end devices running applications from different vendors over the same shared infrastructure poses new challenges to coordinate the applications. In particular, each application has specific constraints in terms of reliability, maximum latency, transmission pattern, and so on. The coordination of the diverse requirements over a single shared infrastructure using an ALOHA-based access is one of the main future challenges for the technology. Therefore, fair spectrum sharing is required beyond the existing duty cycle regulations. Finally, the unplanned and uncoordinated deployment of LoRaWAN gateways in urban regions, along with the deployment of alternative LPWAN solutions (e.g., SigFox), could cause a decrease of the capacity due to collisions and the use of larger SFs (to cope with higher interference levels).

Use Cases

Several application use cases are considered in order to analyze the suitability of LoRaWAN and complement the understanding of the advantages and limitations of the technology when applied to different types of data transmission patterns, latency requirements, scale, and geographic dispersion, among others.

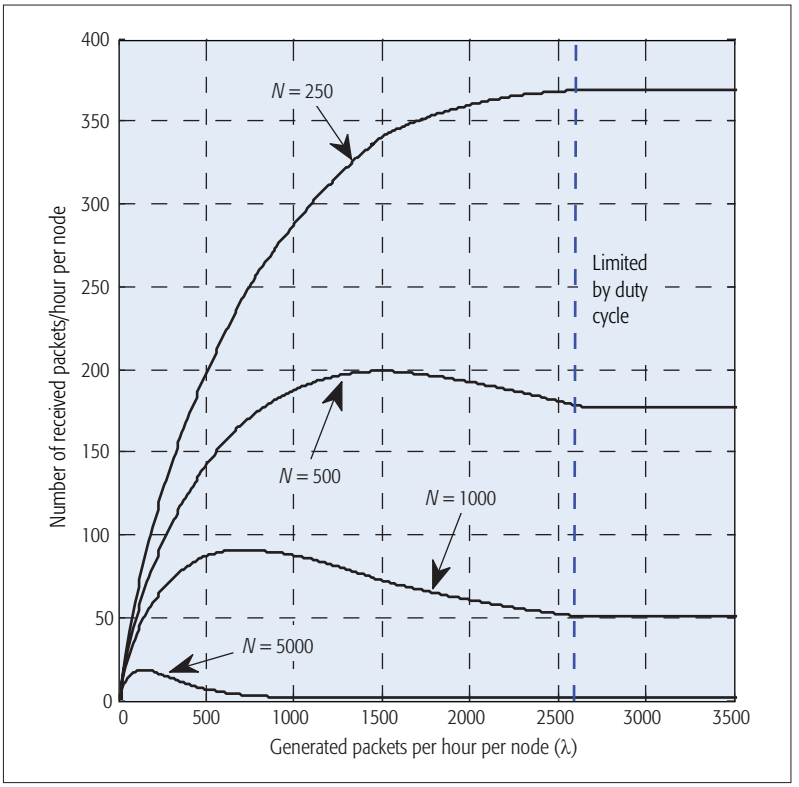

Figure 3. Number of 10 B payload packets received per hour and node for {250, 500, 1000, 5000} end devices and n = 3 channels as a function of packet generation.

Real-Time Monitoring

Agriculture, leak detection, and environment control are applications with a reduced number of periodic/aperiodic messages and relaxed delay constraints. In contrast, the communication range must be long enough to cope with dispersed location of end devices. LoRaWAN has been designed to handle the traffic generated by this type of applications and meets their requirements as long as the deployment of the gateways is enough to cover all end devices.

On the other hand, industrial automation, critical infrastructure monitoring, and actuation require some sort of real-time operation. Real time is understood in general by low latency and bounded jitter, and depends on the specific application. LoRaWAN technology cannot claim to be a candidate solution for industrial automation, considering, for example, that industrial control loops may require response times around 1–100 ms and that even for small packets of 10 B, the time on air with SF = 7 is around 40 ms. As presented in the previous section, due to the MAC nature of LoRaWAN, deterministic operation cannot be guaranteed despite application-specific periodicity as ALOHA access is subject to contention, which impacts network jitter. Despite that, small LoRaWAN networks can deliver proper service to applications that require, for instance, sampling data every second. To do that, two main design considerations should be taken into account:

	250 end devices			500 end devices			1000 end devices			5000 end devices		
Payload (bytes)	10	30	50	10	30	50	10	30	50	10	30	50
Max. throughput per node (packets/hour)	367	217	157	198	117	84	89	53	38	18	10	7.3
Max. throughput per node (bytes/hour)	3670	6510	7850	1980	3510	4200	890	1590	1900	180	300	365
λ of the max. throughput (packets/hour)	2620	1500	1090	1500	870	620	670	390	280	130	70	50
Prob. of successful transmission (%)	14.01	14.47	10.73	13.20	13.45	13.55	13.28	13.59	13.57	13.85	14.29	14

Table 1. Maximum throughput and probability of successful transmission for different deployments (with n = 3 channels and 1 percent duty cycle).

- The spreading factor should be as small as possible to limit both the time on air and the subsequent off-period. In other words, the gateway must be close enough to the end devices.
- The number of channels must be carefully designed and must be enough to:
  - Minimize the probability of collisions (tightly coupled with the number of end-devices)
  - Offer quick alternative channels for nodes to retransmit collided packets, thereby diminishing the impact of the duty cycle

Despite the two aforementioned aspects, latency will not be deterministic.

Metering

The LoRa Alliance is working on standard encapsulation profiles for popular M2M and metering protocols. Keeping an existing application layer allows to keep most of the firmware and ecosystem intact, facilitating migration to LPWAN. These protocols include Wireless M-Bus for water or gas metering, KNX for building automation, and ModBus for industrial automation. It is important to understand that those scenarios range from time-sensitive operation to best effort monitoring. Therefore, it is key to identify in such a diverse ecosystem what the requirements of each application are and if LoRaWAN is the appropriate technology to address them.

Smart City Applications

LoRaWAN has shown key success stories with smart lighting, smart parking, and smart waste collection thanks to their scale and the nature of the data generated by those applications. These encompass periodic messaging with certain delay tolerance. For example, smart parking applications report the status of the parking spots when a change is detected [14]. Parking events are slow; therefore, network signaling is limited to a few tens of messages per day. Analogously, smart waste collection systems and smart lighting actuate and report information in response to a measure with large variation periods. Although latency and jitter are not major issues in these applications, in some of them the triggering factor is simultaneous for a huge number of end devices. For instance, sunset and sundown trigger lighting elements around the whole city, thereby causing an avalanche of messages. LoRaWAN is an appropriate technology for this use case since it handles the wide coverage area and the significant number of users at the expense of an increasing number of collisions, latency, and jitter.

Smart Transportation and Logistics

Transportation and logistics are seen as two major pillars of the expected IoT growth over the next few years thanks to their impact on the global economy. Most applications target efficiency in areas such as public transportation and transport of goods. However, some applications are tolerant to delay, jitter, and unreliability, and others are not.

Different standards have been developed in the 5.9 GHz band for intelligent transportation systems (ITS) based on the IEEE 802.11p standard. The constraints on delay are diverse for different applications, but LoRaWAN, being an LPWAN solution, is not suitable for these applications. On the contrary, solutions such as fleet control and management can be supported by LoRaWAN. Roaming is one of the developments under definition within the LoRa Alliance to enhance mobility. Specifically, a future roaming solution is expected to support back-end to back-end secure connections, clearing and billing between operators, location of end devices (pointed out as an open research challenge below), and transparent device provisioning across networks.

Video Surveillance

The most common digital video formats for IP-based video systems are MJPEG, MPEG-4, and H.264. The bit rate recommended for IP surveillance cameras ranges from 130 kb/s with low-quality MJPEG coding to 4 Mb/s for 1920 × 1080 resolution and 30 fps MPEG-4/H.264 coding. Given that LoRaWAN data rate ranges from 0.3 to 50 kb/s per channel, LoRaWAN will not support these applications.

Open Research Challenges

The effect of the duty cycle stated earlier jeopardizes the actual capacity of large-scale deployments. This has been initially addressed by TheThingsNetwork [15], an interesting global, open, crowd-sourced initiative to create an IoT data network over LoRaWAN technology. The proposed solution defines an access policy, known as the TTN Fair Access Policy, that limits the time on air of each end device to a maximum of 30 s/day. This policy is simple to implement and guarantees pre-defined end-device requirements for a large-scale network (more than 1000 end devices per gateway). However, it fails to provide the network with enough flexibility to adapt to environment and network conditions (i.e. link budget of each end device, number of end devices, number of gateways, etc), as well

as to applications with tight latency or capacity requirements.

At this stage, the optimization of the capacity of the LoRaWAN network, as well as the possibility to perform traffic slicing for guaranteeing specific requirements on a service basis, remain open research issues. From the authors' point of view, the research community will have to address the following open research challenges during the next years.

Explore new channel hopping methods: A pseudo-random channel hopping method is natively used in LoRaWAN to distribute transmissions over the pool of available channels, thereby reducing the collision probability. However, this method cannot meet traffic requirements when there are latency, jitter, or reliability constraints (i.e., downlink acknowledgments for all packets), and it is not able to be adapted according to the noise level of each channel. The design of pre-defined and adaptive hopping sequences arises as an open research issue. From the authors' point of view, the proposed channel hopping sequences should be able to reserve a set of channels for retransmissions of critical packets, both in the uplink and in the downlink (acknowledgment). The design of feasible feedback mechanisms between gateways and end devices must be a key part of the approach in a system where uplink traffic is strongly favored.

TDMA over LoRaWAN: The random nature of ALOHA-based access is not optimal to serve deterministic traffic, which is gaining importance in the IoT ecosystem. Building a complete or hybrid TDMA on top of LoRaWAN opens up new use cases for this technology and provides additional flexibility.The TDMA scheduler should be able to allocate resources for ALOHA-based access and schedule deterministic traffic along time and over the set of available channels. The proposed schedulers should manage the trade-off between resources devoted to deterministic and non-deterministic traffic, meet the regional duty cycle constraints, and guarantee fairness with coexisting LoRaWAN networks.

Geolocation of end devices: The location of end devices is a mandatory requirement for specific use cases, particularly in Industry 4.0. However, GPS-based solutions are not feasible due to cost, and CPU and energy consumption. Currently, interesting works have been initiated to develop time difference of arrival (TDOA)-based triangulation techniques for LoRaWAN. It has been shown that this approach benefits from large SFs and dense gateway deployments.

Cognitive Radio: As pointed out earlier, regulation in ISM bands concerning maximum duty cycle has a significant impact on the capacity of the network. One of the most promising future directions could be the inclusion of cognitive radio in the LoRaWAN standard. In contrast to Weightless-W, LoRaWAN has not been designed to operate in TV white spaces. In the future, the inclusion of cognitive radio in the LoRaWAN standard would be subject to a significant reduction of the energy consumption associated with cognitive radio techniques.

Power reduction for multihop solutions: LoRaWAN is organized with a single-hop star topology for simplicity. As discussed earlier, the impact of high SFs on the capacity of the network is two-fold, since it increases both the time on air and the off period. A two-hop strategy for LoRaWAN networks should be investigated to figure out its potential. Proposals in this direction should consider the reduction of transmitted power and the decrease of the SFs. On the other hand, negative effects such as complexity, synchronization, and increasing power consumption of relays should also be analyzed to thoroughly characterize the trade-off.

Densification of LoRaWAN networks: The proliferation of LPWAN technologies, and particularly LoRaWAN, poses coexistence challenges as the deployment of gateways populate urban areas. Given the random-based access in unlicensed bands of LoRaWAN and its inherent unplanned deployment, the performance achieved in isolated networks is put into question in scenarios with coexisting gateways and limited number of available channels. It is essential to devise coordination mechanisms between gateways from the same or different operators to limit interference and collisions. The coexistence mechanisms encompass coordination and reconfiguration protocols for gateways and end devices.

It is essential to devise coordination mechanisms between gateways from the same or different operators to limit interference and collisions. The coexistence mechanisms encompass coordination and reconfiguration protocols for gateways and end devices.

Conclusions

This article is aimed at clarifying the scope of LoRaWAN by exploring the limits of the technology, matching them to application use cases, and stating the open research challenges. In the low-power M2M fragmented connectivity space there is not a single solution for all the possible connectivity needs, and LoRaWAN is not an exception. A LoRaWAN gateway, covering a range of tens of kilometers and able to serve up to thousands of end devices, must be carefully dimensioned to meet the requirements of each use case. Thus, the combination of the number of end devices, the selected SFs, and the number of channels will determine if the LoRaWAN ALOHA-based access and the maximum duty cycle regulation fit each use case. For instance, we have seen that deterministic monitoring and real-time operation cannot be guaranteed with the current LoRaWAN state of the art.

Acknowledgments

This work is partially supported by the Spanish Ministry of Economy and the FEDER regional development fund under SINERGIA project (TEC2015-71303-R), and by the European Commission through projects H2020 F-Interop and H2020 ARMOUR.

References

[1] L. Labs, "A Comprehensive Look at Low Power, Wide Area Networks for Internet of Things Engineers and Decision Makers," White Paper 2016; http://info.link-labs.com/lpwan-1, accessed Dec. 19, 2016.
[2] C. Goursaud and J. M. Gorce, "Dedicated Networks for IoT: PHY/MAC State of the Art and Challenges," *EAI Endorsed Trans. Internet of Things*, vol. 15, no. 1, Oct. 2015.
[3] X. Xiong *et al.*, "Low Power Wide Area Machine-to-Machine Networks: Key Techniques and Prototype," *IEEE Commun. Mag.*, vol. 53, no. 9, Sept. 2015, pp. 64–71.
[4] R. Sanchez-Iborra and M.-D. Cano, "State of the Art in LP-WAN Solutions for Industrial IoT Services," *Sensors*, vol. 16, no. 5, Feb. 2016, pp. 708.
[5] G. Margelis *et al.*, "Low Throughput Networks for the IoT: Lessons Learned from Industrial Implementations," *IEEE World Forum on Internet of Things*, Dec. 2015, pp. 181–186.
[6] N. Sornin *et al.*, "LoRa Specification 1.0," LoRa Alliance Std Spec., Jan. 2015; www.lora-alliance.org, accessed Dec. 19, 2016.

[7] N. Ducrot *et al.*, "LoRa Device Developer Guide," Orange, Connected Objects and Partnership tech doc., Apr. 2016; https://partner.orange.com/wp-content/uploads/2016/04/LoRa-Device-Developer-Guide-Orange.pdf, accessed Dec. 19, 2016.
[8] A. Minaburo, A. Pelov, and L. Toutain, "LP-WAN Gap Analysis," IETF Std, Feb. 2016; https://tools.ietf.org/html/draft-minaburo-lp-wan-gap-analysis-00, accessed Dec. 19, 2016.
[9] Nokia, "LTE Evolution for IoT Connectivity," white paper, 2016; http://resources.alcatel-lucent.com/asset/200178, accessed Dec. 19, 2016.
[10] T. Watteyne, A. Mehta, and K. Pister, "Reliability Through Frequency Diversity: Why Channel Hopping Makes Sense," *ACM Symp. Performance Evaluation of Wireless Ad Hoc, Sensor, and Ubiquitous Networks*, Oct. 2009, pp. 116–23.
[11] E. C. Committee and Others, "ERC Recommendation 70-03," Troms, Ed., Oct. 2016.
[12] F. C. Commission *et al.*, "FCC Part 15–Radio Frequency Devices, Code of Federal Regulation 47 CFR Ch. 1 (10-1-15 Edition)."
[13] A. F. Molisch, *Wireless Communications*, 2nd ed, Wiley, 2011.
[14] B. Martinez *et al.*, "Lean Sensing: Exploiting Contextual Information for Most Energy-Efficient Sensing," *IEEE Trans. Industrial Informatics*, vol. 11, no. 5, Oct. 2015, pp. 1156–65.
[15] W. Giezeman and J. Stokking, "The Things Network," June 2016; www.thethingsnetwork.org, accessed Dec. 19, 2016.

Biographies

Ferran Adelantado (ferranadelantado@uoc.edu) received his engineering degree in telecommunications (2007) and his Ph.D. degree in telecommunications (2007) from Universitat Politecnico de Catalunya (UPC), Barcelona, Spain, and his B.Sc. in business sciences (2012) from Universitat Oberta de Catalunya (UOC). Currently, he is an associate professor at UOC and a researcher at the Wireless Networks Research Group (WINE). His research interests are wireless networks, particularly 5G, LPWAN, and IoT technologies.

Xavier Vilajosana is Principal Investigator of the Wireless Networks Research Lab at the Open University of Catalonia. He is also co-founder of Worldsensing and OpenMote Technologies. He is an active member of the IETF 6TiSCH WG, where he has authored different standard proposals. He also holds 30 patents. He has been a visiting professor at the Prof. Pister University of California (UC) Berkeley lab. He co-leads Berkeley's OpenWSN project. He was a senior researcher at HP R&D Labs (2014–2016) and a visiting researcher at the France Telecom R&D Labs Paris (2008). He holds Ph.D.(2009), M.Sc., and M.Eng. (2004) degrees from UPC.

Pere Tuset-Peiro [M'12] (peretuset@uoc.edu) is an assistant professor at the Department of Computer Science, Multimedia and Telecommunications, and a researcher at the Internet Interdisciplinary Institute (IN3), both of UOC. He received his B.Sc. and M.Sc. in telecommunications engineering from UPC in 2007 and 2011, respectively, and his Ph.D. in network and information technologies from UOC in 2015. Currently, he has more than 20 high-impact publications and 7 international patents.

Borja Martinez received his B.Sc. in physics and electronics Engineering, his M.Sc. in microelectronics, and his Ph.D. in computer science from the Universidad Autónoma de Barcelona (UAB), Spain. From 2005 to 2015 he was an assistant professor at the Department of Microelectronics and Electronic Systems of UAB. He is currently a research fellow at the Internet Interdisciplinary Institute (IN3-UOC). His research interests include low-power techniques for smart wireless devices, energy efficiency, and algorithms.

Joan Melià-Seguí (melia@uoc.edu), Ph.D. (2011), is a lecturer at the Estudis de Informàtica, Multimèdia i Telecomunicació, and a researcher at IN3, both at UOC. Before, he was a postdoctoral researcher at Universitat Pompeu Fabra and the Xerox Palo Alto Research Center (PARC). He has published more than 30 papers and one patent in the areas of the Internet of Things, intelligent systems, security, and privacy.

Thomas Watteyne (www.thomaswatteyne.com) is a researcher in the EVA team at Inria, Paris, and a senior networking design engineer at Linear Technology/Dust Networks in Silicon Valley. He co-chairs the IETF 6TiSCH working group. He did his postdoctoral research with Prof. Pister at UC Berkeley. He co-leads Berkeley's OpenWSN project. In 2005–2008, he was a research engineer at France Telecom, Orange Labs. He holds a Ph.D. (2008), an M.Sc., and an M.Eng. (2005) from INSA Lyon, France.

Self-Organized Connected Objects: Rethinking QoS Provisioning for IoT Services

Hajar Elhammouti, Essaid Sabir, Mustapha Benjillali, Loubna Echabbi, and Hamidou Tembine

ABSTRACT

The proliferation of connected objects has revolutionized the traditional Internet, giving rise to the emerging Internet of Things (IoT). The IoT ecosystem is very large, and it includes smart interconnections among sensors and devices with applications in both the industrial world and customers' daily lives. As a unified standard for IoT is still under development, many challenges related to IoT must be discussed and addressed, especially those related to energy efficiency. This article tackles the challenge of energy efficiency in IoT from a novel perspective. It shows that instead of maximizing the QoS, which is generally energy costly, better energy efficiency can be achieved by targeting satisfactory QoS levels only. The approach aims to enhance energy efficiency while ensuring a desired QoS threshold. This is supported by a game theoretical solution concept referred to as the *satisfaction equilibrium*. Moreover, as IoT objects require self-configuring techniques to maintain the network scalability and flexibility, this article introduces fully distributed schemes in order to reach efficient satisfaction equilibria in both slow- and fast-fading channel contexts. The proposed schemes can also be adapted to achieve the maximum performance of IoT applications that desire the highest QoS levels. The performance of these algorithms is illustrated through a smart home use case scenario.

INTRODUCTION

Ubiquitous and pervasive connected technology has become an integral part of users' daily lives. Particularly, various studies estimated that the number of connected objects was around a few dozen billions in 2015. This number is expected to experience a many-fold increase by 2020, giving rise to the so-called Internet of Things (IoT) [1].

The large adoption of IoT is fueled by the various applications it covers, including machine-type (MTC), device-to-device (D2D), and vehicle-to-everything (V2X) communications. Figure 1 shows multiple use cases in an urban IoT ecosystem.

However, while facilitating the users' quality of life, IoT is adding a number of new challenges for designers, especially those related to the limited power capacity of the connected objects [2]. As a result, numerous energy-saving techniques are considered in order to improve the lifetime of batteries and thus enhance energy efficiency. The need for high energy performance is enforced by fifth generation (5G) requirements with the targeted 1000× energy efficiency enhancement [3]. Furthermore, the ever increasing number of connected objects has also triggered scalability requirements. Consequently, energy-efficient techniques should be implemented in a fully distributed manner in order to enable self-provisioning capabilities, hence allowing more network flexibility.

Techniques such as sleep mode optimization [4], power control mechanisms [5], adaptive data rates [6], and learning algorithms [7] are used to maximize quality of service (QoS) while minimizing energy consumption. These approaches aim, above all, to achieve the highest QoS levels. However, one cannot separate energy consumption from energy efficiency as they both define a well-known trade-off [8]. In general, in order to reach the highest QoS levels, a significant amount of energy is required. As a result, substantial energy savings can be achieved by targeting satisfactory QoS levels only. This leads to a fundamental question: is it worth maximizing QoS when achieving satisfactory QoS levels is possible and energy-efficient?

This article proposes an alternative energy saving approach that enhances the energy efficiency while ensuring satisfactory QoS levels. The approach is supported by a game theoretical concept, namely, the *satisfaction equilibrium* (SE) [9].

From a game theoretical perspective, two other solution concepts can be considered when dealing with energy efficiency. First, Nash equilibrium (NE) is defined as a strategy profile where no player has incentive to deviate unilaterally. However, it has been shown that the NE may fail to model network performance. For example, when the players' payoffs are described by their individual throughput, by acting selfishly following a Nash behavior, players increase their power, resulting in more interference and wasted resources [10]. Alternatively, in order to support users' QoS requirements, the constrained NE (CNE, also called generalized NE) is introduced [11]. Particularly, constrained games are concerned with payoff maximization (or minimization) subject to coupled and/or orthogonal constraints over the players' strategies and/or payoffs. Hence, at a CNE, each player aims to achieve its optimal utility while satisfying QoS constraints. The CNE is designed to accommodate QoS requirements that the NE fails to model. Nevertheless, from a

The authors tackle the challenge of energy efficiency in IoT from a novel perspective. They show that instead of maximizing the QoS, which is generally energy costly, better energy efficiency can be achieved by targeting satisfactory QoS levels only. The approach aims to enhance energy efficiency while ensuring a desired QoS threshold.

Hajar Elhammouti, Mustapha Benjillali, and Loubna Echabbi are with the STRS Lab, INPT; Essaid Sabir is with the NEST research group, ENSEM, Casablanca, Morocco; Hamidou Tembine is with New York University Abu Dhabi.

Digital Object Identifier: 10.1109/MCOM.2017.1600614

When QoS maximization is targeted, only solutions that meet with the global optimum are permitted. As the maximization assumptions are relaxed, the set of feasible strategies is enlarged, which is mathematically less restrictive in terms of problem resolution.

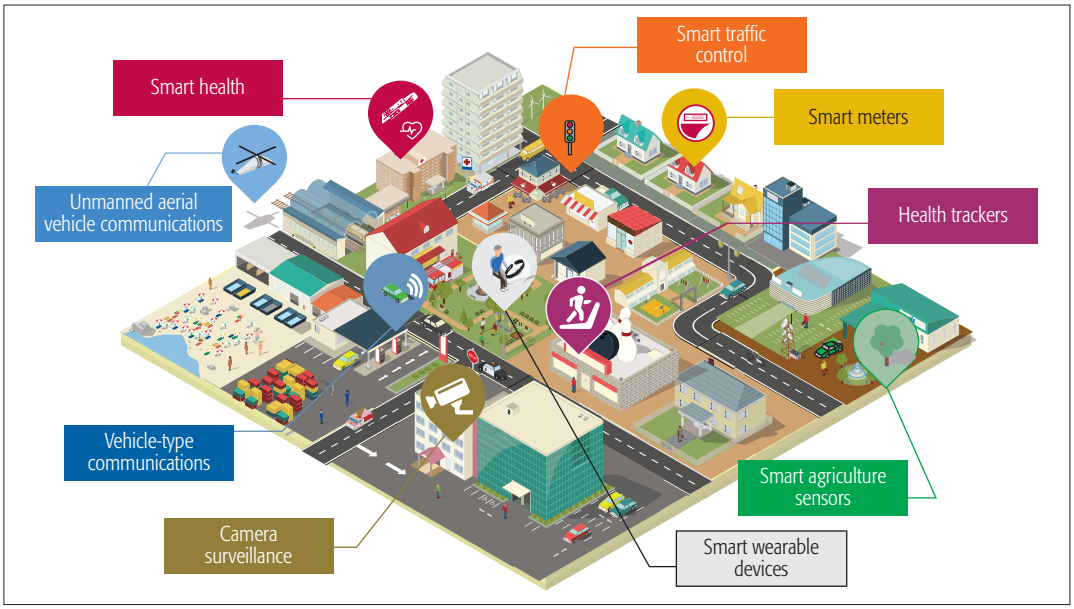

Figure 1. An IoT ecosystem with multiple use cases.

practical point of view, the CNE can be a very restrictive solution that:

- Reduces the set of players' strategies
- Requires costly efforts

In the next section, we motivate the choice of SE as an efficient solution for the IoT ecosystem. The SE and the efficient SE solution concepts are introduced, and their existence and uniqueness are discussed. Fully distributed schemes that reach efficient SE for both slow- and fast-fading channels are described. Besides, in order to reach the maximum performance of the network, an adapted implementation of these schemes is proposed. Finally, a smart home use case scenario is discussed in order to assess the performance of the proposed algorithms.

Why QoS Satisfaction?

Achieving satisfactory QoS is rewarding from both the practical and theoretical viewpoints. First, we highlight the key drivers behind the choice of satisfactory solutions from a practical viewpoint as follows.

To reduce energy consumption: This is the key motivation behind the choice of satisfactory solutions. QoS maximization may require unnecessarily higher energies, whereas achieving an operational network performance that meets users' expectations can result in substantial energy savings.

To adhere to fixed data rates services: Some IoT services such as video surveillance or online gaming need only fixed data rates to work properly. For example, for most video surveillance applications, only 12 to 15 fps is required. Hence, there is no need to maximize data rates since only fixed data rates are used.

To meet users' expectations: The user can be insensitive to small QoS changes that may allow important energy savings. As a consequence, the objective can be to meet a desired QoS level without necessarily exceeding users' expectations. Thus, a good perceived performance could be achieved without targeting optimality in terms of QoS maximization.

To support application-centric networks: Networks are becoming more application-oriented, and a novel applications paradigm has been coined: *app coverage* [12]. App coverage is an application that works properly within a given coverage area. In this specified area, the network is supposed to deliver sufficient performance for a good user experience. Only a required QoS level is achieved within the target area, whereas outside, the application QoS is not satisfied. In order to achieve the target QoS within the tagged area, satisfactory solutions can be envisioned.

Furthermore, from the game-theoretical and mathematical points of view, the choice of satisfactory solutions is endorsed by the following reasons.

Enlarging the set of feasible strategies: When QoS maximization is targeted, only solutions that meet with the global optimum are permitted. As the maximization assumptions are relaxed, the set of feasible strategies is enlarged, which is mathematically less restrictive in terms of problem resolution.

Adaptation to optimization assumptions: In some IoT applications, the full buffer traffic assumption will not be necessary or realistic when tackling the optimization problem. This is the case, for example, for smart meters, which need to report information only at regular or spaced intervals. The full buffer traffic assumption (i.e., users always require the maximization of their QoS) may not hold true. Hence, conveying additional resources by maximizing the QoS all the time would be wasteful.

Satisfactory Power Self-Allocation

As an illustrative example, we consider a power allocation problem that aims to achieve users' throughput satisfaction. We assume a set of heterogeneous devices $\mathcal{N}$ that communicate over the same channel. The channel can be considered as a slow- or fast-fading channel depending on the IoT application (e.g., slow fading for a surveillance camera application and fast fading for V2X communications). Each device selects its transmit power within a bounded power space that we denote by $\mathcal{P}^i$, and has a throughput $r_i(P_i,$

$\mathbf{P}_{-i}$), where $\mathbf{P}_{-i}$ designates the power allocation of devices other than device i. Each device aims to achieve at least a target throughput θ_i.

Satisfaction Equilibrium

In order to meet the respective service requirements, each device should properly adjust its power according to other devices' powers. Hence, an SE $\mathbf{P}^*$ is achieved when each transmitter is satisfied given the powers of the other transmitters. In particular,

$$\forall i \in \mathcal{N},\ r_i(P^*_i, \mathbf{P}^*_{-i}) \geq \theta_i. \quad (1)$$

Existence: The existence of such an equilibrium is conditional:

- First, if the maximum network capacity does not support the overall users' requirements, users' satisfaction cannot be achieved.
- Second, the nature of the power space is inherently related to the SE existence. More precisely, when the set of transmit powers is discrete, the SE may not be achieved. This is because the step probability to meet an SE decreases as the power step size increases.
- Finally, the minimum and maximum allowed transmit powers determine the existence of an SE. In particular, when a transmit power is limited and the QoS requirement is high, satisfaction may not be possible.

Uniqueness: In general, for most QoS functions, when a satisfaction equilibrium exists, it is not necessarily unique. Figure 2 illustrates the feasible powers allowed to device i given fixed transmit powers of its opponents. It can be seen from the figure that the satisfaction of device i can be achieved all along the interval between P_i^{min} and P_i^{max}. Consequently, a continuum of transmit powers can meet device i satisfaction.

Efficient Satisfaction Equilibrium

Among the existing SEs, it would be preferable to select the most efficient. Particularly, an efficient SE (ESE) can be defined as an SE that reduces the overall energy consumption of the analyzed ecosystem, the overall energy consumption being described here by the sum of the transmit powers[1]

$$F(\mathbf{P}) = \sum_{i \in \mathcal{N}} P_i. \quad (2)$$

Existence: When SEs exist, an ESE exists as well. This is due to the convexity of the energy function we consider in Eq. 2, which admits at least one minimizer.

Uniqueness: The uniqueness of an ESE can be solicited as it allows better control of the implemented algorithm's outputs [13]. Intuitively, an ESE can be met when all users' requirements are exactly satisfied, mainly when the inequality of Eq. 1 becomes an equality. This can be proved mathematically when the power space is continuous (interested readers may refer to [14, Proposition 2]). Hence, in order to reach the ESE, a linear system is solved, giving rise to a unique ESE.

How Can QoS Satisfaction Be Reached Efficiently?

Fully distributed schemes are by far the most adapted techniques to achieve better performance in an IoT ecosystem:

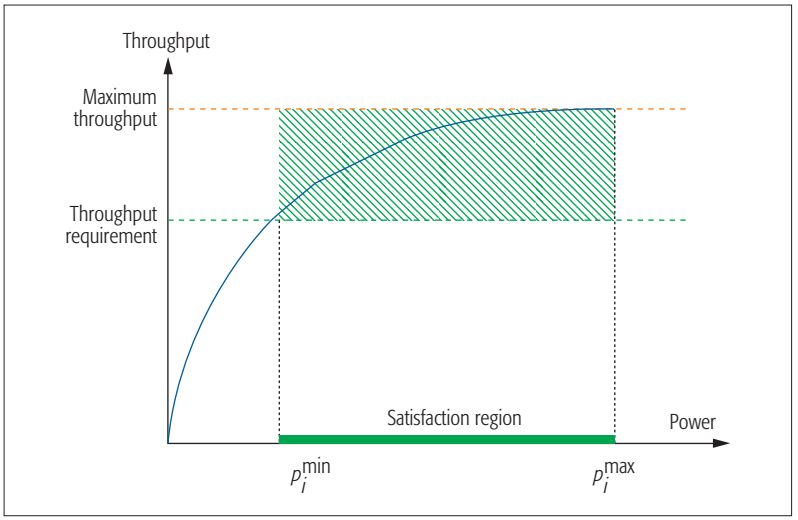

Figure 2. A representation of the throughput of device i given fixed transmit powers of its opponents. This is an illustrative example where the QoS is measured in terms of throughput; other QoS metrics (delay, bit error rate, etc.) could be used interchangeably.

- First, distributed algorithms are important **to ensure the network scalability** since their implementation needs only local information.
- Second, in order **to achieve an efficient satisfaction in a time varying environment** (e.g., radio channel variations), IoT devices should permanently track their satisfactory strategies. This can be achieved by the predictability property of some distributed learning algorithms that enable connected objects to predict their QoS and update their strategies accordingly.

Slow-Fading Channels

When the channel varies within time blocks and the power space is continuous, a distributed learning scheme described by the Banach-Picard algorithm [15] can be used. The algorithm is based on the fixed point theorem. It updates the transmit powers progressively until it reaches the ESE. The Banach-Picard algorithm is known to be convergent with a geometrical rate. Specifically, the algorithm converges within the channels' coherence time, as shown later.

It is important to note that in order to achieve efficient satisfactory throughput for slow-fading channels, connected objects need only to observe their instantaneous throughput without exchanging any information with their neighborhood.

A description of the Banach-Picard algorithm is given in Fig. 3a.

Fast-Fading Channels

The slow fading assumption may not hold for some IoT applications such as V2X communications. Hence, learning algorithms should be adapted to track the time-varying nature of the channels. Mann-iterates [13] is an adaptation of the Banach-Picard algorithm that allows convergence to the ESE in a fast-fading channel context. In contrast to the slow-fading channels case, the learning process cannot be based on the instantaneous throughput. Instead, connected objects should estimate their average throughput and

[1] It is worth noting that the efficiency criterion can be adapted according to the IoT use case. For example, the sum of the transmit powers is meaningful when the ecosystem is considered as a single entity, for example, a smart home where the monthly consumption/bill is considered. On the other hand, when energy fairness is important among appliances, the sum of the powers deviation or a maxmin criterion can be considered as a selection criterion of an ESE.

It is worth noting that each device should use a hybrid implementation of the proposed algorithms. Before switching to one of the presented schemes, the device first listens to the channel, and decides whether it aims at maximizing the QoS or reaching a satisfactory performance.

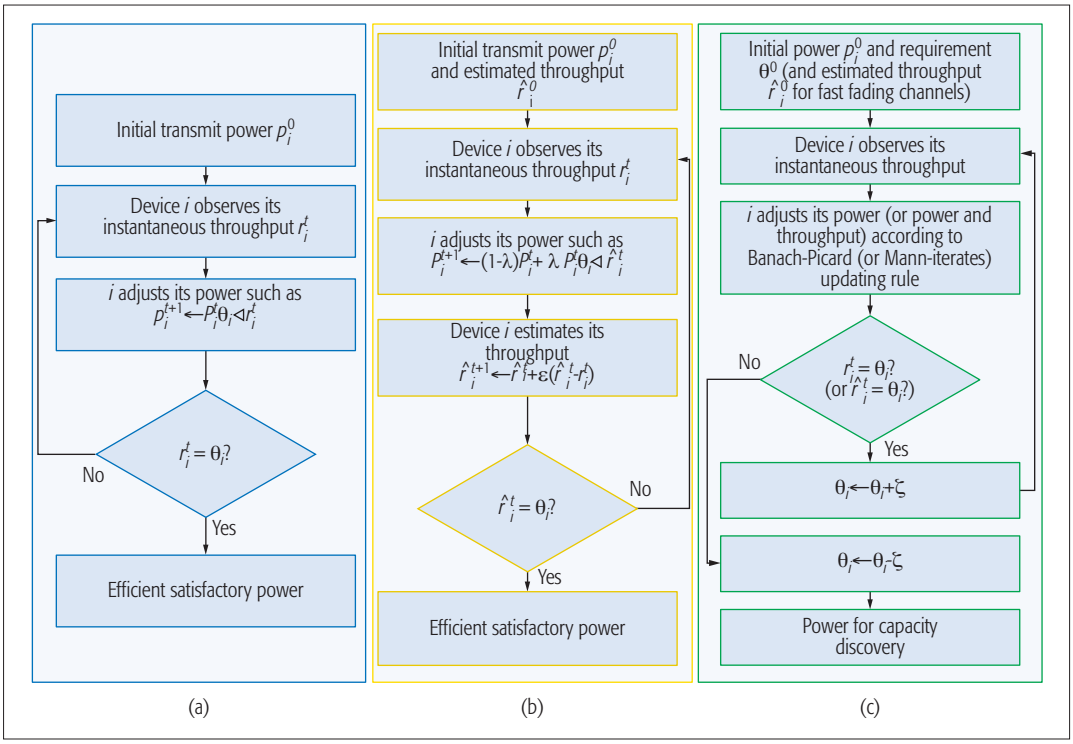

Figure 3. Fully distributed algorithms for efficient satisfactory power allocation and capacity discovery: a) the Bancach-Picard algorithm; b) the Mann-iterates algorithm. Both Banach-Picard and Mann-iterates algorithms can be adapted as described in c) in order to reach the maximum network capacity. The aforementioned algorithms suppose the existence of an ESE; that is, the sum of the requirements is assumed to be supported by the network capacity. The parameters ε and λ should be taken very small in [0.1] in order to ensure the convergence of the Mann-iterates algorithm. ζ is a non-negative constant that represents the requirement adjustment accuracy. A trade-off can be identified between achieving the maximum capacity and the step size ζ. That is, when ζ is very small, the maximum capacity can be reached after a large number of iterations.

adjust their strategies accordingly. As online algorithms show more flexibility, the choice of the best strategy is based on predictability. At each time iteration, the throughput of a given device is updated using a correction rule based on the previous iteration estimation. Consequently, the throughput estimation is more accurate within iterations, getting very close to the instantaneous throughput in a long time run. The Mann-iterates algorithm is presented in Fig. 3b.

Progressive Capacity Discovery

In some scenarios, achieving the best performance is more important than energy consumption. Hence, an adapted version of the aforementioned algorithms, named the progressive capacity discovery algorithm and described in Fig. 3c, can be used. Instead of seeking satisfaction, connected objects can take advantage of the overall network capacity. In particular, when an IoT device reaches its initial requirement, its target is slightly increased, and its power is adjusted accordingly. This process allows stringent QoS applications to improve their performance and avoid any underutilization of the network capacity.

It is worth noting that each device should use a hybrid implementation of the proposed algorithms. Before switching to one of the presented schemes, the device first listens to the channel (in order to detect whether it is a slow- or fast-fading channel), and decides whether it aims to maximize the QoS or reach a satisfactory performance.

Use Case Scenario: A Smart Home

One can identify numerous IoT use cases that require substantial energy savings. Examples include wireless sensor networks (WSNs) where saving battery lifetime is critical. Sensors are generally deployed over inaccessible areas, and therefore, any replacement or recharge of the batteries is impractical most of the time. Another critical example is unmanned aerial vehicle (UAV) networks, commonly known as drone networks. Supplied by limited energy resources, drones require energy conscious behaviors in order to achieve longer flights and better performance.

In this article, we outline another important use case scenario: a smart home. This choice is driven by two main reasons. First, it allows discussion of practical, diverse, and illustrative scenarios that can easily adopt the distributed schemes proposed earlier. Second, although energy efficiency is not critical in a smart home context as energy sustainability can be maintained by smart grids and renewables, the satisfactory QoS approach can be more achievable and cost-efficient.[2]

In the following, we consider a smart home use case, with five main applications/appliances:

- A surveillance camera
- Inband D2D communication
- An online gaming session
- A file upload
- Smart vacuum cleaning

We suppose that the appliances are connect-

[2] Smart grids present many challenges, such as costumers' privacy, renewable availability, and system complexity, that may hamper their wide adoption. Hence, the satisfactory approach can be seen as a good alternative. Another concern related to smart grids is their deployment cost, which is significantly higher than fully distributed schemes locally implemented on smart devices.

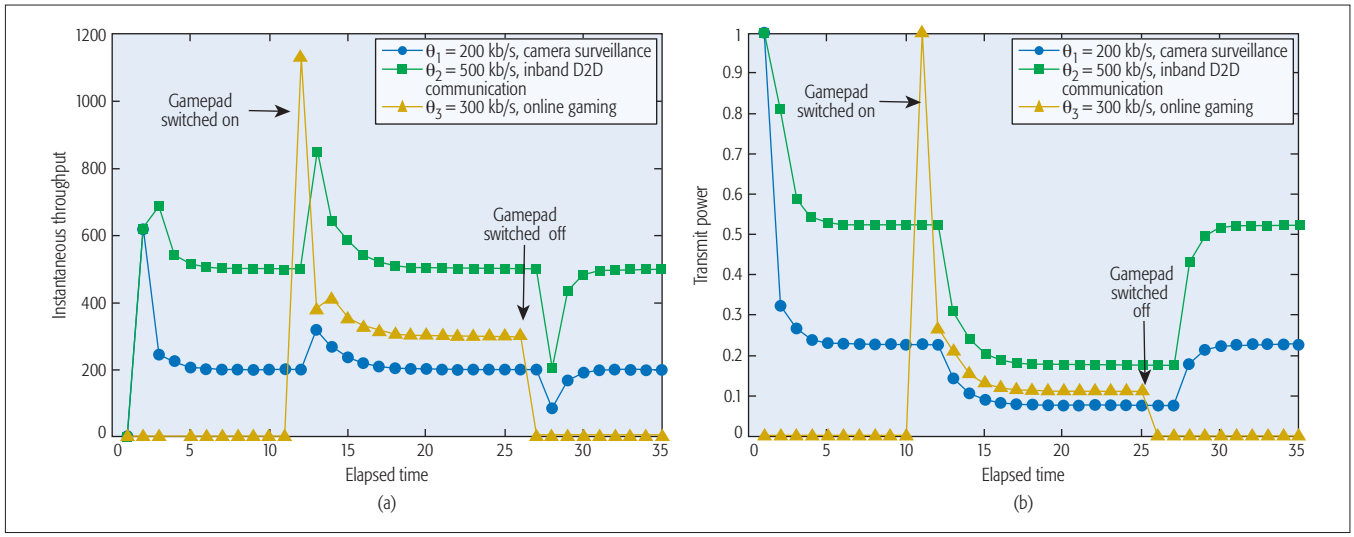

Figure 4. Instantaneous throughput and transmit power for slow-fading channels. Channel gains reflect the path loss, considered equal to $2.5 \cdot 10^{-5}$, and assumed the same for all appliances. The capacity of the network is set to 2000 kb/s.

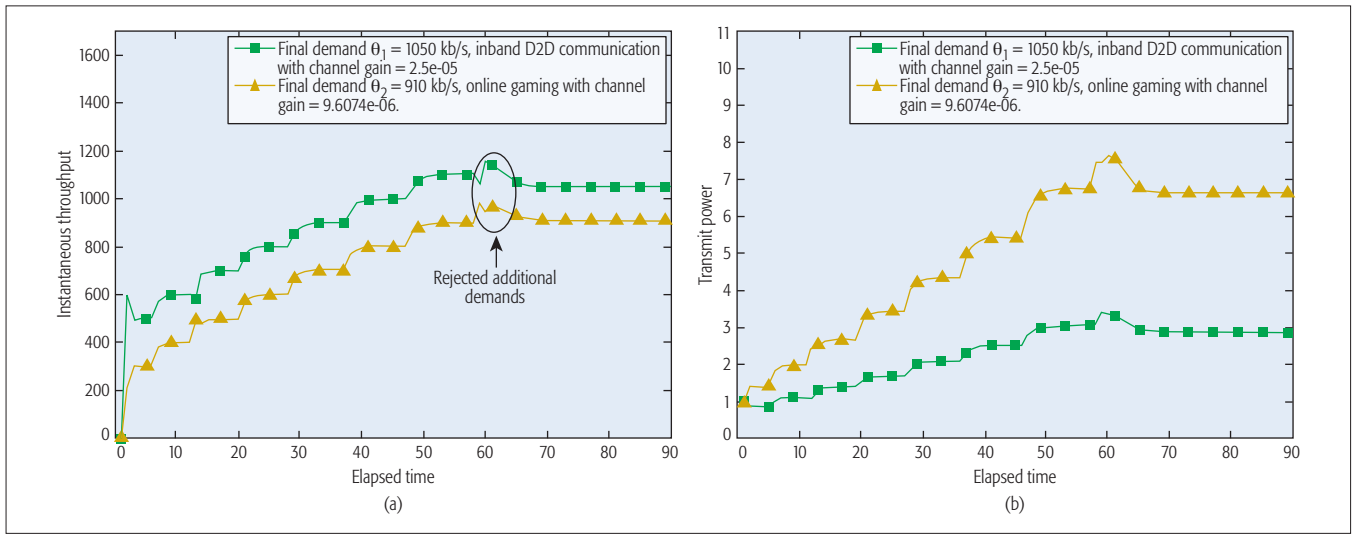

Figure 5. Instantaneous throughput and transmit power using the progressive capacity discovery algorithm for slow-fading channels. The step size is ζ = 100 kb/s. The maximum capacity is 2000 kb/s.

ed to an Internet gateway. We also assume that devices are not active all the time. For example, the smart vacuum cleaner switches off automatically when its task is accomplished. In order to validate our analysis, we present and discuss three illustrative scenarios in the following.

Scenario 1: This scenario covers the slow-fading channel context. It assumes a fixed camera surveillance monitoring and an inband D2D communication, both active when a video game session is started. The D2D communication requires a 500 kb/s data rate, whereas the camera surveillance needs only 200 kb/s. Finally, when the gamepad is connected, 300 kb/s is required to achieve its throughput satisfaction.

Figure 4a shows the instantaneous throughput with respect to time iterations. First, it can be seen from the figure that the Banach-Picard algorithm reaches the satisfactory throughput for both camera surveillance and D2D communication in only a few iterations. Furthermore, we check the behavior of the algorithm when an application is added. Figure 4a shows that when the gamepad is connected (at t = 10), the new demand is quickly satisfied, and the algorithm rapidly returns to the satisfactory throughput levels for the other applications. The same remark holds when the gamepad is switched off. This experience reflects the recovery property of the Banach-Picard algorithm, mainly when a connected object is added or removed. The transmit power is represented in Fig. 4b. It follows from the figure that in order to achieve a higher requirement, more effort is needed, especially when the channel quality is comparable for all appliances. The step size is ζ = 100 kb/s. The maximum capacity is 2000 kb/s.

Scenario 2: The second scenario deals with progressive capacity discovery. It takes into account two stringent applications: inband D2D communication and online gaming. Both applications aim to improve their QoS by exploring the network capacity.

In order to enhance the quality of the D2D communication and the file upload application, the connected devices can target the maximum achievable throughput by implementing the pro-

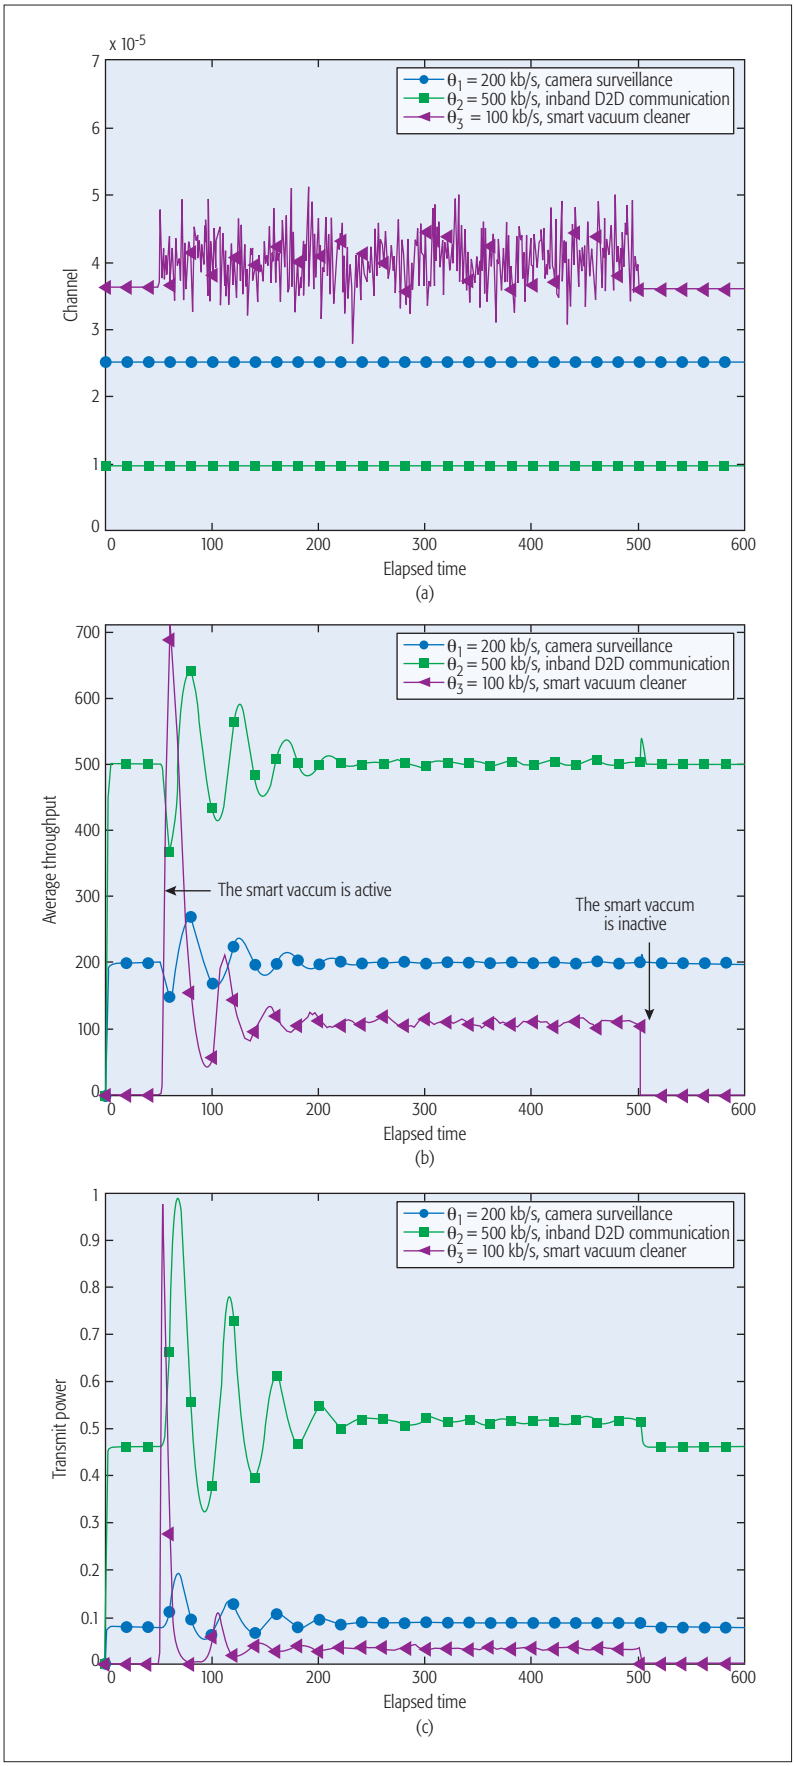

Figure 6. Average throughput, transmit power, and channel state with respect to time for fast-fading channels. The smart vacuum cleaner channel follows an exponential distribution with parameter 1. The network capacity is set to 2000 kb/s.

gressive capacity discovery algorithm. Figure 5 shows the instantaneous throughput and transmit power with respect to time for slow-fading channels. Figure 5a shows that once a user's request is reached, the device increases its demand slightly, and adjusts its power in order to reach the new target (Fig. 5b). When the capacity of the network is reached, devices cannot improve their throughput. Moreover, the figure shows a trade-off between the necessary time and the accuracy to reach the maximum throughput. Clearly, when ζ is large, the time needed to achieve better performance is limited. However, additional demands can be rejected, as shown in Fig. 5a. Unlike when ζ is smaller, the demand adjustment is more accurate, and the time needed to meet the exact maximum throughput is larger. Finally, Fig. 5b highlights a trade-off between the required transmit power and the channel state. Mainly, when the channel state is better, less effort is needed to achieve QoS satisfaction.

Scenario 3: This last scenario considers fast-fading channels. It assumes that the smart vacuum cleaner moves through the house with a target throughput set to 100 kb/s. The smart devices channel gains are represented in Fig. 6a. The smart vacuum cleaner channel follows an exponential distribution with a unit-normalized parameter.

Figure 6b depicts the evolution of the average throughput with respect to time. In order to track the fast varying channel observed by the moving smart vacuum cleaner, the Mann-iterates algorithm is used. The figure shows that the algorithm reaches satisfactory throughput after a few variations around the applications requirements. The transmit power is updated within iterations in order to track the ESE depicted by Fig. 6c.

Open Problems

Although satisfactory solutions can bring important energy savings to the network, they also present many challenges that need to be addressed.

Conditions of existence: Any envisioned satisfaction solution heavily relies on the existence of satisfaction equilibria. Theoretical investigations related to conditions of existence of SEs are still an open research area. The existence of SEs depends on the environment settings, which is very challenging when dealing with random environments. Furthermore, conditions that determine the existence of satisfactory solutions also depend on the studied QoS metrics, the network capacity, and devices parameters, as mentioned earlier.

Latency reduction: Another challenge related to satisfactory solutions' implementation is the time of convergence of distributed algorithms. This is especially important for dense networks where connected objects will take a long time before reaching the optimal solution. Therefore, the reduction of algorithm latency in dense networks still requires particular attention.

Adapted solutions: Other refined solution concepts such as robust SE or long-term SE can be adapted to fast varying contexts, and hence need to be adopted to enhance network performance.

Conclusion

This article presents a novel approach to the management of energy consumption in IoT ecosystems. The idea is to target satisfactory QoS levels

instead of maximizing QoS, which is generally energy consuming. Based on a game theoretical approach, the article introduces the concepts of satisfaction and efficient satisfaction equilibria. The existence and uniqueness of these game theoretical solution concepts are discussed. In order to reach efficient satisfactory solutions, fully distributed schemes adapted for both slow- and fast-fading channels are presented. Finally, the performance of the schemes is illustrated in a smart home use case scenario, and open prolems are discussed.

References

[1] Cisco, "Visual Networking Index," white paper; Cisco.com, Feb. 2016.
[2] A. Sehgal *et al.*, "Management of Resource Constrained Devices in the Internet of Things," *IEEE Commun. Mag.*, vol. 50, no. 12, Dec. 2012, pp. 144–49.
[3] A. Osseiran *et al.*, "Scenarios for Mobile and Wireless Communications: The Vision of the METIS Project," *IEEE Commun. Mag.*, vol. 52, no. 5, Dec. 2014, pp. 26–35.
[4] P. Alvarez et al., "Simulating Dense Small Cell Networks," IEEE WCNC, Doha, Qatar, 3–6 Apr. 2016.
[5] H. Holtkamp *et al.*, "Minimizing Base Station Power Consumption," *IEEE JSAC*, vol. 32, no. 2, May 2014, pp. 297--306.
[6] W. Wu *et al.*, "Energy-Efficient Transmission with Data Sharing," *IEEE INFOCOM*, Hong Kong, 26 Apr.–1 May 2015.
[7] M. Elmachkour *et al.*, "The Greening of Spectrum Sensing: A Minority Game-Based Mechanism Design," *IEEE Commun. Mag.*, vol. 52, no. 12, Dec. 2014, pp. 150–56.
[8] F. Meshkati *et al.*, "An Energy-Efficient Approach to Power Control and Receiver Design in Wireless Data Networks," *IEEE Trans. Commun.*, vol. 53, no. 11, Nov. 2005, pp. 1885--94.
[9] S. M. Perlaza *et al.*, "Quality-of-Service Provisioning in Decentralized Networks: A Satisfaction Equilibrium Approach," *IEEE JSAC*, vol. 6, no. 2, Apr. 2012, pp. 104–16.
[10] F. Meshkati *et al.*, "A Game-Theoretic Approach to Energy-Efficient Power Control in Multicarrier CDMA Systems," *IEEE JSAC*, vol. 24, no. 6, 2006, pp. 1115–29.
[11] E. Sabir *et al.*, "Stochastic Learning Solution for Constrained Nash Equilibrium Throughput in Non Saturated Wireless Collision Channels," *Int'l. ICST Conf. Performance Evaluation Methodologies and Tools*, Pisa, Italy, 20–22 Oct. 2009.
[12] Ericsson, "App Coverage," white paper; ericsson.com, Aug. 2015.
[13] W. R. Mann, "Mean Value Methods in Iteration," *American Mathematical Society*, vol. 4, 1953, pp. 506–10.
[14] H. Elhammouti, E. Sabir, and H. Tembine, "A Satisfactory Power Control for 5G Self-Organizing Networks," arXiv preprint arXiv:1606.07904, 2016.
[15] H. Tembine, *Distributed Strategic Learning for Wireless Engineers*, CRC Press, 2012.

Biographies

Hajar El Hammouti graduated as an engineer from National Telecommunication Institute (INPT), Rabat, Morocco, in 2012. She is currently a Ph.D student at STRS-Laboratory at INPT. Her research interests focus on game theoretical modeling and analysis for self-organized networks.

Essaid Sabir [S'07, M'10, SM'14] received his B.Sc. degree in electrical engineering, electronics, and automation from Mohammed V University, Morocco, in 2004, and his M.Sc. in telecommunications and wireless engineering from National Institute of Post and Telecommunications, Morocco, in 2007. In 2010 he received his Ph.D. degree in networking and computer sciences from the University of Avignon, France. He worked as a lecturer and an assistant professor at the University of Avignon from 2009 to 2012. He is currently a full-time associate professor at the National Higher School of Electricity and Mechanics, Morocco. His current research interests include protocol design for 5G networks, D2D-M2M-IoT and infrastructure-less networking for 5G, cognitive radio, stochastic learning, networking games, pricing, and network neutrality. He serves as a reviewer for prestigious international journals (IEEE, Springer, Elsevier, Wiley, etc.) and is a TPC member for major international conferences (ICC, GLOBECOM, WCNC, etc.). He is a founder and the Vice-Secretary General of the Moroccan Mobile Computing and Intelligent Embedded-Systems Society (Mobitic). As an attempt to bridge the gap between academia and industry, he founded the International Symposium on Ubiquitous Networking (UNet) conference series, co-founded the Casablanca International 5G Summit, and co-founded the International Conference on Wireless Networks and Mobile Communications (WINCOM).

Mustapha Benjillali [S'06, M'09, SM'14] graduated with highest honors as a mobile communications engineer from INPT in 2003. He then joined INRS, Montreal, Canada, where he obtained his M.Sc. and Ph.D. degrees in telecommunications with highest honors in 2005 and 2009, respectively. After a postdoctoral research fellowship with the Electrical Engineering Program at King Abdullah University of Science and Technology (KAUST), Thuwal, Kingdom of Saudi Arabia, he joined the Communication Systems Department at the National Institute of Telecommunications, where he is now an associate professor. His current research interests are in the broad areas of 5G and IoT wireless and mobile communications, green wireless systems, and wireless solutions for smart cities. His focus is on the design of both PHY and MAC layers, closed-form performance analysis, and resource allocation strategies. He is a co-founder and President of the Moroccan Association of Information and Communications Technologies (AMTIC). He is an IBM Academic Initiative member and an active member in IEEE's Green ICT, Smart Cities, and Young Professionals Communities. He serves as a reviewer for many leading international journals, and assumes various TPC Chair and membership roles in many major international conferences. Among his distinctions are a listing as an Exemplary Reviewer of *IEEE Communications Letters* in 2015, a Best Paper Award at IEEE ICC in 2010, a postdoctoral research fellowship from FRQNT Canada in 2009, and the prestigious Alexander Graham Bell Canada Graduate Scholarship (CGS D) from NSERC Canada in 2007.

Loubna Echabbi [S'01, M'05, SM'11] received her Bachelor's degree in applied mathematics from the Faculty of Science and Technology, University Hassan II of Mohammedia, Morocco, in 2000. She then joined PRiSM, University of Versailles, France, where she obtained her M.Sc. and PhD degrees in computer science in 2001 and 2005, respectively. In September 2005, she joined INRIA Lorraine-Nancy as a postdoctoral research fellow. In September 2006, she joined the MASCOTTE project -Sophia Antipolis as a CNRS postdoctoral research fellow. Since December 2006, she is a professor in the computer science department, INPT. Her current research interests include game theory, learning systems, algorithmic design, and applications in telecommunications.

Hamidou Tembine [S'06, M'09, SM] received his M.S. degree in applied mathematics from Ecole Polytechnique, Palaiseau, Paris, France, in 2006 and his Ph.D. degree in computer science from the University of Avignon in 2009. His current research interests include evolutionary games, mean-field stochastic games, and their applications. In December 2014, he received the IEEE ComSoc Outstanding Young Researcher Award for his promising research activities for the benefit of society. He has been the recipient of seven best article awards in the applications of game theory. He is a prolific researcher and has several scientific publications in magazines, letters, journals, and conferences. He is the author of the book *Distributed Strategic Learning for Engineers* (CRC Press, Taylor & Francis 2012), and co-author of the book *Game Theory and Learning in Wireless Networks* (Elsevier Academic Press). He has been co-organizer of several scientific meetings on game theory in networking, wireless communications, smart energy and transportation systems.

Another challenge related to satisfactory solutions implementation is the time of convergence of distributed algorithms. This is especially important for dense networks where connected objects will take a long time before reaching the optimal solution. Therefore, the reduction of algorithms latency in dense networks still requires a particular attention.

Uncoordinated Access Schemes for the IoT: Approaches, Regulations, and Performance

Daniel Zucchetto and Andrea Zanella

The authors provide a comparative overview of the uncoordinated channel access methods for IoT technologies, namely ALOHA-based and LBT schemes, in relation to the ETSI and FCC regulatory frameworks. They also provide a performance comparison of these access schemes, in terms of both successful transmissions and energy efficiency, in a typical IoT deployment.

ABSTRACT

Internet of Things devices communicate using a variety of protocols, differing in many aspects, with the channel access method being one of the most important. Most of the transmission technologies explicitly designed for IoT and machine-to-machine communication use either an ALOHA-based channel access or some type of Listen Before Talk strategy, based on carrier sensing. In this article, we provide a comparative overview of the uncoordinated channel access methods for Internet of Things technologies, namely ALOHA-based and Listen Before Talk schemes, in relation to the ETSI and FCC regulatory frameworks. Furthermore, we provide a performance comparison of these access schemes, in terms of both successful transmissions and energy efficiency, in a typical Internet of Things deployment. Results show that Listen Before Talk is effective in reducing inter-node interference even for long-range transmissions, although the energy efficiency can be lower than that provided by ALOHA methods. Furthermore, the adoption of rate adaptation schemes lowers the energy consumption while improving the fairness among nodes at different distances from the receiver. Coexistence issues are also investigated, showing that in massive deployments Listen Before Talk is severely affected by the presence of ALOHA devices in the same area.

INTRODUCTION

A key element to enable the full realization of the Internet of Things (IoT) vision is the ubiquitous connectivity of end devices, with minimal configuration, such as for the so-called *place-&-play* paradigm [1]. Today, the main three approaches to provide connectivity to the IoT devices are the following.

Cellular systems. The existing cellular networks are a natural and appealing solution to provide connectivity to IoT end devices, thanks to their worldwide established footprint and the capillary market penetration. Unfortunately, current cellular network technologies have been designed targeting wideband services, characterized by a few connections that generate a large amount of data, while most IoT services are expected to generate a relatively small amount of traffic, but from a very large number of different devices. This shift of paradigm challenges the control plan of current cellular standards, which can become the system bottleneck. For these reasons, the IoT and machine-to-machine (M2M) scenarios are considered as major challenges for the next generation of wireless cellular systems, commonly referred to as fifth generation (5G).

Short-range multihop technologies. This family collects a number of popular technologies specifically designed for M2M communications or wireless personal area networks (WPANs). These systems usually operate in the frequency bands centered around 2.4 GHz, 915 MHz, and 868 MHz, although the 2.4 GHz is the most common choice. They are characterized by high energy efficiency and medium/high bit rates (on the order of hundreds of kilobits per second or higher), but limited single-hop coverage area. To cover larger areas, most WPAN technologies provide the possibility to relay data in a multihop fashion, realizing a so-called *mesh network*. Examples of standards in this category are IEEE 802.15.4 [2], Bluetooth Low Energy [3], and Z-Wave, the latter having its physical and data link layers specified in International Telecommunication Union — Telecommunication Standardaization Sector (ITU-T) G.9959 [4].

Low-power wide area (LPWA) networks. A third relevant class in the arena of IoT-enabling wireless technologies consists of the LPWA solutions. According to [5], LPWA technologies will account for 28 percent of M2M connections by 2020. These technologies, specifically designed to support M2M connectivity, provide low bit rates, low energy consumption, and wide geographical coverage. Almost all LPWA technologies operate at frequencies around 800 or 900 MHz, although there are also solutions working in the classic 2.4 GHz industrial, scientific, and medical (ISM) band or exploiting white spaces in TV frequencies. Some relevant LPWA technologies are long-range wide area networking (LoRaWAN), Sigfox, and Ingenu [6].

While cellular systems entail centralized access schemes over dedicated frequency bands, which provide high efficiency, robustness, security, and performance predictability, most WPAN and LPWA technologies operate on unlicensed radio bands, adopting uncoordinated access schemes. The use of unlicensed bands yields the obvious advantage of lowering the operational costs of the network, while the adoption of uncoordinated channel access schemes makes it possible to simplify the hardware of the nodes, thus reducing the manufacturing costs and energy consumption.

Digital Object Identifier:
10.1109/MCOM.2017.1600617

The authors are with the University of Padova

The downside is that the lack of coordination in channel access may yield performance losses in terms of throughput and energy efficiency when the number of contending nodes increases.

To alleviate the problem of channel congestion in the unlicensed bands, radio spectrum regulators have imposed limits on the channel occupation of each device in terms of bandwidth, time, and on the maximum transmission power. However, the Federal Communications Commission (FCC) in the United States and the Conference of Postal and Telecommunications Administrations (CEPT) in Europe have taken different approaches to limit channel congestion: the first imposes very strict limits on the emission power and favors the use of spread spectrum techniques but do not restrict the number of access attempts that can be performed by the nodes [7], while the second limits the fraction of on-air time of a device to be lower than a given *duty cycle*, or imposes the use of Listen Before Talk (LBT) techniques, which are also referred to as carrier sense multiple access (CSMA) protocols [8].[1]

These precautions are actually effective when the coverage range of the wireless transmitters is relatively small (a few meters), as was indeed the case for the first commercial products operating in the ISM frequency bands. However, this condition no longer holds for LPWA solutions, which have coverage ranges on the order of 10–15 km in rural areas and 2–5 km in urban areas, with a star-like topology that can exacerbate the mutual interference and hidden node problems. Furthermore, while short-range communication systems usually support a single or just a few modulation scheme(s) and transmit rate(s), LPWA technologies usually provide multiple transmit rates to optimize the transmission based on the distance to be covered.

Despite these quite radical changes in the transmit characteristics of the recent LPWA technologies with respect to the previous generation of so-called short-range devices (SRDs), the channel access methods and regulatory constraints are still the same. The objective of this study is hence to investigate the performance of well established uncoordinated channel access schemes in this new scenario, characterized by a huge number of devices with large coverage ranges and multirate capabilities. To this end, we first provide a quick overview of the main uncoordinated access schemes used by most common wireless communication technologies considered for the IoT and discuss the regulatory frameworks, with particular focus on the European case. We then compare the performance achieved by two popular uncoordinated access schemes in a typical LPWA network scenario, considering the limits imposed by the regulations. The article is then closed with some final considerations and recommendations.

Uncoordinated Access Techniques for the IoT

Channel access schemes can be roughly divided into two main categories: coordinated and uncoordinated (or contention-based). Coordinated access schemes require time synchronization among the nodes and hence are more suitable for small networks (e.g., Bluetooth) or centrally controlled systems (e.g., cellular), with predictable and/or steady traffic flows (e.g., voice or bulk data transfer). Uncoordinated access strategies, instead, are usually considered for networks with a variable number of devices and unpredictable traffic patterns. In the following, we provide a quick overview of the two main uncoordinated access schemes that are widely adopted by the transmission technologies typically associated with IoT scenarios.

ALOHA-based Schemes

Many protocols for M2M communication are based on pure ALOHA access schemes, according to which a transmission is attempted whenever a new message is generated by the device. This form of channel access may be coupled with a retransmission scheme, according to which a packet is retransmitted until acknowledged by the receiver. However, some IoT services (e.g., environmental monitoring) can tolerate a certain amount of lost messages. In these cases, a retransmission scheme is not needed, allowing for simplification of the device firmware and enabling a significant reduction in energy consumption. For these reasons, ALOHA schemes are widely adopted in M2M communication, such as LoRaWAN and Sigfox. Furthermore, some standards that adopt LBT access techniques optionally provide an ALOHA mode of operation, such as for IEEE 802.15.4.

More sophisticated ALOHA-based protocols can be enabled when nodes are time synchronized, for example, by means of beacons periodically broadcasted by coordinator nodes (e.g., gateways in LoRaWAN). For example, slotted ALOHA divides the time in intervals of equal size, called slots, and allows transmissions only within slots, thus avoiding packet losses due to partially overlapping transmissions. Framed slotted ALOHA (FSA) instead organizes the slots in groups, called frames, and allows each node to transmit only once per frame. The limit of these schemes is that packet transmission time should not exceed the slot duration. A common solution to accommodate uneven packet transmission times is to adopt a hybrid access scheme (HYB) that splits the frame in two parts: the first k slots are used by the nodes to send resource reservation messages to the controller, using an FSA access scheme, while the remaining slots in the frame are allocated by the controller to the nodes, according to the amount of resources required in the accepted reservation messages. The nodes are notified about the allocated resources by a control message that is broadcasted by the controller right after the end of the reservation phase. Variants of these basic mechanisms are currently used in many different protocols (e.g., GSM, 802.11e). However, to the best of our knowledge, the HYB approach has not yet been studied in the M2M scenario.

Carrier Sensing Schemes

When using carrier sensing techniques, each device listens to the channel before transmitting (hence the term Listen Before Talk). The channel sensing operation is typically called clear channel assessment (CCA) and aims to check the occupancy of the channel by other transmitters, in

Channel access schemes can be roughly divided in two main categories: coordinated and uncoordinated. Coordinated access schemes require time synchronization among the nodes and, hence, are more suitable for small networks. Uncoordinated access strategies, instead, are usually considered for networks with a variable number of devices and unpredictable traffic patterns.

[1] The two terms will be used interchangeably in this article.

The use of unlicensed frequency bands by radio emitters is subject to regulations that are intended to favor the coexistence of a multitude of heterogeneous radio transceivers in the same frequency bands, limiting the mutual interference and avoiding any monopolization of the spectrum by single devices.

which case the channel access will be delayed to avoid mutual interference, which may result in so-called *packet collisions*. The LBT schemes can differ in the way the CCA is performed and in the adopted behavior in case the channel is sensed as busy.

The three most common methods to perform the CCA are the following.

- Energy detection (ED). The channel is detected as busy if the electromagnetic energy on the channel in the operational bandwidth is above a given ED threshold.
- Carrier sense (CS). The channel is reported as busy if the device detects a signal with modulation and spreading characteristics compatible with those used for transmission, irrespective of the signal energy.
- Carrier sense with energy detection (CS+ED). In this case, a logical combination of the above methods is used, where the logical operator can be AND or OR.

The IEEE 802.15.4 standard supports all these CCA methods, along with pure ALOHA and two other modes specific for ultra-wideband communications. In an unslotted system, the backoff procedure for the IEEE 802.15.4 CCA mechanism tries to adapt to the channel congestion by limiting the rate at which subsequent CCAs are performed for the same message. If the number of consecutive backoffs exceeds a given threshold, the message is discarded. Details about the CCA procedure in IEEE 802.15.4 networks can be found in [2], together with recommendations about the ED threshold and CCA detection time.

The Regulatory Framework

The use of unlicensed frequency bands by radio emitters is subject to regulations that are intended to favor the coexistence of a multitude of heterogeneous radio transceivers in the same frequency bands, limiting the mutual interference and avoiding any monopolization of the spectrum by single devices. The radio emitters operating in the ISM frequency bands are typically referred to as SRDs. However, ERC Recommendation 70-03, emanated from the CEPT, specifies that "*The term Short Range Device (SRD) is intended to cover the radio transmitters which provide either uni-directional or bi-directional communication which have low capability of causing interference to other radio equipment.*" Despite the name, there is no explicit mention of the actual coverage range of such technologies. Therefore, long-range technologies operating in the ISM bands, such as Sigfox and LoRa, are still subject to the same regulatory constraints that apply to the actual short-range technologies, including IEEE 802.15.4, Bluetooth, IEEE 802.11, and so on.

In the European Union, the European Commission designated the CEPT to define technical harmonization directives for the use of the radio spectrum. In 1988, under the patronage of the CEPT, the European Telecommunications Standards Institute (ETSI) was created to develop and maintain harmonized standards for telecommunications.

In the unlicensed radio spectrum at 868 MHz, ETSI mandates a duty cycle limit between 0.1 and 1 percent over a 1 hour interval for devices that do not adopt LBT [8]. Only very specific applications, such as wireless audio, are allowed to ignore the duty cycle limitation. The duty cycle constraint can be relaxed by employing an LBT access scheme together with adaptive frequency agility (AFA), that is, the ability to dynamically change channel [8]. Devices with LBT and AFA capabilities, in fact, are only subject to a 2.8 percent duty cycle limitation for any 200 kHz spectrum. An example of technology that adopts the LBT approach is IEEE 802.15.4, which, however, does not perfectly match the ETSI specifications, since its channel sensing period is shorter than that mandated by ETSI, which is between 5 ms and 10 ms, depending on the used bandwidth [8]. Instead, the recommendations on LBT sensitivity, which shall be between –102 and –82 dBm, are usually satisfied by commercial transceivers. Due to the adoption by the European Union of a new set of rules for radio equipment, the Radio Equipment Directive (RED) [9], ETSI is reviewing the related harmonized standards. However, devices that are compliant with the previous Radio and Telecommunication Terminal Equipment (R&TTE) Directive [10] could be placed on the market until June 17, 2017. Furthermore, devices that do not satisfy the constraints imposed by the harmonized standards can still be commercialized, but subject to a more comprehensive certification procedure attesting that the device meets the essential requirements of the European Directives [9]. The latest draft version of the ETSI harmonized standards [11] includes some changes in the medium access procedures. In particular, the LBT technique is generalized as a *polite spectrum access* technique, while AFA is no longer required. Furthermore, the LBT ED threshold has been relaxed, while the minimum CCA listening period has been increased.

The agency designated to regulate radio communications in the United States is the FCC, which also grants permits for the use of licensed radio spectrum and emanates regulations for wired communications. The FCC regulation does not impose any duty cycle restrictions to emitters operating in the 902–928 MHz band, but limits the maximum transmit power, for non-frequency hopping systems to –1.25 dBm [7], which is significantly lower than the 14 dBm allowed by ETSI.

Performance Analysis

ALOHA schemes and channel sensing techniques have been comprehensively modeled, and their performance limits in terms of throughput and capacity are well understood (see, e.g., [12, 13], just to cite a couple). However, the use of different spreading techniques, and/or modulation and coding schemes to cope with the interference and to trade transmission speed for reliability, the large coverage range enabled by the LPWA technologies, the total reuse of the same frequency bands by different technologies, and the limitations imposed by the regulations of the channel access raise the question of how effective the classical uncoordinated channel access techniques are to adequately support the expected growth of the IoT services.

In this section we shed some light on these aspects by presenting a simulation analysis of the performance achieved by ALOHA-based (specifically, pure ALOHA and HYB) and LBT access

schemes in the simplest IoT scenario, sketched in Fig. 1: a gateway (GW) receiving packets from a multitude of peripheral devices randomly spread over a wide area. Despite its simplicity, this scenario embodies most of the problems that can be expected in a real IoT deployment based on long-range technologies. In particular, we are interested in investigating how the distance from the gateway may impact the performance experienced by the node, with and without multirate capability and using either ALOHA or LBT techniques. ALOHA-based access schemes, in fact, allow the maximum energy saving in light traffic conditions, since they avoid the (even small) energy cost involved in carrier sensing. On the other hand, nodes farther away from the GW are likely to be more prone to transmission failure due to interference, which, however, can potentially be mitigated by the use of LBT. Furthermore, the adoption of rate adaptation techniques is expected to increase the system capacity by reducing the transmit time of nodes closer to the GW that not only will experience a lower interference probability, but will also have the chance to transmit more packets within the duty cycle limitations. It is hence interesting to investigate how much of such a performance gain will be transferred to the more peripheral nodes, and whether the LBT techniques can further improve performance in a significant manner.

Simulation Scenario

In our simulations we consider a propagation model given by the product of the channel gain, $\gamma(d) = (Ad)^{-\beta}$, which accounts for the power decay with the distance d from the transmitter through the model parameters A and β, and the Rayleigh fading gain, which is modeled as an exponential random variable with unit mean.

We consider a limited set of possible transmission rates, namely $\mathcal{R} = \{0.5, 1, 5, 10, 50, 100\}$ kbit/s, and assume that a packet transmitted at rate $r \in \mathcal{R}$ is correctly decoded if the received signal energy over the total noise energy plus interference energy collected by the receiver during the packet reception time (i.e., the signal-to-interference-and-noise ratio, SINR) is above a certain threshold $\Gamma_{th}(r)$, which is determined from the Shannon channel capacity as

$$\Gamma_{th}(r) = 2^{r/W} - 1 \quad (1)$$

where W is the signal bandwidth.

For the single rate (SR) case, we suppose that all nodes transmit with the lowest bit rate of 500 b/s. For the multirate scenario, instead, we consider a simple rate-adaptation mechanism that keeps a moving-average estimate of the SINR (using a smoothing factor α) and selects the rate R so that the expected outage probability is no larger than $p^* = 0.05$. To improve energy efficiency, furthermore, we assume that no acknowledgment or retransmission mechanism is implemented, so packets that are not successfully received are definitely lost.

The LBT scheme has been implemented based on the IEEE 802.15.4 specifications. The ED CCA threshold has been chosen to match the minimum signal power required to correctly receive a packet transmitted at the basic rate of 500 b/s.

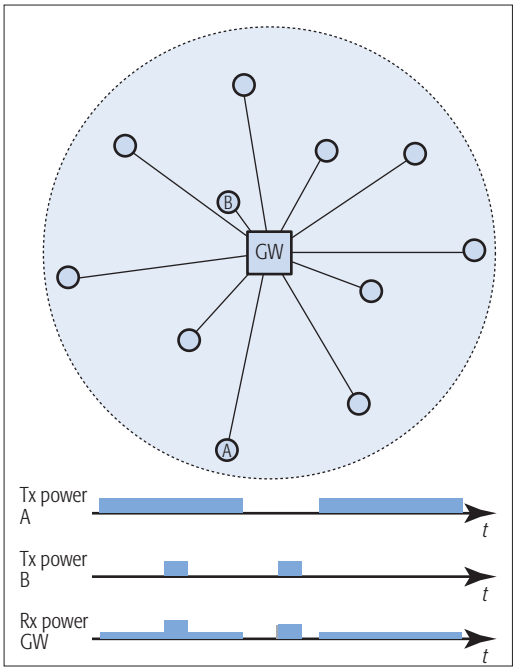

Figure 1. Above: the simulation scenario, with multiple transmitters scattered around the common receiver (GW). Below: an example of signal transmissions by nodes A and B, using different bit rates, and of received signal power at the gateway.

This value is compatible with the limits on the LBT threshold imposed by ETSI [8].

As exemplified in Fig. 1, transmitting nodes are distributed as for a spatial Poisson process of rate λ_s [devices/m^2] over a circle with radius equal to the maximum coverage distance at the basic rate of 500 b/s. Each device generates messages of length L according to a Poisson process of rate λ_t (packets per second). All messages are addressed to the GW, which is placed at the center of the circle.

The setting of all the simulation parameters is reported in Table 1.

Transmission Failure Probability

We define p_{fail} as the probability that a transmitted message (including reservation messages in the case of HYB) is received with SINR below threshold and hence is not correctly decoded. For HYB we also include in the p_{fail} the transmission requests that are not accepted because of lack of slots in the transmission part of the frame. Note that while we consider both the single rate (SR) and rate adaptation (RA) versions of the pure ALOHA and LBT schemes, for the HYB protocol we only consider the RA version, since this access scheme is more effective when packet transmissions have uneven duration. In Fig. 2 we report the failure probability for target nodes placed at increasing distances from the gateway. Red curves with circle markers refer to ALOHA, blue plain curves to LBT, and the green dashed line with diamond markers to HYB. Solid and dashed lines are associated wtih the SR and RA case, respectively.

For the SR case, we can see that the failure probability grows with the distance from the GW, since nodes farther away have less SINR mar-

The adoption of rate adaptation techniques is expected to increase the system capacity by reducing the transmit time of nodes closer to the gateway that not only will experience a lower interference probability, but will also have the chance to transmit more packets within the duty cycle limitations.

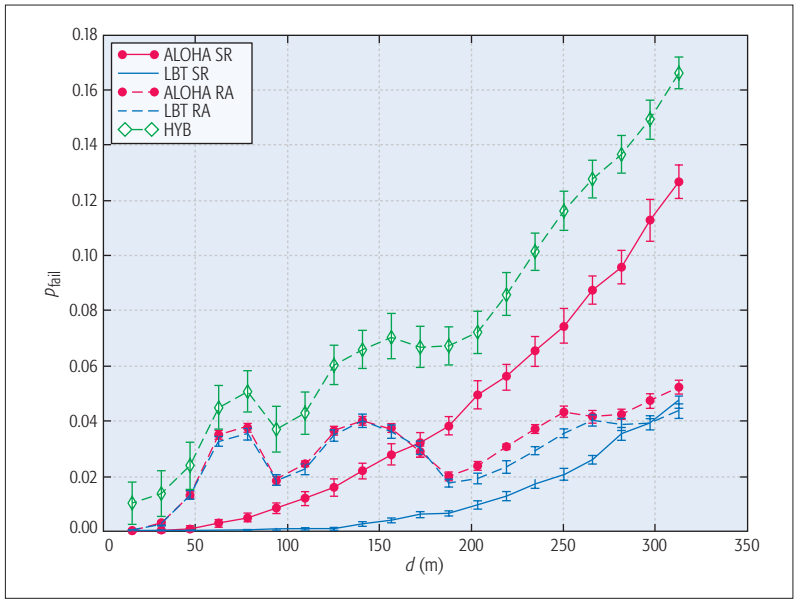

Figure 2. p_{fail} for ALOHA and LBT, for single rate and rate adaptive cases, with 95 percent confidence intervals.

gin for successful decoding and are hence less robust to the interference produced by overlapping transmissions. In this case, carrier sense can indeed improve performance, even if the sensing range does not prevent the hidden node problem.

The downside of using LBT (not reported here for space constraints) is that up to 55 percent of the transmission attempts are aborted in high traffic conditions, because the maximum number of CCAs is reached without finding an idle channel.

The adoption of RA changes significantly the performance, smoothing out the differences between the two access protocols. Indeed, higher bitrates allow the nodes near the receiver to occupy the channel for a lower period of time, thus reducing the probability of overlapping with other transmissions and improving the performance of both access schemes.

Note that the change of rate with the distance is reflected by the oscillation in the failure probability that, however, remains approximately below $1 - p^*$.

Rather interestingly, HYB performs worse than the other schemes. The reason is that in the considered scenario, the transmit time of reservation messages, always sent at the basic rate, is comparable to that of data packets sent at higher rates. Therefore, the reservation channel can become the system bottleneck. The overall channel occupancy of HYB is thus significantly higher than that of the other two schemes, yielding higher failure probability.

Energy Efficiency

Another key performance index in the IoT scenario is the *energy efficiency*, which is here defined as the ratio of the total number of bits successfully delivered to the gateway over all the energy consumed by the node (including channel sensing and failed transmissions).

We modeled the power consumed during a transmission as the sum of a constant term, named circuit power, which represents the power used by the radio circuitry, and a term that accounts for the radiated power, which is called transmission power. When using LBT, we also add the power required to perform the ED CCA. Referring to the data sheets of some off-the-shelf modules,[2] we set the circuit power to 16 dBm, the transmit power to 14 dBm, the receive power to 13 dBm, and the CCA power to 10 dBm [14, 15].

Parameter		Value
Spatial node density	λ_s	10^{-3} nodes/m^2
Packet generation rate	λ_t	0.01 packets/s
Transmission power	P_{TX}	14 dBm
Transmission frequency	f	868 MHz
Path loss coefficient	A	36.36 m^{-1}
Path loss exponent	β	3.5
Packet length	L	240 bit
Transmission bit rates	$\mathcal{R}$	{0.5, …,100 kb/s}
Bandwidth	B_W	400 kHz
Noise spectral density	N_0	$2 \cdot 10^{-20}$ W/Hz
Duty cycle	δ_T	1%
Circuit power	P_c	16 dBm
Sensing time	T_s	0.4 ms
Sensing energy	E_s	3.98 µJ (LBT) 0.2 mJ (LBT+ETSI)
Smoothing parameter	α	0.1
Target outage probability for RA	p^*	0.05
HYB parameters		
Frame duration	T_W	60 s
Number of reservation slots in a frame	N_{RM}	80
Reservation message size	L_{RM}	24 bits
Reservation message transmit rate	R_{RM}	500 bit/s
Beacon duration	T_B	0.12 s
Resource notification message duration	T_{RA}	3.84 s

Table 1. Simulation parameters.

In Fig. 3a we show the energy efficiency for ALOHA and LBT access schemes when varying the distance of the target node from the gateway, in the SR case. We can observe that peripheral nodes exhibit lower energy efficiency because of the larger number of failure transmissions, and that the carrier sensing mechanism can alleviate this problem. The black curve marked with crosses shows the results obtained when using the parameters imposed by ETSI in the CCA procedure. As can be seen, the energy efficiency is slightly lower than that obtained with the parameters adopted by commercial technologies, which may suggest that ETSI recommendations in this regard are possibly too conservative.

The adaptive rate case is shown in Fig. 3b, where we also show the performance achieved by HYB. We can observe that both ALOHA and LBT can reach very high efficiency for nodes near the receiver, since the higher bit rates decrease the transmit energy and the failure probability. It is worth noting that the first factor is dominant for energy efficiency. The benefit transfers to the nodes farther away from the GW, although the performance gain progressively reduces with the distance from the transmitter.

We also observe that for nodes closer to the GW, LBT shows a non-negligible energy efficiency loss with respect to ALOHA, which is even more marked when adopting the ETSI parameters. This is clearly due to the energy cost of the carrier sense mechanism, which takes a time comparable with the packet transmission time when using high

[2] Atmel AT86RF212B, Texas Instruments CC1125 and CC1310, and Semtech SX1272 modules.

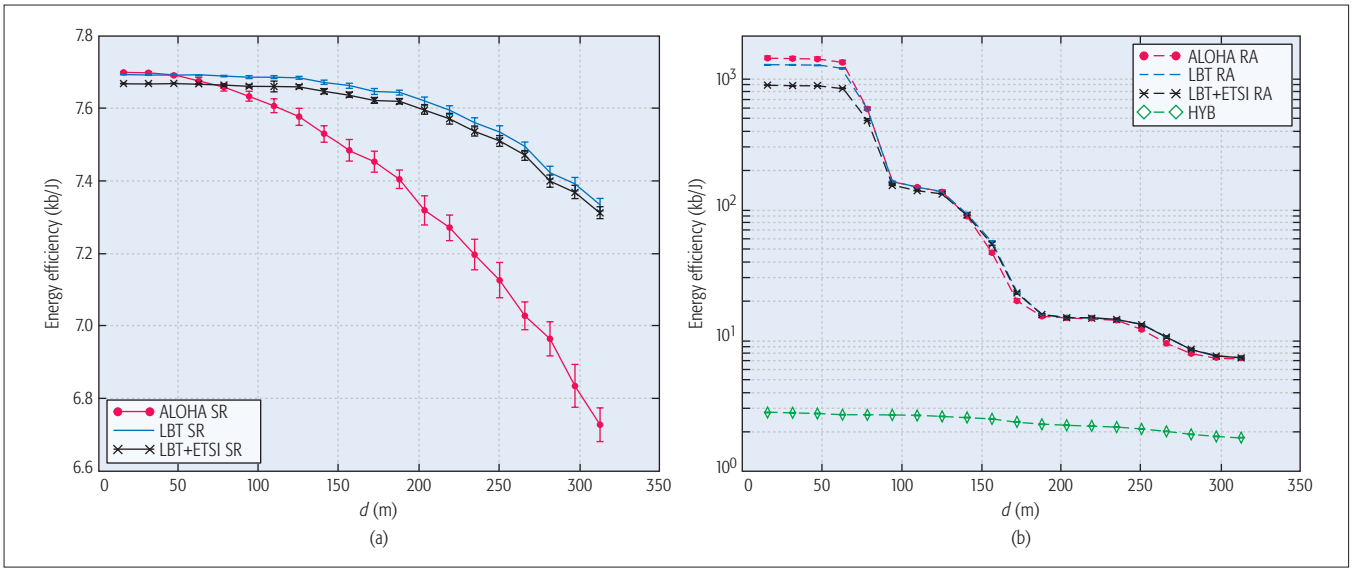

Figure 3. Successfully received bits per unit of consumed energy, with 95 percent confidence intervals: a) single rate case; b) rate adaptive case.

bit rates. Furthermore, as revealed by the analysis of the failure probability, the carrier sense mechanism is not really worthwhile for nodes close to the GW when using RA, also considering that it may yield packet drops due to the impossibility of finding the channel idle within the maximum number of carrier sensing attempts. This problem would be further exacerbated in the case of overlapping cells. Therefore, the use of CCA appears to be fruitless, if not detrimental, for nodes close to the gateway when RA is enabled.

Finally, we observe that the energy efficiency of HYB is the worst, being affected by both the higher failure probability observed in Fig. 2 and the higher energy consumption due to the transmission of resource messages and the reception of beacons. This inefficiency is more marked for nodes near the receiver, where the energy spent on control messages is actually greater than that used for the high-rate transmissions of small data packets.

Coexistence Issues

Another important question regards the coexistence in the same area of nodes using LBT and ALOHA access schemes.

Figures 4 and 5 report the throughput of the two access methods, defined as the overall rate of successful packet transmissions and the energy efficiency. Curves for ALOHA (respectively LBT) have been obtained by fixing the spatial density of this type of node to 0.001 nodes/m^2 and increasing the spatial density of LBT (respectively ALOHA) nodes from 10^{-5} to 10^{-2} nodes/m^2.

Results in Fig. 4 show that the performance of ALOHA nodes is not impacted by an increase in the number of LBT nodes, while the latter suffer strong performance degradation due to the CCA mechanism that aborts a transmission attempt when the channel is sensed busy for a given number of successive attempts. We can also see that the use of multiple transmission rates can only slightly alleviate the problem, but the fragility of the LBT mechanism in the presence of ALOHA traffic still remains. Similar observations can be

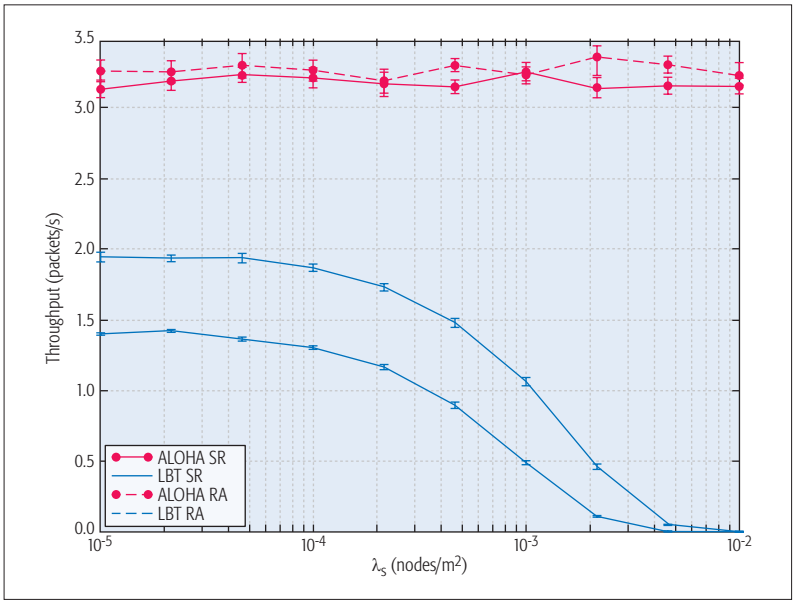

Figure 4. Aggregated throughput for each channel access method in the single and adaptive rate scenarios, with 95 percent confidence intervals.

drawn for the energy efficiency results. In both cases, the use of RA improves energy efficiency quite significantly.

Conclusions

In this work, we present an overview of the three main uncoordinated channel access sensing schemes, pure ALOHA, HYB, and LBT, in an IoT scenario. We compare the performance of these schemes in terms of probability of successful transmission and energy efficiency, by considering the duty cycle limitation for ALOHA, the control packets for HYB, and the CCA procedure for LBT as mandated by the international regulation frameworks.

From this analysis, it appears clear that adding rate adaptation capabilities is pivotal to maintaining a reasonable level of performance when the coverage range and cell load increase. More-

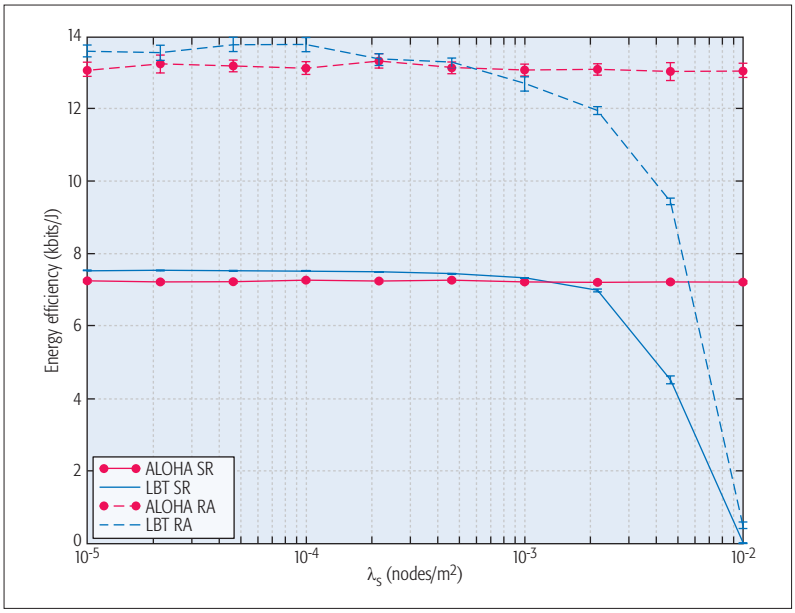

Figure 5. Succesfully received bits per unit of consumed energy for each channel access method, in the single and adaptive rate scenarios, with 95 percent confidence intervals.

over, we observe that LBT generally yields lower transmission failure probability, although packet dropping events may occur because the channel is sensed busy for a certain number of consecutive CCA attempts. This impacts the actual energy efficiency of the LBT access scheme, which may turn out to be even smaller than that achieved by ALOHA schemes. Furthermore, we also observe that LBT performance undergoes severe degradation when increasing the number of ALOHA devices in the same cell, again because of the channel blockage effect caused by the other transmitters. Finally, the HYB scheme proves ineffective in the considered scenario, since the reservation channel becomes the system bottleneck with short data packets. Nonetheless, hybrid solutions that adopt LBT for peripheral nodes and ALOHA for nodes closer to the receiver, or also apply rate adaptation to the reservation phase, can potentially lead to a general performance improvement of the system. This analysis, however, is left to future work.

References

[1] A. Biral *et al.*, "The Challenges of M2M Massive Access in Wireless Cellular Networks," *Digital Commun. and Networks*, vol. 1, no. 1, Feb. 2015, pp. 1–19.

[2] IEEE 802.15.4-2015, "Low-Rate Wireless Networks," Apr. 2016; http://ieeexplore.ieee.org/servlet/opac?punumber=7460873, accessed Nov. 28, 2016.

[3] "Bluetooth Core Specification 4.2," Bluetooth SIG, Dec. 2014; https://www.bluetooth.org/DocMan/handlers/DownloadDoc.ashx?doc id=286439, accessed Nov. 28, 2016.

[4] ITU-T Rec. G.9959, "Short Range Narrow-Band Digital Radiocommunication Transceivers – PHY, MAC, SAR and LLC Layer Specifications," Jan. 2015; http://handle.itu.int/11.1002/1000/12399, accessed Nov. 28, 2016.

[5] "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2015–2020," white paper, Feb. 2016; http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/mobile-white-paper-c11-520862.html, accessed Nov. 15, 2016.

[6] M. Centenaro *et al.*, "Long-Range Communications in Unlicensed Bands: the Rising Stars In the IoT and Smart City Scenarios," *IEEE Wireless Commun.*, vol. 23, no. 5, Oct. 2016, pp. 60–67.

[7] FCC, Code of Federal Regulations, Title 47, Ch. I, Part 15; http://www.ecfr.gov/cgi-bin/text-idx?node=pt47.1.15, accessed Nov. 10, 2016.

[8] ETSI EN 300 220, "Electromagnetic Compatibility and Radio Spectrum Matters (ERM); Short Range Devices (SRD); Radio Equipment to Be Used in the 25 MHz to 1000 MHz Frequency Range with Power Levels Ranging up to 500 mW," Rev. 2.4.1, May 2012; http://www.etsi.org/deliver/etsi en/300200/300299/30022001/02.04.01 60/en 30022001v020401p.pdf, accessed Nov. 28, 2016.

[9] Directive 2014/53/EU, "Radio Equipment Directive," Apr. 2014; http://data.europa.eu/eli/dir/2014/53/oj, accessed Nov. 28, 2016.

[10] Directive 1999/5/EC, "Radio and Telecommunications Terminal Equipment," Mar. 1999; http://data.europa.eu/eli/dir/1999/5/oj, accessed Nov. 28, 2016.

[11] ETSI EN 300 220, Short Range Devices (SRD) Operating in the Frequency Range 25 MHz to 1000 MHz," ETSI Draft Euro. Std., Rev. 3.1.0, May 2016; http://www.etsi.org/deliver/etsi en/300200 300299/30022002/03.01.01 30/en 30022002v030101v.pdf, accessed Nov. 28, 2016.

[12] B. Błaszczyszyn, P. Mühlethaler, and S. Banaouas, "Comparison of Aloha and CSMA in Wireless Ad-Hoc Networks Under Different Channel Conditions," *Wireless Ad-Hoc Networks*, H. Zhou, Ed., InTech, 2012, pp. 3–22.

[13] M. Kaynia and N. Jindal, "Performance of ALOHA and CSMA in spatially distributed wireless networks," *Proc. 2008 IEEE ICC*, May 2008, pp. 1108–12.

[14] I. Ramachandran and S. Roy, "On the Impact of Clear Channel Assessment on MAC Performance," *Proc. IEEE GLOBECOM 2006*, Nov. 2006, pp. 1–5.

[15] L. Negri *et al.*, "Flexible Power Modeling for Wireless Systems: Power Modeling and Optimization of Two Bluetooth Implementations," *Proc. 6th IEEE Int'l. Symp.: A World of Wireless Mobile and Multimedia Networks*, June 2005, pp. 408–16.

Biographies

Daniel Zucchetto (zucchett@dei.unipd.it) received his Bachelor's degree in information engineering in 2012 and his Master's degree in telecommunication engineering in 2014, both from the University of Padova, Italy. Since October 2015 he has been a Ph.D. student in the Department of Information Engineering at the University of Padova. His research interests include low-power wide area network technologies and next generation cellular networks (5G), with particular focus on their application to the Internet of Things.

Andrea Zanella [S'98, M'01, SM'13] (zanella@dei.unipd.it) is an associate professor at the University of Padova. He has authored more than 130 papers, four books chapters, and three international patents on multiple subjects related to wireless networking and the Internet of Things. Moreover, he serves as an Editor for many journals, including the *IEEE Internet of Things Journal* and *IEEE Transactions on Cognitive Communications and Networking*.

Massive Non-Orthogonal Multiple Access for Cellular IoT: Potentials and Limitations

Mahyar Shirvanimoghaddam, Mischa Dohler, and Sarah J. Johnson

ABSTRACT

The Internet of Things promises ubiquitous connectivity of everything everywhere, which represents the biggest technology trend in the years to come. It is expected that by 2020 over 25 billion devices will be connected to cellular networks; far beyond the number of devices in current wireless networks. Machine-to-machine communications aims to provide the communication infrastructure for enabling IoT by facilitating the billions of multi-role devices to communicate with each other and with the underlying data transport infrastructure without, or with little, human intervention. Providing this infrastructure will require a dramatic shift from the current protocols mostly designed for human-to-human applications. This article reviews recent 3GPP solutions for enabling massive cellular IoT and investigates the random access strategies for M2M communications, which shows that cellular networks must evolve to handle the new ways in which devices will connect and communicate with the system. A massive non-orthogonal multiple access technique is then presented as a promising solution to support a massive number of IoT devices in cellular networks, where we also identify its practical challenges and future research directions.

INTRODUCTION

The Internet of Things (IoT) is one of the biggest technology trends, aimed at transforming every physical object into an information source. IoT use cases can be generally divided into two large categories. In massive IoT applications, sensors typically report to the cloud on a regular basis, and the requirement is for low-cost devices with low energy consumption and good coverage. Examples include smart buildings, logistics, tracking, and fleet management. In critical IoT use cases, there are high demands for reliability, availability, and low latency. Critical IoT includes remote health care, traffic safety and control, industrial applications and control, remote manufacturing, training, and surgery. The continued fall in the price, size, and power consumption of autonomous devices capable of sensing and actuating have been the driving force for the increasing popularity of IoT systems and services.

Ericsson forecasted that the IoT will include over 25 billion units installed by 2020, and a large share of these will be applications serviced by short-range radio technologies, such as WiFi, Bluetooth, and Zigbee with limited quality of service (QoS) and security requirements, typically applicable for indoor environments, while a significant proportion will be enabled by wide area networks mostly facilitated by cellular networks [1]. By 2020, IoT product and service suppliers will also generate incremental revenue exceeding $300 billion, mostly in services [1]. As operators are responsible for wireless connectivity on a global scale, they are in an excellent position to participate in the IoT market and capture a share of the added value generated by the emerging IoT applications.

Machine-to-machine (M2M) communications refer to automated data communications among machine type communication (MTC) devices and constitutes the basic communication paradigm in the emerging IoT [2]. The IoT reference model [3] is shown in Fig. 1, which shows that connectivity is the essential part of the entire IoT ecosystem, which can be provided through wired or wireless solutions. Connectivity through cellular networks is facilitated through the Third Generation Partnership Project (3GPP) technologies, including GSM, wideband code-division multiple access (WCDMA), LTE, and the future fifth generation (5G). These technologies operate on licensed spectrum and are primarily designed for high-quality mobile voice and data services. They are now being evolved with new functionality to form attractive solutions for emerging low-power IoT application. 3GPP technologies already dominate many IoT applications that require large geographical coverage and medium-to-high performance requirements [1]. The key challenges for massive IoT deployment in cellular networks include:

- Device cost is the enabler for high-volume mass market applications.
- Battery life must be addressed to reduce the cost of replacing batteries.
- Coverage (deep indoor connectivity and regional coverage) is a requirement for many IoT applications.
- Scalability: The network capacity must be easily scaled to handle millions of devices.
- Diversity: Connectivity should support diverse range of service requirements; for example, alarm signal applications require highly reliable communication with QoS guarantees, while in smart metering applications delays are typically tolerable [3].

The authors review recent 3GPP solutions for enabling massive cellular IoT and investigate the random access strategies for M2M communications, which shows that cellular networks must evolve to handle the new ways in which devices will connect and communicate with the system. A massive non-orthogonal multiple access technique is then presented as a promising solution to support a massive number of IoT devices in cellular networks.

Mahyar Shirvanimoghaddam is with the University of Sydney; Sarah J. Johnson is with the University of Newcastle, Australia; Mischa Dohler is with King's College London.

Digital Object Identifier: 10.1109/MCOM.2017.1600618

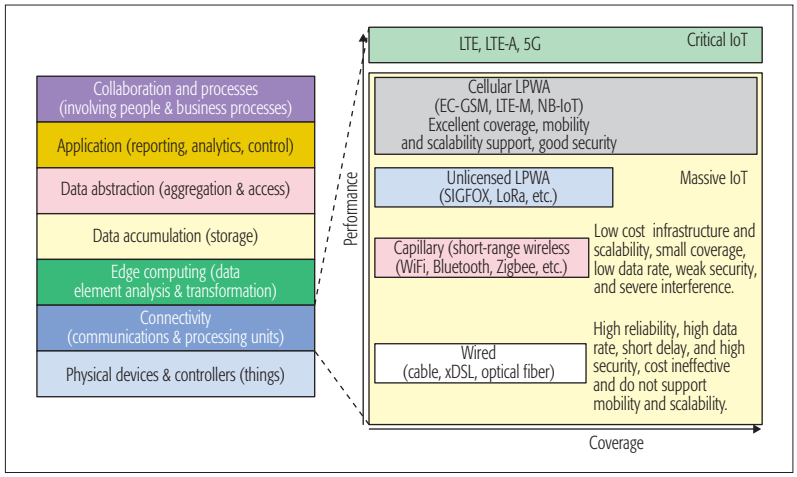

Figure 1. IoT reference model.

3GPP has made significant improvements to meet the requirements of emerging massive IoT applications, which leads to a range of cellular low-power wide area solutions. These new solutions include extended coverage GSM (EC-GSM), which is achieved by defining new control and data channels mapped over legacy GSM; narrowband IoT (NB-IoT), which is a self-contained carrier that can be deployed with a system bandwidth of 200 kHz and is enabled on an existing LTE network; and LTE for MTC (LTE-M), which brings new power-saving functionality to LTE suitable for M2M applications. The most important improvements 3GPP has made so far to enable massive IoT are:

- Lower device cost by reducing peak rate, memory requirements, and device complexity
- Improved battery life up to 10 years by introducing power saving mode and discontinuous reception
- Improved coverage (e.g., 15 dB and 20 dB in link budget on LTE-M and NB-IoT, respectively [4]), which is equivalent to the signal penetrating a wall or floor, enabling deeper indoor coverage

Despite huge efforts of 3GPP toward making M2M communications a reality, there are still open challenges where efficient solutions must be proposed.

As part of massive IoT support, a very large number of MTC devices, a few million devices per square kilometer, must be supported. Toward this, non-orthogonal multiple access (NOMA) has been identified as a key technique in 5G which can enable trillions of MTC devices to communicate with the base station in cellular IoT use cases [3]. Also mentioned in [4], grant-free uplink through resource spread multiple access enables asynchronous, non-orthogonal contention-based access that is well suited for sporadic uplink transmissions of small data bursts common in IoT use cases which are considered as new capabilities for the massive IoT in 5G. The presented NOMA strategy in this article shows how NOMA can be used on top of existing cellular infrastructure to enable massive numbers of devices to share the same radio resources, thus enabling massive IoT applications.

Current Access Techniques

In most existing wireless networks, radio resources, including time and frequency, are orthogonally allocated to different devices for data transmission. The process of devices contacting the base station to request a transmission slot is called the random access (RA) procedure. In the current LTE standard, RA is a four-step handshake process, which is depicted in Fig. 2.

The Random Access Procedure

The devices are first informed of the available physical random access channel (PRACH) resources, comprising a periodic amount of time-frequency resources, through the system information broadcasted by the base station. In the first step of the RA procedure, each device randomly chooses a preamble among an available set of 64 preambles and sends it to the base station. The base station can then detect the transmitted preambles by calculating the cyclic cross-correlation of the set of preambles with the received signal and, using the result, can estimate the transmission time of devices that have selected each preamble. Upon detecting each preamble, the base station sends a random access response (RAR) message, including the information of the radio resources allocated to devices and the timing advance information for all the devices that have selected a specific preamble, to adjust synchronization. Once a device has received the RAR message, it sends a temporary terminal identity through the allocated radio resource to the base station to request a connection. The base station sends information allocating radio resources to each of the devices that have gained access by specifying their terminal identity. Therefore, a connection is established, and the device can start the data transmission in allocated time-frequency slots.

Challenges of Conventional Random Access for Massive Cellular IoT

The random access procedure in current cellular standards is only feasible when the number of devices is small enough that:

- The devices do not unduly interfere with each other in the random access phase.
- There are sufficient radio resources to allocate a separate data channel to each user in the transmission phase.

In the following we outline the most significant challenges of the conventional RA procedure for massive cellular IoT.

Preamble Collision and Overload Problems: Two conditions may arise in the random access procedure that dramatically limit its efficiency:

1. The condition known as preamble collision, where more than one device selects the same preamble in the first step of the RA procedure, causing co-channel interference in the transmission stage.
2. The condition known as overload, which is due to the signal-to-interference-plus-noise ratio (SINR) violation caused by an excessive number of transmissions by other nodes in the same cell or in neighboring cells during the RA procedure.

In the case of collision, the devices will repeat

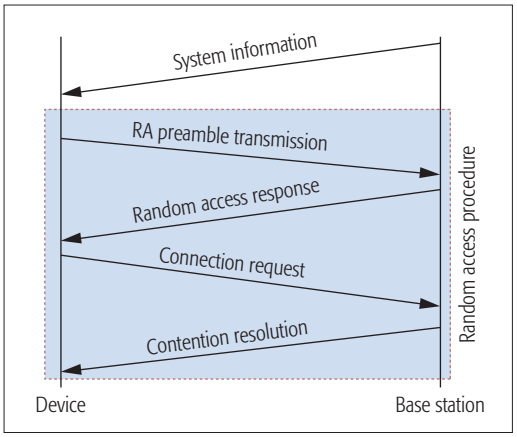

Figure 2. Random access procedure in LTE.

the preamble transmission in the next available RA resource. Frequent preamble retransmissions, however, lead to network congestion, increasing delays, packet loss, high energy consumption, excessive signaling overhead, and radio resource waste [5]. Existing solutions to solve the collision problem use different approaches to delay the retransmission of preambles in order to minimize the collision probability. These include dynamic allocation [6], slotted access, group-based [2], pull-based, and access class barring [7]. Although these approaches can reduce the access collision to a certain degree, they are unable to support a large number of devices in IoT scenarios [8].

Excessive Overhead: Another problem with the connection-oriented communication in the current LTE standard is excessive signaling overhead as significant resources are spent establishing a connection to allow transmission of the very small-sized data (e.g., a few kilobits) typically required for M2M communications, especially when a large number of M2M devices attempt to access cellular networks at the same time [9]. For example, to transmit 100 B of data, approximately 59 B of overhead on the uplink and 136 B on the downlink would typically be required for signaling transmissions [6]. Hybrid schemes are proposed to combine the RA procedure and the data transmission, where the devices will send their messages through the third message of the RA procedure. Data aggregation could also be used for more efficient transmission, but is only applicable for delay-tolerant M2M applications. These strategies are also not scalable and mostly inefficient when the number of devices is very large.

Different QoS Requirements: Many M2M applications have diverse service requirements that must be carefully considered when designing the access techniques. For example, some M2M applications, such as alarm signals, are delay-sensitive, and a very small message must be delivered within 10 ms, while other applications, such as smart metering, are delay-tolerant or have larger packet sizes and can tolerate delays of several hours. Most existing access techniques for M2M communications have not considered QoS requirements of MTC devices and treat all the devices the same. However, this is ineffective for real-world M2M applications and leads to huge radio resource waste and/or service interruptions. QoS can be effectively taken into account when designing the access technology, especially in massive IoT applications, which can significantly reduce the load on the core network and minimize the access delay for delay-sensitive applications.

Coexistence with H2H Devices: In cellular-based M2M applications, MTC devices coexist with H2H devices. Since the number of MTC devices is very large compared to that of H2H, and they use the same radio resources, inefficient design of the transmission procedure for MTC devices may result in huge losses in H2H communication. One could define new PRACH resources for M2M devices to avoid congestion with H2H, or consider dynamic PRACH resource allocation to adjust available resources based on the estimated traffic. Although these approaches can maintain the QoS requirements for H2H devices, many M2M devices may be required to delay their transmissions and wait until an appropriate number of resources become available for M2M devices. This problem will be more challenging when the number of M2M devices is very large and only a limited number of radio resources are available for both H2H and M2M traffic.

Non-Orthogonal Multiple Access for M2M Communications

Multiple access techniques can be classified into orthogonal and non-orthogonal approaches. In orthogonal multiple access (OMA), including time-division multiple access (TDMA), frequency-division multiple access (FDMA), and orthogonal FDMA (OFDMA), signals from different users are not overlapped with each other. Non-orthogonal schemes, in contrast, allow overlapping among the signals in time or frequency by exploiting the power domain, code domain, or interleaver pattern, often providing better performance in comparison to orthogonal schemes in terms of throughput [10]. OMA is a suitable choice for packet domain services with channel-aware time and frequency scheduling [10]. However, further improvements in the system efficiency and QoS required for the mobile cellular networks' 5G and IoT applications necessitates the adoption of NOMA schemes with high throughput efficiency.

The Basic Concept of Uplink NOMA

For simplicity, we assume here two users and a single transmitter and receiver antenna. User i transmits signal s_i with transmission power p_i (Fig. 3a). The received signal at the base station is represented by

$$y = \sqrt{p_1}h_1 s_1 + \sqrt{p_2}h_2 s_2 + w \quad (1)$$

where w denotes the received noise, including inter-cell interference, and h_i is the channel coefficient between user i and the base station. In NOMA, both s_1 and s_2 are sent over the same frequency band in the same time slot and therefore interfere with each other. At the base station, successive interference cancellation (SIC) is implemented, where first s_1 is decoded by treating s_2 as interference. Once the receiver correctly decodes s_1, it subtracts s_1 from the received signal y and then decodes s_2. The receiver decides the order of decoding according to the effective SINR of the users.

Hybrid schemes are proposed to combine the RA procedure and the data transmission, where the devices will send their messages through the third message of the RA procedure. Data aggregation could also be used for more efficient transmission, but is only applicable for delay-tolerant M2M applications.

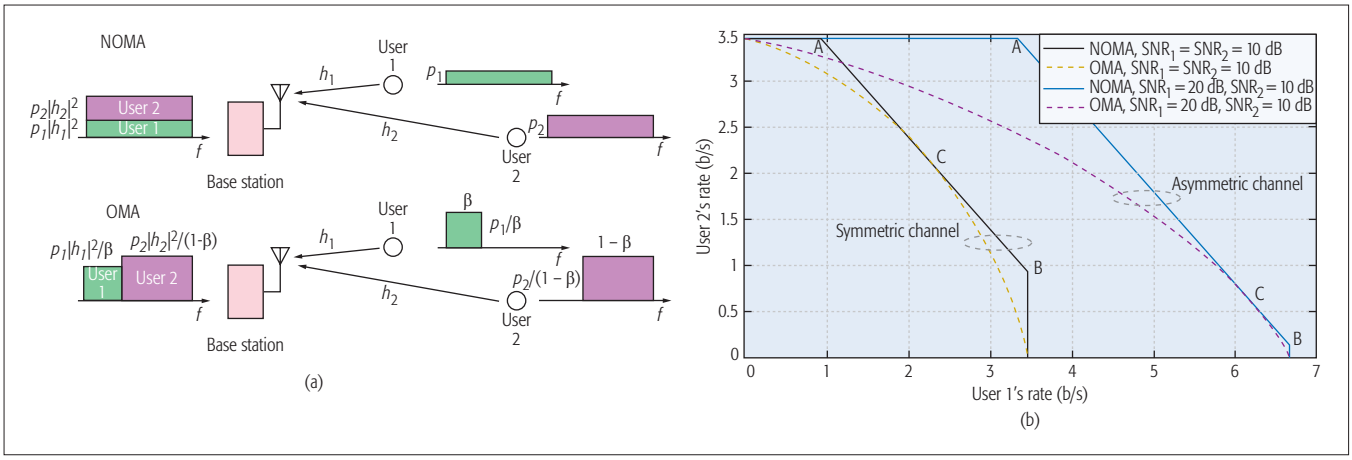

Figure 3. NOMA and OMA and their achievable rate for a two-user scenario: a) NOMA vs. OMA; b) capacity region of the multiple access channel.

From an information-theoretic point of view, NOMA with SIC is an optimal multiple access scheme in terms of the achievable multiuser capacity region in both uplink and downlink [11]. In NOMA, the performance gain compared to OMA increases when the difference in channel gains or path loss between the users and the base station is large. Figure 3b shows the capacity region of a single cell scenario with two users with different signal-to-noise ratios (SNRs), where the total bandwidth is assumed to be 1 Hz [11]. Point A for both cases is achieved when the signal for user 1 is decoded first. Point B is achieved when the base station first decodes user 2. It is important to note that NOMA always achieves the highest sum rate, while OMA has a gap to the maximum achievable sum rate, with the exception of one point, the symmetric case, where the users achieve the same throughput. It is also clear that when the difference between channel gains increases, NOMA brings greater improvements compared to OMA.

The main benefit of NOMA for IoT is the removal of the need for an RA stage and enabling the devices to transmit in the same channels. This leads to efficient use of available radio resources, and solves the signaling overhead problem due to conventional RA strategy in cellular systems.

NOMA for Massive Cellular IoT

In the proposed random NOMA for M2M communications, the devices do not need to perform the RA procedure to access the network. Instead, the RA and the data transmission are combined, and the devices transmit their messages over randomly selected subbands. This is necessary to minimize the overhead, which is critical for many M2M applications with small message sizes. We assume that the total bandwidth is divided into several subbands, and each device that has data to transmit will randomly choose a subband for its data transmission. Multiple subbands have been considered in this article to show how compatible the proposed NOMA is with the existing LTE and LTE-Advanced standards, which are based on OFDMA technologies. The devices then choose a channel code with an appropriate code rate and encode their messages along with their terminal identities. The coded message is then sent over the selected subband by the MTC device. Upon receiving the message, the base station performs SIC to decode the message of each MTC device over each subband. Figure 4 shows the random NOMA for massive cellular IoT, where only four subbands are available.

As the devices randomly choose a subband for their transmissions, the number of MTC devices transmitting over different subbands is not the same. Accordingly, the channel code rate cannot be fixed but instead is adapted to the random activities of devices. Toward this aim, Raptor codes [12], which are rateless and can generate as many coded symbols as required by the base station, can be used. The code structure is random and can be represented by a bipartite graph; then the base station can reproduce the same bipartite graph using a pseudo random generator with the same seed. When more than one device selects the same seed and transmits over the same subband, they will be transmitting using exactly the same code structure; thus, the base station cannot differentiate between them as there is no structural difference between the received codewords. We call this event a collision. But we note that a collision requires both the same subband and same code structure to be chosen, making it far less likely to occur than a traditional RA collision.

In Fig. 5 we show a simple simulation result to show the advantage of NOMA over conventional OMA strategies. As can be seen in Fig. 5a, NOMA outperforms uncoordinated TDMA and FDMA strategies, and can support a larger number of devices, while FDMA and TDMA suffer from high probability of collision. We compare the proposed random NOMA scheme with different numbers of subbands with the access class barring (ACB) scheme [13] when the same number of radio resources are available in Fig. 5b. As can be seen in this figure, NOMA can support a significantly larger number of devices compared to ACB schemes. It is important to note that in the ACB scheme, we assume that a device can successfully deliver its message to the base station in its corresponding data channel when it has successfully completed the RA phase; thus, the only limiting factor for the ACB scheme is the preamble collision in the RA procedure.

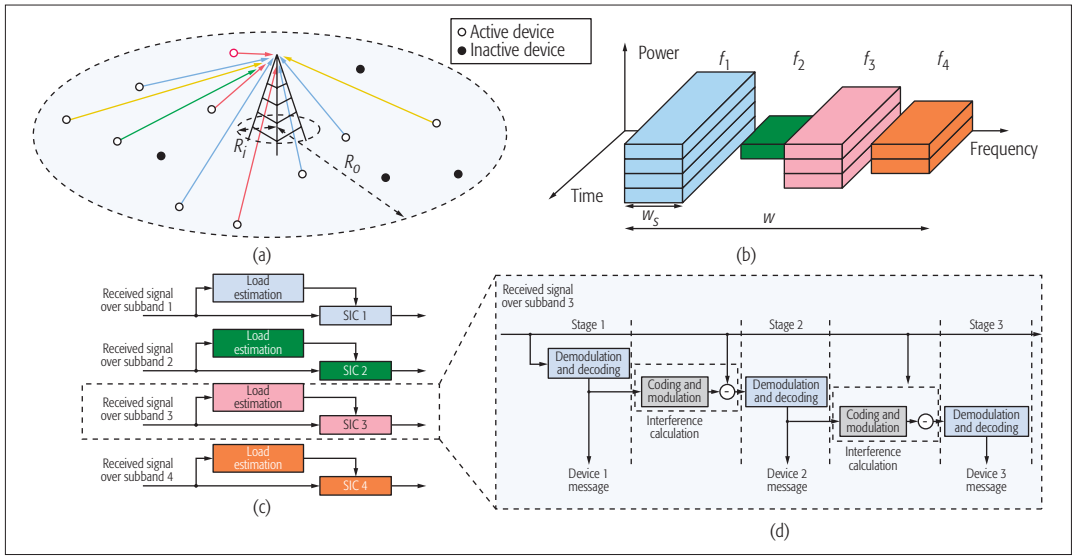

Figure 4. Random non-orthogonal multiple access for cellular M2M communications: a) each MTC device randomly chooses a subband for its data transmission; b) the received signals at the base station from different subbands. Devices perform power control so that the received power from all the devices at the base station is the same; c) the base station performs load estimation and successive interference cancellation over all the subbands; d) the multi-stage structure of SIC for subband 3, where three devices have transmitted their messages over it.

The main benefit of NOMA for IoT is the removal of the need for an RA stage and enabling the devices to transmit in the same channels. This leads to the efficient use of available radio resources, and solves the signalling overhead problem due to conventional RA strategy in cellular systems.

Potentials of NOMA for Massive Cellular IoT

NOMA can bring many benefits to cellular systems, which include but are not limited to the following:

- Effective use of spectrum and higher system throughput are attained through exploiting the power domain and utilizing non-orthogonal multiplexing.
- Robust performance gain is achieved in high-mobility scenarios, where OMA schemes obtains no frequency-domain scheduling gain as channel state information is outdated, but NOMA provides gains in high-mobility scenarios as it relies on the channel state information at the receiver side [11].
- NOMA is compatible with OFDMA and its variants and can be applied on top of OFDMA for downlink and single-carrier FDMA (SCFDMA) for uplink [10].
- NOMA can be combined with beamforming and multi-antenna technologies to improve system performance [11].
- NOMA can easily be combined with radio resource management and random access techniques to solve the collision and overload problem in M2M communications.
- Using clustering and group-based scheduling, NOMA can be used in M2M communications as the multiple access technique to deliver messages of a group of devices to the base station or the cluster head.
- Using NOMA, the RA procedure can be eliminated, and therefore the access delay and signaling overhead will be significantly reduced.

Practical Considerations of Massive NOMA for Massive Cellular IoT and Future Directions

Although NOMA can improve spectrum efficiency and system capacity, there are many practical challenges for this technology to be potentially used in real wireless systems for M2M communications. Here, we outline the main practical consideration of massive NOMA for M2M communications.

Traffic and load estimation at the base station: As the devices randomly select the subbands for their data transmission, the base station needs to accurately estimate the number of devices that are transmitting over each subband. One approach is to perform power control at the devices so that the received power over each subband will be proportional to the number of devices transmitting over that subband. We discuss this in more detail in the following.

Channel estimation and power allocation: As many devices want to simultaneously communicate with the base station, it is almost impossible for the base station to estimate the channel to all of these devices. The problem becomes more challenging when multiple antennas are used in either the devices or the base station. One could consider the channel between each device and the base station as reciprocal in each direction (this is the case in time-division duplexing), so the devices can estimate their channel to the base station using the pilot signal periodically broadcasted by the base station, and then adjust their transmission power so that the received signal power at the base station is the same fixed value for each device. This is beneficial for M2M communications as the devices have small data packets, and the capacity gains of NOMA are not the main advantages of this strategy. Instead, NOMA improves the system throughput through eliminating the RA procedure and enables multiuser detection at the base station.

Synchronization among devices: In random NOMA, the devices will be identified during data transmission, so it is not possible for the base station to determine the timing advance information for every device. One could con-

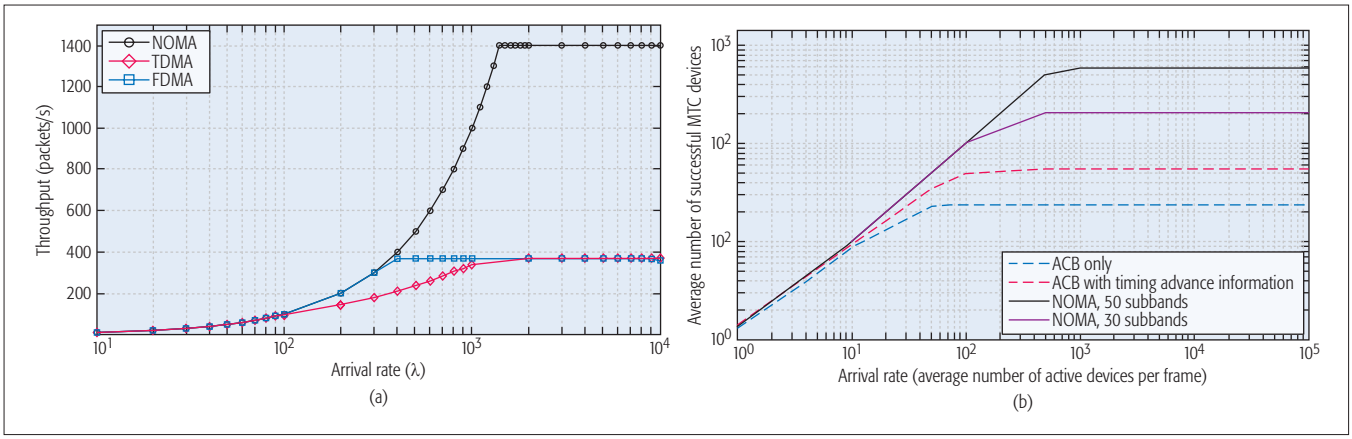

Figure 5. Throughput of NOMA vs. OMA: a) NOMA vs. TDMA and FDMA when only one subband is used for NOMA. Minimum subband bandwidth in FDMA is 100 kHz and minimum time slot duration in TDMA is 1 ms. The total bandwidth is 1 MHz; b) average number of successful MTC devices vs. arrival rate. The message length of MTC devices is 1024 bits, the number of preambles is 64, and a resource block has 1 MHz bandwidth and time duration of 1 ms.

sider that the devices will estimate their timing information from previous transmissions or according to their location information, which is more practical in M2M applications with fixed-location devices. Providing time synchronicity between a large number of devices is a challenging problem and requires major technical efforts.

Proper channel code design: The effective data rate for each device is determined by the number of devices that are transmitting in the same subband. As the devices randomly select a subband for the data transmission, the number of devices that are transmitting over each subband is not known beforehand. This means that the code rate at which each device should transmit its data is not fixed. One approach is to use rateless codes so that the devices will transmit using a rateless code and stop their transmission once they receive an acknowledgment from the base station. This has been investigated in [14], but one should take into account the random structure of rateless codes and think of a way to exchange the random graph structure between the base station and the devices.

Complexity of SIC: In existing orthogonal approaches, the BS needs to perform a separate decoding for each device, but through the proposed NOMA, the same decoder can be used for all the devices in a sequential manner. However, we can also consider a separate decoder for each device at the receiver and decode each device considering the signal from all other devices as additive noise. This may slightly reduce the throughput, but the degradation can be neglected as the achievable rate is mainly determined by the device with the lowest SINR at the BS. Interested readers are referred to [14] for further details on the SIC process.

User fairness: The BS can allocate higher bandwidth for those devices that have low-quality links to the BS so that they can transmit at lower power. This way, devices with low-quality links will transmit with lower power over larger bandwidth and can achieve the same throughput or energy efficiency as devices with high-quality links that are transmitting over smaller bandwidth or with higher power.

Conclusions

This article reviews recent advancements in random access techniques for M2M communications and presents an overview of their benefits and challenges. We have described the basic concept of uplink non-orthogonal multiple access and proposed it as the potential multiple access technology for future cellular systems to accommodate the tremendous growth of M2M applications and traffic. The practical challenges of massive NOMA for M2M communications are also presented, and future research directions are highlighted. Massive NOMA offers high throughput efficiency with simple system structure, which is particularly beneficial for massive IoT applications with low-cost, low-power, and low-complexity devices, and can provide system scalability to support the massive number of devices involved in M2M communications. This technology can easily be adopted by 3GPP technologies for M2M communications to further boost the system performance of current cellular solutions for IoT.

References

[1] Ericsson, "Cellular Networks for Massive IoT," tech. rep. Uen 284 23-3278, Jan. 2016; https://www.ericsson.com/res/docs/whitepapers/wp iot.pdf.

[2] S.-Y. Lien, K.-C. Chen, and Y. Lin, "Toward Ubiquitous Massive Accesses in 3GPP Machine-to-Machine Communications," *IEEE Commun. Mag.*, vol. 49, no. 4, Apr. 2011, pp. 66–74.

[3] 4G Americas, "Cellular Technologies Enabling the Internet of Things," tech. rep., Nov. 2015; http://www.4gamericas.org/files/6014/4683/4670/4G_Americas_Cellular_Technologies_Enabling_the_IoT_White_Paper.pdf, Nov. 2015.

[4] Nokia White Paper, "LTE Evolution for IoT Connectivity," tech. rep. Jan. 2016, http://resources.alcatellucent. com/asset/200178.

[5] M. Hasan, E. Hossain, and D. Niyato, "Random Access for Machine-to-Machine Communication in LTE-Advanced Networks: Issues and Approaches," *IEEE Commun. Mag.*, vol. 51, no. 6, June 2013, pp. 86–93.

[6] H. Shariatmadari *et al.*, "Machine-Type Communications: Current Status and Future Perspectives toward 5G Systems," *IEEE Commun. Mag.*, vol. 53, no. 9, Sept. 2015, pp. 10–17.

[7] S.-Y. Lien *et al.*, "Cooperative Access Class Barring for Machine-to-Machine Communications," *IEEE Trans. Wireless Commun.*, vol. 11, no. 1, 2012, pp. 27–32.

[8] K. Zheng *et al.*, "Challenges of Massive Access in Highly Dense LTE-Advanced Networks with Machine-to-Machine Communications," *IEEE Wireless Commun.*, vol. 21, no. 3, June 2014, pp. 12–18.

[9] D. Wiriaatmadja and K. W. Choi, "Hybrid Random Access and Data Transmission Protocol for Machine-to-Machine Communications in Cellular Networks," *IEEE Trans. Wireless Commun.*, vol. 14, no. 1, Jan. 2015, pp. 33–46.

[10] A. Li *et al.*, "Non-Orthogonal Multiple Access (NOMA) for Future Downlink Radio Access of 5G," *China Commun.*, vol. 12, Supplement, Dec. 2015, pp. 28–37.
[11] K. Higuchi and A. Benjebbour, "Non-Orthogonal Multiple Access (NOMA) with Successive Interference Cancellation for Future Radio Access," *IEICE Trans. Commun.*, vol. 98, no. 3, 2015, pp. 403–14.
[12] A. Shokrollahi, "Raptor Codes," *IEEE Trans. Info. Theory*, vol. 52, no. 6, June 2006, pp. 2551–67.
[13] Z. Wang and V. Wong, "Optimal Access Class Barring for Stationary Machine Type Communication Devices with Timing Advance Information," *IEEE Trans. Wireless Commun.*, vol. 14, no. 10, Oct. 2015, pp. 5374–87.
[14] M. Shirvanimoghaddam, M. Dohler, and S. Johnson, "Massive Multiple Access Based on Superposition Raptor Codes for Cellular M2M Communications," *IEEE Trans. Wireless Commun.*, vol. 16, no. 1, Jan. 2017, pp. 307–19.

Biographies

Mahyar Shirvanimoghaddam [M] (mahyar.shirvanimoghaddam@sydney.edu.au) received his B.Sc. degree with 1st Class Honours from the University of Tehran, Iran, in September 2008, his M.Sc. degree with 1st Class Honours from Sharif University of Technology, Iran, in October 2010, and his Ph.D. degree from the University of Sydney, Australia, in January 2015, all in electrical engineering. He then held a postdoctoral research position with the School of Electrical Engineering and Computing at the University of Newcastle, Australia. Since November 2016, he has been with the School of Electrical and Information Engineering, University of Sydney, Australia, as a Scholarly Teaching Fellow in Telecommunications. His general research interests include channel coding techniques, cooperative communications, compressed sensing, machine-to-machine communications, and wireless sensor networks.

Mischa Dohler [F] (mischa.dohler@kcl.ac.uk) — Fellow of the the Royal Academy of Engineering and the Royal Society of the Arts — is a full professor in wireless communications at King's College London, director of the Centre for Telecommunications Research, co-founder and member of the Board of Directors of the smart city pioneer Worldsensing, and a Distinguished Member of Harvard Square Leaders Excellence. He is a frequent keynote, panel, and tutorial speaker. He has pioneered several research fields, contributed to numerous wireless broadband and IoT/M2M standards, holds a dozen patents, has organized and chaired numerous conferences, has more than 200 publications, and has authored several books. He acts as a policy, technology, and entrepreneurship adviser; and is also an entrepreneur, composer, and pianist with five albums on iTunes; and is also fluent in six languages. He has talked at TEDx. He has had coverage by national and international TV and radio, and his contributions have been featured on BBC News and in *The Wall Street Journal*.

Sarah J. Johnson [M] (sarah.johnson@newcastle.edu.au) received her B.E. degree in electrical engineering and her Ph.D. degree from the University of Newcastle, Australia, in 2000 and 2004, respectively. She held a postdoctoral position at NICTA, Australia's Information and Communications Technology Research Centre, before returning to the University of Newcastle, where she is now an associate professor and Australian Research Council Future Fellow. Her research interests include error correction codes and information theory for multiple-user communication networks. She has authored a book, *Iterative Error Correction* (Cambridge University Press).

Value of Information and Cost of Privacy in the Internet of Things

Damla Turgut and Ladislau Bölöni

Will the Internet of Things happen? Clearly, the hardware and software components comprising the Internet of Things are technologically feasible, but the sweeping adoption we envision might not take place. The success of technological innovations depends on the creation of a business model that both customers and providers perceive as beneficial.

Will the Internet of Things happen? Clearly, the hardware and software components comprising the Internet of Things are technologically feasible, but the sweeping adoption we envision might not take place. The success of technological innovations depends on the creation of a business model that both customers and providers perceive as beneficial.

ABSTRACT

Will the Internet of Things happen? Clearly, the hardware and software components comprising the Internet of Things are technologically feasible, but the sweeping adoption we envision might not take place. The success of technological innovations depends on the creation of a business model that both customers and providers perceive as beneficial. As the recently abandoned Google Glass project shows, privacy concerns can kill an otherwise technologically feasible product. On the other hand, the example of Twitter illustrates that very popular products might fail to make money. Both academic researchers and businesses are becoming increasingly aware that we need to reason about the economic and social implications of provided value and privacy in a rigorous, quantitative way.

In this article we call these quantities *value of information*, which appears to both the service providers and the customers, and *cost of privacy*, which normally is only relevant to the customers. We describe the importance of assessing these values in the context of the Internet of Things, possible directions of their formalization, their relationships to other problems, and related areas as well as future directions of the field.

INTRODUCTION

The life of a citizen of the early 21st century takes place simultaneously in physical and cyber space. We physically move between home, work, shopping centers, and recreation and tourist areas. At the same time, through our smartphones and other mobile devices we also occupy a position in cyberspace, by logging into social networks, connecting to online services, chatting with friends and business partners, or looking up online prices for products we see in brick-and-mortar stores. In some cases our activities in cyberspace have nothing to do with our physical location. However, the overall trend in the case of services like Google Maps, Waze, Yelp, and Foursquare is the strong interdependence of physical and cyber space. This participation in multiple physical and cyber spaces creates unprecedented privacy problems.

Humans living in an urban environment always have had to balance their security and privacy against their public lives. A certain degree of threat is unavoidable when we move in a public space, but we do not choose to barricade ourselves in our homes. Instead, we make an informed decision about the acceptable degree of threat: it is acceptable to walk in a mall on a Sunday afternoon, but it is not acceptable to walk in a crime ridden area after midnight. Similar considerations apply to privacy: by entering a public space, we implicitly accept that we will be seen by other people, will be recognized by acquaintances; and even strangers can make some inferences about our shopping habits based on the stores we frequent. However, we do not normally exhibit our personal details in public spaces — we do not advertise our phone number, bank accounts, and shopping history to strangers.

This negotiation of the benefits and costs of participating in the physical public space is based on our ability to estimate the security and privacy *cost* of our actions. We learn to do this in physical space from a very young age, and we can rely on social and physical signals to perform this estimate. An imposing bank building of marble and steel signals trustworthiness and confidentiality, while a grafitti-ridden, blighted urban area signals potential danger.

In contrast, in cyberspace our ability to estimate our safety and privacy is significantly lowered. For most American citizens, cyberspace is a much less safe and private space than the public physical spaces they visit. Significantly more people are exposed to cyberfraud than physical pickpocketing, and they give up private information more often through online services than by physical communication. The problem is not that there are more bad people in cyberspace than in the real world. Rather, the main problem is that the signals of danger are much less reliable in cyberspace: it is much easier to construct a slick web page for a shady service that collects personal data than to build an imposing physical bank building.

As difficult as these problems were, they were at least ameliorated by the fact that the real world and cyberspace were clearly distinguishable — we learned to act differently in the physical world and in cyberspace. The presence of smartphones brought *entry points to cyberspace* available at every moment of our lives. *The collection of technologies we refer to as the Internet of Things (IoT) will make physical reality and cyberspace essentially indistinguishable.*

Will the Internet of Things (IoT) happen? The software and hardware technologies of the IoT are a direct offshoot of existing research on mobile computing, sensor networks, ad hoc networks, distributed systems, security, machine learning,

Digital Object Identifier:
10.1109/MCOM.2017.1600625

The authors are with the University of Central Florida.

0163-6804/17/$25.00 © 2017 IEEE

big data, and others. While many research problems are open, the technological feasibility of the IoT vision appears guaranteed. The IoT vision, however, assumes widescale adoption by the public of the IoT technologies, and this will only happen if:

a. The customers are persuaded that the IoT devices provide a value that exceeds their physical and privacy costs.
b. The businesses involved in IoT successfully make money.

Both conditions are necessary — there are many examples where the lack of a or b derailed technological visions. For instance, the Google Glass technology (at least in its first version) was abandoned by Google due to the largely negative feedback it received, with most of the feedback centering on the privacy problems it created. Should Google revive this technology, its most likely modifications will be centered on assuaging the privacy concerns. With regard to condition b, naturally, there are many popular technologies that have been abandoned because they did not successfully make money for the businesses that promoted them. Even widely popular technologies such as Twitter are money losing propositions and will fold if an appropriate way to monetize them is not developed.

In this article we concentrate on the value and cost concerning the exchange of data in IoT, while acknowledging that other types of costs (hardware, energy, installation, and maintenance) will also play a role. In particular we can write condition a, describing the customer's benefit, as follows:

$$V_{service} - C_{privacy} - C^{user}_{hardware} - C_{payment} > 0 \qquad (1)$$

That is, the perceived value of the service for the user $V_{service}$ has to cover the cost of lost privacy $C_{privacy}$ and the share of the user in the cost of the hardware and associated services $C^{user}_{hardware}$ and whatever payment the user made for the service $C_{payment}$.

On the other hand, condition b can be described as

$$V_{information} - C^{business}_{hardware} + C_{payment} > 0 \qquad (2)$$

That is, the value of information received $V_{information}$ and direct payments must be higher than the businesses' share of the hardware and maintenance costs $C^{business}_{hardware}$.

The naïve view of such a transaction is that the information received by the provider is necessary for the provision of the service. In this setting, the $V_{information}$ value appears only because the provider commercially exploits the information it receives in the course of service provision. In practice, however, the motivation of the provider is to acquire as much valuable information as possible. Thus, in practice, the only relationship between the value of service received by the user $V_{service}$ and the value of information $V_{information}$ collected by the provider is that the user is willing to enter into a transaction under these terms. This is essentially similar to the pricing of goods under monopolistic competition, and there is a significant wealth of theory that can be used in future research.

Another issue is that the value of information for a business is also determined by the legal and regulatory landscape in which the transaction takes place. Laws and regulations determine both what type of information can be collected and the way in which this information can be used. In general, the Data Protection Directive in the European Union requires more explicit notification about the collected data and puts stricter limits on its use than the currently applicable laws in the United States.

The costs, values, and regulations also depend on the application area. Healthcare and education are two fields that have well established legal norms for privacy and confidentiality from the pre-IoT era. Medical confidentiality dates back thousands of years, being part of the original Hippocratic oath dating from the 5th century BCE. Both medical and educational confidentiality is codified in laws in many countries. These regulations will naturally transfer to e-health and educational IoT applications. Other IoT application areas, such as commerce, sports, and recreation have much fewer legal restrictions on the flow of information. Nevertheless, one of the major challenges of the IoT world is that the abundance of sensors might lead to an information leak among application areas. For instance, fitness sensors capture health related information, and physical location tracking might provide information about educational achievement (e.g., by detecting visits to remedial math classes).

Determining the exact costs and values associated with IoT is not easy. The user's physical hardware costs are spread over many transactions with different providers. The services are rarely paid for in the form of well defined micropayments, as the businesses aim to develop creative pricing schemes that incentivize users to use the service. These models might depend on local cultural norms; for instance, American customers appear to prefer all-included subscriptions, while Europeans prefer metered services. The division of the costs might also be more fine-grained than illustrated above. We separate the cost of service from the cost of hardware as these are normally paid to different recipients, but a more fine-grained model might separate networking and energy costs.

In this article we assume that the primary transaction is between a single user and a single provider. If we allow for group users (e.g., co-owners of a device or groups of users who are pooling together in an Uber vehicle) and group providers (the services of multiple providers need to be integrated in a more complex service), the complexity of the model increases. We need to consider, for instance, how the values are divided across the group members and whether the realizable values are additive, super-additive, or sub-additive. While this opens interesting theoretical research opportunities, the current architecture of web services based on point-to-point REST calls in general enforce discrete one-to-one interactions.

One of the major challenges of the IoT world is that the abundance of sensors might lead to an information leak among application areas. For instance, fitness sensors capture health related information, and physical location tracking might provide information about educational achievement (e.g., by detecting visits to remedial math classes).

Quantifying Value of Information and Cost of Privacy

The formulas introduced in the previous section are a good starting point, but the main challenge is to quantify (i.e., put numbers) on the various values. The simplest values to quantify might be

Users do not act optimally in the privacy-for-service marketplace – sometimes they will decline entire useful service models, while other times give up too much privacy for services of little value to them. Such inefficient marketplaces are disadvantageous for both the service providers and the customers.

the $C^{business}_{hardware}$ and $C^{user}_{hardware}$, as these are actual billable costs. For instance, the user might pay for her wearable devices and IoT components in her home, and the businesses might pay for IoT augmentation of public spaces. Nevertheless, even with these values, things might get more complicated when the IoT devices are shared among multiple users and businesses (as they likely will be).

The value of information to the business $V_{information}$ has seen a significant focus in both research and commercial studies. For instance this is the value for which the participants of the Google Adwords or Bing Ads program are bidding — these systems essentially resell to the advertiser the information that "user X is looking for product Y." Of course, this value depends on the type of information; for instance, the per-click cost can range from about $1 on average to more than $100 for keywords such as "lawyer" or more than $50 for "insurance."

Many of the services in today's mobile economy are "ad-supported" and nominally free for the user, which means $C_{payment} = 0$. By implication, this means that IoT will essentially be a system composed of individual transactions in which the (perceived) cost of privacy is exchanged for the (perceived) value of a service. The central idea is that privacy has a quantifiable value. Participants in the mobile economy do understand the value of privacy, but they are not accustomed to evaluating it on a transaction-by-transaction basis and weighing it against the value of services received. Thus, users do not act optimally in the privacy-for-service marketplace: sometimes they will decline entire useful service models, while at other times they will give up too much privacy for services of little value to them. Such inefficient marketplaces are disadvantageous for both the service providers and the customers.

How do we quantify the cost of privacy? The situations where we can explicitly measure the cost of privacy are rare. An example is the Kindle Special Offers program. Amazon Kindle users can remove the advertising screensaver from Amazon Kindle devices for $20. In our model, this means that customers who take advantage of this consider that the increase of the $C^{user}_{hardware}$ from $80 to $100 is offset by the corresponding decrease of the perceived $C_{privacy}$ (albeit other factors such as convenience might also be a factor). Unfortunately, Amazon does not publish how many people are signing up for the removal of the ads. Another project where the $C_{privacy}$ appears more or less in an explicit form is the Google Contributor program, where people pay $7 or higher for the removal of advertisements from websites, an act that many people associate with the perception of improved privacy (although the exact privacy implications are not clear).

Research Directions in Cost of Privacy for IoT

As the case was made in the previous sections, understanding the cost of privacy and its relationship to the value of information is a critical requirement to the success of the IoT paradigm. This requires a new approach to the problem of privacy. There is an extensive literature on privacy applied in a number of fields, from networking to database queries, a preoccupation going back for decades. The typical definition of privacy was the lack of information leakage. In these models, the ultimate goal is perfect privacy. A typical question addressed, for instance, how a user can avoid the disclosure of her location to another party equipped with a certain set of capabilities. The threat model in this case is that the other party is an *opponent*, who does not offer anything in exchange for the acquired information. On the other side, the user's preferences are also clear: the less disclosure, the better. There are a variety of algorithmic and cryptographic/security techniques. Algorithmic techniques include methods to reduce the quantity or accuracy of data, data anonymization, and distributed architectures [1]. Security-based approaches include secure data sharing approaches [2] and access control techniques.

In the economic model of IoT, however, perfect privacy cannot be the goal as it would reduce or eliminate the profit of the participating companies (or society as a whole). The goal instead should be a *fair trade* — the benefits the users obtain from the services of the IoT system should be commensurate with the information they are willing to give up. For instance, the relationship between the user of a Google product and the company is not antagonistic — it can be more accurately described as a *trade*. In this trade, the user voluntarily gives up some information in exchange for services received. Thus, we can say that the user made the disclosure voluntarily — this does not, however, mean that the trade was an advantageous one. We can say that the user incurs a privacy cost, which requires us to determine the cost of privacy (CoP) of specific chunks of information. Thus, privacy can be seen as a formalizable mathematical value in some situations, but it can also be seen as a tradable economic commodity in others, or simply a value for which users can have more or less predictable preferences. This requires us to reason about both the cost of privacy as well as the value of information (although different researchers might use different terminologies). In the following we discuss some of the research challenges posed by this new approach.

Formal Models of Value of Information and Cost of Privacy

This research direction aims to develop a formal definition of the cost of privacy (CoP), which is mathematically rigorous to allow for formal proofs, but also matches our intuitions behind the concept.

One approach to define CoP is by analogy with the game theoretic concept of value of information (VoI) [3]. The intuition behind the game theoretical definition of VoI is that of the price an optimal player would pay for a piece of information. In recent years, the concept of VoI has been applied to a number of scenarios in wireless networking and mobile computing. A number of recent projects have introduced similar metrics to model situations where one needs to either select a subset of the collected data or choose between transmitting a piece of information or not. Bis-

dikian *et al.* [4] considered the probabilistically defined concept *quality of information* (closely related to VoI) and applied it to sensor networks in military environments. Another approach is that of *pragmatic VoI* as the support the information gives to the decisions and actions of the operator (without assuming an optimal decision maker) [5].

There are obvious parallels between CoP and VoI. VoI attaches a value to an information chunk *acquired* by agent A. In contrast, CoP attaches a value to an information chunk *disclosed* by agent B to agent C. Despite the similarity, there are some important differences, which require careful formal modeling. For the acquisition of information, the benefits of the information are realized instantaneously by agent A. Thus, in the game theoretic sense, VoI depends only on the data chunk, agent A, and the current game. In the case of the CoP, however, B does not incur any immediate costs. The losses suffered by B are more subtle — for instance, in a later situation he might be at a competitive disadvantage vs. C. In the game theoretic sense, the cost of privacy must be defined over a series of games, and it will also depend on the pair of (B,C), rather than only on agent B.

Elicitation-Based Techniques

The first model of determining VoI and CoP is deceptively simple at first sight: let us simply ask the user to assign numbers to these values. The first question raised by this idea is whether the users even think about these values. In the early days of mobile computing, users might have been naive about the amount of data captured by the devices they use, and even today there are situations where a deceptive provider attempts to collect data without acknowledging the fact [6]. However, in recent operating systems, applications are required to disclose that they are collecting a certain type of user data (e.g., location). Indeed, users often choose to not install or unsubscribe from applications whose data collection practices are deemed excessive. We can conclude that the majority of users are aware that there is a CoP associated with these services.

The second question concerns the actual techniques of eliciting the CoP information from users. Fields of science such as psychology, anthropology, market research, and political science have developed many techniques to elicit values from users.

Elicitation interviews ask the subjects to re-enact the specific situation in a laboratory setting either alone or as part of focus groups. For instance, the users might be instructed to imagine that they are in a shopping mall and asked about their perceived CoP. The weakness of this approach is that the artificial setting might influence the users' answers.

The descriptive experience sampling method attempts to elicit users' feedback in the course of their regular day. A user is provided a random timer that, at specific moments, interrupts the user's current activity. At these interruptions, the device records the user's current state and asks hypothetical questions about CoP or VoI. The advantage of this method is that the user is actually part of the current situation, and the lack of preparation might lead to more "honest" answers.

Real-time decision capture uses technological means to investigate the economic decision making process of the subjects at the moment when they are made. While this technique would create the most accurate answers, it requires us to augment the devices through which the actual decisions are being made. In addition, we can only capture whether a user was in favor of or against a certain transaction, rather than the numerical values and costs involved.

The elicited CoP is essentially a perception, which can be affected by many outside factors. For instance, the spread of an Internet meme where the loss of privacy led to significant financial cost can suddenly raise the CoP for all the users it touched. Users might attach little cost to disclosing information that is public knowledge. For instance, the fact that a person is in her workplace from 9 a.m. to 5 p.m., and at home from 6 p.m. to 8 a.m. is well known — disclosing this information has little cost attached to it. Different users might have different privacy requirements [7]: celebrities and political figures might value their privacy more highly than average people. In many cases, the privacy value might be different for different aspects or times. For instance, a doctor might not care about disclosing locations she visits as a private person, but it might be under a confidentiality agreement about house calls made in her professional capacity.

Beyond the actual elicitation of CoP values, this research direction promises to unveil answers to other questions, such as:

Are users acting rationally in their service-for-privacy trades? We expect that, similar to most instances of human economic behavior, the human subjects approximate rational behavior in aggregate but present specific deviations due to cognitive biases.

What is the CoP for specific disclosures and on what does it depend? We expect that the CoP associated by the users with different disclosures depends on the environment and the identity of the service provider. We further conjecture that the perceived trustworthiness of the service provider influences the cost of the privacy [8].

Negotiating the Cost of Privacy with or on Behalf of the User

As discussed above, the CoP and the associated VoI depends on many factors. This not only makes formalization difficult, but also elicitation — users might not be immediately aware of just how much the associated cost will be. An alternative approach would be to discover the cost of privacy iteratively, through negotiation or an auction system. Just like the value of a difficult-to-appraise item can be estimated through an auction, the CoP might be estimated if the user is presented with explicit offers for his data.

Several academic studies analyzed the user's valuation of its privacy through auction-based techniques. For the web browsing model, the authors in [9] found that users allocated about $10 for their browsing history and about $36 for their age and address. In [10] the authors implemented a mobile app where the users could put a price on information recorded by their mobile phones such as their location, applications used, or number of calls made.

Real-time decision capture uses technological means to investigate the economic decision making process of the subjects at the moment when they are made. While this technique would create the most accurate answers, it requires us to augment the devices through which the actual decisions are being made.

The vision of IoT as a truly universal part of human society would require a buy-in from everybody, all the time. Economic forces will ensure that businesses will carefully evaluate the costs of participating in the IoT and the values extracted from it, and they would withdraw from ventures that do not have a positive balance.

In IoT-augmented public spaces there can be many events that lead to information disclosure. Many of these may not be initiated by the user. It is unreasonable to expect the user to enter into a negotiation every time such a disclosure might happen. Ultimately, the appropriate solution would be an intelligent agent that performs these negotiations on behalf of the user, taking into account the preferences and possibly the negotiation strategy of the user. An early example of such a system is the Google Contributor system, which negotiates on behalf of the user for the position of each ad, based on a predefined pool of money. If the user wins the auction, the ad will not be shown. Such a system can be adapted to IoT environments, where the agent acting on the user's behalf can compete against potential buyers of the collected data — if the user wins the auction, the data will remain private.

Conclusions

In the words of Abraham Lincoln, "You can fool all the people some of the time, and some of the people all the time, but you cannot fool all the people all the time." The vision of IoT as a truly universal part of human society would require a buy-in from everybody, all the time. Economic forces will ensure that businesses will carefully evaluate the costs of participating in the IoT and the values extracted from it, and they would withdraw from ventures that do not have a positive balance. In this article, we argue that the same principles apply on the customer side as well — in order to acquire the consent of the customers for extended periods of time, the overall balance of values and costs needs to be positive. The cost of privacy can be a significant part of the customers' costs and can only be ignored for a limited time or for limited groups of uninformed customers. We argue that the overall buy-in can be regulated by seeing the exchange of information in the IoT exchanges as a reciprocally beneficial trade. We feel that the IoT information exchange cannot be prescribed in detail by external authorities, but laws and regulations can establish a safe and predictable playing field in which these trades can take place.

Let us conclude with the insight that the issues discussed in this article only scratch the surface of the challenges brought by IoT. Our focus was on individual IoT transactions involving data interchange, and the benefits and costs that are incurred in an individual transaction. There are significant challenges about the afterlife of that data: the security and trustworthiness of the cloud where the data is uploaded, the rights of businesses to share data, as well as the legal and liability issues stemming from data ownership. The reader should be referred to other papers that cover these issues from several perspectives. Reference [11] provides a thorough overview of the protocols and technologies involved in IoT, and its interrelation with other emerging technologies such as big data, cloud, and fog computing. Reference [12] assembles a strategic IoT research roadmap as seen by European researchers, while [13] looks at a similar problem from a Chinese perspective. The survey in [14] also brings Korean, Indian, and European viewpoints to the IoT research challenges. Finally, [15] looks at the IoT phenomena from the point of view of enterprises and investment opportunities.

References

[1] L. A. Cutillo, R. Molva, and T. Strufe, "Safebook: A Privacy-Preserving Online Social Network Leveraging on Real-Life Trust," *IEEE Commun. Mag.*, vol. 47, no. 12, Dec. 2009, pp. 94–101.

[2] R. Baden *et al.*, "Persona: An Online Social Network with User-Defined Privacy," *ACM SIGCOMM Comp. Commun. Review*, vol. 39, no. 4, 2009, pp. 135–46.

[3] R. A. Howard, "Information Value Theory," *IEEE Trans. Systems Science and Cybernetics*, vol. 2, no. 1, 1966, pp. 22–26.

[4] C. Bisdikian *et al.*, "Building Principles for a Quality of Information Specification for Sensor Information," *Proc. IEEE Int'l. Conf. Info. Fusion*, July 2009, pp. 1370–77.

[5] D. Turgut and L. Bölöni, "IVE: Improving the Value of Information in Energy-Constrained Intruder Tracking Sensor Networks," *Proc. IEEE ICC*, June 2013, pp. 6360–64.

[6] W. Enck *et al.*, "Taintdroid: an Information Flow Tracking System for Real-Time Privacy Monitoring on Smartphones," *Commun. ACM*, vol. 57, no. 3, 2014, pp. 99–106.

[7] C. L. Miltgen and D. Peyrat-Guillard, "Cultural and Generational Influences on Privacy Concerns: A Qualitative Study in Seven European Countries," *Euro. J. Info. Systems*, vol. 23, no. 2, 2014, pp. 103–25.

[8] K. S. Schwaig *et al.*, "A Model of Consumers Perceptions of the Invasion of Information Privacy," *Info. & Mgmt.*, vol. 50, no. 1, 201, pp. 1–123.

[9] J. P. Carrascal *et al.*, "Your Browsing Behavior for a Big Mac: Economics of Personal Information Online," *Proc. 22nd ACM Int'l. Conf. World Wide Web*, 2013, pp. 189–200.

[10] J. Staiano *et al.*, "Money Walks: A Human-Centric Study on the Economics of Personal Mobile Data," *Proc. 2014 ACM Int'l. Joint Conf. Pervasive and Ubiquitous Computing*, 2014, pp. 583–94.

[11] A. Al-Fuqaha *et al.*, "Internet of Things: A Survey on Enabling Technologies, Protocols, and Applications," *IEEE Commun. Surveys & Tutorials*, vol. 17, no. 4, 2015, pp. 2347–76.

[12] O. Vermesan *et al.*, "Internet of Things Strategic Research Roadmap," *Internet of Things: Global Technological and Societal Trends*, vol. 1, 2011, pp. 9–52.

[13] S. Chen *et al.*, "A Vision of IoT: Applications, Challenges, and Opportunities with China Perspective," *IEEE Internet of Things J.*, vol. 1, no. 4, 2014, pp. 349–59.

[14] D. Singh, G. Tripathi, and A. J. Jara, "A Survey of Internet-of-Things: Future Vision, Architecture, Challenges and Services," *IEEE World Forum on Internet of Things*, 2014, pp. 287–92.

[15] I. Lee and K. Lee, "The Internet of Things (IoT): Applications, Investments, and Challenges for Enterprises," *Business Horizons*, vol. 58, no. 4, 2015, pp. 431–40.

Biographies

Damla Turgut [M] (turgut@cs.ucf.edu) is an associate professor of computer science at the University of Central Florida. She received her B.S., M.S., and Ph.D. degrees from the Computer Science and Engineering Department, University of Texas at Arlington. Her research interests include wireless ad hoc, sensor, underwater, and vehicular networks, cloud computing, as well as considerations of privacy in the Internet of Things. She is a member of ACM and the Upsilon Pi Epsilon honorary society.

Ladislau Bölöni [SM] (lboloni@cs.ucf.edu) is a professor of computer science at the University of Central Florida. He holds a Ph.D. in computer science from Purdue University. His research interests include artificial intelligence, machine learning, robotics, and wireless networking. He is a member of ACM and AAAI.

IEEE ComSoc™
IEEE Communications Society

Join our Community!

IEEE

Networking • Conference Discounts • Technical Publications • Volunteer

Special Member Rates

50% off Membership for new members.

Offer valid March through 15 August 2017.

Member Benefits and Discounts

Valuable discounts on IEEE ComSoc conferences

ComSoc members save on average $200 on ComSoc-sponsored conferences.

Free subscriptions to highly ranked publications*

You'll get digital access to IEEE Communications Magazine, IEEE Communications Surveys and Tutorials, IEEE Journal of Lightwave Technology, IEEE/OSA Journal of Optical Communications and Networking and may other publications – every month!

*2015 Journal Citation Reports (JCR)

IEEE WCET Certification program

Grow your career and gain valuable knowledge by Completing this certification program. ComSoc members save $100.

IEEE ComSoc Training courses

Learn from industry experts and earn IEEE Continuing Education Units (CEUs) / Professional Development Hours (PDHs). ComSoc members can save over $80.

Exclusive Events in Emerging Technologies

Attend events held around the world on 5G, IoT, Fog Computing, SDN and more! ComSoc members can save over $60.

If your technical interests are in communications, we encourage you to join the IEEE Communications Society (IEEE ComSoc) to take advantage of the numerous opportunities available to our members.

Join today at www.comsoc.org

ADVANCES IN NETWORK SERVICES CHAIN: PART 1

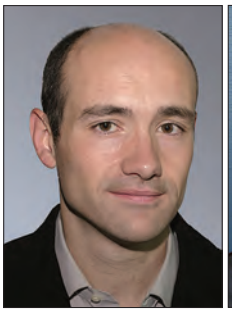
Jordi Mongay Batalla

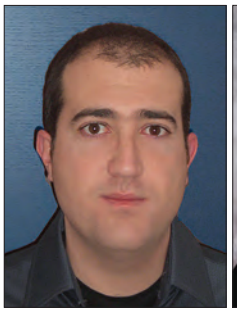
George Mastorakis

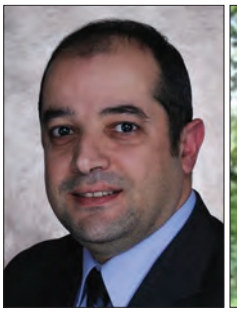
Constandinos X. Mavromoustakis


Ciprian Dobre


Naveen Chilamkurti


Stefan Schaeckeler

Small and large enterprise networks suffer constant pressure updating infrastructure to ever changing trends, techniques, and requirements used in data processing. In a typical enterprise, a data flow arrives on one of its edge routers and then traverses a "service graph" within the enterprise network: it undergoes a series of operations involving Network Address Translation (NAT), access control lists (ACLs), and firewalls before reaching the ultimate destination. The future of enterprise networking is an open software-centric and user-centric approach. As such, enterprises will finally get back complete control over their data. This will be a rocky road, though. Large, globally operating service and content providers push their own data access and governance technologies and/or standards.

Current trends in open software-centric and user-centric networking are mostly based on distributed and decentralized architectures such as software defined networking (SDN) and network functions virtualization (NFV). These are only partial solutions, though. SDN fails to open the platform, and enterprises still cannot completely manage their own data. The situation is the same with NFV, where applications are solely controlled by service providers.

This is why the research is shifting to microservices, containers, and blockchains as the basis for the new networking service chaining:

- Microservices are independently deployable components of an application. Each microservice is designed to scale independent of each other, and service boundaries are clean and open with the possibility to write microservices in different languages.
- Operating-system-level virtualization allows running isolated user space instances such as containers. Containers can be deployed easily and communicate over the network in a distributed fashion.
- Blockchains treat distributed data as large and fully open or permissioned databases where data, once confirmed, cannot be cancelled or modified.

This Feature Topic shows how enterprise applications are built with these technologies and integrate in current networks. The Feature Topic is divided into two parts: The first one focuses on the role of new technologies in network services chains, whereas the second part focuses on the construction of platforms for enterprise networking on the basis of service chains.

The five articles included in the first part propose blockchains, microservices, and containers for providing advanced networking functionalities.

Even though blockchains are universally known as the platform for the digital currency bitcoin, they can be used more generally for services demanding high security and trust. This is the idea behind the article "A Model For Collaborative Blockchain-Based Video Delivery Relying on Advanced Network Services Chains," where Herbaut and Négru adopt blockchains as a virtual (network) function providing negotiations and multimedia service establishment between content providers and occasional end users. In the article "DistBlockNet: A Distributed Blockchains Based Secure SDN Architecture for IoT Networks," Sharma *et al.* incorporate blockchains into an Internet of Things (IoT) network. The DistBlockNet system ensures that end users have (initially non-confident) trust in each other forwarding IoT traffic, avoiding some security threats to IoT constrained devices. Likewise, other technologies for the Internet of Things are presented in "MicroThings: A Generic IoT Architecture for Flexible Data Aggregation and Scalable Service Cooperation." Wang *et al.* propose to decompose IoT implementation in do-it-yourself modules integrated at the sensor level (data aggregation, processing) and at the application level (storage). They base the microthings development on the standardization framework proposed in the article.

Processing big data in a network comes with strict security requirements to counter valid privacy concerns. The simplest solution would be taking the data out of the network; however, this makes user control difficult. New solutions need to be found. In the article "Big Data Orchestration as a Service Network," Liu *et al.* discuss the necessary orchestration tools for bringing big data management into a network service chain.

In the past, microservices have been used for cloud-based health services. Recently, containers are becoming a popular tool for deploying cloud services. Containers make microservices easy to use and particularly interesting for susceptible environments which are, say, under attack. In the article "Security and Privacy for Cloud-Based Data Manage-

ment in the Health Network Service Chain: A Microservice Approach," Esposito *et al.* propose microservices and containers for cloud data management and show the potential gain in security.

The second part of this Feature Topic will show the potential of network service chains in creating novel applications and services across enterprises and society.

BIOGRAPHIES

JORDI MONGAY BATALLA (jordim@interfree.it) is currently the head of the Internet Technologies and Applications Department at the National Institute of Telecommunications. He is also with Warsaw University of Technology. He is an editor of several books and an author or co-author of more than 150 papers published in international journals and conferences in the fields of technologies (radio: 4G and 5G; wired: network services chain, SDN; and applications (Internet of Things, Smart Cities, multimedia) for the future Internet.

GEORGE MASTORAKIS (gmastorakis@ieee.org) obtained his M.Sc. in telecommunications from University College London, United Kingdom, in 2001 and his Ph.D. degree from the University of the Aegean, Greece, in 2008. He currently serves as an associate professor at the Technological Educational Institute of Crete, Greece. His research interests include mobile networks, multimedia applications and services, cognitive radio networks, radio resource management, network management, quality of service, the Internet of Things, and energy-efficient networks.

CONSTANDINOS X. MAVROMOUSTAKIS [SM] (mavromoustakis.c@unic.ac.cy) is currently a professor with the Department of Computer Science at the University of Nicosia, Cyprus, where he leads the Mobile Systems Lab (MOSys Lab., http://www.mosys.unic.ac.cy/) at the Department of Computer Science at the University of Nicosia. He is an active member (Vice-Chair) of IEEE/Region 8 Cyprus Section since January 2016, and since May 2009 he has served as the Chair of the C16 Computer Society Chapter of the Cyprus IEEE Section.

CIPRIAN DOBRE (ciprian.dobre@cs.pub.ro) is a professor at University Politehnica of Bucharest. He leads the activities within the Laboratory on Pervasive Products and Services, and MobyLab. His research interests involve mobile wireless networks and computing applications, pervasive services, context awareness, and people-centric sensing. He is Director or PI for national and international research projects, and has received the IBM Faculty Award, CENIC Awards, and Best Paper Awards. He serves on the Steering and Organization Committees of major conferences.

NAVEEN CHILAMKURTI (n.chilamkurti@latrobe.edu.au) is currently Cybersecurity program co-ordinator in the Department of Computer Science and Information Technology, La Trobe University, Melbourne, Australia. He is also the inaugural Editor-in-Chief of the *International JWNBT*. He has published about 200 journal and conference papers. Some of his publications are in IEEE journals. His current research areas include intelligent transport systems, vehicular security, the Internet of Things, SDN, smart grid, and security in wireless networks.

STEFAN SCHAECKELER (sschaeck@cisco.com) works for Cisco Systems, Inc, San Jose, California. He received his Ph.D. degree from Santa Clara University, California, in 2010. He has published 12 papers in refereed journals and conference proceedings. His works for Cisco on edge, core, and data center routers, helping to pave the road for next generation Internet technologies.

A Model for Collaborative Blockchain-Based Video Delivery Relying on Advanced Network Services Chains

Nicolas Herbaut and Daniel Negru

The authors study how blockchain-powered smart contracts and network service chaining can be exploited to support such novel collaboration schemes. Their findings suggest that the proposed solution can complement existing technologies by supporting a wide range of business cases while at the same time significantly reducing costs.

ABSTRACT

The constant rise of over-the-top video consumption nowadays challenges the current Internet architecture. In this article, we propose a user-centric approach that helps the necessary reshaping of the content delivery ecosystem. We study how blockchain-powered smart contracts and network service chaining can be exploited to support such novel collaboration schemes. Finally, our findings suggest that the proposed solution can complement existing technologies by supporting a wide range of business cases while at the same time significantly reducing costs.

INTRODUCTION

In 2001, Marc Prensky coined the term "Digital Natives," describing what he perceived as a discontinuity in the education world due to the arrival and rapid dissemination of digital technology in a class of age.

Coming to adulthood, this generation is reshaping the TV industry by adopting over-the-top (OTT) services as their primary channel to consume a ubiquitous, on-demand, user-centric entertainment experience. Since the 2010s, this phenomenon has taken off so well that, according to Cisco VNI forecast [1], IP video will reach 82 percent of all IP traffic in 2020. Confronted with the challenges of delivering high-quality content to an ever growing number of users, a new type of architecture started to emerge. This layered delivery architecture promotes a clear separation between:

1. Hardware vendors
2. Content personalization systems
3. Content owners
4. Content providers (CPs)
5. Technical enablers (TEs)
6. Internet service providers (ISPs)
7. End users

In the near future, this model will be strongly challenged, given the current trends toward vertically integrated services. For example Netflix stopped using third party content delivery network (CDN) providers, relying exclusively on its own Open Connect system, making a single company responsible for recommending, selling, producing, owning, and delivering content [2].

In [3], Chuang advocates for future Internet architectures to be "designed for competition," as a means to achieve greater health and sustainability for the network.

The main factor toward ensuring such a design is to *permit different players to express their preferences for a service delivered by various providers*. Following this nomenclature, we identify, in Fig. 1, the six *loci* of competition of the content delivery market. Businesses often span over several competition loci, leaning toward more vertically integrated services. Controlling a certain locus has repercussions on others. For example, an end user cannot choose an alternative TE once he has chosen a CP. This article is focused on the most competition-challenged ones.

To represent the dynamics behind content delivery, Fig. 2 shows the functional interactions between stakeholders. First, the end user initiates a content query; then the CP, TE, and ISP collaborate to run a "content session" representing the actual consumption of the media by the end user. This schema highlights the current status quo in content delivery, but at the same time, it can also serve as the starting point of a more competitive ecosystem design, where:

- The end user expresses his desire to watch a specific content, along with quality of experience (QoE) specifications (e.g., minimal video resolution) in an enriched content query.
- Several CPs read the query and respond with a content offer.
- Several TEs offer their collaboration on the content delivery session, each relying on different technologies and network configurations.

The best content session should be dynamically negotiated between actors providing the desired QoE at the lowest cost.

The cost can be broken down into three parts:

1. The licensing cost charged by the CP to provide access to the content
2. The delivery cost charged by the TE for hosting and delivering the content
3. The network cost charged by the ISP to the TE to transfer the content to the end user

Implementing a trusted, scalable platform able to handle negotiation messages from different stakeholders and process them according to specific business rules can be highly challenging. In this context, the blockchain is perceived as an efficient novel software architecture building block

Digital Object Identifier:
10.1109/MCOM.2017.1700117

The authors are with the University of Bordeaux.

0163-6804/17/$25.00 © 2017 IEEE

that allows reaching a distributed consensus for transactional data without the need for a trusted centralized party [4]. It consists of a read and append-only distributed database that maintains a list of records, called blocks, secured from tampering and revision as each block contains a timestamp and a link to the previous block. Blockchain offers the assurance that data cannot be modified retroactively once recorded. A decentralized consensus can be achieved using specific algorithms such as proof-of-work, proof-of-stake, or Practical Byzantine Fault Tolerance (PBFT). Blockchains can be used in a wide variety of use cases, such as monetary transactions like Bitcoin [5], medical records, and even network control [6].

This article proposes a model for collaborative blockchain-based video delivery. First, a decentralized brokering mechanism is introduced to create content sessions through the collaboration of a CP and a TE. Second, dynamic service chains are exploited in order to benefit from link diversity of different TEs, including user-centric resources.

A Model for Collaborative Video Delivery Based on Blockchain and Network Functions Virtualization Concepts

In this section, we describe a model uising a blockchain to implement a decentralized brokering mechanism enabling a CP and a TE to compete and collaborate for the instantiation of the best content delivery session. After negotiation, this session is implemented as a network service chain the composition of which depends on the underlying technology of its network functions components — legacy CDN server, ISP virtual CDN-as-a-virtual network function (VNF), or user-centric VNF.

A Blockchain-Based Content Delivery Management Mechanism

From Blockchain to Smart Contracts: Current popular implementations of blockchains, such as the one supporting Bitcoin, have been successful at handling simple monetary transactions. However, the lack of native support for advanced programmability encouraged the development of a new generation of blockchain, extending the semantics of transaction through "smart contracts." Written in a Turing-complete language, smart contracts can process data on-chain to implement complex business rules. They can be useful in automating business processes in a trusted way, by allowing all stakeholders to process and validate contractual rules as a group [7].

Implementing Content Delivery Processes with Smart Contracts: We envision a content delivery brokering mechanism as a series of small smart contracts. Each contract has a unique identifier and some data fields, and can perform actions such as creating a new contract or updating the state of the blockchain. Contracts actions are triggered off by on-chain data update (i.e., creation of a new contract) or time.

The proposed model, as shown in Fig. 3, is composed of several blockchains, each one implementing a specific feature used for content distribution, as follows:

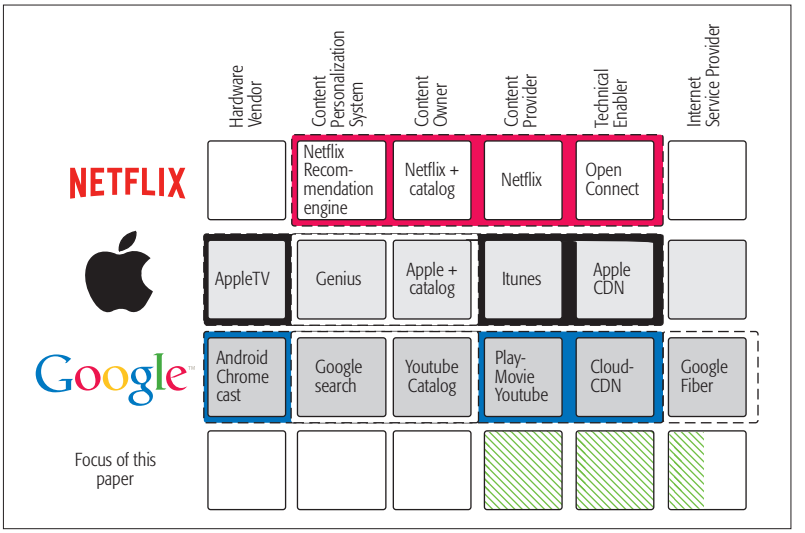

Figure 1. Competition loci in OTT content delivery.

- The **content brokering blockchain** handles the negotiation of the content delivery session. End users, CPs, and TEs publish smart contracts that will be used to determine the best mix for the session.
- The **delivery monitoring blockchain** collects and processes proofs of fulfillment of the delivery contract.
- The **provisioning blockchain** is used by CPs to handle the diffusion of contents on the TE's storage devices.

Content Delivery Session Brokering: Once the query arrives in the blockchain, a content brokering contract (CBC) is created (Fig. 3, step 1) and published. This contract specifies which content *c* to deliver and some user preferences, such as the expected target quality (e.g., 1080p). Then the CPs are notified of the new CBC contract, and use it to create content licensing contracts (CLCs) (step 2). The CLCs specify the price at which each CP is ready to sell content *c* to the end user, a reference to the CBC, and the maximum price for delivery. Next, once the CLCs are visible on the blockchain, TEs respond by publishing *content delivery contracts* (CDCs) (step 3), which specify the cost they are willing to charge for delivering content *c* to the user and the reference to the CLC. Finally, the original CBC collects all the related CDCs and arbitrates toward the cheapest one (step 4). All other contracts are terminated, and the winning contract is used to implement the content delivery. Relevant technical information required to implement the contract, such as content ID, TE ID, and end-user IP, are compiled in a content delivery service description (CDSD) document. We later detail how the contract is used to configure network service chains.

Content Delivery Session Monitoring: Smart contracts can implement a currency system to be used as a collateral means for ensuring the correct execution of the content delivery. Once the collaboration between each actor is formalized in a CDC, the payers (end user and CP) transfer their due payments to the CDC, which behaves as an escrow account. Each partner sends a proof of activity to the delivery monitoring blockchain according to its role in the content delivery. For

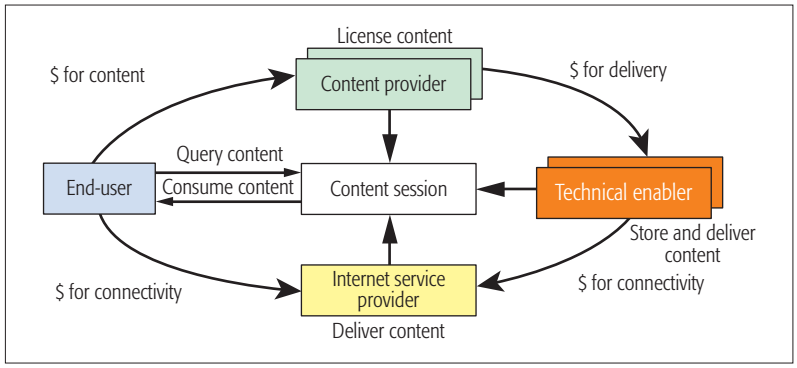

Figure 2. Stakeholders' interactions in the content session.

example, the CP could publish a cryptographic proof (e.g., as in the case of digital rights management) that entitles the TE to deliver the content. The end user publishes a proof of reachability of the content, whereas the TE publishes a proof of transmission. Once all the proofs are collected, the beneficiaries receive their payments. If the contract detects that a party has not fulfilled its duties, penalties can be applied, and some of the initial payment is refunded to payers.

Content Provisioning: Content dissemination throughout the network is key to reducing the price of delivery, and is usually achieved thanks to a resource prediction engine [8]. The blockchain can be used for content provisioning in TEs in two ways. The TEs can audit the content popularity from chain data and proactively decide to *pull* contents from the CP to subsequently sell CDCs. Alternatively, the CPs can *push* contents to the TEs by rewarding them through the blockchain in compensation for storage. CPs benefit from having their content widespread on the network, since more TEs publishing a CDC means more competition and lower price.

Governance Models

As the proposal relies on a fully decentralized agreement conclusion mechanism, we need a way to establish the respective liabilities of stakeholders in case of problems. As smart contracts are not legal contracts in essence, any litigation should be solved by proper prior legal agreements. Several models can be considered:

- Chain of responsibility: Each actor contracts with a supplier, which is liable for the service it provides. CPs are liable toward end users, TEs are liable toward CPs, and TEs are liable toward ISPs. This solution is not very scalable as it implies having thousands of contracts.
- Consortium: Actors create a consortium providing the legal foundations for the service. The consortium manages any liabilities centrally and automatically thanks to the blockchain. This model opposes the decentralization of transactions, but offers a more scalable alternative.
- Decentralized autonomous organization: In this model, legal aspects are directly managed on-chain by an organization the governance of which is defined by the code of smart contracts, bringing full decentralization and automation. However, the legal status of this type of business organization is still unclear.

Our proposal fosters competition by allowing several actors to offer their resources to the system and adjust their prices to match demand. By decoupling the content delivery from the content licensing, we set up a much more diverse ecosystem, by including actual end users (assuming the role of TEs) in the content delivery process. However, constructing content sessions by using third party resources induces a challenge to current Internet architectures. In the next section, we describe how content sessions can be dynamically mapped to network service chains through network softwarization and the use of microservices.

Instantiating the Model through Advanced Dynamic Network Service Chains

Once the brokering of content licensing and delivery is complete, the content session between the TE and the end user is implemented. Content sessions are on-demand, user-centric service chains deployed based on the specifications of the CDSD. The deployment of the service chain is shared between ISPs and TEs, the ISPs being responsible for steering the traffic of the end user to/from the TE domain, while the TEs implement both networking and service configuration of IP endpoints.

TEs implement content delivery in several ways. We detail three complementary approaches, CDN, vCDN, and μCDN, detailed below. Figure 4 shows the deployment of these three types of TE. User 1 gets its content from a CDN, while User 2 uses a vCDN service chain deployed in the ISP network. User 3 retrieves its content directly from user 4's μCDN.

CDN Delivery: This is the classical case used today in OTT scenarios where the content is hosted on servers belonging to another autonomous system (AS) with no possible end-to-end management. The content is delivered in best effort mode from the server, selected by the CDN operator depending on the end user's physical location through dynamic Domain Name Service (DNS) resolution. Little to no collaboration occurs in this scenario, and the server selection made by the CDN operator may not be in line with the traffic engineering objectives implemented by the ISP. To circumvent this issue, the next two sections present the deployment of network service chains where the network is handled by the ISP end to end, allowing the implementation of quality of service policies.

vCDN Delivery: This solution aims at instantiating a CDN as a VNF inside the ISP AS, deployed on a network functions virtualization (NFV) infrastructure point of presence (NFVI-POP). The vCDN can be operated by the ISP [9] or by an external CDN operator leasing the ISP infrastructure. This approach reduces the hop count between client and server, and supports end-to-end management of the service through a service level agreement (e.g., imposing minimal bandwidth or maximal delay). Our previous work [10] described a network service chain that can be used to handle both routing (with the virtual media gateway — VMG) and content distribution (with the Virtual Streamer — vStre). Deploying a service over an NFV infrastructure over multiple data centers [11] usually relies on software defined networking (SDN) for flexibility.

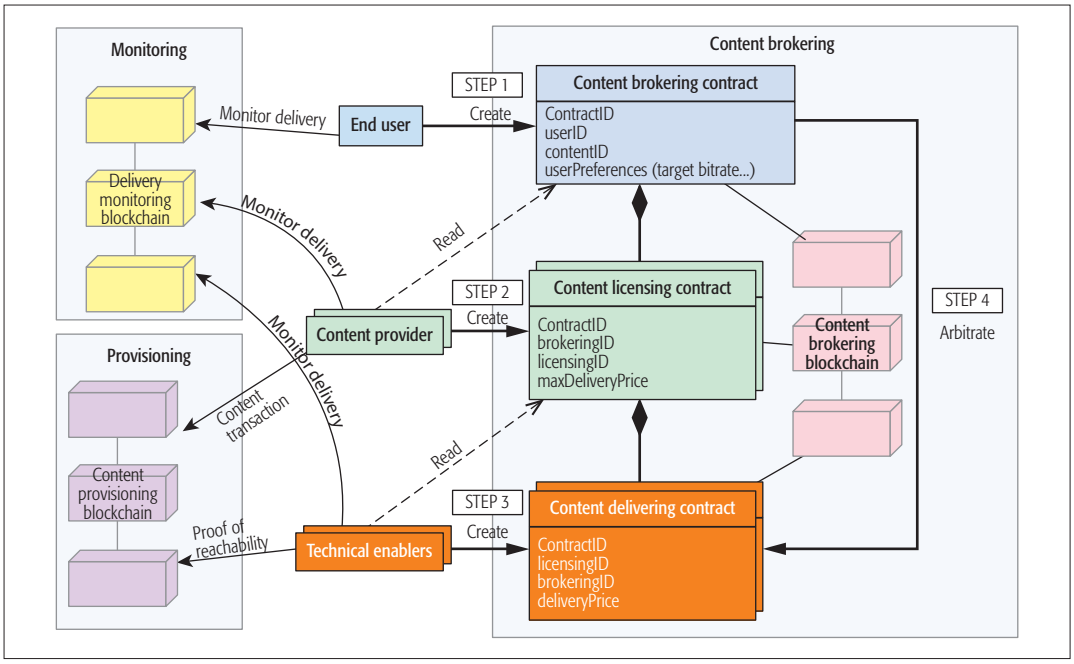

Figure 3. Blockchain-based model for collaborative video delivery.

The Blockchain can be used for content provisioning in TEs in two ways. The TEs can audit the content popularity from chain data and pro-actively decide to pull contents from the CP to subsequently sell CDCs. Alternatively, the CPs can push contents to the TE by rewarding them through the Blockchain in compensation for storage.

µCDN Delivery: Customer premises equipments (CPE) provides plenty of spare system and network resources that can be used for content delivery. With modern operating systems (OSs, e.g., GNU/Linux, Android) they can support the deployment of new services and even VNFs [12]. Their small scale, however, requires downscaling the main concepts behind NFV.

Figure 4 shows the internal microservice architecture used to implement the µCDN. The two key technologies we use to address the above-mentioned challenges are containers and SDN, as follows:

- Containers are lightweight virtualization mechanisms that bundle applications and their dependencies. With their reduced footprint and low CPU overhead, they are often considered in cloud edge architectures [13].
- The SDN-capable software switch deployed on the CPE allows manipulating the containers' connectivity in an OS-independent fashion thanks to the use of standard protocols such as OpenFlow.

Service deployment is triggered by the publication of the CDSD on the blockchain (Fig. 4, 1). It is retrieved by the µOrchestrator module, which spawns content delivery containers running HTTP servers able to stream the content (2a) and configures their network (2b). A CP may choose to use different technologies to license their content, from simple files to more complex DRM-based solutions. Adopting a microservice architecture, our solution keeps these implementation details in the content delivery container and ensures that, regardless of the underlying technology, the CDSD provides all resources needed (e.g., cryptography material). Finally (3a), the µOrchestrator instructs the SDN controller to update the network configuration, which is, in turn (3b), deployed by the software switch so that the connection between the end user and the content delivery container can be established.

Evaluation of the Proposed Model

Network Services Chain Evaluation

We implemented a discrete event simulator with the SimPy Library to emulate content delivery sessions. We simulated 15,000 content session requests spanning over 25 minutes. For every request, each TE that:

1. Stores the content
2. Has enough bandwidth to deliver the content

asks for a delivery price assumed to be proportional to the number of hops between itself and the end user.

The brokered price corresponds to the smallest price demanded by a TE. CDNs were assumed to host the entire content catalog, whereas µCDN and vCDN pulled the content from the CP by auditing blockchain data and downloading the most popular contents. We used a real ISP topology of 2k nodes and 60k edges extracted from the Center for Applied Internet Data Analysis. Six CDNs were placed in a weighted random fashion at the most connected links, which correspond to the Internet exchange points on the operator topology. We then placed 500 service access point nodes representing the user location in the network in a similar way, selecting the least connected links. Finally, 100 vCDNs and 500 µCDNs were randomly distributed among the nodes with connectivity degrees in the middle range (40–90 percent). vCDNs' capabilities were based on common virtual caching appliances' specifications (1 TB of storage supporting 150 Mb/s or 30 concurrent 720p streaming sessions), while µCDN capabilities were based on current CPE specifications (30 GB of storage, 20 Mb/s of upload speed, or 4 concurrent 720p streaming sessions). Contents stored in µCDNs and vCDNs are purged according to a least recently used rule. CDNs were assumed to support a large amount of concurrent connections (2.5 Gb/s or 500 sessions). We assumed content popularity to follow a

The need for performance and Smart Contracts support compelled us to use the open source project Hyperledger-Fabric (www.hyperledger.org/projects/fabric). This Linux foundation project can be used to build Blockchain solutions with a modular architecture to deliver flexibility and scalability.

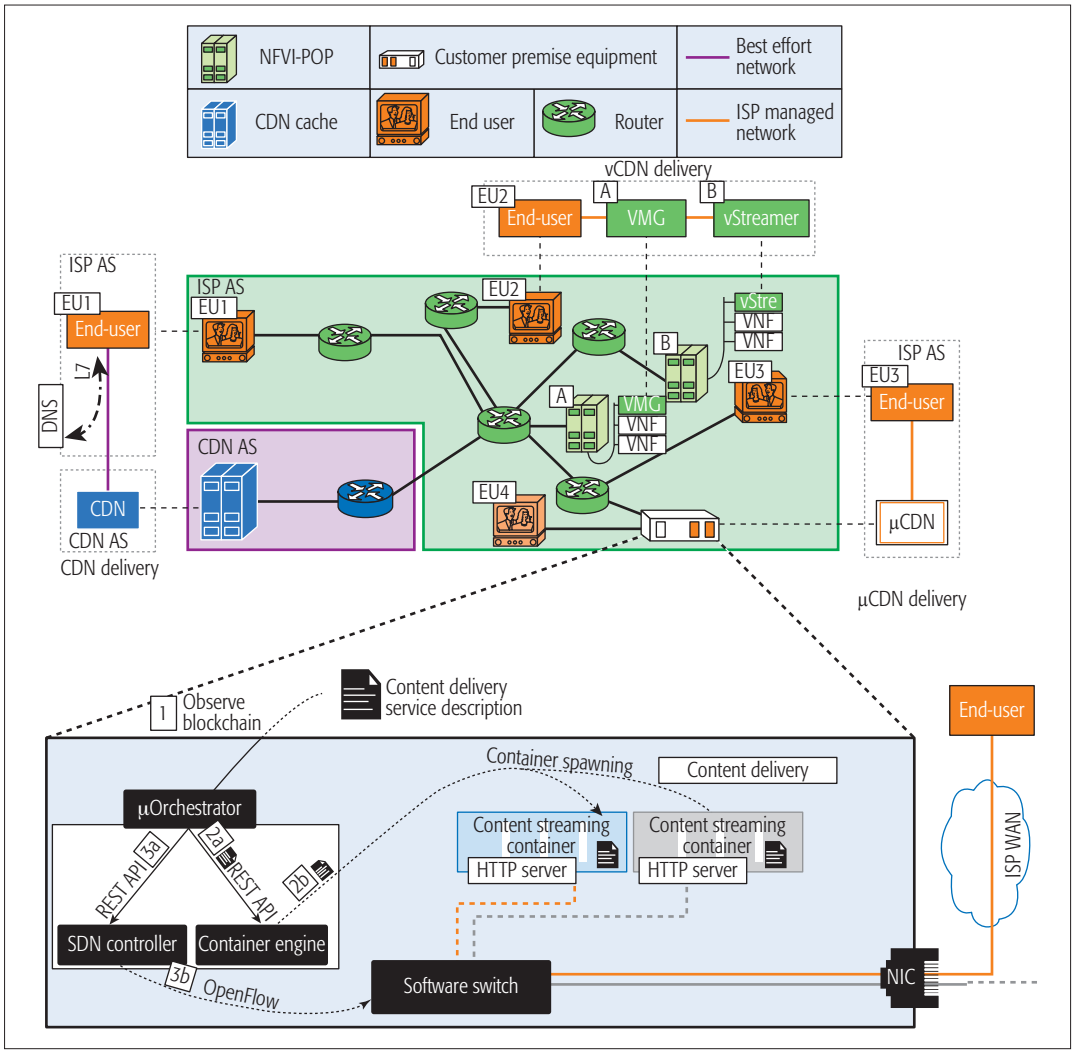

Figure 4. CDN, vCDN, and µCDN services deployed in an ISP network.

Zipf distribution. The hop count is computed from the topology for vCDNs and µCDNs; however, for the CDN, we assume that three additional hops are used within the CDN network between the edge of the ISP network and the final server, corresponding to the average ISP graph distance.

Results of the experiments are presented in Fig. 5. Figure 5a shows the respective shares of TEs. At the beginning we see that every request is served by the CDNs, as they are still the only ones hosting the content. After 2 min, once the popular contents are downloaded by the vCDNs, they also start delivering contents. The reason vCDNs are privileged w.r.t. CDNs is that they are spread more widely in the network, with a smaller average distance to end users. After 3 min, the µCDNs start serving content as well, and their share increases to 12 min, where they become the most used TEs. Again, this can be explained by a denser distribution of µCDNs in the network causing a lower hop count. After the 20 min mark, the shares stabilize. Despite their advantage, µCDNs only absorb half of the content requests. In fact, due to their limited capacity and storage, they are able to store and deliver only very popular contents. vCDNs store both very popular contents and less popular contents and still account for a third of content sessions. Finally, CDNs absorb the long tail of contents that are not popular enough to be stored by other TEs.

Another important benefit of our solution is the hop count reduction. Figure 5b compares the average number of hops between the selected TE and the end user when using all three TE types in conjunction, but also using only some of them. We can see from the figure that using only the CDNs yields a higher hop count, stable over time. When complementing a CDN with vCDNs, the hop count sharply decreases, as content gets stored near the edges of the network, and stabilizes near the four-hop mark. When using both CDN and µCDN, the curve decreases slowly, as contents take more time to be provisioned at the edges. Finally, using all three TEs yields the lowest hop count, with a fast drop at the beginning and a downward trend reaching the lowest value of our experiment.

Blockchain Evaluation

Test Environment: Our goal is to build a system where each content session is brokered on the blockchain. For this reason, its *performance*, measured in terms of number of transactions processed per second, is key to providing the content sessions quickly. At the same time, we envision the number of "clients" (end users, CPs, and TEs)

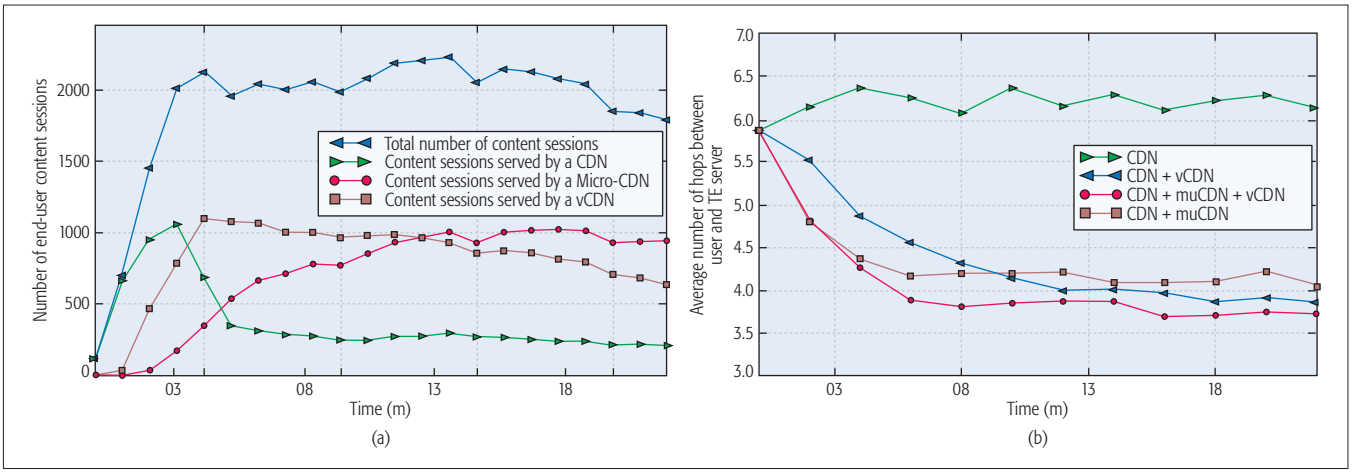

Figure 5. Network services chain evaluation: a) respective TEs share for CDC; b) average price for content delivery.

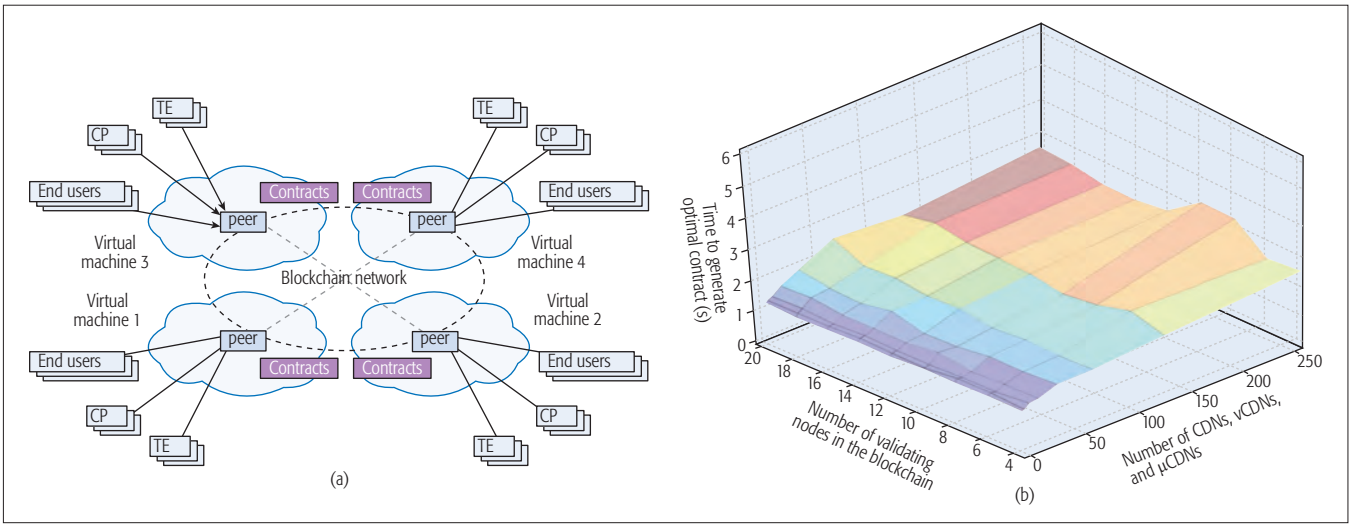

Figure 6. Blockchain evaluation: a) testbed; b) performance and scalability experiment.

using the service to be high, so the blockchain must ensure good *node scalability*. Today, "permissionless" blockchains based on proof-of-work consensus offer great node scalability, but lack the required throughput (e.g., up to 7 tx/s with Bitcoin). On the other hand, blockchains based on advanced Byzantine Fault-Tolerant (BTF) state-machine replication protocols offer excellent performance in terms of throughput and latency but require all nodes to know the IDs of all other nodes [14]. In our case, we used a "permissioned" blockchain as the nodes processing the transactions do not need to be anonymous.

The need for performance and smart contract support compelled us to use the open source project Hyperledger-Fabric (www.hyperledger.org/projects/fabric). This Linux foundation project can be used to build blockchain solutions with a modular architecture to deliver flexibility and scalability. It provides pluggable consensus algorithms (by default PBFT) and simple smart contract implementation in Go or Java.

The critical aspects of the brokering mechanism is the time needed to converge toward the optimal CDC, involving the end user, the CP, and the TE. This delay affects the end-user QoE, as the content delivery session can start only after the optimal CDC is computed. A lot of contracts are published in the blockchain; for example, if we assume that there are 10 CPs and 100 TEs, up to 10 × 100 contracts will be published.

Considering this, the evaluation is focused on the content brokering blockchain as it is the most time-sensitive and subject to scalability issues.

We deployed the solution with Hyperledger-Fabric configured with the PBFT consensus, as shown in Fig. 6a. We then paired end-user applications (publishing CBC), CP applications (reading CBC from the blockchain and responding by publishing CLCs) and TE applications (reading CLCs and publishing CDCs), and the blockchain validating peers, which are the nodes responsible for running the consensus, validating transactions, and maintaining the ledger.

Each user was configured to send 10 requests/min. We then computed the average time needed to obtain the optimal CDC, or convergence time. We varied the number of TE agents, with the number of CPs being fixed at 10.

The results presented in Fig. 6b show that for 50 TEs, the convergence time is below 2 s. This time increases for higher values of TEs, reaching 4 s in the worst case scenario of 250 TEs, which remains acceptable.

On our testbed, the number of nodes increased the convergence time slightly. This is due to the rather good networking performance of our cloud instances, located in the same availability zone. In a production deployment, nodes would not be collocated to improve resiliency, and the performance might be even more impacted.

Discussion on Scalability: The blockchain network is composed of validating nodes that run the smart contracts and append blocks to the chain once consensus is reached. They are also used to query the state of the blockchain by clients. Increasing the number of validating nodes has two antagonistic effects:

1. Each node serves fewer clients, reducing the average number of requests per node.
2. The quorum needed for the consensus is increased, increasing the number of messages shared in the network.

On our testbed, the number of nodes increased the convergence time slightly. This is due to the rather good networking performance of our cloud instances, located in the same availability zone. In a production deployment, nodes would not be collocated to improve resiliency, and the performance might be even more impacted.

The next release of Hyperledger-Fabric and recent research papers such as [15] promote new architectures that support parallelizing the validation of transactions through their endorsement by only a subset of nodes. In this perspective, transactions are managed on sub-chains supporting fine-tuned consensus algorithms, improving scalability.

Conclusions and Future Work

This article proposes a new model for content distribution over the Internet, with a scalable blockchain-based brokering mechanism allowing several providers to collaborate and to provide the requested service through network service chains. On top of reducing the overall delivery cost, it promotes healthy competition, allowing user-centric resources to join the game and paving the way for new business models.

Further work will consist of evaluating the scalability of the solution on a production-ready blockchain architecture. The issues of governance, security, and privacy for end users have not been addressed and need further investigation.

Acknowledgments

We thank Marko Vukolić (IBM Research-Zurich), who provided valuable feedback at early stages of this article. The work performed for this article has been partially funded by the FP7 IP T-NOVA European Project (Grant Agreement N#619520) and the FUI French National Project DVD2C.

References

[1] Cisco Visual Networking Index, Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2015–2020 white paper.
[2] T. Böttger *et al.*, "Open Connect Everywhere: A Glimpse at the Internet Ecosystem Through the Lens of the Netflix Cdn," arXiv preprint arXiv:1606.05519, 2016.
[3] J. Chuang, "Loci of Competition for Future Internet Architectures," *IEEE Commun. Mag.*, vol. 49, no. 7, July 2011, pp. 38–43.
[4] X. Xu *et al.*, "The Blockchain as a Software Connector," *2016 13th Working IEEE/IFIP Conf. Software Architecture*, 2016, pp. 182–91.
[5] S. Nakamoto, "Bitcoin: A Peer-to-Peer Electronic Cash System," 2008.
[6] N. Bozic, G. Pujolle, and S. Secci, "A Tutorial on Blockchain and Applications to Secure Network Control-Planes," *Smart Cloud Networks & Systems*, 2016, pp. 1–8.
[7] R. Hull *et al.*, "Towards a Shared Ledger Business Collaboration Language Based on Data-Aware Processes," *Int'l. Conf. Service-Oriented Computing*, Springer, 2016, pp. 18–36.
[8] Y. Kryftis *et al.*, "Efficient Entertainment Services Provision over a Novel Network Architecture," *IEEE Wireless Commun.*, vol. 23, no. 1, Feb. 2016, pp. 14–21.
[9] P. A Frangoudis *et al.*, "An Architecture for On-Demand Service Deployment over a Telco CDN," *IEEE ICC*, 2016, pp. 1–6.
[10] N. Herbaut *et al.*, "Service Chain Modeling and Embedding for NFV-Based Content Delivery," *IEEE ICC*, 2017.
[11] R. Mijumbi *et al.*, "Network Function Virtualization: State-of-the-Art and Research Challenges," *IEEE Commun. Surveys & Tutorials*, vol. 18, no. 1, 2016, pp. 236–62.
[12] D. Minodier and G. Dalle, "Juniper, Network Enhanced Residential Gateway," tech. rep. TR-317, Broadband Forum, July 2016; https://www.broadband-forum.org/technical/download/TR-317.pdf.
[13] C. Pahl and B. Lee, "Containers and Clusters for Edge Cloud Architectures — A Technology Review," *2015 3rd Int'l. Conf. Future Internet of Things and Cloud*, 2015, pp. 379–86.
[14] M. Vukolic, "The Quest for Scalable Blockchain Fabric: Proof-of-Work vs. BFT Replication," *Int'l. Wksp. Open Problems in Network Security*, Springer, 2015, pp. 112–25.
[15] W. Li *et al.*, "Towards Scalable and Private Industrial Blockchains," *Proc. ACM Wksp. Blockchain, Cryptocurrencies and Contracts*, 2017, pp. 9–14.

Biographies

Nicolas Herbaut (nicolas.herbaut@labri.fr) is a Ph.D. candidate at Bordeaux University doing research on virtualization, network softwarization, and content delivery. He received an M.Sc.Eng. in mathematics from INSA Rouen in 2004, and an M.Econ. in economic analysis from Aix-Marseille University in 2005.

Daniel Negru (daniel.negru@labri.fr) is an associate professor at Bordeaux University, specializing in multimedia and networking. His current activities are focused on video streaming, content delivery, and NFV/SDN. He has participated in numerous collaborative research projects at the European level, and published more than 60 papers in journals and conferences such as *IEEE Communications Magazine*, *IEEE Multimedia*, ICME, GLOBECOM, and ICC.

Technology insight on demand on IEEE.tv

Internet television gets a mobile makeover

A mobile version of IEEE.tv is now available for convenient viewing. Plus a new app for IEEE.tv can also be found in your app store. Bring an entire network of technology insight with you:

- Convenient access to generations of industry leaders.
- See the inner-workings of the newest innovations.
- Find the trends that are shaping the future.

IEEE Members receive exclusive access to award-winning programs that bring them face-to-face with the what, who, and how of technology today.

Available on the App Store

Available in Android Market

Tune in to where technology lives www.ieee.tv

IEEE
Advancing Technology for Humanity

DistBlockNet: A Distributed Blockchains-Based Secure SDN Architecture for IoT Networks

Pradip Kumar Sharma, Saurabh Singh, Young-Sik Jeong, and Jong Hyuk Park

The rapid increase in the number and diversity of smart devices connected to the Internet has raised the issues of flexibility, efficiency, availability, security, and scalability within the current IoT network. These issues are caused by key mechanisms being distributed to the IoT network on a large scale, which has motivated the authors to propose DistBlockNet.

ABSTRACT

The rapid increase in the number and diversity of smart devices connected to the Internet has raised the issues of flexibility, efficiency, availability, security, and scalability within the current IoT network. These issues are caused by key mechanisms being distributed to the IoT network on a large scale, which is why a distributed secure SDN architecture for IoT using the blockchain technique (DistBlockNet) is proposed in this research. It follows the principles required for designing a secure, scalable, and efficient network architecture. The DistBlockNet model of IoT architecture combines the advantages of two emerging technologies: SDN and blockchains technology. In a verifiable manner, blockchains allow us to have a distributed peer-to-peer network where non-confident members can interact with each other without a trusted intermediary. A new scheme for updating a flow rule table using a blockchains technique is proposed to securely verify a version of the flow rule table, validate the flow rule table, and download the latest flow rules table for the IoT forwarding devices. In our proposed architecture, security must automatically adapt to the threat landscape, without administrator needs to review and apply thousands of recommendations and opinions manually. We have evaluated the performance of our proposed model architecture and compared it to the existing model with respect to various metrics. The results of our evaluation show that DistBlockNet is capable of detecting attacks in the IoT network in real time with low performance overheads and satisfying the design principles required for the future IoT network.

INTRODUCTION

According to the recent Gartner's report [1], 1 million new Internet of Things (IoT) devices will be sold every hour, and $2.5 million will be spent per minute on IoT by 2021. We believe that the idea of a distributed IoT network is promising. Meanwhile, software defined networking (SDN) empowers easy management and network programmability [2]. Initially, it brings up some issues of security, performance, reliability, and scalability due to the centralized control architecture. Recently, numerous distributed SDN controllers have been introduced to address these issues [3-5]. Most of the existing work emphasizes the issue of state consistency among multiple controllers. The mapping between the controllers and the forwarding devices is statically configured, which can result in uneven distribution of loads between the controllers and bursting packets breaking down the controller. In addition to these issues, we need a low response time and distributed SDN network with high availability. Some methods try to offer a reliable and scalable solution to the distributed network for management [6–10], but none of them have completely solved this problem. On the other hand, blockchains have recently drawn much attention from interested stakeholders in a wide range of industries [11, 12]. The reason behind this explosion of interest is that with the blockchains technique, we can operate the applications in a distributed manner that could previously run through a trusted intermediary. We can accomplish the same functionality with the same assurance without the need for a central authority. The blockchains technique offers a distributed peer-to-peer network where, without a trusted intermediary, untrusted individuals can interact in a verifiable manner with each other [13, 14].

REQUIRE DESIGN PRINCIPLES FOR SECURELY DISTRIBUTED ARCHITECTURE

In order to design high-performance architecture for the IoT network that is securely distributed in order to deal with current and future challenges and satisfy new service requirements, we need to consider the following design principles based on previous work on designing distributed network architecture, research new network technologies, and investigate new service requirements.

Adaptability: Trends are evolving, and the needs of clients are changing. These changing trends require that the network architecture is improved and is able to adapt to the changing environment. Adaptability is vital to ensure its growth and survival. The network architecture should be able to adapt and have its usage broadened with the increase in clients' needs and demands.

High Availability and Fault Tolerance: The high availability of a network control system is important in the actual operation of the network.

Digital Object Identifier:
10.1109/MCOM.2017.1700041

Pradip Kumar Sharma, Saurabh Singh, and Jong Hyuk Park are with Seoul National University of Science and Technology (SeoulTech); Young-Sik Jeong is with Dongguk University.

Thus, the provisioning of a priori redundancies, the detection of failures, and the invocation of mitigation mechanisms are necessary steps for action.

Performance: Ability to adapt performance linearly. In current IoT environments, it is a common challenge to try and achieve linear performance over a large-scale distributed network architecture.

Reliability: When designing the distributed architecture, reliability is ranked as the highest priority. It measures the correlation between the corresponding performance required and the total performance achieved by the system in all environmental conditions of time and space.

Scalability: Scalability is an essential principle in designing a future-proof distributed network architecture, which not only reduces costs, provides the flexibility to extend the network, and supports unexpected services, but also involves the deployment of new services and meets new market requirements.

Security: Security must be everywhere in a distributed network to build a secure distributed architecture that is provided as a service to protect the confidentiality, integrity, and availability of all connected information and resources. Therefore, securing the network must be one of the objectives of designing new distributed architectures.

Research Contributions: On the basis of the discussion above, the main contributions to the research of this work can be summarized as laid out below:

- We are proposing a distributed secure SDN architecture for IoT using the blockchain technique. When using the proposed architecture, security must automatically adapt to the threat landscape without administrators needing to manually review and apply thousands of recommendations and opinions.
- We are proposing a technique for updating the high-performance availability flow rule tables in the distributed blockchain SDN.
- We have evaluated the performance of our proposed technique and compared it to the existing model with respect to various metrics.

The rest of the article is structured as follows. We discuss the proposed distributed secure architecture used for the IoT using the blockchain technique. We also present the architecture workflow and flow rule tables update technique in the distributed blockchain network. Next, we evaluate the proposed model based on different performance metrics. Finally, we conclude the research.

DistBlockNet Distributed Secure Architecture

According to the analysis in the previous section for rapidly growing IoT networks created by the new communication paradigms, we have observed that the currently distributed network architecture, protocols, and techniques are not designed to meet the required design principles for future challenges and satisfy new service requirements. The speed and complexity of this development exponentially creates new categories of attacks; gathering known and mysterious threats; taking advantage of "zero-day" vulnerabilities; and using malware concealed in websites, documents, networks, and guests. At present, organizations need a single distributed secure architecture that includes powerful network security devices with proactive, real-time protection with high performance to meet the analyzed design principles. In this section, we propose a novel distributed secure SDN architecture called DistBlockNet, its workflow, and a technique for updating high-performance availability flow rule tables in a distributed blockchain network.

DistBlockNet Design Overview

DistBlockNet adopts distributed secure network control in the IoT network by using the blockchain technology concept to improve security, scalability, and flexibility, without the need for a central controller. Figure 1 shows the global and local views of the proposed architecture. In the proposed architecture, all controllers in the IoT network are interconnected in a distributed blockchain network manner so that each IoT forwarding device in the network can easily and efficiently communicate. Each local network view comprises OrchApp, Controller, and Shelter modules. The Shelter and OrchApp modules in each local network handle the security attacks at a different level. OrchApp mainly functions at the management or application layers, the controller-application interface, and the control layer. Shelter operates at the data layer, the controller-data interface, and the control layer. The DistBlockNet architecture provides not only operational flexibility, but also proactive and reactive incident prevention based on the recurring threat landscape by inserting the rapidly changing, dynamic, and high-performance OrchApp and Shelter modules. It offers a network infrastructure that is agile, modular, and secure. Protections must dynamically adapt to the threat landscape without having to include security administrators to manually process a huge number of advisories and approvals. These insurances must coordinate well into the more extensive IoT environment, and the architecture must take on a protective stance that cooperatively leverages both savvy inside and outside sources.

OrchApp: Its prime purpose is to offer programming characterized fortifications and to set out them for execution at the appropriate application layer enforcement points, whether implemented using high-performance as host-based software on mobile devices, in the IoT network or the cloud. Security classifications incorporate access control, data protection, and threat intelligence. Based on the underlying domain knowledge from which security strategy plans are drawn, these methods vary.

Access control implements a security convention model of approved associations among resources and clients in the IoT network, as set up by the management layer. On the other hand, *data protection* focuses on the classification of data rather than on behavior and interaction. The management layer concludes the standards or strategies for data flows in the organization. *Threat intelligence* provides the understanding of threats and their behavior. It is powered by applying collaborative intelligence to real-time threats obtained from different communities.

Protections must dynamically adapt to the threat landscape without having to include security administrators to manually process a huge number of advisories and approvals. These assurances must coordinate well into the more extensive IoT environment, and the architecture must take on a protective stance that cooperatively leverages both savvy inside and outside sources.

To identify the attacks in the security policies resulting from the actual changes made to the system data plan, which is linked to each flowchart and to the topological exchange metadata, the graph builder analyzes the parsed dataset to construct and alter the flow diagrams that are connected to the network traffic.

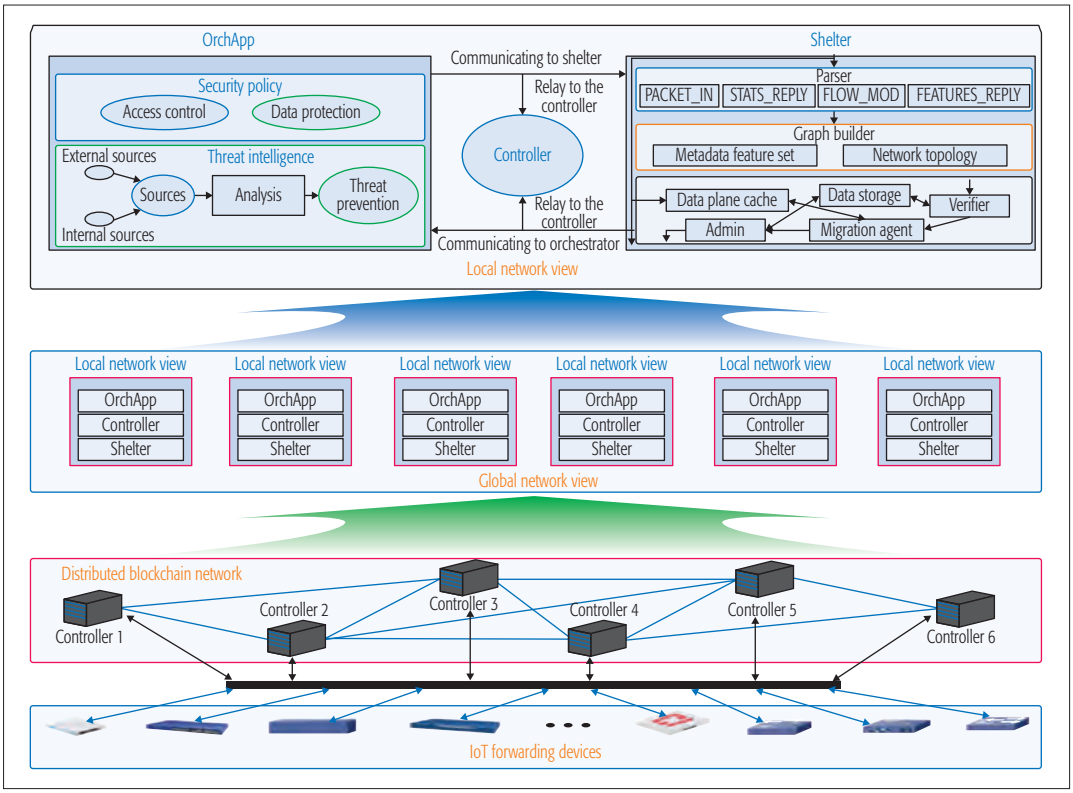

Figure 1. Overview architecture of DistBlockNet.

OrchApp provides the level of adaptability sought to adapt to the new and dynamic threats and modifies the enterprise network configurations. The application layer provides a solid platform that can execute assurances at the application points all throughout the enterprise. Because the protections are software-controlled, critical hardware deployed at these points of the application does not need to be exchanged when a new threat or attack technique is exposed or when new technologies are introduced in the industry. Protections should automatically take place in the threat landscape without requiring manual monitoring of the analysis of a large number of opinions and endorsements. This is accomplished by using an automated threat counteractive action control that works together with the management layer that is only essential for human decision-making for when the threat indicators offer less assurance about recognizing an attack or threat.

Shelter: Attackers often rank in the network to take advantage of the insider's advantage points, and then launch attacks on the internal network. Given that our objective is to assert the appearance of attacks on the network topology and the data plane and the aggression of identity of the flow rules or the strategies within the SDN, our threat system perfectly identifies the scenarios where the antagonist initiates attacks within the SDN. Thus, we designed the SDNs as a non-open system. Removing restrictions on unidentified external communications helps focus our analysis only on OpenFlow control packets or messages within the SDN because the OrchApp handles all of these issues in the DistBlockNet model.

Shelter is composed of a flow control analyzer and packet migration components. The analyzer component takes care of the main functionality of the network infrastructure as soon as the saturation attack has occurred. Whereas, the packet migration component sends a benign network stream to the OpenFlow controller without overloading. As shown in Fig. 1, the module units define the flow analyzer as a control application on the controller platform. Furthermore, the migration agent of the migration component is applied to a controller application between the control plane, the data plane, and an element of the cache data plane.

Parser: The attackers use the subset of OpenFlow messages, such as Packet_in, Flow_Mod, Features_Reply, and Stats_Reply, in order to change the network's view of the controller. Thus, to identify abnormal behavior, we extracted the important metadata by monitoring and parsing incoming packets.

Graph Builder: To identify the attacks in the security policies resulting from the actual changes made to the system data plan, which is linked to each flowchart and to the topological exchange metadata, the graph builder analyzes the parsed dataset to construct and alter the flow diagrams that are connected to the network traffic. Our model retains the flows of logical and physical topologies and Flow_Mod transmission status messages to identify malicious update metadata.

Verifier: We generated path conditions offline and reactive rules online. In order to reduce the overhead at runtime, we processed the path condition generator to navigate the possible paths and to collect all path conditions offline. Online reactive rule generation monitors and assigns the current value of the global variables to the status path. The input variables are symbolized in the path conditions, and the reactive flow rule

dispatcher components are used to parse each status path. It is only with the paths that the final decision is taken into account in processing a small set of modifications to generate a status message. Finally, the reactive flow rules we need are established.

Migration Agent: The migration agent detects attacks and makes the appropriate decisions based on the type of alarms received. In order to generate new rules and migrate the missing table packet in the data cache, it triggers the flow rules of the parser during saturation attacks. It migrates all missing packets to the data plan cache during the generation of flow rules and the update stage. As a result, the controller does not overload itself with the flooding packets. Finally, it processes all the missed packets stored in the cache after the flow rules are updated.

Data Plan Cache: During a saturation attack, it temporarily caches the missing packets. During flooding attacks, most flood packages are redirected to the data plan cache to avoid flooding the controller. By using the classifier, Packet$_i$n generator, and buffer queue, it parses the header of the migrated packets and stores them in the appropriate queue.

Shelter Workflow

As shown in Fig. 1, the Shelter module has three different stages. In the first stage, in order to build a complete network view, Shelter monitors and parses all of the packets communicating with the controller and identifies the appropriate OpenFlow packets. In the second stage, to build an incremental graph network with traffic flow, Shelter analyzes all of these parsed OpenFlow packets to obtain the topological metadata and status of the transmission. Shelter mainly maintains the metadata feature set, the network topological state that is obtained from the OpenFlow packet headers, actual measurements of traffic flow within network connections, and outbound flow path configuration directives, respectively. In the third stage, Shelter allows this metadata to flow against a set of acceptable metadata values collected during the flow period, administrative rules, and strategies. Shelter identifies known attacks through policies specified by the administrator, although it uses precise flow activities obtained over time to detect unplanned and possibly malicious activity.

Shelter does not issue an alarm signal when it detects a new flow behavior. Alternatively, Shelter prompts an alarm signal when it detects untrusted entities that invoke modifiers to the existing flow behavior or where the flow resists any rules or security rules specified by the administrator. Also, Shelter will not raise any alerts on flow reroutes because they are generated by FLOW_MOD messages from the trusted controller. This drastically reduces the alerts that can occur if the recognition of each new behavior is signaled, which is possible in growing networks. However, malicious activity will be noticed by looking back when Shelter later reports authentic activities as being dubious, only to be deemed illegal by the administrator. Shelter can identify such false links by allowing the flow metadata data plan transfer, which collects the flow charts of valid network traffic along a path in the flow graph. Specifically, Shelter applies a custom algorithm to monitor and perceive the bytes of stream statistics by collecting STATS_REPLY messages at each switch in the flow path and determines whether the switches are diverging values of the transmitted byte account.

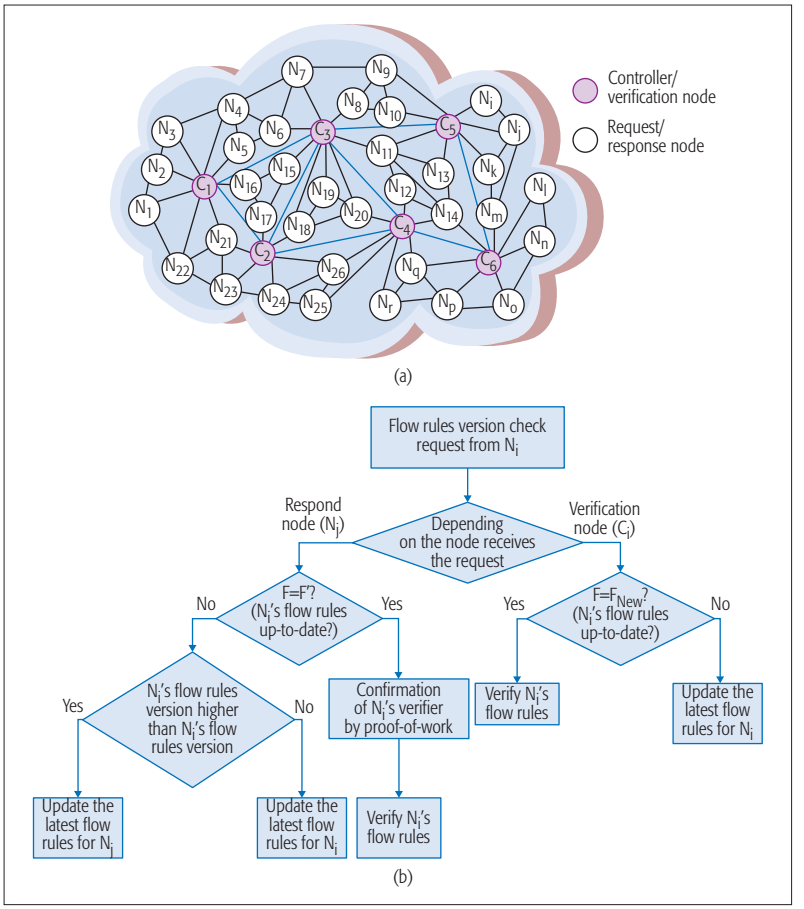

Figure 2. Updating scheme of flow rules table: a) distributed blockchain network; b) flowchart of flow rules table update.

The Updating of Flow Rules Table in the Distributed Blockchain Network

Figure 2a shows the overall DistBlockNet distributed blockchain network. The distributed blockchain network includes the controller/verification and request/response nodes. The verification node denotes the controller in the blockchain network, which maintains the updated flow rules table information in its own database. Request/response nodes are the IoT forwarding devices, which update its flow rules table in a blockchain network. IoT forwarding devices can be a request node or a response node. If a node requests its flow rules table, the node becomes a request node. At the point when a node sends a request message to update its flow rules table, the rest of the other normal nodes are considered to be response nodes from the viewpoint of the requesting node.

Figure 2b shows the DistBlockNet architecture model flow rules update in the distributed blockchain network. When an IoT forwarding device starts its flow rules table update by broadcasting a request packet with a version check, it views it as a request node. Once the version verification request packet is broadcasted in the

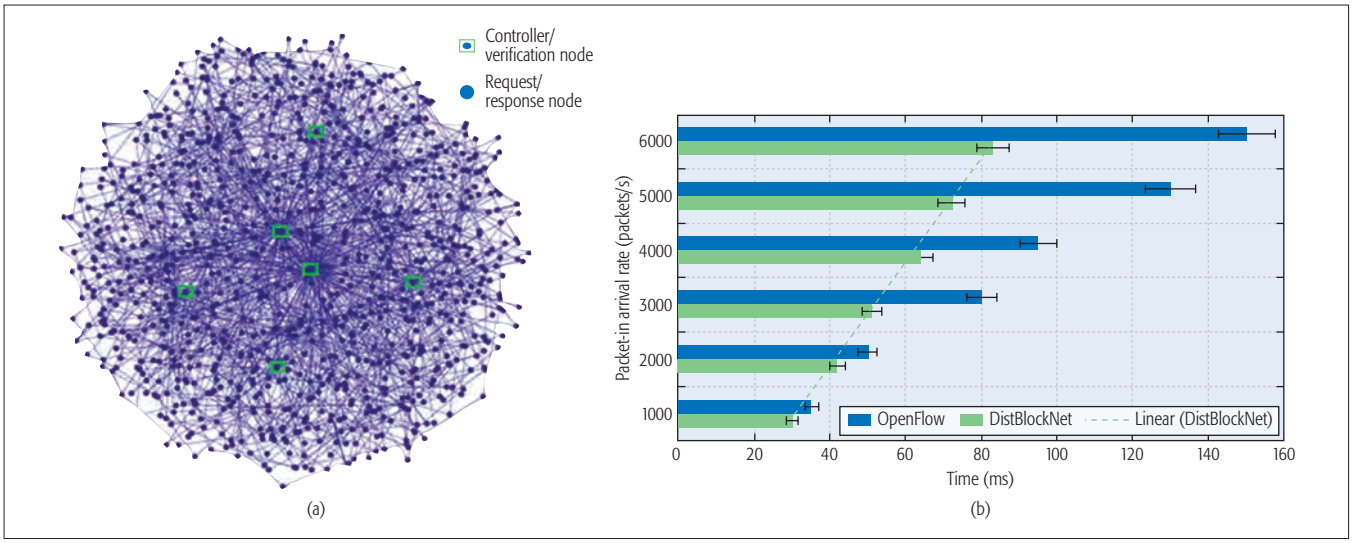

Figure 3. DistBlockNet performance on a large-scale network: a) distributed blockchain network with 6 controllers and 6000 nodes; b) flow table update time vs. packet-in arrival rate.

distributed blockchain IoT network, the rest of the IoT forwarding devices (i.e., response node and all controller/verification nodes) will respond to the request packet of the version verification. The response process varies depending on the node type. In the case of the controller/verification node, it checks whether the request packet node has the up-to-date flow rules table or not. The controller/verification node also checks the integrity of the flow rules table if the requesting node has the up-to-date flow rules table. Otherwise, the controller/verification node sends a response packet with the latest version of the flow rules table to the requested node.

In another case, when the response node receives the request, it checks the request's node version of the flow rules table with its own flow rules table version. If both the request and response nodes have the same version of the flow rules table, the response node requests the other nodes in the distributed blockchain network to verify the hash value of the flow rules table of the requested node. If the response node gets the confirmation of the hash value from the other nodes in the network (i.e., proof-of-work), the response node believes that the flow rules table is correct and sends the responding packet to the requested node. In another case, when the request and response nodes have a different version of flow rules tables, the response node checks whose flow rules table is the latest version. If the response node has the latest version, it will send the response packet to the requesting node with the latest version of the flow rules table. Otherwise, when the response node has a lower version of the flow rules table, it updates its own flow rules table from the request node packet.

Performance Evaluation

In this section, we present the details of the implementation, experimental environment, and evaluation of DistBlockNet. We carried out different experiments to evaluate the scalability, defense effects, accuracy, and efficiency of our proposed DistBlockNet architecture model.

Scalability

To assess the scalability of the DistBlockNet model, large-scale experiments are presented in this subsection with a cluster of 6 Intel i7 3.40 GHz with 16 GB RAM servers. We built a distributed blockchain network with 6 controllers/verifications and 6000 request/response nodes, as shown in Fig. 3a. We used the OpenFlow software switch instead of the OpenVSwitch because when a large number of switches are emulated, OpenVSwitch does not scale well. To compare the performance of the flow rules table update scheme of our proposed DistBlockNet model in a large-scale network, we also built a normally distributed SDN network. Figure 3b shows the result of the flow rules table update time with respect to the packet-in arrival rate in both the DistBlockNet model distributed blockchain network and distributed SDN network. In this experimental result, we observed that our proposed DistBlockNet model constantly performed superior to the distributed SDN network as the rate of the packet-in arrival increased.

Defense Effects

To assess the defense effects of our DistBlockNet model, we evaluated and compared it with an existing OpenFlow network by considering both software and hardware test environments [15]. We used the MININET SDN emulation tool for the software environment. We used the POX controller, OpenFLow switch, and server machines to implement clients and data plane caches in the hardware environment. We used some clients to dispatch a UDP floating attack to the switches. We measured the bandwidth of clients without and with flooding attacks generated by some clients at different speeds to the switch. We evaluated the impact on the bandwidth with and without the DistBlockNet model in both software and hardware environments separately because both environments have different capabilities.

In the software test environment, as shown in Fig.4a, we noticed that the bandwidth starts at 1.9 Gb/s without the presence of any attacks. When we started dispatching flooding attacks, the band-

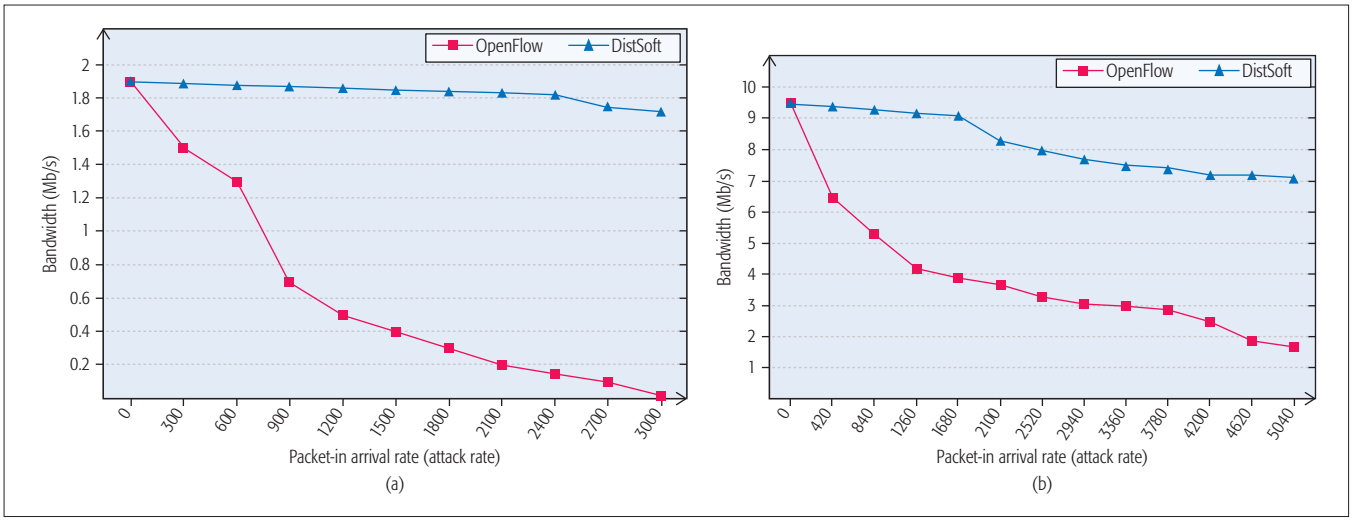

Figure 4. Effects on bandwidth during different attack rate in: a) software environment; b) hardware environment.

width decreased rapidly with an increase in the attack rate. The bandwidth went down to almost half when the packet-in arrival rate reached 800 packets/s. The entire network started malfunctioning when the packet-in arrival rate reached 3000 packets/s. On the other hand, using the DistBlockNet model, the bandwidth started at 1.9 Gb/s without the presence of an attack, and after the packet-in arrival rate reached 3000 packets per second, the bandwidth remained practically unchanged.

Figure 4b shows the results in the hardware test environment. In the hardware test environment, the bandwidth started at 9.5 Mb/s both with and without using the DistBlockNet model with any attack. In this experiment, we noticed that the bandwidth without using the DistBlockNet model went up to half when the attack rate of the packet-in arrival rate reached 1000 packets/s and started malfunctioning when the attack rate reached 5000 packets/s. While using the DistBlockNet model, the bandwidth was maintained above 9 Mb/s until the packet-in arrival reached 1600 packets/s. After that, the bandwidth started to go down because the ternary content addressable memory was not available in our switch. We used the OpenWRT software tool in place of the ternary content addressable memory to execute a flow rule table. Although a software-based flow rule table is not able to achieve a similar level of performance, we still noticed that DistBlockNet conserves resources and provides significant protection.

Accuracy

We evaluated the accuracy rate of the detection of DistBlockNet under two different parameters with one in the real-time identification of attacks and another in the presence various traffic and many distinctive defects in the system. The DistBlockNet model has the ability to identify every attack quickly. In the case of real-time identification, synthetic faults were used in parallel with the suitable traffic with 6K for the Mininet emulsified hosts on our physical testbed. Here we viewed the detection time as the time required for issuing of an alert from the moment when the DistBlockNet model received the offending packet. We used the custom traffic generator, which generates 1500 FLOW_MOD/sec. ARP attacks and fake topology are easily identified when the $PACKET_i N$ messages are processed. The detection times may fluctuate because in order to recognize DDoS/DoS attacks, DistBlockNet occasionally runs the flowchart validator and, as a result, the flow diagram size increases. In another case, we used Mininet to increase the number of hosts to 30K. Then we propelled DDoS, ARP poising, and fake topology attacks throughout the distributed blockchain network. We reiterated each examination more than 15 times and noticed that under the distinctive topologies, DistBlockNet effectively recognized each of the issues.

We ran a pessimistic scenario on the false alerts raised for a given δ using conflicting TCP iperf streams. The fair nature of the TCP will create fluctuations in flow to cause changes in the switches along the flow path, which would raise some precautions. As shown in Fig. 5a, we observed that with the increase of δ, the probability of false alarms occurring decreases.

The recall and precision are zero due to the absence of a true positive. In these experiment results, we observed that at the default value of δ = 1.06, there were 7 alarms out of 10 competing flows over 6 min. We also performed this experiment on our physical testbed and obtained comparable results.

To assess the absence of real alerts for a given δ, we defined the ratio between the number of checks that did not raise alerts to the total number of checks that raised alerts during verification. We evaluated the above metric among the Mininet hosts for controlled flows, which are eight hops apart. As shown in Fig. 5b, we observed that the absence of real alerts during verification increases as δ increases. For a given δ, the recall and precision are identical, which is equal to one minus the probability of the absence of real alerts at every data point.

Overhead Analysis

To evaluate the performance overhead of our DistBlockNet model, we used l2 learning and l3 learning applications and recorded CPU utilization during a flooding attack. We simultaneously

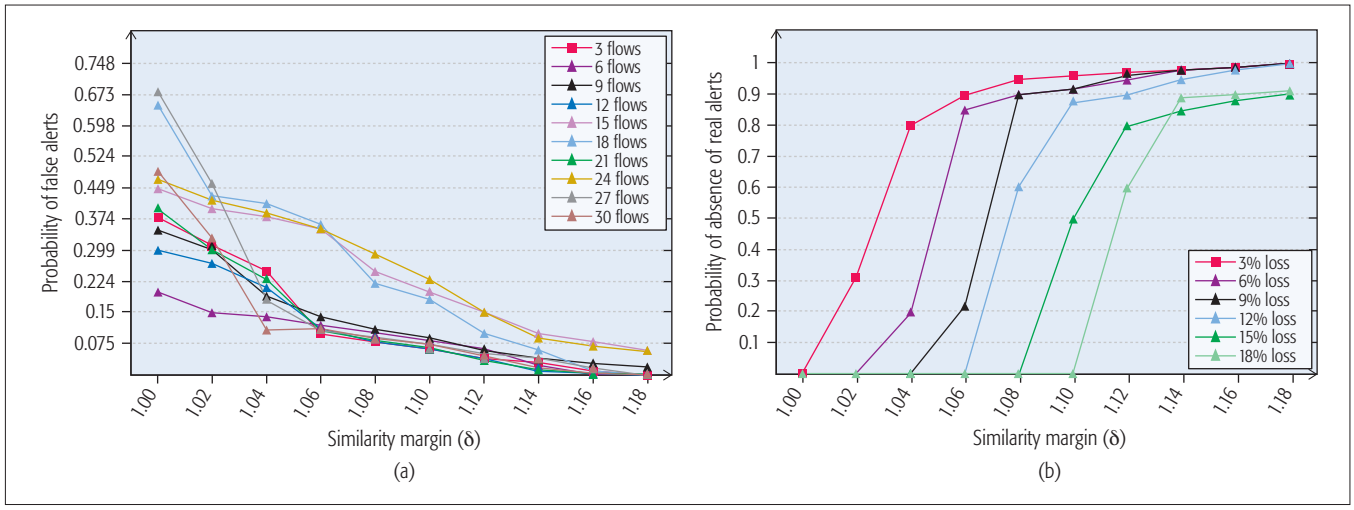

Figure 5. DistBlockNet accuracy rate: a) probability of false alerts with variation in flows and δ; b) probability of absence of real alerts vs. loss rate and δ.

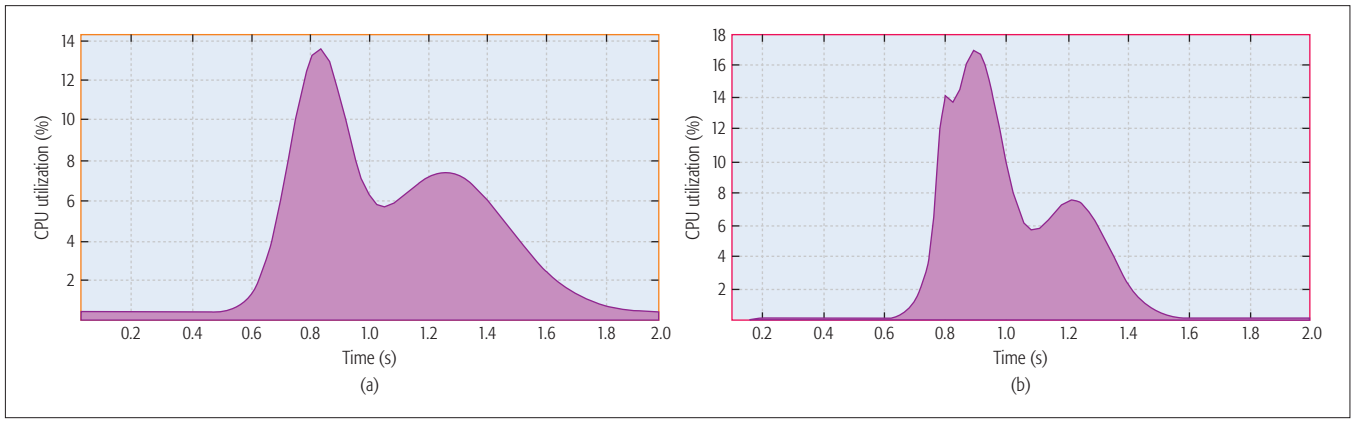

Figure 6. CPU utilization during flooding attack: a) running l2 learning application; b) running l3 learning application.

executed these two applications and used some clients to act as attackers and propelled the saturation attack with a rate of 500 packets/s with a DistBlockNet model in the hardware environment. For each application, we monitored the consumption of the resources. Figure 6 shows the average CPU utilization for all the controllers for different applications with the DistBlockNet model during flooding attacks. The flooding attacks began at about 0.5 s, and we noticed that CPU utilization quickly increased for each application. Then CPU usage started to slowly decrease after we installed the migration rules of flow. Based on the results, we observed that DistBlockNet provides effective protection and creates a more secure distributed network without consuming many resources during a saturation attack.

Conclusion

In this article, based on an analysis of the challenges that large-scale IoT networks face due to new communication paradigms, DistBlockNet, a new distributed secure IoT network architecture consisting of an SDN base network using the blockchains technique, has been proposed to address the current and future challenges and to satisfy new service requirements. DistBlockNet improves a system's performance and capacity. The core role of the DistBlockNet model is to generate and deploy protections, including threat prevention, data protection, and access control, and mitigate network attacks such as cache poising/ARP spoofing, DDoS/DoS attacks, and detect security threats. The DistBlockNet model also focuses on reducing the attack window time by allowing IoT forwarding devices to quickly check and download the latest table of flow rules if necessary. The performance evaluation is based on scalability, defense effects, accuracy rates, and the performance overheads of the proposed model. The evaluation results show the efficiency and effectiveness of the DistBlockNet model and have met the required design principles with minimal overhead.

In the future, we will extend our research work to build a distributed cloud computing architecture with secure fog nodes at the edge of the IoT network.

Acknowledgments

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIP) (No 2016R1A2B4011069)

References

[1] "Top Strategic Predictions for 2017 and Beyond: Surviving the Storm Winds of Digital Disruption," https://www.gartner.com/doc/3471568? ref=unauthreader, accessed May 5, 2017

[2] J. M. Batalla *et al.*, "On Cohabitating Networking Technologies with Common Wireless Access for Home Automation System Purposes," *IEEE Wireless Commun.*, vol. 23, no. 5, Oct. 2016, pp. 76–83.
[3] D. Levin *et al.*, "Logically Centralized?: State Distribution Trade-offs in Software Defined Networks," *Proc. 1st ACM SIGCOMM Wksp . Hot Topics in Software Defined Networks*, Aug. 2012, pp. 1–6.
[4] S. Stefan and J. Suomela, "Exploiting Locality in Distributed SDN Control," *Proc. 2nd ACM SIGCOMM Wksp. Hot Topics in Software Defined Networking*, Aug. 2013, pp. 121–26.
[5] X. Wu *et al.*, "A Multipath Resource Updating Approach for Distributed Controllers in the Software-Defined Network.," *Science China Info. Sciences*, vol. 59, no. 9, Sept. 2016, pp. 92,301–10.
[6] H. Lu *et al.*, "Hybnet: Network Manager for A Hybrid Network Infrastructure," *Proc. Industrial Track, 13th ACM/IFIP/USENIX Int'l. Middleware Conf.*, Dec. 2013, pp. 1–6
[7] D. Drutskoy, K. Eric, and J. Rexford, "Scalable Network Virtualization in Software-Defined Networks," *IEEE Internet Computing*, vol. 17, no. 2, Mar. 2013, pp. 20–27.
[8] Z. Qingyun *et al.*, "On Generally of the Data Plane and Scalability of the Control Plane in Software-Defined Networking," *China Commun.*, vol. 11, no. 2, Feb. 2014, pp. 55–64
[9] Y. Sung *et al.*, "FS-OpenSecurity: A Taxonomic Modeling of Security Threats in SDN for Future Sustainable Computing," *Sustainability*, vol. 8, no 9, Sept. 2016, pp. 919–44.
[10] Q. Vuong, H. M. Tran, and S. T. Le, "Distributed Event Monitoring for Software Defined Networks." *Proc. 2015 Int'l. Conf. Advanced Computing and Applications*, IEEE, Nov. 2015, pp. 90–97.
[11] K. Christidis and M. Devetsikiotis, "Blockchains and Smart Contracts for the Internet of Things," *IEEE Access*, vol. 4, May 2016, pp. 2292–303.
[12] F. Tschorsch and B. Scheuermann, "Bitcoin and Beyond: A Technical Survey on Decentralized Digital Currencies," *IEEE Commun. Surveys & Tutorials*, vol. 18, Mar. 2016, pp. 2084–2123.
[13] X. Xu *et al.*, "The Blockchain as a Software Connector," *Proc. 13th Working IEEE/IFIP Conf. Software Architecture*, Apr. 2016, pp. 1–10.
[14] K. Ahmed *et al.*, "Hawk: The Blockchain Model of Cryptography and Privacy-Preserving Smart Contracts." Univ. MD and Cornell Univ., May. 2015, pp. 1–32.
[15] J. M. Batalla *et al.*, "A Novel Methodology for Efficient Throughput Evaluation in Virtualized Routers," *2015 IEEE ICC*, June 2015, pp. 6899–905.

Biographies

Pradip Kumar Sharma (pradip@seoultech.ac.kr) is a Ph.D. scholar at Seoul National University of Science and Technology. He works in the Ubiquitous Computing & Security Research Group. Prior to beginning the Ph.D. program, he worked as a software engineer at MAQ Software, India. He received his dual Master's degree in computer science from Thapar University (2014) and Tezpur Univerity (2012), India. His current research interests are focused on security, SDN, SNS, and IoT.

Saurabh Singh (singh1989@seoultech.ac.kr) is a Ph.D. scholar at Seoul National University of Science and Technology carrying out his research in the field of ubiquitous security. He holds a strong academic record. He received his Bachelor's degree from Utter Pradesh Technical University and holds a Master's degree in Information Security from Thapar University. His research interests include cloud security, IoT, and cryptography. Finally, he has the experience of being a lab leader of the UCS lab, SeoulTech Korea.

Young-Sik Jeong (ysjeong@dongguk.edu) is a professor in the Department of Multimedia Engineering at Dongguk University, Korea. His research interests include cloud computing, mobile computing, IoT, and wireless sensor network applications. He received his B.S. degree in mathematics and his M.S. and Ph.D. degrees in computer science and engineering from Korea University, Seoul, in 1987, 1989, and 1993, respectively.

James J. (Jong Hyuk) Park (jhpark1@seoultech.ac.kr) received his Ph.D. degrees from the Graduate School of Information Security, Korea University, and the Graduate School of Human Sciences, Waseda University, Japan. He is now a professor at the Department of Computer Science and Engineering, Seoul National University of Science and Technology, Korea. He has published about 200 research papers in international journals/conferences. He has served as Chair and Program Committee member for many international conferences and workshops.

The performance evaluation is based on scalability, defense effects, accuracy rates and the performance overheads of proposed model. The evaluation results show the efficiency and effectiveness of the DistBlockNet model and have met the required design principles with minimal overhead.

MicroThings: A Generic IoT Architecture for Flexible Data Aggregation and Scalable Service Cooperation

Yulong Shen, Tao Zhang, Yongzhi Wang, Hua Wang, and Xiaohong Jiang

Inspired by the Internet architecture over which divergent devices can easily be accessed and a considerable number of applications can be run, the authors propose a generic architecture for IoT. This architecture supports two DIYs: network DIY for data aggregation and application DIY for service cooperation. To connect these two DIYs, a centralized controller has been designed to provide standardized interfaces for data acquisition, organization, and storage, and to support elastic and supportive computing.

ABSTRACT

The Internet of Things has been widely deployed in various areas of daily life through heterogeneous communications protocols. Each unstandardized protocol focuses on a specific IoT communication pattern. Inspired by the Internet architecture over which divergent devices can easily be accessed and a considerable number of applications can be run, we propose a generic architecture for IoT. This architecture supports two DIY areas: network DIY for data aggregation and application DIY for service cooperation. To connect these two DIYs, a centralized controller has been designed to provide standardized interfaces for data acquisition, organization, and storage, and to support elastic and supportive computing. With these properties, divergent devices can coexist in a uniform microworld, and rich services can be developed and provided on demand to interoperate with physical devices. This article discusses the background, design principles, and advantages of the proposed architecture, as well as open problems and our initial solution, which substantiates a novel IoT architecture and new research ground.

INTRODUCTION

Currently, the Internet of Things (IoT) has been widely adopted in various crucial systems such as city sensing, highway transportation, smart communities, and green farming. IoT aims to enable data exchange and smart communication among everyday objects, from watches, cookers, and bicycles, even to humans, plants, and animals [1]. Things can see, hear, and perceive the real-world environment to achieve more comfortable and safer living conditions. For instance, London has deployed all sorts of sensors to improve urban services: traffic prediction, weather forecasts, waste management, and water quality monitoring [2]. In addition to dedicated IoT platforms, mobile devices (e.g., smartphones, wearables, and vehicles) are also utilized as sensing resources.

Sensors in different IoT systems communicate through divergent protocols and standards. These protocols include Message Queuing Telemetry Transport (MQTT), Extensible Messaging and Presence Protocol (XMPP), Constrained Application Protocol (CoAP), Data Distribution Service (DDS), Advanced Message Queuing Protocol (AMQT), and others. The lack of ubiquitous sensor access has resulted in a fragmented IoT ecosystem, which brings about the following two problems. On one hand, IoT services are built on top of application-level message protocols. The design, deployment, and adoption of vendor-dependent IoT elements (sensors, acurators, gateways, etc.) is customized with regard to specific application requirements, making them difficult for an amateur to access or operate, let alone to organize them for customer-oriented application. On the other hand, as devices are controlled by dedicated platforms, the system logics tend to be ossified and hardly meet requirement changes. The differences from one platform to another lead to data disunity and redundancy, which also make it difficult to utilize data from different platforms.

The development of the IoT system is still in its infancy. However, the Internet has exerted significant influence on our lifestyles for a long time. From the hardware perspective, devices are connected to the Internet through unified interfaces, that is, we can communicate with cyberspace via an inexpensive home router after some configurations in wizard webpages. From the software perspective, rich and varied applications flourish on the Internet to satisfy consumers' specific tastes. Service-oriented architecture (SOA) and HTML5 further enrich the web-based application resources through a standardized interface. The design of future IoT should be flexible enough to satisfy the requirements of complex networks and various applications.

There is no doubt that it is necessary to build an adaptive and scalable architecture for sustainable IoT developments, just as the Internet provides today. Some platforms, such as the Eclipse IoT project, Intel IoT framework, and Baidu IoT framework, have been proposed to enable heterogeneous integration with different communication technologies and application protocols. These solutions mainly focus on the design of enhanced protocols or smart gateways for the conversion of multiple baseline protocols. However, as they are strongly coupled with the supported protocols, different platforms may not be compatible with each other. In addition, users have to be familiar with the specifications of each platform, leading to a steep learning curve for the non-expert.

Digital Object Identifier:
10.1109/MCOM.2017.1700104

Yulong Shen, Tao Zhang and Yongzhi Wang are with Xidian University; Hua Wang is with Victoria University; Xiaohong Jiang is with Future University Hakodate.

0163-6804/17/$25.00 © 2017 IEEE

Within the contextual background, this article proposes a generic IoT architecture, called *MicroThings*, to glue all the "Things" fragments into a uniform microworld. MicroThings integrates the application environment and the information aggregation environment with a logically centralized controller. Each of the two environments supports heterogeneous cooperation with a set of uniform interfaces. The controller connects the two environments and allows interoperation among applications and sensing devices.

Figure 1 illustrates the MicroThings architecture. It implements do-it-yourself abilities (DIYs) at two levels. One is the network DIY, which is at the bottom of MicroThings, where multiple sensing devices are accessed via compatible forwarding devices. Consumers are able to customize their IoT networks for data acquisition, transmission, and interaction. The other is the application DIY, which sits on top of the MicroThings, where standardized application programming interfaces (APIs) are provided for the development, deployment, and provision of new applications, as well as the reuse and composition of existing applications. These characteristics empower MicroThings to ubiquitously access various devices and greatly enrich service resources.

In the rest of this article, we first analyze the traditional three-layer IoT architecture, along with the design directions for existing architecture. We then introduce the proposed MicroThings architecture and potential solutions. Case studies are provided to show the benefits of MicroThings. Finally, we conclude the article.

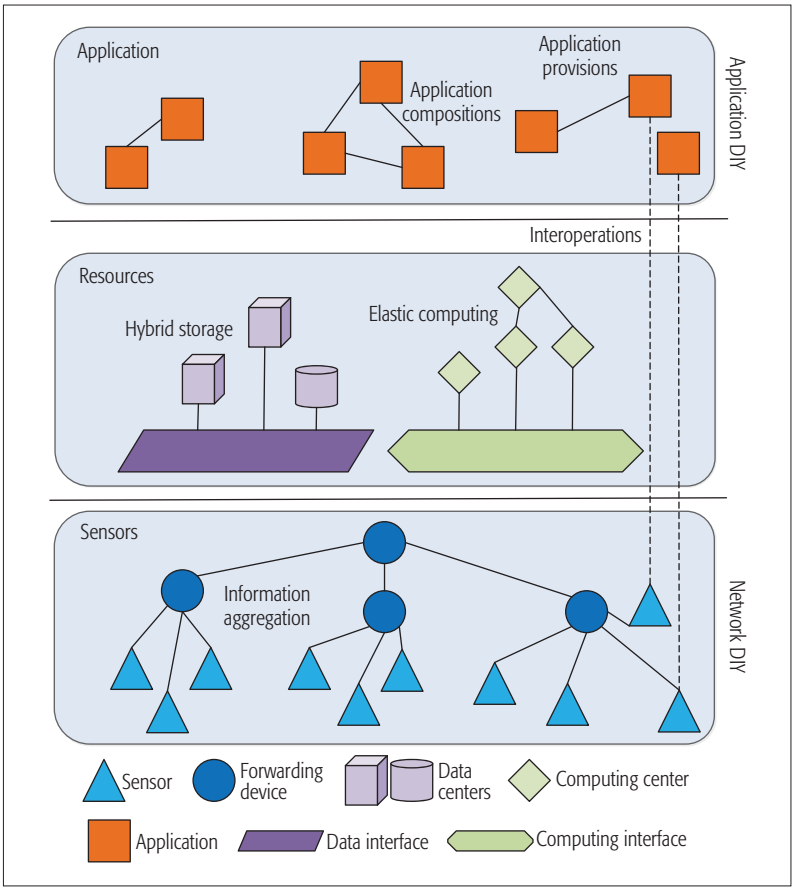

Figure 1. Illustration of the proposed architecture.

IoT: Challenges and Opportunities

State of the Art

Figure 2 illustrates the traditional IoT architecture. This architecture logically consists of three layers: the perception layer, the network layer, and the application layer. Various sensing devices are deployed in the perception layer to collect and aggregate data from the physical world (e.g., temperature and humidity). The network layer completes a wide range of information transmission and exchange on the converged network systems such as cellular networks — second/third/fourth generation (2G/3G/4G), narrowband IoT (NB-IoT), and so on — short-range communication networks (Bluetooth, Wifi, Zigbee, etc.), long-range wireless networks (LoRa, openRF, etc.), and remote communication networks (satellite). The collected data are transmitted to data centers through the network layer for further storage and processing. The application layer processes the collected data and interacts with people for further analysis and decision making. This three-layer architecture has been deployed in different areas such as intelligent homes, smart cities, public security, and green farming.

The general approach to developing IoT systems is to synthetically design these three layers with respect to specific application requirements. More specifically, engineers first select proper sensors to collect required data from the physical world, then design gateways and routers to transmit the data to local or remote data centers, and finally gather data to design various domain-specific applications.

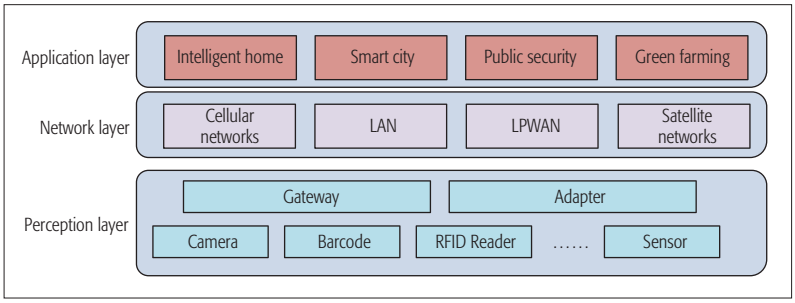

Figure 2. Three-layer Internet of Things architecture.

Problems and Directions

While the domain-specific development of the above three-layer IoT architecture seems quite straightforward, it has several problems that hinder sustainable development. We summarize these issues as follows.

Difficulty with IoT Device Networking: Through split-level developments, the perception layer manages the access and data transmission of sensing devices. However, as devices are closely relevant to the sensing data, frequent reconfigurations of the perception and network layers are needed when the collecting devices change. For example, the improvement in the accuracy of air quality monitoring needs redeployment and reconfiguration of more air sensors. In addition, with the coexistence of different physical interfaces and communication protocols, it is impractical for an ordinary consumer to organize the devices, or handle the heterogeneous security issues [3].

In addition to ubiquitous access, there is the trend toward network automation. It should provide a configurable and programmable interface with which people will need less expertise to be involved in modern intelligent IoT activities.

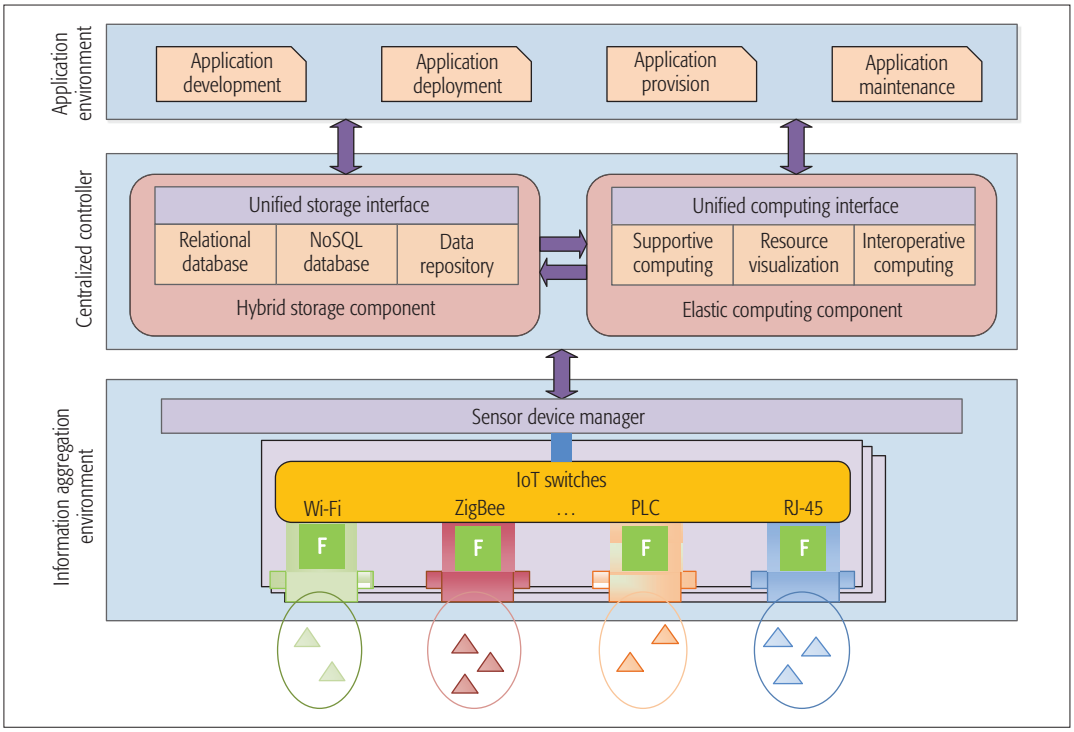

Figure 3. MicroThings architecture.

High Data Storage and Computing Cost: We should build different storage and computing frameworks for different types of data, which means high investment in the IoT application infrastructures. For instance, the system should establish file repositories and NoSQL databases to store unstructured data like real-time videos, and relational databases to store structured data. Even for applications of similar functionality from different stakeholders, the software, hardware, power, and control systems must be deployed independently, leading to a high cost for each IoT system.

Inflexible Application Provisions: The control logics of applications are closely coupled with the platform configurations such as the perceptive sensors, the network structures, and the storage/computing framework. In addition, the problem of broadly reusable applications remains unsolved. For example, existing applications cannot be directly migrated to the new environment without an application market. It requires application re-development or re-configuration for some, or even similar, IoT scenarios, which has led to overlapping investments.

With the development of network technology and the increasing number of applications, industry and academia need to establish a generic architecture that can address the main issues mentioned above. There are some trends apparent in the design of IoT architecture, which we summarize below.

Automation of the ubiquitous network: As the capillary ends, sensors are the data sources that supply blood to the whole IoT system. Since multiple sensing devices provided by different manufacturers have coexisted in the IoT system, the design of new architecture should support the ubiquitous access that is compatible with these different communication protocols and standards. For example, C. Hou *et al.* proposed profile-based access for device integration with IoT-Cloud, which groups the sensors by category and integrates the sensors by the category profile to the IoT-Cloud [4].

In addition to ubiquitous access, there is the trend toward network automation. It should provide a configurable and programmable interface with which people will need less expertise to be involved in modern intelligent IoT activities. For instance, A. Al-Fuquha *et al.* proposed a rule-based gateway for device organizations with divergent protocols which also guaranteed the quality of service (QoS) properties in the IoT systems [5].

The blend of flexible data storage and elastic computing: Flexible data storage provides the ability to properly manage the collected data from multiple data sources, and elastic computing further enhances the ability to process the massive amount data [6–8]. The blend of data storage and computing creates a uniform core that can control the IoT logics. As the linking element between the physical world and the cyber world, this core provides a resource pool with all kinds of operations that can help hide the complexity and the heterogeneity of underlying infrastructures. This trend suggests that IoT architecture should build a control core for interoperation between the southbound devices and the northbound applications.

The reuse of application programming and composition interface means that the northbound environment supports the programming, provision, and sharing of applications to different stakeholders in a unified market [9]. Also, the application environment should be a platform-independent model that can customize and consume existing services to construct composite ones that facilitate reuse, just as SOA) provides.

As the reuse of applications can significantly reduce the development and maintenance costs, it is believed that the scalable service resources should be one of the built-in characteristics of modern IoT architecture.

There are many standardization opportunities that could be implemented to make it easier for service providers and consumers to work with an IoT ecosystem. From the perspective of the IPSO Alliance, the IoT is a system-of-systems, and each system has its own standards. Therefore, supporting interoperability of different systems is still the major hurdle to achieving massive IoT adoption. Standardization organizations such as the International Telecommunication Union Telecommunication Standardization Sector (ITU-T), Third Generation Partnership Project (3GPP), and International Organization for Standardization/International Electrotechnial Commission (ISO/IEC) are leading the IoT standardization progress. As an ongoing standard specification, oneM2M (M2M: machine-to-machine) is devleloping a semantics enabler to bridge the gap between current IoT resources. As well, SmartM2M and Light-Weight M2M (LWM2M) have been developed to provide network-independent services. In addition to the aforementioned standards, some projects such as Eclipse IoT projects, the Intel IoT platform, the Baidu IoT platform, and the Tencent IoT platform are also contributing their efforts to provide uniform open APIs to simplify application development over IoT devices. The general direction is to design a cloud-based architecture to enable intelligent data acquisition and analysis through integrated protocols and standards, and bring uniform access while supporting different interactions between cloud and smart devices. However, these platforms are still excessively complicated for the non-expert, and it is necessary to build a generic architecture with lower entry barrier.

MicroThings: Architecture Overview

In this article, we propose a generic IoT architecture, referred to as MicroThings, which combines the fragmented devices, networks, and applications into a micro IoT world. As illustrated in Fig. 3, the MicroThings consists of two environments connected with a centralized controller.

Information Aggregation Environment: This environment supports the first DIY for operating physical devices. The environment aggregates data from the packages of devices, including the resource-constraint devices and the resource-rich equipment, and multiple multi-mode switches. Each switch integrates a variety of wired and wireless communication interfaces to control the transmission patterns among IoT devices. For prevalent protocols, like MQT and, AMQP, the architecture employs a conversion scheme to address the interoperation issues. In addition, to support more ubiquitous device access, this environment also establishes a JSON encoded HTTP parser to support the extensive accesses of the non-TCP/IP devices and scalable protocols. Thus, consumers can focus only on the connecting devices without worrying about the design specifications.

Moreover, the environment establishes a manager to provide a registration mechanism for the multi-domain connected devices. The manager plays the roles of both virtualized gateway and pattern recommender. The "things" can join or leave MicroThings for public or private use. Thus, MicroThings can recommend suitable applications to involved devices when they talk to each other, for example, recommending smart home applications for air quality devices. On the other hand, devices can be shared on MicroThings for public use, which greatly raises resource utilization ratio and reduces the deployment cost.

Centralized Controller Environment: This environment plays the intermediary role between the information aggregation environment and the application environment. The controller consists of two parts: the hybrid storage component and the elastic computing component. The storage component provides standardized APIs to gather the aggregated data from the southbound environment, and then stores them by category in the corresponding databases. The computing component extracts required information from the storage component, and then prepares the computing resources for further processing. The elasticity ensures that the computing resources can be dynamically allocated to coordinate with the dynamic changes across the entire IoT architecture.

Application Environment: This environment supports the second DIY for the IoT control logics, where applications can be developed with respect to the published requirements. As the underlying differences are hidden from the centralized controller, this environment creates a unified service ecosystem that supports application provision, development, deployment, and maintenance.

Moreover, this environment supports the reuse of applications, where different stakeholders can share the provided applications. Besides, for complex control logics, existing applications can be combined to rapidly create value-added functionalities.

The computing component extracts required information from the storage component, and then prepares the computing resources for further processing. The elasticity ensures that the computing resources can be dynamically allocated to coordinate with the dynamic changes across the entire IoT architecture.

Design Principles and Implementation Solutions

The previous section presented a generic architecture, MicroThings, which supports two kinds of DIYs coordinated with a centralized control. In this section, we illustrate how such architecture has been instantiated. For this purpose, a set of design principles has been proposed as the building blocks for the architecture, and the implementation solutions are given.

Information Aggregation Environment Design

Principle 1: Enabling ubiquitous and on-demand network access for IoT devices: The information aggregation environment lies in the southbound direction of the MicroThings architecture, which is designed to directly interact with the physical sensing devices. On the traditional Internet, communication devices are accessed via a set of uniform standards, such as Ethernet, to transmit information from one point to another. Compared to the traditional Internet, the design of an information aggregation environment for the IoT is much more difficult because of the heterogeneity of devices. Besides, the data aggregated in MicroThings should be identifiable, manageable, and controllable. As devices are abstracted in the

The addressable sensors can directly communicate to our MicroThings switches via the manufactured gateway. Other server-oriented communications, like satellite and wired communications, can be plugged into our architecture for data storage and process. Moreover, different IoT switches can be further placed at different locations for large-scale deployments.

data layer, end users can achieve on-demand network DIY in the MicroThings architecture.

As an initial solution to the above problems, we first designed IoT switches for data device access. An IoT switch is a multi-mode switch consisting of a set of front-end ports and a back-end port. The front-end port supports different communication protocols and standards, including Wi-Fi, ZigBee, PLC, RJ-45, and so on, while the back-end port connects to the sensor manager. To achieve the first objective, an IoT switch is deployed to access the local heterogeneous participant sensors. Therefore, the addressable sensors can directly communicate to our MicroThings switches via the manufactured gateway. Other server-oriented communications, like satellite and wired communications, can be plugged into our architecture for data storage and processing. Moreover, different IoT switches can be further placed at different locations for large-scale deployments.

For the second objective, a device manager is also used for the automatic management of IoT switches. The information that the manager receives contains the identifications and statuses of the sensors. Thus, sensors are registered automatically to provide abstractions when they are active. The centralized controller also uses this information for the abstraction of the data rather than the physical devices.

Centralized Controller Environment Design

The centralized controller is the connection between the southbound environment and the northbound environment. The design of the centralized controller should achieve two objectives: hybrid storage and elastic computing.

Principle 2: Provide a unified framework for heterogeneous data storage and management: The storage component is designed to interact with the southbound information aggregation environment. As the volume of the aggregated data increases rapidly, the data storage component should store the massive data with a high throughput capacity. Data are collected from multiple sources with different structures, and then the storage component categorizes and organizes them in a transparent way [10].

Jiang *et al.* have introduced a solution for data storage by combining multiple databases [11]. They have also built some mapping schemes for database operations. By providing standardized APIs, data can be stored by category automatically. In addition, since the information aggregation environment provides identifiable information for stored data, we propose a method to decide where the data will be stored and when the results will be computed.

Principle 3: Provide elasticity to enable supportive computing in the converged networks: The computing component is designed to interact with the northbound application environment. The design of components should support elasticity to enhance the ability for different heterogeneous data and applications [12, 13].

To achieve the objectives, MicroThings supports the following three features. First, the visualization of computing resources provides computing scalability, which in turn supports the interoperations between the northbound and southbound environments. Second, as IoT is a converged computing platform that supports the next generation communication technology, that is, the fifth generation (5G) mobile network, the edge devices with increasing computing resources can also contribute their capabilities for supportive computing. Finally, the visualization resources can be scheduled in a flexible and interoperative way, which further allows more fine-grained control between services and devices.

Application Environment Design

Principle 4: Provide a whole ecosystem for IoT application development, deployment, and provision: The application environment lies in the northbound of the MicroThings architecture, which is designed to provide applications for multiple stakeholders, such as device providers, software providers, and storage providers. As more and more devices are connected to MicroThings, the application environment should support the whole application production process, including development, deployment, provision, and maintenance [14].

SOA is the most commonly used for service provision via XML-based standards. Constructing an ecosystem for service-oriented application is a general approach to achieve these goals. On one hand, the ecosystem provides all applications with standard interfaces, with which the application can control device logics and interact with other applications. On the other hand, multiple stakeholders can share the developed applications. They can publish requirements to the ecosystem and obtain applications in a pay-for-use manner. Furthermore, the ecosystem supports value-added application provision with model-driven development (MDD) from the composition of existing applications, creating a resource pool of rich and varied applications.

Case Study and Validation

Selected Scenario

To further illustrate the advantages of our MicroThings architecture, we conduct a case study and its analysis in this section. Figure 4 illustrates the selected scenario in which four practical IoT systems are deployed in our MicroThings architecture: highway transportation, safe city, smart community, and green farming.

In the above scenario, these four systems execute the same processing flow. First, sensors are connected to the IoT switches for data collection. Then applications are deployed to process the control logics of the sensors. Finally, the central controller connects the above two parts to automatically control data collection and effectively provide computing resources.

Architecture Validation

Information Aggregation: We consider the deployment of these systems to illustrate the network DIY. The safe city system reflects the need for city sensing. It deploys sensors to observe street views, weather, noise, air condition, and so on. The aggregation data may vary with the sensors' interfaces and functionalities. Consider the following two cases for sensing data aggregation. One is the sensor deployment in a new area. For instance, a modernizing city usually deploys many

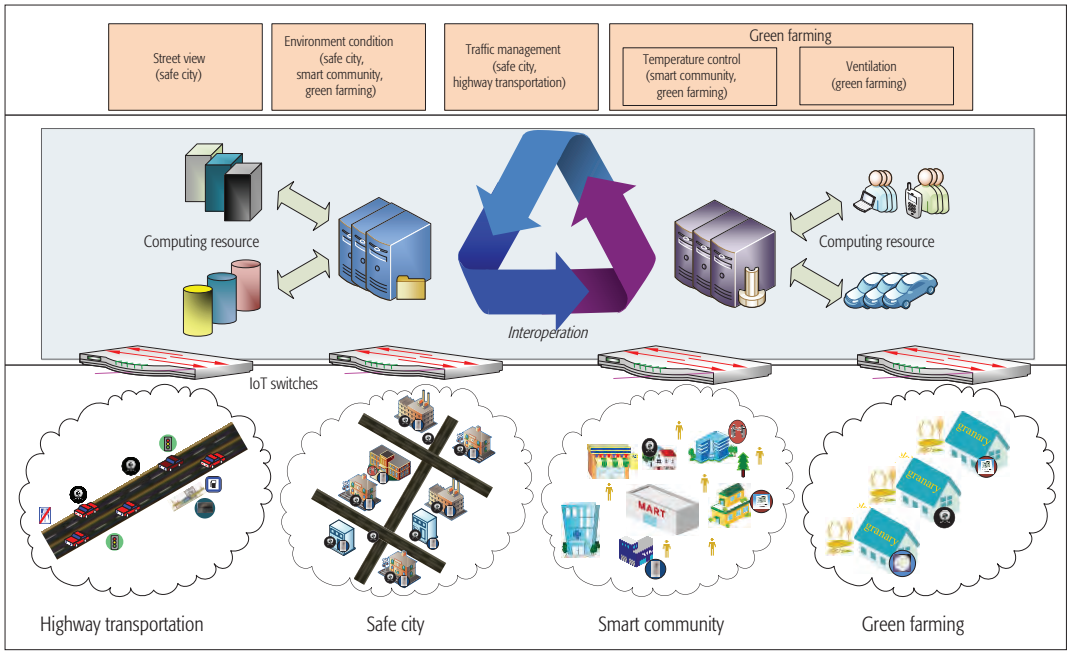

Figure 4. Implementation of multiple IoT systems with MicroThings.

The applications for environmental monitoring can be shared by multiple systems, such as safe city, smart community and green farming. In addition, the applications can be composed together using drag-and-drop way to provide value-added and various applications. MicroThings checks the input-output results of some typical patterns for each composition.

cameras for a street view. Cameras from different vendors can be connected to our IoT switches for data aggregation. Our environment provides a wizard to register the cameras, including the manufacturer, identification, sampling frequency, and so on. If the devices are supported, our MicroThings automatically translates the protocols of these devices and lists the possible communication patterns for the registered devices. The aggregated data will be further processed by the administrator for the city views. Otherwise, for unsupported devices, more configuration information to support JSON-formatted data transmission is required. The other case is sensor maintenance and replacement. For environmental condition monitoring applications, the sensors sometimes need to be replaced for accuracy improvement, and some of the sensors can easily be added or removed only if they are supported by the compatible protocols from the IoT switches.

We then consider another system for highway transportation, where some of the sensors are compatible with sensors in the safe city. The highway transportation shares the sensors to provide dynamic traffic information to enhance the safety of the neighboring city. Thus, the data can be aggregated together to share with other systems.

Hybrid Data Storage: We consider the green farming system. In this case, the environmental data (CO_2, NO_2, temperature) and video data should be stored. The environmental data are stored in structured databases as they are composed of some determined items such as time, value, and location. However, the video data are massive and unstructured. We store them in the NoSQL database. The storage provides standard interfaces to connect with the information aggregation environment and automatically categorizes the environmental and video data. It also supports locating data for further computing such as cooling the granary when the temperature is rising.

Elastic Computing: With the development of computation capability, CPU-controlled devices can contribute their power for computation. Through visualization, intelligent vehicles and handheld cell phones can support computing provision. For instance, sensors on cell phones can perceive temperature and noise data to cooperatively compute the local comfort degree. Similarly, vehicles can also compute the results of traffic flow for various IoT applications.

In addition, the computing component supports interactions between the physical and cyber worlds. Consider the smart community system, where deployed chillers can be automatically adjusted to a comfortable temperature, and the sensors send monitored data to the application for controlling the behavior of the chiller. The adjustment supports interoperation between chillers and control applications. For example, the application can activate more sensors when residents are quite sensitive to the change of the temperature, or the sensors can send control information to reduce the monitoring frequency if the temperature stays constant for a long time.

Application Life Cycle: The application environment supports the whole application life cycle including application development, deployment, and provision. It also supports some advanced features, such as application reuse and composition.

Since environmental monitoring is the most common IoT application, it can be shared by different platforms. In other words, the applications for environmental monitoring can be shared by multiple systems, such as safe city, smart community, and green farming. In addition, the applications can be composed together in a drag-and-drop way to provide value-added and various applications. MicroThings checks the input-output results of some typical patterns for each composition. For instance, a grain preservation application is provided in green farming by composing environment monitoring (for the level of O_2) service and ventilation service.

From these scenarios, we can conclude that

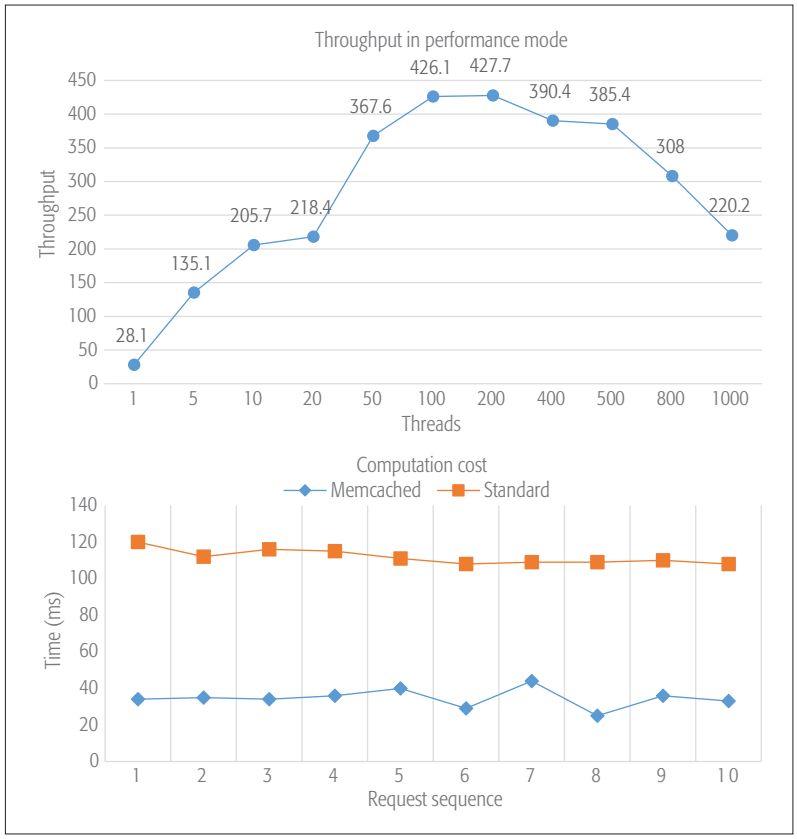

Figure 5. Performance evaluation of MicroThings: a) throughout in performance mode to support elastic computation; b) computation cost for hybrid storage.

the most intuitive advantages of our proposed architecture contain the following three aspects:

1. A unified MicroThings architecture can support multiple IoT systems.
2. Multiple novel systems can be deployed rapidly from existing systems to adapt to different environments.
3. Multiple stakeholders can share the applications run on MicroThings.

Performance Evaluation

We have performed simulation experiments to evaluate the performance of our proposed MicroThings architecture. Our MicroThings runs on a Xen-based cloud with 2-core processors and 4 GB memory. MicroThings supports two running modes, the performance mode and the balanced mode, to collect data from the physical world. The former exploits all the possible resources in the cloud-based centralized controller to handle the sensing requests, while the latter only uses the configured resources. As the latter mode can be considered as one partition of the former one, we only conducted experiments to simulate the requests from our physical devices in the performance mode.

Figure 5 illustrates the results of our evaluation on computation costs and throughput capacity. In MicroThings, our centralized core supports elastic computing using adaptive threads. Figure 5a reveals that our MicroThings adaptively creates close to 200 threads to handle the requests from physical devices under the experimental configuration. As our MicroThings supports a Memcached-enabled hybrid storage framework for ubiquitous data storage, the hot data cached in the Memcached system can be used for massive IoT requests. Compared to the non-cached hybrid framework, our MicroThings can offload 60 percent of requests to the storage system (Fig. 5b). Therefore, the core of MicroThings is an efficient architecture for interoperation between two DIY environments, even for large-scale IoT deployments.

Conclusion

This article introduces a generic Internet of Things architecture, called MicroThings, which lowers the hurdle of IoT development. Specifically, the architecture supports network and application DIYs with the help of the centralized controller. Within this architecture, crowds rather than professional designers can customize their network and application developments. MicroThings consists of four parts: data aggregation, storage, computing, and processing, and provides standardized interfaces. As a result, this architecture unifies fragmented IoT elements into a whole ecosystem, facilitates the control and management of physical devices in the physical world, enriches the application resources in the cyber world, and provides elastic computing and hybrid storage for flexible interoperations between these two worlds.

Acknowledgment

This research was supported in part by China NSFC Grants U1536202, 61373173, 61602365, and 61602364, Shaanxi Science & Technology Coordination & Innovation Project 2016KTZDGY05-07-01, and Fundamental Research Funds for the Central Universities BDY131419.

References

[1] Y. Qin *et al.*, "When Things Matter: A Survey on Data-Centric Internet of Things," *J. Net. Comp. Appl.*, vol. 64, 2016, pp. 137–53.
[2] D. Boyle *et al.*, "Urban Sensor Data Streams: London 2013," *IEEE Internet Comp.*, vol. 17, no. 6, 2013, pp. 12–20.
[3] D Abebe *et al.*, "Lightweight Cybersecurity Schemes Using Elliptic Curve Cryptography in Publish-Subscribe Fog Computing," *Mobile Net. Appl.*, vol. 22, no. 112, 2017, pp. 1–11.
[4] C. Hou *et al.*, "Middleware for IoT-Cloud Integration Across Application Domains," *IEEE Design & Test*, vol. 31, no. 3, 2014, pp. 21–31.
[5] A. Al-Fuqaha *et al.*, "Toward Better Horizontal Integration among IoT Services," *IEEE Commun. Mag.*, vol. 53, no. 9, Sept. 2015, pp. 72–79.
[6] X. Sun *et al.*, "EdgeIoT: Mobile Edge Computing for the Internet of Things," *IEEE Commun. Mag.*, vol. 54, no. 12, Dec. 2016, pp. 22–29.
[7] P. C. Brebner, "Is Your Cloud Elastic Enough? Performance Modelling the Elasticity of Infrastructure as a Service (IaaS) Cloud Applications," *Proc. ACM/SPEC Int'l. Conf. Performance Engineering*, 2012, pp. 263–66.
[8] F. Paraiso *et al.*, "Managing Elasticity across Multiple Cloud Providers," *Proc. 2013 Int'l. Wksp. Multi-Cloud Applications and Federated Clouds*, 2013, pp. 53–60.
[9] O. Krieger *et al.*, "Enabling a Marketplace of Clouds: VMware's vCloud Director," *ACM SIGOPS Operating Systems Review*, vol. 44, no. 4, 2010, pp. 103–14.
[10] A. Botta *et al.*, "On the Integration of Cloud Computing and Internet of Things," *Proc. Int'.l Conf. Future Internet of Things and Cloud*, 2014, pp. 23–30.
[11] L. Jiang *et al.*, "An IoT-Oriented Data Storage Framework in Cloud Computing Platform," *IEEE Trans. Ind. Info.*, vol. 10, no. 2, 2014, pp. 1443–51.
[12] H. L. Truong *et al.*, "Principles for Engineering IoT Cloud Systems," *IEEE Cloud Comp.*, vol. 2, no. 2, 2015, pp. 68–76.
[13] A. Gumaste *et al.*, "Network Hardware Virtualization for Application Provisioning in Core Networks," *IEEE Commun. Mag.*, vol. 55, no. 2, Feb. 2017, pp. 152–59.

[14] A. Bucchiarone *et al.*, "Incremental Composition for Adaptive By-Design Service Based Systems," *Proc. 23rd IEEE Int'l. Conf. Web Services*, 2016, pp. 236–43.

Biographies

YULONG SHEN (ylshen@mail.xidian.edu.cn) is a professor at the School of Computer Science and Technology, Xidian University, China. He is an associate director of the Shaanxi Key Laboratory of Network and System Security and a member of the State Key Laboratory of Integrated Services Networks. He has served on the Technical Program Committees of several international conferences, including ICEBE and INCoS. His research interests include Internet of Things, wireless network, and cloud computing security.

TAO ZHANG (taozhang@xidian.edu.cn) received his B.S degree in computer science from Xi'an University of Post and Telecommunications, China, in 2008. He received his M.S. and Ph.D. degrees in computer science from Xidian University in 2011 and 2015, respectively. Since August 2015, he has been an assistant professor at the School of Computer Science and Technology, Xidian University. His research interests include service-oriented computing, the Internet of Things, and cloud computing security.

YONGZHI WANG (yzwang@xidian.edu.cn) received his B.S and M.S degrees in computer science from Xidian University in 2004 and 2007, respectively. He received his Ph.D. in computer science from Florida International University in 2015. Since August 2015, he has been an assistant professor at Xidian University. His research interests include cloud computing security, network security, big data, and the Internet of Things.

HUA WANG (hua.wang@vu.edu.au) is a full professor at Victoria University. He has more than 10 years of teaching and working experience in applied informatics in both private enterprise and academia. He has expertise in electronic commerce, business process modeling, and enterprise architecture. As Chief Investigator, three Australian Research Council Discovery grants have been awarded since 2006, and 155 peer reviewed scholar papers have been published.

XIAOHONG JIANG [SM] (jiang@fun.ac.jp) is currently a full professor at Future University Hakodate, Japan. His research interests include wireless networks, optical networks, network security, and more. He has published over 280 technical papers at premium international journals and conferences. He was the winner of the Best Paper Award of IEEE HPCC 2014, IEEE WCNC 2012, and the IEEE ICC 2005 Optical Networking Symposium. He is a member of ACM and IEICE.

Big Data Orchestration as a Service Network

Xiao Liu, Yuxin Liu, Houbing Song, and Anfeng Liu

The authors argue that a big data network joint SDN, together with cloud and fog computing platforms, can build a service chain network. In SDN, the purpose is to reduce a large amount of redundant data and response time. They propose a novel Big Data Orchestration as a Service Networking framework, which can dynamically orchestrate big data into services in SDN.

ABSTRACT

This article argues that a big data network joint SDN, together with cloud and fog computing platforms, can build a service chain network. In SDN, the purpose is to reduce a large amount of redundant data and response time. We propose a novel Big Data Orchestration as a Service (BDOaaS) as the networking framework, which can dynamically orchestrate big data into services in SDN. In BDOaaS networking, the data center distributes software to all devices in the distributed network, which can orchestrate big data into services in the distributed network; the services-oriented network model is formed. Thus, the network load and response time is reduced. The BDOaaS framework and various components of BDOaaS as well as operation mechanisms are discussed in detail. Simulation results are presented to show the effectiveness of the proposed BDOaaS framework. In addition, we discuss a number of challenges in implementing the proposed framework in next generation networks.

INTRODUCTION

Nowadays, big data networks and software defined networks (SDNs) are booming [1–3]. The development of big data networks is based on the enhanced ability to collect and process data, which can make human decision making more comprehensive and wise [1, 3]. Big data was developed due to the development of new data acquisition equipment:

1. Pervasive mobile devices such as smartphones are expected to play a significant role, for example, as human-machine interfaces, data sources for environmental monitoring (i.e., sensors for context detection) and user profiling, ubiquitous social networking, and anywhere-anytime-to-anything connectivity [2–5].
2. Sensing devices will be improved. For example, in a manufacturing plant, ubiquitous sensors gather monitor data to support the work [3, 6, 7].

This unprecedented increase in data traffic faces unprecedented challenges for current levels of data transmission:

1. The amount of data cause larger network loads. In big cities, many forms of data are relevant; thus, it is important to propose a method to reduce the amount of data.
2. Due to the longer distance between network edges and core networks, when users request services, the request information is transmitted to the core network, and the service is satisfied.

This requires more time and large delays. Thus, the required time for each service must be reduced. This can cause network congestion and rapid increases in delay, and quality of service (QoS) deteriorates rapidly; thus, the quality of experience (QoE) is worse [8].

In big data networks, there are three different roles. (a) service provider (SP), (b) services customer (SC) or simple user, and (c) big data collector (BDC), which refers to devices or people that can collect data. The aim of the SP is to combine the collected data into advanced services and send services to users after cleaning and refining the data. Usually, the SP is the owner of the data center and distributes the data collection task to the BDC. The BDC collects data through sensing devices and then transmits those data packets to the data center. The SP provides services to users after processing the data. For example, VTrack, which provides omnipresent traffic information, and NoiseTube, which makes noise maps, can be regarded as SPs [7]. Crowd sensing networks (CSNs), which leverage the ubiquity of sensor-equipped mobile devices, are a kind of BDC that can collect information and provide a new paradigm for solving complex sensing applications [2, 9]. For example, users can obtain traffic information or noise map services from VTrack or NoiseTube, respectively. In this application, the SP needs to collect a large amount of data and then provide users with detailed traffic information or noise distribution. In order to improve QoS, it is better to collect a large amount of data. However, a large amount of data can increase the network load.

The SP usually adopts a data center or an information-centric (IC) system, which is located in the cloud, to store data [2, 10]. IC systems and data centers have been extensively studied in recent years [10]. All data need to be sent to the data center, and the BDC for collecting data is at the edge of the network. Because the distance between the cloud and the edge device is usually long, transmitting a large amount of data from the device to the SP may not be feasible or economical. If users are at the edge of the network, the SP may not provide guaranteed low-latency service to users. Thus, the establishment of an effective big data network faces huge challenges.

To address these issues, content-based distribution technology is proposed. In-network caching is used in an information-centric network (ICN) to speed up content distribution and improve net-

Digital Object Identifier:
10.1109/MCOM.2017.1700090

Xiao Liu,Yuxin Liu, and Anfeng Liu are with Central South University; Houbing Song is with Embry-Riddle Aeronautical University.

work resource utilization [2, 10]. In ICNs, requests no longer need to travel to the content source but are served by a closer ICN "content node" along the path, which can speed up the distribution of content [10]. The shortages of this scheme are that it is a kind of improvement strategy and thus will not fundamentally solve the problem because the proportion of caching content on the network is not high. In addition, this method is only suitable for data requests and cannot effectively alleviate the generated network load that the primary data flows upload to the data center.

The main reason for low network performance is the long distance route between BDCs at the network edge and the data center. Fog computing and mobile edge computing (MEC) [5] have been proposed to deploy local data centers closer to end users. Fog (from core to edge) computing, a term coined by Cisco in 2012, is a distributed computing paradigm that empowers network devices at different hierarchical levels with various degrees of computational and storage capability [5]. Compared to cloud computing architecture, which is centralized in nature [5], the scheme provides real-time and low-latency services to billions of IoT devices at the edge of the network [5]. Due to the increase in the fog layer in the system model, a large amount of data and service requests can be processed in the local network but do not need be sent to the data center in the cloud. This layered structure can improve network performance to a certain extent, while this improved performance is achieved by hardware, thus increasing the cost of the SP and decreasing flexibility. Moreover, the local center in the network edge needs to interact with other data centers, which also creates a network load. Therefore, an advanced network is needed that can reduce the network load, making the network more extensible and providing better QoE and QoS.

SDNs have attracted great interest in both academia and industry [2, 11]. SDNs introduce the ability to program the network via a logically software-defined controller. SDNs can be extended to wireless networks, which benefit from the SDN [2, 11].

The rise of SDNs has brought opportunities to solve these problems. This article argues that big data networking jointly with SDN, together with cloud and fog computing platforms, can build a service chain network. Through networks such as ICNs, service requests are processed in less time; caching the same packets at the appropriate time and path is not an easy thing. We propose a novel BDOaaS networking framework that can enable dynamic orchestration of big data into services in SDN to meet the requirements of next generation networks. In BDOaaS, the data center distributes software to all levels of devices in the distributed network, which can orchestrate big data into services in all devices of a distributed network, not only in the data center, thus achieving a services-based network computing model. Orchestrating data as a service is achieved by software stored at all levels of devices in the network. This kind of software is the software for data collection tasks published by the SP, which plays a role in two areas.

Data aggregation. When the BDC submits the data flow to the data center, all devices that pass by call the software for operation data aggregation [3, 7], which can reduce the amount of uplink data flow greatly, thus reducing the network load.

Services build. The data are not only aggregated when the data flow is transmitted to devices but also computed to build the service. Thus, when users request a service, if the services structured by current devices can meet the requests, the results will be immediately returned to the users. Thus, service request flow and services return flow are reduced. Because different devices are in the different network layers, the structured services are different. Devices on the network edge only construct services in the scope of a local small area, and devices near the core network construct services in the greater range of a region. Thus, devices at a lower level can satisfy most service requests by users. Therefore, the system has strong adaptability and expansibility.

Devices on the network edge only construct services in the scope of a local small area, and devices near the core network construct services in the greater range of a region. Thus, devices at a lower level can satisfy most service requests by users. Therefore, the system has strong adaptability and expansibility.

In this article, the proposed BDOaaS has the following characteristics:

- The main target is to orchestrate service; it is different from current networks whose target is data transmission. Moreover, distributed orchestration is operated in all levels of network equipment, and its target is to build a kind of service-based distributed network.
- Data can be transmitted after data has been programmed; thus, the amount of data can be reduced greatly.
- An SP releases software to network equipment, and data services are transformed in the distributed network, thereby greatly reducing the amount of data.

Compared to traditional networks, the main innovations of BDOaaS are as follows.

•A new services-based computing network model is introduced that can run in the cloud and fog computing platform, extending the functionality of big data networks and SDN. In a services-based computing model, the raw data are no longer routed in the network. The two most important orchestration functions are as follows:

–Aggregating data reduce data flow.

–The data are programmed as services.

BDOaaS combines data and software to orchestrate services.

•A new way to reduce network load is conceived. The core idea is that it routes services rather than data. In the previous network, the main function of the network is to receive and forward data, but not take into consideration data content. On the other hand, the BDOaaS can be regarded as a content-based (services-based) network, in which the SP distributes the software to all levels of devices, and then data are orchestrated into services to reduce the network load.

•The BDOaaS networking framework is proposed. The components of the BDOaaS framework are discussed in detail, which has been proven effective by experiments.

However, there are many challenges for providing a better method. For example:

- Methods that are compatible with the current network
- The design of orchestrating software
- Releasing orchestrating software

A detailed discussion is presented in the section on open research challenges.

The BDOaas network framework is useful for networks of large size. For example, for big data networks, the amount of raw data is large. For large cities, the data center can collect thousands of raw data packets each minute.

The rest of this article is organized as follows. We describe an overview of the BDOaaS networking framework. We present the operating mechanism in BDOaaS. Simulation results are presented. Some open research issues are discussed. Finally, we conclude this work.

Overview of Orchestrating Data as Services Networking

Motivations

In the last decade, there has been a rise in a variety of ubiquitous computing and handheld devices [2, 4, 6, 7]. Large amounts of data greatly enhance the human ability to observe the world and make wise decisions; a network with a large amount of data transmission is called a big data network. In the big data network, a large amount of data needs to be transmitted; there are more than 9 billion devices used to generate data according to statistics, with an annual growth rate of 50 percent [5]; the rate of increase of the trunk network is <10 percent [5]. Thus, it becomes an urgent task to establish a new network computing model to cater to the rapid growth in the big data network.

BDOaaS is proposed based on two aspects.

1. A more advanced services-based network framework is needed for the development of big-data-based applications. We note that a large amount of forwarding data in the big data network are relevant; the relevant data always contain the same information. Thus, only some of that information is transmitted to the next devices, and devices can know all the information. Thus, some redundant data can be refined through the appropriate calculation; this operation is called aggregating data (or orchestrating data). For example, a meteorological department (as an SP) releases tasks for gathering haze information in the specified area, and a huge population with smart devices or devices collects the haze information and then submits the haze information to the SP (the SP's data center). In such applications, there are a large number of data packets, which can cause greater network load if all packets are routed to the data center. We note that the content of raw packet is {time, location, content}, and the sensed haze information in small areas is the same at the same time; thus, the data packet can be orchestrated as a new data packet by network devices through the orchestrating software provided by SP. The most simple form is {time, location1, location2,... location*n*, content}, and the amount of data can be reduced comparatively. If the haze information in all locations is the same, the haze information can be expressed as a data packet, such as {time, all-location, content}. Thus, the amount of data can be reduced by an order of magnitude. However, orchestrating software is distributed by specific software released by the SP. The difference with the previous network is that orchestrating software not only resides in network devices but also directly resides in BDC devices; thus, orchestrating data is more effective. For example, if weather conditions or industrial monitoring situations are stable, and the sensed object (e.g., haze information) has not changed for a long time, its report interval can be increased, and it will not have any impact on the application. On the contrary, if the sensed object changes dynamically, the ample data frequency can be increased adaptively, which can be effective to ensure the quality of data acquisition and greatly reduce the network load.

From the viewpoint of building and obtaining services, information based on big data, such as VTrack, NoisTub, and Haze information, due to orchestrating software, resides in the devices of the uplink flow of raw data. The query information of users has already been built; thus, users can send a services request query; the format of Query haze information services, for example, is {time, location, request-content}. If the information requested by users can be satisfied by devices in the routing path, the devices can directly return the results, without routing to the data center. Even if devices cannot satisfy the request, the request information continues routing to a higher level of devices. Because most users' service requests are local, most requests could be returned in the network edge. As a result, the network load and the service request time are reduced, and the user QoE is improved.

2. From the point of view of hardware facilities in the current network, in big data networks, fog computing combined with SDN has the corresponding conditions to implement BDOaaS. This is the case in particular for the fog computing platform, the purpose of which is to move network computing to the network edge; thus, the method is similar to this scheme, where BDOaaS moves the network to the network edge. Thus, the establishment of BDOaaS in the fog computing platform is a natural thing. However, the fog computing platform provides a physical platform, and the implementation aspect is solved by BDOaaS. Therefore, BDOaaS has good significance.

The BDOaas network framework is useful for networks of large size. For example, for big data networks, the amount of raw data is large. For large cities, the data center can collect thousands of raw data packets each minute. If every raw data packet is transmitted to the data center, the network load is 100 MB in a round data collection; the generated traffic flow within one hour is up to 6 GB; however, the capacity of final distribution of the haze map is only a few kilobytes for collecting data in a 10,000-point location in one hour. If the BDOaaS is used, data packets will be merged into one packet in one grid, and the data can be merged among different grids again; the amount of orchestrating data transferred is less than 60 MB in 1 hour, and thus the data transmission can be reduced by a factor of 100.

The Framework of BDOaaS

The framework of BDOaaS is shown as Fig. 1. BDOaaS can be divided into five layers: the data collection layer, fog or edge network layer, core network layer, data center layer, and application layer.

Data collection layer: This is located at the edge of the network and consists of pervasive sensing devices, such as various types of mobile devices, smartphones, and industrial, civil, and public area sensing devices [12, 13]. All devices for collecting data are called BDCs. These devices are the data source of BDOaaS. According to [5], the number of connected devices has exceed-

ed the number of people on Earth since 2011. Connected devices now number 9 billion and are expected to grow more rapidly, reaching 24 billion devices by 2020 [5]. With the increasing number of heterogeneous devices connected to IoT generating large amounts of data [5], using a data-transmission-based network cannot meet the current rapid growth of data. In BDOaaS, the raw data are orchestrated and then form services; thus, the data transmission flow is reduced greatly to meet the development of the network.

Fog or edge network layer: However, in the BDOaaS, those network devices in the fog layer are equipped with the software released by SP and orchestrate data as services. Thus, the devices in the network layer have different functions compared with the previous devices.

•If those devices receive the data that need to be uplink forwarded, those data can be orchestrated by the interface provided by the SP.

•If they receive the service request, they call the interface provided by the SP and then orchestrate those data. If the orchestrated data can meet user requests, the results can be returned to the users; otherwise, the request continues forwarding to the uplink.

•If those devices receive a services return, they not only forward services but also integrate the services and the owner services; thus, the service requests of downlink users can be satisfied locally. Network devices orchestrate each uplink data packet; thus, the local services can be updated, and the downlink services can be orchestrated. Thus, the data that those services include come from larger areas. For example, in NoiseTube applications, the local updated noise distribution may be obtained after network devices orchestrate the received uplink data. Meanwhile, the noise distribution in a wider range can be obtained after the downlink services are orchestrated, which can provide a noise query service in a short path style to help reduce the service request delay and improve user QoE.

Core network: The core network refers to the current backbone network. Its function is similar to the devices in the fog layer. However, these devices are located in the upper layer, which has a larger view and can provide more functionality.

Data center layer: This is large-scale equipment with huge memory capacity and computing ability disposed by the SP; it can analyze and process data in depth. However, the difference from the previous network is that the data center in the previous network is mainly used to store and process data. In the BDOaaS, the SP not only has large-scale software, which processes data locally, but also provides software for orchestrating data to network devices. This software is programmed for specific applications. When the SP publishes the data collection task, it releases specific software to the network devices. Thus, in BDOaaS, the data center not only has a function for storing and processing data but also for strutting and releasing software, so the BDOaaS forms a service-oriented network that combines data and SDN.

Application layer: This refers to users for the application of services.

These five components in different situations have multiple roles. In the data collection layer, a smartphone can be used as a data collection device (i.e., a BDC). If devices can request services, they can be called users. In addition, there is no clear boundary between the fog layer and the core network layer. In addition, some big data centers in the network edge can be regarded as data centers. Figure 1 shows a typical network structure diagram. Compared to the core network, the fog layer mainly relies on distributed computer resources near local devices. Operations regarding data, data processing, and applications in the fog layer are more dependent on local devices rather than servers. Unlike the core network, almost all devices are stored in the core network.

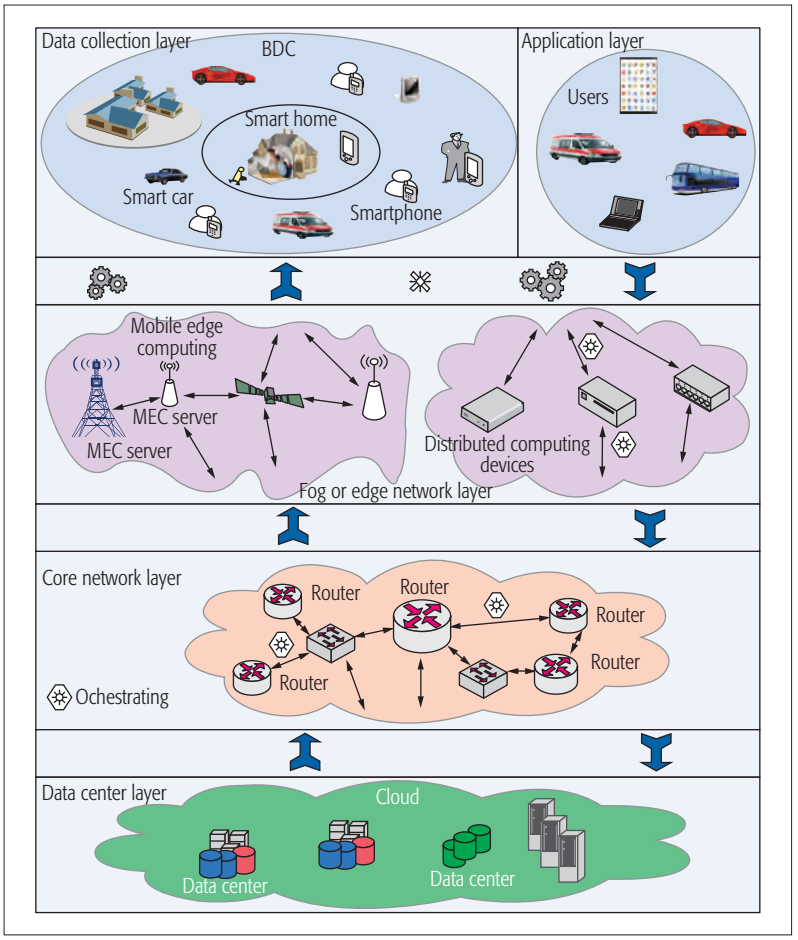

Figure 1. The framework of BDOaaS.

The Running Mechanism in BDOaaS

There are four kinds of flows in BDOaaS. These four types of flow are as follows.

Data or services flow: The content of the route is the data collected by the BDC, and the flow direction is from devices on the network edge to the data center. In BDOaaS, because network devices in the routing path orchestrate data, data are transformed into services. The transformed services will be transmitted to the data center; this is called the data or services uplink route.

Services request flow: The user sends a request to the SP to ask for a service. The flow direction is from the user to the SP (i.e., the data center); this is called the service uplink route.

Service return flow: After the SP or network

The main idea of BDOaaS is that the SP releases program codes for orchestrating data to all network devices, and then refines the forwarding data to reduce the amount of data. Its goals are to reduce the network load and enhance user QoS and QoE.

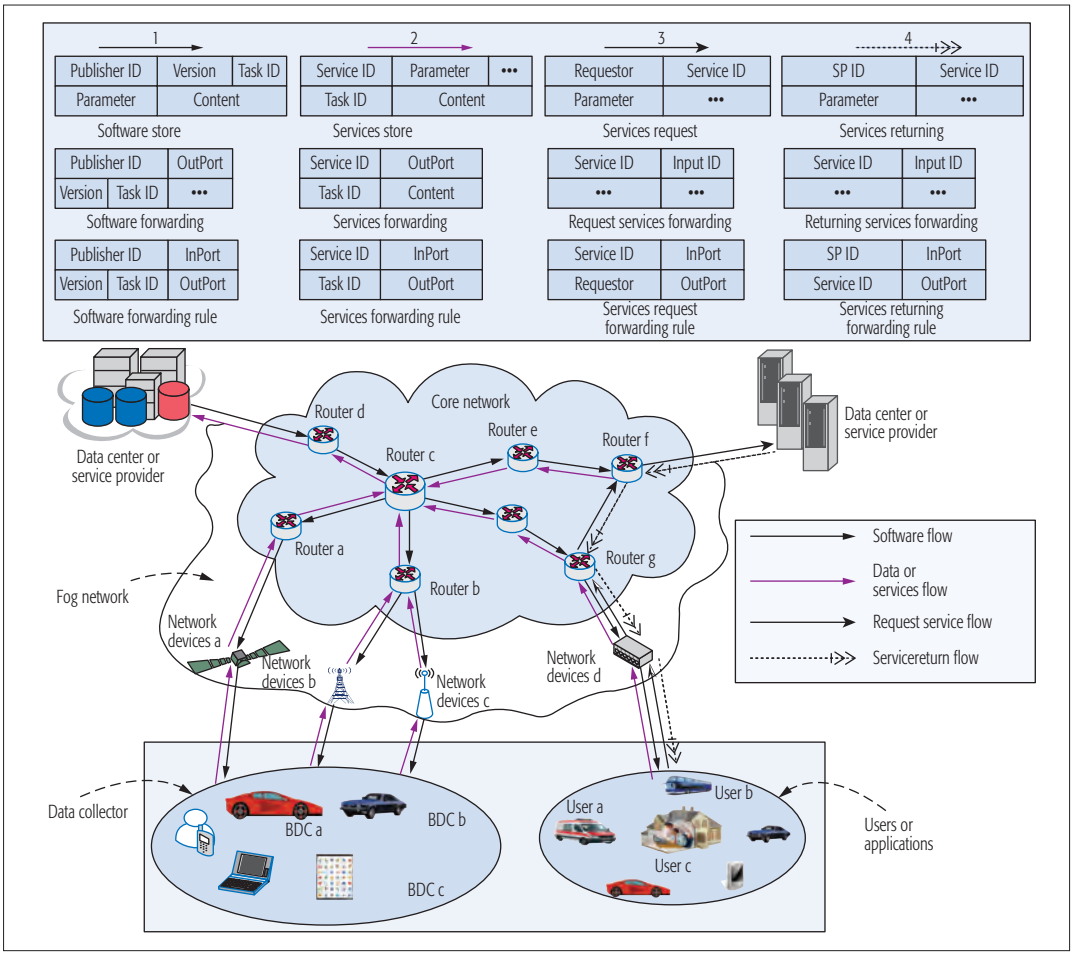

Figure 2. The running mechanism in BDOaaS.

devices receive a service request from users, if the services orchestrated by devices can meet the request of a user (the SP must be able to meet the user request), the device will send the service flow to the user, which is called the service downlink route.

Software release flow. The network devices can work effectively only when the devices receive software orchestrated by the SP; thus, the SP needs to release the software to the network devices, referred to as software flow.

The running mechanism in BDOaaS can be illustrated in Fig. 2. These four flows can show the operating mechanism.

1. First, a BDC obtains data by sensing and measuring the surrounding environment and then submits those data to the network; the operation is shown in flow 2 of Fig. 2 (data or services flow). BDC a submits its sensing data to network device b, and network device b receives the collected data from other BDCs at the same time. Thus, network device b orchestrates data through software provided by the SP. The orchestrated data may be services or aggregated data. Then the data or service flow continues flowing to the data center; this route is routed through router b, router c, router e, and router f. Each network device can orchestrate the data or services in the routing path; the route data are refined after devices orchestrate the data once. Finally, the data center obtains all information.

2. Services request flow. This is shown as Fig. 2; user a requests a service from the SP. After network device d receives the request, it checks whether the service request can be satisfied; if it can be satisfied, device b returns a service result or continues forwarding the service request to the SP. All devices (e.g., router g, router f) in the routing path check whether the service meets the request; if the request service can be satisfied, the service result can be returned, or the service request continues to be forwarded. If no devices in the routing path can satisfy the service request, the SP must be able to meet the service request of the user.

3. Services return flow. After network devices satisfy the service request of the user, the service results can be returned along the original routing path.

4. Software flow. The SP (residing in the data center) customizes software for orchestrating data to services. The customized software is spread to devices in the designated area (including the wireless network and the BDC) to complement BDOaaS.

The main idea of BDOaaS is that the SP releases program codes for orchestrating data to all network devices, and then refines the forwarding data to reduce the amount of data. Its goals are to reduce the network load and enhance user QoS and QoE. First, the SP releases the program code for orchestrating data to all levels of network devices (Fig. 2) for special functions, and all devices receive software. In the process of uploading

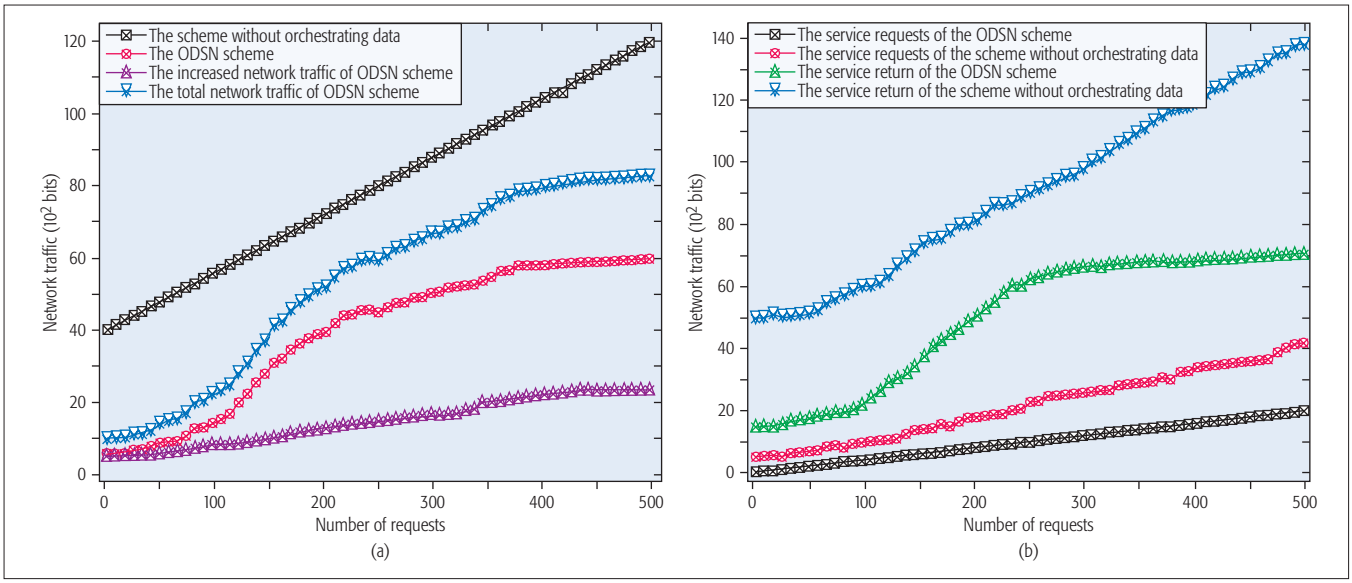

Figure 3. The data traffic of BDOaaS.

data from the BDC to the SP, when all devices receive data from downstream devices, the received data can be fused (or the services can be re-orchestrated) according to received software; the amount of data can be reduced greatly while ensuring complete information. Then the fused data continue to be transmitted to the data center. When users request services, if devices near users can orchestrate this service according to the received software, it returns the service to the users. Thus, the delay can be reduced.

Simulation Results

The main goals of this article are:

- Reduce the amount of redundant data.
- Speed up the response of the service request.

The model for orchestrating sensed data is as follows. If current device A receives the raw data p_a and receives the raw data p_b from sensing devices, the results for aggregating data are $\xi(A, B) = \max(p_a, p_b) + (1 - \beta_{A,B})\min(p_a, p_b)$, where

$$\beta_{A,B} = \frac{1}{e^{(d_{A,B}^2/\alpha)}}, d_{A,B}^2$$

is the distance between devices A and B, and α is a constant coefficient. If device A receives the aggregated data δ_b from device B, as long as one data packet in the received packet is not raw data, the result for aggregating data is $\xi(A, B) = \max(p_a, p_b) + (1 - \beta_{A,B})\min(p_a, p_b)$, where η is the forgetting factor, a decimal number less than 1. The parameters are p_a = 30bits, η = 0.8, α =2.

In the simulations, we consider one eNodeB, one WiFi router, and one WiMAX router, each of which has 10 users in the network. The kind of requirement service of users is 100. Assume that there are 100 popular content requests, and there are a total of 1000 requests for popular content, caching, and computing from all 30 users randomly. Assume the transmission delay is a value in the range of [10, 20] ms per hop. The storage space of devices is 500 bits. We consider the average response delay of users' requests and the average network traffic as our performance metrics.

The traffic generated with the number of requests is shown in Fig. 3. If the method for orchestrating data is not used, the generated traffic is given in the black line of Fig. 3. It can be seen that the generated traffic is basically linear growth with the number of requests. However, if the method for orchestrating data is adopted, the data may be uploaded after a large amount of data is orchestrated, so the growth rate is moderate. However, in BDOaaS, due to the release software, the increased flow for releasing software is given in Fig. 3a. In BDOaaS, the total traffic is shown in the blue line of Fig. 3a. A comparison of generated traffic for service requests is given in Fig. 3b. The generated traffic for service requests is small, and for return services, it is big. This shows the effectiveness of the BDOaaS.

Figure 4 shows the average number of hops to complete a full service and delay.

- Comparison of services request hops. Obviously, in BDOaaS, most services requests will be satisfied at the edge of the network; thus, the hop number in BDOaaS is far lower than that of other networks.
- Comparison of delay. Delay is related to network hop count and network congestion. If the hop count is large, the delay is higher. Obviously, the hop count in BDOaaS is smaller than the hop counts in other networks. Thus, the delay in BDOaaS is small.

Open Research Challenges

Despite the potential of BDOaaS networking, there are some research challenges that must be addressed. In this section, we address some of these challenges, which help to promote the study of BDOaaS.

Compatibility

The most important issue is compatibility for BDOaaS. In the proposed BDOaaS, we assume that each network device can receive software released by the SP, and can orchestrate data and services. However, designing current network devices should be as simple as possible. In particular, for the backbone of the high-speed network, complex data processing and comput-

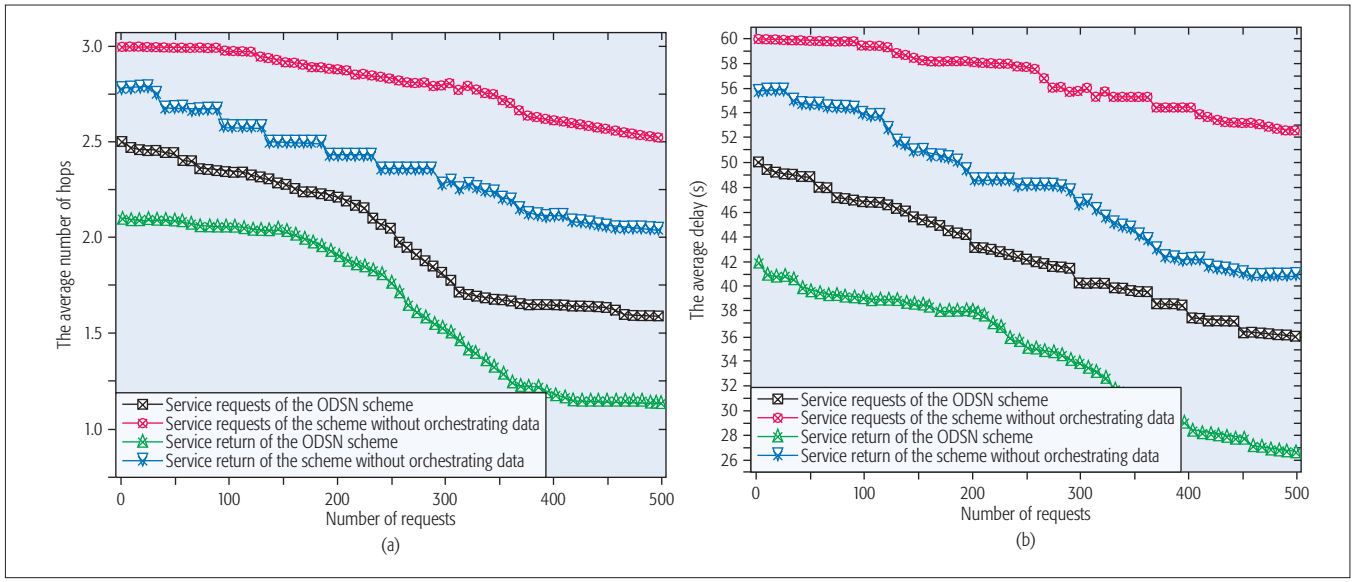

Figure 4. The quality of service of BDOaaS.

ing will greatly reduce the speed of forwarding data. Thus, the compatibility of BDOaaS with the current network is an important issue. In the current network, first, some specialized devices for orchestrating services can be deployed in the network edge (also in the core network), so BDOaaS can be achieved to some extent. Second, the network uses a unified format; the devices can process data packets quickly. Third, the data have been pre-processed by source collector devices. Due to the large number of devices of the BDC, the pre-processed function on the BDC devices will not reduce network performance but can effectively reduce the required computing for orchestrating data in network devices. With the development of hardware for network devices, the enhanced ability for computing and processing data will help the BDOaaS become more achievable.

The Issue of Orchestrating Software

The design of the orchestrating software is another important challenge. The orchestrating software in BDOaaS requires the following characteristics:

1. Small size. Software can quickly spread to all levels of devices in the network and lead to small network loads. Small-sized software can be stored in network devices with limited storage space.

2. The calculation is simple, and the speed of calculation is fast, to make it possible to adapt to the network devices with high-speed forwarding data. Meanwhile, upgrading old software is important in the BDOaaS.

The Routing Mechanism for Services

The current network is based on IP routing. Although there is a routing strategy based on content [14], there are still many issues worth studying. In the BDOaaS, the best routing method is based on services routing; this method is regarded as special routing based on content. Because each service has a service ID, it is more complex to design routing based on services than on the current IP. In the BDOaaS, releasing the orchestrating software is another important issue. Such software usually needs to spread to devices in a specific geographic area. The previous broadcast and multicast mode are not suitable for a route whose sensed target area is a geographic region. For routing that releases software to mobile devices, the design faces challenges. How to detect and update software is another important issue. Due to the large number of devices, the version number of software received by multiple devices is different; thus, how to ensure consistency among different software versions is a new issue.

In fact, the BDOaaS, as a new network, faces many challenges. Security is a permanent theme in the network and is even more important in the BDOaaS. Most services data based on big data are from the data collected by smartphones, and there are therefore greater threats to people's privacy, security, and so on [13]. In the network, designing incentive mechanisms for collecting data in addition to the allocation mechanisms of the network source face challenges in the BDOaaS.

Conclusion

In the future, the ability to obtain data will grow rapidly, and the growth of network transmission capacity will be slow; the contradiction between the two will become increasingly prominent. We believe that a fundamental solution is to construct a network that transmits information after refining data in the future network [15]. In fact, networks that transmit information instead of raw data have a variety of implementations and face many challenges. In this article, we propose a novel BDOaaS networking framework that can dynamically orchestrate big data to services in SDN, reducing the amount of data and network delay and thus improving user QoE. In BDOaaS, orchestrating data as services or orchestrating services as services is a novel idea for constructing future networks.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China (61379110, 61073104, 61572526, 61272494, 61572528) and the National Basic Research Program of China (973 Program) (2014CB046305)

References

[1] E. Zeydan *et al.*, "Big Data Caching for Networking: Moving from Cloud to Edge," *IEEE Commun. Mag.*, vol. 54, no. 9, Sept. 2016, pp. 36–42.

[2] R. Huo *et al.*, "Software Defined Networking, Caching, and Computing for Green Wireless Networks," *IEEE Commun. Mag.*, vol. 54 no. 11, Nov. 2016, pp. 185–93.

[3] L. Kong *et al.*, "Embracing Big Data With Compressive Sensing: A Green Approach in Industrial Wireless Networks," *IEEE Commun. Mag.*, vol. 54, no. 10, Oct. 2016, pp. 53–59.

[4] Y. Liu *et al.*, "APMD: A Fast Data Transmission Protocol with Reliability Guarantee for Pervasive Sensing Data Communication," *Pervasive and Mobile Computing*, DOI:10.1016/j.pmcj.2017.03.012, 2017.

[5] S. Sarkar and S. Misra, "Theoretical Modelling of Fog Computing: A Green Computing Paradigm to Support IoT Applications," *IET Networks*, vol. 5, no. 2, 2016, pp. 23–29.

[5] A. Liu *et al.*, "Distributed Multi-representative Re-Fusion Approach for Heterogeneous Sensing Data Collection," *ACM Trans. Embedded Computing Systems*, 2017, vol. 16, no. 3, 73, DOI: http://dx.doi.org/10.1145/2974021.

[6] X. Liu *et al.*, "Large-Scale Programing Code Dissemination for Software Defined Wireless Networks," *Computer Journal*, DOI: 10.1093/comjnl/bxx014, 2017.

[8] H. Li *et al.*, "Engineering Searchable Encryption of Mobile Cloud Networks: When QoE Meets QoP," *IEEE Wireless Commun.*, vol. 22, no. 4, Aug. 2015, pp. 74–80.

[9] J. Wang, C. Hu, and A. Liu, "Comprehensive Optimization of Energy Consumption and Delay Performance for Green Communication in Internet of Things," *Mobile Info. Systems*, vol. 2017, article ID 3206160, DOI: 10.1155/2017/3206160.

[10] C. Fang *et al.*, "A Survey of Green Information-Centric Networking: Research Issues and Challenges," *IEEE Commun. Surveys & Tutorials*, vol. 17, no. 3, 2015, pp. 1455–72.

[11] K. Xu *et al.*, "Toward Software Defined Smart Home," *IEEE Commun. Mag.*, vol. 54, no. 5, May 2016, pp. 116–22.

[12] X. Liu *et al.*, "Trace Malicious Source to Guarantee Cyber Security for Mass Monitor Critical Infrastructure," *J. Computer and System Sciences*, 2016, DOI: 10.1016/j.jcss.2016.09.008.

[13] I. M. Butun *et al.*, "Cloud-Centric Multi-Level Authentication as a Service for Secure Public Safety Device Networks," *IEEE Commun. Mag.*, vol. 54, no. 4, Apr. 2016, pp. 47–53.

[14] J. M. Batalla *et al.*, "ID-based Service-Oriented Communications for Unified Access to IoT," *Computers & Electrical Engineering*, vol. 52, 2016, pp. 98–113.

[15] Y. Kryftis *et al.*, "Efficient Entertainment Services Provision over a Novel Network Architecture," *IEEE Wireless Commun.*, vol. 23, no. 1, Feb. 2016, pp. 14–21.

Biographies

Xiao Liu received M.Sc. degrees from Central South University, China, in 2017. Currently she is a Ph.D candidate with the School of Information Science and Engineering of Central South University, China. Her research interests are crowd sensing networks, wireless sensor networks, and wireless security.

Yuxin Liu is currently a student in the School of Information Science and Engineering of Central South University. Her research interests are services-based networks, crowd sensing networks, and wireless sensor networks

Houbing Song [M'12, SM'14] (Houbing.Song@erau.edu) received his Ph.D. degree in electrical engineering from the University of Virginia, Charlottesville, in 2012. In 2017, he joined the Department of Electrical, Computer, Software, and Systems Engineering, Embry-Riddle Aeronautical University, Daytona Beach, Florida, where he is currently an assistant professor and the founding director of the Security and Optimization for Networked Globe Laboratory (SONG Lab, www.SONGLab.us). He is an Associate Technical Editor of *IEEE Communications Magazine*. He is a Senior Member of ACM. His research interests lie in the areas of communications and networking, cyber-physical systems, cybersecurity, and big data analytics.

Anfeng Liu received his M.Sc. and Ph.D degrees from Central South University in 2002 and 2005, both in computer science. He is currently a professor with the School of Information Science and Engineering of Central South University. He is also a member (E200012141M) of the China Computer Federation (CCF). His major research interests are service computing and wireless sensor networks.

In the future, the ability to obtain data will grow rapidly, and the growth of network transmission capacity will be slow; the contradiction between the two will become increasingly prominent. We believe that a fundamental solution is to construct a network that transmits the information after refining data in the future network.

Security and Privacy for Cloud-Based Data Management in the Health Network Service Chain: A Microservice Approach

Christian Esposito, Aniello Castiglione, Constantin-Alexandru Tudorica, and Florin Pop

The authors deal with healthcare-related data management and exchange, and they propose security and privacy requirements together with a novel microservices approach. They investigate how cloud computing can be adopted within healthcare systems.

ABSTRACT

New and useful tools that facilitate the harmonization and interconnection of health data have become a requirement for monitoring and preventing illness and also for sharing medical knowledge. Nowadays, cloud-based solutions can support collaborative data science platforms and deliver all types of processing operations through the network service chain. In this article we deal with healthcare-related data management and exchange, and we propose security and privacy requirements together with a novel microservices approach. We investigate how cloud computing can be adopted within healthcare systems. The interoperability of existing technologies will improve the quality of life and the efficiency of healthcare systems by making them more personalized and centered on patients, together with reducing operational costs and medical errors. To have a socially acceptable health network service chain, the security and privacy issues need to be analyzed and addressed. We explore the the security and privacy requirements and implications, and also discuss existing methods, and in the end propose an architecture of a secure manager for cloud-based healthcare-related data management and exchange.

INTRODUCTION

Healthcare is a data-intensive activity, where medical personnel require a huge amount of up-to-date healthcare-related data of their patients and their medical history, in order to make the correct medical decisions and to offer the best care. The lack of such knowledge on patients to be treated forces healthcare providers to conduct a series of tests on the patients that can be unnecessary, costly, and bothersome for the patients. In the past, all healthcare-rated data were archived internally by each healthcare provider based on paper documents. Therefore, the medical personnel could only access the patient data held by their membership provider. Moreover, retrieving valuable information from these traditional paper-based archives is not a simple task, and paper documents are easily vulnerable to manipulation, destruction, and loss. With the proliferation of information and communications technology (ICT) in many aspects of our modern society, the advent of computers and electronic documents within healthcare was unavoidable due to its benefits [1], and we are witnessing the increasing proliferation of ICT-based solutions for healthcare-related data management and exchange (HDME). Unfortunately, while this has allowed several of the problems (characterizing the traditional paper-based management of healthcare-related data) to be resolved, it has also paved the way for novel kinds of issues closely related to the security of these solutions from external and internal attacks as well as the protection of patient privacy from data breaches. Therefore, the effective use of ICT in approaching HDME requires novel solutions for offering security and privacy by considering the peculiarities of the healthcare domain and its legal framework of reference.

MOTIVATION

Within the healthcare domain, we have multiple actors that have to exchange data among each other so that patients are able to perceive a suitable degree of quality in the received healthcare. Each of these actors generate data as a side consequence of their work and must receive data in order to perform their duties; Fig. 1 shows the data flows among these actors (indicated by red lines in the figure). Clearly, such a figure does not have the ambition to deal exhaustively with the matter specifying all the possible data usages within a healthcare provider and the aspects of data management in the current healthcare practice, but rather aims to highlight the complexity in the current healthcare scenario, and present all the different enabling technologies for the vision of pervasive and efficient flow of data in medical practice. The figure distinguishes among two different kinds of data flows. The first one is for primary uses (with continuous lines), and is represented by the exchange of medical, clinical, and health information from the patient to the medical personnel (e.g., general practitioners, pathologists, laboratory staff), and vice versa, in order to support the medical personnel in their work. The second one is for secondary uses (with dashed lines), and is represented by data gathered by the provider administration in order to collect the costs and performance, and make the best administrative decisions, or by academic researchers to

Digital Object Identifier:
10.1109/MCOM.2017.1700089

Christian Esposito and Aniello Castiglione are with the University of Napoli "Federico II" and the University of Salerno.
Constantin-Alexandru Tudorica and Florin Pop are with the University Politehnica of Bucharest.

0163-6804/17/$25.00 © 2017 IEEE

investigate disease spreading, statistics related to cure effectiveness over the ill population, or by policy makers in order to tailor investments and regulations to the providers' needs and evaluation. Currently, there is a huge debate on several aspects presented in the figure, for example, the most effective way of dealing with exchanging healthcare information across multiple different infrastructures and devices.

Nowadays, hosting, running, and maintaining a data center to support HDME for a given healthcare provider imply very high costs, spanning from those to buy all the needed commodities and hardware to those for energy to run the machines and the cooling systems, to the salary of the technical staff in charge of administrating the data centers, and to the maintenance to substitute failed components and/or upgrade obsolete hardware. Such costs are starting to be important expenditure terms for the overall budget of healthcare providers, since the amount of data to be stored is progressively increasing according to an exponential trend as computer-based communication and document dematerialization are increasingly being applied and enforced within the medical practice, taking the place of paper-based communication and documents [2]. Cloud computing has the evident benefit of improving the flexibility, maintainability, and scalability of the underlying IT infrastructures, paving the way for novel ways of healthcare-related data fruition by means of mobile computing and easy sharing of medical data within a country or across borders. In addition, alongside the introduction of cloud computing within healthcare, we are witnessing the progressive application of microservices, which is one of the latest architectural trends in software engineering.

Microservices are starting to meet considerable application within the context of cloud computing, thanks to the increasing use of containers in such a computing environment for software development and deployment [3]. Specifically, the service-oriented software hosted within the cloud and in charge of HDME can consist of a set of small applications that are characterized by a precise and hardened interface, the possibility of being deployed independently, and ease of integration thanks to a proper middleware for communication and deployment purposes. The use of microservices has important consequences on the provided security and privacy degree; in fact, a microservice approach for processing healthcare-related data is superior to a monolithic approach [4]. If a certain service (e.g., a service that processes images) has a security vulnerability that is exploited, its reach is limited by the sandboxing done by the container system, at the system and network levels (e.g., no access to disk, no way to create a connection to the outside). By also limiting at the communication level what a microservice can communicate with other microservices, the attack surface is greatly reduced. Also, by splitting the application into microservices, each microservice is easier to review from a security perspective. Microservices that do not touch sensitive data can be migrated off of a private cloud into a public cloud if there is a need for scaling or for minimizing costs.

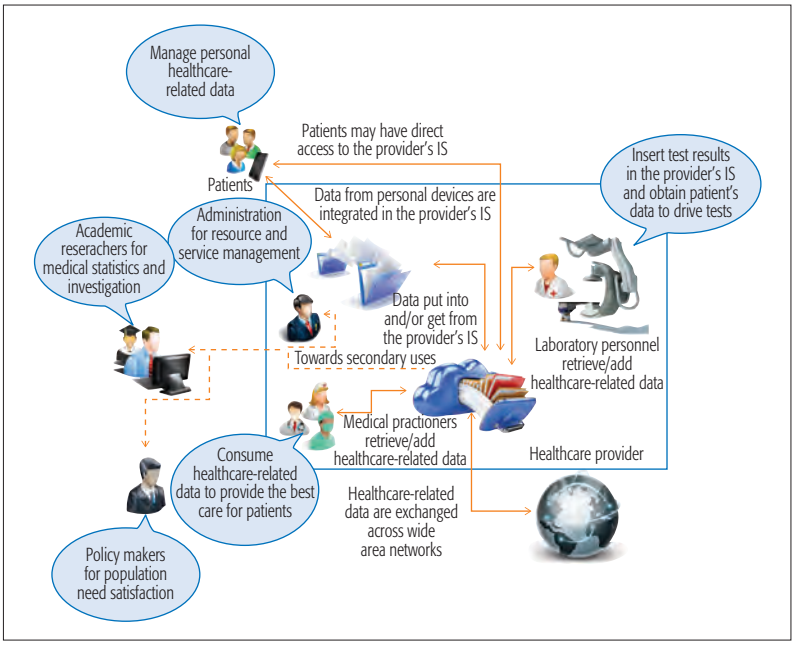

Figure 1. Overview of the communication flows within the healthcare domain throughout the information system (IS) of a provider.

Issues

Security and privacy are critical issues for healthcare providers, considering that the data stored and exchanged by them may contain very sensitive information. More precisely, digitized medical records are open to potential abuses and threats (loss, leakage, theft, misuse, etc.) [5]. The evolution from hospital-size health information systems (HISs) to regional and trans-regional systems, traversing multiple heterogeneous communication networks characterized by different protocols, standards, and conventions, increases the complexity of the security issues to be addressed. This motivates the increasing pressure and interest in properly dealing with security and privacy in the data flows in Fig. 1. Moreover, the possibility of using microservices for architecting the service-oriented software hosted within the cloud further exacerbates the security and privacy issues in the envisioned scenario due to the well-known problem referred to in the literature as the *sharding pattern* [6]. In fact, the driving idea of microservices, that is, distributing the complexity of an application into multiple simple units of computation, has the side-effects of segmenting the application data into multiple horizontal partitions or shards and exhibiting decentralized data governance.

Conclusions

The work described in this article is related to the ongoing debate on the use of cloud-based HDME and the novel security and privacy implications brought by the use of microservices architectures within the mentioned scenario. The contribution is twofold. On one side, we specify the unique security and privacy requirements in cloud-based HDME. On the other side, we propose a set of security and privacy enhancement means to mitigate the vulnerability of data management in a healthcare provider and highlight a possible realization of a robust, multi-layered, and holistic framework for data protection in healthcare.

Cloud-Based Healthcare-Related Data Management and Exchange

All healthcare providers have already abandoned the traditional paper-based characterization of healthcare-related data and moved to their digital representation in the so-called electronic medical records (EMR) [7], which are a digital version of the paper charts in the clinician's office that contains the medical and treatment history of the patients in one practice. However, nowadays it is common for patients to receive healthcare from multiple providers, even if they are not in the same geographical area. Such patient mobility can be motivated by economic reasons (in certain countries treatments are cheaper than in other ones) or the quality of received services (a certain hospital can be specialized in a specific treatment so as to be preferred by patients over other ones close to their place of residence). When a patient receives healthcare at another provider, its updated anamnesis must be recorded, so it is crucial to share patient healthcare-related data across providers. Such a demand promoted the design and realization of the electronic health record (EHR) and the systems supporting their exchange (EHR-Ss) [8], which go beyond the data encompassed in the EMR by including a more comprehensive patient history and are designed to allow information to be shared among multiple providers involved in a patient's care. The EMR and EHR-S represent the core technologies for HIS in hospitals and support the medical personnel in their daily activities; therefore, their interoperation and effectiveness is a crucial issue of great concern. In addition, the need to reduce patient hospitalization and the increasing application of mobile computing and "bring your own computer" in the fruition and access of EHR have paved the way for the advent of a third element in the HIS panorama, personal health records (PHR) and systems (PHR-S) supporting their management [9]. A PHR-S offers the possibility to patients of transparent and instant access to their healthcare information, and also the continuous monitoring of patients' vital signs and their fast delivery to medical staff in order to realize timely alert notification systems.

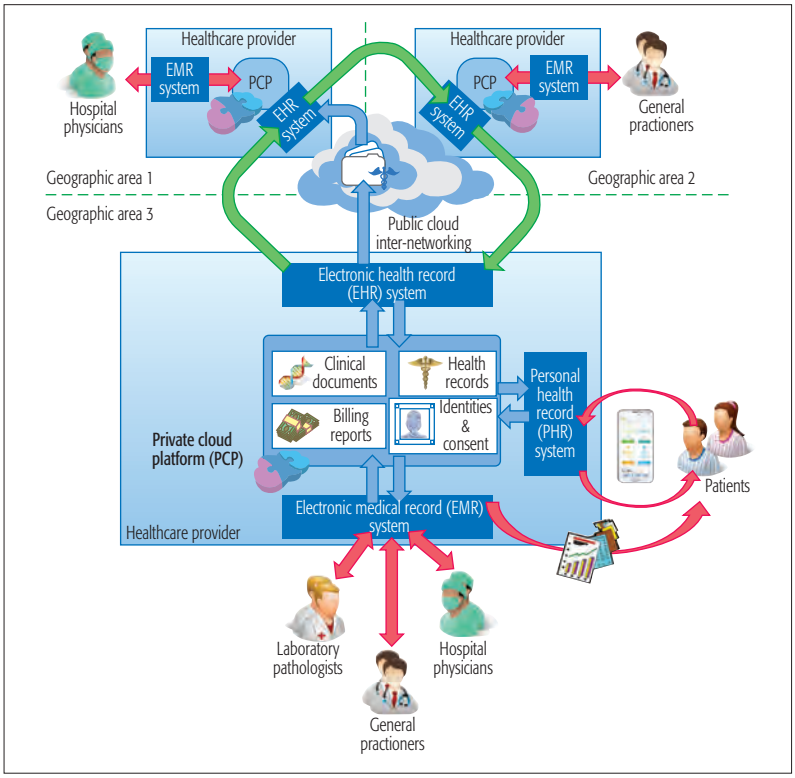

Figure 2. Schematic architecture of cloud-based healthcare-related data management and exchange.

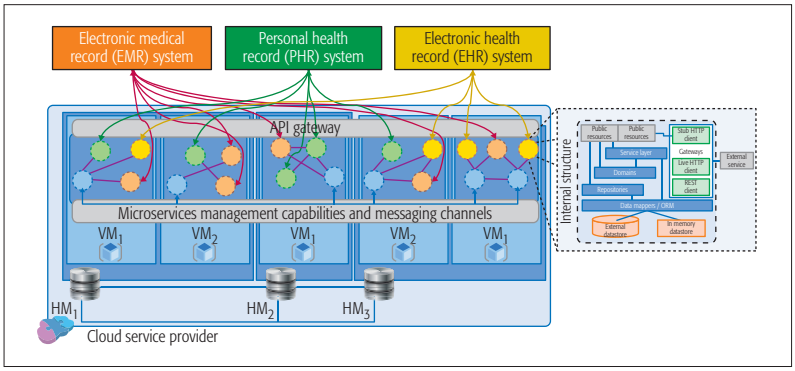

Figure 3. Scheme of the software for HDME as a set of microservices.

Figure 2 illustrates how cloud computing can be adopted within the healthcare domain for medical data management by gluing together EMR-, EHR- and/or PHR-S: each healthcare provider can be equipped with a cloud platform, which may be private, for the storage and sharing of medical data among patients and medical personnel. Such a platform may host services for the management of the identities of all the involved users, and the patient consent on the management of their medical data [10].

The current trend for designing the software hosted by the cloud is to adopt a microservice approach, by segmenting the overall application into multiple units of logic and execution and spreading the application data across these units. Figure 3 represents a microservice architecture for the envisioned software hosted on the cloud for gluing together the eHealth solutions in the figure. In this figure we can see that the microservices are running within containers hosted in virtual machines running in host machines composing the cloud service platform. The platform for managing the microservices have a low-level set of functionalities for their management, deployment, and communication by means of messaging channels. In addition, such a platform also offers high-level functionalities in the so-called application programming interface (API) gateway, providing a seamless view of the microservice architecture. Specifically, the microservice system is accessed through an API gateway that handles security, transformation, and orchestration. Its role is to authenticate and encrypt API requests made to the microservice system, transform requests from other users into requests for microservices, which might mean dealing with different formats of data exchange (e.g., from XML to JSON) or just dealing with protocols (converting from a REST call to a message put on a queue), and aggregate a result from multiple calls to microservices into an API response. Each microservice has an internal structure, as shown in the figure, made of publicly accessible resources and some that are kept privately. Moreover, each microservice has a set of gateways to have access to externals services

Req. ID	Name	Issues	Description
01	Flexible access control	Healthcare-related data may travel through the EHR solutions of different providers, or be access by medical personnel belonging to different providers, each with given access control policies.	Multiple heterogeneous access control methods should coexist so that a user with a security claim can access the data of various healthcare providers.
02	Privacy preserving HDME and shard consistency	Healthcare-related data are considered particularly critical since they can expose sensitive information about the patients and/or personal or financial data useful for fishing or stealing attacks.	Healthcare-related data are particularly critical and must be protected when in motion from/to the cloud and when at rest within the cloud without implying excessive performance and availability side-effects.
03	Breach notification and documentation	Within a solution for EHR sharing and storage, it is probable to have data breaches by taking advantage of emergency situations or an-expected vulnerability of the system.	It is needed to equip a cloud-based HDME solution with proper ways to identify the occurrence of data breaches, and properly report and document them so as to cope with the legal requirements and eventual trials in court.

Table 1. Summary of the identified requirements for HDME.

and resources. Looking at their application logic, it is structured according to the MVC pattern, with distinct application layers for the implementation of the service logic, the proper acquisition and distribution of information with other microservices, and the internal management of the data. In the figure we can distinguish four types of microservices: the ones responsible for interacting with the EMR/PHR and EHR services, and the ones in charge of implementing the integration functionalities of our cloud-based HDME that glues together the three health information systems.

Security and Privacy Means

Requirements

Unfortunately, lessons learned from past experiences highlight that securing such a huge and highly critical ICT infrastructure presents multiple severe challenges mainly concerning the overall communication reliability, integrity, and confidentiality and large-scale identity management. These are combined with all the complexities and technological oddities related to the integration of multiple wired and wireless communications mechanisms as well as communication protocols and standards. Moreover, according to the principles of "privacy by design," a proactive approach that aims to prevent certain events before leading to data breaches must be preferred to a reactive one, which offers remedies for resolving data breaches. Moreover, end-to-end life cycle protection should be enforced by applying security and privacy for the entire life cycle of the health-related data, from start to finish. Last, particular respect for patient privacy must be considered by offering appropriate notice and empowering for data protection with a patient-centric approach. We have summarized the main key requirements for cloud-based HDME in Table 1.

The key concern within healthcare is related to the violation of the privacy of such sensitive information as health-related data. To protect the envisioned cloud-based HDME from external attacks, a proper authorization scheme is needed to be coupled with the solution shown in Fig. 3. An aspect that is evident from Fig. 2 is that the evolution from hospital-size eHealth systems to regional and trans-regional systems federated by means of clouds [11] implies that healthcare-related data must traverse information systems and networks managed by different organizations, belonging to several distinct countries, each having its own laws and privacy enforcement rules. Furthermore, eHealth systems were developed to deal with specific requirements of organizations and providers that use them; thus, integration problems may arise. Among these problems, the most demanding one is related to the coexistence of multiple heterogeneous access control models for organizations integrated by the cloud. Specifically, an ICT infrastructure traditionally adopts a single and precise access control model, which has been agreed and specified by the organization owning it. When data must traverse across multiple organization boundaries, they are subject to multiple authorization means, without the possibility of imposing a unique access model within the overall integrated regional and/or trans-regional EHR. There is a requirement for a flexible access control model that lets the heterogeneous ones coexist and interoperate. Such a problem has been left unaddressed by the IHE and HL7 standards for access control within healthcare, since they are limited in introducing means to model and deploy an access control solution, and guarantee technological and syntactical interoperability of the security claims, leaving interoperability unattended at the semantic and process levels.

In order to protect a solution for HDME from internal attacks, a proper approach is to adopt cryptographic primitives in order to achieve a privacy preserving HDME for data in motion and at rest. The sharding problem introduced by the microservices further exacerbates such a problem, since the data at rest is scattered in multiple locations of the infrastructure and must be kept consistent in case of changes. For this reason, it is important to identify a proper cryptographic primitive to cope with the needs and issues imposed by the sharding problem caused by the microservice approach. In general, encryption is already provided by many cloud service providers (CSPs) as a means to offer some privacy guarantees, allowing users to keep a sort of control on their own data, thus leaving the CSP, assumed to be untrustworthy, unable to access the outsourced data. However, such solutions are not suitable for the sharding problem, since they do not allow computations to be made unless the encrypted data are previously decrypted, leaving the microservices unable to use the encrypted shards. More-

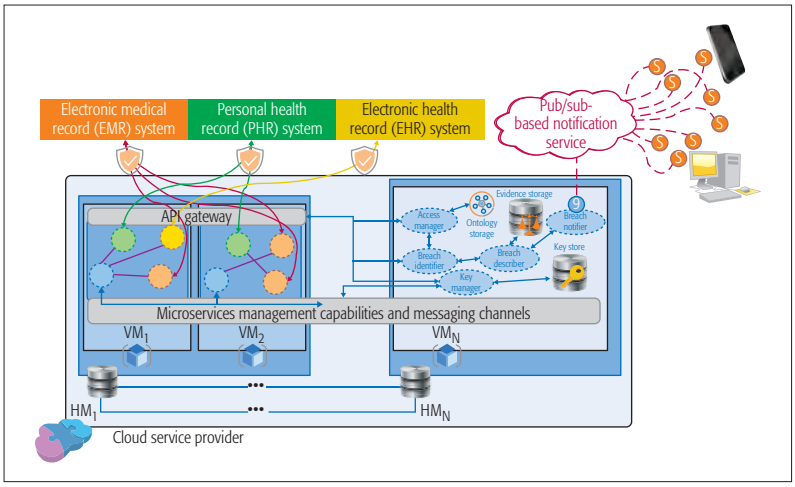

Figure 4. Proposed architecture of a secure manager for cloud-based HDME.

over, the healthcare domain is characterized by stringent timeliness and availability requirements: during an emergency, a doctor needs to quickly retrieve a datum of interest, and such a datum must be provided even if the doctor is not able to expose valid credentials or security claims, according to a so-called break-the-glass action [12] in emergency situations.

Figure 4 highlights the key elements of the security enforcement solution for our envisioned cloud-based HDME, with the intent of having proper means to deal with the requirements in Table 1. Specifically, the API gateway is properly equipped with HTTP handlers for semantic access control, so as to realize a flexible authorization solution for the microservices hosted within the cloud, for the identification of breaches and their documentation, and for applying encryption to the exchanged messages with the external services. In addition, the microservices management capabilities are augmented with the possibility of applying encryption for the data at rest. Our security manager is composed of the proper interconnection of microservices hosted in a dedicated virtual machine (VM), encompassing the following elements:

- A microservice acting as access manager and interacting with proper storage for the semantic representation of access control models.
- A key manager in charge of dealing with the generation or revocation of the cryptographic keys stored in a proper manner.
* A breach identifier that collects all the possible hints on breaches that have occurred, passes them to a proper microservice responsible for filling up an audit log related to these hints and properly stored in a reliable and secure manner.
- A last microservice responsible for notification duties.

Semantic Access Control

The problem of having multiple access control models that must coexist cannot be resolved by imposing a single standard model. On the contrary, the interoperability of the access control models can be achieved by resolving the differences of the terms used to express the access control rules and policies of heterogeneous organizations. Due to the scale of the healthcare organizations integrated by our proposed cloud-based HDME solution depicted in Fig. 3, it is not efficient to have a manual resolution of the mentioned issues done by one or more experts, so automatic tools are required. To this aim, it is necessary to have a formal description of each access control model so that a computer program can resolve the differences and automatically match the different models. In [13], we have used ontologies for the description of access control models by specifying subjects, and their attributes and contexts, and our novelty with respect to the available literature on this topic is that we do not impose a given model, such as role-based access control or policy-based access control, but offer freedom to specify any possible access control model. This is a first level of interoperability that our solution can provide. Moreover, each model can use a given term to indicate the elements of the models; for a concrete example, a subject can be indicated as citizen or user of a given municipal service. An ontological representation of this ters is able to relate terms that are syntactically different but share the same semantic.

Figure 5 provides an example of the ontological representation of the set of users that are interested in requesting services and data provided by the EHR solution of a given healthcare provider (named domain ontology), coupled with a description of the access control rules adopted by the healthcare provider (indicated as control ontology) and a specification of the the patient consent to share the healthcare-related data through cloud-based HDME (referred to as consent ontology in the figure) by considering the semantic modeling of citizen consent in [14]. This has the benefit of making the interpretation of the policies across different organizations much easier. Based on such a model, it is possible to specify the set of security policies to be checked in order to allow a received request or to deny it. In our work we have used ASK forms provided by SPARQL as a means to express access rules, and for concrete examples, Fig. 5 contains three SPARQL predicates representing the following three policies:

1. Obtain access to the events related only to the patients currently treated.
2. Administrative staff can gain access to the events generated in the given laboratory only within the accounting period.
3. A doctor can obtain all the events containing the results of medical tests.

When an access control scheme is modeled as an ontology, it is possible to exchange it with other organizations and have an automatic mapping among the heterogeneous ontology schemes. Matching diverse ontologies is still an open issue in the current literature, and a survey on this topic is available in [15], but in our work, we have adopted a simple approach based on the semantic similarity of the terms composing two diverse ontologies.

Privacy Preserving Data Storage

There is strong interest in outsourcing health-related data to cloud storage in order to optimize the increasing costs related to data management. The evidence of the high data breaches occurring

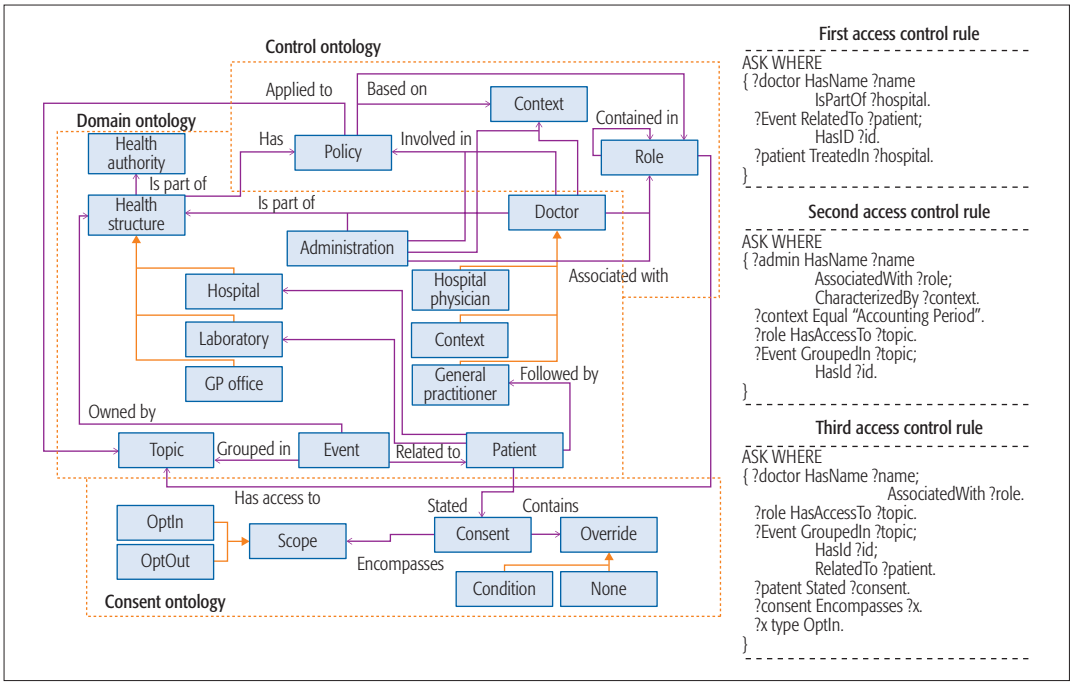

Figure 5. Example of an ontological representation of a pharmaceutical benefits advisory committee model for an healthcare-organization, and a specification of access rules as SPARQL predicates.

at third parties collecting and controlling massive amounts of personal data question the validity of such a deployment model and call for efficient solutions in order to have privacy-preserving cloud storage. Within the context of academic research, several approaches have been proposed to deal with the privacy preservation problem and current cloud service providers: blockchain for healthcare-related data management and homomorphic encryption for microservices.

Data Breach Notification

Data breaches can occur in several ways. Some examples may include lost or stolen devices, databases storing personal information hacked or illegally accessed by malicious users, employees accessing or disclosing personal information outside the requirements or authorization of their employment, an individual deceiving an agency or organization into improperly releasing the personal information of another person, and so on. For this reason, organizations that process personal data must take appropriate actions against unauthorized or unlawful processing, accidental loss, destruction, or damage of personal data. Data breaches can be identified by means of techniques for assessing the authenticity and integrity of exchanged data, such as digest, digital signatures or watermarking in the case of multimedia data. In addition, emergency access to the EHR system due to a break-the-glass action should be marked as a potential breach to be checked after its occurrence.

The key part of several legal frameworks, especially in the European context, is the proper notification of the eventual occurrence of data breaches in the healthcare domain. Within our solution a widely known technology for publish/subscribe services, such as the ones based on the OMG DDS standard, has been identified to be integrated in our solutions of reference, in order to implement the mentioned data breach notification service. This is motivated by the intrinsic decoupling and scalability guarantees that this kind of middleware is able to provide and the rich set of quality of service (QoS) that is able to tune in order to achieve resilient and secure notification.

Data breaches can be identified by means of techniques for assessing the authenticity and integrity of exchanged data, such as digest, digital signatures or watermarking in case of multimedia data. In addition, an emergency access to the EHR system due to a break-the-glass action should be marked as a potential breaches to be check after its occurrence.

Experimental Evaluation

An experimental campaign has been conducted to quantify the performance costs for data exchange though different domains. We considered several HISs to exchange healthcare documents with and without using the proposed security solution. We send a query to the regional HIS, then delegate to another HIS. The aim of the experimental evaluation was to show the feasibility of the new proposed system. We measure the response time (in milliseconds) for queries and document retrievals (Fig. 6). The average time to perform retrieval is increased by approximately 25 percent in the case of encryption. The queries among regions have a measured response time increased by 10–75 percent in the case of encryption. We can conclude that the protection cost of the created infrastructure supports the security and privacy of exchanged medical data, the main requirements imposed by our proposed solution. In this approach, the packet loss and any delays were not considered because we used reliable TCP communications.

Lessons Learned

In this article we investigate the problems of security and privacy caused by using cloud computing for the management and exchange of healthcare-related data among healthcare providers. The microservice approach copes with these problems, due to the provided isolation, but there are still some problems left unaddressed. We describe three main ones, and propose promising solutions

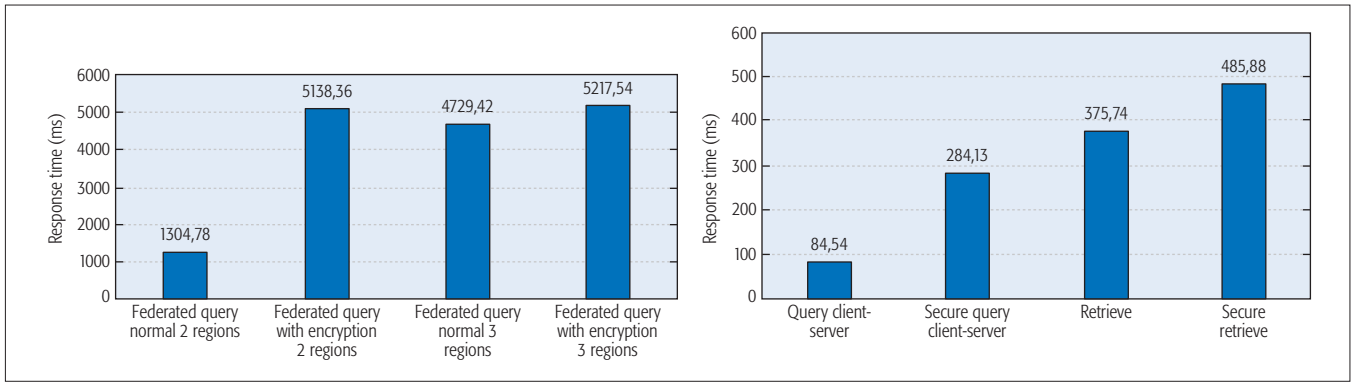

Figure 6. Performance of test scenarios.

for their resolution within the context of microservices hosted in the cloud for managing and distributing healthcare-related data. We present in detail the security and privacy means by describing the requirements, a novel model of access control based on semantics, privacy preserving of data storage, and data breach notification. The experimental results only show the feasibility of the new proposed system in a specific scenario.

Acknowledgments

The authors would like to acknowledge the contribution of the CA COST Action CA15140-ImAppNIO, and the following projects: DataWay (PN-II-RU-TE-2014-4-2731) and MobiWay (PN-II-PT-PCCA-2013-4-0321). We would like to thank the reviewers for their time and expertise, constructive comments, and valuable insights.

References

[1] T. Schabetsberger *et al.*, "From a Paper-Based Transmission of Discharge Summaries to Electronic Communication in Healthcare Regions," *Int'l. J. Medical Informatics*, vol. 75, no. 3–4, 2006, pp. 209–15.
[2] J. M. Batalla *et al.*, "Efficient Media Streaming with Collaborative Terminals for the Smart City Environment," *IEEE Commun. Mag.*, vol. 55, no. 1, Jan. 2017, pp. 98–104.
[3] C. Chilipirea *et al.*, "Cloud Elasticity: Going Beyond Demand as User Load," *Proc. 3rd Int'l. ACM Wksp. Adaptive Resource Management and Scheduling for Cloud Computing*, 2016, pp. 46–51.
[4] I. Nadareishvili *et al.*, *Microservice Architecture: Aligning Principles, Practices, and Culture*, O'Reilly Media, Inc., 2016.
[5] Y. Nikoloudakis *et al.*, "A Fog-Based Emergency System for Smart Enhanced Living Environments," *IEEE Cloud Computing*, vol. 3, no. 6, 2016, pp. 54–62.
[6] C. H. Costa *et al.*, "Sharding by Hash Partitioning — A Database Scalability Pattern to Achieve Evenly Sharded Database Clusters," *Proc. 17th Int'l. Conf. Enterprise Info. Systems*, Apr. 2015.
[7] M. Steward, "Electronic Medical Records," *J. Legal Medicine*, vol. 26, no. 4, 2005, pp. 491–506.
[8] K. Häyrinena, K. Sarantoa, and P. Nykänenb, "Definition, Structure, Content, Use and Impacts of Electronic Health Records: A Review of the Research Literature," *Int'l. J. Medical Informatics*, vol. 77, no. 5, May 2008, pp. 291–304.
[9] I. Iakovidis, "Towards Personal Health Record: Current Situation, Obstacles and Trends in Implementation of Electronic Healthcare Record in Europe," *Int'l. J. Medical Informatics*, vol. 52, no. 1–3, Oct. 1998, pp. 105–15.
[10] C. Esposito *et al.*, "Interconnecting Federated Clouds by Using Publish-Subscribe Service," *Cluster Computing*, vol. 16, no. 4, 2013, pp. 887–903.
[11] N. Sultan, "Making Use of Cloud Computing for Healthcare Provision: Opportunities and Challenges," *Int'l. J. Info. Management*, vol. 34, no. 2, Apr. 2014, pp. 177–84.
[12] R. Gajanayake, R. Iannella, and T. R. Sahama, "An Information Accountability Framework for Shared eHealth Policies," *Data Usage Management on the Web: Proc. WWW2012 Wksp.*, 2012, pp. 38–45.
[13] C. Esposito, A. Castiglione, and F. Palmieri, "Interoperable Access Control by Means of a Semantic Approach," *Proc. AINA Wksps.*, 2016, pp. 280–85.
[14] E. Coiera, and R. Clarke, "e-Consent: The Design and Implementation of Consumer Consent Mechanisms in an Electronic Environment," *J. American Medical Informatics Association*, vol. 11, no. 2, 2004, pp. 129–40.
[15] E. Rahm, and P.A. Bernstein, "A Survey of Approaches to automatic Schema Matching," *The VLDB J.*, vol. 10, no. 4, 2001, pp. 334–50.

Biographies

Christian Esposito (esposito@unisa.it) received his Ph.D. degree in computer engineering from the University of Napoli "Federico II," Italy, where he is an adjunct professor. He is also a research fellow and adjunct professor at the University of Salerno, Italy. He regularly serves as a reviewer, organizer, andGuest Editor for international conferences and journals. His research interests include reliable and secure communications, large-scale distributed systems, positioning systems, multi-objective optimization, and game theory.

Aniello Castiglione [M] (castiglione@ieee.org) is an adjunct professor of computer science at the University of Salerno and the University of Naples "Federico II." His research interests include security, communication networks, information forensics and security, and applied cryptography. He has a Ph.D. in computer science from the University of Salerno. He is on several Editorial Boards of high rank international journals. He is a member of several associations, including ACM.

Constantin-Alexandru Tudorica (constantin.tudorica@cti.pub.ro) is a Ph.D. student in the Computer Science Department, University Politehnica of Bucharest. His research interests are in the area of scheduling and resource management in distributed computing, with a special focus on a microservice approach to current systems. His technical specialties are C/C++, PHP, MySQL, algorithms, distributed algorithms, and large-scale distributed algorithms. He currently works on understanding the performance of data center resource managers.

Florin Pop (florin.pop@cs.pub.ro), professor, Ph.D., Habil., received his Ph.D. in computer science at the University Politehnica of Bucharest in 2008. His main research interests are large-scale distributed systems, adaptive and autonomous methods, multi-criteria optimization methods, grid middleware, prediction methods, self-organizing systems, data retrieval and ranking techniques, contextualized services in distributed systems, and evaluation using modeling and simulation. He is an active reviewer and Guest Editor for several journals.

NETWORK TESTING AND ANALYTICS

Ying-Dar Lin | Irena Atov | Erica Johnson

This Series continues to receive and report interesting works on network testing as well as network analytics. On network testing, we see increasingly sophisticated tools or even a framework of tools that can test beyond what existing commercial tools could do and render results with good insights. They have the potential to further push the state of the art in network testing. On network analytics, insightful observations from real data collected from real systems are becoming the essential theme for submissions accepted here. We shall continue to solicit interesting works in both tracks.

For this September issue we received 12 submissions and accepted three after two rounds of review. The first article reports a benchmarking framework for network functions virtualization (NFV), while the second article presents a tool coordinating multiple wireless technologies. Both of them are on a framework integrating several tools. The third article is on network analytics instead of network testing. It demonstrates how machine learning can be applied to network failure diagnosis.

NFV maintains a fast-growing pace toward the realization of virtual network functions (VNFs) in virtualized execution environments with the reliability and performance expectations of existent telco/carrier services. R. V. Rosa, C. Bertoldo, and C. E. Rothenberg present the Gym open source framework, a prototyped skeleton of components well suited for customized development of arbitrary VNF testing methodologies, aiming to support different stages of a VNF life cycle in the spirit of DevOps. After introducing the architecture and presenting the design and implementation choices, Gym is validated through a series of automated benchmarking tests targeting different vIMS deployments, instantiated using OpenStack image flavors as a comparative factor. Gym extensibility is evidenced through the easy support of custom tools to stimulate the vIMS components and measure the target performance metrics. The obtained performance profiles compare the overall percentage of vCPU consumption of each vIMS component associated with the overall rate of accepted calls. Particularly, the results illustrate the saturation of vCPU resources in the case of vIMS deployment with m1.small flavor. However, the expected efficiency gains do not sustain (approximately from 75 to 90 percent) when comparing an m1.medium with an m1.large vIMS instance, exhibiting high variability and confirming the challenges of VNF benchmarking.

The paradigm shift toward the Internet of Things and the growing interest of the fifth generation (5G) community in the unlicensed spectrum result in ever increasing contention for the shared industrial, scientific, and medical (ISM) bands. Coexistence of these heterogeneous wireless networks will be a huge challenge, as they have to cooperate to share the same spectrum efficiently. The architecture presented by P. Ruckebusch *et al.* facilitates building solutions for coordinating medium access in such networks, ultimately allowing them to coexist. For this purpose, it offers wireless researchers and developers the ability to control wireless devices using a unified set of programming interfaces (UPIs) in a context-aware manner. An open source implementation is available for a diverse set of devices and operating systems (Linux, Contiki, Windows), currently supporting both standardized technologies (IEEE 802.11, IEEE 802.15.4, LTE) as well as advanced reconfigurable radio systems (IRIS, GNURadio, WMP, TAISC). In terms of memory overhead, only 4 kB ROM and 1 kB RAM is required to enable control on a resource-constrained Contiki device. The showcase experiments demonstrate that with minimal complexity, UPIs allow coordinating medium access on a millisecond level between IEEE 802.15.4 Contiki nodes and IEEE 802.11 Linux nodes, effectively enforcing a cross-technology time-division multiple access (TDMA) scheme.

In root cause analysis of network failures using machine learning and summarization techniques, J. Navarro *et al.* propose a novel approach for performing root cause analysis on networks. Exploiting the capabilities of machine learning, they present a new method of analyzing a network and extracting information about it with little to no user input required. It does so through the creation of the Influence Matrix, a new, efficient approach for summarizing a large number of models. As a validation example, exploring a dataset containing more than 20,000 events over a 200-day timespan, divided into 548 different types of events (of which 27 are critical and 65 are of major severity) obtained from a banking services network, the authors show how their proposal is able to reason and find important facts, such as the need to improve the event policy recollection; the appearance of a small number of significant event-type clusters; the pairing mode of work of virtual machines; chains of events that lead to critical situations in the network; and especially dangerous events that influence the appearance of a large number of errors. All of these findings are easy to grasp and require no extensive knowledge of the studied network.

BIOGRAPHIES

YING-DAR LIN [F] (ydlin@cs.nctu.edu.tw) is a Distinguished Professor at National Chiao Tung University, Taiwan. He received his Ph.D. in computer science from the University of California Los Angeles in 1993. He is the director of the Network Benchmarking Lab, which reviews network products with real traffic and is an approved test lab of the Open Networking Foundation (ONF). He is an IEEE Distinguished Lecturer and an ONF Research Associate. He co-authored *Computer Networks: An Open Source Approach* (McGraw-Hill, 2011).

IRENA ATOV [SM] (i.atov@ieee.org) received her Ph.D. in electrical engineering from RMIT University, Australia, in 2003. She is currently a principal architect at Microsoft in their Skype for Business Core Engineering Group. Previously, she has worked in academia in both teaching and research roles, consulted for industry through her own company, and worked for Telstra in Melbourne, Australia as program director of Network Analytics and Resilience. Her research has led to the development of several commercial IT software products.

ERICA JOHNSON (erica.johnson@iol.unh.edu) combines business acumen and an in-depth understanding of complex networking technology to direct the University of New Hampshire InterOperability Laboratory (UNH-IOL). In recognition of her ability to drive technical innovation, *Fierce Telecom* named her in the publication's 2011 Women in Wireline. She serves as an IPv6 Ready Logo Regional Officer, IPv6 Forum Fellow, and USGv6 Test Program lead. She received her Bachelor of Computer Science and M.B.A. from UNH in 2001 and 2011, respectively.

Take Your VNF to the Gym: A Testing Framework for Automated NFV Performance Benchmarking

Raphael Vicente Rosa, Claudio Bertoldo, and Christian Esteve Rothenberg

The authors introduce Gym as their proposed testing framework and methodology for automated NFV performance benchmarking. They present their design principles and the outcomes from a practical validation on a vIMS scenario. A discussion of the lessons learned and the overall NFV performance testing landscape are further contributions of this article.

ABSTRACT

A VNF is a software entity to be run in diverse execution environments with variable configuration options and capabilities (e.g., hardware acceleration) impacting performance. NFV resource multiplexed infrastructures can impose hard-to-predict relationships between VNF performance metrics (e.g., latency, frame loss), the underlying allocated resources (e.g., units of vCPU), and the overall system workload. Characterized by many-fold platform configuration and environment variables, the evolving scenario of NFV calls for adequate testing methodologies embracing modern continuous development and integration practices and leveraging open source tools and mindset. To this end, we introduce Gym as our proposed testing framework and methodology for automated NFV performance benchmarking. We present our design principles and the outcomes from a practical validation on a vIMS scenario. A discussion of the lessons learned and the overall NFV performance testing landscape are further contributions of this article.

INTRODUCTION

As network functions virtualization (NFV) matures through the realization of proof of concept implementations (http://nfvwiki.etsi.org/index.php?title=PoCs_Overview, accessed Jan. 6, 2017) identified challenges toward wider rollouts include the need for carrier-grade testing and operational standards to match the service continuity and performance predictability levels of current physical infrastructures [1, 2]. Being pure software entities, virtual network functions (VNFs) lend themselves to continuous deployment and integration following agile DevOps methodologies. As illustrated in Fig. 1, software-oriented processes applied to NFV call for automated testing practices spanning platform portability, functional correctness, and performance benchmarking for each candidate VNF version before making it available for deployment. A single line of code change passing all functional tests could also undermine the VNF performance for specific workloads and platforms — a risk that calls for standardized testing methods [3–6] toward adequate VNF benchmarks (e.g., [7–9]).

The heterogeneity of NFV infrastructure (NFVI) environments include diverse virtualization options and system capabilities (e.g., hardware [HW] offloading, kernel bypassing) for varying workloads and diverse resource sharing conditions (e.g., co-located VNFs trashing shared CPU memory caches) [1]. Such a multi-dimensional testing landscape with multiple configuration knobs (Fig. 2) introduces unprecedented challenges toward useful performance profiles delivering valuable assessments for different stakeholders at different stages, for example, during VNF development, for pre-deployment NFVI validation, or even for service level agreement (SLA) compliance at runtime.

In our initial work on VNF benchmarking as a service (VBaaS) [3], we introduced the problem statement of VNF benchmarking based on "trust, but verify" principles in seeking standardized performance testing allowing proper evaluation of candidate platforms and locations to host (chains of) VNFs with respect to target key performance indicators (KPIs). In this article, we revisit our functional and architectural vision of VBaaS based on the prototype development and practical evaluation of Gym, the proposed testing framework that allows automated performance benchmarks of NFV embodiments. We advocate for a framework that defines a minimum set of standardized interfaces while allowing user-defined tests along a catalog of reusable VNF testing procedures and reports with widely and well defined system configuration descriptors, workload parametrization (linking to specific traffic generation tools and their parameters), KPI computation, along with all supporting code and data expected from a standardized and reproducible benchmarking methodology.

Outcomes of automated performance tests can be used as inputs of NFV orchestrator (NFVO) embedding algorithms (cf. [4]) and/or parameters to support business decisions such as pricing and allocation of resources to fulfill SLAs. As noted by the vision behind NFV-VITAL [9], standardized characterization of VNF performance enables analyzing optimal sizing and configuration of VNFs in order to automatically:

- For a given resource configuration, estimate the VNF capacity.
- For a given workload, determine optimal resource configuration.
- Evaluate different operating system (OS) virtualization/HW alternatives and compute

Digital Object Identifier:
10.1109/MCOM.2017.1700127

The authors are with the University of Campinas.

0163-6804/17/$25.00 © 2017 IEEE

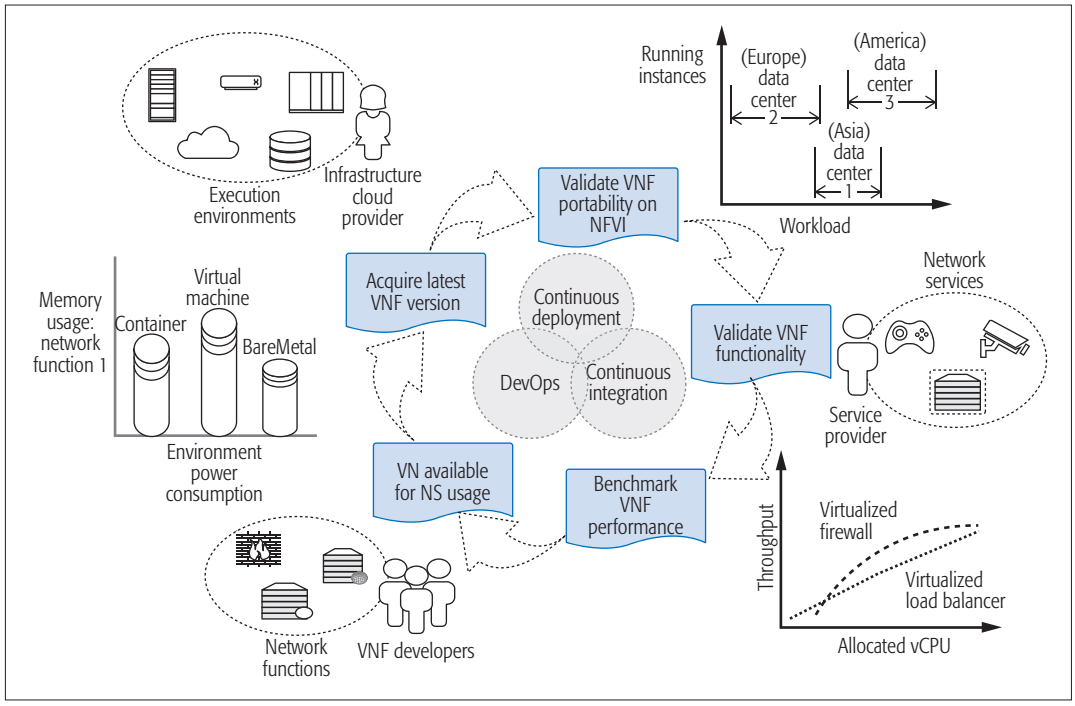

Figure 1. Gym motivation: big picture of VNF benchmarks as part of rapid service processes though automation, regression, and performance testing.

system overhead associated with dynamic scaling (up/out/down/in).
- Fine-tune VNF implementation and performance debugging (i.e., "if you cannot measure it, you cannot improve it" — William Thomson, known as Lord Kelvin).

The article is organized as follows. The next section presents the approach and design of the Gym framework and presents a generic workflow to illustrate the main functionalities. Then we put Gym into practice by performing benchmarking tests on an open source IP Multimedia Subsystem (IMS) implementation. Then we have an open discussion on the achievements and limitations, considering aspects of the vIMS experiments as well as more general aspects and challenges of VNF testing framework design and implementation. Following that, we discuss the vibrant related work before the final remarks and conclusions.

Gym: From Design to Implementation

Taking root in the former design efforts of VBaaS [3], our early envisioned abstractions evolved into the framework implementation baptized as Gym. Our approach is based on the development of a skeleton of software components delivering the abstractions and toolset in support of practical methodologies to validate, benchmark, and dimension VNFs [6]. Gym is mainly characterized by:
- Modular architecture with standalone programmable components
- Simple messaging system following generic remote procedure call (RPC) guidelines
- Extensible set of testing tools and target metrics
- Rich test definition through dynamic compositions of modules
- Flexible methods for output processing and results visualization

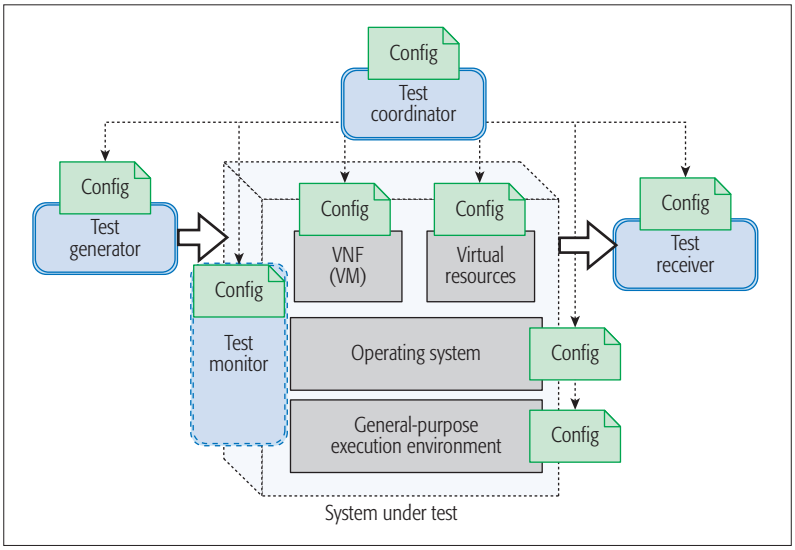

Figure 2. VNF under test scenario illustrating the multiple configuration knobs and the diverse multiple system and platform variables involved.

As shown in Fig. 1, Gym aims at introducing new opportunities to different NFV actors. VNF developers can rely on the framework to add automated, repeatable VNF performance profiling to their agile continuous integration and DevOps practices. Service providers might enhance offered quality of service (QoS) with tested-deployed scenarios (e.g., varying workloads at multiple sites), containing transparent sets of operational VNF metrics, targeting continuous deployment. Cloud/infrastructure providers, when extensively testing VNFs in their execution environments, can use Gym to implement SLA compliance methods to increase the infrastructure reliability and operational efficiency (e.g., energy consumption).

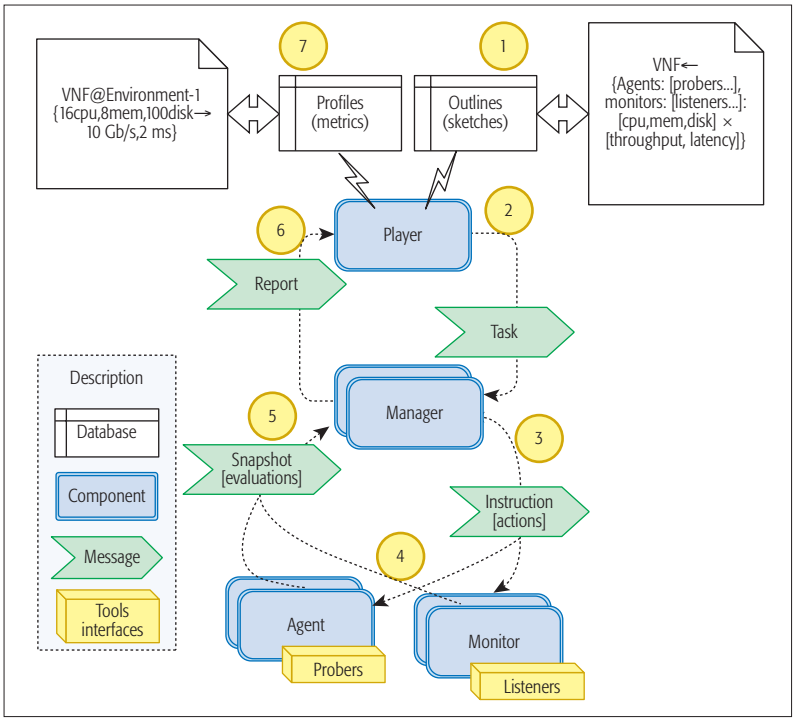

Figure 3. The Gym architecture is based on four main components (Agent, Monitor, Manager, Player), allowing flexible workflows and embodiments as illustrated by the various message exchanges, interfaces, databases, and tools.

Conceptual Ideas and Guiding Principles

Design for modularity is one of the main guiding principles of Gym to allow independent software components to be orchestrated on demand based on well defined testing objectives without compromising customization and overall extensibility. To address the heterogeneous and complex set of requirements and capabilities of NFV instantiations, the framework offers a high degree of freedom through user-defined composition of sets of tools and evaluation models using simple description formats. Gym's overall principles, enunciated below, come later in further discussion when evaluating a VNF benchmarking use case. The proposed guiding principles to design and build a performance testing framework can be compounded in multiple practical ways for multiple VNF testing purposes:

- Comparability: Output of tests shall be simple to understand and process, in a human-readable format, coherent, and easily reusable (e.g., inputs for analytic applications).
- Repeatability: Test setup shall be comprehensively defined through a flexible design model that can be interpreted and executed by the testing platform repeatedly but supporting customization.
- Configurability: Open interfaces and extensible messaging models shall be available between components for flexible composition of test descriptors and platform configurations.
- Interoperability: Tests shall be ported to different environments using lightweight components.

Architecture

The system architecture of Gym is illustrated in Fig. 3 and comprises the following four main modules.

Agent: Provides extensible interfaces for testing tools (e.g., iperf, ping), named *probers*, to create stimulus in order to collect network and host performance metrics. Agents enable both local (e.g., CPU and disk I/O benchmarks) and distributed (e.g., end-to-end latency/throughput between Agents) measurements, and expose modular application programming interfaces (APIs) for flexible extensibility (e.g., new probers). Agents receive *instructions* from a Manager defining sets of *actions* to consistently configure and run prober instances, parse the results, and send back *snapshots* containing output evaluations of the probers' actions.

Monitor: Performs internal and external instrumentation of VNFs and their execution environments in order to extract passive metrics using monitoring tools (e.g., top, tcpdump) interfaces, named *listeners*. Monitors can work jointly with Agents' workloads, for instance, when the VNF throughput shall be correlated with the vCPU utilization. Similar to the Agent, Monitors interact with the Manager by receiving instructions and replying with snapshots. Different from the generic VNF prober approach of the Agent, Monitors may listen to particular metrics according to capabilities offered by VNFs and their respective execution environment (e.g., CPU cycles of DPDK-enabled processors).

Manager: Responsible for:

- Keeping a coherent state and consistent coordination of the managed components (Agents and Monitors), and their features and activities
- Interacting with the Player to receive tasks and decompose them into a concrete set of instructions
- Processing snapshots along proper aggregation tasks into reports back to the Player

Player: Defines a set of user-oriented, northbound interfaces abstracting:

- Metric extraction descriptors, named *Sketches*, according to the requirements and settings of probers/listeners
- VNF testing *Outlines* containing one or more sketches with their configurable parameters. A Player might store different outlines, and trigger their execution when receiving a testing Layout request that might reference one or more parametrized *outlines*, which are decomposed into a set of *tasks* orchestrated by Managers to obtain the *reports*. Interfaces are provided for storage options (e.g., database, spreadsheets) and visualization of the extracted reports into *profiles*.

Two relevant terms deserve further explanation.

Outline: Used as input by the Player module, it defines how to test one or more VNF types following a particular syntax in YAML to express structural settings (e.g., Agents/Monitors topology) and functional properties (e.g., probers/listeners parameters), named *sketches*.

Profile: Refers to the formatted outcome composed by the outputs of an outline execution and the requested Layout scenario. A profile represents a mapping between virtualized resources (e.g., vCPU, memory) in a given environment and VNF performance/benchmarking metrics (e.g., throughput, latency between in/out or ports),

abstracting VNF allocation with certain resources to deliver a unifying metric for a given (predictable/measured) performance quality.

Messaging System and Workflow

Gym core components communicate through a Representational State Transfer (REST) API using generic remote procedure calls (RPCs) with custom JSON message formats. In the following, we describe a generic workflow based on request-reply message exchanges and pairwise component interactions represented as numbered (1 to 7) circles in Fig. 3.

1. The first step consists of a user defining the composition of the VNF testing Outline through sketches containing the structural and functional requirements to express target performance metrics to generate a VNF profile.

2. The Player processes the parametrized outline considering the features offered by the associated Manager(s). The output is a workflow of tasks, in sequence or parallel, submitted to a selected Manager that satisfies (i.e., controls a matching set of Agents/Monitors) the outline requirements. Based on input variables, an outline can be decomposed into different sets of tasks with the corresponding high-level probers/listeners parameters.

3. The Manager decomposes tasks into a coherent sequence of instructions to be sent to Agents and/or Monitors. Inside each instruction, sets of actions define parametrized execution procedures of probers/listeners. Sequential or parallel tasks may include properties to be decomposed into different sets of instructions, for instance, when sampling cycles might define their repeated execution.

4. By interpreting action into a prober/listener execution, an Agent or Monitor performs an active or passive measurement to output metrics via a pluggable tool. A VNF developer can freely create a customized prober or listener to interface her tests and extract particular metrics. An interface of such a tool is automatically discovered by an Agent/Monitor and exposed as "available" to Managers and Players along the corresponding execution parameters and output properties.

5. After computing the required metrics, a set of evaluations (i.e., parsed action outputs) integrate a so-called snapshot sent from an Agent/Monitor to the Manager. A snapshot associated with a specific task is received from the Agent/Monitor that received the corresponding instruction. An evaluation contains timestamps and identifiers of the originating prober/listener, whereas a snapshot receives an Agent/Monitor unique identifier along with the host name information.

6. After processing all the instructions' related tree of snapshots, the Manager composes a report as a reply to each task requested by the Player. The Manager can sample snapshots in a diverse set of programmable methods. For instance, a task may require cycles of repetition, so the correspondent snapshots can be parsed and aggregated in a report through statistical operations (e.g., mean, deviation, confidence intervals).

7. Finally, the Player processes the report following the profile metrics definition, as established initially during the outline decomposition. While the profile contains filtered evaluation metrics and parameters, snapshots can be aggregated/sampled into a report. Results can be exported in different file formats (e.g., csv, json, yaml) or saved into a database for further analysis and visualization. For instance, in our current Gym prototype we integrate two popular open source components, the Elasticsearch database and the Kibana visualization platform — tools providing high flexibility in querying, filtering, and creating different visual representations of the extracted profiles.

To address the heterogeneous and complex set of requirements and capabilities of NFV instantiations, the framework offers a high degree of freedom through user-defined composition of sets of tools and evaluation models using simple description formats.

A Day in the Gym: vIMS Performance Testing

We now exercise the Gym framework in the case study of IMS telecom network functions [10] in the scope of European Telecommunications Standards Institute (ETSI) NFV ISG proofs of concepts ("Multi-vendor on-boarding of vIMS on a cloud management framework"; http://www.etsi.org/technologies-clusters/technologies/nfv/nfv-poc, accessed Jan. 6, 2017) and also target related work on VNF benchmarking [9]. Understanding vIMS performance for a given resource allocation can be relevant for NFVI configuration and setting adequate system parameters in light of potential price and SLA offerings in multi-vendor cloud infrastructures. Throughout the practical evaluation of the chosen use case, we assess Gym functionalities and further discuss the development of testing methodologies following the architectural design aspects introduced by Gym.

Prototype Implementation

Gym provides a framework for VNF testing and features off-the-shelf Linux tools to generate generic metrics. More specifically, the implemented Monitor component includes the following.

Host listener: A listener process to record resource metrics such as the utilization of CPU, memory, disk, and network. Based on the psutil cross-platform library (https://pypi.python.org/pypi/psutil, accessed on 2017-06-01) and parametrized by interval (sampling rate) and duration, the host listener is able to extract multiple host runtime metrics (approximately 80+).

User-defined, VNF-specific metrics are supported through component extensions. For our vIMS benchmarking purposes, the following prober was added to the deployed Agent.

SIPp Prober: A prober integrated to interface the SIPp (http://sipp.sourceforge.net, accessed Jan. 6, 2017) open source SIP traffic generator used to stress the vIMS deployment based on the four-step SIP registration procedure [11], herein called *transaction*. The Gym outline and profile describing the SiPp prober input parameters and output metrics are shown in Fig. 4 and publicly available (https://github.com/intrig-unicamp/gym, accessed Jan. 6, 2017).

Scenario and Testbed

The vIMS under test is the Clearwater open source project (http://www.projectclearwater.org/, accessed Jan. 6, 2017), which provides virtual machine (VM) form factors for the different IMS network functions, namely Edge Proxy/P-CSCF ("Bono"), SIP Router I/S-CSCF ("Sprout"), and the HSS Cache ("Homestead"). We opt for OpenStack as the virtualized infrastructure man-

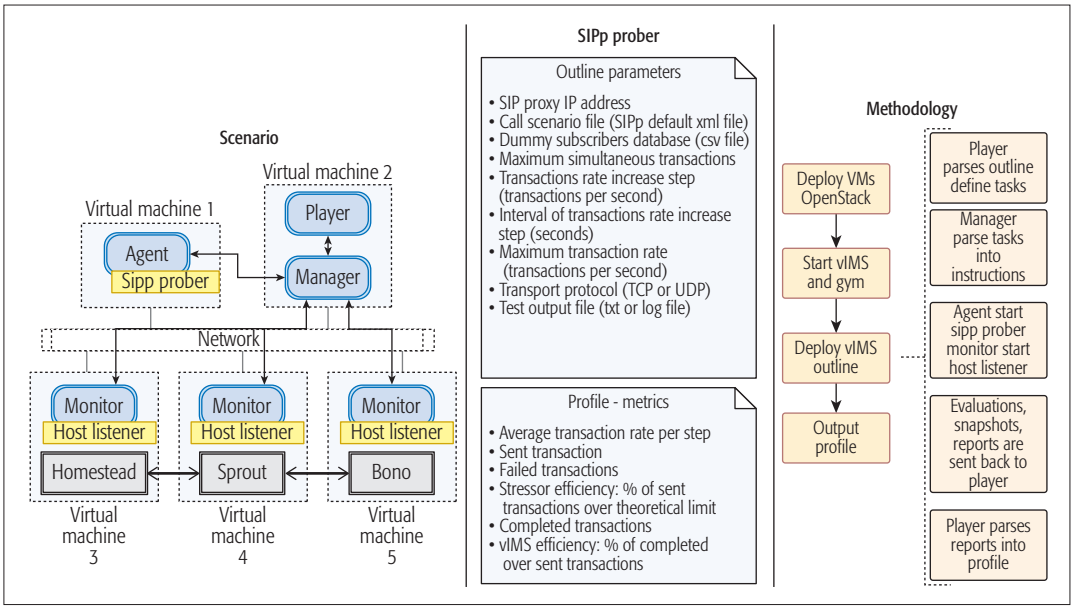

Figure 4. Testbed, SIPp prober outline parameters/profile metrics, and experimental methodology.

VM flavor	Transactions			Efficiency	
	Sent	Failed	Ack	SIPp prober	vIMS
m1.small	5433	206	4227	48%	77%
m1.medium	10924	49	10187	99%	92%
m1.large	10337	232	9170	93%	88%

Table 1. Efficiency of vIMS and SIPp prober per OpenStack Flavor.

ager (VIM) to deploy and host the vIMS subject of the Gym benchmarking experiments.

Figure 4 shows the experimental scenario featuring a total of five VMs interconnected through a layer 2 network. Gym and vIMS main components constitute a simple topology: a Monitor instance runs (as an independent daemon) inside each of the three vIMS main network function VMs; one Agent interfaces the *SIPp prober* in another VM; and the *Manager* and Player run in the fifth VM. The three Clearwater vIMS network functions (Bono, Sprout, and Homestead) run on OpenStack compute nodes based on Linux Ubuntu Server 14.04.3 with different default flavors: m1.small (1 vCPU/2 GB RAM), m1.medium (2 vCPUs/4 GB RAM), and m1.large (4 vCPUs/8 GB RAM).

Experimental Evaluation

We follow the experimental workflow shown in the methodology description of Fig. 4. For each VM flavor, 10 20-s-long benchmarking runs were executed with the Monitor host listener capturing metrics every second. To analyze SIPp prober vs. host metrics and derive a benchmarking behavior according to an IMS stress ladder workload [10], the following SIPp Outline parameters were defined:

- Transaction rate increase step: 100 transactions/s
- Interval of transaction rate increase step: 2 s
- Maximum transaction rate: 1000 transactions/s
- Maximum simultaneous transactions: 1000
- Transport protocol: UDP

Table 1 summarizes the observed results, with the first grouped columns presenting the overall amount of transactions *Sent*, *Failed* (vIMS could not answer/complete), and *Ack* (completed) by the SIPp prober. Note the existence of *Delayed/Queued* transactions, neither Ack nor Failed, corresponding to pending events vIMS could not provide an answer to during the experiments' runtime. In terms of efficiency, Table 1 presents the amount of transactions Sent by the SIPp Prober divided by the transactions that could be completed in theory, as well as vIMS efficiency, representing the amount of completed transactions (Ack) over the amount of those Sent. The observed results point to performance issues when using the m1.small VM configuration.

The explanation behind the system limits faced by m1.small is presented by the Monitor and can be observed by the resource metrics collected by the Monitor. Figure 5 presents the overall system CPU percentage usage, mean, and 95 percent confidence intervals of each monitored vIMS component. The solid line represents the SIPp prober output in terms of average transaction rate of SIP registration attempts following the "Stressful Ladder" as per the outline parameters. Note that the SIPp tool follows a congestion-avoidance-like behavior by automatically falling back according to the fail rate.

We can observe that in all three VM flavors, the I/S-CSCF VNF (Sprout) accounts for the largest CPU consumption, as expected from the central signaling function of the IMS core network. In the case of the m1.small VM (Fig. 5a) configuration, Sprout suffers from CPU over-consumption and saturates at around 500 transactions/s, whereas m1.medium and m1.large reach around 800 to 900 transactions/s.

We may conclude that, in general, scaling up virtual resources leads to higher vIMS efficiency. However, despite running experiments in an arguably well controlled environment with low background interference, the metrics pres-

ent high variability. Curiously, m1.large does not consistently surpass m1.medium, as one might expect from a resource-richer configuration. This observation only confirms the practical challenges behind VNF benchmarking.

Discussion

While the presented experimental work features a limited combination of Gym, OpenStack, and vIMS using a simple VM testbed, there are sufficient partial results for a rich discussion and a critical analysis on different aspects.

VNF Testing Challenges: As exemplified by our vIMS benchmarking efforts, designing and implementing a generic VNF testing framework is subject to multiple challenges requiring further investigation:

- Consistency: Naturally, our first insight goes to question if a VNF, when deployed in a certain execution environment, delivers a given performance described in its extracted profile, especially, when tested and put in production using multiple virtualization technologies and concurrent system workloads.
- Stability: VNF performance measurements need to present consistent results over different scenarios. Consequently, we would like to answer if test descriptors transparently handle service/resource definitions and metrics of VNFs placed in heterogeneous environments.
- Goodness: A VNF might be tested with different allocated resources and stimuli, unlike the possibilities of production environments. Crucially, we would like to comprehend how well testing results, and associated stimuli, correspond to VNF measured performance when running in execution environments under real workloads.

Resource Optimization and SLAs: Greatly facilitated by an automated approach, our experimental evaluation unveils some patterns that could be useful for optimized configurations by sizing the VMs according to their resource demands (e.g., Sprout ⇒ m1.large, Bono ⇒ m1.small). Further optimization and resource allocation strategies (e.g., CPU pinning) should also be investigated to derive more complete scalability recommendations in addition to wider experimentation with realistic workloads, altogether yielding more reliable vIMS performance profiles. We look for a better understanding on how much VNF testing profiles could be part of network SLAs; for instance, complementing continuous monitoring, or an optimized process in terms of performance and cost reduction.

Comparability, Repeatability, and Interoperability of All: Results are extracted based on the vIMS profile by the Gym Player. The obtained metrics can be exported in different file formats or committed to a database for comparison purposes. The same outline can be used by Gym to run performance tests in other virtual environments and extract the same type of metrics. This allows Gym users to continuously execute the same pattern of benchmarking in their own deployments with the ability to customize and debug their tests. Gym can be deployed in heterogeneous environments as the main requirements sit on plain Linux and Python support. All components, together with the developed SIPp prober, can be reused to perform benchmarking tests in case of alternative virtualization (e.g., containers) or bare metal deployments.

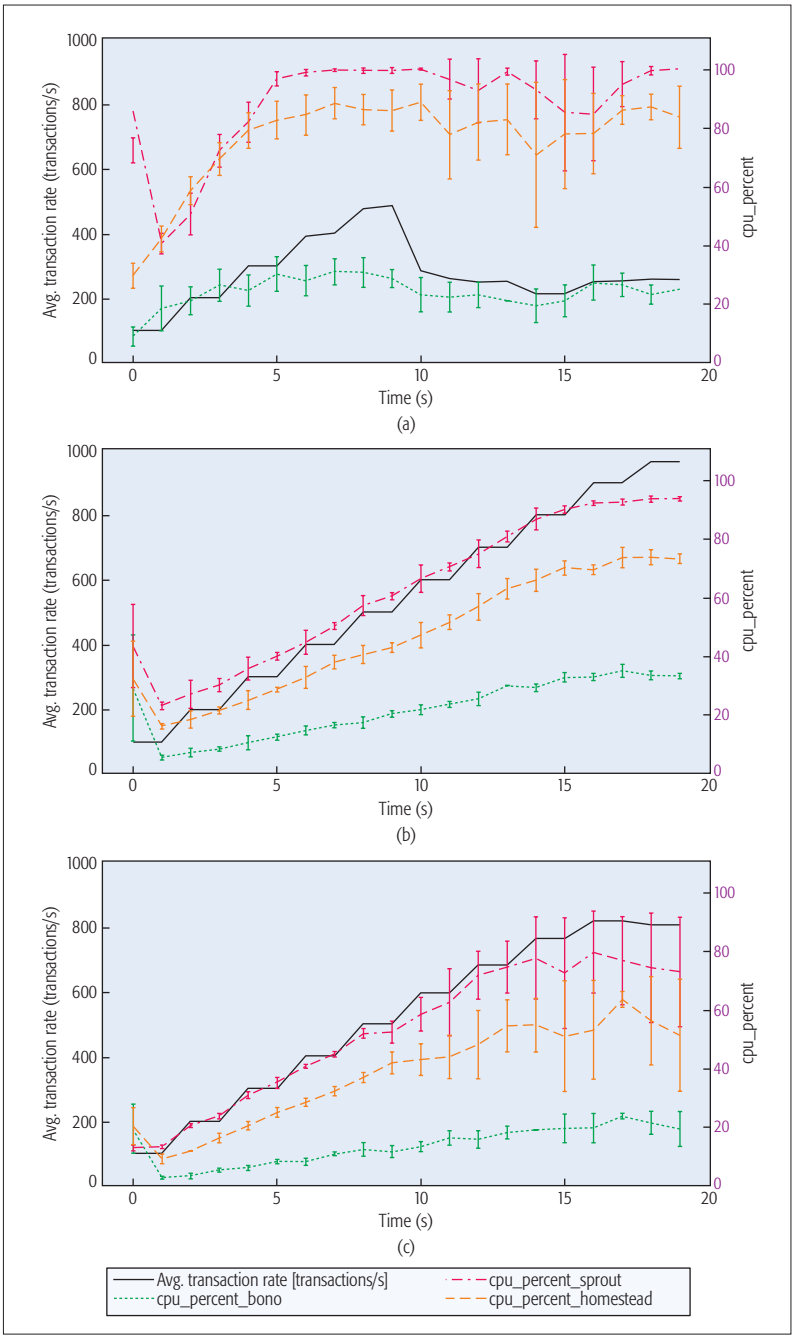

Figure 5. Average transaction rate (transactions per second) vs. CPU usage (percent): a) VNFs over m1.small — raw data at https://plot.ly/~bertoldo/848; b) VNFs over m1.medium — raw data at https://plot.ly/~bertoldo/830; c) VNFs over m1.large — raw data at https://plot.ly/~bertoldo/856.

Customization and Configurability: Gym provides a skeleton of components well suited for customized development of arbitrary VNF testing methodologies. Using sketches and outline to benchmark the vIMS offered unfettered choices for customized methodologies based on specific topologies, workloads, metric extractions, and so on. The composition of an outline is a recipe that, when interpreted by the Player component,

We envision user-contributed extensions in Gym to support different testing tools, to evaluate newborn VNFs, and to allow users to build and replicate tests by reporting profiles and maintaining common repositories for reproducible research practices involving VNF testing and analytics.

guides the architectural and functional definitions of VNF tests leveraging Agent/Monitor features. The presented use case exemplifies Gym extensibility by showing a SIPp prober easily integrated to drive the vIMS benchmarks.

NFV Orchestration: In line with our initial VBaaS vision [3], Gym was developed agnostic to any particular NFVO. We envision life cycle management interfaces in Gym to provide workflows for flexible VNF testing. NFVO would consume APIs exposed by Gym to extract desired VNF profiles and explore them in decision making processes of VNF allocation in terms of target host and resource allocation.

Related Work

In line with our initial theoretical vision on VNF benchmarking, [3, 7] propose a structured approach to develop benchmarking methodologies tailored to VNF. The use of performance profiles in support of NFV DevOps workflows was recently proposed [8] to support management and orchestration decisions leveraging offline profiling of complex service chains.

While developing Gym, we sought alignment [4, 5] with ongoing work in the Internet Engineering/Research Task Forces (IETF/IRTF), where, under the umbrella of "Considerations for Benchmarking Virtual Network Functions and Their Infrastructure" [12], relevant guidelines are being discussed toward standardized VNF benchmarking. Likewise, Gym was influenced by related efforts at the ETSI ISG NFV Testing Group [6] defining requirements and recommendations for VNFs and NFVI validation. Gym shares similarities to NFV-VITAL [9] with regard to the overall problem statement and framework approach as well as our so-called vIMS efficiency metric, which could be used in auto-scaling strategies after detecting saturation in vIMS transactions per unit of time.

A number of open source projects sprint common abstractions for benchmarking VNFs and the underlying infrastructures. Most closely related to Gym, OPNFV incubated projects include:

- Yardstick (https://wiki.opnfv.org/display/yardstick, accessed Jan. 1, 2017), targeting infrastructure compliance when running VNF applications
- QTIP (https://wiki.opnfv.org/display/qtip/Platform+Performance+Benchmarking, accessed Jan. 6, 2017), providing definitions toward platform performance benchmarking
- Bottlenecks (https://wiki.opnfv.org/display/bottlenecks, accessed Jan. 1, 2017), proposing a framework to execute automatic methods of benchmarks to validate VNFs deployment during staging

Compared to Gym, these efforts are very much tied to their choice of technologies, compromising portability and repeatability due to the focus on supporting OPNFV developments without broader aspirations of generic VNF testing tools.

The extensible and modular approach of Gym through outlines and profiles allows embracing such standalone projects by integrating them as new probers and listeners along user-defined metrics and testing workflows. One such candidate open source tool we intend to support in the Gym framework through Agent extensions is Network Function Performance Analyzer (NFPA) [13], which was also born to address the frustrating landscape of benchmarking comparison of network functions over varying software/hardware systems. Last but not least, an inspiring independent related open source effort is ToDD (https://github.com/toddproject/todd, accessed Jan. 6, 2017), which walks in the direction of an on-demand extensible framework for distributed testing of network capacity and connectivity but without focus on NFV or complex workflows.

Conclusions and Future Work

The software nature of VNFs and the multi-dimensional and time-varying aspects of heterogeneous virtualized environments call for adequate methods to assess the infrastructure capabilities with regard to target performance levels. As an evolution of our initial VNF benchmarking vision [3], this article introduces the Gym testing framework along with the principles behind our open source implementation.

The vIMS test deployment serves as a practical validation of the current Gym prototype, illustrating both its potential and the wider open challenges of automated performance benchmarking in NFV. In spite of the identified limitations, we conclude that Gym offers a meaningful apparatus to express VNF testing abstractions that can certainly be explored in continuous development and integration methodologies. On a recent proof-of-concept evaluation of Gym [14], we used Open vSwitch (OVS) as the VNF under test.

Multiple directions overtake our future work, many of them driven by 5G realization efforts leveraging NFV and software defined networking (SDN) [15]. Examples include benchmarking tests using OpenAirInterface (http://www.openairinterface.org/, accessed Jan. 6, 2017) components and NFPA [13] as traffic generator and metric collector. As an architectural framework and an open source project, Gym is still very much in its infancy. We expect Gym to keep evolving, not only in terms of low-level debugging but broadly driven by the community. We envision user-contributed extensions in Gym to support different testing tools, to evaluate newborn VNFs, and to allow users to build and replicate tests by reporting profiles and maintaining common repositories for reproducible research practices involving VNF testing and analytics. Furthermore, we foresee multiple research opportunities when applying a Gym-like approach to NFVs at runtime in support of resource orchestration and business-oriented decisions.

Acknowledgments

This research was partially supported by FAPESP grant #14/18482-4 and by the Innovation Center, Ericsson S.A., Brazil, grant UNI.58.

References

[1] P. Veitch, M. J. McGrath, and V. Bayon, "An Instrumentation and Analytics Framework for Optimal and Robust NFV Deployment," *IEEE Commun. Mag.*, vol. 53, no. 2, Feb. 2015, pp. 126–33.
[2] R. Mijumbi *et al.*, "Network Function Virtualization: State-of-the-Art and Research Challenges," *IEEE Commun. Surveys & Tutorials*, vol. 18, no. 1, 1st qtr. 2016, pp. 236–62.
[3] R. V. Rosa, C. E. Rothenberg, and R. Szabo, "VBaaS: VNF Benchmarkas- a-Service," *Proc. 2015 4th Euro. Wksp. Software Defined Networks*, Sept. 2015, pp. 79–84.
[4] —, 2015, VNF Benchmark-as-a-Service, Internet draft; https://www.ietf.org/archive/id/draftrorosz- nfvrg-vbaas-00.txt, accessed Jan. 6, 2017.

[5] —, 2016, VNF Benchmarking Methodology, Internet draft; https://tools.ietf.org/id/draft-rosabmwg- vnfbench-00.html, accessed Jan. 6, 2017.
[6] ETSI GS NFV-TST, "ETSI GS NFV-TST 002 V1.1.1 — Report on NFV Interoperability Testing Methodology," Oct. 2016; http://www.etsi.org/deliver/etsi_gs/NFVTST/001_099/002/01.01.01_60/gs_NFV-TST002v010101p.pdf, accessed Jan. 6, 2017.
[7] J. Blendin *et al.*, "Towards a Structured Approach to Developing Benchmarks for Virtual Network Functions," *Proc. 2016 5th Euro. Wksp. Software Defined Networks*, Oct. 2016.
[8] M. Peuster and H. Karl, "Understand Your Chains: Towards Performance Profile-Based Network Service Management," *Proc. 5th Euro. Wksp. Software Defined Networks*, Oct. 2016.
[9] L. Cao *et al.*, "NFV-Vital: A Framework for Characterizing the Performance Virtual Network Functions," *Proc. 2015 IEEE Conf. Network Function Virtualization and Software Defined Networks*, Nov. 2015, pp. 93–99.
[10] D. Thissen, J. Miguel, and E. Carl, "Evaluating the Performance of an IMS/NGN Deployment," *Proc. 2nd Wksp. Serv. Platforms, Innov. Res. New Infrastructures Telecommun.*, 2009; http://subs.emis.de/LNI/Proceedings/Proceedings154/gi-proc- 154-224.pdf, accessed Jan. 6, 2017.
[11] ETSI, "ETSI TS 186 008-2 V2.1.1 — IMS Network Testing," 08/2013; http://www.etsi.org/deliver/etsi_ts/186000_186099/18600802/ 02.01.01_60/ts_18600802v020101p.pdf, accessed Jan. 6, 2017.
[12] A. Morton, "Considerations for Benchmarking Virtual Network Functions and Their Infrastructure," 2016, Internet draft; https://datatracker.ietf.org/doc/draftietf- bmwg-virtual-net/, accessed Jan. 6, 2017.
[13] L. Csikor *et al.*, "NFPA: Network Function Performance Analyzer," *Proc. 2015 IEEE Conf. Network Function Virtualization Software Defined Networks*, Nov. 2015, pp. 15–17.
[14] R. V. Rosa and C. Rothenberg, "Taking Open vSwitch to the Gym: An Automated Benchmarking Approach," to appear, IV Workshop pre IETF/IRTF, July 2017.
[15] D. Kreutz *et al.*, "Software-Defined Networking: A Comprehensive Survey," *Proc. IEEE*, vol. 103, no. 1, 2015, p. 63.

Biographies

Raphael Vicente Rosa (raphaelvrosa@dca.fee.unicamp.br) is currently pursuing his thesis on multi-domain distributed NFV as a Ph.D. student at the University of Campinas, Brazil. During the last two years, he worked as a visiting researcher at Ericsson Research Hungary, where he contributed to the EU-FP7 Unify project and developed activities within the H2020 5G Exchange project. His main interests are in state-of-the-art SDN and NFV research topics.

Claudio Bertoldo (bertoldo@dca.fee.unicamp.br) has an M.Sc. from the University of Campinas He has been involved with next generation fixed and mobile broadband networks since 2007, working at telecommunications companies such as Telefónica and Huawei, and also at several startups. His research interests include network functions virtualization and next generation mobile networks.

Christian Esteve Rothenberg is an assistant professor in the Faculty of Electrical & Computer Engineering at the University of Campinas, where he received his Ph.D. and currently leads the Information & Networking Technologies Research & Innovation Group. His research activities span all layers of distributed systems and network architectures, and are often carried out in collaboration with industry, resulting in multiple open source projects in SDN and NFV, among other scientific results.

WiSHFUL: Enabling Coordination Solutions for Managing Heterogeneous Wireless Networks

Peter Ruckebusch, Spilios Giannoulis, Domenico Garlisi, Pierluigi Gallo, Piotr Gawowicz, Anatolij Zubow, Mikoaj Chwalisz, Eli De Poorter, Ingrid Moerman, Ilenia Tinnirello, and Luiz DaSilva

The paradigm shift toward the Internet of Things results in an increasing number of wireless applications being deployed. Since many of these applications contend for the same physical medium, there is a clear need for beyond-state-of-the-art solutions that coordinate medium access across heterogeneous wireless networks.

ABSTRACT

The paradigm shift toward the Internet of Things results in an increasing number of wireless applications being deployed. Since many of these applications contend for the same physical medium (i.e., the unlicensed ISM bands), there is a clear need for beyond-state-of-the-art solutions that coordinate medium access across heterogeneous wireless networks. Such solutions demand fine-grained control of each device and technology, which currently requires a substantial amount of effort given that the control APIs are different on each hardware platform, technology, and operating system. In this article an open architecture is proposed that overcomes this hurdle by providing unified programming interfaces (UPIs) for monitoring and controlling heterogeneous devices and wireless networks. The UPIs enable creation and testing of advanced coordination solutions while minimizing the complexity and implementation overhead. The availability of such interfaces is also crucial for the realization of emerging software-defined networking approaches for heterogeneous wireless networks. To illustrate the use of UPIs, a showcase is presented that simultaneously changes the MAC behavior of multiple wireless technologies in order to mitigate cross-technology interference taking advantage of the enhanced monitoring and control functionality. An open source implementation of the UPIs is available for wireless researchers and developers. It currently supports multiple widely used technologies (IEEE 802.11, IEEE 802.15.4, LTE), operating systems (Linux, Windows, Contiki), and radio platforms (Atheros, Broadcom, CC2520, Xylink Zynq,), as well as advanced reconfigurable radio systems (IRIS, GNURadio, WMP, TAISC).

INTRODUCTION

The paradigm shift toward the Internet of Things (IoT) will result in an increasing number of interfering devices that operate in the unlicensed spectrum, especially given the recent interest of the 5G community in also using the same ISM bands. Coexistence will be a huge challenge, as many heterogeneous networks have to cooperate to share the same spectrum efficiently. To this end, advanced coordination techniques must be developed that allow mitigating cross-technology interference.

Currently, multiple custom tools are used to configure and monitor wireless networks, and each type of device requires a different toolset. For this reason, controlling a heterogeneous set of wireless devices is cumbersome at least and often demands considerable effort to get acquainted with the different hardware platforms and corresponding configuration tools.

The proposed control architecture offers the possibility to create and test coordination techniques while minimizing the complexity and implementation overhead, thereby fostering innovations in a challenging research domain. For this purpose, it relies on the following key enablers.

Unified Programming Interfaces (UPIs): These allow reconfiguring various features of the network stack and monitoring its state without the need to have deep knowledge of the software and hardware particularities of each platform. The UPIs enable the design of technology-independent control programs (CPs) on top of different hardware and software platforms.

Context-Aware Execution of UPIs: This enables the definition of exactly where, when, and how a UPI call must be executed. It also allows a particular configuration value on a group of nodes to be changed at a specific time in a synchronized manner.

Connector Modules: These transform each UPI call into one or more platform-specific calls, thereby hiding the complexity of the underlying tools and/or APIs.

Hierarchical Control: This enables the creation of multi-level control loops spanning multiple and possibly heterogeneous networks. Hierarchical control allows CPs to delegate control among each other and to create custom control flows.

The UPIs and the control architecture are integrated in several federated wireless experimentation facilities. They are offered as an open source tool to the research community and have already been successfully deployed both inside and outside testbed facilities. In this article, a high-level

Peter Ruckebusch, Spilios Giannoulis, Eli De Poorter, and Ingrid Moerman are with Ghent University;
Domenico Garlisi, Pierluigi Gallo, and Ilenia Tinnirello are with the University of Palermo;
Piotr Gawłowicz, Anatolij Zubow, and Mikołaj Chwalisz are with Technische Universität Berlin; Luis DaSilva is with Trinity College Dublin.

Digital Object Identifier:
10.1109/MCOM.2017.1700073

overview of the architecture is given together with the results of several experimental showcases.

The implemented showcases demonstrate that the proposed architecture simplifies control of standardized technologies, while still offering advanced control of future reconfigurable radio systems.

Representative Use Case

The difficulty of efficiently managing coexisting wireless networks increases significantly when multiple technologies are considered. As a representative use case, this article considers an example where coexistence between IEEE 802.11 WiFi and IEEE 802.15.4 time slotted channel hopping (TSCH) is managed by separating them in the frequency and time domains. As such, different frequencies and time slots must be allocated to networks that are in each other's interference range. To realize this, advanced monitoring, coordination, and configuration techniques are required. Moreover, it must be possible to exchange control messages and maintain some level of synchronization between the different devices.

Building such a system is a nontrivial task and requires the use of different domain-specific expertise: Linux and WiFi management tools on one hand, and embedded OS (Contiki, openWSN, etc.) and programming knowledge on the other hand. Moreover, to apply the same solutions to different technologies (e.g., Bluetooth) or different operating systems (Windows, Unix, TinyOS, etc.) would require re-implementing the same control logic all over again.

The proposed architecture aims to facilitate control in all aforementioned scenarios by providing the necessary building blocks. First, UPIs allow reuse of the same control logic in different setups. Second, the context-aware execution of UPIs supports building solutions that require fine-grained control. Third, the connector modules simplify the process of extending the architecture toward new technologies and platforms.

Related Work

Control Architectures

The need for fine-grained control of communication networks is becoming increasingly apparent. This is well demonstrated by the interest of the scientific community in solutions that enable software defined networking, (SDN). OpenFlow[1], for instance, is a good example of an SDN enabler because it allows researchers to control routing without knowing the internals of vendor-specific implementations. OpenFlow, however, focuses on controlling the forwarding rules between devices (switches, routers, and wireless access points) connected by means of pre-installed links (usually wired).

Recently, a number of solutions have been proposed that enable software defined wireless networks (SDWNs) such as 5G-EmPOWER [2], OpenSDWN [3], and Sensor OpenFlow [4]. The latter two focus on enabling SDWN in a single technology (i.e., IEEE 802.11 and IEEE 802.15.4 respectively). 5G-EmPOWER is broader in scope and provides programming abstractions for managing both WiFi access points and LTE eNodeBs. However, not a single architecture exists today that can facilitate true cross-layer control (from the PHY layer up to the network layer and in some cases up to the presentation layer of the OSI model) in a unified way across multiple wireless technologies. Our proposed WiSHFUL architecture aims to go further by providing abstractions for any device and wireless technology. Furthermore, to the best of our knowledge, our architecture is the first to include reconfigurability of the medium access control (MAC) and PHY layers, which strongly affect the link availability and capacity. As such, the WiSHFUL architecture addresses this gap by offering full-stack cross-layer and cross-network control of reconfigurable wireless networks.

The WiSHFUL architecture was first conceptually presented in[5, 6]. Now we focus on the novel features such as context-aware execution and hierarchical control that allowed us to implement and evaluate the experimental showcases, illustrating how to build cross-technology coordination solutions.

Federation of Experimentation Facilities

Since most SDN solutions have been evaluated in wireless testbeds, the federation of (wireless) testbeds [7, 8] has gained much attention over the past few years. Federated testbeds aim to accelerate experimental research by providing easy reservation of experiment time slots as well as the corresponding access to resources (radios, spectrum monitoring, mobile robots, etc.) residing in different testbeds. Despite the clear progress that has been made, executing an experiment still requires manual combination and integration of different vendor- or technology-specific tools to reconfigure and monitor the devices under test. This imposes a huge burden on the experimenters since they need deep knowledge of the tools at hand, even for setting up a novice experiment. The proposed WiSHFUL architecture builds further on top of testbed federation tools to support easy experimentation using heterogeneous systems to a user base with a diverse skill set.

Reconfigurable Radio Systems

The proposed architecture supports commonly used operating systems (Linux, Contiki) for standard wireless technologies (IEEE 802.11, IEEE 802.15.4). In addition, the architecture also supports emerging state-of-the-art standards, such as European Telecommunications Standards Institute (ETSI) reconfigurable radio systems (RRSs) [9], and novel RRSs that allow more fine-grained control over the radio than is possible with typical off-the-shelf radio chips. Currently, four advanced open RRSs are supported: Wireless MAC Processor (WMP) for IEEE 802.11 radios [10], Time-Annotated Instruction Set Computer (TAISC) for IEEE 802.15.4 radios [11], and GNU radio and Implementing Radio in Software (IRIS) for software defined radios (SDRs)[12].

These novel architectures allow the design of state-of-the-art techniques [13] for managing coexistence between devices. For instance, they enable the separation of medium access in the time domain, effectively allowing the enforcement of a cross-technology time-division multiple access (TDMA) scheme. However, although they are very flexible, several of these frameworks

Despite the clear progress that has been made, executing an experiment still requires manual combination and integration of different vendor or technology specific tools to reconfigure and monitor the devices under test. This imposes a huge burden on the experimenters since they need deep knowledge of the tools at hand, even for setting up a novice experiment.

The control plane extensions offered via the UPIs allow optimization of the QoS in all networks under control, not only by considering node-local and in-network optimizations but also by taking into account the cross-technology interaction between the different networks.

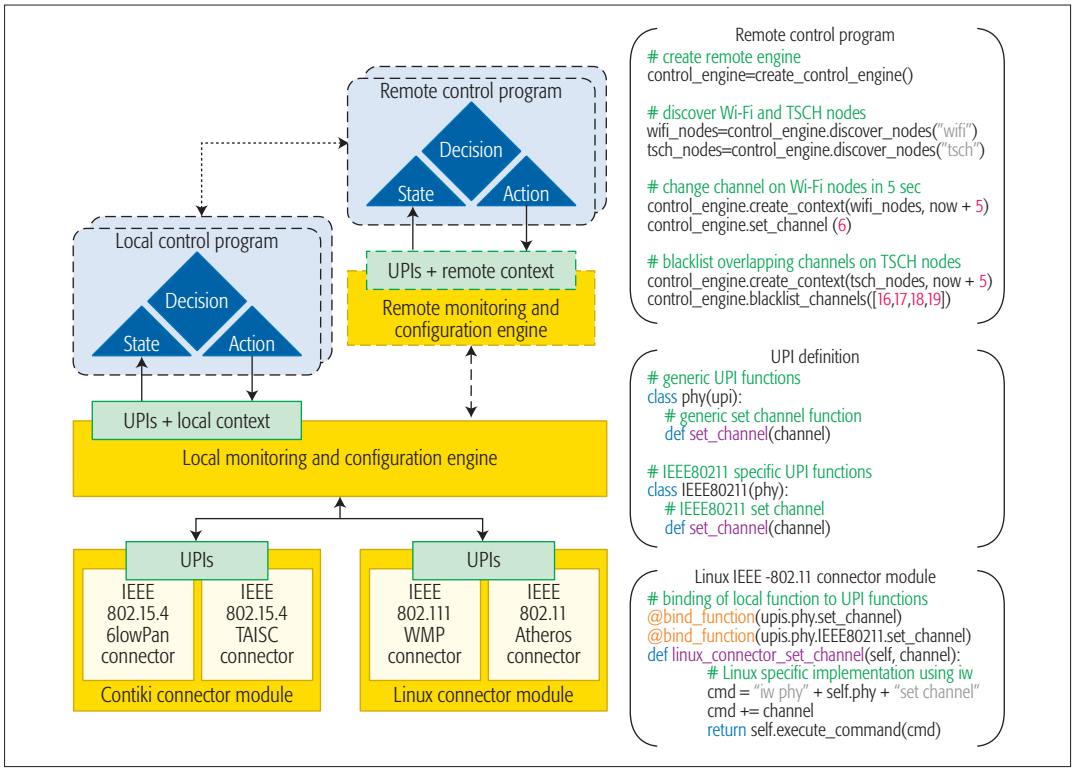

Figure 1. A high-level overview of the WiSHFUL architecture (left) and example code snippets (right). The architecture features both local and remote control, as well as context-aware execution. For each platform and technology, connector modules adapt generic UPI calls to platform-specific calls. The upper code snippet demonstrates the use of UPIs in a remote control program. The lower code snippet illustrates how generic UPI calls are mapped to platform-specific calls.

lack proper documentation and require learning yet another programming language and programming framework, thereby imposing a steep learning curve on wireless researchers and developers before they can be used. The availability of simple, cross-technology WiSHFUL UPIs remedies these shortcomings and allows integration of these advanced platforms with traditional radio platforms.

WiSHFUL Architecture and Concepts

To lower the threshold for building coexistence solutions, a novel control architecture was designed and created within the WiSHFUL project. The left side of Fig. 1 illustrates the main architectural blocks discussed in this section. The simplified code snippets on the right exemplify a remote control program (upper), UPI definition (middle), and a connector module (lower).

Control Programs

The control programs (CPs, top of the figure) execute the user-defined control logic. They build up a view on the network state by collecting monitoring information that can be used to drive decisions leading to configuration actions. For this purpose they use a set of UPIs in a particular execution context.

The control programs can be used locally, on the node and/or remotely within a subnet of nodes or across different networks. Control programs can be simple rule-based scripts, but can also comprise more intelligent components, allowing a fully self-organizing network to be built.

By allowing interactions between control programs (dotted arrows), it is possible to implement a hierarchical control logic where local CPs execute time-sensitive control loops, while remote CPs gather information from and make decisions on a group of nodes.

The upper code snippet demonstrates how a remote control program uses the UPIs to configure the WiFi network on a particular IEEE 802.11 channel and blacklist the overlapping IEEE 802.15.4 channels in the TSCH network. The example also illustrates how an execution context can be attached to a UPI function.

Unified Programming Interfaces

The UPIs (green blocks) provide generic hooks, which enables controlling the behavior of the network stack on a heterogeneous set of nodes by exposing common functions to monitor and configure networked devices in any layer of the protocol stack (i.e., from PHY to application). Both request (pull) and event-based (push) UPIs, are provided for monitoring the state and performance of the network.

There is a two-tier unification for protocol control interfaces:

- Unification across different platforms and implementations (e.g., the same IEEE 802.11 parameters provided in an identical way for Windows and Linux platforms)
- A unification across technologies and protocols with similar behavior, such as carrier sense multiple access (CSMA) parameters for both IEEE 802.11 and IEEE 802.15.4

The UPIs also include meta-information that allows reasoning on logical connections between

different implementations (e.g., set_channel in IEEE 802.11 and IEEE 802.15.4). The example snippet in the middle illustrates the two-tier unification of UPIs for the set_channel function.

The UPIs focus on common control functions that are found in most typical radio platforms and networking standards. For control features that are not yet supported across multiple technologies, we offer the possibility to support them as technology-/platform-specific APIs in an intuitive manner.

Monitoring and Configuration Engine

The monitoring and configuration engines (MCEs, dark yellow blocks) implement the core WiSHFUL services required for controlling one or more wireless nodes. Since the nodes have diverse capabilities and can reside in different networks, providing such services is a nontrivial task. The MCEs provide the following core WiSHFUL services.

Remote Execution: UPIs can be executed both locally and remotely on one or more nodes using remote procedure calls.

Context-Aware Execution: It is possible to specify exactly how (blocking or non-blocking), where (one or more nodes in the same or different networks), and when (exact time or relative delay) UPI functions are executed.

User-Defined Control Flows: The architecture allows establishing a dedicated control channel between CPs, thereby enabling custom interactions. In addition, control logic can be injected on the fly, allowing delegation of control between CPs.

Support services: These include node discovery and time synchronization, which work across different networks and on platforms with different capabilities.

More details concerning the discussed services can be found in [14].

Connector Modules

The connector modules (light yellow blocks) transform the generic UPI calls to platform-specific calls. They are implemented on each platform and for each technology. In most cases they are a simple wrapper around existing configuration tools such as netlink and iw. In other cases custom extensions are required to enable the functionality of UPIs.

The connector modules are dynamically loaded by the MCE based on the platforms and technologies used in the setup. This implies that the set of active UPIs changes over time and can be tailored toward the specific needs of a solution.

The example in the lower code snippet illustrates how the Linux iw command is wrapped in the platform-specific set_channel function. This function is then bound to both the generic and IEEE 802.11 UPI function set_channel.

UPI Enabled Control Plane in Wireless Experimentation Facilities

The control plane extensions offered via the UPIs allow optimization of the quality of service (QoS) in all networks under control, not only by considering node-local and in-network optimizations but also by taking into account the cross-technology interaction (e.g., interference) between the different networks.

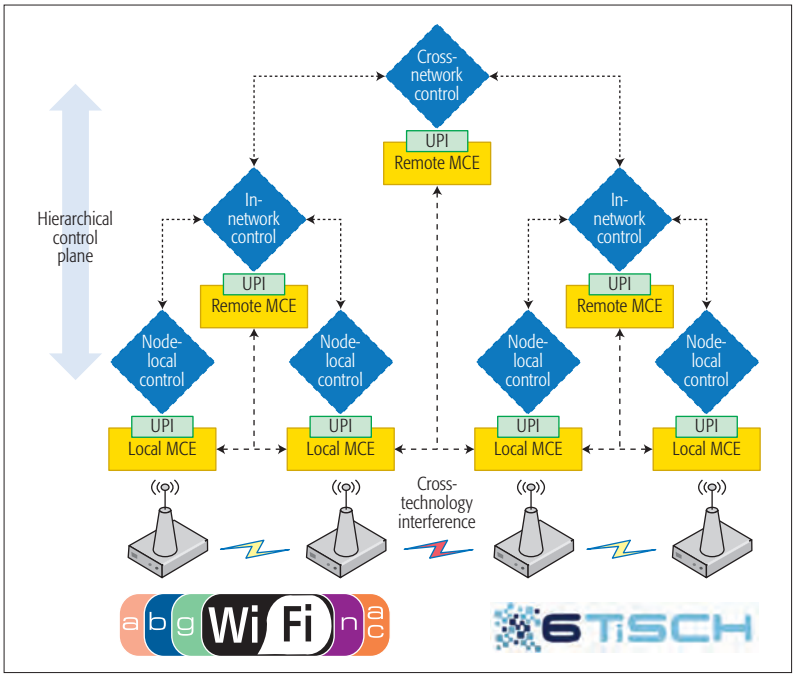

Figure 2. Illustrates the possibility to build a hierarchical control plane using the WiSHFUL architecture. Two types of control flows are enabled: 1) UPI based, between control programs and UPIs; or 2) User defined, between control programs.

Figure 2 demonstrates how a hierarchical control plane can be built using the WiSHFUL architecture. The control programs (blue shapes) can be executed on different logical levels, allowing the placement of delay-sensitive operations close to the hardware while maintaining a broader network-wide or cross-network view on a higher level. The figure depicts three logical levels of control: node-local, in-network, and cross-network. Each level can directly use the UPIs (dashed arrows) or delegate control to another level (dotted arrows). For instance, a cross-network control program can directly monitor single devices or delegate monitoring processes to the local level and work on aggregated values to reduce the amount of data to be transferred over the network.

UPI Control Channels

Two types of control channels can be employed to enable monitoring and configuring nodes across different networks. Besides the default UPI control channel, that is, between a (local or remote) control program invoking UPIs, and the node through the MCEs, it is also possible to set up communication channels between control programs of different levels (node-local, in-network, and cross-network). These communication channels can be used to share information and delegate control functionality between different control programs.

This enhances the flexibility in creating the control programs because researchers can, for instance, choose to aggregate monitoring information on the node-local level and only forward information in a custom format. It is also possible to execute certain configuration tasks node-locally on the fly triggered by a central control program.

UPI Multi-Level Control Loops

The ultimate goal of the UPIs is to enable the creation of multi-level control loops that can span between different networks. In each level, a control program uses UPIs to monitor the network performance and state. Based on this information, the CPs can decide to change the network behavior by executing configuration commands employing UPIs. The types of control loops made possible by the proposed architecture are presented below:

Technology	Operating system	Hardware platform	Hardware driver
IEEE 802.11	Linux, Windows	Atheros, Broadcom	Ath9k, NDIS driver, WMP
IEEE 802.15.4	Contiki, TinyOS	MSP430, ARM-Cortex-M	Contiki TinyOS, TAISC
LTE	Linux	Linux Server Femtocell	SIRRAN EPC LTE 245F
SDR	Linux, Windows	USRP, Xilinx ZedBoard	Iris, LabView, GNU radio

Table 1. Main overview of supported technologies, operating systems, hardware platforms, and drivers.

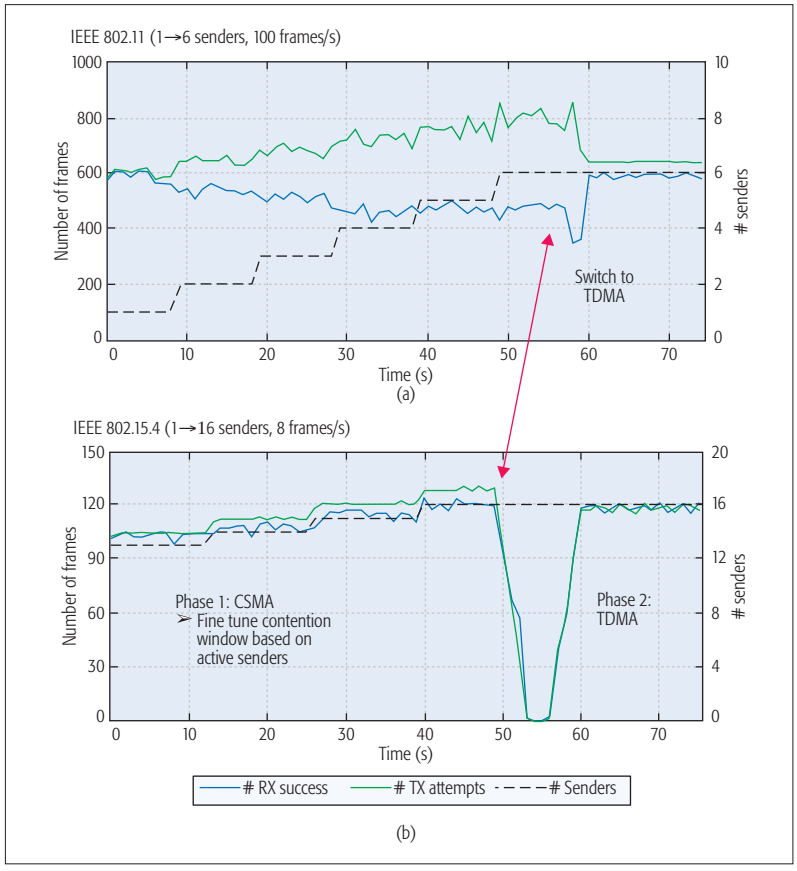

Figure 3. The graphs show the number of received frames (blue) vs. the number transmitted frames (green) for an increasing number of senders (black dashed). This experiment was conducted on both IEEE 802.11 nodes (upper chart) and IEEE 802.15.4 nodes (lower chart).

Node-Local Control Loop: The first level provides the possibility to create a node-local control loop where local decisions are made based on information observed locally via the UPIs or received from other control programs via a user-defined control channel. The node-local reconfiguration always uses the UPIs directly. This local approach is efficient to implement quick reactions to the rapidly changing context. The delay of a local UPI call is usually on the order of microseconds, depending on the complexity and CPU speed.

In-Network Control Loop: The second level enables control of all nodes in a logical network (i.e., the nodes are in the same "subnet" and use the same technology). Now network-wide monitoring drives decisions, and configuration settings are changed on a single node or a group of nodes in the network. The information can be retrieved using UPIs remotely or from the node-local CPs. Similarly, network reconfiguration commands can be done remotely, using the UPIs, or via control delegation. The delay of a UPI call inside a network is typically on the order of milliseconds, depending on the network latency and bandwidth.

Cross-Network Control Loop: In many cases, control is required across network and technologies (e.g., interference avoidance between different technologies in the industrial, scientific, and medical — ISM — band). For this purpose, the architecture allows the creation of a cross-network control loop that regulates the medium access between different networks. The interactions are similar to the in-network control loop except that they can now span multiple networks. The typical delay of a UPI call across different networks is on the order of 100 ms and is mainly influenced by the latency of the backbone network.

Supported Experimentation Facilities

The WiSHFUL architecture is currently fully supported in the imec iLab.t, TU Berlin TWIST, the Rutgers University ORBIT lab, and the TCD Iris wireless experimentation facilities. Table 1 lists the communication technologies, operating systems (OSs), hardware platforms, and drivers controlled using the UPIs. With minimal effort, UPI support can be given to experiment facilities that use (a subset of) the technologies listed below. Support for other technologies such as Bluetooth, LoRa, and SigFox is planned in the near future.

In terms of memory overhead, the full WiSHFUL framework requires only 0.75 percent of the 512 kB ROM and 3 percent of the 32 kB RAM on the employed embedded Zolertia Remote Cortex-M3 devices, making it feasible to support WiSHFUL even on constrained devices.

In-Band vs. Out-of-Band Control Channels

To support solutions beyond experimentation, the control channels can be set up both out-of-band and in-band. The in-band control channel shares the (wireless) communication channels of the devices with the data flows, while the out-of-band control channel uses the backbone network provided by the experimental facilities for transferring control flows. Using the latter approach, it is possible to separate the control flows physically from the data flows, thereby allowing the evalu-

ation of control strategies without impacting the applications.

In real-life deployments (when no testbed backbone is available), however, only in-band control channels can be employed, introducing overhead and impacting the performance of the network. The WiSHFUL architecture supports in-band control channels and allows the impact of the control flow overhead to be evaluated.

Experimental Showcases

In this section, the strengths of the WiSHFUL architecture are demonstrated by listing results that were obtained when conducting several advanced wireless experiments. Without the presented architecture, a deep knowledge of the particular details of each platform and related tools would have been required. Thanks to the WiSHFUL architecture, each showcase only required creating a generic control program, which could then be used repeatedly during experimental validation and evaluation.

The showcases are grouped and discussed by topic. The results shown in this section were obtained on the imec w.iLab.t testbed using 32 RM-090 (MSP430-CPU-based) sensors equipped with a CC2520 IEEE 802.15.4 radio, running Contiki/TAISC and 8 embedded Linux devices equipped with a Broadcom IEEE 802.11b/g card running WMP.

Load- and Topology-Aware MAC Adaptations

This showcase illustrates how the UPIs can be used to apply the same MAC adaptations on two different platforms and technologies, investigating their applicability in a heterogeneous setup and evaluating the differences between technologies. It is important to note that in both cases, the same control programs were used.

Figure 3 compares the overall network throughput (blue line is RX throughput, green line is TX attempts, black dashed line is number of senders) for both technologies in two phases. Initially, a CSMA with collision avoidance (CSMA/CA) protocol with a contention window optimization algorithm is applied, and in a second phase, a TDMA protocol is activated. In this experiment, the active traffic flows were increased gradually by activating the senders one by one up to a pre-defined maximum, after which TDMA is activated.

The applied algorithm adapts the CSMA/CA contention window based on the number of active traffic flows in the network. It can be expected that after a while, applying this technique does not yield a higher RX throughput, and packet loss starts to increase due to collisions. At this point, it is more efficient to switch to a TDMA protocol. The exact tipping point depends on many factors such as number of senders and application data rate. Figure 3 shows a snapshot of such a tipping point during an experiment.

Coexistence of Heterogeneous Technologies

This showcase demonstrates that the WiSHFUL architecture can be used to implement advanced strategies to solve the use case presented earlier, that is, coexistence between IEEE 802.11 WiFi and IEEE 802.15.4 TSCH. This showcase exploits the hierarchical control features as well as the built-in

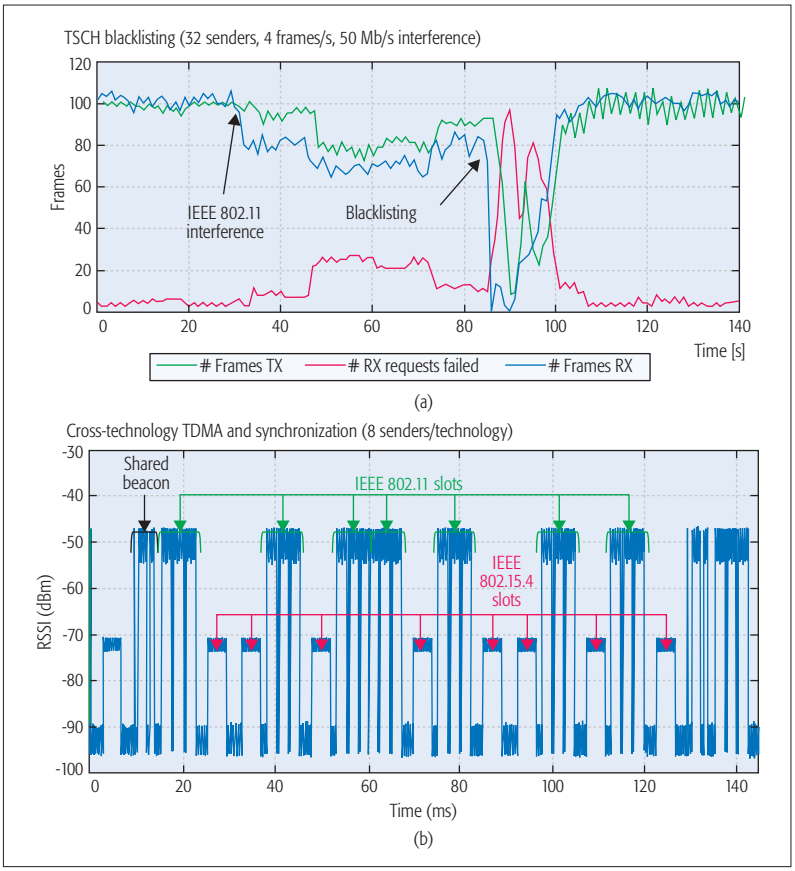

Figure 4. Results from two experiments that evaluate coexistence strategies. In the first experiment (top), the channel blacklisting features of the TSCH MAC is used to avoid channels with high IEEE 802.11 interference. The second experiment (lower part) illustrates a solution where a TDMA schedule and synchronization are shared across heterogeneous technologies.

synchronization support. Moreover, it also illustrates how the architecture supports both standardized platforms and technologies, as well as state-of-the-art frameworks.

Two different approaches were evaluated. The first solution uses the standard channel blacklisting feature in IEEE 802.15.4e TSCH to avoid channels used by the IEEE 802.11 WiFi network. The second solution uses a state-of-the-art implementation where a time-slotted MAC (TDMA) is applied in both networks based on a shared synchronization beacon and TDMA schedule.

The upper part of Fig. 4 shows the overall network throughput in the blacklisting scenario (blue line is RX throughput, green line is TX attempts, red line is TX request fails). The results clearly show that the throughput of the IEEE 802.15.4 nodes drop when there is IEEE 802.11 interference. This is mainly due to synchronization loss caused by interfered beacons. After the blacklisting of interfered IEEE 802.15.4 channels, the throughput stabilizes again to the level before adding IEEE 802.11 interference.

The lower part of Figure 4 shows an energy plot obtained by a USRP device operating in energy detection mode while testing the second solution. The results clearly demonstrate that an IEEE 802.15.4 network can be synchronized using a cross-technology beacon sent by a TDMA MAC implementation of an IEEE 802.11 network. The IEEE 802.15.4 nodes use energy detection to

Foremost, the WiSHFUL architecture offers a unified set of programming interfaces on top of a heterogeneous set of technologies, platforms, and protocol stacks, thereby drastically reducing the time and complexity typically required to build innovative solutions.

search for a particular beacon pattern transmitted by the IEEE 802.11 access point. The WiSHFUL architecture allows distribution of both the beacon pattern and cross-technology TDMA scheme among both IEEE 802.11 and IEEE 802.15.4 nodes, enabling separation of both networks in the time domain. A more detailed discussion of this particular experiment can be found in [13].

Conclusions

In the context of 5G, coexistence is a huge challenge, as many heterogeneous networks will have to cooperate to share the same spectrum efficiently. To this end, solutions are required that allow detailed network insights, fine-grained network control and management, and so on. The WiSHFUL framework offers the possibility to create and test such solutions while minimizing the complexity and implementation overhead, thereby fostering innovations in a challenging research domain.

Foremost, the WiSHFUL architecture offers a unified set of programming interfaces on top of a heterogeneous set of technologies, platforms, and protocol stacks, thereby drastically reducing the time and complexity typically required to build innovative solutions.

Furthermore, the architecture offers the possibility to execute control logic on different hierarchical levels (i.e., node-local, in-network, and cross-network) in a context-aware manner. This enables defining exactly where, when, and how UPIs are used. The presented cross-technology TDMA scheme fully exploits these features in order to synchronize and coordinate medium access between IEEE 802.11 and IEEE 802.15.4 nodes while retaining the ability to reschedule the slot allocation within the TDMA superframe at runtime.

The design of the architecture also incorporates the possibility of extensions toward new platforms and technologies. This requires only the creation of connector modules implementing and/or extending the UPIs for the particular platforms or technologies. For instance, adding support for controlling LTE networks was not a huge effort, allowing future investigation of 5G challenges such as coexistence between LTE and other technologies in the ISM band. Finally, all solutions are publicly available as open source implementations on https://github.com/wishful-project .

Acknowledgment

This work was supported by the European Commission Horizon 2020 Programme under grant agreement n645274 (WiSHFUL) and the FWO SBO "SAMURAI: Software Architecture and Modules for Unified RAdIo control'" project.

References

[1] N. McKeown *et al.*, "OpenFlow: Enabling Innovation in Campus Networks," *SIGCOMM Comp. Commun.*, vol. 38, no. 2, Mar. 2008, pp. 69–74.
[2] R. Riggio *et al.*, "Programming Abstractions for Software-Defined Wireless Networks," *IEEE Trans. Network Service Mgmt.*, vol. 12, no. 2, June 2015, pp. 146–62.
[3] J. Schulz-Zander *et al.*, "OpenSDWN: Programmatic Control over Home and Enterprise WiFi," *Proc. 1st ACM SIGCOMM SOSR*, Santa Clara, CA, June 17–18, 2015.
[4] T. Luo, H. P. Tan and T. Q. S. Quek, "Sensor OpenFlow: Enabling Software-Defined Wireless Sensor Networks," *IEEE Commun. Lett.*, vol. 16, no. 11, Nov. 2012, pp. 1896–99.
[5] C. Fortuna *et al.*, "Wireless Software and Hardware Platforms for Flexible and Unified Radio and Network ControL," *Proc. 2nd EuCNC*, Paris, France, June 29–July 2, 2015, pp. 712–17.
[6] N. Kaminski *et al.*, "Unified Radio and Network Control across Heterogeneous Hardware Platforms," Proc. ETSI Wksp. Future Radio Technologies: Air Interfaces, Sophia Antipolis, France, Jan. 27–28, 2016, pp. 1–10.
[7] M. Berman *et al.*, "GENI: A Federated Testbed for Innovative Network Experiments," *Computer Networks*, vol. 61, 2014, pp. 5–23.
[8] W. Vandenberghe *et al.*, "Architecture for the Heterogeneous Federation of Future Internet Experimentation Facilities," *Proc. Future Network & Mobile Summit*, Lisbon, Portugal, July 2013, pp. 1–11.
[9] Y. Jin *et al.*, "ETSI Reconfigurable Radio System: Standard Architecture and Radio Application," *Proc. 7th ICTC*, Jeju Island, Korea, Oct. 19–21, 2016, pp. 1094–97.
[10] Tinnirello *et al.*, "Wireless MAC Processors: Programming MAC Protocols on Commodity Hardware," *Proc. 31st IEEE INFOCOM*, Orlando, FL, Mar. 25–30, 2012, pp. 1269–77.
[11] B. Jooris *et al.*, "TAISC: A Cross-Platform MAC Protocol Compiler and Execution Engine," *Computer Networks*, vol. 107, 2016, pp. 315–26.
[12] P. D. Sutton *et al.*, "Iris: An Architecture for Cognitive Radio Networking Testbeds," *IEEE Commun. Mag.*, vol. 48, no. 9, Sept. 2010, pp. 114–22.
[13] P. Ruckebusch *et al.*, "Cross-Technology Wireless Experimentation: Improving 802.11 and 802.15. 4e Coexistence," *Proc. 17th IEEE WoWMoM*, Coimbra, Portugal, June 21–24, 2016, pp. 1–3.
[14] P. Gawowicz *et al.*, "UniFlex: A Framework for Simplifying Wireless Network Control," *Proc. IEEE ICC*, Paris, France, May 21–25, 2017.

Biographies

Peter Ruckebusch (peter.ruckebusch@ugent.be) received his M.Sc. in computer science from Hogeschool Ghent Faculty Engineering, Belgium. Since 2011 he has been a Ph.D. student at the University of Ghent, IMEC, IDLab, in the Department of Information Technology (INTEC). He has been collaborating in several national and European projects. His research topics are situated in the low end of IoT, mainly focusing on reconfigurability and reprogrammability aspects of protocol stacks for constrained devices in IoT networks.

Spilios Giannoulis [M] (spilios.giannoulis@ugent.be) received his M.Sc in electrical and computer engineering (2001) and Ph.D. (2010) from the University of Patras. Since 2015 he has been a postdoctoral researcher at the University of Ghent, IMEC, IDLab, INTEC. He is involved in several EU projects. His main research interests are mobile ad hoc networks, wireless sensor networks, especially flexible and adaptive MAC and routing protocols, QoS provisioning, and cross-layer and power-aware architecture design.

Eli De Poorter (eli.depoorter@ugent.be) received his M.Sc (2006) in computer science engineering and Ph.D. (2011) from the University of Ghent. He is now a professor at INTEC, University of Ghent. He is currently also coordinating several national and international projects. His main research interests include wireless network protocols, network architectures, wireless sensor and ad hoc networks, future Internet, self-learning networks, and next-generation network architectures.

Ingrid Moerman (ingrid.moeman@ugent.be) received her M.Sc. in electrical engineering (1987) and Ph.D. (1992) from the University of Ghent, where she became a part-time professor in 2000. She is also a staff member at IDLab-UGent-IMEC, where she coordinates research activities on mobile and wireless networking. Her research interests include IoT, LPWAN, cooperative networks, cognitive radio networks and flexible hardware/software architectures for radio/network control and management. She has long experience in coordinating national and EU research funded projects.

Domenico Garlisi (domenico.garlisi@dieet.unipa.it) received his M.Sc. in telecommunication engineering (2010) and Ph.D. (2014) from the University of Palermo. Since May 2015, he has been working as a researcher for Consorzio Nazionale Interuniversitario per le Telecomunicazioni (CNIT). He is also working as a collaborator researcher and software developer for TTI-Lab (DEIM) at the University of Palermo on smart grid and smart mobility. His current research interests include performance evaluation and medium access control in wireless LANs, including mesh network.

Pierluigi Gallo (pierluigi.gallo@dieet.unipa.it) received his M.Sc. with distinction in electronic engineering (2002) and Ph.D (2006) from the University of Palermo, where he has been an assistant professor since 2010. His work and interests focus on wireless networks, particularly on the MAC layer and its localization applications. He has contributed to several national and European research projects including ITEA POLLENS, IST ANEMONE, IST PANLAB II, ICT FLAVIA, H2020 CREW, and WiSHFUL.

Ilenia Tinnirello (ilenia.tinnirello@dieet.unipa.it) received her M.Sc. degree in telecommunications engineering (2000) and Ph.D. (2004) from the University of Palermo, where she is currently an associate professor. Her research activities have been focused on wireless networks, and in particular on the design and prototyping of protocols and architectures for emerging reconfigurable wireless networks. She has been involved in several European research projects, among them FP7 FLAVIA, H2020 WiSHFUL, Flex5gWare, and Symbiote.

Piotr Gawowicz (gawowicz@tkn.tu-berlin.de) received his M.Sc in electronics and telecommunications from AGH University of Science and Technology, Krakow, Poland, in 2014. Currently, he is working as a researcher at TKN Group at the Technische Universität Berlin, Germany. He is the author or co-author of several technical papers. He has been involved in national and European research projects. His research interests include software defined networking, wireless networks, and simulation tools.

Anatolij Zubow (zubow@tkn.tu-berlin.de) received his M.Sc. in ccomputer science (2004) and Ph.D.(2009) from Humboldt University Berlin. He is a senior researcher at the Telecommunication Networks Group at the Technische Universität Berlin since March 2013, where he is coordinating the research activities in the areas of cognitive radio, wireless access networks, and software-defined networking. In the past he did research in the area of wireless ad hoc mesh and self-organized networks.

Mikolaj Chwalisz (chwalisz@tkn.tu-berlin.de) received his M.Sc. in electrical and computer engineering from the Warsaw University of Technology and in computer engineering from the Technische Universität Berlin in 2011, where he has been a Ph.D. student since. His research focus is on coexistence and cooperation of heterogeneous wireless networks. He is actively involved in European and German funded projects working on experimentally driven solutions and testbed orchestration.

Luiz A. DaSilva [FM] (dasilval@tcd.ie) received his M.Sc. in electrical engineering (1988) and his Ph.D. (1998) from the University of Kansas. Since 2014, he has been a professor at Trinity College Dublin. His research focuses on distributed and adaptive resource management in wireless networks, and in particular radio resource sharing and the application of game theory to wireless networks. He is leading research projects funded by the National Science Foundation, the Science Foundation Ireland, and the European Commission.

Root Cause Analysis of Network Failures Using Machine Learning and Summarization Techniques

José Manuel Navarro González, Javier Andión Jiménez, Juan Carlos Dueñas López, and Hugo A. Parada G.

The authors propose an offline method based on machine learning techniques for the automatic identification of dependencies between system events, enhanced with summarization, operations on graphs, and visualization that help network operators identify the root causes of errors. They illustrate it with examples from a corporate network.

ABSTRACT

Root cause analysis includes the methods to identify the sources of errors in a network. Most techniques rely on knowledge models of the system, which are usually built by using network operators' expertise. This presents problems related to knowledge extraction, scalability, and understandability. We propose an offline method based on machine learning techniques for the automatic identification of dependencies between system events, enhanced with summarization, operations on graphs, and visualization that help network operators identify the root causes of errors. We illustrate it with examples from a corporate network.

INTRODUCTION

Root cause analysis (RCA) is key to prevent failures and errors in networks and systems and to ensure downtime is kept to a minimum when it happens. It deals with the identification of the causes of errors, playing an integral part in network management systems along with monitoring, prediction and reparation.

As stated in [1], the main techniques for RCA include searching the issued problem into a knowledge base — usually organized into sets of rules — to obtain its cause. Thus, it is possible to find case-based reasoning, where the base is a library of errors described by their cause and past solutions; model-based systems that hold models of behavior for each device in the network; and collections of system events labeled either symptoms or problems in a correlation matrix (codebook). Whenever a problem happens, information about the system state is checked against the knowledge base to make a decision to solve the failure.

Creating the knowledge base is usually difficult, as it mainly relies on network operators' knowledge and is based on rules defined by them, extracted from their experience working on the system. Thus, it is expensive to create and maintain, and completely specific to the system for which it was created. Even so, this kind of approach was good enough for small networks and static systems. Current networks are large and extremely complex, which renders any system based primarily on human knowledge incapable of completely covering all the changing network properties and features [1–3].

This article presents an offline method aimed at creating the knowledge base automatically, based on the application of machine learning algorithms on a historic dataset of events that happened in the network, enhanced with summarization, visualization, and reasoning processes that help network operators to extract information from complex systems, delve into the reasoning information to validate their domain knowledge, and establish and confront hypotheses about the root causes of errors.

STATISTICAL METHODS FOR RCA

Bayesian networks (BNs) are the base of most RCA methods used so far [2]. They model system features such as events, conditions, and metrics, as nodes in the BN, and their dependencies are expressed as conditional probability. Then probabilistic inference is used within the dependency model to obtain the most probable root causes of errors. This method scales well and can adapt to dynamic environments, but it is not without hindrances: it suffers from performance issues (although there are efforts that partially solve this [1]), do not perform well on high-dimensional data, and are extremely dependent on prior knowledge and the distributions chosen to model the data.

Just to alleviate these problems, research has been varied in scope and techniques used. For instance, in [4] the authors tackle the problem of data variety, which is one of the issues the size of current networks brings. To address this, the authors propose to aggregate information into message templates and effectively reduce the number of event categories to deal with. Then frequent patterns are found through data mining, which creates a database for a case-based reasoning-like system.

Some of the other open issues on RCA have also been recently explored through improvements or combinations of methods. A perfect example of these efforts is the work found in [1], which implements an RCA system that solves the efficiency problem by joining BNs with case-based reasoning: the network is only used with new cases. The same authors previously presented in [2] a different solution for the same problem, this time by clustering the network in different groups, which allows them to shrink the nodes to analyze, effectively reducing the inference time.

Digital Object Identifier:
10.1109/MCOM.2017.1700066

The authors are with Universidad Politécnica de Madrid.

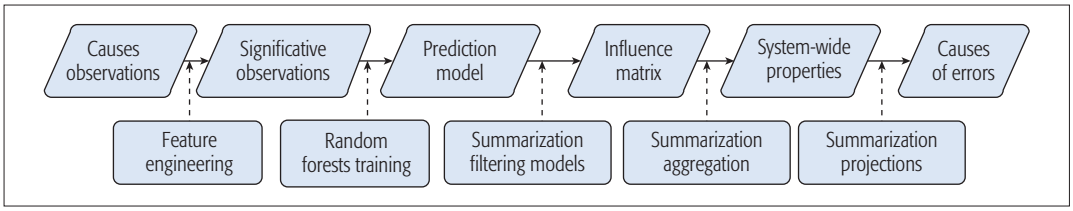

Figure 1. Proposed methodology.

All in all, the problem of creating the knowledge base in an understandable form remains; hence, machine learning algorithms (those that learn from data to find patterns and predict new data) are also having their share of the spotlight. For example, Random Forests, the algorithm we use in our proposal, has already been applied successfully in [3], where the authors feed it with two different sources of data: first, network operators label network data where they find anomalies; second, they run a series of anomaly detectors previously proposed in the literature over the same data. From these two sources, they create a labeled dataset, which is then fed into the Random Forests algorithm, which then constitutes an anomaly detector on its own. While this approach is thorough and very interesting, it assumes that the collaboration and knowledge of network operators and managers is available.

Last, one emergent technique that could become very important in the next years for understanding complex systems is summarization, understood as expressing a system in reduced, but valuable, information. One of the most used techniques for this task is clustering, a data mining technique based on grouping the data based on their features. This way, all the data can be summarized just by giving some cluster statistics. An interesting usage of summarization is shown in [5], applied on social network graphs.

Despite the efforts in autonomic management [6], the response to critical situations is set by human judgement — at least in the industrial setting — so understanding the root causes of errors and expressing them in a clear way is of paramount importance for the practical application of RCA methods.

A Method for Machine-Learning-Based RCA

We propose a method that allows for the construction of an offline diagnosis helping system based on a machine learning algorithm (in our example, Random Forests, although almost any algorithm can be used) and the analysis of its results by summarization and graph theory, which does not require deep knowledge about the modeled network and allows for the inclusion of several temporal scale ranges through windowing on the feature engineering phase.

Figure 1 sketches our proposal's workflow: the upper boxes represent the intermediate and final information obtained by the application of the operations in the lower part. It is organized as a pipeline of information, starting from system observations and finishing with causes of errors.

While the usual purpose of machine learning models is predicting new events, we intend to use them for information extraction. We do so by using a concept named "variable importance," a measure of how much a certain feature is contributing to accurately predicting the objective variable on a model. We combine all the extracted importances into the cornerstone of our method: the influence matrix, which represents in a compact way the information obtained by the application of the machine learning method. This measure of influence of an event over another one points to the concept of causality. Although it cannot be said that an observation of system behavior is the cause of another one, if the machine learning method determines a certain degree of influence, we can use it to represent how the appearance of an event contributes to the appearance of another one. Once we get this matrix of influence-causality, we can consider the value in each cell the weight of the vertex from the event in the column to the event in the row, and study it both as a matrix and as a directed graph.

Application of Random Forests for Diagnosis

A suitable dataset for a machine-learning-based RCA method must contain information about the network state in different points of its life; the larger observation period, the better statistical properties and therefore better predictive and diagnostic capabilities will be achieved. Usually, networks run under monitoring of network management systems, which are able to obtain information about the managed network, store it, and perform analysis on it. In order to abstract network state from details, these management systems often label system operations with events that are later presented to operators. This method can be applied to any event log as long it contains event traces (i.e., any data following a schema similar to "Event Type-Timestamp" would be a proper candidate for system observations in Fig. 1). As an illustration of this concept, from now on we refer to a dataset obtained from an industrial network composed of 21 devices under a unique network management system that renders the event identifier, the date when this event is registered, generating the agent or name of the device the event happened on, its type, description, and severity categorized by the monitoring tool, divided into "critical," "major," "minor," and "blank." Some of these critical events are, for instance:

- A device has stopped responding to polls and external requests.
- A chassis and all its blades have stopped responding to polls and external requests.
- A device is reporting a critical threshold violation regarding resource usage.
- A virtual machine is running fewer software instances than allocated.

In current networks, it is extremely important to be able to pinpoint not only if a failure will happen but where it will do so, as they are usually very large and complex. As our dataset contains information regarding each event's location, we define a new event type, consisting of the concatenation of the event type and the device in which it happened.

For a time range of 206 days, 21,442 events happened on the system, divided into 566 different types of events, of which 1628 events are critical (in 27 types), 1539 marked as major (65 types), 1614 minor (17 types), and many blank events (16,661 divided into 457 types). Critical and major events receive most of our analysis efforts, as they are the events that affect the system the most. Furthermore, we try to predict not only which event happened, but in which device it did so.

Feature Engineering

In current networks, it is extremely important to be able to pinpoint not only if a failure will happen but where it will do so, as they are usually very large and complex. As our dataset contains information regarding each event's location, we defined a new event type, consisting of the concatenation of the event type and the device in which it happened. Thus, the value to predict with our models is whether a certain event will happen on a certain device. Not only does this allow us to locate failures, but as a single model is created for each type of event, each model's prediction accuracy can be tested on its own. This permits us to discard models that do not perform well, as it would not be wise to trust information extracted from them. Additionally, we transformed our input data using a standard practice in the field of online failure prediction, data windowing [7]. This technique uses the following concepts.

Observation Window: a period that is observed for events when predicting an event (the events that happen in this period are the input for the predictive method). For our case, we used the number of events of each type that happened in each window.

Prediction Window: a period, after the observation window that is observed for the predicted event to happen. It creates the expected output of the predictive method. We coded it as a binary variable, indicating whether the objective event happened in the prediction period or not.

To create the final datasets that the Random Forests algorithm uses, each moment of time when an event happened is examined: an observation window is created for it, and the event in question is assigned to the prediction window. There are as many different datasets as events to predict, with equal observation windows and unique prediction windows. This process transforms a sequence of events into a proper tabular dataset, where each row contains variables used to predict (observation window) and the variable to predict (the prediction window). Based on the datasets and talks with the data providers, we settled on using 5 minutes for both observation and prediction as the great majority of events are separated by less than 5 minutes of time, with more than 80 percent of data separated by 60 seconds or less. Applying the windowing process to the dataset and for each event to predict, we obtained 501 (one per each unique event on every prediction window) tables of 10,661 rows (one per each observation-prediction window tuple) and 550 columns (one per each unique feature on every observation window). Each row of each table contained the input variables (how many times each different event had happened on the observation window) and the output variable, a Boolean value indicating whether the objective event had happened in the prediction window.

Thus, the feature engineering phase has two steps: first, a new event tag is created combining the event label and the device on which it happened, and then the data are windowed, divided in different observation tables composed of observation and prediction windows. This set of tables conforms the "significative observations" box of Fig. 1.

Random Forests Training

The training phase is now carried out using a machine learning algorithm. There are several algorithms able to yield some measure of influence between variables that we could use. Between the most usual ones (e.g., regression analysis, naive Bayes classifiers, neural networks, or even BNs), we settled on Random Forests, which are combinations of decision trees. We chose the Random Forests algorithm because it does not assume any data distribution, inherently captures interactions and performs feature selection, can work with data with a lot of variables, and is not hard to tune. It is based on recursive splits of data into subsets, trying to maximize the separation between different classes. In our case, the classes are the appearance or absence of the objective event on the prediction window.

We trained 501 Random Forests models (as many as the data tables we had created). To do so, we used the R programming language, as well as the randomForest and caret (classification and regression training) packages. We decided to use a 10-fold cross validation scheme, with the SMOTE algorithm [8] for class imbalance and the Area Under the ROC Curve [9], a very common metric in machine learning that ranges from 0 (worst performance) to 1 (perfect predictor) as a performance metric. The first thing to check is that the models were correctly created. If there were not enough original samples in the dataset of a certain event, for example, the model would not be able to be trained correctly. From the 501 generated models, our dataset rendered 351 correct and 150 incorrect models. To ensure the quality of the extracted information, we imposed a 0.8 threshold on the AUC score, which filtered out 19 models. This value could be tuned to the necessities of the problem at hand. The final number of models we used to create the system influence matrix was 332, so the resulting matrix contained 332 rows (one for each used model) and 548 columns (each one corresponding to the number of events that influence each model).

The influence matrix contains as many rows as the number of trained models and as many columns as the number of distinct events present in the observation windows. Regarding its actual content, the value in each cell contains the measure of influence of each event on the apparition of each other. For Random Forests, this variable importance of each feature is calculated as the difference in prediction accuracy between training it normally and retraining the model when the variable is replaced by random values. The rationale behind this process is that if the influence of an event is not important, randomly permuting

its values should not affect the predictor's performance. If it changes, we can measure how much, and that gives us an intuition about its overall importance in the prediction process.

Analysis of System-Wide Properties

A first visual assessment of the influence matrix can be useful in order to find general patterns. One way to visualize large matrices is to draw a heatmap of them, where each row represents an event model and contains the influence values of the others on it. Both rows and columns have the events in the same order; events in the same device appear together ordered by their severity. Vertical lines show the influence that a certain event has on all the models, and horizontal lines indicate how the event they represent is influenced by the others (Fig. 2).

From this visual analysis, we conclude that:

- There is a large part of the matrix (almost 75 percent) with insignificant values (yellow); we interpret that only a fourth of events convey information useful for RCA. Empty rows signal events for which no cause can be discovered. Empty columns are for events that do not add information. This suggests that, on average, models are influenced by a small number of events and vice versa.
- There seems to appear a principal diagonal, which suggests that the appearance of most events is influenced by the appearance of the same event at a time lower than the observation window. Perhaps these represent alarms on certain conditions of the network, which, once they are raised, generate events periodically.
- The switches seem to have a strong influence on themselves and on several virtual machines. The virtual machines work in groups. The router has around five events that strongly influence everything else, something that indicates its importance in the network.

Direct manipulation of the influence matrix allows reasoning on events, aggregated by severity, rendering a summarized table (Table 1) whose rows and columns are severity levels, and each cell contains the median proportion of events of the row's severity affected by events of the column's severity.

We draw two important conclusions from this analysis: Critical events affect, overall, the largest proportion of events, which is consistent with the fact that this kind of event is supposed to cause a large disruption in the system. Critical events only seem to be affected, mostly, by other critical events. This is a potential issue in terms of network resiliency, which would point to a need for an event policy recollection improvement.

A different, complementary analysis on the influence matrix is to consider it as a graph where nodes are types of events (type-device, in fact), and edges are influence values. This kind of analysis permits obtaining an "influence topology" of the network, even if no topology information is present on the dataset. This can either confirm that the actual topology dictates most of the devices' behavior (e.g., separated devices do not influence each other), or lead to discovering some cascading problem events or situations. While

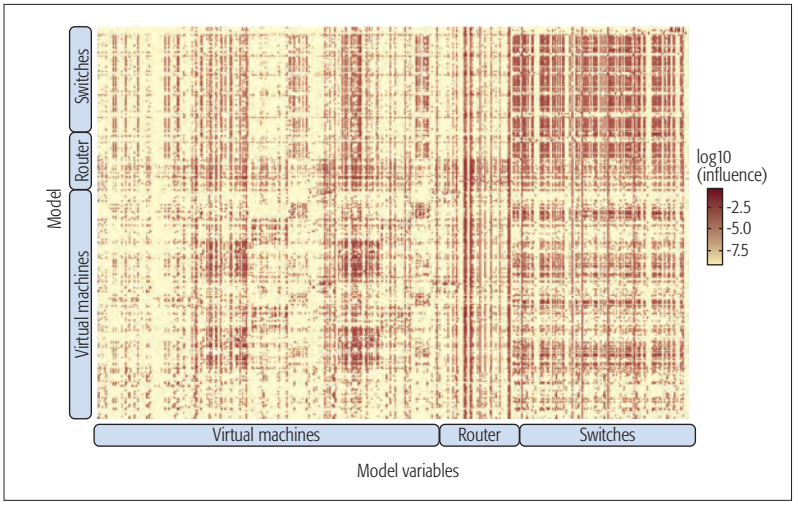

Figure 2. Influence matrix depicted as a heatmap.

	Blank	Minor	Major	Critical
Blank	0.13	0.38	0.11	0.52
Minor	0.14	0.57	0.14	0.29
Major	0.11	0.29	0.14	0.33
Critical	0.10	0.30	0.10	0.80

Table 1. Aggregated matrix for severity of events.

such a graph can be directly plotted, for most systems this would lead to an entangled mess, with no clear knowledge or relation present. Thus, we filtered it by edge magnitude (influence magnitude between events), solving the visualization problem and displaying only the strongest relations in the system. The selected edge magnitude threshold can be varied depending on its size, the effect magnitude distribution, and the desired strength of the relationships shown. We propose to use percentile analysis to set this value, as it allows the user to abstract from absolute values. Using this, we constructed two different graphs: first, we used the top 1 percent influences on the influence matrix. Analyzing it using Gephi and the algorithms it implements, we extracted some graph metrics from it that give us some insights on the system: it is composed of 381 nodes and 1614 edges, forming two disjoint graphs. Its average node degree, the average number of edges (input or output) per node, is 4.2, meaning that, on average, each event is related to around 4 different events.

Employing the visualization capabilities of Cytoscape (http://www.cytoscape.org/, accessed 15 June 2017), we plotted Fig. 3, which shows the graph, along with the different communities we found applying the algorithm proposed in [10]. Thus, for example, there is a group of events related to alerts and threshold violations, another one related with devices connection problems; maintenance operations; execution of Open Shortest Path First (OSPF) in the router; and the stopping of Java management agents (these were obtained inspecting the different communities' members). This visualization shows the two different disjoint graphs (the second one being the

small community on the lower left part) and how, generally, critical events are associated with large clusters of blank and major events, suggesting that they cause a large disruption in the system. There also seems to be an overall separation between hardware devices, where switches are mainly influenced by other switches, and virtual machines are influenced by the router or another virtual machine.

Even though this rich analysis can give a broad overview of the system, we can leverage the information contained in our data to view it from a different angle. Specifically, summarizing all the device information while excluding the events allows us to study the general device behavior on the network. As this resulting matrix is much smaller, we can visualize its top 10 percent influences as a graph (Fig. 4). It shows very clearly the hardware division we spoke about previously: switches interact mainly with each other, and the router nfluences events on virtual machines. As a novelty, this graph unveils how virtual machines are normally paired with one another, probably indicating that they are being executed on the same physical device. Also, the router is only influenced by itself and, in turn, influences both switches and virtual machines. This figure also offers a clearer view of something that the previous one and the heatmap hinted: the switches

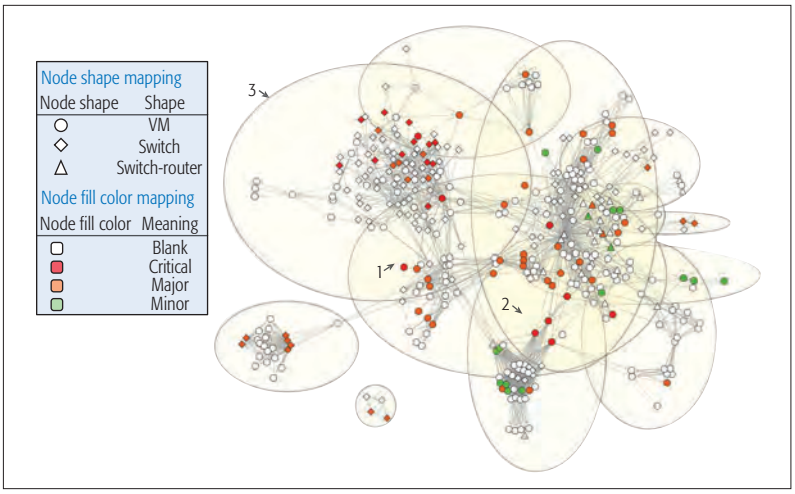

Figure 3. Influence graph of the top 1 percent influences, including communities.

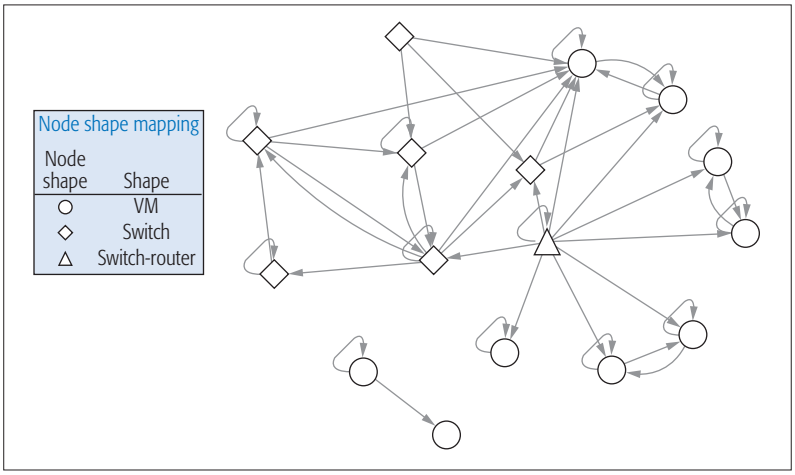

Figure 4. Summarized devices graph.

form a highly interconnected mesh of devices, suggesting that a failure on one of them may render the whole network unusable, whereas failures on virtual machines do not seem to propagate to each other except to their pair.

Identification of Causes of Errors

Once a system-wide characterization has been made, if necessary, specific events can be studied. As an example, focusing on certain critical events, there are several interesting observations we can make looking at the graph in Fig. 3 and their event descriptions:

- From the isolated critical event on the left (arrow 1), we find that losing contact with the virtual machine manager precludes a virtual machine stopping to respond to requests.
- From the group of three critical events on the lower part of the graph (arrow 2), we observe that a critical event (a software alarm) on any of the three virtual machines causes a plethora of blank, minor, and major events.
- On the top left, a large cluster of critical events (arrow 3) on switches, alarming of loss of contact with them, is surrounded by a myriad of management events and to a loss of contact on a specific virtual machine.

Overall, these observations show how, usually, critical events are highly disruptive situations with a lot of blank events associated with them. But on a system of this scale, just inspecting the graph can be cumbersome in order to find specific event causes or influences. To do so, we can perform a search on the matrix for all the associated nodes of a specific event. This gives us a smaller, clearer graph that can be leveraged to analyze an event in depth. An example of this highly focused analysis is shown in Fig. 5, which contains the strongest influences on a critical event in a specific virtual machine reporting a CPU threshold violation. The dashed lines show a chain of events of normal behavior: when a software is started or some process changes on the virtual machine, there is a sudden spike of CPU signaled by first a major event and then a critical one. On the contrary, the arrowed line shows a strange situation: a chain of a blank and a major event on the router indicating that it is not responding to Simple Network Management Protocol (SNMP) messages is influencing the critical threshold violation and not the major one. This fact tells us that most of the time the major violation happened, the router event was not an influence on it. Thus, there is an anomalous situation, different from the expected behavior of the virtual machine, that causes its CPU to spike when the router suffers this specific event. This would need further research in collaboration with the network managers.

Last, we would like to propose another highly summarized index to analyze the system: a normalized index that indicates how "harmful" an event is for the system. We use the percentage of critical and major events that are strongly influenced by each event. It is a similar approach to the severity table we studied before, but this time it is specific for each event, and we only take into account values above a certain threshold to ensure that the influence they have on events is

important enough. Our analysis shows that out of 481 possible events, 8 have a danger index higher than 0.5, which means that they influence more than half of the 55 critical and major events in the system. Four of them happen on a switch, three on the router, and only one on a virtual machine, and they all are subproducts of the network monitoring system, which suggests that there does not seem to be any event that causes widespread disruption on the network with a constant, high-level influence.

Conclusions

Root cause analysis is a key part of networks. Current solutions based on Bayesian networks suffer some limitations, such as a high dependency on previous knowledge or some limitations when representing temporal relations. As an alternative, we propose a method that allows users to decide which kind of relationships and modeling they want to use for their analyses, while offering a highly compact matrix that contains the network's behavior. We have shown several applications and summarization tools that can be used or applied to it, including several novel approaches, such as expressing the influence matrix as a directed graph or proposing an event danger index, which allows network managers to quickly assess how the network behaves. Our methodology forms a toolset of analysis that can extract useful data from networks even when little to no information is available apart from the dataset. As an example, we applied it to a real dataset obtained from a corporate network, to discover several nontrivial facts about it at a system overview level, such as the distinct influence separation based on hardware, where the switches are mainly influenced only by themselves while the virtual machines are paired. We also analyze the network at a specific level, such as observing how critical events cause major disruptions in the system, and studying the influences of one single critical event, discovering an anomalous influence caused by the router. We have thus shown how our proposal not only allows for probable cause extraction but also serves as a guide for further diagnosis research on the network. It can be applied to any network dataset that can be expressed as combinations of events and timestamps, and it can work with large networks, due to the summarizations included in the workflow. Furthermore, in large-scale networks, matrix factorization techniques such as singular values decomposition could be applied to the influence matrix in order to subdivide the problem into smaller, independent problems, allowing our proposal to work in a modular, efficient way. Regarding future work, we would like to apply our methodology to larger and more diverse datasets, such as network anomaly data.

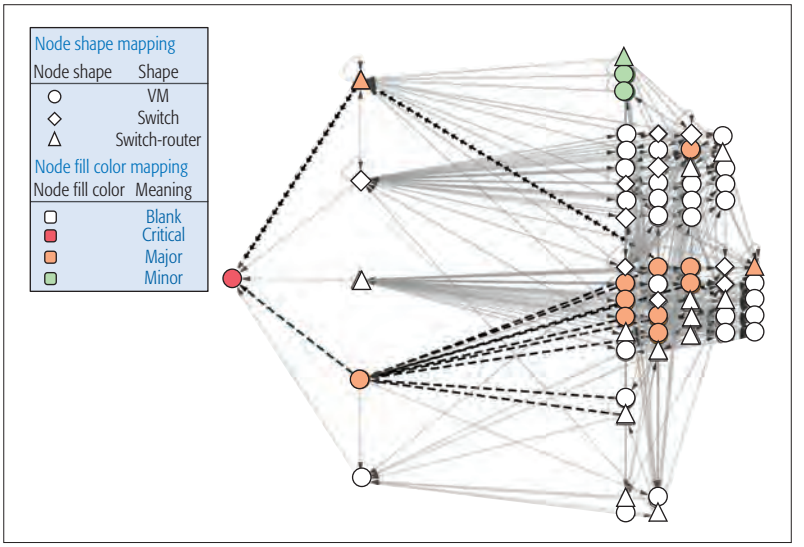

Figure 5. CPU threshold violation critical event chain influences.

References

[1] L. Bennacer *et al.*, "Self-Diagnosis Technique for Virtual Private Networks Combining Bayesian Networks and Case-Based Reasoning," *IEEE Trans. Autom. Sci. Eng.*, vol. 12, no. 1, Jan. 2015, pp. 354–66.

[2] L. Bennacer *et al.*, "Scalable and Fast Root Cause Analysis using Inter Cluster Inference," *Proc. IEEE ICC*, 2013, pp. 3563–68.

[3] D. Liu *et al.*, "Opprentice: Towards Practical and Automatic Anomaly Detection Through Machine Learning," *Proc. 2015 ACM Conf. Internet Meas. Conf.*, 2015, pp. 211–24.

[4] T. Kimura *et al.*, "Spatio-Temporal Factorization of Log Data for Understanding Network Events," *Proc. IEEE INFOCOM*, 2014, pp. 610–18.

[5] J. Li et al., "Personalized Influential Topic Search via Social Network Summarization," *IEEE Trans. Knowl. Data Eng.*, vol. 28, no. 7, July 2016, pp. 1820–34.

[6] R. Quinn *et al.*, "KnowNet: Towards a Knowledge Plane for Enterprise Network Management," *Proc. NOMS 2016*, 2016, pp. 249–256.

[7] N. Chawla and K. Bowyer, "SMOTE: Synthetic Minority Over-Sampling Technique Nitesh," *J. Artif. Intell. Res.*, vol. 16, no. 1, 2002, pp. 321–57.

[8] F. Salfner, M. Lenk, and M. Malek, "A Survey of Online Failure Prediction Methods," *ACM Comput. Surv.*, vol. 42, no. 3, Mar. 2010, p. 1–42.

[9] T. Hastie, R. Tibshirani, and J. H. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed., Springer, 2009.

[10] V. D. Blondel *et al.*, "Fast Unfolding of Communities in Large Networks," *J. Stat. Mech. Theory Exp.*, vol. 10008, no. 10, Mar. 2008, p. 6.

Additional Reading

[1] H. Yan *et al.*, "G-RCA: A Generic Root Cause Analysis Platform for Service Quality Management in Large IP Networks," *IEEE/ACM Trans. Networking*, vol. 20, no. 6, Dec. 2012, pp. 1734–47.

Biographies

José Manuel Navarro González (josemanuel.navarro@upm.es) obtained an M.Sc. in telecommunications engineering from Universidad Miguel Hernndez de Elche, Spain, in 2013. He is a Ph.D. candidate at Escuela Técnica Superior de Ingenieros de Telecomunicación at Universidad Politécnica de Madrid (UPM), Spain. His research interests are distributed systems management through applied machine learning and the Internet of Things.

Javier Andión Jiménez [M] (j.andion@upm.es) is a telecommunications engineer at UPM. He is pursuing a Ph.D. in proactive failure prediction with machine learning at Escuela Técnica Superior de Ingenieros de Telecomunicación, UPM. He has been involved in the organization of several workshops and congresses related to artificial intelligence, robotics, and programming challenges.

Juan Carlos Dueñas López [SM] (juancarlos.duenas@upm.es) obtained a Degree in telecommunications engineering from UPM in 1991and a Ph.D. from UPM in 1994. He is currently deputy vice-president for research at UPM.

Hugo A. Parada G. (hugo.parada@upm.es) is an assistant professor at UPM's Escuela Técnica Superior de Ingeniera y Sistemas de Telecomunicación. His current areas of research include machine learning, big data, and cloud computing. He received his Ph.D. in telecommunications engineering from UPM in 2010

RADIO COMMUNICATIONS: COMPONENTS, SYSTEMS, AND NETWORKS

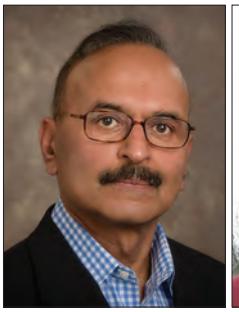

Amitabh Mishra

Tom Alexander

Multiple-input multiple-output (MIMO) technology is virtually taken for granted today in all areas of wireless communication, ranging from consumer Wi-Fi products to Long Term Evolution (LTE) cellular base stations. Decades of research and development have produced improvements in space-time coding and channel estimation that bring achievable MIMO capacity ever closer to the theoretical channel capacity. Since channel capacity in most RF environments generally increases with the number of antennas, the trend over the last decade has been to place an ever increasing number of antennas on communication terminals. Unfortunately, this is counterbalanced with the space, power, and aesthetic limitations of user equipment; nobody really appreciates a handset that resembles a porcupine and runs hot enough to cook an egg.

Of late, therefore, most of the increase in antennas occurs at the base stations or Wi-Fi access points, rather than at the mobile terminals. This leads to an asymmetric communications architecture that is nevertheless quite useful: rather than increasing the capacity of an individual terminal-to-base-station link, the larger number of base station antennas is utilized to support simultaneous communications with multiple terminals. This is commonly referred to as multi-user MIMO (MU-MIMO) and has been posited since the beginning of MIMO research, although it is only fairly recently that advances in coding and digital signal processing (DSP) implementations have made this practical.

Massive MIMO is a logical evolution of this trend. Rather than a dozen or fewer antennas, massive MIMO assumes hundreds or even thousands of antennas at the base station, each individually driven by a separate RF chain. In fact, the number of RF channel multipaths is sparse relative to the number of antennas. This supports the ability of massive MIMO to provide significant resistance to small-scale fading and open up interesting new techniques for multi-user communication. The notion of building, synchronizing, and processing several hundred microwave RF chains is daunting at first glance, but so was MU-MIMO several years ago. Testbeds exceeding 100 antennas have now been operated and show impressive system capacity levels, nearly 150 b/s/Hz. Massive MIMO is an even better complement to millimeter-wave (mmWave) technology. The dimensions of antennas and RF components shrink tremendouslyat mmWave frequencies, making massive MIMO reasonably practical, and at the same time the ability of massive MIMO to overcome path losses and interference becomes essential for achieving reliable mmWave links.

However, massive MIMO has its own share of problems. Achieving many of its benefits requires very precise and continuously updated knowledge of the RF channel for every link. Spatial correlation between such large numbers of antennas in close proximity is significantly increased, causing issues with accurate channel estimation. Pilot contamination from nearby terminals or base stations is difficult to avoid. And of course, achieving the truly large numbers of independently driven antennas at reasonable cost is still an unsolved problem. Nevertheless, massive MIMO has been identified as a key technology and research area for 5G wireless, and there is a great deal of ongoing work in this field.

We present an article in this issue of the Radio Communications Series, "Hybrid Beamforming for Massive MIMO: A Survey," which deals with a particularly interesting area where massive MIMO gains can be exploited. Rather than the pure digital beamforming used with "traditional" MIMO, massive MIMO allows a hybrid of analog and digital beamforming to be applied. Combining analog beamformers with digital beamformers in a coordinated fashion allows radical simplification of the digital portion of the system while retaining most of the benefits of the large number of antennas. Of course, accurate channel state information (CSI) is key to achieving this. The article provides a taxonomy of hybrid beamforming structures in terms of CSI, and then explores various dimensions of the hybrid beamforming techniques that have been proposed. It then surveys scenarios for applying hybrid beamforming to mmWave communications and notes some key issues. Guidelines are provided to adapt various techniques to different situations.

We are grateful to our reviewers for helping us select and improve the quality of the papers we publish in this series, as well as to the many authors who submit their work for publication. The support and encouragement of the Editor-in-Chief and the publication staff are invaluable.

IEEE Access®

The **journal** for rapid **open access** publishing

Best New Journal in STM 2015

Become a published author in 4 to 6 weeks.

Get on the fast track to publication with the multidisciplinary open access journal worthy of the IEEE.

IEEE journals are trusted, respected, and rank among the most highly cited publications in the industry. IEEE *Access* is no exception with a typical **one-third** acceptance rate. Even though it provides authors faster publication time, every submitted article still undergoes extensive peer review to ensure originality, technical correctness, and interest among readers.

Published only online, IEEE *Access* is ideal for authors who want to quickly announce recent developments, methods, or new products to a global audience.

Publishing in IEEE *Access* allows you to:

- Submit multidisciplinary articles that do not fit neatly in traditional journals
- Reach millions of global users through the IEEE *Xplore*® digital library with free access to all
- Integrate multimedia with articles
- Connect with your readers through commenting
- Track usage and citation data for each published article
- Publish without a page limit for **only $1,750** per article

Learn more about this award-winning journal at:
www.ieee.org/ieee-access

IEEE
Advancing Technology for Humanity

14-PUB-196 11/15

Hybrid Beamforming for Massive MIMO: A Survey

Andreas F. Molisch, Vishnu V. Ratnam, Shengqian Han, Zheda Li, Sinh Le Hong Nguyen, Linsheng Li, and Katsuyuki Haneda

Hybrid multiple-antenna transceivers, which combine large-dimensional analog pre/postprocessing with lower-dimensional digital processing, are the most promising approach for reducing the hardware cost and training overhead in massive MIMO systems. The article provides a comprehensive survey of the various incarnations of such structures that have been proposed in the literature.

ABSTRACT

Hybrid multiple-antenna transceivers, which combine large-dimensional analog pre/postprocessing with lower-dimensional digital processing, are the most promising approach for reducing the hardware cost and training overhead in massive MIMO systems. This article provides a comprehensive survey of the various incarnations of such structures that have been proposed in the literature. We provide a taxonomy in terms of the required channel state information, that is, whether the processing adapts to the instantaneous or average (second-order) channel state information; while the former provides somewhat better signal-to-noise and interference ratio, the latter has much lower overhead for CSI acquisition. We furthermore distinguish hardware structures of different complexities. Finally, we point out the special design aspects for operation at millimeter-wave frequencies.

INTRODUCTION

Multiple-input multiple-output (MIMO) technology, that is, the use of multiple antennas at transmitter (TX) and receiver (RX), has been recognized since the seminal works of Winters, Foschini and Gans, and Telatar, as an essential approach to high spectral efficiency (SE). In its form of multi-user MIMO (MU-MIMO), it improves SE in two forms:

- A base station (BS) can communicate simultaneously with multiple user equipments (UEs) on the same time-frequency resources.
- Multiple data streams can be sent between the BS and each UE.

The total number of data streams (summed over all UEs in a cell) is upper limited by the smaller of the number of BS antenna elements, and the sum of the number of all UE antenna elements.

While MU-MIMO has been studied for more than a decade, the seminal work of Marzetta introduced the exciting new concept of "massive MIMO," where the number of antenna elements at the BS reaches dozens or hundreds. Not only does this allow increasing the number of data streams in the cell to very large values, it also simplifies signal processing, creates "channel hardening" such that small-scale fading is essentially eliminated, and reduces the required transmission energy due to the large beamforming gain; see, for example, [1] for a review. Massive MIMO is *beneficial* at centimeter-wave (cmWave) frequencies, but is *essential* in the millimeter-wave (mmWave) bands,[1] since the high free-space path loss at those frequencies necessitates large array gains to obtain sufficient signal-to-noise ratio (SNR), even at moderate distances of about 100 m.

However, the large number of antenna elements in massive MIMO also poses major challenges:

- A large number of radio frequency (RF) chains (one for each antenna element) increases cost and energy consumption.
- Determining the channel state information (CSI) between each transmit and receive antenna uses a considerable amount of spectral resources.

A promising solution to these problems lies in the concept of *hybrid* transceivers, which use a combination of analog beamformers in the RF domain, together with digital beamforming in the baseband, connected to the RF with a smaller number of up/downconversion chains. Hybrid beamforming was first introduced and analyzed in the mid-2000s by one of the authors and collaborators in [2, 3]. It is motivated by the fact that the number of up-downconversion chains is only lower-limited by the number of data streams that are to be transmitted, while the beamforming gain and diversity order is given by the number of antenna elements if suitable RF beamforming is done. While formulated originally for MIMO with arbitrary number of antenna elements (i.e., covering both massive MIMO and small arrays), the approach is of interest in particular to massive MIMO. Interest in hybrid transceivers has therefore been revived over the past three years (especially following the papers of Heath and co-workers, e.g., [4]), where various structures have been proposed in different papers. Thus, the time seems ripe for a review of the state of the art, and a taxonomy of the various transceiver architectures (often simplified to provide computational or chip-architectural advantages) and algorithms. The current article aims to provide this overview, and point out topics that are still open for future research.

This survey covers hybrid beamforming structures using instantaneous or average CSI in the following two sections. A special structure incor-

A version of this article with additional references can be found at arxiv.org/abs/1609.05078.

[1] In a slight abuse of notation, we denote 1–10 GHz as "centimeter waves," and 10–100 GHz as "millimeter waves."

Digital Object Identifier: 10.1109/MCOM.2017.1600400

Andreas F. Molisch, Vishnu V. Ratnam, Shengqian Han, and Zheda Li are with the University of Southern California, Los Angeles; Shengqian Han is also with Beihang University; Sinh Le Hong Nguyen, Linsheng Li, and Katsuyuki Haneda are with the Aalto University School of Electrical Engineering; Linsheng Li is presently with Huawei Helsinki.

porating switches between the analog and digital parts is then described. Following that, we clarify constraints at mmWave bands due to propagation conditions and hardware imperfections. A summary and conclusions round up the article.

Hybrid Beamforming Based on Instantaneous CSI

Figure 1 shows block diagrams of three hybrid beamforming structures at the BS, where we assume a downlink transmission from the BS (acting as TX) to the UE (RX). The classification is applicable to both cmWave and mmWave bands. At the TX, a baseband digital precoder $\mathbf{F}_{BB}$ processes N_S data streams to produce N_{RF}^{BS} outputs, which are upconverted to RF and mapped via an analog precoder $\mathbf{F}_{RF}$ to N_{BS} antenna elements for transmission. The structure at the RX is similar: an analog beamformer $\mathbf{W}_{RF}$ combines RF signals from N_{UE} antennas to create N_{RF}^{UE} outputs, which are downconverted to baseband and further combined using a matrix $\mathbf{W}_{BB}$, producing signal $\mathbf{y}$ for detection/decoding.[2] Hence, we use terms "beamformer" and "precoder/combiner" interchangeably hereinafter. For a full-complexity structure, each analog precoder output can be a linear combination of *all* RF signals (Fig. 1, *A*). Complexity reduction at the price of somewhat reduced performance can be achieved when each RF chain can be connected only to a subset of antenna elements, as in Fig. 1, *B*. Different from structures *A* and *B*, where baseband signals are jointly processed by a digital precoder, structure *C* employs the analog beamformer to create multiple "virtual sectors," which enables separated baseband processing, downlink training, and uplink feedback, and therefore reduces signaling overhead and computational complexity [5].

Even assuming full-instantaneous CSI at the TX, it is very difficult to find the analog and digital beamforming matrices that optimize, for example, the net data rates of the UEs [6]. The main difficulties include:

- Analog and digital beamformers at each link end, as well as combiners at the different link ends, are coupled, which makes the objective function of the resulting optimization non-convex.
- Typically, the analog precoder/combiner is realized as a phase-shifter network, which imposes additional constraints on the elements of $\mathbf{W}_{RF}$ and $\mathbf{F}_{RF}$.
- Moreover, with finite-resolution phase shifters, the optimal analog beamformer lies in a discrete finite set, which typically leads to NP-hard integer programming problems.

Two main methodologies are explored to alleviate these challenges and achieve feasible near-optimal solutions.

Approximating the Optimal Beamformer

For single-user MIMO (SU-MIMO), we start with optimum beamforming for the fully digital case with $N_{RF}^{BS} = N_{BS}$ and $N_{RF}^{UE} = N_{UE}$, where the solution is known (dominant left/right singular vectors of a channel matrix $\mathbf{H}$ from singular value decomposition [SVD]). Then one approach (e.g., [2]) is based on eigen decomposition, while another (e.g., [6]) finds an (approximate) optimum hybrid beamformer by minimizing the Euclidean distance to this fully digital one. The objective function of the approximation problem is still non-convex, but much less complex than the original one. For sparse channels (as occur in mmWave bands), minimizing this distance provides a quasi-optimal solution. In non-sparse channels, such as usually occur at cmWave bands, an alternating optimization of analog and digital beamformers can be used. A closed-form solution for each of the alternating optimization steps can be developed for the reduced-complexity structure, while for the full-complexity structure, the non-convex problem can be expanded into a series of convex sub-problems by restricting the phase increment of the analog beamformer within a small vicinity of its preceding iteration.

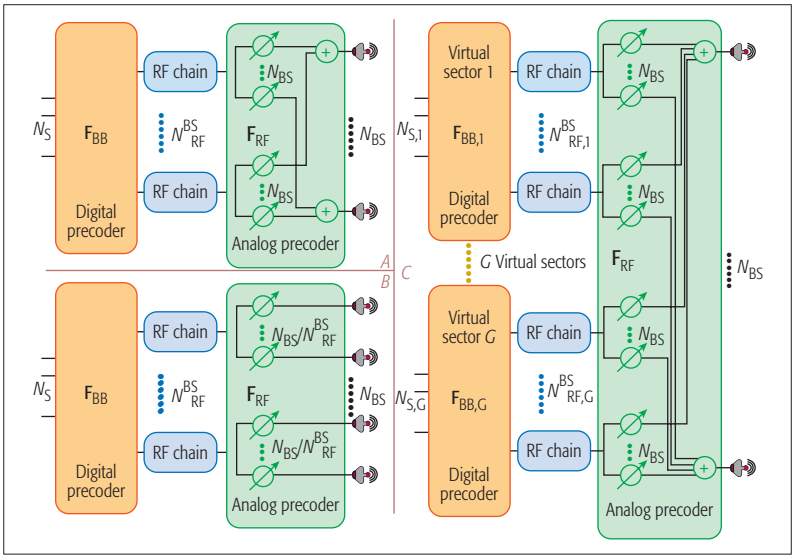

Figure 1. Block diagrams of hybrid beamforming structures at the BS for a downlink transmission, where structures *A*, *B*, and *C* denote the full-complexity, reduced-complexity, and virtual sectorization structures, respectively.

Figure 2 compares the performance of the three structures for downlink transmission of single-cell MU massive MIMO. The full-complexity structure of Fig. 1, *A*, performs the same as the fully digital structure when the number of RF chains is no smaller than the number of users (or streams). Performance loss of structure *B* is rather large for the considered MU case, although it is much smaller for SU-MIMO (not shown here). For structure C, the employed algorithm (JSDM, discussed later in this article) divides the users into four or eight groups, which might lead to a performance floor due to inter-group interference; note that the significantly reduced training overhead of JSDM is not shown here; this is discussed later.

Decoupling the Design of the Analog and Digital Beamformers

One of the main challenges in hybrid beamformer design is the coupling among analog and digital beamformers, and between the beamformers at TX and RX. This motivates decoupling the beamformer designs for reducing the problem complexity. By assuming some transceiver algorithms, optimization of beamforming matrices can be solved sequentially. For example, in order to maximize the net rate for SU-MIMO, one can eliminate the impact of the combiner on the pre-

[2] Obviously, a UE with a single antenna element is a special case.

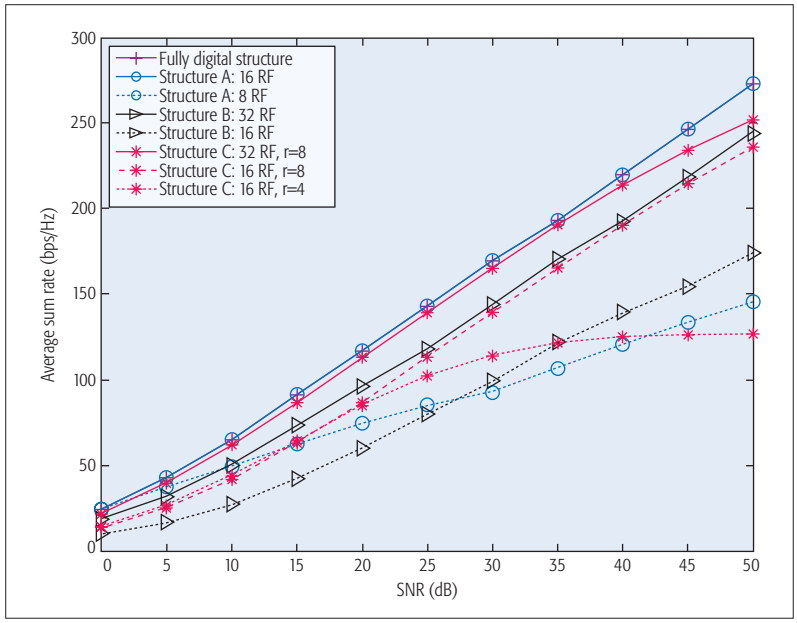

Figure 2. Performance comparison of the three hybrid structures with MU-MIMO; $N_{BS} = 64$, $N_{UE} = 1$, 4 groups of users located in a sector with mean directions [−45°, −15°, 15°, 45°], and each group has 4 users. AoDs of MPCs concentrate around the mean directions of each group with 10° AoD spread. {This analysis assumes ideal hardware and typical channel conditions for cmWaves.)

coder by assuming a fully digital minimum mean square error (MMSE) receiver. Further decoupling of the analog and digital precoder is possible by assuming that the digital precoder is unitary. Subsequently, $\mathbf{F}_{RF}$ is optimized column by column by imposing the phase-only constraint on each antenna. With the known analog precoder, a closed-form expression of the digital precoder can then be obtained.

Alternatively, some simple heuristic decoupling beamforming strategies have been explored. For example, the element-wise normalized conjugate beamformer can be used as the analog precoder, with which the asymptotic signal-to-interference-plus-noise ratio (SINR) of hybrid beamforming is only reduced by a factor of $\pi/4$ compared to fully digital beamforming when letting the number of antenna elements and streams, N_{BS} and N_S, go to infinity while keeping N_{BS}/N_S constant.

Extending to the situation where the UE is also equipped with a hybrid structure for MU-MIMO, one can first construct the RF combiner by selecting the strongest receive beams from the Fourier codebook to maximize the Frobenius norm of the combiner-projected channel. Then the same normalized eigenbeamformer is implemented as the analog precoder on the effective channel. In the baseband, the BS performs block diagonalization (BD) over the projected channel to suppress inter-user interference.

Wideband Hybrid Beamforming

The above discussion focused on narrowband (i.e., single-subcarrier) systems. In wideband orthogonal frequency-division multiplexing (OFDM) systems, however, analog beamformers cannot have different weights across subcarriers; for strongly frequency-selective channels, such beamformers extending over the whole available band adapt to the average channel state.

Frequency-domain scheduling was believed unnecessary for fully digital massive MIMO systems because the sufficiently large number of antennas can harden the channels and provide sufficient spatial degrees of freedom for multiplexing UEs [1]. However, under practical constraints on array size (e.g., according to 3GPP LTE Release 13), frequency-domain scheduling is still necessary for hybrid transceivers [7]. With frequency-domain scheduling, UEs are served on different subcarriers, making the existing narrowband hybrid precoders no longer applicable. Existing works have studied the joint optimization of wideband analog precoder and narrowband digital precoders, aimed at minimizing the BS transmit power or maximizing the sum rate of UEs.

Another important issue in the existing design of hybrid beamforming is control signaling coverage. While narrow analog beams are preferred for user-specific data transmission, wide beams are preferred for broadcasting control signals to all UEs. This problem may be solved, for example, by splitting signaling and data planes so that they are transmitted at different carrier frequencies.

Impact of Phase-Only Constraint and the Number of RF Chains

Hybrid beamforming does not necessarily have inferior performance to fully digital beamforming. Analog beamforming can be implemented by means of phase shifters together with variable gain amplifiers. In this case, analog beamforming can provide the same functionality as digital beamforming, and combine desired multipath components (MPCs) (and suppress interfering MPCs) to the same degree as linear digital processing. Thus, in a narrowband massive MIMO system, with full-instantaneous CSI at the TX, this hybrid beamforming can achieve the same performance as fully digital beamforming if $N_S \le N_{RF}$ [2]. A similar result can be obtained for a wideband system, where the number of RF chains of the hybrid structure should be not smaller than $\min(N_{BS}, N_{S,wb})$ with $N_{S,wb}$ denoting the total number of data streaming over all subcarriers [7].

Since two phase-only entries for the analog precoder are equivalent to a single unconstrained (amplitude and phase) entry, fully digital performance can be achieved with phase-only hybrid structures if $N_{RF}^{BS} \ge 2N_S$ in narrowband systems [2].

Hybrid Beamforming Based on Averaged CSI

Average CSI based Hybrid Beamforming

A major challenge for the beamformers discussed previously is the overhead for acquiring CSI at the BS. Information-theoretic results taking training overhead into account show that for time-division duplexing (TDD) systems, the spatial multiplexing gain (SMG) of massive MIMO downlinks with fully digital structure equals $M(1 - M/T)$, where $M = \min(N_{BS}, K, T/2)$, $K = N_S$ is the number of single-antenna users, and T is the number of channel uses in a coherence time-frequency block [5]. In frequency-division duplexing (FDD) systems, the overhead is even larger, since both downlink training and uplink feedback for each antenna are required. In addition to the coherence time, the

frame structure of systems may provide additional constraints for the pilot repetition frequency and thus the training overhead.

It is evident that for any massive MIMO systems relying on full CSI between all antenna elements of the BS and UEs, the maximal achievable SMG is limited by the size of the coherence block of the channel because N_{BS} and K are generally large. This necessitates the design of transmission strategies with reduced-dimensional CSI to relieve the signaling overhead. Specifically, a number of papers have considered analog beamforming based on slowly varying second order statistics of the CSI at the BS (a two-stage beamformer, with the first analog stage based on the average CSI only, followed by a digital one adapted to instantaneous CSI). The beamforming significantly reduces the dimension of the effective instantaneous CSI for digital beamforming within each coherent fading block by taking advantage of a small angular spread at the BS. Such structures work robustly even with analog beamformers, which cannot usually adapt to varying channels as quickly as digital beamformers.

Hybrid beamformers using average CSI for the analog part were first suggested in [3], which also provided closed-form approximations for the optimum beamformer in SU-MIMO systems. For the MU case, [5] proposed a scheme called "joint spatial division multiplexing" (JSDM), which considered a hybrid-beamforming BS and single-antenna UEs; to further alleviate the downlink training/uplink feedback burden, UEs with similar transmit channel covariance are grouped together, and inter-group interference is suppressed by an analog precoder based on the BD method. Specifically, using the Karhunen-Loeve representation, the N_{BS}-by-1 channel vector can be modeled as $\mathbf{h} = \mathbf{U}\Lambda^{1/2}\mathbf{w}$, where $\mathbf{w} \in \mathbb{C}^{r\times 1} \sim \mathcal{CN}(0, \mathbf{I}_r)$, Λ is an r-by-r diagonal matrix, which aligns eigenvalues of channel covariance $\mathbf{R}$ on its diagonal, $\mathbf{U} \in \mathbb{C}^{M\times r}$ indicates the eigenmatrix of $\mathbf{R}$, and r denotes the rank of the channel covariance. Dividing UEs into G groups and assuming that UEs in the same group g exhibit the same channel covariance $\mathbf{R}_g$ with rank r_g, the JSDM analog precoder is

$$\mathbf{F}_{RF} = [\mathbf{F}_{RF,1}, \ldots, \mathbf{F}_{RF,G}] \text{ with } \mathbf{F}_{RF,g} = \mathbf{E}_g\mathbf{G}_g.$$

By selecting $r_g^* \le r_g$ dominant eigenmodes of $\mathbf{R}_g$, denoted by $\mathbf{U}_g^*$, JSDM builds the eigenmatrix of the dominant interference to the gth group: $\Xi_g = [\mathbf{U}_1^*, \ldots, \mathbf{U}_G^*]$. Then $\mathbf{E}_g$ consists of the null space of Ξ_g, and $\mathbf{G}_g$ consists of dominant eigenvectors of $\mathbf{E}_g^\dagger\mathbf{R}_g\mathbf{E}_g$. This creates multiple "virtual sectors" in which downlink training can be conducted in parallel, and each UE only needs to feed back the intra-group channels, leading to the reduction of both training and feedback overhead by a factor equal to the number of virtual sectors.

In practice, however, to maintain the orthogonality between virtual sectors, JSDM often conservatively groups UEs into only a few groups, because UEs' transmit channel covariances tend to be partially overlapped with each other. This limits the reduction of training and feedback overhead. Once grouping UEs into more virtual sectors violates the orthogonality condition, JSDM is not able to combat the inter-group interference. Eliminating overlapped beams of UEs in different groups is a heuristic approach to solve this problem. In [8], JSDM is generalized to support non-orthogonal virtual sectorization, and a modified MMSE algorithm is proposed to optimize the multi-group digital precoders to maximize the lower bound of the average sum rate.

Two UE grouping methods have been proposed as extensions to JSDM: K-means clustering and fixed quantization. In the large antenna limit, the number of downlink streams served by JSDM can be optimized given the angle of departure (AoD) of MPCs and their spread for each UE group. To reduce the complexity of JSDM, in particular due to SVD, an online iterative algorithm can be used to track the analog precoder under time-varying channels. When considering single-antenna UEs, a Fourier codebook-based analog precoder, and a zero-forcing (ZF) digital precoder, the performance of JSDM can be further improved by jointly optimizing the analog precoder and allocation of RF chains to groups based on second order channel statistics. This principle can be extended to multicell systems, where an outage constraint on the UEs' SINR can be considered.

Decoupling of Analog and Digital Beamformers

Different from the previous section, where both analog and digital beamformers are based on instantaneous CSI, now analog and digital beamformers are based on the average CSI and the instantaneous effective CSI, respectively. Thus, to find the optimal beamformers, one needs to first design the digital beamformer for each snapshot of the channel and then derive the analog beamformer based on their long-term time-average, making their mathematical treatment difficult. Decoupled designs of the analog and digital beamformers therefore make the optimization problem simpler and practically attractive. For SU-MIMO where a UE is equipped with a single RF chain and multiple antennas, the optimal analog combiner is intuitively the strongest eigenmode of the UE-side channel covariance. However, when there are more RF chains at the UE, the strongest eigenmodes are not always the optimal combiners since they may be associated with a single transmit eigenmode of the BS-side channel covariance. For MU-MIMO with multiple RF chains at both link ends, the digital beamformer design needs to consider the UE-level spatial multiplexing and inter-user interference suppression, which will affect the analog beamformer design. In [8], the optimality (in the sense of maximizing the so-called intra-group signal-to-inter-group interference-plus-noise ratio) of decoupling analog and digital beamformers is shown under the Kronecker channel model.

Full-Dimensional MIMO in 3GPP

While the Third Generation Partnership Project (3GPP) standard does not prescribe particular transceiver architectures, hybrid digital-and-analog structures have motivated the design of CSI acquisition protocols in Release 13 of LTE-Advanced Pro in 3GPP, especially the non-precoded and beamformed pilots for full-dimensional (FD) MIMO. The non-precoded beamformer is related to the reduced-complexity structure B in

While the 3GPP standard does not prescribe particular transceiver architectures, hybrid digital-and-analog structures have motivated the design of CSI acquisition protocols in Release 13 of LTE-Advanced Pro in 3GPP, especially the non-precoded and beamformed pilots for FD MIMO.

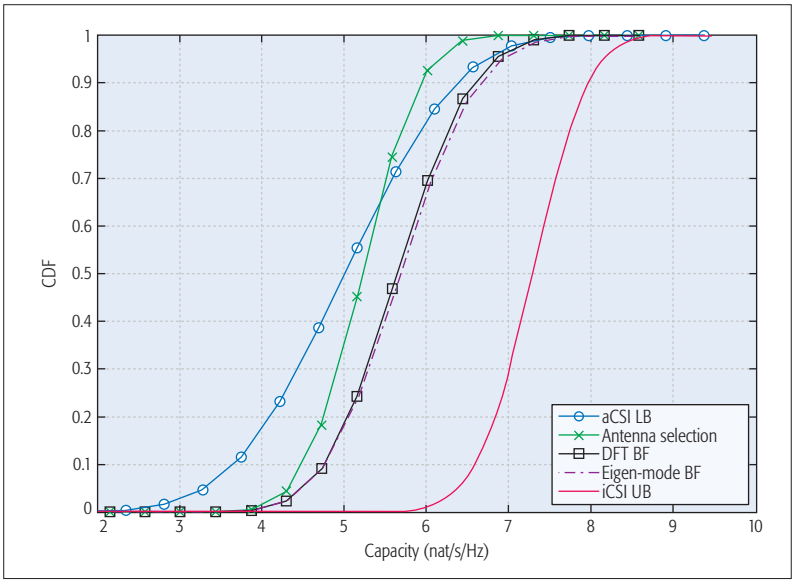

Figure 3. Performance of different analog precoders in a hybrid TX with selection. We consider an SU-MIMO system at cmWaves with ideal hardware conditions, where the RX has full complexity with $N_{UE} = N_{RF}^{UE} = 2$ and the TX has a switched hybrid beamforming structure with $N_{BS} = L = 10$, $N_{RF}^{BS} = 2$. The channels are Rayleigh distributed in amplitude, doubly spatially correlated (both at TX and RX), and follow the Kronecker model of spatial correlation $[R_{BS}]_{ij} = [R_{UE}]_{ij} = 0.5^{|i-j|}$; "aCSI LB" and "iCSI UB" refer to optimal unconstrained precoding with average CSI [3] and with instantaneous CSI, respectively. The RX SNR is 10 dB.

Fig. 1, where a (possibly static) analog precoder is applied to a subset of an antenna array to reduce the training overhead. The beamformed approach may assume the full-complexity structure *A* in Fig. 1, where analog beamformers are used for downlink training signals. The BS transmits multiple analog precoded pilots in different time or frequency resources. Then user feedback indicates the preferred analog beam; given this, the user can further measure and feed back the instantaneous effective channel in a legacy LTE manner. These approaches can, under some circumstances, reduce the overhead in average CSI acquisition, and generally perform well for SU-MIMO but may suffer large performance degradation for MU-MIMO unless the average CSI of all users is fed back. Recently proposed hybrid CSI acquisition schemes in 3GPP combine the above two approaches. First, the BS sends non-precoded pilots to estimate the average CSI at users. Then, based on the analog (non-codebook-based) or digital (codebook-based) feedback of the average CSI from users, the BS determines the analog beamformer and next sends beamformed pilots. These hybrid schemes essentially enable the form of beamforming discussed above in this section, namely, adaptation of the analog beamformer based on long-term statistics, which is then followed by the digital beamformer based on instantaneous effective CSI. Increasing the array size further motivates studies to reduce the training and feedback overhead through, for example, aperiodic training schemes. The JSDM-based structure C in Fig. 1 that separates a cell into multiple "virtual sectors" is one approach to reduce the overhead significantly by simultaneous downlink training and uplink feedback across virtual sectors.

Hybrid Beamforming with Selection

A special class of hybrid systems involves a selection stage that precedes (at the TX) or succeeds (at the RX) the analog processing, called hybrid beamforming with selection hereinafter. The up-converted data streams at the TX pass through the analog precoder $\mathbf{F}_{RF}$, as discussed. However, unlike conventional hybrid beamforming, the number of input ports of the analog block is $L \geq N_{RF}^{BS}$ (and typically, $L = N_{BS}$). A selection matrix $\mathbf{S}$, realized by a network of RF switches, feeds the data streams to the best N_{RF}^{BS} out of the L ports for transmission. The premise for such a design is that, unlike switches, analog components like phase shifters and amplifiers might not be able to adapt to the quick variation of instantaneous channels over time. Therefore, $\mathbf{F}_{RF}$ is either fixed or designed based on average channel statistics as described earlier, and $\mathbf{S}$ picks the best ports for each channel realization, thus making the effective analog processing more channel adaptive. The switching networks are also advantageous over full-complexity analog beamforming in terms of their cost and energy efficiency (EE). Although we focus on the TX for brevity, a switched analog combiner may also be implemented at the RX.

Design of Analog Precoding/Combining Block

The simplest "hybrid beamforming with selection" performs the antenna selection and omits the analog precoding. However, significant beamforming gains can be achieved by introducing analog precoding before the selection, to take advantage of the spatial MPCs. Such an architecture performs signal processing in the beam-space. While $\mathbf{F}_{RF}$ may be designed by discrete Fourier transform (DFT), its performance can be improved by eigenmode beamforming based on the TX correlation matrix [3]. To reduce the CSI feedback overhead for FDD systems, $\mathbf{F}_{RF}$ in the conventional hybrid beamforming can be chosen from a set of a predetermined codebook of matrices. By regarding the codebook entries as realizations from switch positions, this design can be interpreted as hybrid beamforming with selection. The codebook design is discussed, for example, in [9]. The performance of some of these analog precoders is compared in Fig. 3.

Design of Selection Matrix

Since complexity of searching for the best ports in the analog block is exponentially increasing with N_{RF}, many algorithms have been proposed to reduce it. Several greedy algorithms have been proposed to leverage diversity and spatial-multiplexing gains in an MU scenario. Restricted selection architectures allow each RF chain to choose from a subset of the analog ports, thereby reducing both search and hardware complexities. Iterative algorithms with varying search complexities from a linear to sub-exponential order have been proposed. An alternative technique that does not use the instantaneous CSI is called eigen-diversity beamforming [10]. It draws the selection matrix for each channel realization from an optimized probability distribution, thereby leveraging the temporal diversity.

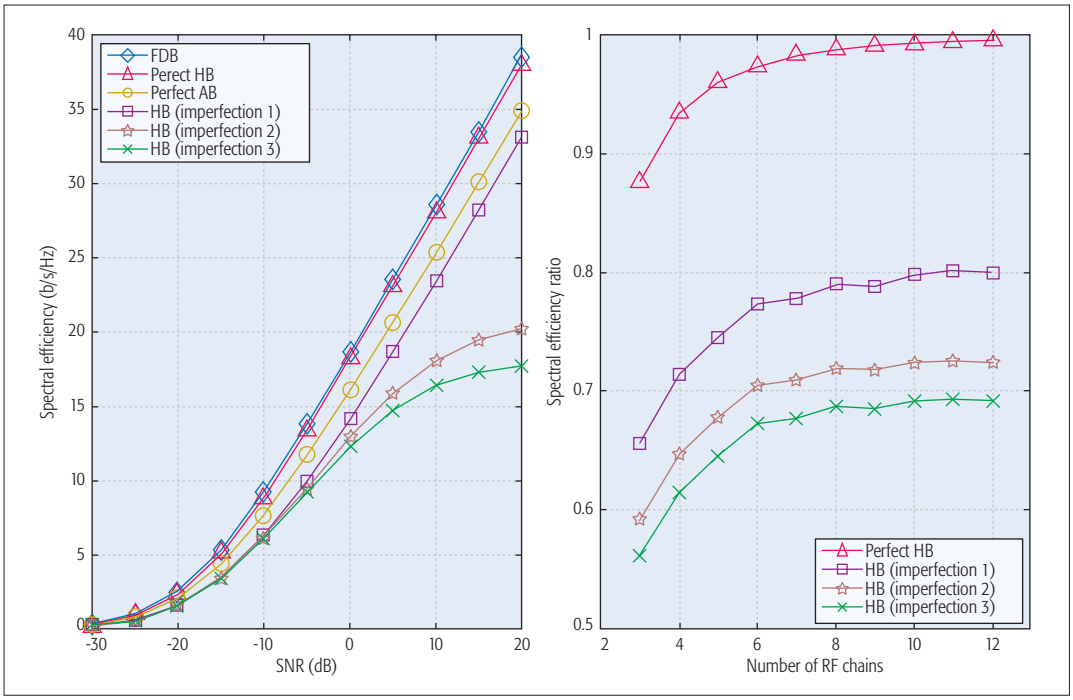

Figure 4. SE comparison of fully digital beamforming (FDB), Hybrid beamforming with perfect RF hardware (Perfect HB), analog-only beamsteering with perfect RF hardware (Perfect AB), and HB with three different cases of RF hardware imperfection: case 1 considers quantization error caused by 6-bit phase shifters; cases 2 and 3 additionally consider residual transceiver impairments at the BS and at both BS and UE, respectively. The spatially sparse precoding [6] is used in the HB. We assume that $N_{BS} = 64$, $N_{UE} = 16$, $N_S = 3$, the radio channel has 3 multipath clusters, and each has 6 rays, as representative of mmWave channels. The residual transceiver impairments at TX and RX are characterized by error-vector magnitude of –20 dB. In the left subfigure, $N_{RF}^{BS} = N_{RF}^{UE} = 6$. In the right subfigure, the SE of the HB with different RF hardware assumptions normalized to the FDB is characterized at SNR = 0 dB and $N_{RF}^{BS} = N_{RF}^{UE}$.

The presence of RF transceiver imperfections degrades SE in various ways. For example, it is harder to accurately generate desired transmit signals when higher beamformer gain is aimed for; non-linear distortion at the RX depends on the instantaneous channel gain and hence the SNR.

Hybrid Beamforming at mmWave

Hybrid beamforming architectures and algorithms in the cmWave band described in the previous sections can in principle be used at mmWave frequencies. In practice, however, propagation channel and RF hardware aspects are significantly different in those bands, and hence novel hybrid beamforming techniques taking into account the practicalities are needed. At mmWave frequencies, the multipath channel experiences higher propagation loss, which needs to be compensated by gain from antenna arrays at the TX, RX, or both. While such arrays have reasonable physical size thanks to short wavelengths, fully digital beamforming solutions become infeasible, and hybrid beamforming becomes harder due to power- and cost-related RF hardware constraints. Moreover, mmWave channels may be sparser, such that fewer spatial degrees of freedom are available. The sparsity can be exploited for optimizing channel estimation and beam training.

Hybrid Beamforming Methods Exploiting Channels' Sparsity

Exploiting the channels' sparsity, the simplest form of hybrid beamforming in SU-MIMO systems focuses array gains to a limited number of multipaths in the RF domain, while multiplexing data streams and allocating powers in baseband. This hybrid architecture is asymptotically optimum in the limit of large antenna arrays [11].

For systems with practical sizes of arrays, which have, for example, 64 to 256 elements for the BS and under 20 elements for the UEs, hybrid beamforming structures are highly desirable. In addition, reduction of the hardware and computational complexity is of great interest. For those purposes, a number of hybrid beamforming methods have been proposed for mmWave SU-MIMO channels that can be categorized into the use of codebooks, spatially sparse precoding, antenna selection, and beam selection.

Use of Codebooks: While having the same principle as the schemes described earlier, the codebook-based beamforming does not directly estimate the large CSI matrix at the RX, but instead performs downlink training using pre-defined beams and then only feeds back the selected beam IDs to the transmitter. To further reduce the complexity of beam search and feedback overhead for large antenna systems, a codebook for full-complexity hybrid architecture can be designed to exploit the sparsity of mmWave channels. Each codeword is constructed based on the Orthogonal Matching Pursuit (OMP) algorithm to minimize the MSE with the pre-defined ideal beam pattern.

Spatially Sparse Precoding: This method finds the approximation of the unconstrained (i.e., fully digital) beamformer as described earlier; at mmWave bands with electrically large arrays and a small number of dominant multipaths, the approximation can be made sufficiently close to the optimal precoder by using a finite number of antenna elements in the array [6]. The multipath sparsity restricts the feasible analog precoders $\mathbf{F}_{RF}$ to a set of array response vectors, and the baseband precoder optimization can be translated into a matrix

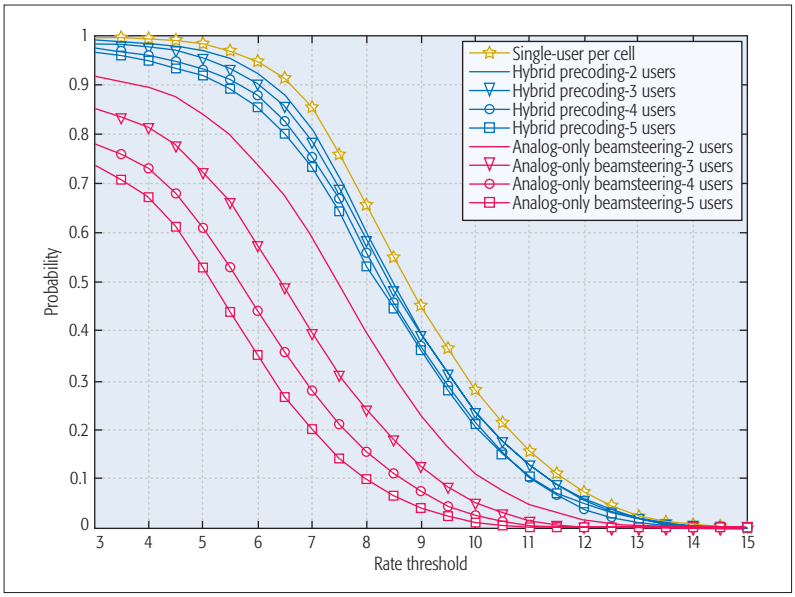

Figure 5. Comparison of achievable rates for hybrid precoding and analog-only beamsteering, from [14]. A single-path model is assumed between the BSs and UEs, and each link is assigned a line-of-sight or non-line-of-sight condition based on a blockage model, that is, the second reference in [14]. Each UE is associated to the BS with the least path loss and the BS randomly selects n = 2, ... ,5 associated UEs to be simultaneously served.

reconstruction with the cardinality constraint on the number of RF chains. The near-optimal solution of $\mathbf{F}_{BB}$ can then be found using sparse approximation techniques (e.g., OMP). The SE comparison of this method with unconstrained fully-digital beamforming and analog-only beamsteering with perfect transmit CSI is shown in Fig. 4.

While the general structure is the same as the one for cmWaves described earlier, in sparse mmWave channels, fast and greedy antenna subset selection [12] performs as robustly as exhaustive antenna search. Hybrid antenna selection can outperform a sparse hybrid combiner with coarsely quantized phase shifters in terms of power consumption when both have the same SE performance. There is still a large gap in SE between the hybrid combiner with switches and a fully digital one with ideal phase shifters.

Beam Selection: Another hybrid beamforming structure is based on continuous aperture phased (CAP)-MIMO transceivers. It uses a lens antenna instead of the phase shifters or switches for RF beamforming, and realizes the beamspace MIMO (B-MIMO) [13] similarly to the spatial DFT with selection discussed previously. An electrically large lens antenna is excited by a feed antenna array beneath the lens. The feed array is called a beam selector since the lens antenna produces high-gain beams that point at different angles depending on the feed antenna. The CAP-MIMO can efficiently utilize the low-dimensional high-gain beamspace of the sparse multipath channel by selecting a couple of feed antennas using a limited number of RF chains, like the spatially sparse precoding.

Hybrid Beamforming in mmWave MU Scenarios

Hybrid beamforming is also a promising solution for mmWave MU-MIMO systems. The hybrid structure at the BS can transmit multiplexed data streams to multiple UEs; each UE can be equipped with an antenna or an antenna array with fully analog beamforming. Figure 5 shows achievable rates of hybrid beamforming in MU multi-cell scenarios [14]. Consider UEs with a single RF chain and many antennas, which distributively select the strongest beam pair to construct analog beamformers. Thanks to the ZF digital precoding at the BS mitigating the inter-user interference, the hybrid structure significantly outperforms the analog beamsteering approach.

The hybrid beamforming based on beam selection and B-MIMO concept can also be extended to MU-MIMO systems with linear baseband precoders. While its effectiveness (compared to its full complexity counterparts) has been demonstrated in mmWave channels, many system and implementation aspects of hybrid beamforming in mmWave MU-MIMO systems, including multi-user scheduling, and 2D and 3D lens array design, are still open for further research.

Impact of Transceiver Imperfections

The presence of RF transceiver imperfections degrades SE in various ways. For example, it is harder to accurately generate desired transmit signals when higher beamformer gain is aimed for; nonlinear distortion at the RX depends on the instantaneous channel gain and hence the SNR. Due to the transceiver imperfections being more pronounced at mmWaves, the SE and SNR of hybrid precoder/combiners no longer scale well with the number of RF chains. Figure 4 compares the SE of spatially sparse hybrid precoding, including RF imperfections, to that from fully digital beamforming based on SVD. The aggregate impact of the transceiver imperfections is modeled as a Gaussian process. The coarsely quantized phase shifters and the transceivers' imperfections significantly degrade the SE. Knowledge of transceiver imperfections at mmWaves is essential for analyzing the scalability of the SE in the large MIMO regime.

Spectral-Energy Efficiency Trade-off

Finally, we discuss a relationship between EE and SE of hybrid beamforming structures at mmWaves based on [15]. The hybrid structure B in Fig. 1 was studied, where the BS uses a sub-array with N_{BS}/N_{RF}^{BS} antennas to serve each user individually. Figure 6 shows the EE-SE trade-off, indicating an optimal number of RF chains achieving the maximal EE for any given SE.

Conclusion

Hybrid beamforming techniques were invented more than 10 years ago, but have seen a dramatic uptick in interest in the past 3 years due to their importance in making massive MIMO systems cost- and energy-efficient. They use a combination of analog and digital beamforming to exploit the fine spatial resolution stemming from a large number of antenna elements, but keep the number of (expensive and energy-hungry) RF up/downconversion chains within reasonable limits. This article categorizes the hybrid beamforming according to:

- Amount of required CSI (instantaneous vs. average) for the analog beamformer part
- Complexity (full complexity, reduced complexity, and switched)

- Carrier frequency range (cmWave vs. mmWave, since both channel characteristics and RF impairments are different for those frequency ranges)

It is clear that there is no single structure/algorithm that provides the "best" trade-off between complexity and performance in all those categories, but rather that there is a need to adapt them to application and channel characteristics in every design.

Acknowledgment

The financial support of the Academy of Finland and the National Science Foundation through the WiFiUS project "Device-to-Device Communica tions at Millimeter-Wave Frequencies" is gratefully acknowledged.

References

[1] E. Larsson *et al.*, "Massive MIMO for Next Generation Wireless Systems," *IEEE Commun. Mag.*, vol. 52, no. 2, Feb. 2014, pp. 186–95.
[2] X. Zhang, A. Molisch, and S.-Y. Kung, "Variable-Phase-Shift-Based RF-Baseband Codesign for MIMO Antenna Selection," *IEEE Trans. Signal Processing*, vol. 53, no. 11, Nov. 2005, pp. 4091–4103.
[3] P. Sudarshan *et al.*, "Channel Statistics-Based RF Pre-Processing with Antenna Selection," *IEEE Trans. Wireless Commun.*, vol. 5, no. 12, Dec. 2006, pp. 3501–11.
[4] A. Alkhateeb , "MIMO Precoding and Combining Solutions for Millimeter-Wave Systems," *IEEE Commun. Mag.*, vol. 52, no. 12, Dec. 2014, pp. 122–31.
[5] A. Adhikary *et al.*, "Joint Spatial Division and Multiplexing — The Large-Scale Array Regime," *IEEE Trans. Info. Theory*, vol. 59, no. 10, Oct. 2013, pp. 6441–63.
[6] O. El Ayach *et al.*, "Spatially Sparse Precoding in Millimeter Wave MIMO Systems," *IEEE Trans. Wireless Commun.*, vol. 13, no. 3, Mar. 2014, pp. 1499–1513.
[7] L. Kong, S. Han, and C. Yang, "Wideband Hybrid Precoder for Massive MIMO Systems," *Proc. 3rd Global Conf. Signal and Info. Processing*, Orlando, FL, Dec. 2015, pp. 305–09.
[8] Z. Li, S. Han, and A. F. Molisch, "Hybrid Beamforming Design for Millimeter-Wave Multi-User Massive MIMO Downlink," *IEEE ICC '16*, Kuala Lumpur, Malaysia, May 2016.
[9] S. Hur *et al.*, "Millimeter Wave Beamforming for Wireless Backhaul and Access in Small Cell Networks," *IEEE Trans. Commun.*, vol. 61, no. 10, Oct. 2013, pp. 4391–4403.
[10] J. Choi, "Diversity Eigenbeamforming for Coded Signals," *IEEE Trans. Commun.*, vol. 56, no. 6, June 2008, pp. 1013–21.
[11] O. El Ayach *et al.*, "The Capacity Optimality of Beam Steering in Large Millimeter Wave MIMO Systems," *Proc. 13th Wksp. Signal Processing Advances in Wireless Commun.*, Cesme, Turkey, June 2012, pp. 100–04.
[12] R. Mendez-Rial et al., "Channel Estimation and Hybrid Combining for mmWave: Phase Shifters or Switches," *Proc. Info. Theory Appl. Wksp.*, San Diego, CA, Feb. 2015, pp. 90–97.
[13] J. Brady, N. Behdad, and A. M. Sayeed, "Beamspace MIMO for Millimeter-Wave Communications: System Architecture, Modeling, Analysis, and Measurements," *IEEE Trans. Antennas Propag.*, vol. 61, July 2013, pp. 3814–27.
[14] A. Alkhateeb, G. Leus, and R. Heath, "Limited Feedback Hybrid Precoding for Multi-User Millimeter Wave Systems," *IEEE Trans. Wireless Commun.*, vol. 14, no. 11, 2015, pp. 6481–94.
[15] S. Han, C.-L. I, Z. Xu, and C. Rowell, "Large-Scale Antenna Systems with Hybrid Analog and Digital Beamforming for Millimeter Wave 5G," *IEEE Commun. Mag.*, vol. 53, no. 1, Jan. 2015, pp. 186–94.

Biographies

Andreas F. Molisch [F'05] is the Solomon-Golomb — Andrew-and-Erna Viterbi Chair Professor at the University of Southern California (USC), Los Angeles. His research interests include wireless propagation, multi-antenna systems, ultrawideband communication and localization, and wireless video distribution. He has published 4 books, more than 500 journal and conference papers, and more than 80 patents. He is a Fellow of the National Academy of Inventors, AAAS, and IET, and a member of the Austrian Academy of Sciences.

Vishnu V. Ratnam received his B.Tech. degree from the Indian Institute of Technology Kharagpur, where he graduated as the Salutatorian for the class of 2012. He is currently pursuing a Ph.D. degree in electrical engineering at USC. His research interests include the design and analysis of low-complexity transceivers for large antenna and ultra-wideband systems, and resource allocation problems for multi-antenna networks. He is a recipient of the ICUWB 2016 Best Student Paper Award.

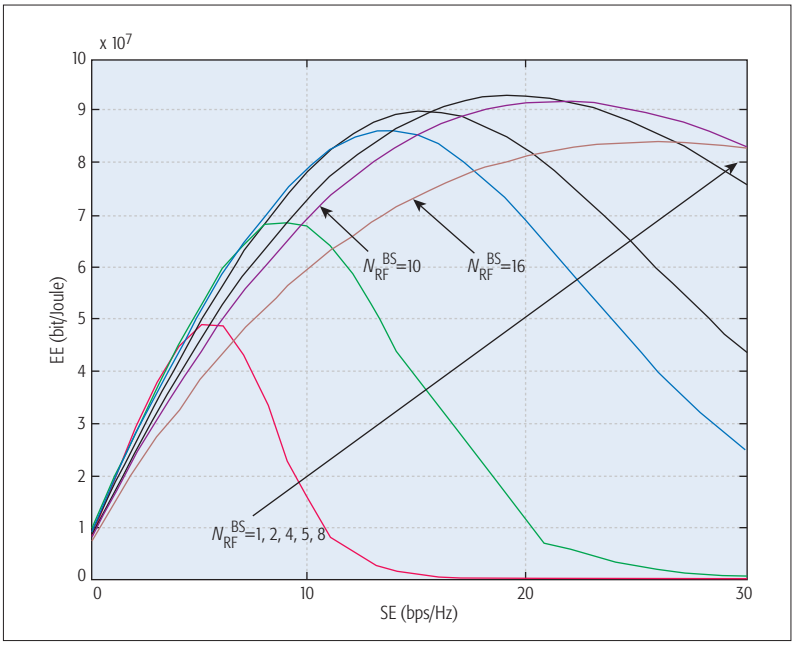

Figure 6. EE-SE relation of mm-wave massive MIMO system [15]. The hybrid transceiver follows Structure B of Fig. 1; N_{BS} = 800, system bandwidth is 200 MHz, noise power spectrum is 10^{-17} dBm/Hz, average channel gain is –100 dB, the efficiency of power amplifier is 0.375, the static power consumptions for each RF chain and each antenna are both 1 Watt, and the other fixed power consumption is 500 Watt.

Shengqian Han [S'05, M'12] received his B.S. and Ph.D. degrees from Beihang University, Beijing, China, in 2004 and 2010, respectively. He is currently a lecturer at Beihang University. From 2015 to 2016, he was a visiting scholar at USC. He is an Associate Editor for the *EURASIP Journal on Wireless Communications and Networking*. His research interests include full-duplex communications, multiple-input multiple-output systems, and energy-efficient transmission.

Zheda Li received his B.S. degree in communication engineering from Beijing University of Posts and Telecommunications, China, in 2010 and his M.S. degree (with honors) in electrical engineering from USC in 2012. He is currently pursuing a Ph.D. degree in the WiDeS group of USC. His research interests include massive multiple-input multiple-output systems and millimeter-wave communications.

Sinh Le Hong Nguyen [S'10, M'13] received his Ph.D. degree in electrical engineering from Concordia University, Canada, in 2013. From 2013 to 2014, he was a postdoctoral fellow at McGill University, Canada. He is now with Aalto University, Finland, as a postdoctoral researcher. He has participated in several industrial collaborative and EU-funded projects on 5G technologies. His current research interests include channel modeling and signal processing for mobile communications, compressive sensing, large-scale networks, and the Internet of Things.

Linsheng Li [S'09, M'14] received his Ph.D. degree in electrical engineering from Southeast University, Nanjing, China, in 2014. From September 2014 to July 2016, he was a research engineer with the Department of Radio Science and Engineering at Aalto University. Since August 2016, he has been an antenna engineer with the Research Center of Huawei, Helsinki. His main research interests include sub-6 GHz and millimeter-wave antennas for 5G and future wireless communication.

Katsuyuki Haneda [S'03, M'07] is an associate professor at Aalto University. He is the recipient of the best paper awards at IEEE VTC-Spring 2013 and EuCAP2013 as a first author, and at a number of international conferences and journals as a co-author. He has been an Editor of *IEEE Transactions on Antennas and Propagation* and *IEEE Wireless Communications*. His research covers wireless and electromagnetic design and modeling for cellular, medical, and emergency applications.

ACCEPTED FROM OPEN CALL

In-Flight Broadband Connectivity: Architectures and Business Models for High Capacity Air-to-Ground Communications

Ergin Dinc, Michal Vondra, Sandra Hofmann, Dominic Schupke, Mikael Prytz, Sergio Bovelli, Magnus Frodigh, Jens Zander, and Cicek Cavdar

The authors investigate A2G architectures in terms of economic and technical perspectives, and propose business models by identifying new roles and positioning them in the A2G business ecosystem. In addition, they provide an extensive summary of the state-of-the-art and future improvements for A2G communications.

ABSTRACT

In-Flight Broadband Connectivity (IFBC) is a significant open market for mobile network operators, considering that more than 3.3 billion passengers were served by airlines in 2015. On-board broadband services are provided via air-to-ground (A2G) connectivity through direct A2G communication (DA2GC) and satellite A2G communication (SA2GC). Available on-board connectivity systems have significant limitations: high latency in SA2GC and low capacity in DA2GC. The customer expectancy is multi-Mb/s connections in every seat, which leads to capacity requirements of Gb/s to the aircraft. Creation of high capacity IFBC requires a collaborative interaction between different industry partners. For this reason, we investigate A2G architectures in terms of economic and technical perspectives, and propose business models by identifying new roles and positioning them in the A2G business ecosystem. In addition, we provide an extensive summary of the state-of-the-art and future improvements for A2G communications.

INTRODUCTION

Today, users demand high speed broadband connectivity regardless of their location and time. To this end, in-flight connectivity has recently attracted significant research attention from both industry and academia. While passengers tend to use their own devices, and expect to directly access the Internet at high performance, in-flight broadband connectivity (IFBC) solutions are only partially able the meet this demand. Hence, IFBC creates large-scale market opportunities for the mobile network industry, considering more than 3.3 billion passengers were served in 2015 [1].

Some airlines are currently offering on-board WiFi services based on satellites. Satellite A2G communications (SA2GC) is a natural choice for transcontinental flights. However, satellite connection is not a long-term solution for the in-flight connectivity market due to long transmission latencies. On the other hand, continental flights have a significant share of the airlines market. More than 800 million passengers travelled within Europe in 2015 [1]. Therefore, direct A2G communications (DA2GC) has a growing customer base. The main advantage of next-generation DA2GC will be a new LTE service in the cabin, easy login and high-sustaining bit rates. With current satellite-based solutions, passengers are required to connect to WiFi for in-flight connectivity. With DA2GC, users can maintain their cellular connection (LTE, and 5G in the future) without any connection break. DA2GC is the only alternative to provide applications with quality-of-service (QoS) requirements such as video call, streaming and phone calls, due to latency problems of SA2GC. Although DA2GC ground stations can be placed in petroleum platforms and islands, DA2GC will have limited transcontinental coverage. Hence, a full-scale A2G connectivity solution requires a hybrid network via DA2GC and SA2GC, as shown in Fig. 1.

Gogo Inc. has already deployed more than 200 DA2GC ground stations across the US and Canada based on CDMA2000 [2]. However, this service has low data rates due to bandwidth limitations (up to 9.8 Mb/s/cell). In addition, Deutsche Telekom and Inmarsat are deploying the European Aviation Network (EAN) by installing 300 ground stations in Europe to provide A2G connectivity up to 75 Mb/s/cell [3]. LTE-based trials in Europe typically can achieve 26-30 Mb/s average data rate in the forward link (ground-to-aircraft) [4, 5]. Since customers expect multi-Mb/s on-board connection, IFBC systems require Gb/s links to the aircraft [6]. To provide these data rate levels, DA2GC requires more spectrum, increased spectral efficiency, and improvements by communication techniques as provided in 5G, which is discussed later.

The IFBC market also has significant challenges in terms of business modeling [5]. To provide IFBC in the European airspace, at least 48 states (including non-EU states) with different frequency regulations need to participate. Thus, multiple operators in different countries are required to work together to provide on-board connectivity via DA2GC. In [5], some initial sketches of business models are proposed, but not analyzed in any great depth. The authors distinguish between "ecosystem-models," with a multitude of business players that interact to provide passenger service, and "all-in-one"-type models, where a single player dominates the provisioning of the service. This could either be an operator providing connectivity or an in-flight entertainment provider. We see

Digital Object Identifier:
10.1109/MCOM.2017.1601181

Ergin Dinc, Michal Vondra, Jens Zander and Cicek Cavdar are with KTH Royal Institute of Technology; Sandra Hofmann and Dominic Schupke are with Airbus; Mikael Prytz and Magnus Frodigh are with Ericsson Research; Sergio Bovelli is with Airbus Defence and Space.

0163-6804/17/$25.00 © 2017 IEEE

it as less likely that a single player will be able to dominate the service, and therefore this article focuses on the ecosystem-type business models.

Figure 2 shows the business canvas for the A2G operator. The A2G operator is an intermediary between the ground network and the passenger network through terrestrial and satellite operators. The business canvas is highly utilized in describing, analyzing, developing and revising business models [7]. As illustrated in the business canvas for A2G operators, multiple business activities will be shared among the key players: airline, content provider, cabin system operator, passengers' home operator, terrestrial operator and satellite operator. In the IFBC market, the airline and the home operator are the front-end players that are directly in touch with the customers. Hence, the front-end players will be responsible for customer relationship management, represented with the green color in Fig. 2. Technical responsibilities (the red part of the business canvas) will be shared among the back-end-players, i.e., the operators. The cabin system operator will provide WiFi and LTE connections in th eaircraft. Connectivity with the ground network will be provided by an intermediary A2G operator through terrestrial and satellite operators. The content provider will provide content for the passengers. Eventually, the network will create value for passengers, i.e., in-flight entertainment, represented by the blue region in Fig. 2. Business models aim to minimize costs such as infrastructure and frequency spectrum, and maximize revenues via higher ticket price and on-board fee.

The main contributions of this article are threefold. First, we provide an extensive survey of the available A2G systems and their future. Second, the players and their activities in the business-ecosystem are defined. In addition, we propose A2G architectures and investigate their feasibility. Finally, three business models are proposed to manage the value and money exchange between the players. For the ecosystem-type business models, it is critical for companies to find a value and revenue generating business model. Hence, we investigate the business network while including both technical and business perspectives.

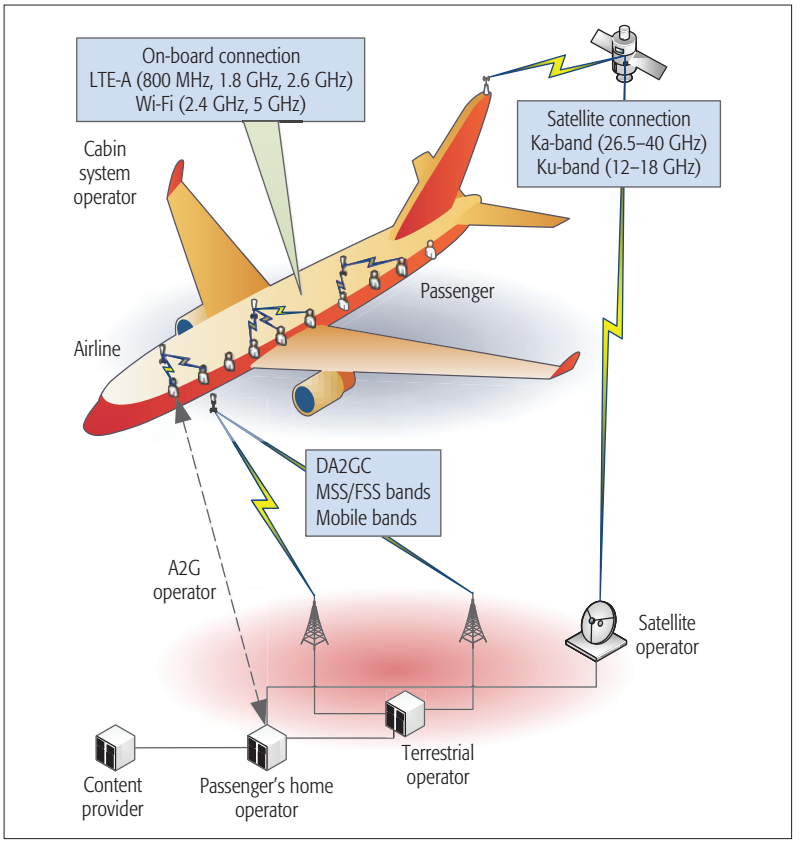

Figure 1. A2G communication chain.

The A2G Market: Today and Future

In the following subsections we summarize the current and future technologies for SA2GC and DA2GC.

SA2GC

Almost all commercially available A2G systems utilize satellite-based solutions to provide IFBC. Connection with satellites is provided with an antenna placed on top of the aircraft. SA2GC operators provide WiFi connectivity for passengers on-board, and use different business models.

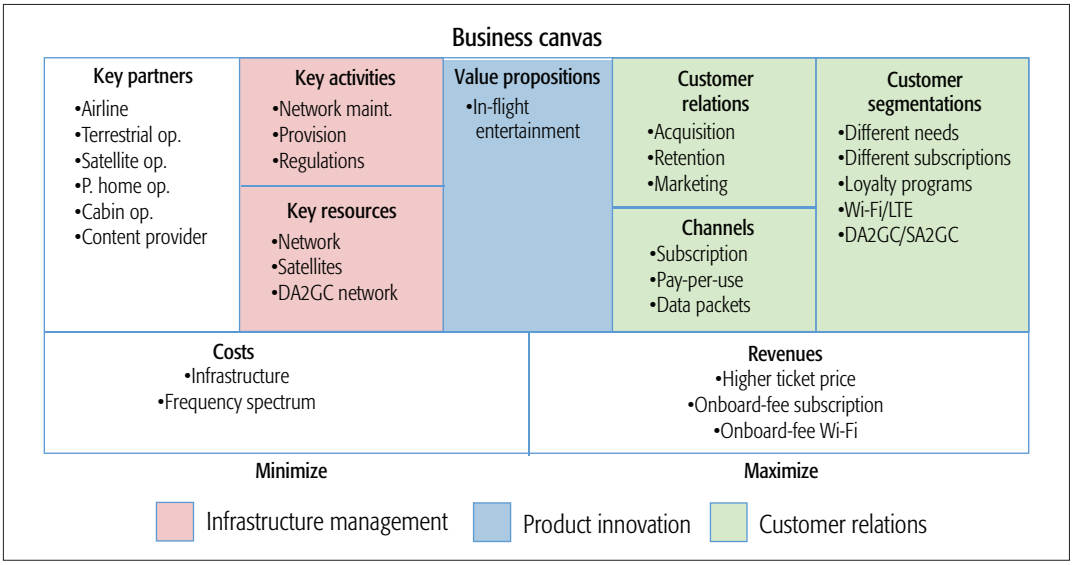

Figure 2. Business canvas for A2G operator.

Spectrum sharing with mobile satellite services as complementary ground component and fixed satellite services as moving platforms may be promising for DA2GC. The FCC in the US also considers possible frequency sharing between DA2GC and FSS in 14–14.5 GHz band. To summarize, the discussion about the DA2GC frequency spectrum is an important open issue.

Parameter	Requirements
Data rate	Download: 15 Mb/s/active user Upload: 7.5 Mb/s/active user
Latency	10 ms
Mobility	1000 km/h (max.)
Aircraft density	60/18000 km^2
Traffic density	Download: 1.2 Gb/s/aircraft Upload: 600 Mb/s/aircraft

Table 1. NGMN's 5G DA2GC KPI requirements [6].

Some airlines prefer to offer the service for free to acquire more customers. Others are charging an additional fee for the service. Some SA2GC operators offer subscriptions and limited data plans. Despite the market penetration, current SA2GC via geostationary-orbit (GEO) satellites has limitations in transmission latency around 500 ms (round-trip-time (RTT)).

GEO satellite operators generally uses Ku-band satellites due to their availability and wide coverage. There are several players who utilize Ku-band satellites for SA2GC such as Gogo-2Ku [8] and Panasonic [9]. SA2GC with Ku-band can provide capacity levels up to 70 Mb/s per aircraft, and with high-throughput-satellites (HTS) the achievable data rate can reach up to 100 Mb/s with frequency re-use and spot-beam technologies [8, 10]. Broadband low-earth-orbit (LEO) satellite initiatives, e.g., OneWeb, could be an alternative solution for low latency and high capacity A2G connectivity with their close Earth orbit (≈ 1200–1500 km) [11]. However, the first LEO broadband satellite system will not be operational before 2022. Thus, we envision that IFBC will be provided by SA2GC via GEO satellites and, where possible, by DA2GC via ground base stations in near future.

DA2GC

DA2GC utilizes ground base stations to connect the aircraft and the ground network. This way, latency problems of on-board broadband services can be alleviated because cell range will be between 50–100 km based on inter-site distance (ISD), and 5–10 ms RTTs can be achieved [6]. Compared with GEO (36000 km and 500 ms RTT) and LEO satellites (1500 km and 30 ms RTT [11]), DA2GC provides significant improvement in latency, and enables IFBC to offer applications with QoS requirements such as video calls. However, for seamless connection in transcontinental flights, DA2GC and SA2GC will complement each other such that SA2GC will provide connectivity, where DA2GC is not available or too congested.

In the DA2GC market, the most significant player is Gogo Inc. with its ATG-4 product that operates at 850 MHz with 4 MHz bandwidth based on the EV-DO CDMA2000 standard [2]. Gogo ATG-4 can achieve up to 9.8 Mb/s per cell and provides on-board connectivity for flights in North America with more than 200 ground base stations. However, Gogo suffers from low bandwidth levels, so the company is moving to satellite-based solutions to provide high data rate levels. Deutsche Telekom and Inmarsat are deploying EAN, which is a hybrid SA2GC/DA2GC connectivity by using S-band frequencies (2x15 MHz) [3]. EAN will initially have 300 ground stations to provide DA2GC coverage for Europe, and provide up to 75 Mb/s per cell. This capacity will be shared by the number of aircraft in the cell, and then the resulting capacity per aircraft is shared by the passengers on-board. Chinese Government entities are also performing tests at 1785-1805 MHz (20 MHz bandwidth) for TD-LTE technology, and providing coverage with more than 17 base stations in China's air routes with the CDMA EV-DO standard [4].

For the IFBC market, European airspace is an important open market currently serving more than 800 million passengers per year [1, 4]. However, it is also a challenging environment due to the different regulations by different countries. For this reason, a report [4] was published on the frequency regulations and company trials for broadband DA2GC services in Europe. Based on this report, Deutsche Telekom, Nokia and Airbus have tested LTE-based ground stations having 100 km ISD [4, 5]. According to their results, an A2G link at 2.6 GHz (bandwidth 2 × 10 MHz) typically provides 26–30 Mb/s in the forward link (ground-to-aircraft) and 17 Mb/s in the reverse link (aircraft-to-ground) with less than 60 ms latency for an aircraft at 10 km with 800 km/h speed. However, DA2GC requires increased spectrum resources to provide high achievable data rates to be an alternative solution for SA2GC.

Frequency Regulations for DA2GC: An ECC report [4] describes the frequency designation discussions and possible regulations to make use of the current spectrum in Europe. To solve the bandwidth problem, spectrum repurposing/transferring is also proposed by the ECC. For DA2GC at 5855–5875 MHz and 1900–1920 MHz, there are some regulatory efforts to provide harmonization in European states [12, 13]. However, these bandwidths cannot provide the data rates that can be an alternative solution for the SA2GC currently having 70–100 Mb/s. Thus, spectrum sharing with mobile satellite services (MSS) as a complementary ground component and fixed satellite services (FSS) as moving platforms may be promising for DA2GC. The FCC in the US is also considering possible frequency sharing between DA2GC and FSS in the 14–14.5 GHz band [4]. To summarize, the discussion about the DA2GC frequency spectrum is an important open issue.

Toward 5G: The Next Generation Mobile Networks (NGMN) Alliance proposed the key performance indicators (KPI) for future (2020+) 5G IFBC [6]. Based on their estimations, each user will have 15/(7.5) Mb/s download/(upload) speeds on average, so that 1.2/(0.6) Gb/s download/(upload) speed is required per aircraft with the assumption of 20 percent active users per aircraft and 400 passengers in each aircraft. To achieve these data rates, DA2GC systems require increased spectrum, increased spectral efficiency and improved network management with 5G technologies (Table 1).

Millimeter wave (mmWave) frequencies have been attracting significant research atten-

tion due to the available amount of bandwidth of ≈500 MHz and more. With low wavelengths of mmWave, large antenna arrays can be realized to provide high array gains to compensate for the high path losses. Large antenna arrays also enable advanced antenna techniques such as multi-user beamforming and interference cancellation. Hence, high spectral efficiencies can be maintained with mmWave systems by modulation schemes such as 256QAM, 1024QAM, and 4096QAM (8, 10, and 12 bits/symbol, respectively). However, for utilizing mmWave frequencies, a feasibility analysis for DA2GC is required considering, for example, the effects of rain and atmospheric attenuations.

5G will provide advanced network coordination techniques such as cooperation of terrestrial and satellite networks, advanced resource allocation, mobile edge cloud and virtualization. Some content, e.g., a live football match, can be multicast to passengers. In addition, the cabin operator may utilize smart caching techniques to offload some traffic from A2G links by edge cloud functionalities. Network virtualization/slicing techniques will enable onboard IoT services for non-critical applications such as cargo monitoring, CCTVs and temperature monitoring.

Open Problems and Challenges for DA2GC: The most critical challenge in DA2GC are the frequency regulations as outlined above. Depending on the regulations, an aircraft may need to utilize multiple bands and roam between terrestrial networks and satellite networks to provide seamless connectivity. In addition, the DA2GC ground station deployment problem imposes a significant open research issue to reduce the cost of providing IFBC. To this end, DA2GC ground stations can be placed in existing base station towers to utilize existing fiber and grid infrastructure. Furthermore, flight corridors can be exploited to reach a cost-effective deployment for DA2GC. The interference between ground and in-cabin LTE networks is another open research problem. Both networks may experience high interference when the aircraft is flying close to the ground (especially <3000 m). For this reason, the in-cabin network cannot use the licensed ground LTE spectrum for low altitudes. To avoid this problem, the cabin system operator can utilize license assisted access (LAA) based LTE standards to provide seamless cellular connectivity in all phases of the flight.

Business Modeling: The Players

We envision that the future IFBC will be provided by the cooperation of different players in a business eco-system. Thus, this section includes all players and definitions of their roles. However, some players can combine multiple roles in the chain. The economical relationships between these players are covered later.

Passenger

The main purpose of the A2G chain is to provide IFBC for passengers. Passengers may be charged for this service by the airline via higher ticket prices and/or their home operator via subscription/pay-per-use. The IFBC market has a growing customer base with more than 3.3 billion passengers worldwide and 800 million passengers in Europe in 2015 [1].

Airline

The airline is not a direct player in the technical part of the A2G business. The role of the airline in the A2G chain is to provide hosting for the cabin system operator's equipment: DA2GC/SA2GC antennas and in-cabin network equipment. However, the airline is a front-end player in the market, thus they will take advantage of the service by charging higher ticket prices and/or acquiring more customers via new service offerings. In Europe, there are currently 387 airlines with 6,586 aircraft in service and 7,560,360 flights in 2015 [1].

Cabin System Operator

Management of the in-cabin network will be performed by the cabin system operator. The cabin system operator will provide two types of services: WiFi for non-QoS-guaranteed services, and LTE for QoS-guaranteed services and operator services. Any additional price will be charged by the cabin system operator such as WiFi only services for non-SIM devices. The cabin system operator is also a customer of the terrestrial and satellite operators who buys SA2GC and DA2GC services, respectively. This way, the cabin system operator will work with multiple terrestrial operators located in different countries. For these reasons, the cabin system operator becomes a new player in the A2G chain unlike the available in-flight Internet services (e.g., Gogo), where the terrestrial operator also performs as the cabin system operator.

Satellite and Terrestrial Operators

The satellite and terrestrial operators provide backhaul connection between the cabin system operator and the ground network. The satellite operator provides A2G connectivity for non-QoS applications via satellites and connectivity to the evolved packet core (EPC) of the passengers' home operator. The terrestrial operator provides DA2G connectivity and backbone connectivity to the EPCs of the passengers' home operator and A2G operator for the applications with QoS requirements. For DA2GC coverage in Europe, approximately 1300 and 320 ground stations are required for 100 km and 200 km ISDs, respectively. (This calculation is based on dividing the European continent area into circular cell areas.) DA2GC requires increased spectrum and spectral efficiency with 5G to provide high sustaining bit rates. In the business modeling, we assume that DA2GC can provide high sustaining bit rates and be utilized as the main A2G channel for continental flights.

Passengers' Home Operator

The passengers' home operator provides on-board connectivity services via the A2G chain. The A2G operator can be considered as a roaming partner for the passengers' home operator, and the connection between passengers and their home operator as a tunnel connection. Therefore, the passengers' home EPC provides home subscriber server and authorization-authentication-accounting. The passengers' home operator is a front-end player and is directly in touch with the end-users. The home operator will offer on-board subscriptions and pay-per-use deals to their customers. These services increase the connectivity

The main purpose of the A2G chain is to provide IFBC for passengers. Passengers may be charged for this service by the airline via higher ticket prices and/or their home operator via subscription/pay-per-use. The IFBC market has a growing customer base with more than 3.3 billion passengers worldwide and 800 million passengers in Europe in 2015.

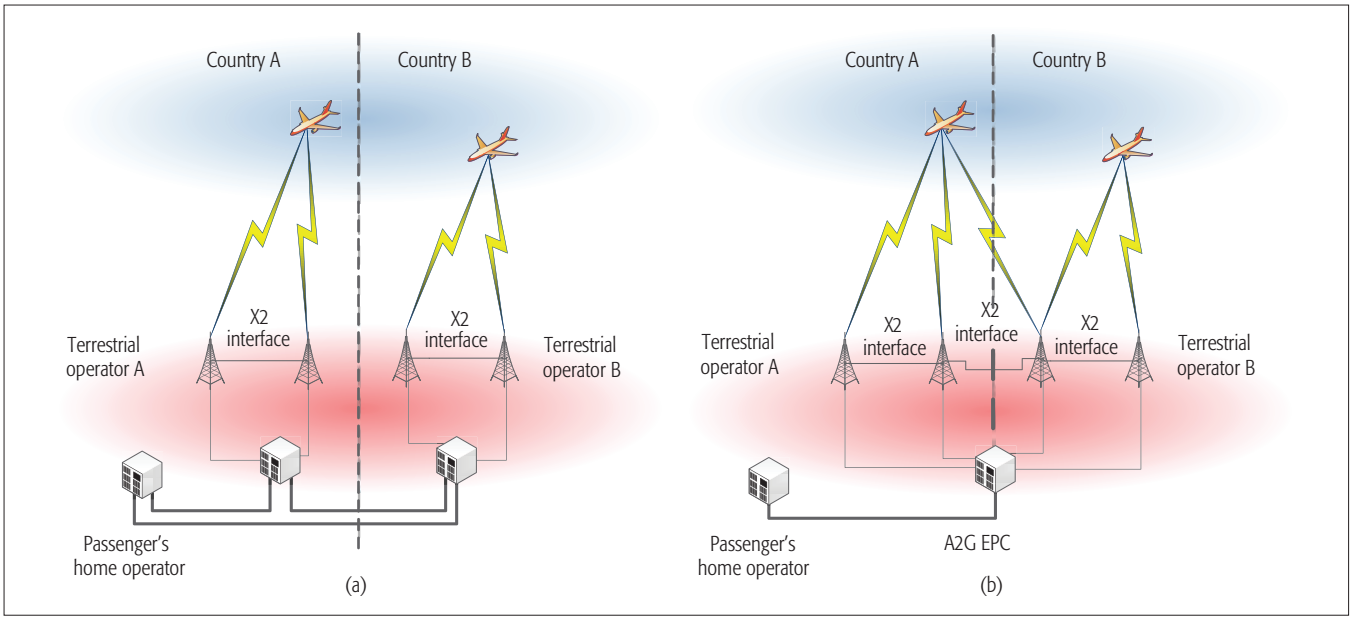

Figure 3. A2G architectures: a) business entity; b) one A2G EPC.

time of subscribers, and they are more expensive than connectivity on the ground. Thus, IFBC will improve the home operator's average revenue per-user (ARPU). The home operator can also utilize in-flight connectivity service for advertisement campaigns such as "services even in the sky."

The passengers' home operator can also act as a terrestrial operator by deploying DA2GC ground stations. Since the home operator has nationwide coverage in their country of operation, capital expenditures (CAPEX) of DA2GC may be lower by using their existing network. Mobile network operators (MNO) having IFBC will have a competitive advantage in the market; thus, this competition will push all MNOs to join the A2G eco-system to increase their revenues by providing service and/or becoming a terrestrial operator.

A2G Operator

The A2G operator is an entity that manages the A2G connection for the cabin system operator via terrestrial and satellite operators depending on type of data traffic and location of the aircraft. It is a consortium of all involved terrestrial and satellite operators. The cabin system operator and the passengers' home operator can optionally be part of the consortium. The A2G operator acts like a virtual operator/customer, who buys services and capacity (radio/backhaul) from satellite and terrestrial operators. The cabin system operator acts like a virtual operator/customer of the A2G operator. The A2G operator is a roaming partner for the passengers' home operator.

The A2G consortium is required for several reasons. Considering Europe, an aircraft will pass through multiple countries' airspace, and each country has different frequency regulations, different home operators (more than 100 MNOs) and different terrestrial operators. Therefore, the cabin system operator and home operator need to make tens of separate agreements with terrestrial operators in every country. This condition will create challenges for newcomers trying to enter the business. To avoid such problems, one unified contact point for all partners can be realized with the A2G consortium. This way, different operators can handle the frequency regulations in their countries, and the new home and cabin system operator can enter the market with an agreement to all partners through the A2G operator.

There are different possibilities for the A2G architecture in terms of the role of the A2G operator.

Business Entity: The A2G operator can be assumed to be a business entity, and its only role is to manage the interaction between terrestrial operators in different countries. As shown in Fig. 3a, the A2G operator will not own any network equipment, and different terrestrial operators will be connected with the home operator's EPC through different links. In this architecture, the A2G operator will operate as a clearing house in which the interaction between home operators and multiple terrestrial operators will be managed through a single contract. However, this architecture significantly limits the capabilities of the system due to the limited possibility of advanced cooperation between terrestrial networks.

One A2G EPC: In this case, the A2G operator will own the network infrastructure. Terrestrial networks in different countries belonging to different companies will be connected to a shared A2G EPC, as shown in Fig. 3b. This way, different terrestrial networks can employ advanced cooperation techniques such as seamless connection through the borders, efficient scheduling and resource allocation, and coordinated multi-point techniques. One A2G EPC will also facilitate newcomers to enter the market because there is no need to build a new network for ground communication. A new terrestrial operator can utilize the existing communication networks; thus, this architecture provides low CAPEX. In this architecture, passengers will be connected to the ground network through the A2G operator's packet data network gateway. Therefore, there will be policy exchange between the A2G operator and the home operator through the home and visited network's policy and charging rules functions.

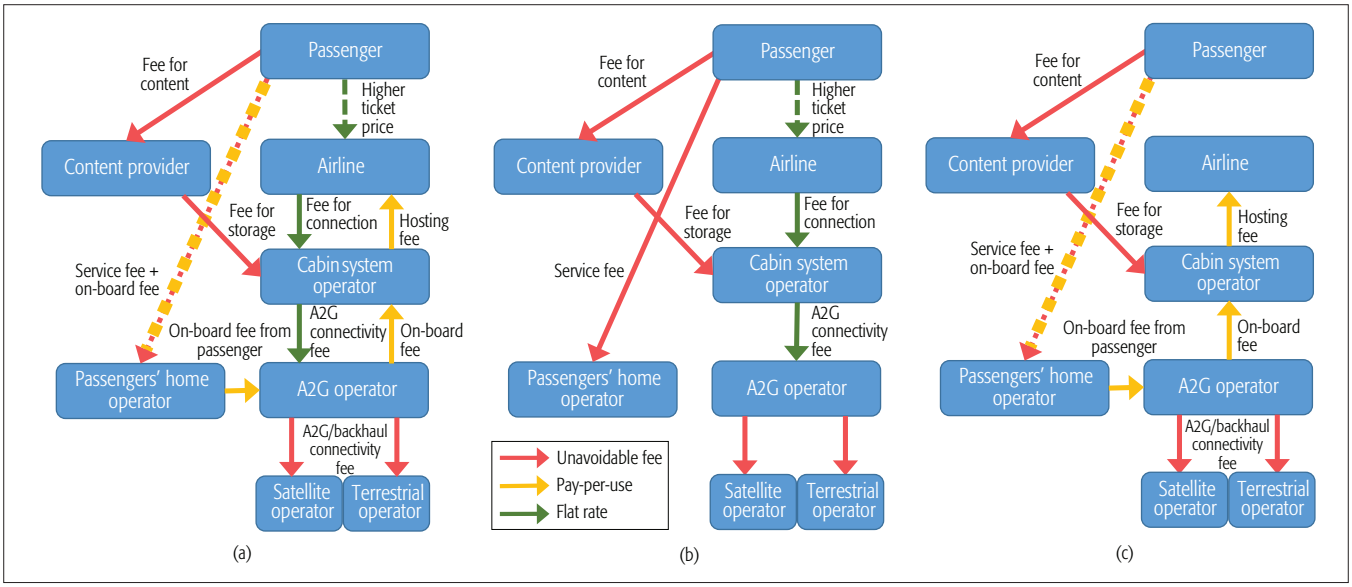

Figure 4. Proposed business models: a) cash flow; b) free services; c) on-board fee.

New Network: The last architecture is building a completely new network that is owned by a single A2G operator acting as both a terrestrial and cabin system operator, as in the Gogo model. However, this structure requires extremely high investment for a single organization, and it is economically ineffective.

One A2G EPC architecture will be promising for the A2G chain because it is economically effective compared to the new network architecture, and provides superior performance compared to the business entity architecture through advanced cooperation techniques.

Content Provider

The content provider offers content for passengers such as movies and music. Special offers can be provided for the passengers through special agreements with the cabin system operator and the passengers' home operator. In addition, the content provider may offer tailored content available off-line that will be stored in the cabin system operator's equipment for a storage fee. With A2G communication, the content provider can increase their revenues since the time users spend online will increase.

A2G Business Models

The A2G market is a collective business ecosystem consisting of many players, and creates value through interactions instead of stand-alone strategies [14, 15]. The proposed value comes from the IFBC, where passengers can use their own devices and reach content in the Internet. In this section, we propose and investigate three business models to analyze the value and cash flow among the players.

Cash Flow

In this business model, every service is charged by its provider, as shown in Fig. 4a, and this model is called "cash flow." Passengers may be charged by their home operator via subscriptions or pay-per-use deals, by airlines via higher ticket prices, and by content providers via content fees. The passengers' home operator receives service and on-board fees from passengers, and pays fees for the extended coverage to its roaming partner, i.e., the A2G operator. Since the creation of this service requires a new network with its own CAPEX, the home operator can charge an on-board fee for their extended coverage. Hence, this on-board fee is not a roaming fee that will be abolished in EU states in 2017.

The airline can charge passengers higher ticket prices and can charge the cabin system operator to host equipment because the extra weight in the aircraft will increase the costs of flights. However, the airline will also pay the cabin system operator for the connectivity. The content provider receives a content fee from the passengers. On the other hand, the content provider will pay a storage fee to the cabin system operator for content available off-line that is stored in the cabin system operator's equipment. The cabin system operator receives a connection fee from the airline, but they will pay for both A2G connection and equipment hosting. In addition, the cabin system operator may also provide WiFi services for passengers, especially for non-SIM devices, but this flow is omitted for the sake of simplicity. The A2G operator receives fees from the home operator and the cabin system operator for the connectivity, but pays for the A2G and backhaul connection, and the on-board fee to the cabin system operator that comes from the users. The terrestrial and satellite operators receive a connectivity fee from the A2G operator, but they pay for the infrastructure and frequency spectrum.

Figure 4a shows the cash flow business model. Red arrows represent fees that are always present whether the service is exploited or not. Since the resources of the satellite and terrestrial operators are allocated for the A2G operator, the connectivity fee is unavoidable. The content provider and the home operator will charge customers for their subscriptions. However, on-board connectivity depends on whether users exploit the service or not. Therefore, the on-board fee charged by the home operator to the passengers is pay-per-use and represented with a yellow arrow. In the same way, the hosting fee and the on-board fee to the

Full scale A2G communication requires a hybrid solution based on SA2GC and DA2GC. SA2GC will provide transcontinental connectivity, and DA2GC will provide applications with QoS requirements in continental flights. In near future, the coexistence of A2G connectivity via LEO satellites and 5G mmWave frequencies has very high potential to meet the latency and data rate requirements of the IFBC.

A2G operator and cabin system operator also depend on the amount of data transferred. Green arrows represent the price charged to the airline and have flat rates such as a fee for connection charged by the cabin system operator and the A2G connectivity fee charged by the A2G operator. The dashed green arrow between the passenger represents the price that may not be exploited because some airlines may try to attract more customers by offering this service for free.

Free Services

The second business model is "free services" in which some of the services are provided for free for passengers, as shown in Fig. 4b. In this model, the price for in-flight connectivity is free for passengers; however, the cost of the service may be reflected in the ticket price by the airline. With this model the airline can attract more passengers and increase their customer base. In the free services model, the airline is the entity that distributes the income to the other partners. The airline pays a connectivity fee to the cabin system operator. The cabin system operator pays an A2G connectivity fee to the A2G operator, and the A2G operator pays the satellite and terrestrial operators. The passengers still should pay a fee to their home operator and content provider for their normal subscriptions. The free services model is especially promising for big airlines to promote their brands. Since this service introduces a new cost to the airline, this model would be undesirable for low-cost airlines. Some low-cost airlines may still use this model, and compensate the cost of this service by the increase in the number of passengers without charging higher ticket price.

On-Board Fee

The primary objective of low-cost airlines is to provide the lowest possible price for plane tickets, and the market for low-cost airlines is highly competitive. In the free services and cash flow models, the airline will likely compensate the costs of the service by charging higher ticket prices. For these reasons, we propose the "on-board fee" business model in which airlines do not charge any fee for in-flight connectivity, as shown in Fig. 4c. Rather, this service is sold through the passengers' home operator via subscriptions or pay-per-use deals. With this model, the airline can keep their costs the same while offering in-flight connectivity. In this model, the revenue is collected by the home operator, and the fees are pay-per-use depending on the utilization of the network. The home operator pays the on-board connectivity fee to the A2G operator. The cabin system operator charges the A2G operator for providing networking, and pays a hosting fee to the airline based on the amount of usage. This way, low-cost airlines can make income and offer a new service without adding an additional cost to their system.

Conclusion

In this article, we propose new architectures and ecosystem-type business models for the A2G operator. Since different companies have different goals, multiple business models will coexist in the market. Full-scale A2G communication requires a hybrid solution based on SA2GC and DA2GC. SA2GC will provide transcontinental connectivity, and DA2GC will provide applications with QoS requirements in continental flights. In the near future, the coexistence of A2G connectivity via LEO satellites and 5G mmWave frequencies has very high potential to meet the latency and data rate requirements of the IFBC.

Acknowledgment

This study is supported by EIT Digital ICARO-EU (Seamless Direct Air-to-Ground Communication in Europe) Activity.

References

[1] Air Transport Action Group, "Aviation: Benefits Beyond Borders," July 2016.
[2] Gogo Inc., ATG-4, "What is It, and How Does It Work," available, http://concourse.gogoair.com/technology/gogo-atg-4-work; accessed on July 19, 2017.
[3] Inmarsat, "The European Aviation Network," available: http://www.inmarsat.com/wp-content/uploads/2016/01/Inmarsat_European_aviation_network_April_2016_EN_Low-Res.pdf; accessed on July 19, 2017.
[4] CEPT ECC Report 214, 6 Jun. 2014. "Broadband Direct-Air-to-Ground Communications (DA2GC)."
[5] Alcatel Lucent, "Using Air-to-Ground LTE for In-Flight Ultra-Broadband," Strategic White Paper, 2015, available: http://www.tmcnet.com/tmc/whitepapers/documents/whitepapers/2015/11529-using-air-to-ground-lte-in-flight-ultra.pdf; sccessed on July 19, 2017.
[6] NGMN Alliance, "5G White Paper," Feb. 2015.
[7] A. Osterwalder and Y. Pigneur, *Business Model Generation: A Handbook for Visionaries, Game Changers, and Challengers*, John Wiley & Sons, 2010.
[8] A. Ghezzi, M. N. Cortimiglia, and A. German Frank, "Strategy and Business Model Design in Dynamic Telecommunications Industries: A Study on Italian Mobile Network Operators," *Technological Forecasting and Social Change*, vol. 90, Jan. 2015, pp. 346–54.
[9] A. Laya *et al.*, "Business Model as Relational Aggregator: Exploring Business Relationships," INDEK Working Paper Series 2016/14, Department of Industrial Economics and Management, Royal Institute of Technology, Jan. 2016.
[10] C. McLain *et al.*, "High Throughput Ku-Band Satellites for Aeronautical Applications," *Proc. MILCOM*, 2012, pp. 1–6.
[11] V. Velivela, "Small Satellite Constellations: The Promise of Internet for All," ORF Issue Brief, iss. 107, Sept. 2015.
[12] CEPT ECC Decision (15)03, 3 July 2015, "The Harmonised Use of Broadband Direct Air-to-Ground Communications (DA2GC) Systems in the Frequency Band 5855–5875 MHz."
[13] CEPT ECC Decision (15)02, 3 July 2015, "The Harmonised Use of Broadband Direct Air-to-Ground Communications (DA2GC) Systems in the Frequency Band 1900–1920 MHz."
[14] B. R. Elbert, White Paper, "Aeronautical Broadband for Commercial Aviation: Evaluating the 2Ku Solution, Application Technology Strategy," L.L.C., Oct. 2014.
[15] S. Panthi, C. McLain, and J. King, "The eXConnect Broadband Aero Service," *Proc. AIAA Int'l. Commun. Satellite Systems Conf.*, Oct. 14–17, 2013, Florence, Italy.

Biographies

Ergin Dinc (ergind@kth.se) received his B.Sc. degree in electrical and electronics engineering from Bogazici University, Istanbul, Turkey, in July 2012. He received his Ph.D. degree in electrical and electronics engineering from Koc University, Istanbul, Turkey in June 2016. After receiving his Ph.D., he started working as a postdoctoral researcher at KTH-The Royal Institute of Technology, Stockholm, Sweden. Currently, he is a research associate in molecular communications and nanonetworks at The University of Cambridge, Cambridge, UK. His research interests include direct-air-to-ground communication, beyond-line-of-sight communication, 5G wireless communication, nanonetworks and molecular communication.

Michal Vondra (mvondra@kth.se) received the B.Sc. degree in electronic engineering and M.Sc. and Ph.D. degrees in telecommunication engineering from the Czech Technical University in Prague in 2008, 2010, and 2015, respectively. Presently he is a postdoc researcher with the Department of Communication Systems, KTH Royal Institute of Technology (Wireless@KTH), Stockholm, Sweden. He has been actively involved in several national and international projects funded by the European Commission such as FREEDOM, or TROPIC. In 2014, he spent six months on

an internship with University College Dublin, Ireland. His area of research interest covers topics toward 5G mobile networks, such as mobility management in wireless networks, direct air-to-ground communication, and intelligent transportation system.

SANDRA HOFMANN (sandra.s.hofmann@airbus.com) received her M.Sc. degree in electrical engineering from Technische Universitat Munchen, Germany in 2016. In the same year she joined Airbus, Munich, Germany, where she is pursuing a Ph.D. degree in the area of wireless communications. Her current focus is on providing high capacity communication for aerial vehicles using 5G.

DOMINIC A. SCHUPKE (dominic.schupke@airbus.com) is with Airbus in Munich, Germany, working in the area of research and innovations for wireless communications. Prior to that he was with NSN, Siemens, and the Institute of Communication Networks at Munich University of Technology (TUM). He received his Dipl.-Ing. degree from RWTH Aachen in 1998 and his Dr.-Ing. degree from TUM in 2004. He has over 15 years experience in the area of communication networks, especially their design and optimization. Since April 2009 he has taught the course 'Network Planning' at TUM. He is the author/co-author of more than 120 journal and conference papers. His research interests include network architectures and protocols, routing, recovery methods, availability analysis, critical infrastructures, security, virtualization, network optimization, and network planning. He is a Senior Member of IEEE, and member of the IEEE Communications Society, VDE/ITG, and VDI.

MIKAEL PRYTZ (mikael.prytz@ericsson.com) is a research leader and head of the Network Control and Management unit at Ericsson Research, focusing on 4G and 5G mobile network architectures and protocols (radio and core), network control and automation, and end-to-end performance. He has more than 15 years of experience in fixed and mobile communications research and engineering at Ericsson's research and services units, as well as at international research programs, including EU projects Ambient Networks, E3, and QUASAR. He holds a Ph.D. in optimization and systems theory from the KTH Royal Institute of Technology in Stockholm, Sweden, and an MS in operations research from Stanford University in Palo Alto, in the US.

SERGIO BOVELLI (sergio.bovelli@airbus.com) received a laurea degree (M.S.) in electronic engineering and a Ph.D. degree in information engineering from the University of Perugia, Italy, in 2001 and 2006, respectively. His main activities are in the area of wireless communication, in particular for aeronautics and space applications. He was with Airbus Innovation Works, Munich, Germany, from 2005 to 2016. He is now responsible within Airbus for regulatory affairs. He served on the Technical Program Committee of multiple International Conferences, and is co-author of several patents for communication systems.

MAGNUS FRODIGH (magnus.frodigh@ericsson.com) is the research area director for network architecture and protocols at Ericsson Research, responsible for research in network architecture and protocols covering radio networks, transport networks and core networks including network management. He joined Ericsson in 1994 and has since held various key senior positions within Ericsson's Research & Development and Product Management focusing on 2G, 3G, 4G and 5G technologies. He holds a master of science degree from Linkoping University of Technology, Sweden and a Ph.D. in radio communication systems from the Royal Institute of Technology in Stockholm, Sweden. Since 2013 he has been an adjunct professor at the Royal Institute of Technology in wireless infrastructures.

JENS ZANDER (jenz@kth.se) received his M.Sc. and Ph.D. in electrical engineering from Linkoping University, Sweden in 1979 and 1985, respectively. Since 1989 he has been a full professor at the KTH Royal Institute of Technology in Stockholm, Sweden. In 2001 he became co-founder and scientific director of Wireless@KTH, the KTH Center for Wireless Systems. Since 2013 he has been the Dean of the KTH School of ICT. He has authored close to 300 scientific papers and several textbooks on radio communication and radio resource management. He is on the board of directors of the Swedish National Post and Telecom Agency (PTS), and he is a member of the Royal Academy of Engineering Sciences. He was the Chairman of the IEEE VT/COM Swedish Chapter (2001-2005) and General Chair for the Crowncom 2012 and IEEE DySPAN 2015 conferences. He is one of the organizers of the Johannesberg Summits on 5G and Future Wireless Systems. His current research interests include architectures, resource and flexible spectrum management regimes and economic models for future wireless infrastructures.

CICEK CAVDAR (cavdar@kth.se) is a senior researcher in the Communication Systems Department at the KTH Royal Institute of Technology. She has been leading a research group in the Radio Systems Lab composed of eight researchers focusing on the design and planning of intelligent network architectures, direct air to ground communications and IoT connectivity platforms. She finished her Ph.D studies in computer science at the University of California, Davis in 2008, and at Istanbul Technical University, Turkey in 2009. After receiving her Ph.D., she worked as an assistant professor in the Computer Engineering Department, Istanbul Technical University. She has chaired several workshops on green mobile broadband technologies and green 5G mobile networks last few years co-located with IEEE ICC and Globecom. She served as the chair of the Green Communication Systems and Networks Symposium at ICC 2017 in Paris. At the Wireless@KTH research center, she has been leading EU EIT Digital projects such as "5GrEEn: Towards Green 5G Mobile Networks" and "Seamless DA2GC in Europe." She is serving as the leader of the Swedish cluster for the EU Celtic Plus project SooGREEN "Service Oriented Optimization of Green Mobile Networks."

ACCEPTED FROM OPEN CALL

Big Data Enabled Mobile Network Design for 5G and Beyond

Shuangfeng Han, Chih-Lin I, Gang Li, Sen Wang, and Qi Sun

The authors propose a mobile network architecture enabled by big data analytics, which is capable of efficient resource orchestration, content distribution, and radio access network optimization. The protocol stack configuration at each access point and the processing optimization of each layer are presented. Key physical layer design including reference signals and frame structure are discussed.

ABSTRACT

Mobile communication networks are more and more characterized by the integration of distributed and centralized computing and storage resources. Big data capability thus available throughout such networks will not only deliver enhanced system performance, but also profoundly impact the design and standardization of the next-generation network architecture, protocol stack, signaling procedure, and physical-layer processing. In this article, a mobile network architecture enabled by big data analytics is proposed, which is capable of efficient resource orchestration, content distribution, and radio access network optimization. The protocol stack configuration at each access point and the processing optimization of each layer are presented. Key physical layer designs including reference signals and frame structure are discussed. Moreover, utilizing signals in the transform domains, such as delay, Doppler, and angle, may bring enlarged coherence time of the effective channels. It enables much simpler physical layer design, and effectively bridges the latency gap between big data cloud computing and real-time network optimization.

INTRODUCTION

This is the era of big data [1], which is widely recognized as one of the most powerful enablers to promote productivity, improve efficiency, and support innovation. Driven by the increasing demands of the mobile Internet and the Internet of Things, future wireless networks, e.g. 5G [2], are expected to provide a great variety of services in diversified scenarios. However, the correspondingly generated pervasive and exponentially increasing wireless data traffic in the anticipated ultra-dense 5G networks [3] poses imminent challenges to all aspects of wireless system design.

With powerful data acquisition platforms of operators and Over-The-Top (OTT) companies, almost all wireless data can be collected from users, the radio access network (RAN), the core network (CN), and service providers. With efficient storage management of the collected data, wireless big data analytics [4–6] will further process these data and help obtain the key features of each type of data, and label them with sufficiently fine granularity. In contrast to traditional data processing at the base stations, big data platforms are capable of both predictive and prescriptive analytics with powerful machine learning techniques, such as support vector machine and deep learning.

Wireless big data roughly includes four categories, as shown in Table1. The application data is about the features of wireless applications, e.g., service types and content popularity. User data is the profile of user behaviors, such as location, mobility, and preference. These two categories are mainly obtained via big data analytics, for example, analyzing the user plane packets. The user's location and mobility can also be obtained via the Global Positioning System (GPS) or network measurements. Note that big data analytics and traditional data processing may be combined. For example, to obtain information about service types, big data analytics may be needed to characterize each type beforehand, with the period of hours, days or weeks. Then for each new packet, the packet inspection at the CN or RAN side may well succeed in a real-time manner. Network data includes the network configurations such as the coordinates of each base station (BS), antenna number, bandwidth, key performance indicators such as traffic load and outage, user equipment (UE) capability, and various signaling interactions between the UE and the network. Link data is about the wireless links between the users and the BSs, which is generally obtained via downlink and uplink measurements.

Note that the above measurement based network and link data has long been used for traditional network optimization. However, this is confined to per radio link/user/cell optimization and simple inter-cell coordination, for example, link adaptation based on instantaneous channel information and inter cell interference cancellation (ICIC) based on statistical interferences. Minimization of drive test (MDT) measurements have been utilized to optimize network coverage, capacity and reliability, for example, for self-healing self-organizing networks with local cooperation architecture [7]. However, this may not be cost effective due to excessive time (money), human resources, and limited measurement results (generally confined to outdoor environments). Besides, storage and processing of the huge amount of data collected is very challenging to the traditional network. With a big data platform, the mobile network is capable of globally

Digital Object Identifier:
10.1109/MCOM.2017.1600911

The authors are with the China Mobile Research Institute.

Data category	Contents	Big data vs. traditional data
Application data	Content popularity, service type, and so on.	Big data analytics helps to obtain these data, which can be conveniently used for network optimization. Traditional network optimization generally did not use application and user data.
User data	User preference, location, mobility, user behavior, and so on	
Network data	Cell configurations, downlink and uplink signal strength, traffic load, outage rate, inter/intra-cell interference, signaling, UE capability, and so on	Big data analytics is able to analyze the network wide data and provides global optimization solutions. Traditional network optimization is generally confined to per radio link, per user, per cell optimization, or simple inter-cell coordination.
Link data	Physical channel information such as path loss, shadowing, channel statistics, and so on	

Table 1. Wireless big data.

optimized resource allocation and network operation based on the downlink and uplink measurements (not necessarily MDT results). Furthermore, big data analytics helps obtain application data and user data, which are generally unavailable and thus rarely used in traditional network optimization.

A framework for applying big data analytics in mobile cellular networks was presented in [5], where big data of signaling, traffic, location and radio waveforms were analyzed. The characteristics of wireless big data were investigated in [6], where big data-driven 5G network optimization was proposed. A signaling-based intelligent network optimization scheme was introduced in 4G Long Term Evolution (LTE) network according to [8]. The authors in [9] envisioned a big-data aware wireless network with better wireless service quality and new mobile applications. The work in [10] showed how big data analytics could help optimize wireless/wired edge caching to maximize energy efficiency and spectrum efficiency.

With wireless big data available to the central unit (CU), the operation of Cloud RAN is optimal, since it is feasible to globally optimize resource allocation and scheduling. However, as the antenna number and bandwidth become larger at the distributed units (DUs), especially in the millimeter wave (mmWave) bands, the burden on the fronthaul is tremendous if the quantized in-phase and quadrature signals are to be fed back to the CU for central processing. Also, a lower latency is highly motivated in some 5G use cases such as vehicle-to-vehicle communication, online gaming and virtual reality. A Cloud RAN structure with highly centralized processing is not efficient. The Next Generation Fronthaul Interface (NGFI) [11], which adopts a two-layer architecture (CU and DU) to conveniently support flexible CU/DU function split, is an elegant extension of Cloud RAN operation. Currently, the Third Generation Partnership Project (3GPP) is actively discussing and evaluating functional split options by moving full or partial baseband functions from the CU to the DU. It is anticipated that future network architecture may adopt this design methodology.

So far, the research interests regarding wireless big data are generally related to mobile network resource allocation and network operation. In this article, we strive to investigate the potential impacts of wireless big data on the design of network architecture, protocol stack, signaling procedure, and physical layer technologies, which are the fundamental elements of the wireless communication systems.

In the following section, a big data enabled wireless communication network architecture is presented, where the big data platform collects all the required data and outputs resource orchestration, content distribution strategy, and RAN optimization. Then we detail a big data based RAN stack configuration; stack processing optimization and signaling procedure are presented, which significantly reduce operation complexity and latency. Physical-layer procedure optimization is then discussed, which begins with an analysis of how big data enables low overhead reference signal design. Also, how transform domain data processing alleviates the necessity of adaptive scheduling with physical-layer processing is analyzed. The big data enabled flexible frame structure design is further presented. This article is summarized in the final section.

Network Architecture Enabled by Big Data

A big data enabled network architecture is shown in Fig. 1, including different layers from the access network through to the core network, which then connects to the Internet Protocol (IP) backbone network. The distributed RAN with integrated BSs coexists in this architecture with Cloud RAN, which has a flexible functional partition between the CU and the DU. Each network node, for example, the core network gateway (CN-GW), CU, DU, integrated BSs, and data-only BSs (with low data rate fronthaul to only provide local data service), can be equipped with proper storage and processing capabilities.

The big data platform is responsible for processing the large volume of data and providing useful information for resource optimization and RAN optimization. It can be deployed at either the CN side or the RAN side. For RAN optimization, the big data platform tends to be deployed at the RAN side, for example, in the CU. For CN optimization, the platform needs to be deployed at the CN. By leveraging data mining and data analytics, big data technology can help predict user mobility, traffic behavior, network loading fluctuation, channel variations, and link and system level interference. This enables efficient resource assignment, flexible network capability distribution, flexible protocol stack configuration and optimization at each node, signaling proce-

When the traffic is low or the traffic trend is of low predictability, it is more attractive to place content storage and computing at a higher level in a cost-efficient way. However, when the traffic trend can be accurately predicted, it is beneficial to improve user's experience by moving the corresponding content storage and computing functions to network nodes that are closer to users.

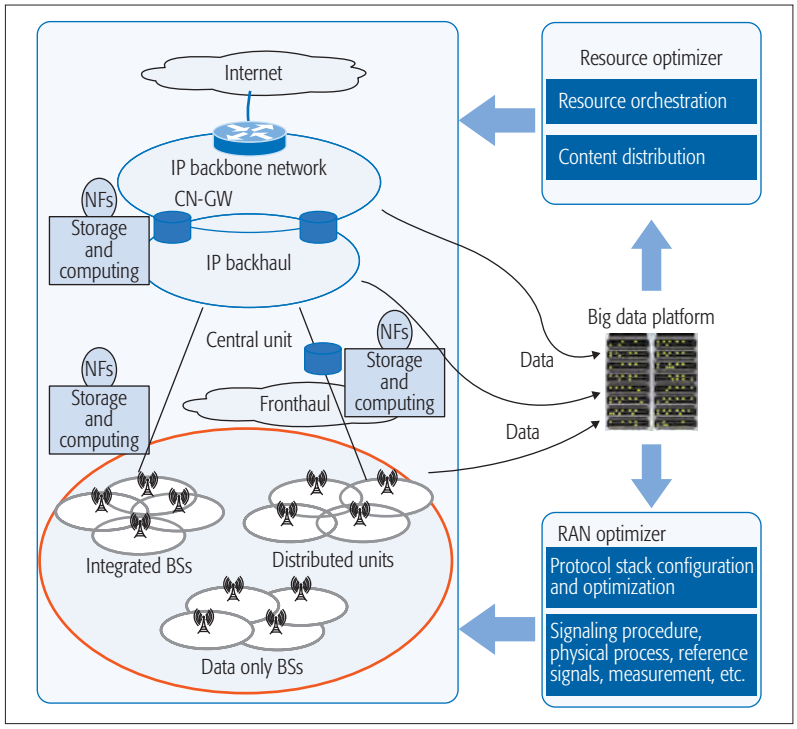

Figure 1. Big data enabled mobile network architecture.

dure and physical layer optimization. Big data changes the network design from the reactive BS-centric approach to the proactive user-centric paradigm.

According to the service requirements and network conditions analyzed by the big data platform, the network functions (NF) implementation can be flexibly either centralized or distributed at the edge. Specifically, the new features in big data enabled networks are listed as follows.

On Demand Resource Orchestration

In conventional networks, a large amount of resources may be wasted in low-traffic scenarios. By utilizing wireless big data, it is possible to predict users' future service requests, location and mobility, and the network conditions. With this useful information, the proper resources can be provisioned, guaranteeing high resource efficiency and reducing network cost by avoiding overprovision. For example, the processing resources of the CU can be properly configured to achieve more pooling gain according to the predicted traffic fluctuation. In addition, some parts of access resources can be turned off if no traffic is predicted in the corresponding coverage area.

Network slicing [12] is a key new feature of 5G networks for highly diversified 5G services. In order to guarantee Quality of Service (QoS) of each specific network slice, some relatively static isolation of radio resources may be adopted, which may lead to insufficient resource utilization. Different slices may be of diverse characteristics, for example, distinct peak or idle traffic patterns. The usage of radio resources for different slices may compensate for each other, for example, a slice is at its busy hour, and another slice is at its idle hour. Such traffic patterns of different slices can be identified via big data technologies, leading to better utilization of radio resources among slices.

Flexible Content Distribution

With the assistance of big data technologies, the network is able to predict users' traffic patterns with more accuracy. With adequate storage and computing capabilities, the network edge nodes could pre-fetch the predicted popular content beforehand during idle hours, instead of fetching content upon users' request via potentially overloaded backhaul links.

Obviously, the closer the content is to users, the shorter the responsive latency is. However, a higher level of storage and computing means more pooling gain and less maintenance cost. Different traffic patterns, such as content popularity distributions, may require content deployment in different levels of network nodes. For example, when the traffic is low or the traffic trend is of low predictability, it is more attractive to place content storage and computing at a higher level in a cost-efficient way. However, when the traffic trend can be accurately predicted, it is beneficial to improve the user's experience by moving the corresponding content storage and computing functions to network nodes that are closer to the users.

Protocol Stack Configuration and Optimization

With flexible function split between the CU and the DU, the corresponding protocol stack configuration is required. Also, there may be the scenario where dual connectivity (DC) or carrier aggregation (CA) is implemented in the network. The protocol stacks at the CU, DU, integrated BS and data-only BS need to be flexibly configured. For example, with ideal fronthaul, media access control (MAC) functions can be configured at the CU, optimally allocating resources between different DUs. Big data analytics help optimize the protocol stack configuration and stack processing in various scenarios.

Signaling Procedure Optimization

Signaling procedure, as specified in various standards, stipulates the signaling flows between the user equipment (UE), enhanced Node B (eNB), mobility management entity (MME), serving gateway (S-GW), packet gateway (P-GW), policy and charging rules function (PCRF), and home subscriber server (HSS), to take the terminologies of LTE as an example. With big data analytics, many signaling procedures can be simplified, when the data are available at the network. The simplified procedure may bring much reduced operation complexity and latency, which is urgently needed in, for example, ultra-low latency applications.

Physical Layer Procedure Optimization

When the protocol stack is configured and the processing at each layer is optimized based on service type and scenario, the application data will be transmitted step by step through the stacks. Traditional physical layer processing, such as synchronization, modulation and coding, multiple access, multiple antenna precoding, duplex mode selection, numerology configuration, reference signals, measurements and feedback, and power control and so on, can also be significantly facilitated via big data technology.

Protocol Stack and Signaling Procedure Optimization

Protocol Stack Configuration

The traditional protocol stack for the normal integrated BS is shown in Fig. 2a, where data processing is through the Packet Data Convergence Protocol (PDCP), radio link control (RLC), MAC, and physical (PHY) processing, with the management of radio resource control (RRC). While in the DC, as shown in Figs. 2b and 2c, the PDCP function resides at the Master eNB (MeNB). The Secondary eNB (SeNB) is only responsible for the RLC, MAC, and PHY. For CA operations, joint MAC scheduling is feasible, leaving the SeNB responsible only for PHY processing, as shown in Figs. 2d and 2e. For both DC and CA, there is only one RRC at the MeNB. As shown in Fig. 2f, when the data-only BS is deployed with low data rate fronthaul to provide local data services, TCP and IP processing can be conveniently omitted. The PDCP processing can be further simplified, for example, IP header compression is no longer needed. With big data analytics, BSs may receive recommendations on the operation mode, for example, coordinated multiple point (CoMP), DC or CA. The protocol stack is also configured accordingly. For example, with information about the users' geographical distribution and/or the inter-cell interferences, CoMP schemes may be adopted in a certain area for better performance.

For the CU/DU structure, stack partition between the CU and the DU can also be optimally configured via big data analytics based on service type, fronthaul capability, frequency bands, user mobility, quality of experience, and so on. For example:

- For low latency services, fewer functions will be allocated to the CU, for example, the full eNB protocol stack at the DU.
- For high frequency bands such as mmWave, more processing is required at the DU to alleviate the burden of fronthaul from extremely high data rate. One stack configuration for case 1 and 2 is shown in Fig. 2g, where RRC and PDCP reside at the CU, while RLC, MAC and PHY are handled at the DUs.
- For low frequency bands where severe inter-cell interference exists, more processing will be motivated at the CU for efficient interference mitigation.
- With ideal fronthaul, MAC functions can be configured at the CU, optimally allocating resources between different DUs. While with non-ideal fronthaul, MAC functions at the central unit are not necessary, since cross-cell fast scheduling and resource allocation may not be supported well due to much longer fronthaul latency. In this case, only RRC and PDCP can be allocated at the CU. One stack configuration for cases 3 and 4 is shown in Fig. 2h, where RRC, PDCP, RLC and MAC reside at the CU, while PHY is handled at the DUs. A hybrid CU/DU protocol stack is shown in Fig. 2i, where the CU is controlling the DUs with different protocol stack configurations.
- For high mobility users, integrated BSs with a full protocol stack are not preferred. Cloud RAN architecture with RRC at the CU is more suitable.

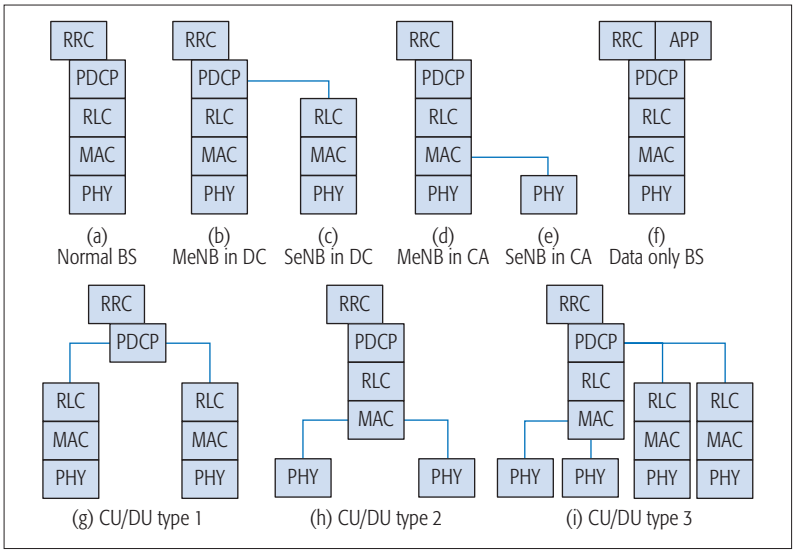

Figure 2. Protocol stack configuration.

Protocol Stack Processing Optimization

Reconfigurable Compression and Encryption in PDCP: After the protocol stack is configured, data processing can also be further optimized based on big data analytics. For instance, a robust header compression mechanism (ROHC) is utilized to handle user plane data flow which has large packet headers. However, the incurred delay takes up 20.01 percent of the L2 total delay. Big data analysis can be used to identify data packets or data flows with the same service types. The identified latency-insensitive IP packets can then be aggregated into a large data packet that shares one IP header, leading to much reduced ROHC processing delay.

Traditional user plane and control plane packets all need to be encrypted when passing through the PDCP layer. The delay caused by the ciphering process takes up 59.16 percent of the L2 total time delay. If their corresponding service types can be analyzed by big data, differentiated ciphering processing can be selected. If some services are not private, they do not need ciphering, thus reducing processing delay and complexity. For example, e-commerce transactions and news browsing definitely have distinct privacy requirements, so ciphering over the air should be adaptable for different service categories to avoid unnecessary overhead. Besides, powerful big data analysis is capable of identifying potential security attacks, monitoring and eliminating potential dangers, and may effectively reduce the necessity of data ciphering.

Context-aware user privacy protection has also been investigated in [13], where the quality of protection (QoP) middleware in the mobile phone controls how upper-layer applications access sensor data. Accordingly, the service provider offers personalized services based on users' sensing data. With powerful big data analytics, the privacy level of each mobile application can be configured automatically based on comprehensive analysis of service types and users' preferences, not necessarily relying on the assistance of QoP middlewares.

Optimized Transmission Mode in RLC: Tradi-

Big data analysis can be used to identify data packets or data flows with the same service types. The identified latency-insensitive IP packets can then be aggregated into a large data packet that shares one IP header, leading to much reduced ROHC processing delay.

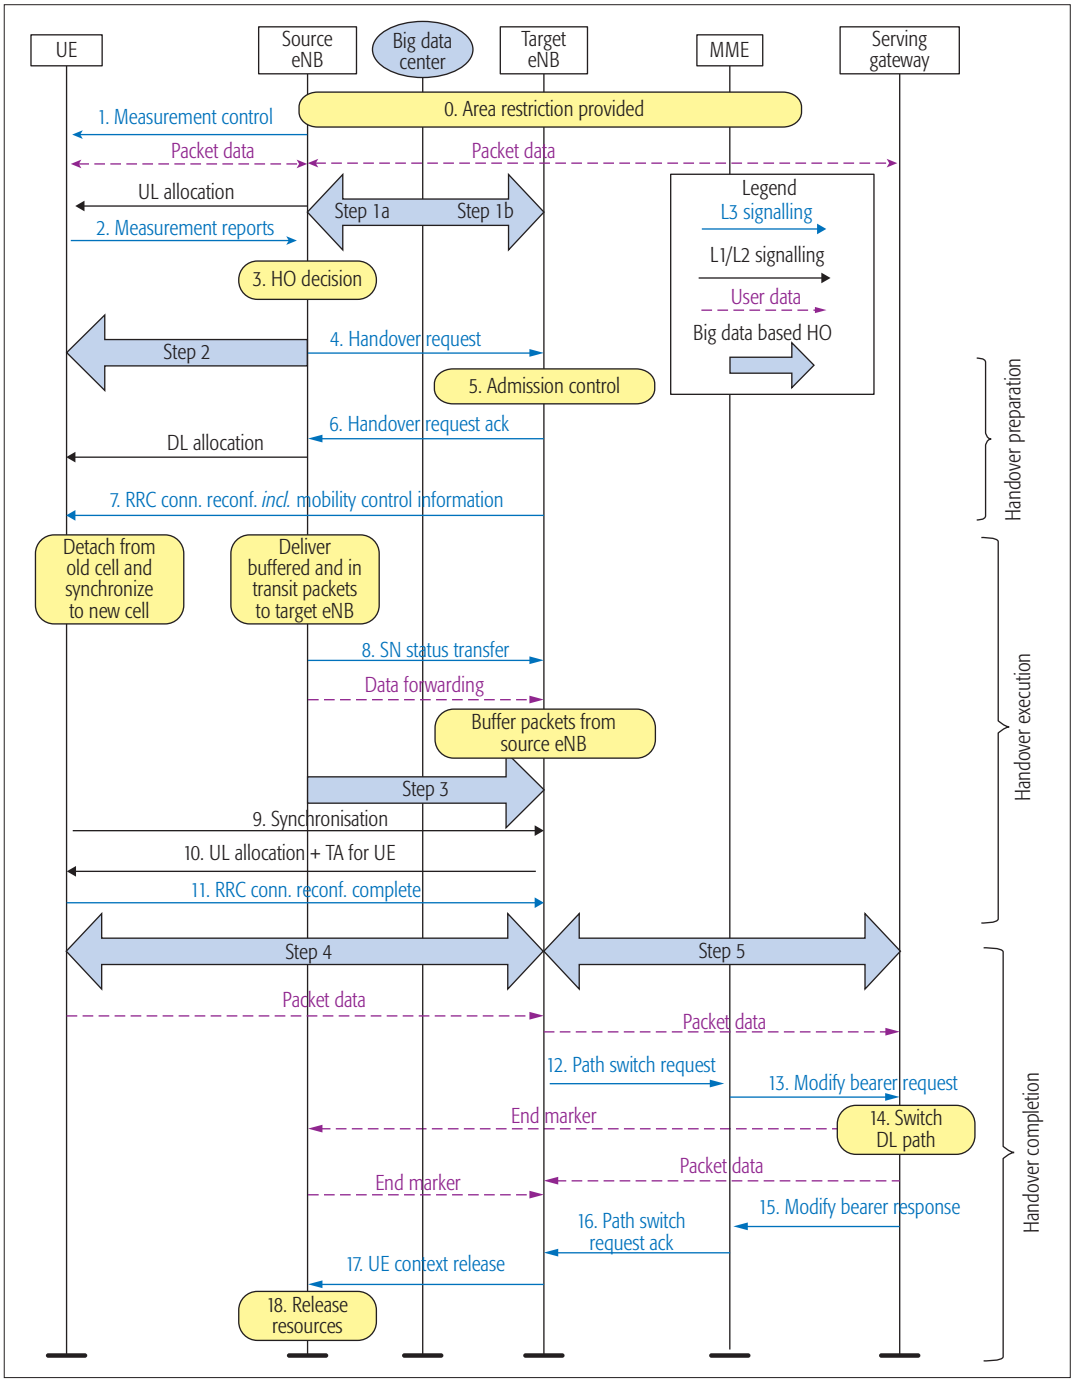

Figure 3. Handover procedure.

tional RLC modes are configured by CN according to service types. Audio and video flow services basically use unacknowledged mode (UM), but acknowledged mode (AM) can be used if delay requirements are not very strict, under which the reliability will be greatly improved. UM mode transmission is more suitable for small-packet traffic as well, which usually has a small number of segments and generally does not need Automatic Repeat reQuest (ARQ). Big data analysis is capable of identifying traffic delay sensitivity and accurate identification of small packet traffic, and brings much-reduced delay and processing complexity.

MAC Hybrid ARQ (HARQ): Through statistical analysis of channel and traffic characteristics, a maximum retransmission number can be dynamically configured by HARQ. This will reduce overhead and improve resource utilization.

Signaling Procedure Optimization: Taking handover as an example, we will investigate how big data helps simplify system operation and improve performance. The basic procedure of the X2 handover [14] is illustrated in Fig. 3 (not including the gray color procedures with arrows and the big data center), which is very complicated. The handover is generally triggered by the eNB based on the measurement report feedback from UEs if certain criteria of the measured channel conditions in the adjacent cells are met. The measurement process is complicated and the signaling overhead is large.

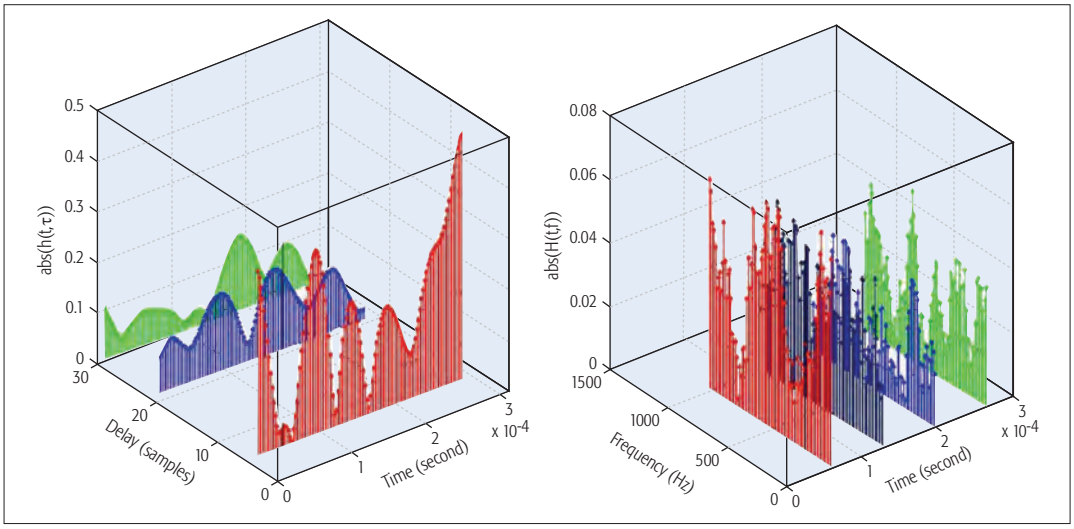

Figure 4. a) Time domain channel; b) frequency domain channel.

The availability of CSI at BS and UE is crucial for wireless communication systems. To this end, reference signals are widely adopted in various wireless standards for estimation of fast variation of wireless channels. For time-varying channels, the wide-sense stationary uncorrelated scatter model is widely used in theoretical analysis.

Based on big data analytics, one possible handover approach is proposed as follows (also shown in Fig. 3 with bold gray arrow lines).

Step 1a: The big data center determines when each UE should handover to which target eNB based on accurate prediction of UE location and movement. It sends a handover command to the source eNB, with information anout the target eNB.

Step 1b: Meanwhile, the big data center sends a handover command to the target eNB, possibly with the sequence number (SN) status.

Step 2, 3 and 4 (step 7 to step 11 in the traditional handover): These can be the same as the traditional handover.

Step 5: Path switch (step 12 to step 16 in the traditional handover) and context release can be made before steps 2, 3 and 4 finish, especially when the big data center is sure of the success of the handover.

Compared with the traditional handover procedure, big data based handover has the following advantages:

- Signaling overhead is reduced, e.g. handover request, acknowledgment (ACK), admission control and possibly SN status transfer can be omitted.
- UE measurement and feedback efforts can be reduced.
- Handover interruption time can be reduced, due to reduced signaling procedure and due to possible concurrent UE access to the target eNB and path switch process.
- Ping-pong handover can be effectively avoided.

This methodology can be extended to many other signaling procedures. For example, in the attach procedure, if it is not the first attach, and the time interval between the current attach and the previous attach is small (e.g. several minutes), the identity, authentication and security procedures can be neglected or significantly simplified.

RAN Assisted Application Optimization

Obviously, there is a mismatch between millisecond-timescale variation of the radio information (e.g. UE buffer status on the RAN side, radio data rate, UE channel quality, cell congestion level) and more-than-second-timescale application layer adjustments. If some useful RAN information is provided to the application layer after it is processed by data technologies (not necessarily big data processing), the application behavior may be optimized to better match the radio channel variations, leading to better user quality of experience and higher network performance. Taking TCP congestion window adjustment as an example, radio information can be obtained from the BS, and then processed to produce application-level control indication, which can be utilized by the TCP sender to adjust its congestion window.

High-Efficiency Physical Layer Operation

Channel State Information Acquisition and Feedback

The availability of channel state information (CSI) at the BS and the UE is crucial for wireless communication systems. To this end, reference signals are widely adopted in various wireless standards for estimation of fast variation of wireless channels. For time-varying channels, the wide-sense stationary uncorrelated scatter (WSSUS) model is widely used in theoretical analysis. Typical channel responses in a multiple path environment in the time domain (i.e., $h(t, \tau)$, where, t and τ denote time and delay spread, respectively), and responses in the frequency domain (i.e., $h(t, f)$, where f denotes frequency) are depicted in Figs. 4a and 4b, respectively, which are very dynamic.

The density of reference signals in frequency division duplex (FDD) should be high enough to capture the channel's characteristics in both the time and frequency domains. Generally, this density is selected according to a comparatively high mobility and large delay spread, especially for the cell common reference signals. In cases where most users within the cell are moving slowly, a high density of the reference signals leads to unnecessary overhead which is over provisioning. The situation is more problematic with a large number of antennas at the BS, since the overhead of reference signals scales with the number of antenna ports.

One efficient way to minimize the overhead

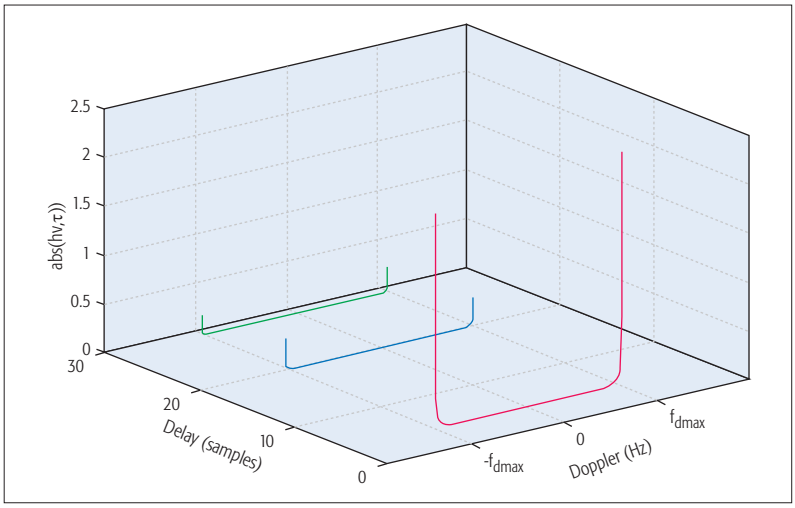

Figure 5. Effective channels in the delay and Doppler domains.

of the reference signals is to resort to wireless big data analytics. With the prediction of user mobility and wireless channel statistics, the BS is able to schedule common reference signals accordingly. If there are no high-mobility users, very sparse reference signals may suffice. If high-mobility and low-mobility users coexist, the BS can schedule dense reference signals in some frequency bands and sparse reference signals in other bands, thus effectively minimizing the overall overhead.

In time division duplex (TDD) systems, the CSI can be obtained by exploiting the channel reciprocity using uplink pilots. However, due to limited training sequences, pilot contamination may dramatically degrade the downlink performance since the uplink channels are contaminated by the inter-cell interference from the adjacent cells. In order to reduce pilot contamination, with the help of big data, joint user and pilot scheduling schemes can coordinate the non-orthogonal pilot sequence among users of weak mutual interference, and arrange orthogonal pilot sequences to users with strong mutual interference.

Transform-Domain Signal Processing to Facilitate Semi-Static Scheduling

The dynamic channel variations in the time and frequency domains necessitate dense reference signals, fast feedback, adaptive modulation and coding, and fast scheduling, thus posing tough challenges for wireless communication system design. Fortunately, via Fourier transformation with respect to variable t, the dynamic $h(t, \tau)$ can be transformed to a stationary $h(v, \tau)$ in delay and Doppler domains, where v is the Doppler frequency. An example of $h(v, \tau)$ is shown in Fig. 5. In contrast to the time and frequency domain channel responses, the delay and Doppler domain channels are more stable, depending merely on the multi-path channel structure (angular distribution and the power delay profile) and mobility. Therefore, it is almost static if the channel structure and mobility do not change.

The direct impact of transform-domain signal processing is the alleviated difficulty in tracking the time-varying fading. This is particularly useful in high-speed train communications. The significantly increased coherence time of the effective channel in transform domains brings abundant opportunities for the simplification of wireless systems in both standardization and implementation. For example, reference signals can be designed with very low overhead and the channel feedback need not to be fast. Channel coding schemes can also be simplified. The well studied AWGN codes may perform sufficiently well over the effective channel, thus alleviating the burden of traditional adaptive modulation and coding. Moreover, it enables FDD massive multiple input multiple output (MIMO) in moving applications due to easy CSI estimation. Most importantly, the slow variation of the effective channels in transform domains can significantly facilitate big data analytics, since analyzing the statistical channels may be enough for satisfactory PHY and MAC operations.

Flexible Frame Structure Configuration

Since there may be many use cases emerging in 5G and beyond, it is very important for operators to deploy one network to support all use scenarios and use cases. Toward this end, it is critical to adopt one unified and flexible air interface framework to meet diverse requirements of the key 5G scenarios, for example, enhanced mobile broadband, ultra-reliable and low latency communications and massive machine type communications. The unified framework of the software defined air interface (SDAI) [15] may include a flexible frame structure, waveform, duplex mode, multiple access, MIMO, coding and modulation, and corresponding layer 2 and layer 3 signaling. For efficient operation of SDAI, the key is a flexible frame structure, for example, the numerology can be dynamically configured; the time resource within each subframe can be flexibly partitioned between downlink and uplink; the duration of the transmit time interval is adaptive; and the period of uplink feedback of ACK/NACK is configurable.

The practical implementation of a flexible frame structure at each BS is very difficult. For example, flexible downlink and uplink transmission may cause severe inter-cell and intra-cell interference. However, mitigating the interference, especially the cross-link interference, can be very challenging. Another example is that the frequent uplink feedback for the latency-sensitive service may cause severe interference to the downlink of the latency-non-sensitive service. Thanks to big data technologies, a lot of useful information can be utilized to optimize the frame structure, for example, UE service types, quality of experience, location, traffic volume, mobility, channel information and inter-user interference.

Conclusions

This article has discussed how wireless big data can impact wireless communication network design, from the perspectives of network architecture, protocol stack, signaling procedure, and physical layer operations. A big data enabled network architecture was proposed, along with design considerations on protocol stack configuration, simplified signaling procedure like handover, simplified stack processing in the PDCP, RLC, and MAC layer, low overhead reference signals, and flexible frame structure. The potential impact of transform-domain signal processing on system design is also discussed, which facilitates the application of wireless big data. Wireless big

data, which is generated in mobile networks and is seemingly a burden to the network, nevertheless can be eventually transformed into a blessing, enabling much simplified network operations and standardization.

Acknowledgment

The authors would like to thank the 5G team members within the Green Communication Research Center at the China Mobile Research Institute, particularly Zhiming Liu, Wei Zhou, Jun Zuo, Guozhen Xu, Ailing Wang, Yami Chen, Siming Zhang and Jiqing Ni.

Reference

[1] S. Yu *et al.*, "Networking for Big Data: A Survey," *IEEE Commun. Surveys Tutorials*, vol. 19, no. 1, First Qrtr. 2017, pp. 531–49.
[2] C. Wang *et al.*, "Cellular Architecture and Key Technologies for 5G Wireless Communication Networks," *IEEE Commun. Mag.*, vol. 52, no. 2, Feb. 2014, pp. 122–30.
[3] M. Feng *et al.*, "Cooperative Small Cell Networks: High Capacity for Hotspots with Interference Mitigation," *IEEE Wireless Commun.*, vol. 21, no. 6, Dec. 2014, pp. 108–16.
[4] J. Liu, F. Liu, and N. Ansari, "Monitoring and Analyzing Big Traffic Data of a Large-Scale Cellular Network with Hadoop," *IEEE Network*, vol. 28, no. 4, July/Aug. 2014, pp. 32–39.
[5] Y. He *et al.*, "Bigdata Analytics in Mobile Cellular Networks," *IEEE Access*, vol. 4, Mar. 2016, pp. 1985–96.
[6] K. Zheng *et al.*, "Big Data-Driven Optimization for Mobile Networks toward 5G," *IEEE Network*, vol. 30, no. 1, Jan. 2016, pp. 44–51.
[7] W. Wang and Q. Zhang, "Local Cooperation Architecture for Self-Healing Femtocell Networks," *IEEE Wireless Commun.*, vol. 21, no. 2, Feb. 2014, pp. 42–49.
[8] S. Bi *et al.*, "Wireless Communications in the Era of Big Data," *IEEE Commun. Mag.*, vol. 53, no. 10, Oct. 2015, pp. 190–99.
[9] C.-L. I *et al.*, "On Big Data Analytics for Greener and Softer RAN," *IEEE Access*, vol. 3, Mar. 2015, pp. 3068–75.
[10] D. Liu, *et al.*, "Caching at the Wireless Edge: Design Aspects, Challenges and Future Direction," *IEEE Commun. Mag.*, vol. 54, no. 9, Sept. 2016, pp. 22–28.
[11] C.-L. I *et al.*, "Rethink Fronthaul for Soft RAN," *IEEE Commun. Mag.*, vol. 53, no. 9, Sept. 2015, pp. 82–88.
[12] X. Zhou *et al.*, "Network Slicing as a Service: Enabling Enterprises' Own Software-Defined Cellular Networks," *IEEE Commun. Mag.*, vol. 54, no. 7, July 2016, pp. 146–53.
[13] W. Wang and Q. Zhang, "Toward Long-Term Quality of Protection in Mobile Networks: A Context-Aware Perspective," *IEEE Wireless Commun.*, vol. 22, no. 4, Aug. 2015, pp. 34–40.
[14] Y. Li, B. Cao, and C. Wang, "Handover Schemes in Heterogeneous LTE Networks: Challenges and Opportunities," *IEEE Wireless Commun.*, vol. 23, no. 2, Feb. 2016, pp. 112–17.
[15] C.-L. I *et al.*, "New Paradigm of 5G Wireless Internet," *IEEE JSAC*, vol. 34, no. 3, Mar. 2016, pp. 474–82.

Biography

Shuangfeng Han (hanshuangfeng@chinamobile.com) is a senior project manager in the Green Communication Research Center of the China Mobile Research Institute. He is also vice chair of the wireless technology work group of China's IMT-2020 (5G) promotion group. His research interests are mainly focused on 5G wireless communication systems, including massive MIMO, flexible duplex, non-orthogonal multiple access schemes, energy efficiency and spectrum efficiency co-design, joint design of virtual reality applications and wireless networks, high spectrum efficiency technologies for high speed trains, and the application of wireless big data analytics. He graduated from Tsinghua University in Beijing in 2006, majoring in information and communication systems. He is an inventor of over 60 patents in the 4G and 5G areas, and he is the author of over 30 peer-reviewed conference and journal publications. He has been reviewer for various IEEE journals and conferences.

Chih-Lin I (icl@chinamobile.com) received her Ph.D. degree in electrical engineering from Stanford University. She has been working at a number of world-class companies and research institutes leading their R&D activities, including AT&T Bell Labs, Director of AT&T HQ, Director of ITRI Taiwan, and VPGD of ASTRI Hong Kong. She received the *IEEE Transactions on Communications* Stephen Rice Best Paper Award. She is a winner of the CCCP National 1000 Talent Program, and the 2015 Industrial Innovation Award of the IEEE Communication Society for Leadership and Innovation in Next-Generation Cellular Wireless Networks. In 2011, she joined China Mobile as its chief scientist of wireless technologies, established the Green Communications Research Center, and launched the 5G Key Technologies R&D. She is spearheading major initiatives including 5G, C-RAN, high energy efficiency system architectures, technologies and devices, and green energy. She was an area editor of *IEEE/ACM Transactions on Networking*, an elected board member of the IEEE Communications Society, chair of the ComSoc Meetings and Conferences Board, and founding chair of the IEEE WCNC Steering Committee. She has been a professor at NCTU, an adjunct professor at NTU, and is currently an adjunct professor at BUPT. She is the chair of FuTURE 5G SIG, an executive board member of GreenTouch, a Network Operator Council founding member of ETSI NFV, a steering board member of WWRF, the ComSoc representative on the IEEE 5G Initiative, a member of IEEE ComSoc SDB, SPC, and CSCN-SC, and a Scientific Advisory Board member of Singapore NRF. Her current research interests center around "green, soft, and open."

Gang Li (ligangyf@chinamobile.com) received his B.A. degree in telecommunication engineering and M.E. degree in automation engineering from Sichuan University. After graduation, he worked for Lucent Technologies for four years as a team leader and software developer for the core network. He is now a senior researcher at the Green Communication Research Center of the China Mobile Research Institute, working on the key technologies of next generation 5G wireless communication systems. His research interests include radio access network architecture optimization, service-aware signaling/control redesign, and radio and core network convergence.

Sen Wang (wangsenyj@chinamobile.com) received the B.S.E. degree from Information Engineering University, Zhengzhou, China, in 2005, and the Ph.D. degree in information and communication engineering from Beijing University of Posts and Telecommunications (BUPT), Beijing, China, in 2013. After graduation in January 2013, he joined the Green Communication Research Center (GCRC), China Mobile Research Institute, as a project manager. His research interests include 5G air interface technologies, especially on massive MIMO, new waveform and multiple access, radio resource allocation and performance evaluation for future cellular networks. Together with his colleagues, he has authored 10 IEEE journal and conference papers in IEEE JSAC, *IEEE Communications Letters*, ICC, and WCNC in these areas. He also serves as a principle member of the IMT-2020 requirements and evaluation WG.

Qi Sun (sunqiyjy@chinamobile.com) received her Ph.D. degree in information and communication engineering from Beijing University of Posts and Telecommunications in 2014. After graduation, she joined the Green Communication Research Center of the China Mobile Research Institute. Her current research interests include non-orthogonal multiple access, cross-layer resource allocation for wireless networks, and big data enabled wireless network design.

Wireless big data, which is generated in mobile networks and is seemingly a burden to the network, nevertheless can be eventually transformed into a blessing, enabling much simplified network operations and standardization.

ACCEPTED FROM OPEN CALL

Nonlinear Self-Interference Cancellation for Full-Duplex Radios: From Link-Level and System-Level Performance Perspectives

Min Soo Sim, MinKeun Chung, Dongkyu Kim, Jaehoon Chung, Dong Ku Kim, and Chan-Byoung Chae

The authors explore several nonlinear digital self-interference cancellation techniques. They then propose a low-complexity pre-calibration-based nonlinear digital self-interference cancellation technique. They discuss issues about reference signal allocation and the overhead of each technique.

ABSTRACT

One of the promising technologies for Long Term Evolution is full-duplex radio, an innovation that is expected to double spectral efficiency. To realize full-duplex in practice, the main challenge is overcoming self-interference, and to do so, researchers have developed self-interference cancellation techniques. Since most wireless transceivers use power amplifiers, especially in cellular systems, researchers have revealed the importance of nonlinear self-interference cancellation. In this article, we first explore several nonlinear digital self-interference cancellation techniques. We then propose a low-complexity pre-calibration-based nonlinear digital self-interference cancellation technique. Next, we discuss the issues about reference signal allocation and the overhead of each technique. For performance evaluations, we carry out extensive measurements through a real-time prototype and link-/system-level simulations. For link-level analysis, we measure the amount of cancelled self-interference for each technique. We also evaluate system-level performance through 3D ray-tracing-based simulations. Numerical results confirm the significant performance improvement over a half-duplex system even in interference-limited indoor environments.

INTRODUCTION

As a solution to the tremendous expansion of mobile traffic, researchers have been developing, over the past several years, fifth generation (5G) wireless communication/Long Term Evolution (LTE). One of the main requirements for this service is to provide a 1000-fold improvement in throughput over current 4G mobile networks such as LTE-Advanced [2]. To achieve this requirement, researchers have strived to accomplish the following: improving spectral efficiency (bits per second per Hertz), expanding system bandwidth (Hertz), and/or increasing throughput per area (bits per second per square meter). Several promising technologies have been developed to improve spectral efficiency such as massive multiple-input multiple-output (MIMO), 3D beamforming, and in-band full-duplex radios.

In-band full-duplex radios (full-duplex radios hereafter) simultaneously transmit and receive on the same frequency band [1, 3–11]. Full-duplex systems are expected, by definition, to double the spectral efficiency of half-duplex systems. Current commercial systems, however, have engaged with little attention to this type of system due to its propensity for self-interference. Self-interference is the phenomenon of a signal, transmitted from a transmitter, being received by its own receiver while that receiver is trying to receive a signal sent from another device (signal of interest). The self-interference, which is generally far stronger than the signal of interest, makes it impossible for a device to decode the signal of interest. Today, mobile networks, in order to avoid self-interference, operate in half-duplex. Frequency-division duplex (FDD) systems prevent self-interference by allocating different frequency bands for uplink and downlink. Time-division duplex (TDD) systems transmit and receive at different times. To deal with the self-interference issue in full-duplex systems, researchers have developed several self-interference cancellation techniques, the objective of which is to mitigate or cancel the self-interference to noise level.

Applying full-duplex technology, due to its high transmit power, is a challenge in cellular networks such as LTE. Typical full-duplex systems provide self-interference cancellation of approximately 120 dB [5, 6]. For a macrocell base station (BS) whose transmission power is up to 46 dBm, considering a noise level of –90 dBm, it is not feasible to suppress self-interference to the noise level. Therefore, applications of full-duplex radio in LTE systems could be limited to relay systems or small cell systems where transmission power is at most 23 dBm (Fig. 1). To apply full-duplex to these scenarios, the nonlinearity of power amplifiers becomes a bottleneck of self-interference cancellation. A power amplifier's nonlinearity is one of the imperfections of RF transceivers; other imperfections include in-phase and quadrature (I/Q) imbalances and phase noise. Since self-interference with these RF imperfections are not sufficiently suppressed in the analog domain, digital self-interference cancellation should share the burden. Most of the digital processing in conventional communication systems, however, is designed linearly. Therefore, special techniques that can handle this nonlinearity are required for the processing to perform self-interference cancellation well.

This article was presented in part at the IEEE Globecom FDWC Workshop 2016 [1].

This work was in part supported by the MISP under the "ICT Consilience Creative Program" (IITP-2017-2017-001015), the ICT R&D Program of MSIP/IITP (2015-0-00294), and LG Electronics.

Digital Object Identifier: 10.1109/MCOM.2017.1600264

Min Soo Sim, MinKeun Chung, Dong Ku Kim, and Chan-Byoung Chae are with Yonsei University; Dongkyu Kim and Jaehoon Chung are with LG Electronics.

In this article, we investigate several techniques that can cancel nonlinear self-interference. We introduce the concepts, compare reference signal allocations, and evaluate link- and system-level performance. To the best of our knowledge, this is the first attempt to investigate *system-level performance of full-duplex radios based on measured data*.

The rest of this article is organized as follows. First, we introduce several self-interference cancellation techniques including isolation, active analog cancellation, and linear digital cancellation. We next detail representative nonlinear digital self-interference cancellation techniques. Then we investigate link- and system-level performance evaluations, and our conclusions are given.

An Overview of Self-Interference Cancellation Techniques

Most prior work on full-duplex radios has focused on designing a wireless transceiver that can perform a sufficient amount of self-interference cancellation. Included in such work are the following:

1. Designing an antenna and its configuration that lessen coupling between a transmitter and a receiver
2. Suppressing self-interference by mimicking the analog self-interference signal and subtracting from it
3. Canceling digitalized self-interference by modeling self-interference and subtracting from it

Figure 2 illustrates a schematic of several self-interference cancellation technologies. The ultimate goal is to maximize total cancellation amount by integrating and optimizing these techniques. In this section, to show how self-interference cancellation works in full-duplex systems, we briefly introduce some of the conventional self-interference cancellation techniques.

Analog Self-Interference Cancellation

Analog self-interference cancellation suppresses self-interference in the analog domain, that is, before the signal passes through an analog-to-digital converter (ADC). The main role of analog self-interference cancellation is to take up a portion of the total cancellation amount and to make sure that the residual self-interference can be canceled out in the digital domain. Since the cancellation performance in the digital domain is limited by the dynamic range of an ADC, analog self-interference cancellation should suppress certain portions of self-interference. According to tunability or adaptability, analog self-interference is categorized, in this article, as isolation (a.k.a. passive analog self-interference cancellation) and active analog self-interference cancellation.

Isolation: Isolation suppresses self-interference signals in the analog domain without any adaptive tuning. Special antenna structures and configurations, and passive devices are employed to weaken, passively, self-interference signals. A basic method for isolation is antenna separation, which causes path loss between a transmit antenna and a receive antenna [7]. All transceivers that use different antennas for transmitters and receivers can obtain self-interference suppression gain. A circulator was employed for the system in which transmitters and receivers share antennas [5]. A circulator is a three-port device that transfers a transmit signal from a transmit port to an antenna and a receive signal from an antenna to a receive port while the signal from the transmit port is blocked from the receiver port. With a dual-polarized antenna, a leakage from a transmitter to a receiver can be prevented despite the short physical distance [3]. To obtain extra isolation, a technique called antenna cancellation was proposed, which uses a symmetric antenna configuration and a π-phase shifter to take advantage of destructive interference [8].

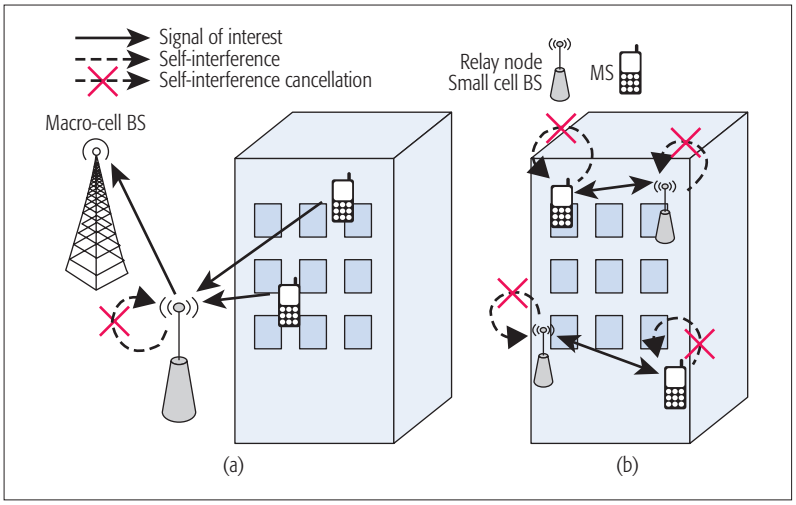

Figure 1. Scenarios of full-duplex-radio-based LTE systems: a) a full-duplex LTE relay system; b) a full-duplex LTE small cell system.

Active Analog Self-Interference Cancellation: Active analog self-interference cancellation, through adaptive tuning and algorithms, aims at the dominant components of self-interference. The dominant components here represent a line-of-sight component of a system with separate transmit and receive antennas, or a leakage from a circulator of a system with a single transceiver antenna. A typical solution of active analog self-interference cancellation is a tunable RF circuit [4 5]. The circuit, which consists of several taps with RF components such as attenuators, phase shifters, and delays, uses a replica of the transmit signal as an input, and tries to mimic the self-interference signal. Since the circuit directly exploits the transmit signal with RF imperfections, which in the digital domain are in fact hard to handle, it would have a better chance of suppressing the self-interference in analog cancellation. To follow the time-varying characteristic of the self-interference channel, a real-time adaptive algorithm and extra reference signal are needed to control RF components. Note that minimizing a noise of the RF circuit, which can be achieved by employing passive devices that generate virtually no noise, is essential for stable cancellation [11]. Analog cancellation that generates an RF signal with an auxiliary transmit chain was also introduced in [7, 10], where the generated signal was combined with the received signal to suppress self-interference. With a multiple-antenna system, self-interference can be suppressed by transmit beamforming [11].

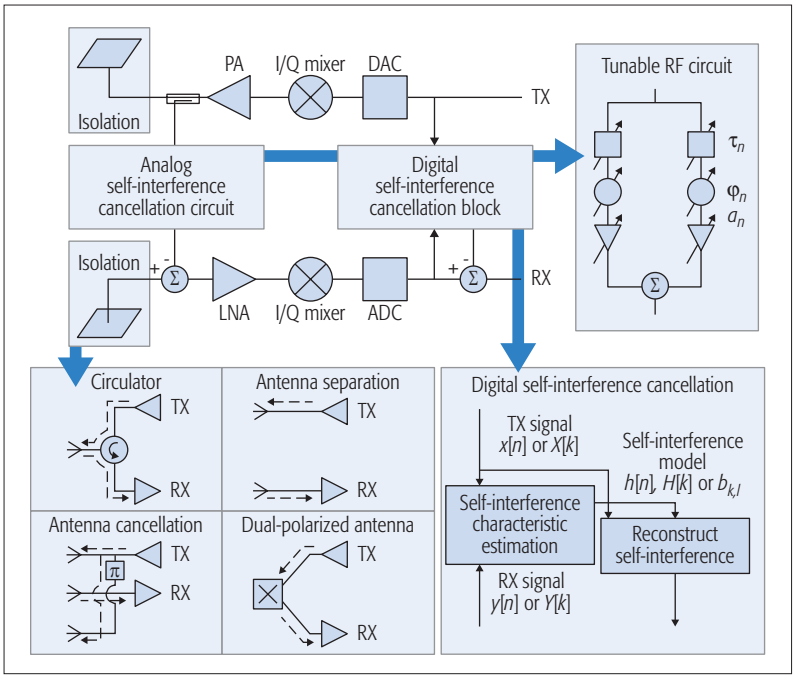

Figure 2. A full-duplex transceiver with self-interference cancellation techniques. Representative isolation, active analog self-interference cancellation, and digital self-interference cancellation schemes are illustrated.

Linear Digital Self-Interference Cancellation

Reconstruct the self-interference cancellation and subtract cancels out the residual self-interference from analog self-interference cancellation. Generally, a line-of-sight component or a direct leakage of a circulator is suppressed by analog self-interference cancellation, but non-line-of-sight components (reflections) are not. Therefore, these components should be removed by digital cancellation. There are three steps for digital cancellation:

1. Set a model of a self-interference signal
2. Estimate the channel
3. Reconstruct the self-interference signal and subtract it from a received signal

Note that a reference signal is not only for self-interference cancellation, but also for signal of interest demodulation.

Due to its simplicity, many researchers use a linear model for a wireless channel. Therefore, self-interference channels are assumed to be linear, and linear self-interference cancellation becomes the basic technology of digital self-interference cancellation. In this article, we introduce two kinds of linear estimation methods and the following reference signal allocations. For simplicity, we assume orthogonal frequency-division multiplexing (OFDM) with the extended cyclic prefix (CP).

Linear Digital Cancellation in the Time Domain: The first method is to estimate the self-interference channel in the time domain [5]. Consider a full-duplex system that adopts, as shown in Fig. 3a, one OFDM symbol as a reference signal. In this case, the self-interference channel is obtained by the least squares method in the time domain. Note that this method should exploit the reference signal allocated in all subcarriers. To overcome fast-fading channel, however, the reference signals should be repeated, which creates a tremendous overhead for the reference signal. Furthermore, since the channel-estimation part includes pseudo-inverse operations, and the self-interference reconstruction part includes convolution operations, this method has high computational complexity.

Linear Digital Cancellation in the Frequency Domain: It is more natural to allocate the reference signal as shown in Fig. 3b [3], which is a structure similar to that used in the LTE standard. This reference signal is used for both the self-interference link and the desired link. With this pattern, the self-interference channel can be calculated by the least squares method in the frequency domain. In other words, the channel is obtained by dividing a received signal passed through a fast Fourier transform (FFT) by the reference signal before an inverse FFT (IFFT) block. This method has relatively low complexity since it only requires element-wise multiplications and divisions, and interpolations.

Nonlinear Self-Interference Cancellation

The need for cancelling nonlinearity arises when we attempt to apply full-duplex radios to systems with high transmission power. Suppose that current LTE systems support a transmission power of 23 dBm. With analog cancellation that suppresses self-interference by approximately 60 dB [3, 5, 6], digital cancellation has the burden of self-interference to cancel of 50 dB (≈ 23 dBm – (–90 dBm) –60 dB). Digital cancellation, however, is limited by intermodulation distortion (IMD) caused by a power amplifier. Therefore, several techniques, called nonlinear digital self-interference cancellation, have been proposed to cancel the self-interference with IMD.

Nonlinearity Models

As explained earlier, the first step in digital self-interference cancellation is determining what model to adopt. For decades, a challenging problem has been overcoming the nonlinearity of power amplifiers. Several theoretical approaches for modeling the nonlinearity have been developed such as the Wiener model and the Hammerstein model. In this article, we introduce the parallel Hammerstein model adopted in prior work [5, 6]. The parallel Hammerstein model is expressed as:

$$y[n] = \sum_{k=0}^{K-1} \sum_{\ell=0}^{L-1} b_{k,\ell} \left| x[n-\ell] \right|^{2k} x[n-\ell],$$

where $x[n]$ and $y[n]$ are the power amplifier's time domain complex baseband input and output signals, $\{b_{k,\ell}\}$ are the coefficients of the model, and $2K - 1$ and L represent the order of the model and the number of the model's taps, respectively. Note that the number of $\{b_{k,\ell}\}$ is KL. This model is constructed of odd-order terms of the input signal because in wireless communication systems the only thing considered is a passband signal near a center frequency. It can be inferred that the model becomes memoryless if $L = 1$ and linear if $K = 1$.

Expressed in terms of the complex baseband signals, the parallel Hammerstein model can be affected by I/Q imbalances. The effect of I/Q imbalances can be avoided by employing the real-valued model introduced in [11]. The authors in [6] introduced the polynomial basis function,

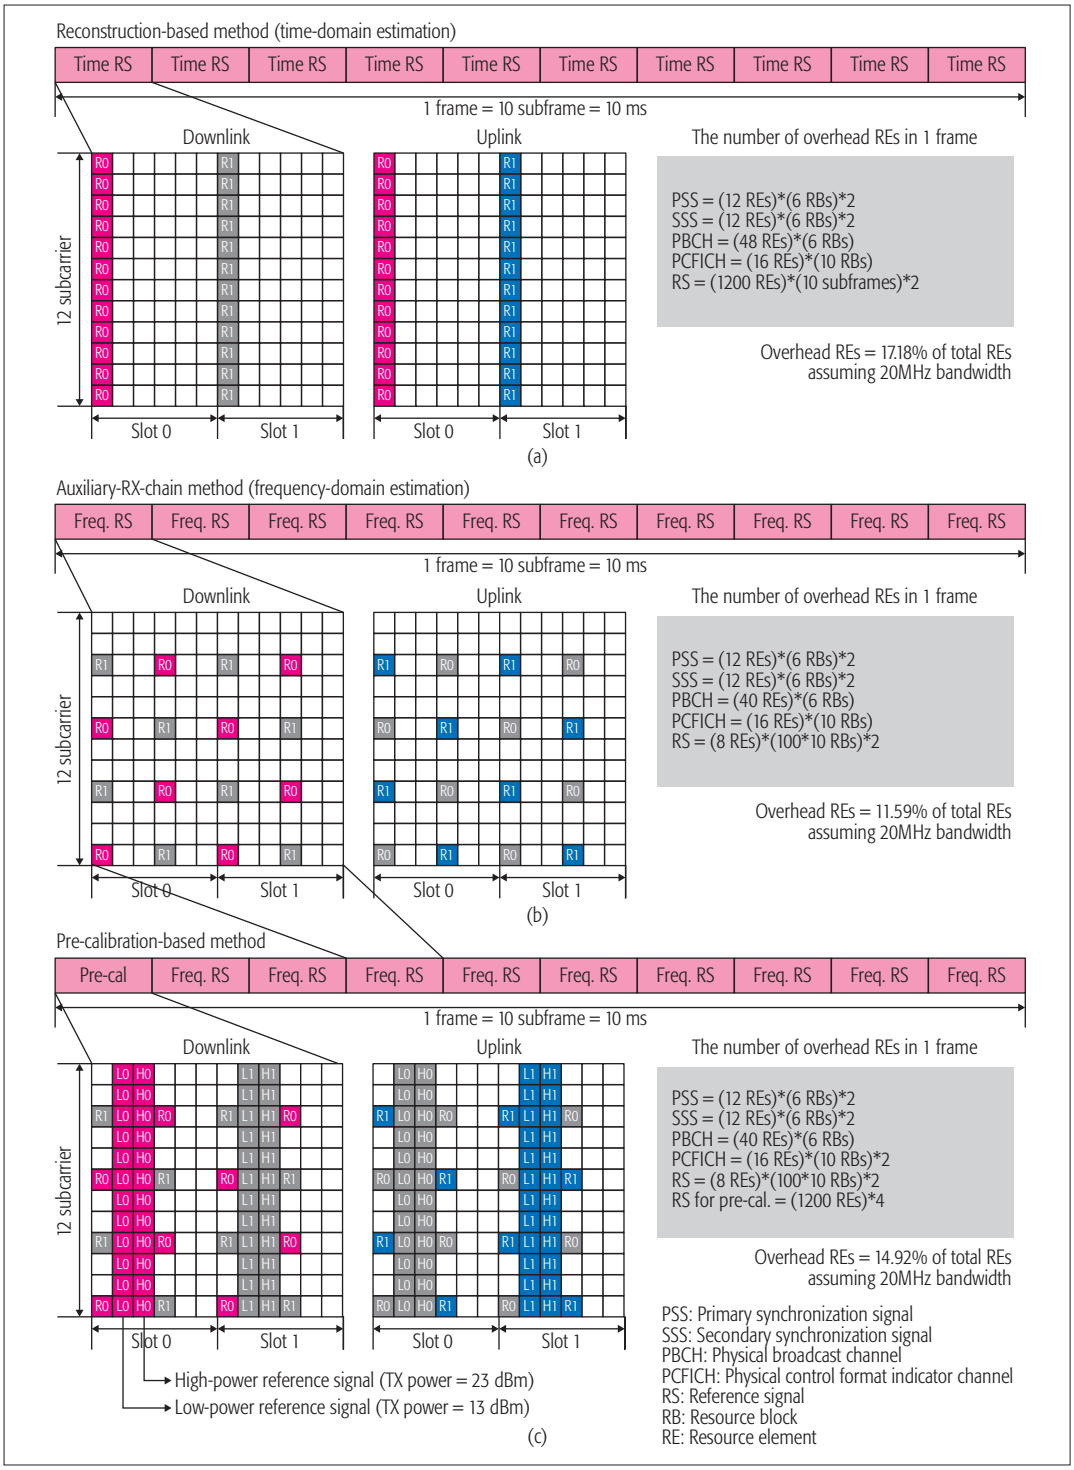

Figure 3. The frame structure, reference signal allocation, and overhead in terms of resource-element ratio based on the LTE standard for each nonlinear digital self-interference cancellation: a) the reconstruction-based method; b) the auxiliary-receive-chain method; c) the pre-calibration-based method.

Due to its simplicity, many researchers use a linear model for a wireless channel. Therefore, self-interference channels are assumed to be linear, and linear self-interference cancellation becomes the basic technology of digital self-interference cancellation.

which includes the effect of I/Q imbalances. In this article, we assume that the system follows the parallel Hammerstein model, which can be directly applied to LTE systems.

Conventional Nonlinear Cancellation Techniques

In contrast to cancelling linear components, canceling nonlinear components calls for extra resources such as hardware, pilot overhead, and/or computational complexity. In this section, we introduce two representative conventional nonlinear digital cancellation techniques, and introduce their strengths and weaknesses. Figure 4 shows the schematic block diagrams and the number of parameters constructing each model. $x'[n]$ is the measured signal via an auxiliary receive chain. L_{ch} is the delay spread of a self-interference channel, L_{pa} is the number of memory taps of a power amplifier, and N_{sub} is the number of used subcarriers.

A pre-calibrator estimates the nonlinearity of a power amplifier, and modifies the input signal of the power amplifier to linearize the output signal of the power amplifier. Put simply, since the signal with a high amplitude is saturated by the power amplifier, a pre-calibrator strengthens the high-amplitude-signal more than the low-amplitude-signal.

Reconstruction-Based Method: This method follows a typical digital cancellation methodology: estimate and reconstruct the nonlinear self-interference signal, and subtract it from the received signal [5, 6]. To estimate the nonlinearity — similar to linear digital cancellation in the time domain — the reference signal allocation shown in Fig. 3a and the least squares method are adopted. Two signals used to estimate coefficients are the signal before passing a power amplifier and the received signal. Unlike linear channel estimation, the model of self-interference signal is reformulated with odd-order terms, and the number of the coefficients increases K-fold. The larger number of coefficients causes more complex computation to estimate them. Furthermore, if the number of coefficients is greater than the number of the samples of the reference signal, it makes the least squares problem something that should be avoided: an undetermined system.

The major weakness of this method is that the nonlinearity model consists of both a power amplifier and a wireless self-interference channel. This is why the number of the coefficients is abnormally large. Figure 4a shows the number of parameters of the power amplifier, wireless channel, and joint nonlinearity model. The taps of the nonlinearity model are caused not only by the memory effect of the power amplifier, but mostly by the reflections of the wireless channel. Therefore, the number of coefficients is large, and real-time cancellation becomes difficult due to computational complexity. Furthermore, the wireless channel makes the total nonlinear model time-varying. In other words, the nonlinear model needs to be estimated more frequently; thus, the overhead of the reference signal increases. One might argue that the overhead can be reduced by utilizing the data signal, which is always known to the transceiver, as a reference signal. However, since the reference signal is also needed for another node that is trying to receive the signal, we can continue as if there is no extra reference signal for self-interference cancellation.

Auxiliary-Receive-Chain Method: The nonlinear self-interference can be cancelled by using an auxiliary receive chain [9]. Through the auxiliary receive chain, as illustrated in Fig. 4b, the signal distorted by a power amplifier is directly obtained rather than estimated or reconstructed. This distorted signal is used to estimate the self-interference channel and to reconstruct the self-interference signal. Recall that conventional linear digital self-interference cancellation methods obtain the reference signal before a power amplifier distortion. The limitation of such methods is that the self-interference cannot be reconstructed accurately since it is assumed to pass through a linear system, which is in fact nonlinear. On the other hand, by using the power-amplifier-distorted signal as the reference signal, it can be assumed that this signal passes through a linear system: the wireless channel. The estimation and reconstruction schemes are the same as in linear self-interference cancellation, and both the time-domain and frequency-domain processing can be applied. Therefore, there is no need for extra computational resources compared to linear cancellation. The only drawback of this method is the need of the extra receive chain.

The Proposed Pre-Calibration-Based Cancellation Method

To cancel nonlinear self-interference, we propose a pre-calibration-based cancellation technique that linearizes a transmitter and cancels self-interference with linear-only cancellation at a receiver. A pre-calibrator, as illustrated in Fig. 4c, estimates the nonlinearity of a power amplifier, and modifies the input signal of the power amplifier to linearize the output signal of the power amplifier. Put simply, since the signal with high amplitude is saturated by the power amplifier, a pre-calibrator strengthens the high-amplitude signal more than the low-amplitude signal. Then the combined system of the pre-calibrator and the power amplifier is linearized. As IMD — which limits the linear self-interference cancellation — decreases, the linearized self-interference signal can be cancelled by the linear digital self-interference cancellation technologies. Furthermore, a receiver can achieve higher signal-to-noise ratio (SNR) with improved transmit error vector magnitude (EVM).

Frame Structure of Reference Signal: A frame structure of a reference signal, as illustrated in Fig. 3c, consists of two parts. The first part is for calculating a pre-calibrator. This reference signal is for measuring the characteristic of the power amplifier. Since the characteristic of the power amplifier is more static than that of a wireless channel, this reference signal can be allocated with a long period compared to the coherence time of the wireless channel. For instance, in this article, we allocate this reference signal every 10 ms (1 frame in the LTE standard), as illustrated in Fig. 3c. A detail of this reference signal is explained in a later section. The second part is for data transmission and linear cancellation. Due to its low complexity and low overhead, as explained earlier, the reference signal allocation shown in Fig. 3b and linear cancellation in the frequency domain are adopted.

Calculating a Pre-Calibrator: To measure the output of the power amplifier without an extra wire, we propose the following technique. To model the power amplifier, two signals are needed: the input signal, and the output signal of the power amplifier. Note that what we need is the output signal of the power amplifier that does not pass through the wireless channel. Therefore, we decide to exploit the accurately estimated wireless channel and remove its effect from the received signal by equalization. We suggest that the reference signal contains a low-power signal and a high-power signal, as illustrated in Fig. 3c. The low-power reference signal (denoted by L0/L1) aims for accurate wireless channel estimation. Thus, to avoid distortion, it operates on a linear region (low-power) of the power amplifier, and to lessen estimation errors, it occupies all subcarriers. The high-power reference signal (denoted by H0/H1) is for estimating the power amplifier's distortion. This signal experiences a nonlinear region (high-power) of the power amplifier and undergoes sufficient distortion. Then, using the precisely estimated wireless channel, the output of the power amplifier can be calculated.

There are several conventional methods to calculate a pre-calibration function. The method used in this article is to estimate coefficients of

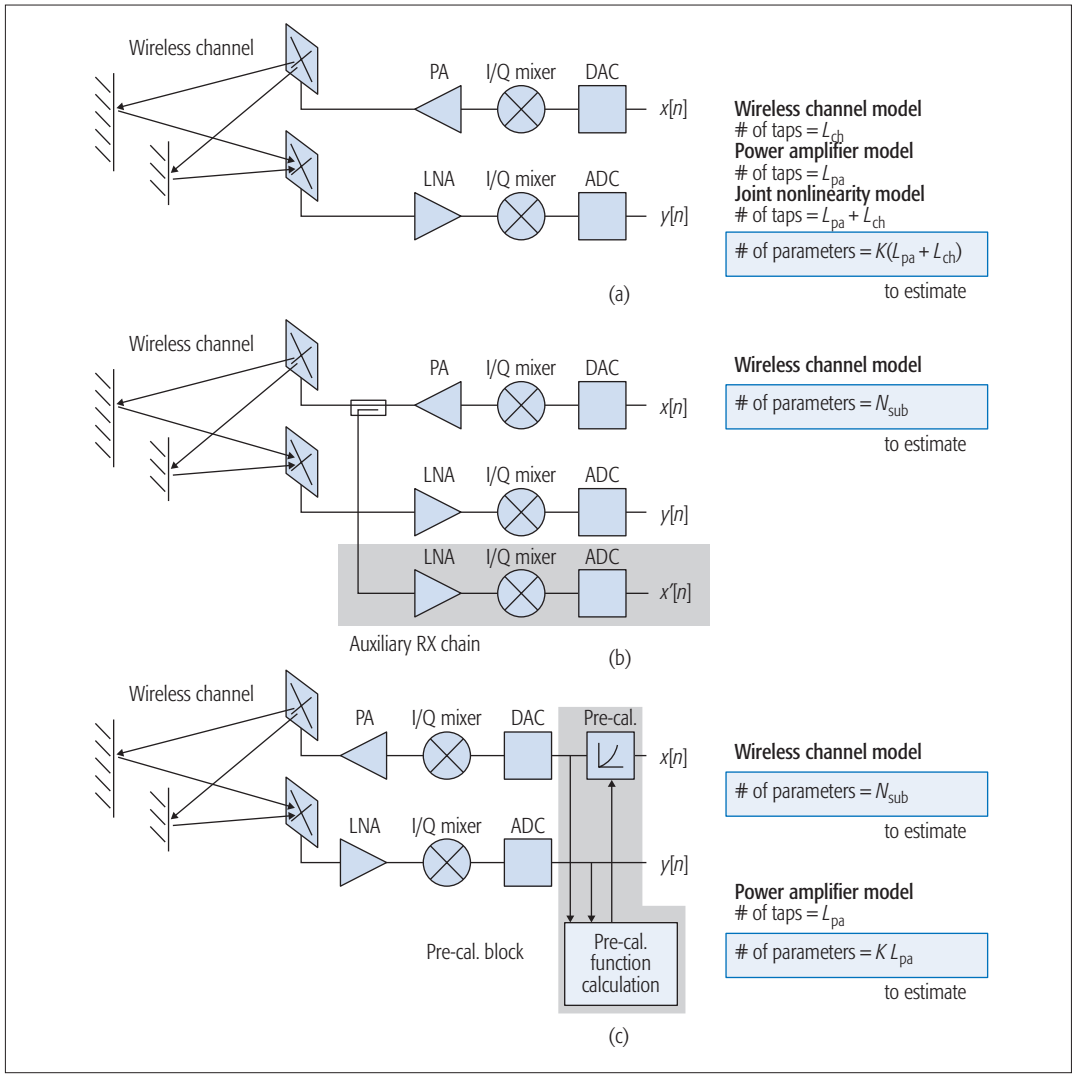

Figure 4. The schematic block diagram of each nonlinear digital self-interference cancellation method, and the number of parameters constructing each model: a) the reconstruction-based method; b) the auxiliary-receive-chain method; c) the pre-calibration-based method.

a polynomial-based pre-calibration function. It is just like estimating a nonlinearity model, but uses the input of the power amplifier as output and, vice versa, to estimate a reversed function. The pre-calibration function is obtained from the reversed function of the power amplifie scaled by the power amplifier gain.

Strengths of Proposed Method: The proposed method has two main strengths compared to the conventional methods. First, there is no need for an extra receive chain or wire. Techniques for power amplifier linearization have been applied to most conventional radios with high power. One of the effective technologies to reduce the nonlinearity of a transmitter is digital pre-distortion (DPD) [12]. However, most DPD systems need a secondary receive chain to estimate the output of the power amplifier [12]. In [10], a full-duplex system that measures the effect of the power amplifier through a cable and applies it to an auxiliary transmit signal was proposed. With full-duplex systems, since a receiver is able to sense a transmit signal without saturation, a pre-calibration function can be calculated without an extra receive chain and cable.

Second, complex and time-varying nonlinear cancellation is not required. Contrary to the reconstruction-based method, the pre-calibration-based method takes advantage of the nonlinear model of the power amplifier without the effect of wireless channels. It is therefore expected that the coefficients $b_{k,\ell}$ are more static and fewer in number, which lessens the burden of computational complexity and the reference signal overhead. For example, with the proposed reference signal illustrated in Fig. 3c, as shown in Fig. 4, KL_{pa} coefficients need to be estimated every 10 ms, while the reconstruction-based method estimates $K(L_{pa} + L_{ch})$ coefficients every 1 ms.

Performance Evaluation

We evaluate the nonlinear digital self-interference cancellation techniques from two points of view: link-level evaluations for measuring nonlinearity of a power amplifier and cancellation amounts, and system-level evaluations for throughput gain. A basic criterion of self-interference cancellation technologies is the amount of cancellation. This is important because it guarantees the feasibility of full-duplex radios.

There are several conventional methods to calculate a pre-calibration function. The method used in this article is to estimate coefficients of a polynomial-based pre-calibration function. It is just like estimating a nonlinearity model, but uses the input of the power amplifier as output and vice versa to estimate a reversed function.

The link-/system-level analyses were based on computer simulations. To obtain a realistic power amplifier model, we measured the characteristics of a power amplifier via an LTE-based software-defined radio platform that we developed. To model a 3D indoor environment and obtain path-losses between nodes, a 3D ray-tracing tool was employed.

Simple link-level evaluations, however, do not fully explain why nonlinear digital cancellation is needed despite its high complexity. Most prior work has tried to suppress self-interference to the same level as the noise floor, but it is still questionable whether this is actually necessary in practice where there is other interference. Therefore, we also carry out, with a 3D ray-tracing tool, system-level simulations of multi-BSs and multi-mobile stations (MSs) in an indoor environment. The link-/system-level analyses were based on computer simulations. To obtain a realistic power amplifier model, we measured the characteristics of a power amplifier via an LTE-based software-defined radio platform that we developed. To model a 3D indoor environment and obtain path losses between nodes, a 3D ray-tracing tool was employed.

Link-Level Self-Interference Cancellation Performance

To simulate nonlinear digital cancellation, we assume a single-input single-output (SISO) wireless system based on the LTE parameters given in Fig. 5a. To exploit a realistic model for a power amplifier, we measure the nonlinearity of a power amplifier[1] through the PXIe software-defined radio platform[2] introduced in [3]. This platform generates an LTE-based signal with given physical-layer parameters. As shown in Fig. 5b, the input and output (distorted) signals of the power amplifier are measured via the PXIe platform. A 30 dB attenuator is employed to receive the output signal of the power amplifier without distortion from the receiver of the PXIe platform. The GUI of the LTE downlink framework, which is shown in Fig. 5d, shows the received signal has passed through the power amplifier. Figure 5c shows the measured input and output signals of the power amplifier, and a parallel-Hammerstein-model-based approximated power amplifier model for simulations. In this simulation, the power amplifier was modeled as the third order model, which showed the greatest similarity with the measured data. Tending to make the residual self-interference channel frequency selective [7], analog self-interference cancellation is simulated by employing a longer-than-normal delay spread for the residual self-interference channel and attenuating the channel 50 dB.

The three nonlinear digital cancellation techniques introduced earlier were evaluated, and compared to frequency-domain linear digital cancellation. The reference signal of the reconstruction-based method was determined to be repeated every subframe (12 OFDM symbols) considering the Doppler effect and the reference signal overhead, as illustrated in Fig. 3a. For the linear cancellation part of the auxiliary-receive-chain method and the pre-calibration-based method, the frequency-domain-based linear cancellation was adopted, and the reference signal pattern followed the cell-specific reference signal in LTE, as seen in Fig. 3b. The orders of the reconstructed parallel Hammerstein model and the pre-calibration function are 7 and 5, respectively.

We simulated the amount of cancellation with different coherence times. Figures 5e and 5f show how much self-interference power was reduced when the coherence times of the self-interference channel were 7.14 ms and 142.86 ms, respectively. These coherence times were calculated from the Doppler effect, assuming that the center frequency was 2.52 GHz, and MSs were vehicles (60 km/h) or pedestrians (3 km/h). Figure 5g shows the amount of self-interference cancellation with a static self-interference channel.

When the self-interference channel was static, the performance of the reconstruction-based method, as Fig. 5g shows, was about 62 dB — the best performance among the three nonlinear cancellation methods. This occurred because this method can reconstruct self-interference accurately (only limited by noise), but the other two methods were limited by channel estimation error from the interpolating process in the frequency domain. On the other hand, it is observed that its cancellation performance was 43 dB when the coherence time was 142.86 ms, and when the coherence time was 7.13 ms, even worse than that of linear self-interference cancellation. Therefore, we conclude that in practice it is not proper to cancel self-interference by reconstructing, due to its weakness in the time-varying channel.

The auxiliary-receive-chain method and the pre-calibration-based method are robust to fading channel. Figure 5e shows that even though the coherence time was small, these methods could cancel nonlinear self-interference by about 43 dB. This result is quite obvious because the reference signals exploit the very same structure of LTE. The performance of these two methods degrades as coherence time decreases because of an increase in the channel estimation error from the interpolating process in the time domain. The auxiliary-receive-chain method behaves as a performance upper bound of the pre-calibration-based method because the pre-calibrator does not linearize ideally. One might argue that the performance of the pre-calibration-based method is not as stable as that of the auxiliary-receive-chain method. This is because the pre-calibration function generator, which calculates in every frame, fails to generate the well-operating pre-calibration function in some frames. By applying adaptive polynomial filters [13] or orthogonal polynomials [14], we can have a better chance of improving performance. We leave this for future work.

System-Level Throughput Gain

We investigated the need for nonlinear digital self-interference cancellation technologies from a system throughput perspective. For system-level evaluations, we used Wireless System Engineering (WiSE), a 3D ray-tracing tool developed by Bell Labs [15, 16]. For an indoor environment, we modeled the building structure of Veritas Hall C of Yonsei University in Korea, shown in Fig. 6a. Each BS is modeled with the measured radiation pattern of the dual-polarized antenna introduced in [3, 17] (Fig. 6c). On the third floor, there were uniformly distributed MSs, each of which was equipped with an isotropic antenna; these were associated with the BS that provided the strongest downlink power. Figure 6b illustrates how five BSs were deployed on the third floor, and how each cell's coverage was determined. Note that due to its radiation pattern, each BS has its own direction, represented as an arrow in Fig. 6b. For evaluations, we assumed that there were five MSs,

[1] Mini-Circuits ZVA-183W+ Super Ultra Wideband Amplifier, http://www.minicircuits.com/pdfs/ZVA-183W+.pdf

[2] With this platform, we developed and demonstrated the real-time full-duplex SISO and MIMO systems at IEEE GLOBECOM 2014 and IEEE GLOBECOM 2015, respectively. Full demo video clips are available at http://www.cbchae.org/

Parameter	Value
Center frequency	2.52 GHz
Bandwidth	20 MHz
FFT size	2048
Used subcarrier	1200
CP length	512
TX power	23 dBm
Power amplifier gain	29 dB
Power amplifier P1dB	28 dBm
Analog cancellation	50 dB
Noise floor	-90 dBm
Resolution of the ADC	14 bits

(a)

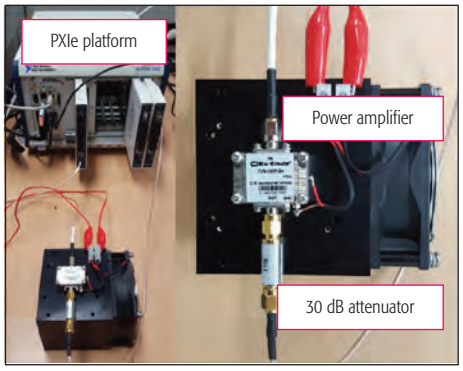

(b)

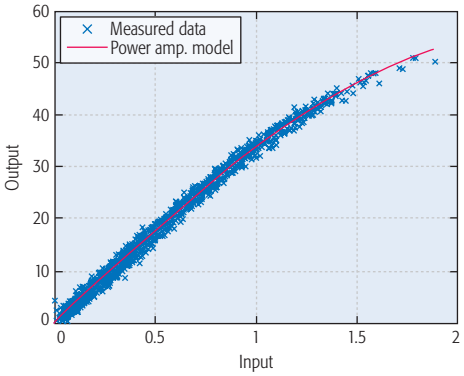

(c)

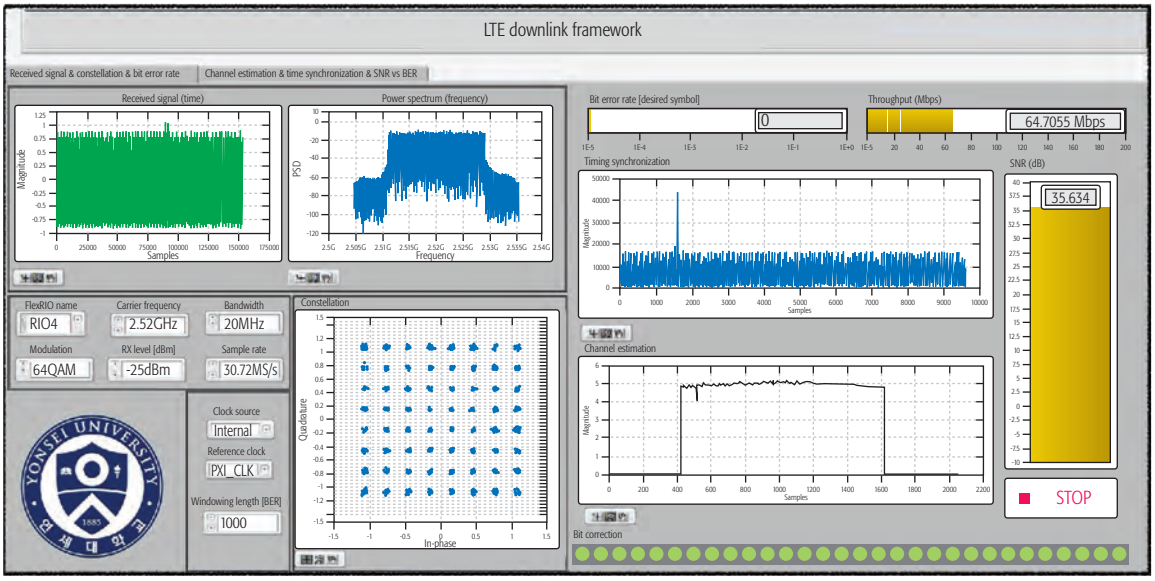

(d)

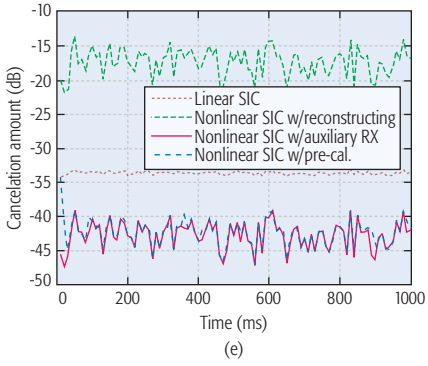

(e)

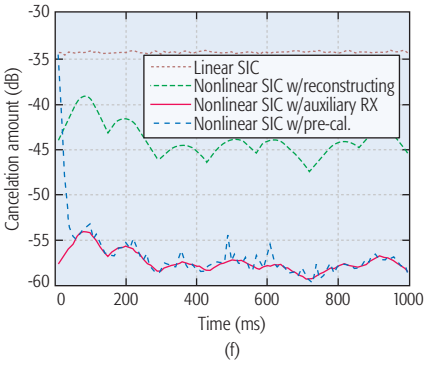

(f)

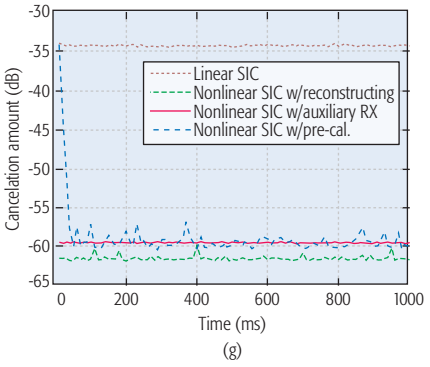

(g)

Figure 5. a) The parameters for link- and system-level simulations; b) on the left side is a setup for measuring the characteristic of the power amplifier, and on the right side is the used power amplifier and the attenuator; c) measured input and output signals of the power amplifier, and an approximated model of the power amplifier; d) the GUI of the LTE downlink framework used for measuring the model of the power amplifier; e, f) the results of the link-level simulations with the coherence time of 7.13 ms and 142.86 ms; g) the result of the link-level simulation with a constant wireless channel.

which were located in different cells' coverage, and adopted the parameters in Fig. 5a.

In system-level simulations, we investigated the throughput of a half-duplex downlink system based on FDD LTE and full-duplex systems with different self-interference cancellation levels. For cancellation performance, we exploited the average of the link-level simulation results with the coherence time of 142.86 ms, and combined them with the analog cancellation of 50 dB given in Fig. 5a. The exact cancellation amounts were 84.28 dB for linear cancellation, 93.03 dB for reconstruction-based nonlinear cancellation, 107.17 dB for auxiliary-receive chain nonlinear cancellation, and 106.20 dB for pre-calibration-based nonlinear cancellation. With the

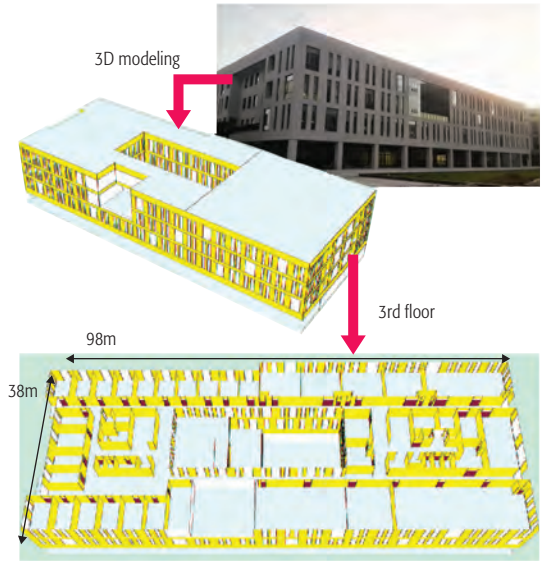

(a)

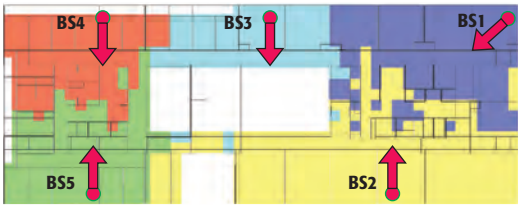

(b)

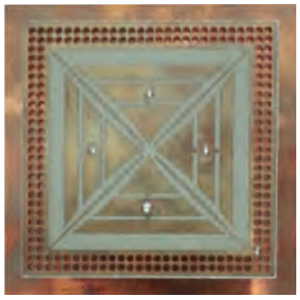

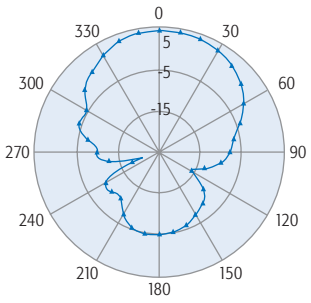

(c)

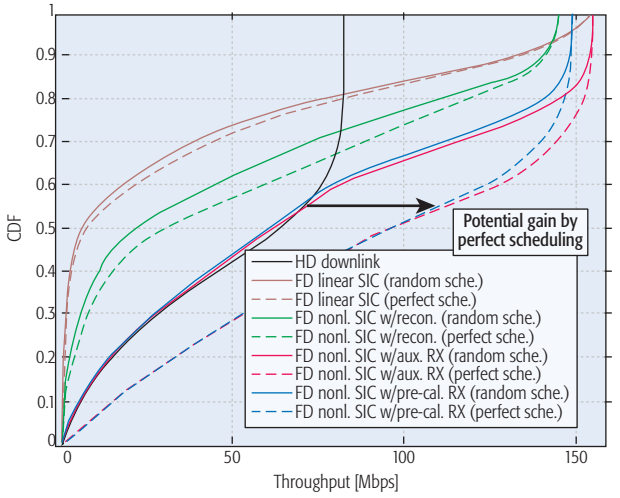

(d)

[Mbps]	HD downlink	FD linear SIC (random sche.)	FD linear SIC (perfect sche.)	FD nonl. SIC w/recon. (random sche.)	FD nonl. SIC w/recon. (perfect sche.)	FD nonl. SIC w/aux. RX (random sche.)	FD nonl. SIC w/aux. RX (perfect sche.)	FD nonl. SIC w/pre-cal. (random sche.)	FD nonl. SIC w/pre-cal. (perfect sche.)
Average	51.80	35.28	36.98	48.10	53.67	72.45	92.47	70.23	89.96
Median	63.50	5.89	7.37	21.50	29.73	61.63	97.76	60.25	95.68
10-percentile	5.35	0.13	0.17	0.63	1.08	4.72	14.62	4.74	14.95
90-percentile	82.16	133.16	134.17	140.32	141.23	153.96	154.44	148.19	148.64

(e)

Figure 6. a) Topology for system-level performance evaluations; b) BS deployments and cell coverage; c) dual-polarized antenna and radiation pattern; d) CDF of the system throughput; e) throughput results. Here, "random scheduling" means random user selection, and "perfect scheduling" means Genie-aided ideal user selection.

received power of signal of interest and the summation of the received power of interference and noise, the signal-to-interference-plus-noise ratio (SINR) of each node was calculated. Finally, a system bandwidth of 20 MHz and the overheads introduced in Fig. 3 were applied. Note that there was also the overhead for the extended CP. We also simulated the full-duplex systems without considering MS-to-MS interference as upper bounds, which could be achieved by Genie-aided perfect scheduling.

Figure 6d illustrates the results of the ergodic throughput of the half-duplex system and the full-duplex systems with different self-interference levels. The solid lines of full-duplex indicate results with MS-to-MS interference, and the dashed lines indicate those without it. Figure 6e gives the representative values of each case such as average, median, 10-percentile, and 90-percentile. The result implies that the throughput of a full-duplex system depends heavily on the performance of self-interference cancellation. If the full-duplex system employs only linear self-interference cancellation, the average throughput is lower than that of the half-duplex system, and only about 20 percent of MSs can experience the benefit of full duplex. On the other hand, with the well-performing nonlinear self-interference cancellation techniques such as the auxiliary-receive-chain method and the pre-calibration-based method, the average throughput increased by approximately 35 to 40 percent over that of a half-duplex system. Even though self-interference is cancelled out sufficiently, we can observe that for certain portions of MSs, half-duplex outperforms full-duplex. Furthermore, if the interference between MSs was avoided somehow, the average throughput could be improved by up to 79 percent. From this, we have the insight that once self-interference cancellation techniques guarantee certain performances, a critical issue in full-duplex research will become user allocation and user scheduling.

Conclusion

Full-duplex radio is expected to play a major role in enhancing the spectral efficiency in 5G wireless communications/LTE. In this article, we have investigated two existing nonlinear digital cancellation techniques and proposed a low-complexity pre-calibration-based technique. Link-level and system-level performance have been analyzed through a real-time software-defined radio platform and 3D-ray-tracing-based simulations of an indoor environment. The results of our analysis confirm a significant performance enhancement even in interference-limited environments. We expect our study to provide insights into developing practical cellular systems based on full-duplex radios.

Acknowledgment

The authors would like to thank Mr. Y.-G. Lim for helpful discussions on WiSE simulations.

References

[1] M. S. Sim *et al.*, "Low-Complexity Nonlinear Self-Interference Cancellation for Full-Duplex Radios," *Proc. IEEE GLOBECOM FDWC Wksp.*, Dec. 2016.
[2] 5G Forum, "5G Vision, Requirements, and Enabling Technologies," v. 2.0, Mar. 2016.
[3] M. Chung *et al.*, "Prototyping Real-Time Full Duplex Radios," *IEEE Commun. Mag.*, vol. 53, no. 9, Sept. 2015, pp. 56–63.
[4] M. Jain et al., "Practical, Real-Time, Full Duplex Wireless," *Proc. ACM MobiCom*, 2011, pp. 301–12.
[5] D. Bharadia, E. McMilin, and S. Katti, "Full Duplex Radios," *Proc. ACM SIGCOMM*, Aug. 2013, pp. 375–86.
[6] M. Heino *et al.*, "Recent Advances in Antenna Design and Interference Cancellation Algorithms for In-Band Full Duplex Relays," *IEEE Commun. Mag.*, vol. 53, no. 5, May 2015, pp. 91–101.
[7] M. Duarte *et al.*, "Design and Characterization of a Full-Duplex Multiantenna System for WiFi Networks," *IEEE Trans. Vehic. Tech.*, vol. 63, no. 3, Mar. 2014, pp. 1160–77.
[8] E. Aryafar et al., "MIDU: Enabling MIMO Full Duplex," *Proc. ACM MobiCom*, Aug. 2012, pp. 257–68.
[9] E. Ahmed and A. M. Eltawil, "All-Digital Self-Interference Cancellation Techniques for Full-Duplex Systems," *IEEE Trans. Wireless Commun.*, vol. 14, no. 7, July 2015, pp. 3519–32.
[10] A. K. Khandani, "Full-Duplex (Two-Way) Wireless: Antenna Design and Signal Processing," tech. rep.; http://cst.uwaterloo.ca/reports/antenna design.pdf.
[11] Y. Hua *et al.*, "Radio Self-Interference Cancellation by Transmit Beamforming, All-Analog Cancellation and Blind Digital Tuning," *Signal Processing.*, vol. 108, Mar. 2015, pp. 322–40.
[12] W.-J. Kim *et al.*, "Digital Predistortion Linearizes Wireless Power Amplifiers," *IEEE Microwave Mag.*, vol. 6, no. 3, Sept. 2005, pp. 54–61.
[13] V. J. Mathews, "Adaptive Polynomial Filters," *IEEE Signal Processing. Mag.*, vol. 8, no. 3, July 1993, pp. 10–26.
[14] R. Raich and G. T. Z. Zhou, "Orthogonal Polynomials for Complex Gaussian Processes," *IEEE Trans. Signal Processing*, vol. 52, no. 10, Oct. 2004, pp. 2788–97.
[15] R. A. Valenzuela, D. Chizhik, and J. Ling, "Measured and Predicted Correlation between Local Average Power and Small Scale Fading in Indoor Wireless Communication Channels," *Proc. IEEE VTC-Spring*, May 1998, pp. 2104–08.
[16] J. Jang *et al.*, "Smart Small Cell for 5G: Theoretical Feasibility and Prototype Results," *IEEE Wireless Commun.*, vol. 23, no. 6, Dec. 2016, pp. 124–31.
[17] T. Oh *et al.*, "Dual-Polarization Slot Antenna with High Cross Polarization Discrimination for Indoor Smallcell MIMO Systems," *IEEE Ant. and Wireless Prop. Lett.*, vol. 14, Feb. 2014, pp. 374–77.

Biographies

Min Soo Sim [S'14] received his B.S. degree from the School of Integrated Technology, Yonsei University, Korea, in 2014. He is now with the School of Integrated Technology at the same university and is working toward a Ph.D. degree. He was the recipient of the Silver Prize in the 22nd Humantech Paper Contest. His research interest includes emerging technologies for 5G communications. He was the recipient/co-recipient of two Silver Prizes in the 22nd/23rd Humantech Paper Award.

MinKeun Chung [S'11] received his B.S. degree from the School of Electrical and Electronic Engineering, Yonsei University, Korea, in 2010. He is now working toward a Ph.D. degree under the joint supervision of Prof. D. K. Kim and Prof. C.-B. Chae. He did his graduate internship in the advanced wireless research team at National Instruments, in Austin, Texas, in 2013. His research interests include the design and implementation of architectures for next-generation wireless communication systems.

Dongkyu Kim [M'13] received his B.S. degree in electrical engineering from Konkuk University, Seoul, South Korea, in 2006, and his M.S. and Ph.D. degrees in electrical and electronic engineering from Yonsei University in 2008 and 2013, respectively. From 2013 to 2014, he was a postdoctoral researcher with the Information and Telecommunication Laboratory, Yonsei University. Since 2014, he has been with LG Electronics Inc. as a senior researcher, where he was involved in development of advanced wireless technologies, including 5G mobile communications and 3GPP standards for future wireless systems. His current research activities are focused on future wireless communication, including flexible and full-duplex radio, V2X, massive MIMO, and mmWave technologies.

Jaehoon Chung is with LG Electronics Inc., where he has been involved in development of advanced wireless technologies, including 5G mobile communications and 3GPP standards for future wireless systems. His current research activities are focused on future wireless communication, including flexible and full-duplex radio, V2X, massive MIMO, and mmWave technologies.

Dong Ku Kim [SM'15] received his B.S. from Korea Aerospace University in 1983, and his M.S. and Ph.D. from the University of Southern California, Los Angeles, in 1985 and 1992, respectively. He worked on CDMA systems in the cellular infrastructure group of Motorola at Fort Worth, Texas, in 1992. He has been a professor in the School of Electrical and Electronic Engineering, Yonsei University, since 1994. Currently, he is a Vice President for Academic Research Affairs of the Korean Institute of Communications and Information Systems (KICS). He has been a Vice Chair of the Executive Committee of the 5G Forum since 2013. He received the Minister Award for Distinguished Service for ICT R&D from the Ministry of Information, Science, and Future Planning in 2013, and the Award of Excellence in leadership of 100 Leading Core Technologies for Korea 2020 from the National Academy of Engineering of Korea. He received the Dr. Irwin Jacobs Academic Achievement Award 2016 from Qualcomm and KICS.

Chan-Byoung Chae [SM'12] is an associate professor in the School of Integrated Technology, Yonsei University. Before joining Yonsei University, he was with Bell Labs, Alcatel-Lucent, Murray Hill, New Jersey, as a member of technical staff, and Harvard University, Cambridge, Massachusetts, as a postdoctoral research fellow. He received his Ph.D. degree in electrical and computer engineering from the University of Texas at Austin in 2008. He was the recipient/co-recipient of the IEEE INFOCOM Best Demo Award (2015), the IEIE/IEEE Joint Award for Young IT Engineer of the Year (2014), the KICS Haedong Young Scholar Award (2013), the *IEEE Signal Processing Magazine* Best Paper Award (2013), the IEEE ComSoc AP Outstanding Young Researcher Award (2012), the IEEE Dan. E. Noble Fellowship Award (2008), and two Gold Prizes (1st) in the 14th/19th Humantech Paper Contest. He currently serves as an Editor for *IEEE Transactions on Wireless Communications*, the *IEEE/KICS Journal on Communications Networks*, and *IEEE Transactions on Molecular, Biological, and Multi-Scale Communications*.

Full-duplex radio is expected to play a major role in enhancing the spectral efficiency in 5G wireless communications/LTE Evolution. The results of our analysis confirmed a significant performance enhancement even in interference-limited environments. We expect our study to provide insights into developing practical cellular systems based on full-duplex radios.

ACCEPTED FROM OPEN CALL

Buffer-Aided Relay Systems under Delay Constraints: Potentials and Challenges

Deli Qiao and M. Cenk Gursoy

The authors consider the buffer-aided relay systems under delay constraints specified by limitations on delay violation probability. They discuss the necessary research tools and several results on different relay networks including two-hop and three-node relay channels. They outline challenges associated with the buffer-aided relaying under delay constraints and topics for future research.

ABSTRACT

Relay transmissions can help improve system coverage and throughput, and hence information-theoretic analysis of relay channels has been at the forefront of research for decades. While providing powerful results, information-theoretic studies generally assume no buffer at the relay, and the performance is limited by the worst conditions of the transmitting and receiving channels of the relay. Recently, it has been shown that the achievable throughput can be significantly improved with the introduction of a buffer-aided relay model. Specifically, buffer-aided relaying can provide significant gains in terms of throughput, diversity, and signal-to-noise ratio without delay constraints. This article considers buffer-aided relay systems under delay constraints specified by limitations on delay violation probability. The necessary research tools and several results on different relay networks including two-hop and three-node relay channels are discussed. Certain challenges associated with buffer-aided relaying under delay constraints and topics for future research are outlined.

INTRODUCTION

Cooperative communication is one fundamental component of current and next generation wireless systems [1]. Cooperative communication is achieved via some nodes acting as relays such that they can share their resources with other nodes and help each other's communication. Generally, the relay nodes receive, process and retransmit the information sent by the source node, and hence improve transmission between the source and destination. Information-theoretic analysis of relay channels has been at the forefront of research forefront, and has revealed the benefits of relay systems in terms of capacity and coverage. Different relay policies have been proposed such as decode-and-forward (DF) relaying, in which the relay decodes the source message from its received signal, re-encodes it into a new codeword, and transmits it later; and amplify-and-forward (AF) relaying, where the relay amplifies its received signal and transmits it immediately. However, for such relay policies, performance of the relay is limited by the worst channel between the transmitting and receiving channels of the relay. In practical systems, due to various factors, such as mobility, changing environment and multipath fading, the transmitting and receiving channels of the relay vary significantly with time, thus hindering the system from utilizing the resources efficiently.

Recently, studies on relay systems involving the buffer at the relay have yielded relay policies that can take advantage of the buffer at the relay and further improve system performance, and introduced new degrees of freedom for system design [2]. This is generally due to the fact that the relay can store the information such that data transmission at the relay will not be limited by the unfavorable channel conditions of its transmitting or receiving channels [3–6]. For example, for a multirelay network, a buffer-aided half-duplex relay selection protocol, where the link with the largest weighted instantaneous channel capacity is selected, can significantly increase the throughput [4].

However, the buffer at the relay can introduce additional delay to the communication between the source and destination. In current wireless systems, multimedia traffic such as mobile video and voice over IP (VoIP), has surged significantly, and delay has become an important consideration. Guaranteeing the delay constraints of various applications is a key design issue [7]. Thus, the benefits of buffer-aided relay under delay constraints need further investigation for practical applications. Because of the time-varying nature of wireless channels, the deterministic delay constraints for multimedia wireless traffic are usually hard to guarantee, or when it is possible, requires the system to operate pessimistically and achieve low performance, for example, the delay limited capacity of Rayleigh fading channel is zero. In [4–6], average delay constraints are incorporated, and it has been shown that comparable performance with that of delay-unconstrained systems can be achieved with the proposed relay strategies, which indicates that buffering relays help improve the throughput under the average delay constraints.

In this article, we follow a different approach and consider statistical delay constraints, the study of which has led to favorable and powerful techniques for provisioning and characterization of delay guarantees for wireless traffic [7]. The basic approach in the study of performance under statistical delay constraints inherits from the large deviation techniques such that we can guarantee that the delay violation probabilities decay exponentially with the delay bound. Then, an approximate probability density function of the random asymptotic delay can be obtained, which can be

Digital Object Identifier:
10.1109/MCOM.2017.1600213

Deli Qiao is with East China Normal University; M. Cenk Gursoy is with Syracuse University.

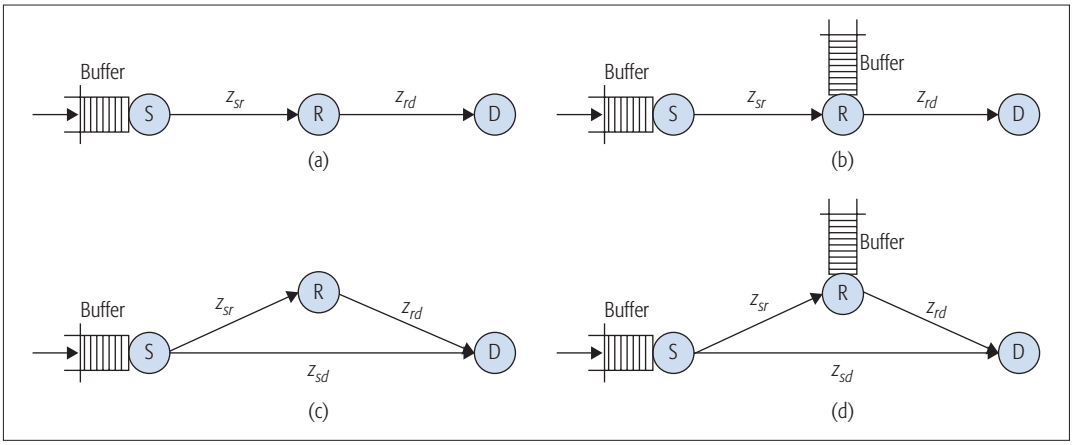

Figure 1. System Model. The left figures refer to the conventional relaying systems without buffer, and the right figures are the buffer-aided relay systems. z_{sr}, z_{rd}, and z_{sd} represent the channel gain of the links **S-R**, **R-D**, and **S-D**, respectively: a) two-hop channel; b) buffer-aided two-hop channel; c) three-node relay channel; d) buffer-aided three-node relay channel.

Most existing studies consider half-duplex relaying due to the difficulties in realizing full-duplex relaying. Fortunately, recent progress in full-duplex relaying has enabled the full-duplex radio to be applicable. Therefore, studies on full-duplex relaying have revived recently, and we should also consider full-duplex relaying protocols.

extended to analyze the asymptotic delay of concatenated queues. Note that with the probability density function of each random variable describing the delay experienced in each buffer, we can characterize the end-to-end delay behavior of the relay networks composed of concatenated queues of the source and relay nodes through the summation of the random variables of the buffers. We can identify and analyze the throughput and relay policies of the relay networks under delay constraints by employing the notions of effective bandwidth [8], which identifies the minimum bandwidth to support a given arrival process, and effective capacity [9], which is the dual concept of effective bandwidth characterizing the maximum arrival rate that can be supported by a given service process of one queue.

In addition to the buffer-aided relay, another research direction taking advantage of the memories across the networks is caching, where the data is stored in the cache of transmitters and receivers [10]. In this setting, part of the files can be transmitted to the users during *idle* hours, stored at the users and exploited to reduce the required data rate during *busy* hours. While physical layer analysis is important for buffer-aided relay systems, the application layer analysis involving file placement, file popularity and cache size are more relevant to caching studies. In this article, we discuss buffer-aided relay networks under delay constraints, and introduce statistical delay provisioning to identify and analyze the throughput and relay policies of the relay wireless networks. Using the results of effective bandwidth and effective capacity theory, we propose the statistical delay trade-off for two concatenating queues, and develop a method for characterizing the performance of the buffer-aided relay systems under end-to-end delay constraints. We also evaluate the system performance of two relay models, namely two-hop and three-node relay channels. We show that buffer-aided relaying is still beneficial under certain delay constraints. Note that the method is generic, and hence can be used to deepen our understanding of buffer-aided relay networks under delay constraints in different applications, such as device-to-device, sensor, and cognitive radio networks.

The organization of this article is as follows. First, we present the typical buffer-aided relay systems. Then, we elaborate on the statistical delay constraints in buffer-aided relay systems. We then consider two relay models, the two-hop channel and the three-node relay channel, and show the methodology to derive their throughput under statistical delay constraints. We also provide several performance evaluations of the discussed relay channels. Afterward, we discuss certain practical challenges associated with the practical use of buffer-aided relay under delay constraints, and some future research topics. Finally, we conclude this article.

Buffer-Aided Relaying

Two-Hop Channel

As shown in Figs. 1a and 1b, we consider a two-hop channel model in which the communication between the source node **S** and the destination node **D** is aided by an intermediate relay node **R**. It is assumed that there is no direct link between the source **S** and the destination **D**. We assume that there are buffers at the source and relay nodes with individual buffer constraints. The channels of the different links experience block fading such that the channels stay constant in one block and change independently from one block to another. It is assumed that the arrival data sequences are divided into frames of duration of one block and stored in the buffers before they are transmitted over different links. In the receiving channel, the relay receives the information and stores it in the buffer, which can be retransmitted to the destination in the transmitting channel. For full-duplex relaying, the relay can transmit and receive simultaneously, while for half-duplex relaying, the relays can only transmit or receive in each time and frequency band (TFB). Most existing studies consider half-duplex relaying due to the difficulties in realizing full-duplex relaying. Fortunately, recent progress in full-duplex relaying has enabled the full-duplex radio to be applicable [11]. Therefore, studies on full-duplex relaying have revived recently, and we should also consider full-duplex relaying protocols.

We assume full-duplex DF relaying. In the absence of buffer at the relay as shown in Fig. 1a,

With the aid of large deviation techniques, we can characterize the tail distribution of the queue length. When the effective bandwidth of the arrival processes and the effective capacity of the departure processes are equal, the queue is stable and the tail distribution of the stationary queue length decays exponentially.

the instantaneous rate of the system is confined by the minimum value of the capacities of the links **S-R** and **R-D**. The system is limited by the worst channel condition of the receiving channel z_{sr} and transmitting channel z_{rd} of the relay. The achievable throughput of the system is then given by the expectation of the minimum value. On the other hand, suppose that the buffer can store the received information for transmission later as shown in Fig. 1b. The receiving and transmitting channels of the relay can always work at their instantaneous capacities. The achievable throughput is now decided by the minimum value of the average throughput of the links. Note that the average of the smaller value of two random variable is always less than the expectation of each random variable, and hence we have larger achievable throughput in buffer-aided relay systems. This demonstrates that buffer-aided relaying can improve the achievable throughput of the two-hop channels.

Three-Node Relay Channel

As depicted in Figs. 1c and 1d, in the case of a three-node relay channel, we consider the scenario where the communication between the source **S** and the destination **D** is assisted by a relay node **R**. Now, in the presence of the direct link between the source **S** and the destination **D**, the source **S** can send information directly to the destination **D**. Again, it is assumed that there are buffers at the source and relay nodes with individual buffer constraints. The relay node employs a DF scheme. Still, buffer-aided relaying can improve the achievable throughput of the three-node relay channel. It is remarkable that buffer-aided relaying leads to significant performance gains in cooperative communication networks with time varying channels without delay constraints.

Statistical Delay Constraints

One-Hop Delay Constraints

We first consider the delay constraints of the single queue case. We assume a first-in first-out (FIFO) queue of infinite size, and consider statistical delay constraints imposed as limitations on delay violation probability. To capture the queue dynamics, we employ the notion of effective bandwidth, which identifies the minimum bandwidth to support a given arrival process. With the aid of large deviation techniques, we can characterize the tail distribution of the queue length. When the effective bandwidth of the arrival processes and the effective capacity of the departure processes are equal, the queue is stable and the tail distribution of the stationary queue length decays exponentially. The decay rate is denoted by $\theta \geq 0$ such that

$$\lim_{Q_{\max}\to\infty} \frac{\log \Pr\{Q > Q_{\max}\}}{Q_{\max}} = -\theta \quad (1)$$

where Q is the stationary queue length. Note that for large queue length $Q_{\max}$, the queue overflow probability is approximated by $e^{-\theta Q_{\max}}$. Hence, while larger θ corresponds to stricter queueing constraints, smaller θ implies looser queueing constraints. It is worth mentioning that to ensure a stable queue, the average arrival rate of the queue should be less than the average service rate. We enforce that this condition is satisfied for all queues of the buffer-aided relay networks. Otherwise, there is overflow in queues, and the relay network cannot support any non-zero throughput without violating the delay constraints.

Then, the delay violation probability can be approximated by $e^{-J(\theta)D_{\max}}$ for large delay threshold $D_{\max}$, where $J(\theta)$ is the delay exponent depending on θ, the arrival processes and the departure processes jointly. Statistical delay constraints generally require the delay violation probability to be less than a certain target value $\varepsilon > 0$. Note that the higher $J(\theta)$ is, the smaller delay violation probability is, corresponding to the stricter delay constraints. Note also that while the delay constraint is defined for large delay, applications to a small delay regime such as voice transmissions with strict delay constraints still hold. On the other hand, it has been shown that for a given signal-to-noise ratio (SNR) value and channel conditions, larger $J(\theta)$ results in smaller achievable throughput [12]. Therefore, we would like to have as small $J(\theta)$ as possible to achieve the largest throughput while satisfying the delay constraints.

Two-Hop Delay Constraints

With the above characterizations, we now consider the delay constraints of two concatenated queues, which is the case for the relay networks under consideration. We assume that the buffers are saturated, and hence the source **S** and the relay **R** always attempt to transmit. We need to guarantee that the data transmission with the largest *end-to-end* delay should satisfy the statistical delay constraints, that is, consider data flow over multiple buffers. For the buffer-aided relay systems mentioned above, we need to characterize the *end-to-end* delay of data flows over two concatenated buffers. We can determine the statistical delay through two steps. First, consider the two queues with arbitrary statistical queueing constraints for queue 1(**S**) and queue 2(**R**). Under the assumption of stable queues, we can characterize the delay exponents by J_1 for the source queue and J_2 for the relay queue, respectively. Then, the end-to-end queueing delay violation probability for data going through both queues can be characterized by considering the cumulative distribution of the sum stationary delays. Through simple computations, we can derive the set of all potential delay exponents J_1 and J_2 satisfying the statistical delay constraint imposed as limitations on the maximum delay violation probability specified by ε. Since we would like to have J_1 and J_2 as small as possible, we can express the set of J_1 and J_2 that satisfy the statistical delay constraints in equality, and if deeming J_2 as a function of J_1, we can demonstrate that such J_2 is convex in J_1. This shows with the above characterization that we must have larger J_1 if J_2 is decreased, and vice versa [12].

Figure 2 illustrates the potential pairs of J_1 and J_2 satisfying the statistical delay constraints. We assume $\varepsilon = 0.05$ and $D_{\max} = 1$ sec. In the figure, we plot J_2 as a function of J_1 for those pairs satisfying the statistical delay constraints, and note that only (J_1, J_2) in the gray region can be acceptable to satisfy the statistical delay constraints. As can be seen from the figure, the curve given by the lower boundary yields the potential set of (J_1, J_2) for characterizing the throughput of the

buffer-aided relay systems under statistical delay constraints. For other pairs of (J_1, J_2), either the statistical delay constraints cannot be satisfied, or one of the queues is subject to a more stringent constraint than necessary, decreasing the achievable throughput.

Throughput and Relay Policy

We assume full-duplex decode-and-forward (DF) relaying. Now that the relay is equipped with buffer, it can decode the received information from the source and store the information in the buffer for later transmission. We assume that the transmission power levels at the source and relay are fixed and hence no power control is employed (i.e., nodes are subject to short-term power constraints). We also assume that the channel capacity for each link can be achieved, that is, the service processes are equal to the instantaneous Shannon capacities of the links, and hence no outage occurs. Note that to achieve the capacity of each link, the channel state information (CSI) of each link should be available at both the transmitter and receiver of the link. In addition, for the three-node relay channel, we will have a multi-access channel considering the data reception at the destination. Specifically, the CSI of the link **S-R** is available at **S** and **R**, while the CSI of the links **R-D** and **S-D** (if it exists) is available at all nodes. We need to assure that the average rate of the link **S-R** is smaller than the average rate of the link **R-D**. Otherwise, the queue at the relay cannot be stable.

In the presence of the link **S-D** or multiple relays, selection relaying, in which the source chooses the relay for data reception when its channel is beyond a certain threshold, can greatly affect the performance of a relay network. In buffer-aided relay systems, new relay policies have been proposed to improve the throughput and/or diversity such as weighted max link relay selection [4], where the relay with the strongest weighted receiving channel or transmitting channel is selected for data reception or transmission, respectively. Note that these relay protocols are designed for buffer-aided cooperative communications without delay constraints. Incorporating the statistical delay constraints, we expect that new relay protocols can further improve the system performance.

In the following, we only consider the full-duplex DF relaying. In the case of half-duplex DF relaying, for any fixed time allocation for the receiving and transmitting channels at the relay, we can also characterize the effective capacity of the system since delay exponents for the links can be expressed similarly. For adaptive link selection, the authors in [13] have characterized the effective capacity of a half-duplex two-hop channel with the same queueing constraints for the source and relay buffers, but the analysis for the general statistical delay constraints still needs further investigation.

Two-Hop Channels

If there is no direct link between the source and the destination, the source can only select the relay for data reception, and the received signals at the destination all come from the relay. Note that due to the buffer at the relay, the transmission rate from the source will not be limited by the minimum rate of the link **S-R** and the link **R-D**. Instead, each link can work at its capacity.

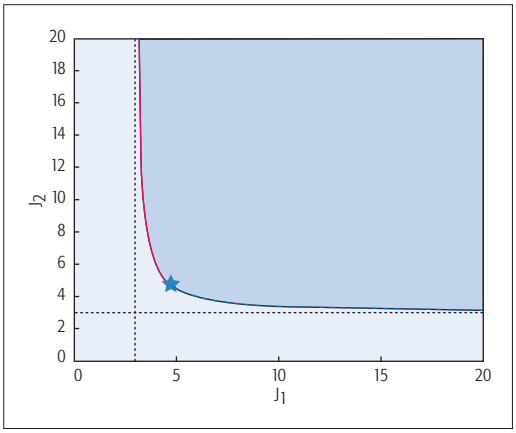

Figure 2. J_2 vs. J_1 satisfying the statistical delay constraints.

In order to satisfy the statistical delay constraints, we need to consider all possible delay exponent pairs J_1 and J_2 associated with the two buffers characterized earlier. Since we assume that each link operates at its capacity, we can derive the statistical queueing constraints reaching the statistical delay exponents for the different links. Given the statistical queueing constraints of the two-hop channel, we can express the effective capacity of the two-hop channel with the method provided in [14]. More specifically, the effective capacity is the maximum constant arrival rate such that the exponential decay rates of the queues are no less than the previously dictated values, respectively.

We employ the branch-and-bound method to characterize the maximum effective capacity. Specifically, we consider three sets of (J_1, J_2) specified by the lower boundary curve of the region (J_1, J_2) such that the statistical delay constraints can be satisfied, that is, $J_1 = J_2$, $J_1 < J_2$, and $J_1 > J_2$, and find the maximum effective capacity achieved in each set. Then, the maximum effective capacity of the system will be given by the maximum one of the obtained three values. We start from the point $J_1 = J_2$, for example, the star point in Fig. 2. For this case, we can obtain the associated effective capacity (buffer-aided symmetric). Then, depending on the obtained results with $J_1 = J_2$, we can choose to iterate over (J_1, J_2) in the set of $J_1 < J_2$ or $J_1 > J_2$, for example, the dashed or solid curve in Fig. 2. Now, we can derive the maximum effective capacity of the two-hop channel satisfying the statistical delay constraints (buffer-aided optimal (asymmetric)). As it turns out, in the presence of statistical delay constraints, considerable performance improvement can still be achieved by equipping the relays with buffer [12].

Three-Node Relay Channels

If there is a direct link between the source and the destination, the source can select the destination or the relay for data reception, and the received signal at the destination consists of the signals from the source and the relay. For this case, we assume that successive decoding of the received signal can be performed at the destination when the destination is selected for data reception.

If the relay is equipped with a buffer, it can store the received information and forward to the destination at a later time. In the absence of delay constraints, the relay selection policy based on max channel gain (MCG) or max received SNR can provide considerable performance gain.

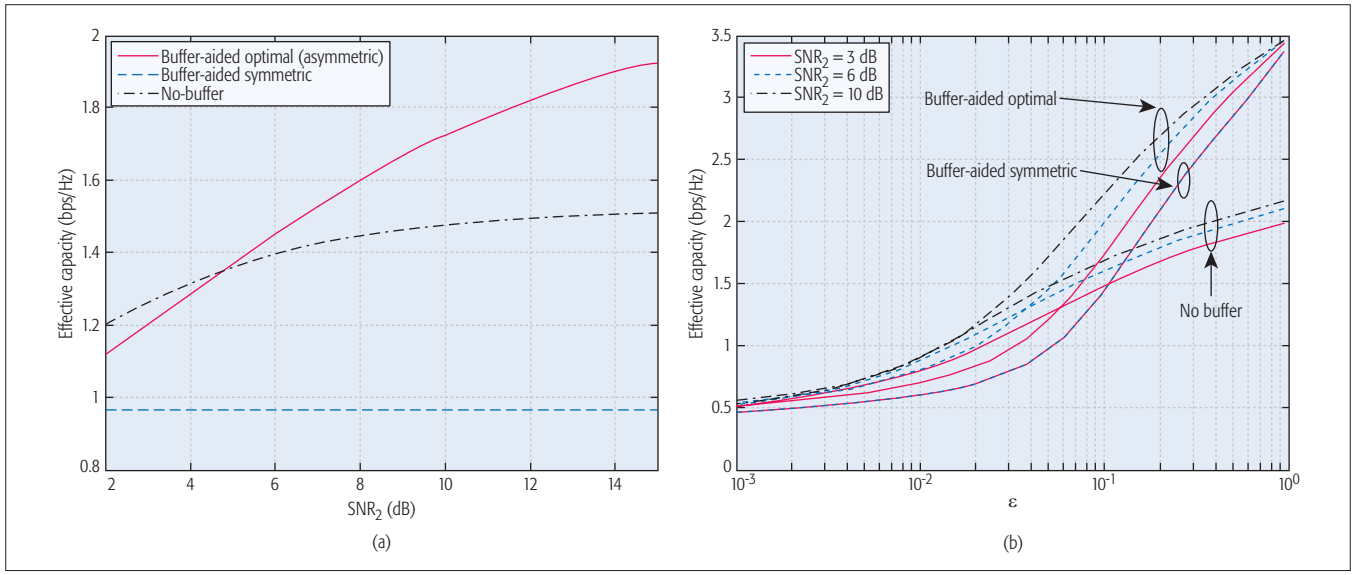

Figure 3. Effective capacity of a two-hop channel. The performances of the buffer-aided optimal, buffer-aided symmetric and no buffer DF relay strategies are plotted: a) effective capacity in SNR of the relay; b) effective capacity vs. ε.

Again, if the relay is equipped with a buffer, it can store the received information and forward to the destination at a later time. In the absence of delay constraints, the relay selection policy based on max channel gain (MCG) or max received SNR can provide considerable performance gain. This is because for each block, the service rate of the queue at the source can be maximized, and hence the average service rate of the queue can be maximized leading to larger average throughput.

Now, under the statistical delay constraints, we should again consider all possible delay exponent pairs J_1 and J_2. In order to derive the maximum effective capacity of the relay system, we need to characterize the effective bandwidth of the departure processes of the queue at the source for any arbitrary relay selection policy. By viewing the relay system as an in-tree network [8], we can treat the buffer at the source as the *predecessor* of the buffer at the relay, or equivalently, the buffer at the relay is the *successor* of the buffer at the source. Taking advantage of the effective bandwidth theory, we can express the effective bandwidth of the departure processes of the queue at the source for any arbitrary relay selection policy accordingly. This characterization enables us to obtain the effective capacity of the relay systems following steps similar to the previous discussion about the two-hop channels.

For relay systems under delay constraints, we can also modify the relay protocol proposed in the absence of delay constraints. In particular, to take into account the delay constraints for relay selection, we can select the relay for data reception when the instantaneous rate of the link **S-R** is greater than the effective capacity of the link **S-D**, that is, max delay exponent (MDE), instead of the instantaneous rate of the link **S-D**. The maximization problem can be expressed as

$$\max_{g} J_1\left(\theta_1, g\left(z_{sd}\right)\right). \tag{2}$$

where θ_1 is the statistical queueing constraints at the source, and $g(z_{sd})$ denotes the relay selection policy such that when the channel of the source-relay links is stronger than the threshold function g, the relay is selected for data reception. In this protocol, the delay exponent of the buffer at the relay J_1 is maximized assuming that the statistical queueing exponent θ_1 is given [15]. Now, given the delay exponent, the statistical queueing exponent can be minimized, which in turn yields the largest effective capacity. Note that the proposed delay constraint based relay selection policy can further increase the achievable throughput, as will be validated by the numerical results in the following.

Evaluation Results

In Figs. 3a and 3b, we evaluate the system performance of the two-hop channels, where we assume that the **S-R** and **R-D** links experience Rayleigh fading and have the same channel conditions. We assume the SNR at the source is 0 dB, $\mathbb{E}\{z_{sr}\} = \mathbb{E}\{z_{rd}\} = 16$. Figure 3a shows the effective capacity as a function of SNR of the relay node, SNR2, for the two-hop channel. The effective capacity of a two-hop system increases with the SNR of the relay, and can achieve significant performance improvement for large SNR of the relay. And, in all cases, the achievable rate of asymmetric delay constraints is greater than the one achieved with symmetric delay constraints at the two buffers. Figure 3b plots the effective capacity as ε varies. Buffer-aided relay can be helpful even in the presence of end-to-end delay constraints for certain cases. Also, we can find that for larger SNR of the relay, SNR2, the buffer at the relay can help improve the achievable rate at a smaller ε, that is, under more stringent delay constraints.

Figures 4a and 4b depict the effective capacities of the three-node relay channel for the different relay policies as a function of SNR of the relay node, SNR_r, and ε, respectively. We assume that all links experience Rayleigh fading, and the links **S-R** and **R-D** have the same channel conditions while the link **S-D** has the worst channel conditions. We assume that the SNR at the source is 0 dB, $\mathbb{E}\{z_{sr}\} = \mathbb{E}\{z_{rd}\} = 16$, and $\mathbb{E}\{z_{sd}\} = 1$. Figure 4a shows that buffering relay can still help improve

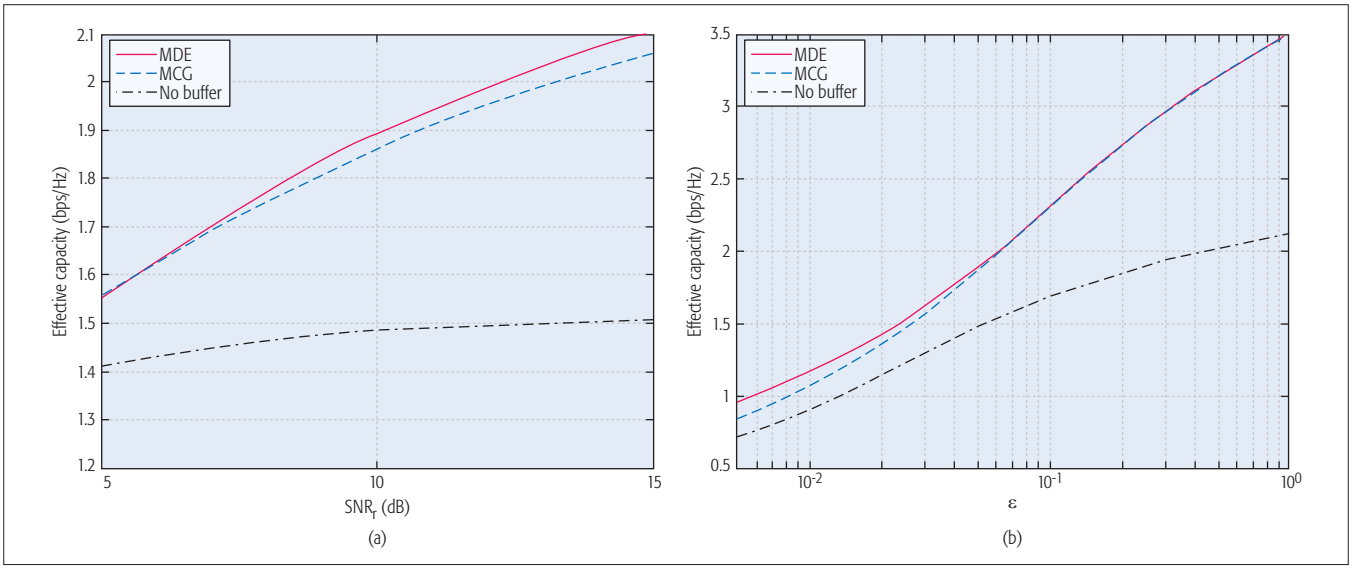

Figure 4. Effective capacity of a three-node relay channel. The performances of MDE, MCG, and no buffer with DF relaying policies are presented: a) effective capacity as a function of the SNR of the relayl b) effective capacity as a function of α.

the system throughput under statistical delay constraints. In addition, the proposed relay policy (MDE) can achieve better performance than the one without delay constraints (MCG). Figure 4b shows that the performance improvement of the proposed relay policy can be enlarged at smaller ε, that is, under more stringent delay constraints. We can also find that buffering relay still helps improve throughput for a wide range of statistical delay constraints in the presence of source-destination link.

Challenges and Future Research

With the introduction of effective bandwidth and effective capacity theory, the analysis of buffer-aided relay systems can be simplified, and several closed-form results can be obtained. We now discuss several challenges to be addressed for practical implementation and some interesting directions for future research extensions.

Multiple Relay Systems: In the future fifth-generation systems where femto- and pico-cells are deployed, multiple nodes will cooperate together, in which case extensions to the multiple relay systems will be mature and can give us valuable insights.

Power and Rate Allocation: In the absence of delay constraints, the power and rate adaptation policies have been developed for different relay networks based on channel capacity. Generally, the throughput optimal policies can be derived by solving the optimization problems built upon the explicit expression of throughput and power constraints. However, under the statistical delay constraints, the throughput of the relay systems cannot be expressed explicitly, in which case the optimization problems become difficult if not intractable to formulate. Consequently, the design of throughput or energy efficiency optimal resource allocation policies is highly challenging.

Finite Considerations: In a practical system, the buffer size, codeword length and constellation size can be finite. However, due to the implicit assumption of infinite queue length, the effective capacity model cannot be adopted directly to address the scenarios of finite buffer. Hence, new approaches or variations of the effective capacity theory should be proposed. In addition, while channel capacity is assumed for different links, the potential impact and throughput loss due to the finite codeword length or/and finite constellation size has not been investigated yet.

Caching: The delivery phase of caching is very similar to the buffer-aided relay systems, although the buffer at the destination is also incorporated for analysis. Joint design of caching with the buffer-aided relaying will yield interesting insights for the caching systems with delay constraints for the delivery phase.

Queue-Aware Policies: It has been reported that scheduling policies incorporating the queue length information can further improve the effective capacity. In such cases, scheduling and relay policies of buffer-aided relay systems should be designed to take into account the queue length. However, analysis based on the effective capacity will become extremely complicated due to the nature of derivation of the effective bandwidth of the service processes.

Information-Theoretic Security: In addition to the delay constraints, security is also an important issue in wireless communication systems. For the current studies of security problems of relay systems, secure communications are generally limited not only by the eavesdropper channel but also the worst channel condition of the link that is not overheard. Due to the buffer-aided relay, separations of the different links can potentially improve the secrecy rate achieved by the relay systems.

Multicarrier Systems: In the fourth and fifth generations of cellular systems, multicarrier modulations such as orthogonal frequency division multiplexing (OFDM) are fundamental components. In such scenarios, relay policies and resource allocation in buffer-aided relay systems will be interesting.

Conclusion

This article has discussed buffer-aided relay systems, and introduced a method to characterize the maximum constant arrival rate of such systems under delay constraints, namely the effec-

Although the effective capacity approach discussed can provide preliminary results on buffer-aided relay systems, shedding light on the benefits of buffer-aided relay, there are still several challenges for future research, such as efficient resource allocation policies and performance analysis in the finite buffer/delay regimes.

tive capacity. We have shown that buffer-aided relay can still provide considerable performance gains compared to conventional relay systems under statistical delay constraints. Although the effective capacity approach discussed can provide preliminary results on buffer-aided relay systems, shedding light on the benefits of buffer-aided relay, there are still several challenges for future research, such as efficient resource allocation policies and performance analysis in the finite buffer/delay regimes. Overall, buffer-aided relaying under delay constraints is an important research area in cooperative communications, and many interesting problems are still open.

Acknowledgments

This work has been supported in part by the National Natural Science Foundation of China (Nos. 61671205 and 61672238), the Shanghai Sailing Program (No. 16YF1402600), and the National Science Foundation (NSF) (grant CCF-1618615).

References

[1] Q. Li *et al.*, "Cooperative Communications for Wireless Networks: Techniques and Applications in LTE-Advanced Systems," *IEEE Commun. Mag.*, vol. 50, no. 4, Apr. 2012, pp. 22–29.
[2] N. Zlatanov *et al.*, "Buffer-Aided Cooperative Communications: Opportunities and Challenges," *IEEE Commun. Mag.*, vol. 52, no. 4, Apr. 2014, pp. 146–53.
[3] T. Charalambous *et al.*, "Modeling Buffer-Aided Relay Selection in Networks with Direct Transmission Capability," *IEEE Commun. Lett.*, vol. 19, no. 4, Apr. 2015, pp. 649–53.
[4] N. Zlatanov, V. Jamali, and R. Schober, "Achievable Rates for the Fading Half-Duplex Single Relay Selection Network Using Buffer-Aided Relaying," *IEEE Trans. Wireless Commun.*, vol. 14, no. 8, Aug. 2015, pp. 4494–4507.
[5] V. Jamali *et al.*, "Achievable Rate of the Half-Duplex Multi-Hop Buffer-Aided Relay Channel with Block Fading," *IEEE Trans. Wireless Commun.*, vol. 14, no. 11, Nov. 2015, pp. 6240–56.
[6] V. Jamali, N. Zlatanov, and R. Schober, "Bidirectional Buffer-Aided Relay Networks with Fixed Rate Transmission-Part II: Delay Constrained Case," *IEEE Trans. Wireless Commun.*, vol. 14, no. 3, Mar. 2015, pp. 1339–55.
[7] X. Zhang, W. Cheng, and H. Zhang, "Heterogeneous Statistical QoS Provisioning over 5G Mobile Wireless Networks," *IEEE Wireless Netw.*, vol. 28, no. 6, Nov./Dec. 2014, pp. 46–53.
[8] C.-S. Chang, *Performance Guarantees in Communication Networks*, New York: Springer, 1995.
[9] D. Wu and R. Negi "Effective Capacity: A Wireless Link Model for Support of Quality of Service," *IEEE Trans. Wireless Commun.*, vol. 2, no. 4, July 2003, pp. 630–43.
[10] M. A. Maddah-Ali and U. Niesen, "Fundamental Limits of Caching," *IEEE Trans. Inform. Theory*, vol. 60, no. 5, May 2014, pp. 2856–67.
[11] Z. Zhang *et al.*, "Full Duplex Techniques for 5G Networks: Self-Interference Cancellation, Protocol Design, and Relay Selection," *IEEE Commun. Mag.*, vol. 53, no. 5, May 2015, pp. 128–37.
[12] D. Qiao and M. C. Gursoy, "Statistical Delay Trade-Offs in Buffer-Aided Two-Hop Wireless Communication Systems," *IEEE Trans. Commun.*, vol. 64, no. 11, 2016, pp. 4563–77.
[13] K. T. Phan, T. Le-Ngoc, and L. B. Le, "Optimal Resource Allocation for Buffer-Aided Relaying with Statistical QoS Constraint," *IEEE Trans. Commun.*, vol. 64, no. 3, Mar. 2016, pp. 959–72.
[14] D. Qiao, M. C. Gursoy, and S. Velipasalar, "Effective Capacity of Two-Hop Wireless Communication Systems," *IEEE Trans. Inform. Theroy*, vol. 59, no. 2, Feb. 2013, pp. 873–85.
[15] D. Qiao, "Effective Capacity of Full-Duplex Buffer-Aided Relay Systems with Selection Relaying," *IEEE Trans. Commun.*, vol. 64, no. 1, Jan. 2016, pp. 117–29.

Biographies

Deli Qiao received his Ph.D. degree in electrical engineering from the University of Nebraska-Lincoln, Lincoln, NE, in 2012. Currently, he is on the faculty at the East China Normal University. From 2012 to 2014, he worked at Huawei Technologies Inc., LTD as a research engineer, focusing on 5G techniques, especially massive MIMO. His research interests are in the general areas of wireless communications, information theory, and signal processing.

M. Cenk Gursoy received his Ph.D. degree in electrical engineering from Princeton University in 2004. He is currently an associate professor in the EECS Department at Syracuse University. His research interests are in the general areas of wireless networks and information theory. He has been serving as an editor of *IEEE Transactions on Green Communications and Networking*, *IEEE Transactions on Communications*, and *IEEE Transactions on Vehicular Technology*. He received an NSF CAREER Award in 2006.

ACCEPTED FROM OPEN CALL

Defense Mechanisms Against DDoS Attacks in SDN Environment

Kübra Kalkan, Gürkan Gür, and Fatih Alagöz

ABSTRACT

SDN is a pivotal technology that relies on the fundamental idea of decoupling control and data planes in the network. This property provides several advantages such as flexibility, simplification, and lower costs. However, it also brings several drawbacks that are largely induced by the centralized control paradigm. Security is one of the most significant challenges related to centralization. In that regard, DDoS attacks are particularly pertinent to the SDN environment. This article presents a concise survey on solutions against DDoS attacks in software-defined networks. Moreover, several mechanisms are analyzed, and a comparative classification is provided for rendering the current state of the art in the literature. This analysis will help researchers to address weaknesses of these solutions and thus mitigate such attacks using more effective defense mechanisms.

INTRODUCTION

The number of network devices connected to the Internet is increasing at an accelerated pace. Not only the proliferation of mobile devices, but also emerging technologies such as the Internet of Things (IoT) will multiply the number of network devices with advanced anytime-anywhere services. Accordingly, growing network sizes will result in more complex networks and thus more challenging requirements. However, existing network technologies and infrastructure are not designed to enable such convoluted systems.

In order to design future networks that can satisfy these burgeoning demands and requirements, several approaches have already been proposed, and software defined networking (SDN) is considered as a vital solution, among others. The salient characteristic of SDN is the decoupling of data and control planes in network devices. In traditional networks, routers apply all high-level routing algorithms and decide where data packets should be forwarded. In SDN, decision and forwarding functions are separated. The hard decision process is provided by the SDN controller, whereas data forwarding is handled by switches. Since decision algorithms do not run on network devices, simpler hardware can be utilized rather than more sophisticated routers. Therefore, simplification of network devices and central management are the most attractive SDN properties for network operators. Since there will be an enormous number of network devices, they are expected to enjoy substantial cost savings due to these characteristics. Besides, each router has its own security, link failure, and forwarding mechanisms in traditional networks. If any of these mechanisms needs to be updated, each network device should handle these tasks individually. However, all these issues/problems can be managed centrally in SDN.

Despite the fact that SDN has obvious advantages such as simplification and network flexibility, there are some important challenges that are worth consideration such as reliability, scalability, latency, and controller placement. However, there is a limited industry and research community effort on security issues of SDN, although it comes with new vulnerabilities and attack vectors for malicious agents. SDN's vulnerabilities stem from its two inherent properties: the ability to control the network by means of software and the centralization of network intelligence in the controller. These features result in several trust issues and single-point management failures. In order to provide trust, authorization policies and authentication mechanisms should be applied. Single-point management failure can be attained by compromising the availability of the SDN controller. One of the most common ways that can be utilized for such malicious impact is distributed denial of service (DDoS) attack.

The simple strategy behind a denial of service (DoS) attack is to deny the use of system resources to legitimate users and degrade the system availability. The fundamental mechanism for DoS attack execution is to send a flood of superfluous network traffic to the target so that it cannot respond to genuine requests for services or information. If multiple sources are used by the attacker(s), it is called DDoS attack, which is much more catastrophic than DoS. For instance, a giant botnet (Mirai IoT botnet) made up of hijacked IoT devices such as cameras, lightbulbs, and thermostats launched a 600+ GB/s DDoS attack against a security blogger in September 2016 such that the global service provider had to cancel the hosting account since it was extremely costly to defend that website. Considering its design pillars and characteristics, DDoS attacks are also intrinsically viable for SDN.

One of the drawbacks of SDN architecture regarding DDoS attacks is the limited passive capabilities of switches. Since in general they send all packets with unknown flows to the controller, their medium becomes attractive for DDoS

The authors present a concise survey on solutions against DDoS attacks in software-defined networks. Moreover, several mechanisms are analyzed, and a comparative classification is provided for rendering the current state of the art in the literature. This analysis will help researchers to address weaknesses of these solutions and thus mitigate such attacks using more effective defense mechanisms.

This work was supported in part by the Scientific and Technical Research Council of Turkey (TUBITAK) under grant number 117E165 and in part by the Turkish State Planning Organization (DPT) under the TAM Project, number 2007K120610.

The authors are with Bogazici University

Digital Object Identifier: 10.1109/MCOM.2017.1600970

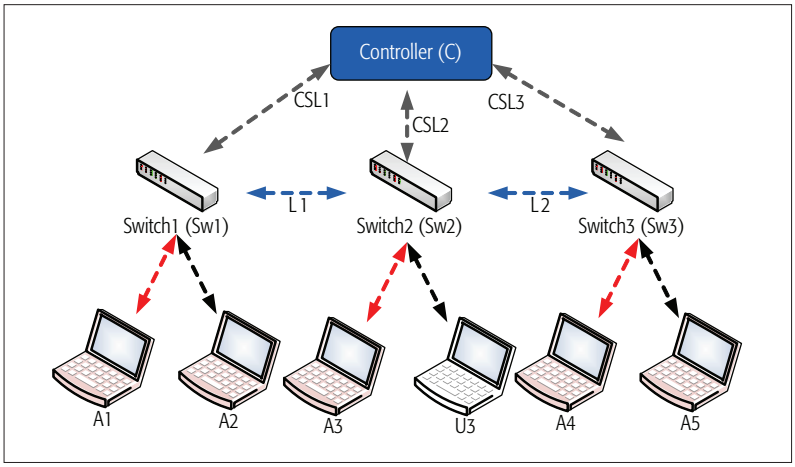

Figure 1. A sample SDN topology with DDoS attackers.

attacks. Moreover, they do not possess sufficient resources for very large volumes of traffic. On the controller side, because of the central management property, DDoS attack can cause catastrophic results if the controller is saturated with attack traffic.

In the literature, early connections between SDN and DDoS were constructed for providing defense mechanisms against DDoS attacks by utilizing SDN [1]. However, it was realized that SDN has its own vulnerabilities that can attract DDoS attacks, and it also needs security for itself. For this reason, DDoS and SDN related works can be categorized as *SDN for security* and *security for SDN*. The latter category especially focuses on providing security against DDoS attacks in SDN and is a more recent research topic. It is still an immature area since there is no outstanding solution, and proposed solutions exhibit various drawbacks. To the best of our knowledge, this is the first survey that focuses only on security for SDN in terms of DDoS attacks and classifies these works according to several aspects. Our main goal is to construct an illustrative categorization of existing defense mechanisms in SDN and make it easier for network experts/practitioners to adopt appropriate mechanisms for various contexts and circumstances. Also, we establish a baseline view for researchers to propose a new defense mechanism in that setting.

DDoS Attack Scenarios in SDN

Conventional DDoS attacks such as UDP, ICMP, TCP SYN flooding, NTP amplification, and ping of death are also viable in SDN. Since SDN infrastructure suggests a centralized management for network flows, SDN is very attractive for DDoS attackers. The baseline network flow policy of SDN suggests that when a packet comes from an unknown IP to a switch, it is forwarded to the controller. Then the controller sends a flow rule to the switch for this IP. If attackers send a large number of packets from several IPs, each packet will be forwarded to the controller. Then a huge number of attack packets will make the system unavailable for legal users. A sample topology for SDN is illustrated in Fig. 1. Some DDoS attack scenarios in SDN can be scripted as follows.

Scenario 1: The controller can be the target for the attack. Attacker(s) can generate traffic with spoofed IP addresses. In this attack, attacker(s) are under switch Sw1, which needs to send all packets to the controller as they are coming from unknown IP(s). Then the link between Sw1 and the controller called CSL1 is potentially congested. (Target: CSL1, Attacker: A1)

Scenario 2: The system resources of the controller can be the target for attackers. This time, attackers are under different switches managed by the same controller. Since attack traffic is coming from several switches, the attack load is divided, and it is intrinsically difficult to detect. This type of attack is named blind DDoS attack in the literature [2]. (Target: C, Attacker: A1, A2, and A3)

Scenario 3: Switch memory can be the target for attackers. A switch has limited memory, while it needs to store a new entry for each unmatched flow. Although this requirement can be relaxed with sophisticated flow table designs, the switch will suffer from inflated table size in general. When an attacker generates new flows, the controller sends new entries to the switch. If the attacker uses all available table entries, the legal traffic from new IPs cannot be served. Besides, the target switch can also be unreachable in a more sophisticated way. It can be blocked by intercepting the links to this switch. As an attack example in our sample topology, A1 and A2 try to block L1, whereas A4 and A5 use resources of L2. Then Sw2 will be unreachable. Each traffic flow does not have large bit rate; thus, it is not easy to detect. However, as a result it makes the target unavailable. A concrete example, Crossfire attack, is explained in [3]. If we suppose that the target area is the network under Sw2, the target links are L1 and L2 for our sample topology. (Target: Sw2, Attacker: A1, A2, A4, and A5)

Scenario 4: A link between switches can also be the target. This attack can be facilitated by communication between attackers under different switches. For instance, A1 and A3 can use all available communication resources between Sw1 and Sw2. This type is called Coremelt attack in the literature [4]. (Target: L1, Attacker: A1 and A3)

Scenario 5: A legal user under a switch can be the victim of an attacker (e.g., a server in a cloud-computing environment). The attacker can reside on the same switch or another switch. If the controller or the switch cannot detect it, the server's resources will be depleted.(Target: U1,Attacker: A1 or A2 or A3 or A4 or A5)

Compared to traditional networks, SDNs are prone to such DDoS attacks due to three inherent dynamics of attacks, as seen in these scenarios.

Propagation of Attacks: Simpler processing pipelines in switches compared to conventional routers render SDN more prone to attack propagation. The controller has to step up to hinder such dynamics. However, it is nontrivial to integrate and enable relevant capabilities in SDN controllers.

Aggregation of Attacks: Attack traffic in terms of payload and control traffic aggregates toward the SDN controller. This amplifier effect boosts the attack impact and thus severity.

Widespread Impact of Attacks: A DDoS attack can rapidly affect the entire network by crippling controller(s). This phenomenon is magnified via the two characteristics described above.

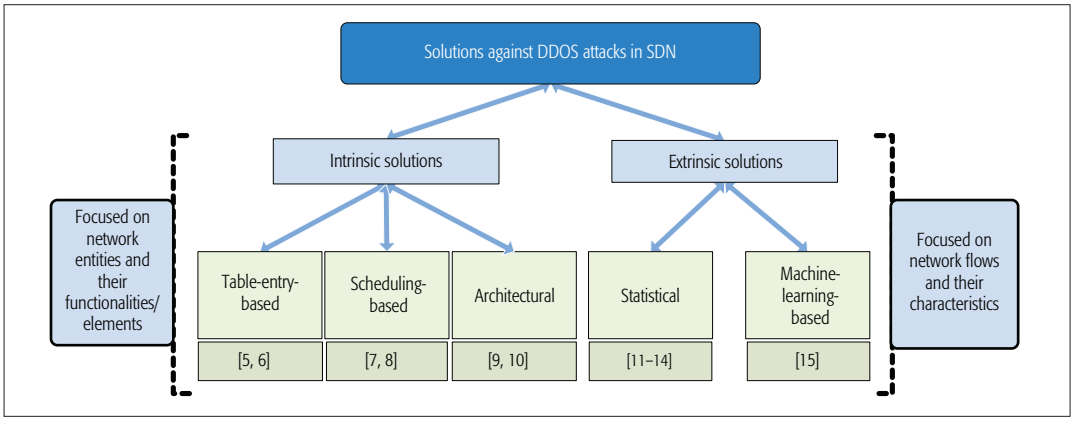

Figure 2. Classification of solutions against DDoS attacks in SDN.

In our case, some solutions are related to structural attributes of SDN environment, whereas others are mostly related to the properties of network flows. For this reason, we propose to classify identified mechanisms as intrinsic vs. extrinsic solutions.

While these aspects make SDN DDoS attacks more challenging, they also provide opportunities for potential defense approaches, which are listed and described in the next section.

Solutions against DDoS Attacks in the SDN Environment

In order to cope with various DDoS scenarios in SDN environments, several solutions are proposed in the literature. In fact, it is a novel research topic in which almost all mechanisms have been formulated in the last few years. In this section, these solutions are analyzed to examine their properties. Since all models have their own pros and cons, it is not possible to state that one of these mechanisms is a superior solution. For this reason, security practitioners need to choose the appropriate one(s) according to their requirements. In order to provide a clear way to analyze and decide, classifications of these methods in terms of several aspects are also provided in this section. For that purpose, we elaborate on two dichotomies: one focusing on which elements they rely on (network elements vs. flows) and another focusing on their defense functionalities.

Solutions in the literature can be classified according to whether they are intrinsic or extrinsic. A property that is inherited and essential is named intrinsic, whereas a property that varies depending on exterior factors is called extrinsic. In our case, some solutions are related to structural attributes of the SDN environment, whereas others are mostly related to the properties of network flows. For this reason, we propose to classify identified mechanisms as intrinsic vs. extrinsic solutions. This classification is illustrated in Fig. 2.

Intrinsic Solutions

Intrinsic solutions can be further categorized as table-entry-based, scheduling-based, and architectural. Table-entry-based models propose solutions related to the limited table size of switches. Each unknown flow needs a new entry in switch memory. This becomes a bottleneck during a DDoS attack, which contains packets with different IP addresses. For instance, [5, 6] suggest some solutions for this problem. In [5], the impact of a DDoS attack in SDN is presented. Their results highlight the importance of managing the flow tables. They conclude that table entry replacement policies should use multiple parameters such as number of packets, generation date, and utilization properties of a flow entry, rather than using just one parameter such as earliest expiration time. Besides, a controller should have an intermediate buffer module, which stores the flow entries temporarily and manages the replacement of flow entries. These suggestions can also be utilized as DDoS mitigation methods. Similarly, Katta *et al.* [6] presents a solution for a general attack scenario related to the memory of switches. Switches can allow a limited number of entries in their tables due to resource constraints on memory capacity. Proper update policy of these entries is essential against DDoS attacks since attack packets are also dropped or forwarded according to these entries. Their work proposes a rule update mechanism for switch tables. Although their idea is not specifically proposed targeting DDoS attacks, it is beneficial for DDoS mitigation.

Scheduling-based solutions are implemented on the controller. These models suggest that it is essential to protect the controller since it is the core of the system in SDN. In order to provide this capability, such models deal with scheduling assignment of tasks from switches. The approach in [7] provides scalability, whereas [8] provides isolation of the attack traffic with the help of scheduling algorithms. In [7], Hsu *et al.* proposed a hash-based mechanism that operates in the controller to increase scalability of the network. Their work performs hash-based round-robin scheduling for assigning the incoming packets from a crowded switch to several queues in the controller. In this model, the controller can still serve the switch even if it has a high amount of traffic due to flash crowd or DDoS attack.

This model does not have a detection mechanism and acts in the same way for flash crowd. This situation results in unproductive service for DDoS attack packets in the controller. In [8], Lim *et al.* suggest that the most essential aim of a defense mechanism is to provide the controller's work continuity in case of an attack since the controller's failure leads to the entire SDN being unavailable. They leverage a scheduling-based scheme that contains most of the attack traffic at attack ingress switches so that the SDN as a whole can continue normal operation. If one switch is infected by DDoS, normally the controller cannot serve other users. In order to prevent this problem, they create different queues for each switch. Actually, this model realizes the opposite of the

Facilitating switches with some discreet intelligence features does not compromise the main paradigm of SDN. However, a system designer should be cautious for this trade-off of the "pure SDN" paradigm vs. "capable" switches. Over-empowering switches can harm the fundamental SDN concept of simple and lean data-forwarding network nodes.

idea that was proposed in [7]. The goal of [7] is to provide scalability for the controller, whereas [8] is aimed at the continuity of the controller's operation. The conspicuous drawback of this latter work is that it considers flash crowd and DDoS attacks in the same manner. It does not employ an actual detection mechanism. Moreover, since it does not have packet-based analysis, it does not differentiate among legal and attack packets coming from the same switch.

Architectural solutions are related to the hierarchies and roles of network elements for facilitating solutions against DDoS attacks. References [9, 10] suggest decoupling monitoring and controlling properties of the controller. They also propose a master that manages these controller functionalities. The proposed model in [9] suggests that the controller's functions of monitoring and controlling should be decoupled. Communicating with switches for control messages and simultaneously monitoring traffic in the network increases the overhead on the controller. Some attack types can be resolved without having access to all packets in the network. These are called "low resolution attacks" such as DDoS attacks, whereas others cannot be detected without accessing all packets such as Address Resolution Protocol (ARP) caching and poisoning. In their architecture, there is an orchestrator that decides which modules should be activated according to suspected attack types as low or high resolution. Then the network monitor and controller modules participate in detection and mitigation according to the orchestrator commands. In this model, after an attack is detected, it just rate-limits the traffic rather than completely eliminating attack traffic. It is not an apt method for mitigation since the system is still occupied by attack packets.

Another model that suggests splitting the controller's duties is presented in [10]. Similar to [9], they decouple application monitoring and packet monitoring in the controller. They also propose that the controller should be distributed for security and load balancing reasons. In the case of a main controller failure, the secondary controller takes charge as the main one. They also suggest another hierarchical model in which the controller is separated into two layers consisting of a delegator and lower-level controllers that perform specific tasks assigned by the delegator. However, they do not describe a detection module, and their model considers flash crowd and DDoS attack cases identically which is not desirable.

Extrinsic Solutions

In contrast to intrinsic solutions, extrinsic solutions do not deal with network entities and modules of SDN. Instead, their solutions are related to the properties of the flows. Some works focus on identifying the malicious flows to defend against DDoS attacks. These solutions can be grouped into two categories: statistical solutions and machine-learning-based solutions.

References [11–14] employ statistical models, which generate baseline profiles with some statistical characteristics that are collected during an attack-free period. Subsequently, attack packets are eliminated by comparing them to these profiles. In the FlowFence mechanism [11], network switches monitor traffic and detect congestion by monitoring the bandwidth usage. When congestion occurs, the switch notifies the controller. Then, the controller requests statistics from every switch that sends flows to the congested link. It determines badly behaving flows that consume more bandwidth, and sends commands to switches to rate-limit them. Their results suggest that it is a fast and simple solution which prevents starvation. However, this mechanism is specifically a rate-limit mechanism that does not halt an attack completely.

Avant Guard [12] is a framework to improve the security and resiliency against DDoS and scanning attacks with greater involvement of switches. It introduces two modules on switches: connection migration (CM) and actuating triggers (ATs). CM proxies and classifies TCP SYN requests. If they are regarded as legitimate, they are authorized and migrated to the real target. This mechanism provides protection against SYN Flood attacks. For instance, if an enormous volume of requests is realized, an AT module triggers an event in the controller to insert a flow rule into the flow table automatically so that response time is reduced. The most conspicuous side-effect of this mechanism is the performance penalty. Since it utilizes CM, each flow needs to be classified. Besides, this module can only defend against a single type of DDoS attack.

There are several works that utilize entropy in traditional networks, but there are few such works in SDN. One of them is [13], which proposes an entropy-based lightweight DDoS flooding attack detection model running in the OpenFlow edge switch. In this mechanism, entropy is calculated for a destination IP address. If entropy decreases under a threshold, DDoS attack is detected. It enables determining the victim, but it is not possible to dissociate the legal packets from the attack ones. This scheme achieves a distributed anomaly detection in SDN and reduces the flow collection overload on the controller at the expense of more complicated switches.

In [14], it is explained that the communication between SDN controller and switches can also be the decisive system element regarding attack impact. Scotch [14] is a mitigation method that scales up the SDN control plane using vSwitch-based overlay. Their experimental results suggest that the bottleneck is the communication from the switch to the controller in the case of DDoS attacks or flash crowds. Thus, they suggest that when a physical switch is overloaded, new flows will be tunneled to multiple vSwitches. The most important drawback of this model is that it does not have an actual detection mechanism. It acts identically for flash crowds and DDoS attacks. Besides, it drops packets when they exceed a threshold. This means that legitimate packets cannot be preserved in the case of flash crowds and DDoS attacks.

Another subcategory of extrinsic solutions is machine-learning-based models. Similar to the statistical mechanisms, the machine-learning-based solution in [15] trains its defense mechanism with attack-free packets, then classifies attack packets with the help of machine learning algorithms. In [15], Kokila *et al.* propose a detection mechanism using a support vector machine (SVM) classifier. Their model deals with the attack scenario whose target is the controller. They also tried several machine learning techniques and concluded that SVM outperforms others. However, it takes more

Solution property	Capable switches	Dumb switches
Detection	[13]	[15]
Mitigation	–	[5–8, 10, 14]
Detection and mitigation	[11, 12]	[9]

Table 1. Classification of solutions according to defense functionality and switch intelligence.

time compared to other techniques. Besides, this model only deals with detection and does not give any mitigation solutions.

Defense Functionalities and Switch Intelligence

Another classification focusing on functionality and switch characteristics is illustrated in Table 1. Some of these solutions are dealing with detection or mitigation, whereas others focus on both mitigation and detection. Network administrators can determine their needs and then choose from these models in the corresponding category. As a general guideline, if the network infrastructure has sufficient resources, it is recommended to employ both detection and mitigation. Moreover, security practitioners can integrate disparate detection and mitigation models if their requirements cannot be met with an existing model. Table 1 also shows switch intelligence level in these proposed solutions. Some of the works in the literature suggest to bring some intelligence to SDN switches. The motivation for this approach is based on the idea of keeping flows in the data plane as much as possible with higher processing locality. When more flows are forwarded to the controller, it is more prone to attacks by malicious users. If switches are more capable of flow decisions, it is safer for the controller. Facilitating switches with some discreet intelligence features does not compromise the main paradigm of SDN. However, a system designer should be cautious for this trade-off of "pure SDN" paradigm vs. "capable" switches. Over-empowering switches can harm the fundamental SDN concept of simple and lean data forwarding network nodes.

Conclusion

In this article, we present a thorough treatment on solutions against DDoS in the SDN environment. First, we explain some DDoS scenarios that can be viable in SDNs. Then we give overall descriptions of several DDoS defense techniques and their comparative classifications. According to this analysis, it is realized that there is limited research work in machine-learning-based solutions, which provides a research direction for SDN security research. In order to maintain flows as much as possible in the data plane, some models suggest integrating some intelligence capabilities to switches. This establishes another research direction, which dwells on the trade-off between conforming to the fundamental SDN paradigm and augmenting switches with more functionalities.

References

[1] Q. Yan *et al.*, "Software-Defined Networking (SDN) and Distributed Denial of Service (DDoS) Attacks in Cloud Computing Environments: A Survey, Some Research Issues, and Challenges," *IEEE Commun. Surveys & Tutorials*, vol. 18, no. 1, 2016, pp. 602–22.

[2] D. Ma, Z. Xu, and D. Lin, "Defending Blind DDoS Attack on SDN Based on Moving Target Defense," *Proc. Int'l. Conf. Security and Privacy in Commun. Systems*, Springer, 2014, pp. 463–80.

[3] M. S. Kang, S. B. Lee, and V. D. Gligor, "The Crossfire Attack," *Proc. 2013 IEEE Symp. Security and Privacy*, 2013, pp. 127–41.

[4] A. Studer and A. Perrig, "The Coremelt Attack," *Proc. Euro. Symp. Research in Computer Security*, Springer, 2009, pp. 37–52.

[5] N. Tri, T. Hiep, and K. Kim, "Assessing the Impact of Resource Attack in Software Defined Network," *Proc. 2015 Int'l. Conf. Info. Networking*, 2015, pp. 420–25.

[6] N. P. Katta, J. Rexford, and D. Walker, "Incremental Consistent Updates," *Proc. 2nd ACM SIGCOMM Wksp. Hot Topics in Software Defined Networking*, 2013, pp. 49–54.

[7] S.-W. Hsu *et al.*, "Design a Hash-Based Control Mechanism in vSwitch for Software-Defined Networking Environment," *Proc. 2015 IEEE Int'l Conf. Cluster Computing*, 2015, pp. 498–99.

[8] S. Lim *et al.*, "Controller Scheduling for Continued SDN Operation under DDoS Attacks," *Electronics Letters*, vol. 51, no. 16, 2015, pp. 1259–61.

[9] A. Zaalouk *et al.*, "OrchSec: An Orchestrator-Based Architecture for Enhancing Network-Security Using Network Monitoring and SDN Control Functions," *Proc. IEEE Network Operations and Mgmt. Symp.*, 2014, pp. 1–9.

[10] D. Chourishi *et al.*, "Role-Based Multiple Controllers for Load Balancing And Security in SDN," *Proc. 2015 IEEE Canada Int'l. Humanitarian Technology Conf.*, 2015, pp. 1–4.

[11] A. F. M. Piedrahita *et al.*, "Flowfence: A Denial of Service Defense System for Software Defined Networking," *Proc. Global Info. Infrastructure Networking Symp.*, 2015, Oct. 2015, pp. 1–6.

[12] S. Shin *et al.*, "Avant-Guard: Scalable and Vigilant Switch Flow Management in Software-Defined Networks," *Proc. 2013 ACM SIGSAC Conf. Computer & Commun. Security*, ACM, 2013, pp. 413–24.

[13] R. Wang, Z. Jia, and L. Ju, "An Entropy-Based Distributed DDoS Detection Mechanism in Software-Defined Networking," *Proc. 2015 IEEE Trustcom/BigDataSE/ISPA*, vol. 1. 2015, pp. 310–17.

[14] A. Wang *et al.*, "Scotch: Elastically Scaling up SDN Control-Plane using vSwitch Based Overlay," *Proc. 10th ACM Int'l. Conf. Emerging Networking Experiments and Technologies*, 2014, pp. 403–14.

[15] R. Kokila *et al.*, "DDoS Detection and Analysis in SDN-Based Environment Using Support Vector Machine Classifier," *Proc. 2014 IEEE Sixth Int'l. Conf. Advanced Computing*, 2014, pp. 205–10.

Biographies

Kübra Kalkan (kubra.kalkan@boun.edu.tr) received her M.S. and B.S. degrees in computer science and engineering from Sabanci University in 2011 and 2009, respectively. She got her Ph.D. degree in computer engineering from Bogazici University, Turkey, in 2016. Her research interests include network security, computer networks, and wireless networks.

Gürkan Gür [SM] (gurgurka@boun.edu.tr) received his B.S. degree in electrical engineering in 2001 and Ph.D. degree in computer engineering in 2013 from Bogazici University. His research interests include cognitive radios, green wireless communications, network security, and information-centric networking.

Fatih Alagöz (fatih.alagoz@boun.edu.tr) is a professor in the Computer Engineering Department of Bogazici University. He obtained his B.Sc. degree in electrical engineering in 1992 from Middle East Technical University, Turkey, and his D.Sc. degree in electrical engineering in 2000 from George Washington University. His research areas include cognitive radios, wireless networks, and network security.

In order to maintain flows as much as possible in the data plane, some models suggest integrating some intelligence capabilities to switches. This establishes another research direction, which dwells on the trade-off between conforming to the fundamental SDN paradigm and augmenting switches with more functionalities.

ACCEPTED FROM OPEN CALL

Integrating Events into SOA for IoT Services

Yang Zhang, Jun-Liang Chen, and Bo Cheng

In IoT scenarios, there are smart devices hosting web services and also very simple devices with external web services. Without unifying the access to different kinds of devices, the construction of IoT service systems would be cumbersome. In the authors' work, integrating distributed events into SOA is the basic principle.

ABSTRACT

In IoT scenarios, there are smart devices hosting web services and also very simple devices with external web services. Without unifying the access to different kinds of devices, the construction of IoT service systems would be cumbersome. In our work, integrating distributed events into SOA is the basic principle. The data accessing capability of physical entities is separated from their actuation capability, which acts as a foundation for ultra-scale and elastic IoT applications. We then establish a distributed event-based IoT service platform to support IoT service creation and allow for the hiding of service access complexity, where the IoT services are event-driven, and impedance matching between service computation and event communication is the design goal. The coordination logic of an IoT service system is extracted as an event composition, which supports the distributed execution of the system with scalability. We finally implement applications over the platform to show its effectiveness and applicability.

INTRODUCTION

In the work of [1], it was advocated that, in Internet of Things (IoT) applications, real-world devices should provide their functionality via SOAP-based web services or RESTful application programming interfaces (APIs), enabling other components to interact with them dynamically. The functionality provided by these devices was called real-world services. They then designed a series of discovery, query, and selection schemes for these real-world services. Unfortunately, there are lots of sensors and actuators in use that are just very simple devices without the ability to provide web service interfaces. We endeavor to cope with different kinds of physical devices by integrating distributed events into service-oriented architecture (SOA) for service provision, as follows:

- The data provision functionalities of different kinds of physical devices are unified as universal IoT services, and their actuation capabilities are separated from these services and often localized. These differences are hidden behind an event-driven service infrastructure for transparent service interactions.
- There are often real-time requirements in IoT applications. The event-based communication fabric should cooperate with IoT service systems over it in order to satisfy these requirements. It is demanded that there be a service infrastructure to integrate the cooperation mechanism, the event-based communication fabric, and service environments.
- Physical entities often have their own locations such that an IoT service system over them should be distributed. The event-based IoT service infrastructure should support its distributed execution with consistency.

In some existing works, a communication foundation was optimized for IoT applications. For example, in the GridStat project [2], the publish/subscribe-based communication foundation was specially designed for smart grids, supporting, for instance, different receiving rates for the same event type. This work focused on redesigning the underpinning communication fabric to uphold real-time coordination among heterogeneous IoT services. It did not shed light on upper-level applications. We believe that connecting together all things in our environment requires the re-portrayal of the communication foundation and upper-level applications at the same time.

In contrast to GridStat, which rethought the communication foundation, there is some work that attempted to redesign the upper-level applications to accommodate the underpinning. In the work of [3], a business process was decomposed into different types of activities that were distributed in a distributed event-based system (DEBS) to improve the service system's scalability. In the work of [4], a business process fragmentation method was proposed based on distributed events, whereby each fragment acted as a distributed execution unit connected to some DEBSs. Partitioning the business processes was the sole focal point of their attention, and they did not discuss how to realize cooperation between the business processes and DEBSs.

In our work, we integrate distributed events into SOA to build an event-driven SOA (EDSOA) infrastructure for distributed IoT services. We adopt service-oriented principles to design our publish/subscribe middleware, and its network operations work on the service protocol (SOAP) to support service routing such as addressing service endpoints and delivering service invocations. For the upper-level applications, a flexible service process utility, together with service runtime environments, is built to support the distributed orchestration of IoT services, where distributed resource pools support the separation of the data accessing capability from the device actuation capability, and a business process can be partitioned and distributed by enacting its coordination logic as event compositions. Cross-layer design is also carried out besides redesign of each layer.

Digital Object Identifier:
10.1109/MCOM.2017.1600359

The authors are with Beijing University of Posts and Telecommunications.

EDSOA Platform for IoT Services

We propose a platform to support the scalable creation of event-driven IoT services based on IoT resource models, called the *EDSOA platform*. Our EDSOA platform consists of three parts: a distributed resource pool, a DEBS-based service environment, and a flexible service process utility.

There are lots of sensors and actuators in IoT applications, and most of them are simple devices without the ability to provide web service interfaces. Only some of the embedded devices host low-level and generic services. We use distributed resource pools as a base to unify interactions with different physical devices, as follows:

•Physical entities in the physical world are introduced into the information world. The digital entities corresponding to physical entities are modeled as IoT resources managed by IoT services.

•The capability to access the attributes of IoT resources is separated from the capability to actuate the IoT resources, where the former is distributed for concurrent accessing, and the latter is often deployed on one site (in an embedded device or on a nearby delegated server) for safe actuation. The soundness of such a separation lies in the fact that the actuation of physical systems is often local, and should be rigidly checked and executed according to local signals that are not all known to remote services.

•The capability to access the resource attributes is constructed based on freshly updated sensor data, where each pool management service at each of the different sites builds IoT resource instances in a real-time tuple space, and then they are refreshed by updating events received from multicast notifications.

•The publish/subscribe paradigm is used for multiple pool management/updating services to obtain recent resource attributes at the same time by multicast means and active notifications. Furthermore, the actuation capability is location-transparent based on the publish/subscribe paradigm so that remote services can use the local capability of a resource actuation service without searching for its endpoint.

The DEBS-based service environment uses distributed events as the primary mechanism to describe service interfaces, run IoT services, and support the implementation of functional reactive service functions. In this environment, the service-oriented publish/subscribe middleware acts as the distributed service bus core to route services, and it works on SOAP messages with web service notification (WSN) [5] as interfaces. In the original WSN specification [5], active push primitives were not defined; here we extend the middleware by adding endpoint management and push primitives.

A flexible service process utility is established to support implementing service coordination logic as independent event composition blocks, such that the process can be partitioned, and each sub-service could be deployed near (or in) the managed physical devices (in order to satisfy some of the real-time requirements encountered). In addition, the partitioned process fragments are flexibly executed under the condition of consistency.

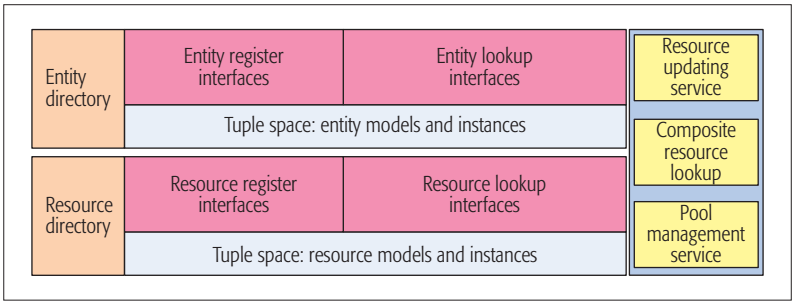

Figure 1. Resource pool.

Distributed Resource Pool

Our distributed resource pool is established over IoT resource models, and we discuss these first. In our previous work [6], an IoT resource is defined as an informational description of one sensor or physical entity in the physical world, which includes a life cycle model and an object model. The object model says which attributes the resource has as well as the relations between them. The life cycle model describes the possible ways in which the resource might progress, as well as the event-driven transitions between lifetime states.

Given the resource models, the resource pool can be designed, as illustrated in Fig. 1. Tuple space is used to store the resource instances created from their resource models, where only the most recent attributes of the IoT resources are kept, and the distributed independent tuple spaces are connected by our distributed DEBS-based service environment. In the pool, there are basic resource register interfaces and lookup interfaces. The resource updating service subscribes to data update events in the DEBS-based service environment, and the publish/subscribe service bus pushes these events from sensors to all subscribing services so that each updating service refreshes the resource instances by using the newly received resource attributes to replace the ones in the pool. Thus, IoT resource attributes and states can be locally accessed by the other services without being aware of whether they originate from smart devices or from simple devices, that is, *unifying the access to the attributes of IoT resources, and separating such access from the capabilities of actuating resources.*

Such a construction of resource pools is based on the assumption of weak data consistency. This assumption means that the attributes of distributed resource instances in different locations are not required to have the same value at every instant, but these values should stay very close over a given time slice. The soundness of this assumption relies on the fact that the IoT resources' attributes and states are periodically and continuously updated from physical sensors, and partial data loss will not impair the eventual data consistency. We establish distributed resource pools according to this assumption to support access to localized IoT resource instances.

When services attempt to access some remote resources, the pool-related services look up the resource/entity directory to find the nearest pool having the resource models and their recent instances, and then instantiates the resource models in its local pool and subscribes to the resource

We have designed a service-oriented publish/subscribe middleware, based on which we endeavor to establish an event-driven service bus to connect together wide-area heterogeneous IoT resources and services for transparent interactions without reference to service locations and online status, called Unified Message Space (UMS).

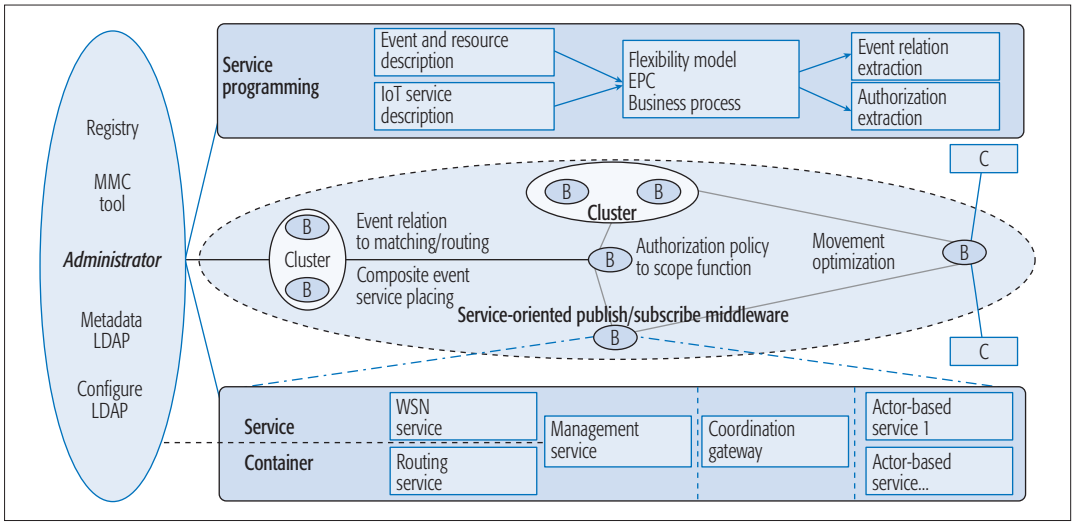

Figure 2. Unified message space.

updating events according to the models and the topic names.

DEBS-Based Service Environment

We have designed a service-oriented publish/subscribe middleware, based on which we endeavor to establish an event-driven service bus to connect together wide-area heterogeneous IoT resources and services for transparent interactions without reference to service locations and online status, called *Unified Message Space (UMS)*. It also supports cooperation between the underlying event communication fabric and the upper-level service systems, while service programming, service registry, service deployment, service running, and service routing are its basic functionalities. Our UMS consists of four parts: a service-oriented publish/subscribe middleware, multiple service containers, multiple service programming environments, and an administration center, called *Administrator*, illustrated in Fig. 2.

The service-oriented publish/subscribe middleware acts as an event-driven service bus core within our UMS. It is combined with service environments, including service containers and programming tools, to build a service bus wherein event names can be used to describe IoT service interfaces for bridging service capabilities and event communication capabilities; also, event schemas are used to self-describe events for realizing interoperability. Distributed events are the basic mechanism in our UMS, and these are named in order that they can be shared among different IoT services. A topic name is assigned to one particular kind of event, and multiple topic names form a name tree. Based on the topic names and consumers' subscriptions, the events can be pushed to consumers:

The upper-level applications and the underlying communication fabric are able to cooperate via our UMS. For example, the two problems described below are solved by this cooperation.

•Even if they subscribe to the same event topic, two different IoT services may have different requirements for event receiving rates. One service may try to receive all events over a period of time, and another may try to obtain only some parts of the event flow. They will adjust their receiving rates to within the allowed range to avoid congestion in the communication fabric, while our UMS can optimize its performance by appropriately splitting event flows for the two different subscribers.

•In our UMS, event delivery reliability cannot be realized by the retransmission method because events are delivered by multicast communication. Some cross-layer cooperation is needed. If an IoT service requires the events with a particular topic name to be completely reliably delivered, it actively requests the lost parts, and the UMS, having cached multiple event instances, will then directly respond to the request for lost event instances by the ones in caching.

From the above discussion, we know that, compared to existing work [3, 4], our two-layer redesign not only improves the performance of UMS, but also provides opportunities to overcome some difficulties such as reliable event delivery. Our main goal is to comprehensively explore the appearing relationship among events, IoT resources, IoT services and the underlying event delivery fabric to realize impedance matching between service computation and event communication. The design tenet is to establish an integrated architecture for a distributed IoT service platform by two-layer redesigns, where one event name is used to describe a service interface linking to a reactive service function. The same name is also used to represent an event flow for matching and routing in the communication fabric, and the same cross-layer entity could be utilized to address some hard issues. For example, we discuss how to protect IoT services and their interacting events by integrating the access control into event routing. The event-driven IoT services are decoupled from each other and do not have their event access requests sent to publishers, so the classic request interception mechanism is unavailable for authorization policy enforcement. In addition, the variety and complexity of the service provision required to access different kinds of physical devices are further hidden by some gluing architectural blocks such as resource pools — which unify resource access.

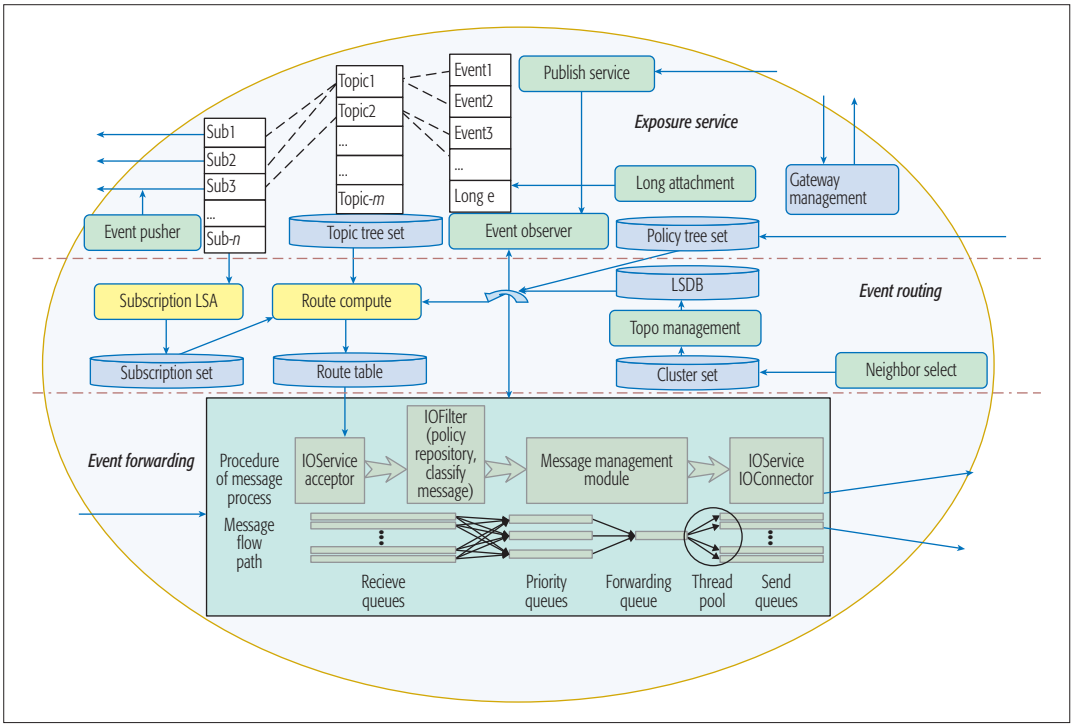

Figure 3. The service-oriented publish/subscribe middleware.

The middleware's key functionality is service routing, which is realized by a routing layer combined with an event forwarding layer. In the routing layer, all hosts, routers, and clients are organized into clusters by a neighbor selection algorithm.

UMS

An actor model [7] is used to abstract the event-driven service programming whereby a service listens to incoming messages according to their name strings, and then reacts by posting their processed results using other name strings. There are three types of actors. The first kind of actor is a *functionality unit*, which is managed by the second kind of actor, called *chatroom*. The chatroom represents a service consisting of multiple functionality units, and acts as a service coordination layer to coordinate itself to participate in a composite system. The third kind of actor is a *service composition*, and this role can also be assumed by a chatroom. The Scala software [8] is extended to realize actor-based services and maintain service running, where service interface descriptions, event descriptions, and resource descriptions are created, stored, and referenced.

The service-oriented publish/subscribe middleware consists of a communication capability exposure layer, an event routing layer, and a central management layer, as illustrated in Fig. 3. The communication capability exposure layer provides local clients with a set of service interfaces to use and configure communication capabilities, which is an extension of WSN [5] called *WSN service*. The event routing layer provides vertical layer APIs to the WSN service, and provides peer-to-peer topology — maintaining functionality, link state update functionality, and event forwarding functionality to other routing layers, called *routing service*. The central management layer stores configuration information, meta-data about topic trees and authorization policies, and runtime information concerning network nodes; it also provides data access APIs to all network nodes.

All services, including the *WSN service* and *routing service*, are deployed in service containers. The middleware's configuration and management functionalities are also deployed in the service container as a *management service*, which is implemented based on a Lightweight Directory Access Protocol (LDAP) specification. When a network node starts, it registers to the *administrator service* to get the configuration information, topic trees, authorization policies, and online nodes' service endpoints.

Service Routing and Access Control

The publish/subscribe middleware in Fig. 3 provides a set of service interfaces whereby its clients can express their interest in events, publish their own events according to predefined topics and event schemas, and reconfigure the middleware to satisfy their special requirements. The middleware's key functionality is service routing, which is realized by a routing layer combined with an event forwarding layer. In the routing layer, all hosts, routers, and clients are organized into clusters by a neighbor selection algorithm. Although the nodes are grouped into clusters, the clusters themselves and their connections often change with time elapsed — clusters disappear, and links change among clusters. Topology maintenance is carried out by the periodic and event-triggered advertising of link states. The clients' subscriptions are also disseminated into each cluster, and the topology advertisement and subscription advertisement are often aggregated to reduce communication cost. In the event forwarding layer, event matching and multiple-priority queue maintenance are the two prime functionalities, and events are forwarded to neighboring clusters by looking up routing tables from the routing layer.

Figure 4 illustrates an example of path computation, where the overlay network includes clusters G_1, G_2, ..., G_9 in its left column without authorization constraints: G_1 publishes events with topic t, G_2, G_5, G_7, G_8, and G_9, all subscribe

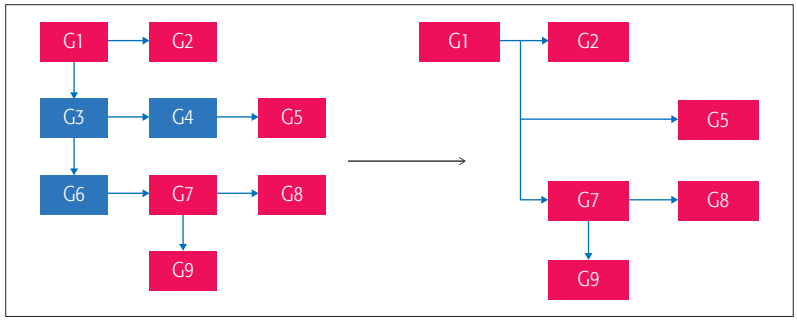

Figure 4. Multicast tree for one topic.

to topic *t*. The authorization policy applied in the figure's right column defines that G_3, G_4, and G_6 have no rights to visit the events with topic name *t*, so the multicast tree is changed.

In traditional publish/subscribe systems, brokers evaluate their own authorization policies against clients' requests for network operations, such as registering topics and reading events, while the invocation protection requirements of IoT services over them are not considered. Because our event-driven decoupled IoT services indirectly and anonymously interact, the classical assumption that access to services and their interacting events is controlled by an omniscient monitor performing perfect surveillance on requests will become impossible. In our work, the cross-layer design is then intended to integrate the upper-layer service access control into the bottom-layer event routing, which addresses the hard issue of losing omniscient monitors. The basic idea is as follows:

A security policy is efficiently embedded into an event, an independent meaningful entity in the communication fabric, such that the policy can be evaluated independent of its security context. Broker nodes at the edge of publish/subscribe middleware then make decisions on event delivery according to the embedded policies. In addition, the broker network topology of publish/subscribe middleware can be translated into a graph with each node being guarded by multiple Boolean values, which are the result of comparison between the subscribers' attributes and the publishers' policies. We then compute many-to-many multicast trees over the Boolean-value-labeled graph, which does not route an event to unauthorized services and vice versa.

Distributing IoT Service Process with Scalability

The flexible service process utility implemented in our platform supports the realization of service scalability, where the coordination of IoT services is modeled by an event-driven process chain (EPC) graphical language [9]. The scalability of EPC-based IoT processes is achieved by distributing it but at the same time avoiding inconsistencies from this distribution.

An EPC-based IoT service process is fragmented into multiple process fragments and multiple coordination logics represented by event relations. In order to execute the process fragments with distribution, we propose an execution model for the IoT service whereby each service is armed with a coordination layer, the service is executed as usual, and the coordination layer mediates the local service and the coordination logic in the process [6], illustrated in Fig. 5.

In order to avoid state inconsistency, a special service is used to explicitly extract event relations from different process fragments and translate them into composite events. That is to say, the special service — complex event processing (CEP) [10] — handles the event relations and publishes a new event with a new name to represent the result of event relation processing. Such a result is also an event in its own right, called a *composite event*. This special service has multiple instances during runtime, and each instance copes with a part of the event relations to avoid bottleneck effects. The service coordination layer subscribes to the composite event by its topic name. When the composite event arrives at the coordination layer, the layer enables the corresponding service interfaces.

Here, an online controller is used in the service process environment, as shown in Fig. 5, where the IoT service system is allowed to be adapted at runtime for accommodating the varying environment, and the controller checks the service trace just before process activities are actually executed. If the execution of an activity violates some predefined service properties, the online controller lets the CEP service switch off the publication of the composite events that drive the activity. Otherwise, the controller lets the CEP service switch on the publication of the composite events. Introducing CEP services and online controllers into the process environment makes it possible for IoT service processes to run flexibly with scalability.

Compared to activity-distribution process decomposition methods [3, 4], our method is to explicitly separate coordination logic from a business process, where the issue of distributing sequential process activities (a kind of glitch avoidance problem in event-driven applications) is addressed by separating the sequence relation into event causality [6]. Event relation is enacted into composite events, which act as a consistent result for distributed components without additional agreement protocols.

Applications and Experiments

We have implemented our EDSOA platform, and have deployed on it a coal mine monitor-control system (CMCS) application. In Shanxi province, the Sanyuan coal company has multiple coal mines, where each coal mine has a local CMCS, and all coal mines' CMCSs and the company headquarters' CMCS collaborate to keep the coal mining operations safe and efficient. In each coal mine, there are a great many meter devices that generate real-time metered data collected by the programmable logic controller (PLC) and/or receive actuation instructions from the PLC.

We have also conducted experiments on the CMCSs to measure the throughput and jitter performance of service bus. We set a simple and symmetric topology used to test the basic performance. At the horizontal axis in Fig. 6, the total number of events injected into the system is gradually increased through the different experimental steps. By the left vertical axis, the number of events processed per second by the system is given to represent the system's throughput at each experimental step. From the figure,

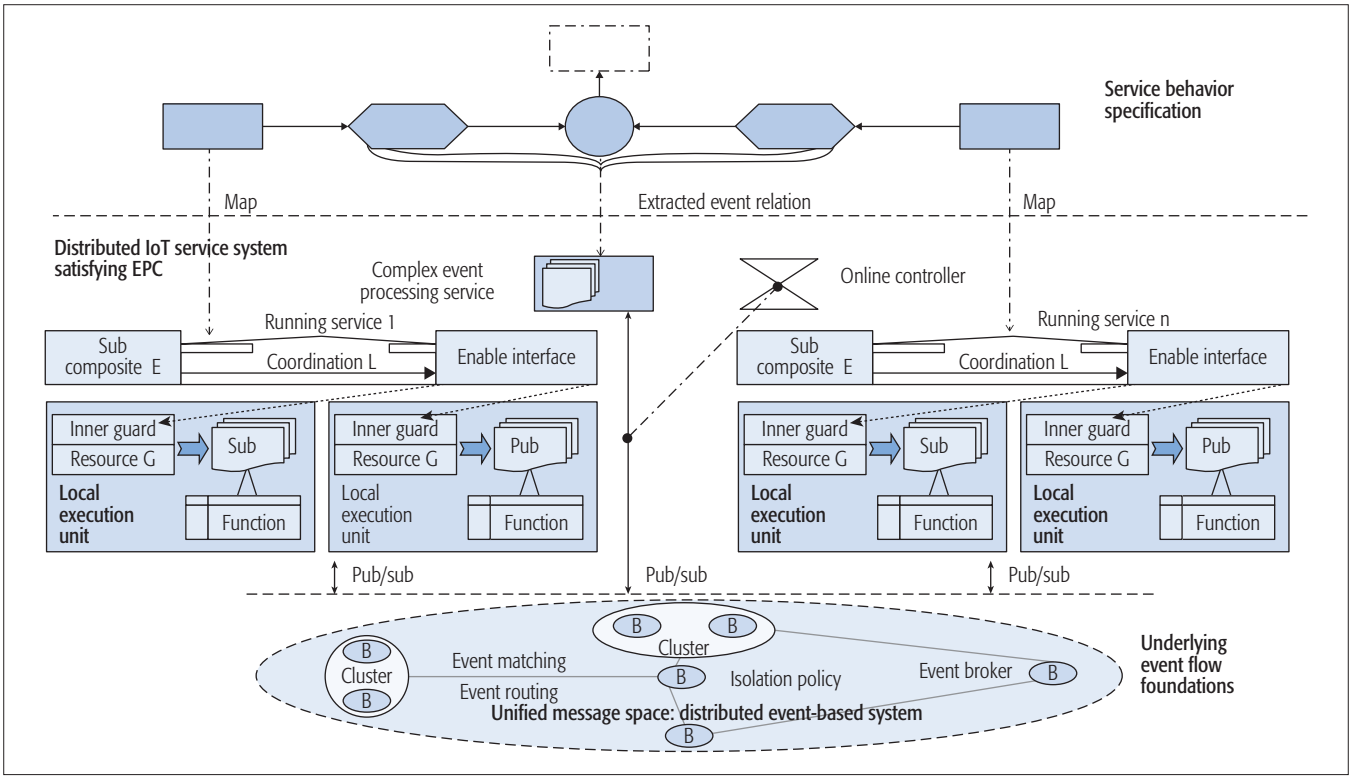

Figure 5. Service execution and distribution.

we can see that the throughput increases as the total number of injected events and their speed of publication increases; however, at the 2400k horizontal point, the throughput reaches its maximum, and then this decreases as the total number of injected events and their speed of publication continues to increase. The increasing slope of throughput is not sharp, and the system runs stably with little jitter.

Related Work

Our research team discussed using EDSOA in IoT [11, 12]. The technology surveys in [13, 14] also considered using the event-driven methodology for IoT applications. But these papers did not discuss why such an architecture is applicable and how it could be used.

The PLAY project [15] tried to build an ultra-scale federated service platform based on dynamic and complex event interaction patterns; this was similar to our work here. However, although our work has a similar goal to PLAY, the design and implementation of ours are different.

There are two kinds of existing related work. One is to redesign the publish/subscribe middleware to accommodate the applications. The other is to redesign the applications to utilize the functionalities of middleware to realize scalability and high performance. In [2], the GridStat project proposed the publish/subscribe-based WAMS-DD infrastructure for a smart grid, where different event consumers can express different receiving rates for one kind of event. It did not consider adjusting the IoT services to accommodate the WAMS-DD.

In the work of [3, 4], redesigning applications to accommodate publish/subscribe middleware was considered. They both adopted the decomposing method to partition an upper-level application to avoid the bottleneck effect of centrally executing business processes. They did not consider the consistency issue involved with running different partitioned units, which is our focus here, with modeling coordination logic as event compositions.

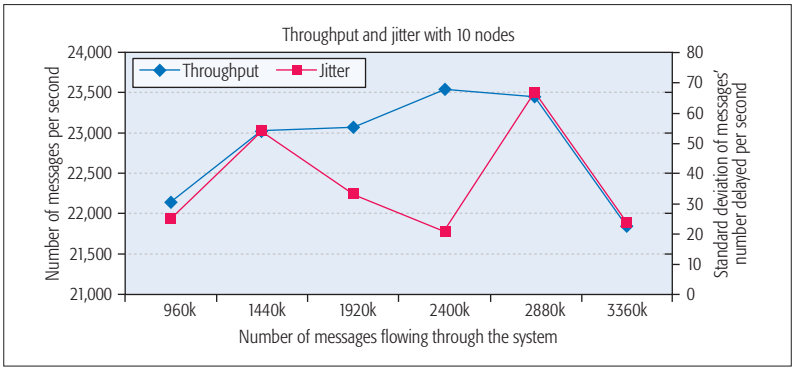

Figure 6. Throughput and jitter with 10 nodes.

Conclusions

In this article, we first design a service-oriented publish/subscribe middleware to support service interaction and service routing with the communication capabilities being exposed as services for upper-level applications to use. We then try to establish an event-driven service infrastructure based on this middleware to converge IoT services and IoT data with scalability, where the distributed IoT resource pool supports the separation of the attribute accessing capabilities and the resource actuation capabilities, the DEBS-based service environment supports event-based service programming and running, and the service pro-

In our work, we integrate distributed events into SOA to build an EDSOA (Event-Driven SOA) infrastructure for distributed IoT services. We adopt service-oriented principles to design our publish/subscribe middleware, and its network operations work on the service protocol (SOAP).

cess utility supports separating coordination logic from service computation logic for distributed execution of IoT service processes. Applications and experiments have shown our solution's effectiveness and applicability.

Acknowledgment

This work is supported by the National Natural Science Foundation of China (no. 61372115, no. 61132001), the National Grand Fundamental Research 973 Program of China under Grant no. 2013CB329102, and the National High-Tech R&D Program of China (863 Program) under Grant no. 2013AA102301.

References

[1] D. Guinard *et al.*, "Interacting with the SOA-Based Internet of Things: Discovery, Query, Selection, and On-Demand Provisioning of Web Services," *IEEE Trans. Services Computing*, vol. 3, no. 3, 2010, pp. 223–35.
[2] D. E. Bakken *et al.*, "Smart Generation and Transmission with Coherent, Real-Time Data," *Proc. IEEE*, vol. 99, issue 6, 2011, pp. 928–51.
[3] G. L. Li, V. Muthusamy, and H. A. Jacobsen, "A Distributed Service-Oriented Architecture for Business Process Execution," *ACM Trans. the Web*, vol. 4, no. 1, article 2, 2010, pp. 2:1–33.
[4] P. Hensa *et al.*, "Process Fragmentation, Distribution and Execution Using an Event-Based Interaction Scheme," *J. Systems Software*, vol. 89, Mar. 2014, pp.170–92.
[5] OASIS, "OASIS Web Services Notification (WSN)," OASIS Completed Version, Oct. 2006.
[6] Y. Zhang and J. L. Chen, "Constructing Scalable IoT Services Based on Their Event-Driven Models," *Concurrency and Computation: Practice and Experience*, vol. 27, issue 17, 2015, pp. 4819–51.
[7] G. A. Agha, "ACTORS: A Model of Concurrent Computation in Distributed Systems," *The MIT Press Series in Articial Intelligence*, MIT Press, 1986.
[8] M. Schinz and P. Haller, "A Scala Tutorial for Java Programmers," Scala Language v. 1.3, Jan. 16, 2014.
[9] J. Mendling, *Detection and Prediction of Errors in EPC Business Process Models*, dissertation, Vienna Univ.of Economics and Business Administration, 2007.
[10] A. Lundberg, "Leverage Complex Event Processing to Improve Operational Performance," *Business Intelligence J.*, vol. 11, no. 1, 2006, pp. 55–65.
[11] D. Zhu *et al.*, "Towards a Flexible Event-Driven SOA Based Approach for Collaborating Interactive Business Processes," *IEEE Int'l. Conf. Services Computing*, Washington, DC, June 2011, pp. 749–50.
[12] S. Zhao, *Research on Internet of Things Resource Management Platform and Service Provision Framework*, dissertation, Beijing Univ. Posts and Telecommun., 2014.
[13] L. Atzori, A. Iera, and G. Morabito, "The Internet of Things: A Survey," *Computer Networks*, vol. 54, issue 15, 2010, pp. 2787–2805.
[14] J. Gubbi *et al.*, "Internet of Things (IoT): A Vision, Architectural Elements, and Future Directions," *Future Generation Comp. Syst.*, vol. 29, issue 7, 2013, pp. 1645–60.
[15] N. Stojanovic, "Pushing Dynamic and Ubiquitous Interaction between Services Leveraged in the Future Internet by Applying Complex Event Processing," FP7 PLAY Project 258659, June 2010.

Biographies

Yang Zhang (YangZhang@bupt.edu.cn) received his Ph.D. degree in computer applied technology from the Institute of Software, Chinese Academy of Sciences in 2007. His research interests include service-oriented computing, the Internet of Things, and service security and privacy. He currently works at the State Key Laboratory of Networking and Switching Technology, Beijing University of Posts & Telecommunications (BUPT), China. He leads a team doing scientific research on the theoretic foundation of EDSOA for IoT services.

Jun-Liang Chen (chjl@bupt.edu.cn) graduated from the Department of Telecommunications, Shanghai Jiaotung University in 1955 and received his Doctor of Engineering degree from the Moscow Institute of Electrical Telecommunications, former Soviet Union, in 1961. He is a professor and director of the Research Institute of Networking Technologies of BUPT. He is a member of both the Chinese Academy of Sciences and the Chinese Academy of Engineering.

Bo Cheng (ChengBo@bupt.edu.cn) received his Ph.D. degree in computer applied technology from the University of Electronic Science and Technology of China in 2006. His research interests include service-oriented computing, the Internet of Things, and mobile services. He currently works at the State Key Laboratory of Networking and Switching Technology, BUPT.

ACCEPTED FROM OPEN CALL

Alternative Networks: Toward Global Access to the Internet for All

Jose Saldana, Andrés Arcia-Moret, Arjuna Sathiaseelan, Bart Braem, Ermanno Pietrosemoli, Marco Zennaro, Javier Simó-Reigadas, Ioannis Komnios, and Carlos Rey-Moreno

ABSTRACT

It is often said that the Internet is ubiquitous in our daily lives, but this holds true only for those who can easily access it. In fact, billions of people are still digitally disconnected, as bringing connectivity to certain zones does not make a good business case. The only solution for these unsatisfied potential users is to directly undertake the building of the infrastructure required to obtaining access to the Internet, typically forming groups in order to share the corresponding cost. This article presents a global classification and a summary of the main characteristics of different Alternative Network deployments that have arisen in recent years with an aim to provide Internet services in places where mainstream network deployments do not exist or are not adequate solutions. The "Global Access to the Internet for All" Research Group of the Internet Research Task Force, where all authors actively participate, is interested in documenting these emerging deployments. As an outcome of this work, a classification has converged by consensus, where five criteria have been identified and, based on them, four different types of Alternative Networks have been identified and described with real-world examples. Such a classification is useful for a deeper understanding of the common characteristics behind existing and emerging Alternative Networks.

INTRODUCTION

It is often claimed that the Internet is a part of our daily lives, but the reality is that in 2016 there were only around 3 billion Internet users in the world, out of a population of over 7 billion people. The reasons behind this lack of usage cannot be entirely attributed to limitations of infrastructure, as global satellite and mobile data coverage are widely available. It is estimated that over 5.5 billion of the world's population have access to 3G communications, yet 2.5 billion are not using the Internet [1]. Even though factors such as the lack of relevant content and inadequate digital skills among those offline are also responsible for this situation, it is widely acknowledged that the main reason for this gap is cost [2].

In this context, finding alternative deployment models that may reduce the cost of communications is a matter of urgent concern, as highlighted by the numerous relevant initiatives worldwide, including the Global Connect Initiative (https://share.america.gov/globalconnect/); Internet for All (https://www.weforum.org/projects/internet-for-all); 1 World Connected (http://1worldconnected.org/); and the UN Internet Governance Forum 'Policy Options for Connecting and Enabling the Next Billion' framework (http://www.intgovforum.org/cms/policy-options-for-connection-the-next-billion). We note that a number of sessions at the latest Internet Governance Forum, convened by the United Nations in Guadalajara, Mexico, in December 2016, (https://igf2016.sched.com/) were devoted to the discussion of this issue.

The present article addresses this gap by presenting a survey of different alternative models identified through a consensus process achieved by the Internet Research Task Force (IRTF) Global Access to the Internet for All (GAIA) Research Group. This consensus crystallized in a Request for Comments on "Alternative Network Deployments" that this article summarizes [3]. Alternative Networks are considered those that share some of the following characteristics:

- They have a relatively small scale.
- They may follow de-centralized approaches.
- The investment in infrastructure may be low, and may be shared by independent users, commercial and non-commercial entities.
- Users may be involved in the design, deployment, maintenance and daily operation of the network.

In particular, we explain the criteria and present a classification of Alternative Networks into four distinct types, detailing the main characteristics of each one, as well as the technologies they rely on through real life examples. To the best of our knowledge, this classification does not exist in the literature and provides a guide to people interested in non-traditional deployments, ranging from researchers to community members, and a set of references for further research into each of them.

In the next section, the key challenges that Alternative Networks aim to solve are discussed. We then present the classification criteria, detail the classification of Alternative Networks into four distinct types, and refer to emerging types of networks. Finally, findings are summarized.

It is often said that the Internet is ubiquitous in our daily lives, but this holds true only for those who can easily access it. In fact, billions of people are still digitally-disconnected, as bringing connectivity to certain zones does not make a good business case. The only solution for these unsatisfied potential users is to directly undertake the building of the infrastructure required to obtaining access to the Internet, typically forming groups in order to share the corresponding cost.

Jose Saldana is with the University of Zaragoza; Andrés Arcia-Moret and Arjuna Sathiaseelan are with the University of Cambridge; Bart Braem is with the University of Antwerp; Ermanno Pietrosemoli and Marco Zennaro are with The Abdus Salam International Centre for Theoretical Physics; Javier Simó-Reigadas is with the Universidad Rey Juan Carlos; Ioannis Komnios is with Democritus University of Thrace; Carlos Rey-Moreno is with the University of the Western Cape.

Digital Object Identifier: 10.1109/MCOM.2017.1600663

Alternative Networks have emerged to provide Internet services to areas not covered by traditional operators due to high cost or challenges that commercial networks are ill-equipped to solve. Such challenges range from privacy concerns to limited power resources or lack of technical expertise.

Challenges that Alternative Networks Aim to Solve

Alternative Network deployments are nowadays present in every part of the world. Even in high-income countries, they are being built as an alternative to commercial networks managed by traditional network operators. Alternative Networks have emerged to provide Internet services to areas not covered by traditional operators due to high cost or challenges that commercial networks are ill equipped to solve. Such challenges range from privacy concerns to limited power resources or lack of technical expertise. In this work, we do not aim at providing an exhaustive list of these challenges. Instead, we focus on two key aspects that trigger the development and deployment of Alternative Networks: the digital divide and the differentiation of areas based on geography and user density.

Digital Divide

According to the ITU's report *ICT Facts and Figures* 2016 [1], half the people on Earth are still disconnected from the Internet. Furthermore, the connected population is unevenly distributed: while 84 percent of households are connected in Europe, in the African region only 15.4 percent of households are connected. The digital divide between "Global North" and "Global South" is based on information and communication technologies (ICT) factors such as:

- The availability of both national and international bandwidth.
- The difficulty to pay for the services and the devices required to access the ICTs.
- The instability (or lack) of power supply.
- The scarcity of qualified staff.
- The existence of a policy and regulatory framework that hinders the development of Alternative Network deployment models, favoring instead state monopolies or entrenched incumbents.

The uneven digital development state of a country may produce another form of inequality, which involves infrastructures, the ICT sector, digital literacy, legal and regulatory framework, as well as content and services. In this context, the concept of digital divide refers to the limitation or the total absence of one or more of these dimensions. This divide constitutes a new inequality vector that may simultaneously generate progress for some, while creating economic poverty and exclusion for others, as happened during the Industrial Revolution. It is undeniable that mobile network operators have certainly contributed to lowering the divide, but at the same time the model they follow for increasing connectivity has some restrictions that result in a limitation of the development outcomes. Furthermore, a significant part of the costly bandwidth may be spent on updates, advertising and other data not contributing to development or economic inclusion.

Thus, prices are still unaffordable to many people, as they may constitute an exaggerated percent of an individual's income, hindering one's willingness to invest in communications. Furthermore, the cost of prepaid packages, which are the most suitable option for informal economies, is high when compared with the rate of post-paid subscribers.

In this context, in November 2015, the World Summit of the Information Society called upon governments, the private sector, civil society and international organizations to work actively to bridge the digital divide by achieving "a people-centered, inclusive and development-oriented Information Society," allowing access for everyone to information and knowledge to achieve sustainable development and improve the quality of life.

Alternative Networks can be seen as a way for civil society and local stakeholders to become more active in the promotion of affordable alternatives to connect themselves to the Internet. Additionally, these networks can enhance other dimensions of digital development, such as increased human capital and the availability of localized content and services, fulfilling the specific needs of each local community.

Adverse Geography and Low User Density

The digital divide presented in the previous subsection is present in different countries, but also among different regions within a country. Such is the case for rural inhabitants, who represent more than half of the world's population. The disposable income of citizens in rural areas, with many surviving on a subsistence economy, is typically lower than those inhabiting urban areas. Additionally, a significant percentage of the disconnected population is located in geographies difficult to access and/or exposed to extreme weather conditions, sometimes even lacking electrical infrastructure. From a networking point of view, customers in rural zones are spread over a wider area and are typically located farther from the Internet access point compared to urban users.

As an example, Fig. 1a shows the mobile network coverage map in Johannesburg, while Fig. 1b depicts the two different zones, which can be defined as urban/suburban and rural. Figure 1a highlights the coverage variation between the connected urban/suburban areas in color and the rural areas with no coverage in white. The latter create an ideal niche for the deployment of Alternative Networks.

In rural areas, low population density discourages telecommunications operators from providing the services offered in urban areas due to lack of profitability. This situation has motivated residents and stakeholders of certain rural areas to become the owners of an Alternative Network deployment. The cost of the required wireless infrastructure to set up a network, including a proper power supply (e.g., via solar energy), is within the affordability range of many rural individuals or small communities. This means that they can share the cost of the infrastructure and the Internet gateway and access the network via inexpensive wireless devices. Some examples are presented in [4, 5].

Classification of Alternative Networks

The discussion within the GAIA Research Group started with the identification of the criteria to be used in the classification of the different types of Alternative Networks. Only then could we build a coherent classification of the existing networks, which have been divided into community networks, wireless internet service providers, shared infrastructure model, and crowdshared approaches. This section explains both the criteria and the classification,

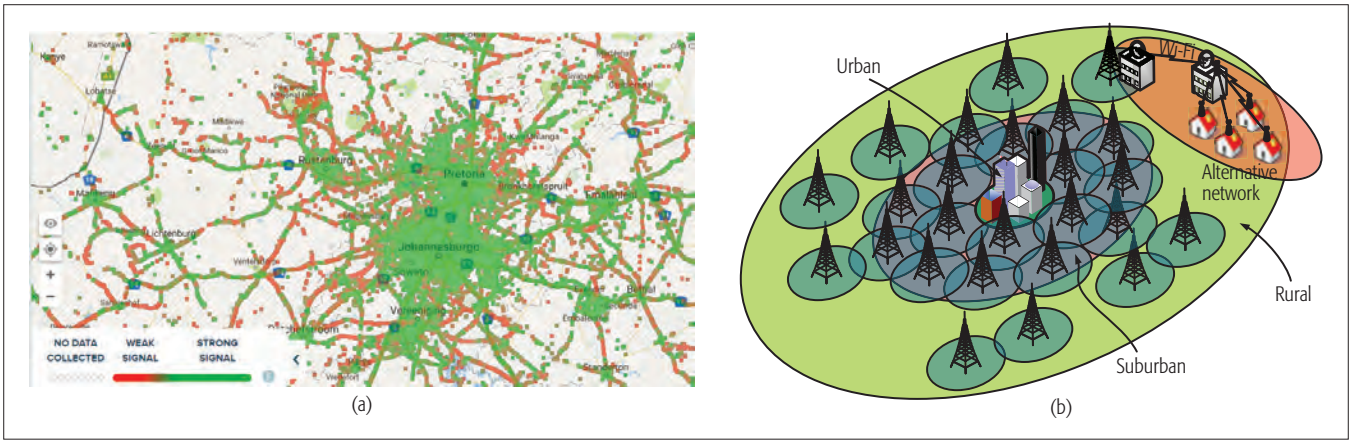

Figure 1. Urban, suburban and rural zones' network coverage: a) 2G/3G/4G mobile network coverage map in Johannesburg (http://opensignal.com/) (white color: no coverage; red color: weak signal; green color: strong signal); b) a typical Alternative Network deployment for an under-served rural area.

along with real examples. We conclude with a discussion of emerging Alternative Networks.

Classification Criteria

After a detailed study of existing deployments, and a long discussion within the IRTF Research Group, five criteria that differentiate existing Alternative Networks have been identified. We note that the criteria are not "fully orthogonal," as is obvious from the description of the different network types. In particular, the classification criteria include the following.

Entity Behind the Network: The entities or individuals that start, manage and push the network can be a public stakeholder, a community of users, or even a private company. Each of these entities can build and manage a network on their own or collaborate with each other, sharing network resources (e.g. "crowdshared" approaches). In Fig. 2, we depict the three possible promoting entities and showcase where the different types of Alternative Networks (detailed in the next subsection) fall.

Purpose: The purpose and benefits of Alternative Networks can be classified depending on their economic, political, social or technological objectives. Both the society as a whole and specific actors can enjoy the benefits provided by these networks, such as:

- Extending coverage to under-served areas (users and communities).
- Providing affordable Internet access for all.
- Reducing the initial capital expenditures (CAPEX) for the network, end user, or both.
- Providing additional sources of capital beyond the traditional carrier-based financing.
- Reducing ongoing operational costs (OPEX) such as backhaul, power provisioning or network administration.
- Reducing hurdles to adoption as digital literacy or literacy in general.
- Leveraging expertise and having a place for experimentation and teaching, including research purposes.
- Sharing connectivity, resources and local content.

As far as users are concerned, other underlying motivations may be present:

- Their desire for affordable sharing of Internet connectivity.
- The experience of becoming active participants in the deployment and management of a real and operational network.
- Raising awareness of political debates around issues like network neutrality, anti-censorship and more.

Administrative Model: The administrative model can either be centralized, where a single entity plans and operates the network, or non-centralized, where the network is managed following a distributed approach, in which a whole community may participate, including the enhancing of the network by the addition of new users.

Technologies Employed: Alternative Networks employ a variety of technologies to achieve connectivity, including optical fiber, femtocells, variations of WiFi, WiMAX and dynamic spectrum access solutions. Figure 3 depicts these technologies and the type of Alternative Networks where they are usually employed. Other options may exist, but the most common ones have been included in the figure.

Optical fiber has been used in cases where national service providers decline to bring connectivity to isolated villages, so the community decides to build their own fiber network. Such examples include Lowenstedt in Germany and parts of Guifi.net in Spain, which consists of more than 33,000 nodes [6].

Licensed mobile spectrum has also been exploited through the use of femtocells, i.e., small, low-power cellular base stations. Even though the paradigm of femtocells was conceived to improve indoor coverage, it has proven to be a feasible solution for bringing 3G coverage to under-served rural areas with low population density, as the number of users and the covered area are small enough to be managed by a low-cost femtocell. Moreover, if the community already owns an IP network for other purposes, sharing that infrastructure with the 3G operator as a low-cost backhaul may dramatically reduce the costs for the operator and make the service sustainable for small communities that could not be served otherwise [7].

IEEE 802.11 (WiFi) is by far the most popular standard in Alternative Networks; its different variants (a/b/g/n/ac/ad/af) use unlicensed bands, thus defying spectrum costs. The medium access control (MAC) is based on carrier sense

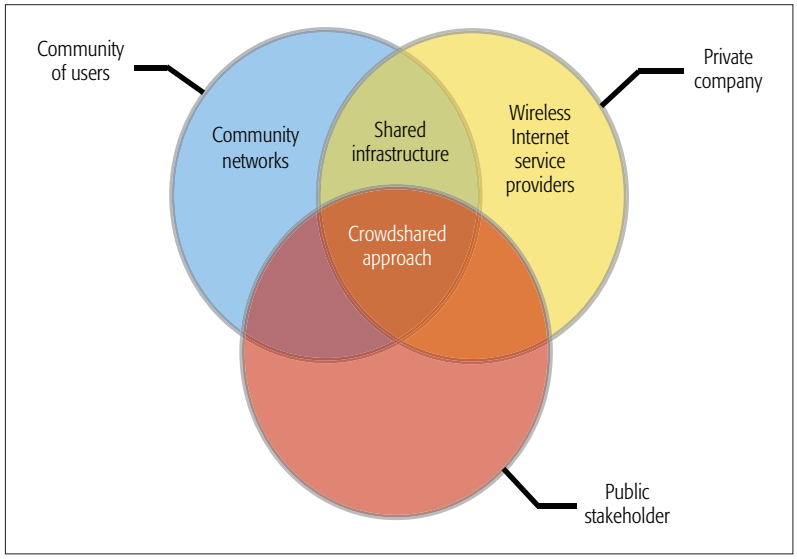

Figure 2. Entity behind the network and type of Alternative Network.

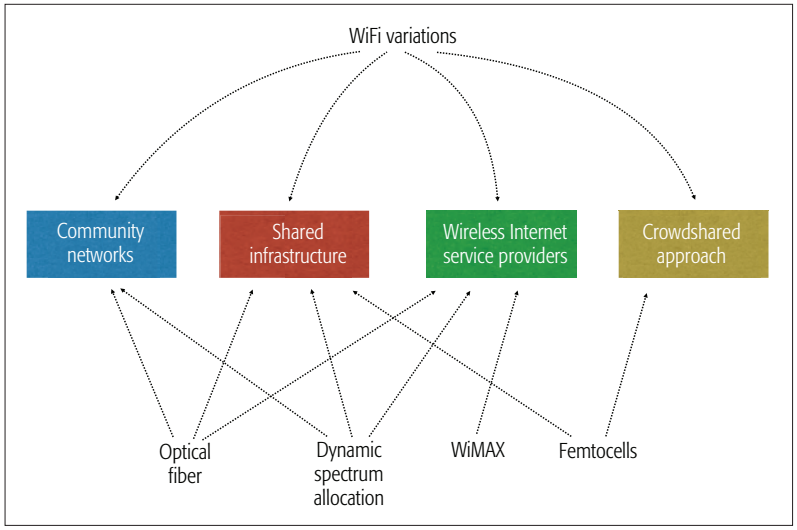

Figure 3. Employed technologies and type of Alternative Network.

multiple access with collision avoidance (CSMA/CA), and was designed for short distances, so modifications of MAC parameters are required for long distance links. Some of these modifications (e.g., WiFi over long distance (WiLD)) [8] are frequently employed in Alternative Networks. However, a modified contention MAC is still inefficient at long distances. Many manufacturers have developed alternative time division multiple access (TDMA) MAC protocols for long-distance 802.11-based products that can be activated as a CSMA/CA replacement on a per-link basis. As a result, low-cost equipment using these techniques can achieve high throughput even at distances beyond 100 kilometers.

WiMAX systems (IEEE 802.16-compliant) over non-licensed bands have also been employed in certain cases. WiMAX can enable usage at distances up to 50 km while achieving high spectral efficiency [9].

Finally, nowadays there is an increasing interest in exploiting TV white spaces in regions where parts of the VHF and UHF spectra are unused, by means of dynamic spectrum access solutions. There are emerging technologies that detect those unused fragments of the spectrum by jointly sensing and querying spectrum databases, so they can be leveraged by secondary users with no harmful interferences to primary users. Cognitive radio techniques permit the dynamic adaptation of the transmission power, modulation and frequencies, as required by these solutions. The two dominant standards for TV white spaces are IEEE 802.11af (specifically adapted from 802.11) and IEEE 802.22, designed for long-range rural communication.

Typical Scenarios: Based on the challenges described above, Alternative Networks can be found in urban/suburban and rural areas of both "Global North" and "Global South" countries, although some types of networks are more typical in certain zones.

Common Types of Alternative Networks

Having defined the classification criteria, we present a classification of Alternative Networks. Four different types of networks have been identified, explained in detail below, including some real-world examples for each one. Table 1 summarizes the characteristics of each type of network.

Community Networks: Community Networks are large-scale, self-managed networks that are built and organized in a non-centralized and open manner. As participation in a Community Network is open, they grow organically, since new links are created every time a host is added. This is done via the sharing of an open peering agreement among all members, with the common objective of freely connecting them and increasing network coverage. In this sense, members of a Community Network are not only users, but active contributors to the network. In most cases, members keep ownership of the part of the infrastructure they have contributed to build. Thus, the network presents a high degree of heterogeneity with respect to the devices used in the infrastructure and its management. This results in increased entropy, as different protocols (e.g., routing) may be used in different parts of the networks. However, on the positive end, it allows the increase of the network size without incurring in major costs. One example that represents this model is Guifi.net [6], which has shown an exponential growth rate in the last decade, both in the number of nodes and end users. Figure 4a shows the structure of this network around Barcelona, Spain. As can be seen, the network covers both urban and rural areas, usually connected through long-distance links (the so called community mesh approach). In networks covering remote rural areas, tree and mesh topologies are frequent because they follow available terrestrial infrastructures such as rivers or roads that connect villages to the closest well-connected city. Figure 4b depicts a real network supernode used in Guifi.net, built using typical common off-the-shelf equipment, such as a Raspberry Pi.

Given that the ownership of the network is open and non-centralized, Community Networks incentivize the transfer of knowledge in order to maintain and expand the existing infrastructure. Another characteristic resides in the way community members organize themselves not only to control the usage of the network, but its opera-

	Entity behind the network	Purpose	Administrative model	Technologies employed	Typical scenarios
Community networks	Community of users	All goals mentioned above	Non-centralized	WiFi variations Optical fiber Dynamic spectrum allocation	Urban/suburban and rural
Wireless Internet service providers (WISPs)	Private company	To extend coverage to underserved areas To reduce CAPEX To provide additional sources of capital	Centralized	WiFi variations Optical fiber WiMAX Dynamic spectrum allocation	Suburban and rural
Shared infrastructure model	Community of users Private company	To eliminate a CAPEX barrier for operators To decrease the OPEX being supported by the community To extend coverage to underserved areas	Non-centralized	WiFi variations Optical fiber Femtocells Dynamic spectrum allocation	Rural in "Global South" countries
Crowd shared approach	Community of users Private company and Public stakeholder	To share connectivity and resources	Non-centralized	WiFi variations Femtocells	Urban/suburban and rural

Table 1. Alternative Networks: characteristics and classification.

tion as well, as certain tasks like IP addressing and routing require a minimum governance infrastructure. This participatory model has proven to be effective in connecting sparse populations, which is key for the enhancement and extension of digital Internet rights. This participatory model also plays a role in the range of services offered by a Community Network, which can be used as a backhaul for services that are either completely free or commercial, depending on the preferences of their members.

Wireless Internet Service Providers (WISPs): Wireless Internet Service Providers (WISPs) are commercial entities that use wireless technologies in order to create the infrastructure required to provide Internet and/or Voice over IP (VoIP) services. They are common in areas not covered by traditional operators. WISPs mostly employ wireless point-to-multipoint links using unlicensed spectrum. However, these bands face challenges in some places, either for the overcrowding of such spectrum, which compromises the quality of service, or where the regulatory framework forbids its use. In these cases, WISPs are resorting to the use of licensed frequencies.

Local companies operate most WISPs, responding to a perceived market gap. Nevertheless, a non-negligible number of WISPs, such as AirJaldi in India, have expanded from local service into multiple locations. For the past decade, most WISPs using cloud-managed solutions have been in the "Global North" markets. In 2014, a similar cloud-managed service initiative, aimed at the "Global South" markets, appeared; Everylayer uses a proprietary cloud-based platform to coordinate low-cost WiFi and fiber optic high-speed last mile connections.

Shared Infrastructure Model: Because of the low returns expected, operators may be reluctant to deploy network infrastructures in large, sparsely populated areas. This happens when the usual model is followed, in which a mainstream operator deploys and owns the infrastructure, or rents it to/from other companies. However, if a community of users already owns a network infrastructure (e.g., connecting a public building, a medical dispensary, and so on), it can be shared with an operator, resulting in a win-win scenario. On the one hand, the operator significantly reduces their initial investment, as CAPEX is mainly associated with the deployment of the access network, in exchange for a small increase in the OPEX caused by the renting of the infrastructure. On the other hand, the users gain access to telecommunications services, and get some income from the operator, which can be used for maintaining and improving their network. Although this kind of win-win situation could happen in any country, it is typically found in rural areas of the "Global South" where no universal service regulations are in place. In cases where incumbent operators were reluctant to deploy rural infrastructures because they did not find it profitable to serve small rural communities, communities or their local institutions deployed their own infrastructures, often with public funds or support from development agencies.

One example of this model is the deployment of 3G infrastructure in rural areas where a broadband community network was already in place. In these cases, placing a femtocell in close proximity to the community and sharing the Internet backhaul connection benefits both the users (by obtaining low-cost 3G coverage) and the operator (by avoiding the costs of deploying new infrastructure). Real use cases have been described in the European Commission FP7 TUCAN3G project, which deployed experimental testbeds in two regions in the Amazon forest in Peru [9]. In these networks, the operator and several rural communities cooperated to provide services through rural networks built up with WiLD links.

Crowdshared Approach: This type of Alternative Network corresponds to a set of WiFi routers whose owners share common interests (e.g., sharing connectivity, resources or peripherals) regardless of their physical location. Crowdshared approaches conform to the following idea. A home router hosts two wireless networks, one for serving the owner, and another for public (shared) access, offering a small fraction of the bandwidth to any user of the service in the immediate area (some examples are described in [10]). A governmental initiative corresponds to the networks

Figure 4. (a) Structure of Guifi.net around Barcelona (Spain); (b) junction box of a guifi.net supernode (Spain), including a Mikrotik wireless router, a Raspberry Pi used as proxy, and the router of the operator connecting to the Internet.

created and managed by city councils (e.g., [11]), which act as virtual network operators (VNOs). Other entities that act as VNOs can be grass root user communities, charities, content operators or smart grid operators.

Similarly, some companies (e.g., FON and Vodafone) also promote the use of WiFi routers with dual access (a dedicated WiFi network for the owner, and a shared one for public access). After having a community of users sharing their routers, the members of this community can share their connection and, in turn, get access to all other community resources. In some cases, the owners of the Internet connection can benefit from the temporary leasing of their equipment to nomadic users that connect to WiFi access points. Some other users outside the community can pay passes to gain network access.

Traditional network operators have a financial incentive to lease out the unused capacity at a lower cost to the VNOs, producing revenues for both the VNOs and the sharers [12]. Thus, an incentive structure is created for all actors: end users get money for sharing their network, and network operators are paid by the VNOs, who in turn accomplish their socio-environmental role. Some mainstream operators ship their routers with pre-installed crowdsharing functionality to ease the community formation process.

Emerging Alternative Networks

In addition to the aforementioned classified types of networks, Alternative Networks can also emerge as side-effects of other activities. Some networks that were started by academic entities as research testbeds [13] resulted in non-centralized networks partly governed by regional entities [14].

In a similar way, some rural electric cooperatives have ended up providing broadband access to their users through fiber [15]. These cooperatives started in the 1930s with the aim of providing electric power to the dwellers of remote farms in some zones of the United States. Nowadays, the problem is quite similar, but related to connectivity instead of electricity: investors may be reluctant to deploy an infrastructure to serve a limited number of users. Certain cooperatives installed fiber for running smart grid applications, but later noticed that the same fiber can be used to connect their customers to the Internet.

More recently, the challenge of Internet access provisioning for remote areas has proved fertile ground for innovation. A decade ago, research on delay-tolerant networking led to the creation of DakNet, a network that provides Internet connectivity in a delay-tolerant fashion using buses as mechanical backhaul. Along the same lines, low altitude satellites, drones and balloons are nowadays being considered as means to provide Internet access to remote areas, but these solutions are still at the research level and have not yet been deployed in a real functional Alternative Network.

Conclusions

This article has presented a global classification and a summary of the main characteristics of Alternative Network deployments, which have arisen with the aim of getting more people connected to the Internet. In particular, we have identified five classification criteria and proposed a classification of Alternative Networks into four distinct types. For each type, we detail the main characteristics, describe the technologies they rely on and present real-life examples. To the best of our knowledge, this is the first time a classification of non-traditional network deployments has been proposed. It has been elaborated within the Global Access to the Internet for All Research Group of the IRTF. Its objective is to act as a guide for researchers and community members interested in alternative deployments, and it can help them identify common characteristics of these networks.

Acknowledgments

We would like to acknowledge the contributions of those who have participated in the discussions in the GAIA IRTF list, providing ideas based on their experience and knowledge about Alternative Networks. Arjuna Sathiaseelan and Andrés Arcia-Moret were funded by the EU H2020 RIFE proj-

ect (Grant Agreement no: 644663). Jose Saldana was funded by the EU H2020 Wi-5 project (Grant Agreement no: 644262). Javier Simó-Reigadas was funded by the EU FP7 Research Project TUCAN3G IST-601102 STP and the TEC2013-41604-R Project, funded by the Spanish Ministry of Economy, Industry and Competitiveness.

References

[1] International Telecommunications Union, "ICT Facts and Figures 2016," available: http://www.itu.int/en/ITU-D/Statistics/Documents/facts/ICTFactsFigures2016.pdf, accessed 22 Dec. 2016.
[2] Alliance for Affordable Internet, "Affordability Report 2015-2016," available: http://a4ai.org/affordability-report/, accessed 22 Dec. 2016.
[3] J. Saldana *et al.*, RFC 7962, "Alternative Network Deployments, Taxonomy, characterization, technologies and architectures," Working Group Document in the IRTF GAIA (Global Access to the Internet for All) group, Aug. 2016; available: https://www.rfc-editor.org/info/rfc7962.
[4] E. Pietrosemoli, M. Zennaro, and C. Fonda, "Low Cost Carrier Independent Telecommunications Infrastructure," *Proc. 4th Global Information Infrastructure and Networking Symposium*, Choroni, Venezuela, 2012.
[5] B. Bernardi, P. Buneman, and M. Marina, "Tegola Tiered Mesh Network Testbed in Rural Scotland," *Proc. ACM Workshop on Wireless Networks and Systems for Developing Regions (WiNS-DR '08)*, ACM, New York, NY, USA, 2008, pp. 9–16.
[6] L. Cerda-Alabern, "On the Topology Characterization of Guifi.net," *Proc. IEEE 8th Int'l. Conf. Wireless and Mobile Computing, Networking and Communications (WiMob)*, 2012, pp. 389–96.
[7] K. Heimerl *et al.*, "The Village Base Station," *ICTD 2013*, Cape Town, South Africa, 2013.
[8] J. Simo-Reigadas *et al.*, "Modeling and Optimizing IEEE 802.11 DCF for Long-Distance Links," *IEEE Trans. Mobile Computing*, vol. 9, no. 6, 2010, pp. 881–96.
[9] J. Simo-Reigadas *et al.*, "Sharing Low-Cost Wireless Infrastructures with Telecommunications Operators to bring 3G Services to Rural Communities," *Computer Networks*, Elsevier, 2015.
[10] A. Sathiaseelan *et al.*, "A Feasibility Study of an In-the-Wild Experimental Public Access WiFi Network," *Proc. Fifth ACM Symposium on Computing for Development (ACM DEV 5)*, San Jose, 2014, pp. 33–42.
[11] T. Heer *et al.*, "Collaborative Municipal Wi-Fi Networks — Challenges and Opportunities," *Proc. 8th IEEE Int'l. Conf. Pervasive Computing Communications Workshops (PERCOM Workshops)*, 2010, pp. 588–93.
[12] A. Sathiaseelan and A. J. Crowcroft, "LCD-Net: Lowest Cost Denominator Networking," *ACM SIGCOMM Computer Communication Review*, 2013.
[13] V. Samanta *et al.*, "Metropolitan Wi-Fi Research Network at the Los Angeles State Historic Park," *J. Community Informatics*, North America, 2008.
[14] C. Rey-Moreno *et al.*, "A Telemedicine WiFi Network Optimized for Long Distances in the Amazonian Jungle of Peru," *Proc. 3rd Extreme Conf. Communication: The Amazon Expedition, ExtremeCom '11 ACM*, 2011.
[15] C. Mitchell, "Broadband at the Speed of Light: How Three Communities Built Next Generation Networks," Washington, DC: Institute for Local Self-Reliance and the Benton Foundation, 2012.

Biographies

Jose Saldana (jsaldana@unizar.es) received his B.S. and M.S. in telecommunications engineering in 1998 and 2008, respectively. He received his Ph.D. degree in information technologies in 2011 from the University of Zaragoza, where he is currently a research fellow in the Department of Engineering and Communications. His research interests focus on quality of service in real-time multimedia services, VoIP and networked online games, traffic optimization and resource management in wireless LANs.

Andrés Arcia-Moret (andres.arcia@cl.cam.ac.uk) has been an associate professor in computer systems engineering at the University of Los Andes, Mérida, Venezuela since 2002. He obtained his undergraduate degree (with honors) and his Master's in computer science at the University of Los Andes. He also holds a Ph.D. in computer engineering from the Institut Mines-Telecom (Telecom Bretagne), France. He has been a guest researcher at the IRISA/CNRS in Rennes (France), in the Guglielmo Marconi Laboratory at the International Centre for Theoretical Physics in Trieste, Italy, and in the Computer Laboratory of the University of Cambridge, UK. His research is aimed at wireless networks, alternative network deployments and systems design.

Arjuna Sathiaseelan (arjuna.sathiaseelan@cl.cam.ac.uk) is a senior research associate at the Computer Laboratory, University of Cambridge. He leads the Networking for Development (N4D Lab). The research group conducts research on novel Internet architectures for improving and reducing the cost of Internet access. He is the Chair of IRTF Global Access to the Internet for All (GAIA) research group and a member of the Internet Research Steering Group (IRSG). He is on the Access Advisory Panel of the United Nations Foundation's $75 million Digital Impact Alliance (funded by the Melinda and Bill Gates Foundation, USAID and SIDA). He is also on the advisory board of Ubuntu Power, a social enterprise focused on providing affordable off-grid energy and Internet to under-served communities; and Ensemble, a social business incubator in the Democratic Republic of Congo. He is also on the advisory board of the EU NETCOMMOMS project.

Bart Braem (bart.braem@imec.be) received his Master's degree in computer science at the University of Antwerp (magna cum laude). In September 2005, he joined the IDLab research group at the University of Antwerp, where he defended his Ph.D. thesis on wireless body area networks. Currently a senior researcher at imec in the IDLab research group, he is continuing research on complex, chaotic networks while working on European and regional projects for community networks and smart cities.

Ermanno Pietrosemoli (ermanno@ictp.it) was a professor of telecommunications at the Universidad de los Andes in Venezuela for 30 years. He was one of the founders of Escuela Latinoamericana de Redes (EsLaRed), an organization that has been promoting ICT in Latin America since 1992, and is currently its President. The Internet Society recognized Eslared's efforts with the 2008 Jonathan Postel award. After retirement, Ermanno was a consultant in several organizations, and since 2010 he has been a full-time member of the Telecommunications/ICT4D Laboratory at the International Centre for Theoretical Physics (ICTP) in Trieste, Italy. He has taught courses in wireless data communications and done deployments in many countries, his research interest being mainly in affordable telecommunications systems. He has published several papers and is one of the authors of the book *Wireless Networks in the Developing World*. Ermanno obtained his MSc degree from Stanford University and his EE from the Universidad de los Andes.

Marco Zennaro (mzennaro@ictp.it) is a researcher at the Abdus Salam International Centre for Theoretical Physics in Trieste, Italy, where he coordinates the wireless group of the Telecommunications/ICT4D Laboratory. He received his Ph.D. from the KTH-Royal Institute of Technology, Stockholm, and his M.Sc. degree in electronic engineering from the University of Trieste. He is a visiting professor at the KIC-Kobe Institute of Computing, Japan. His research interest is in ICT4D, the use of ICT for development, and in particular he investigates the use of IoT in developing countries. He has given lectures on wireless technologies in more than 30 countries.

Javier Simó-Reigadas (javier.simo@urjc.es) received the telecommunications engineering degree and the Ph.D. degree from Polytechnic University of Madrid, Spain, in 1997 and 2007, respectively. He was a researcher with the EHAS Foundation between 2003 and 2005 in the field of rural broadband networks for developing countries. Since 2005, he has been an associate professor with the Department of Signal Theory and Communications with Rey Juan Carlos University. His main fields of research are broadband wireless in rural networks supporting multimedia services.

Ioannis Komnios (ikomnios@ee.duth.gr) holds a Ph.D. in computer networks from the Democritus University of Thrace, Greece, and a M.Sc. and five-year Diploma in electrical and computer engineering from the same department. Since 2008, Ioannis has been work package leader and technical leader in several H2020, FP7 and ESA-funded projects, mainly in internetworking in information-centric and delay/disruptive tolerant environments. In 2016, Ioannis joined EXUS Software Ltd. in London, UK, as a senior research consultant.

Carlos Rey-Moreno (crey-moreno@uwc.ac.za) is a post-doctoral fellow in the Computer Science Department at the University of the Western Cape. Fron 2007 to 2011, he was a researcher at the EHAS Foundation working on rural broadband telemedicine networks in Spain, Peru and Malawi. Since 2012, he has been at UWC, where he has been instrumental in the co-creation of Zenzeleni Networks–Mankosi. In studying how to scale Zenzeleni Networks, he has become one of the most knowledgeable people about the community networks in Africa.

Alternative Networks can also emerge as side-effects of other activities. Some networks that were started by academic entities as research testbeds resulted in non-centralized networks partly governed by regional entities. In a similar way, some rural electric cooperatives have ended up providing broadband access to their users through fiber.

ACCEPTED FROM OPEN CALL

M2M Communications in 5G: State-of-the-Art Architecture, Recent Advances, and Research Challenges

Yasir Mehmood, Noman Haider, Muhammad Imran, Andreas Timm-Giel, and Mohsen Guizani

The authors outline the state-of-the-art architecture, recent advances, and open research challenges in communication technologies for M2M. Also, an overview of considerable architectural enhancements and novel techniques expected in 5G networks is presented followed by the resultant services and benefits for M2M communications.

ABSTRACT

M2M communication offers ubiquitous applications and is one of the leading facilitators of the Internet of Things paradigm. Unlike human-to-human communication, distinct features of M2M traffic necessitates specialized and interoperable communication technologies. However, most existing solutions offering wired or wireless connectivity have limitations that hinder widespread horizons of M2M applications. To cope with the peculiar nature of M2M traffic, the evolving 5G system considers the integration of key enabling networking technologies for ubiquitous connectivity and guaranteed QoS. This study outlines the state-of-the-art architecture, recent advances, and open research challenges in communication technologies for M2M. Also, an overview of considerable architectural enhancements and novel techniques expected in 5G networks is presented, followed by the resultant services and benefits for M2M communications.

INTRODUCTION

The exponential growth in new machine-to-machine (M2M) and Internet of Things (IoT) applications, autonomous technologies, and spectrum demand challenges conventional network capabilities in capacity and data rates. Future forecasts suggest a dramatic increase in mobile Internet use, which in turn requires more spectrum. The standardization bodies, telecom industry, and academia have initiated efforts for achieving the fifth generation (5G) milestone[1]. Moreover, M2M and IoT communications offering multi-fold services are potentially considered very important to the future success of 5G networks. The M2M-driven economic growth in the future depends on advancements in conventional communication systems to effectively relay the massive M2M traffic.

Existing communication technologies employ both licensed (e.g., cellular networks) and unlicensed (e.g., WiFi) spectrum for data transmission. These technologies are not capable of offering sufficient radio resources for massive M2M traffic. Moreover, radio resource allocation is becoming challenging to support traditional mobile traffic in the presence of massive M2M devices. In the past, notable research has been carried out to design and propose efficient resource allocation and utilization techniques. The key 5G enabling technologies to meet future requirements, such as massive multiple-input multiple-output (MIMO) systems, small cells, millimeter-wave (mmWave), and cognitive radios, will be used to considerably elevate capacity, resource, speed, and reliability constraints. Native support for M2M communications is considered as one of the five main disruptive technologies of 5G [2]. Considering enough enhancements in cellular and WiFi technologies exclusively, there is a very limited possibility of a further breakthrough in these technologies. However, the convergence of cellular and WiFi seems to be a promising solution to overcome shortcomings of each technology and integrate advantages of both. Such innovative integration would yield positive results for M2M devices to successfully integrate into the future communication system.

Upcoming 5G networks are particularly underlined to enrich future cellular M2M communications with increased throughput, reduced end-to-end (E2E) delay, wide coverage, increased battery operating time, and support for an enormous number of devices per cell. Therefore, several standardizations bodies such as the Third Generation Partnership Project (3GPP), European Telecommunications Standards Institute (ETSI), and IEEE are working to develop cutting-edge technologies for this new generation of M2M/IoT communications. Moreover, efforts from various industries and standardization bodies are still going to decide and bring forth adequate solutions to fully cater to the performance demands of M2M/IoT.

The remainder of this article is organized as follows. A generic overview of M2M architecture, major ongoing standardization activities, and notable projects are given in the following section. Then we outline 5G architecture and its expected enhancements for supporting M2M communications. M2M applications and service provisioning in the context of 5G are then described. Then we highlight open research challenges in accommodating M2M traffic in a 5G network. Finally, we draw conclusions.

Yasir Mehmood is with the University of Bremen; Noman Haider is with the University of Technology Sydney; Muhammad Imran is with King Saud University; Andreas Timm-Giel is with the University of Technology Hamburg; Mohsen Guizani is with the University of Idaho.; The corresponding authors are Yasir Mehmood and Mohsen Guizani.

Digital Object Identifier: 10.1109/MCOM.2017.1600559

M2M Architecture and Standards

To ensure global connectivity and interoperability among M2M devices, actuators, and other computational elements, the main standardization bodies are actively involved in developing new standards, protocols, and open interfaces. This section highlights a generic overview of widely used ETSI M2M architecture followed by the key ongoing standardization activities for promoting M2M communications [3].

ETSI M2M Architecture

ETSI proposed the basic M2M network architecture by defining its service requirements, which have been adopted worldwide for realizing M2M communications. According to ETSI specifications [4], the proposed M2M architecture consists of the following principle components, as shown in Fig. 1.

M2M Device Domain: The fundamental task of an M2M device is to fetch and send data to nearby devices or infrastructure. The collected information includes sensing of internal temperature, humidity level, position and speed of a vehicle, and fuel consumption. The devices are generally connected to local area networks (LANs) for transmitting data to the backend server or combining data through an M2M communication domain. Examples of LANs include power line communication, short-range device (SRD), ultra-wideband (UWB), ZigBee, meter bus (M-BUS), wireless meter bus, Bluetooth, and cellular gateways/relays [5].

M2M Communication Domain: The data collected from M2M devices is forwarded to gateways such as roadside units and cellular base stations, called eNBs. The communication networks enable devices to communicate with the application servers via wired or wireless connectivity. Examples of communication networks include xDSL, SRDs, UWB, WiMAX, Global System for Mobile Communication (GSM), Universal Mobile Telecommunications System (UMTS), LTE/LTE-Advanced, and satellite communications. In addition, the M2M gateway is identical to a bridge, ensuring connectivity and communication between two or more devices and communication networks. The gateways forward the collected data from devices to the back-end servers.

M2M Server/Application Domain: The M2M server/application domain comprises a middleware layer where the received packets from M2M devices pass through several application services and later are used by the related business processing bodies.

M2M Standardization Activities

Several standard bodies such as 3GPP, ETSI, IEEE, and the World Wide Web Consortium (W3C) have started putting efforts individually as well as collectively in supporting massive deployment of M2M applications. This section summarizes the key ongoing standardization activities in the scope of M2M communications.

Individual Efforts: ETSI has put notable efforts in developing E2E network architecture, protocols, and interfaces for supporting M2M communications. 3GPP has actively worked on featuring M2M communications across several releases

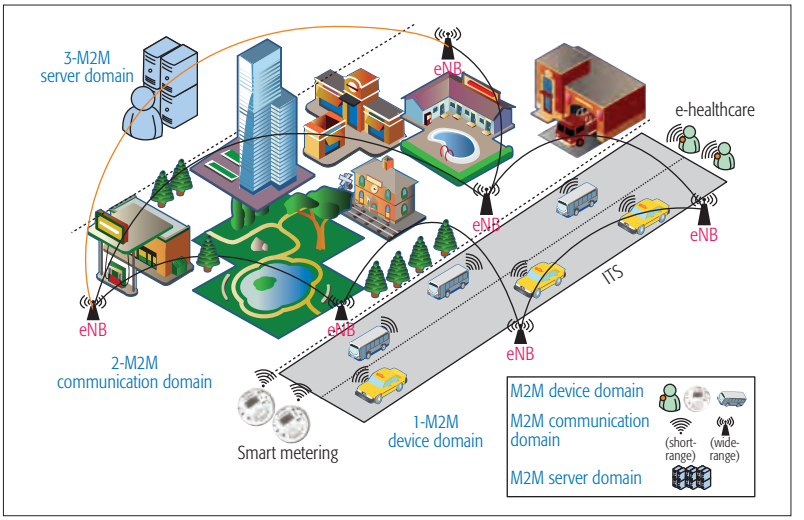

Figure 1. M2M communications architecture proposed by ETSI.

of LTE/LTE-Advanced and contemporary mobile technologies. For instance, 3GPP in its Release 12 addressed problems, such as excessive signaling overhead, and the cost and power consumption of M2M devices, and subsequently proposed several solutions such as power saving modes, and a new low-complexity device category called Cat-0. In addition, based on existing wireless communication technologies, 3GPP (http://www.3gpp.org/news-events/3gpp-news/1785-nbiotcomplete, accessed June 10, 2016) in its Release 13 further standardizes narrowband IoT (NB-IoT) to provide better network coverage for M2M/IoT by further reducing the bandwidth to 200 kHz, and supporting physical resource block level throughput, more devices, low energy consumption, and coverage extension by 20 dB [6].

Collaborative Efforts: M2M communications demand interoperable architectures and platforms that enable interworking of multiple standards. For instance, interoperability, scalability, and connectivity to massive M2M devices are crucial requirements in the future. To address these issues, 3GPP in cooperation with ETSI and other renowned telecom standardization bodies are focusing on supporting M2M communications in upcoming 5G networks. As a result, enhancements in E2E network architecture and extension of radio spectrum are under consideration and are discussed later. The, W3C Web of Things (WoT) is also actively involved in standardizing interoperable platforms for M2M/IoT communications. W3C is working in collaboration with other M2M/IoT standards bodies and alliances worldwide to develop application program interfaces (APIs) independent of platforms. Additionally, ForestGreen, a W3C Interest Group, offers a forum for discussing various technical aspects and open markets of IoT and W3C-based applications. Besides, a global initiative, the oneM2M partnership project, was launched by ETSI in 2012 along with eight of the world's well-known standardization bodies, six global forums and standard development organizations, and over 200 companies from all industrial sectors.

EU Sponsored Projects: Besides the individual and collaborative efforts discussed in previous section, several European Union (EU) sponsored

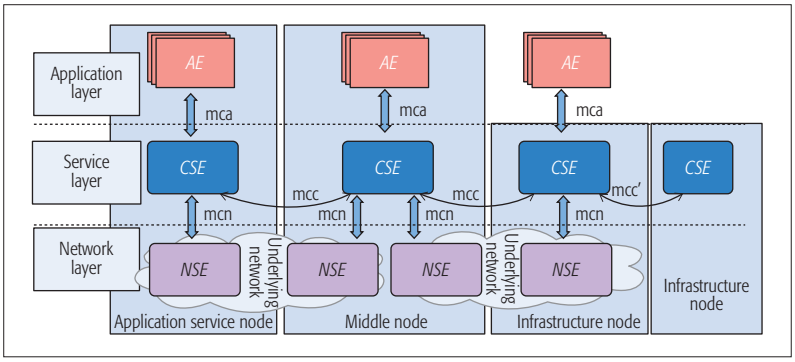

Figure 2. oneM2M functional architecture.

projects have also addressed M2M/IoT related issues. For instance, the METIS (2012-2015) was sponsored by EU, extensively addressed several M2M related issues. As a result, significant enhancements were proposed to support multiple random access and control procedures, low signaling overhead, radio resource efficiency using user multiplexing, as well as low power consumption with device sleep mode. Moreover, EXALTED (2010-2013) was another EU sponsored project that aimed to deliver secure, cost- and energy-efficient M2M communications. Consequently, EXALTED laid the foundation of a new scalable network architecture under the name of LTE-M networks, particularly in mesh networks. Approximately 1000 devices were supported per cell with the proposed architecture and schemes by assuming latency requirements greater than 100 ms. Furthermore, LOLA proposed efficient spectrum utilization techniques by targeting a particular set of applications with low latency requirements such as gaming. LOLA also delivered inputs to 3GPP directly by designing schedulers, framing, adaptive modulation and coding schemes, as well as Acknowledgment feedback.

Furthermore, several ongoing Fifth Generation Public Private Partnership (5G-PPP) projects sponsored by EU are also addressing M2M issues. For instance, METIS-II (https://metis-ii.5g-ppp.eu/about-metis/?doingwp-cron=1485738639.1112298965454101562500, accessed January 30, 2017) is aimed at the designing of holistic spectrum management architecture and air interface framework, agile resource management techniques, and cross-layer and cross-interface access methods. As a result, strict quality of service (QoS) and resource requirements of M2M will be addressed. Since 5G networks are expected to support trillions of devices, CogNet (https://5g-ppp.eu/cognet/, accessed January 30, 2017), a 5G-PPP sponsored project, aims to introduce machine learning for designing autonomic network management to ensure a self-administering and self-managing network. Moreover, FANTASTIC (http://fantastic5g.eu/, accessed January 30, 2017) aims to support massive M2M traffic. Therefore, its major objective is developing a scalable multi-service air interface with high capacity and wide coverage. Moreover, the air interface will be highly efficient in the context of radio resource utilization and power consumption. Additionally, the 5G-PPP SESAME project (https://5g-ppp.eu/sesame/, accessed February 01, 2017) proposes a novel small cell concept by introducing multi-operator capabilities through integrating a virtualized execution platform. This will significantly help to support M2M communication on a large scale, and ensure the interoperability among M2M devices with reduced power consumption.

oneM2M Architecture Reference Model

oneM2M is the global initiative[7] that covers architecture, requirements, API specifications, privacy and security solutions, as well as an interoperability framework for M2M/IoT technologies. Unlike ETSI M2M architecture, the principle objective of oneM2M is to deliver a generic service layer framework for M2M communications. oneM2M architecture is based on ETSI M2M architecture along with the additional service layer for applications to share common infrastructure, environments, and network elements in order to ensure interoperability among M2M/IoT devices. Furthermore, oneM2M aims to deliver open interfaces that ensure a wide range applicability of the entire M2M/IoT ecosystem. As a result, the intended design of the framework will be applied to several hardware- and software-based networks worldwide. Figure 2 illustrates the proposed oneM2M functional architecture, which consists of three key components: the application entity (AE), the common service entity (CSE), and the network service entity (NSE); we briefly discuss them below.

Application Entity: The AE resides within the application layer and implements an application layer service logic. Each service logic can reside in several M2M devices and/or several times on a single device. The AE identity (AE-ID) represents an execution entity of each application layer service logic. The major AE examples include an instance of remote monitoring of temperature, pulse rate, and blood sugar in e-healthcare scenarios, smart metering in smart city/smart living applications, and tracking of fleet management in modern transportation and logistics. Furthermore, the reference point mca allows the communication flow between the AE and the CSE. Using the mcn reference point, the AE uses the services offered by the CSE, and gives an opportunity for the latter to communicate with the former. Furthermore, both the AE and the CSE may or may not be placed inside the same physical entity.

Common Service Entity: In oneM2M architecture, an instantiation of the set of common service functions illustrates a CSE in M2M communication environments. The service functions are unveiled to other entities within the oneM2M architecture using several reference points such as mcs and mcc. Like the AE, each CSE is also identified with an identity called the CSE-ID. Moreover, mcn reference points are used to access the underlying NSE. The CSE service functions include device and data management, subscription management of M2M service, as well as location-based services. Reference point mcc is responsible for communication flows between several CSEs. Thus, mcc enables a CSE to interact and use services offered by another CSE (Fig. 2).

Network Service Entity: The NSE offers services to CSEs from the underlying network. Examples include device management, location services, and device triggering operations. Ref-

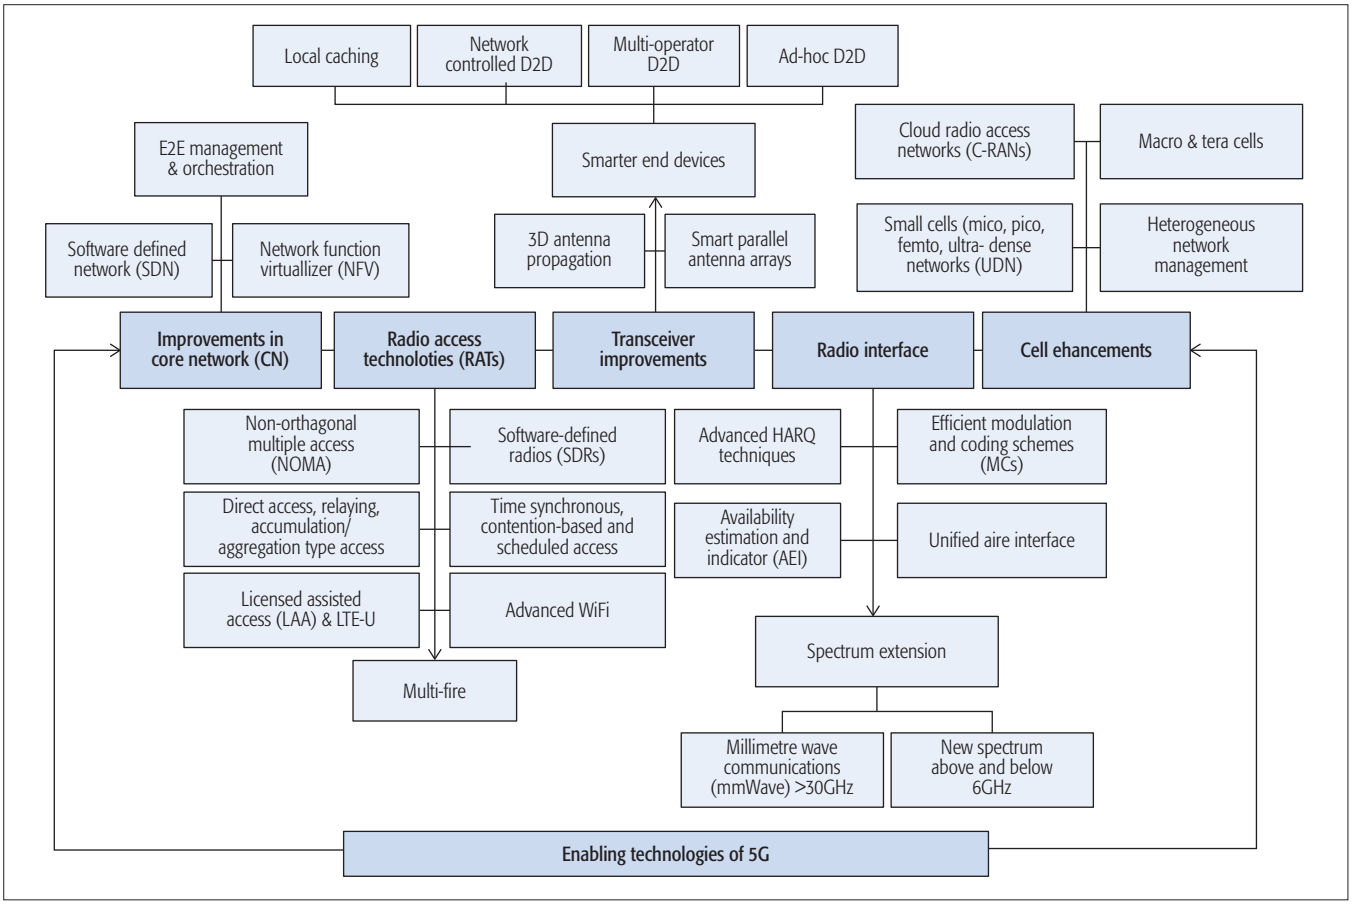

Figure 3. Enabling 5G technologies.

erence point mcn offers communication opportunities between CSEs and NSEs. Additionally, mcn enables CSEs to use services offered by NSEs except transport and connectivity services.

5G Architecture and Capabilities

The architecture and capabilities of next generation cellular and wireless networks are driven by the demands and requirements of human-to-human and M2M traffic.

Architectural Overview

5G networks are expected to provide support for all types of communication with programmable system protocols that can be tuned according to user requirements. Unlike the previous generation of networks, where control and processing tasks are heavily assigned to the infrastructure side; 5G aims to balance this factor by employing architectural changes from cell-centric to device-centric design [2]. The reasons for this architectural shift include the variety of communication methods like multiple radio access technologies (multi-RATs), half-duplex/full-duplex mode, above/below 6 GHz spectrum option, and 3D MIMO.

The fundamental design aspect of 5G is to bring forth a unified solution in terms of both hardware and software for end users and network operators, appearing as one transparent system integrated with legacy and novel technological components providing a seamless user experience. Figure 3 manifests legacy and new technological advancements that are expected to be an integral part of 5G.

Major Enhancements in 5G

The challenging design objective of 5G is to provide connectivity to an enormous number of devices with diverse characteristics and application requirements. For example, mission-critical M2M applications require minimum latency, while applications like smart metering are delay-tolerant. Such a broad range of applications with diverse QoS requirements leads to significant enhancements in 5G [8]. The following major enhancements are foreseen in upcoming 5G networks, which bring compatible and integrated support for M2M communications.

Network Virtualization: The virtualization of network resources has proven its significance in the context of operational efficiency, resource management, and cost reduction for back-end management in routing and switching. Due to a great deal of versatility and dynamic resource management requirements, network virtualization is a good candidate for integration of key enabling technologies in 5G while minimizing the equipment cost and reducing the implementation complexity. A feasible proposal is to implement 5G network functions as customizable software components using network functions virtualization (NFV). This, in turn, helps to implement future advancements by updating the only API of related software components. The virtualization concept at the core and fronthaul level is a major technology enabler for support of resource and energy constrained massive M2M devices. The idea is to empower the backhaul system while reducing the computation and processing load at the end

5G has native ubiquitous connectivity for massive M2M communications and IoT. Technologies like cell enhancements, centralized resource allocation, and new spectrum extends the potential of 5G for M2M. For instance, D2D communication will be a major communication trend especially in wearable devices and smart sensor environments for relaying.

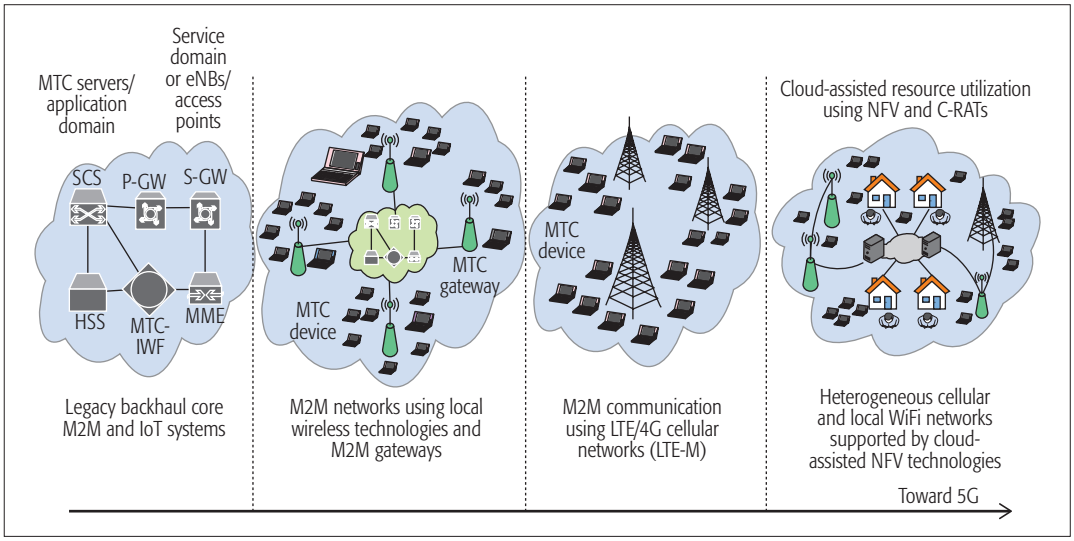

Figure 4. Evolution of MTC over different types of networks toward 5G.

device due to requirements of small-sized, energy-constrained, and dedicated IoT devices.

Cell Enhancements: Cell enhancements will play a pivotal role in serving ultra-dense networks (UDNs) where a plethora of devices will be operating within a single cell. For instance, phantom, micro, pico, and femtocells are used as underlay networks for data transmission or reporting by M2M devices along with macrocells. Nevertheless, an interesting concept of phantom cells for segregation of control and data planes has also shown potential performance improvements. The key idea is to use macrocells for control signaling over microwave or cellular frequencies, while microcells will be used for data services operating on high frequencies. This fusion of high and low frequencies will bring forth a great improvement in terms of spectrum efficiency, network capacity, and coverage for M2M communication.

Extension of Radio Spectrum: One of the prominent features of 5G is employing newly available spectrum in licensed and unlicensed bands. Moreover, end devices are now able to utilize inter-band and intra-band frequencies from different bandwidths using the carrier aggregation framework proposed in 3GPP Release 10. 3GPP started a study item in 2014 for LTE deployment in unlicensed bands with a major challenge of fair and peaceful coexistence with incumbents operating in the same band like WiFi, IEEE 802.11x. This research introduced licensed-assisted access (LAA), which requires control signals from traditional licensed bands in macrocells to use unlicensed bands for data services in small micro/picocells. In addition, coordinated multi-point among multiple base stations and direct device-to-device (D2D) communication opened up another major research area in 5G focusing on advanced interference mitigation and avoidance techniques. These techniques consider optimized cell association and power control methods for multi-tier networks, paving the road for integration of M2M systems.

Figure 3 contains highlights of some of the important enhancements being considered in 5G. Among other design considerations, IoT, M2M, and cyber-physical systems are very crucial stakeholders of this system. In the next section, we focus on different aspects of 5G to serve billions of M2M/IoT devices.

M2M Communication in 5G Networks

5G has native ubiquitous connectivity for massive M2M communications and IoT. Technologies like cell enhancements, centralized resource allocation, and new spectrum extend the potential of 5G for M2M. For instance, D2D communication will be a major communication trend, especially in wearable devices and smart sensor environments for relaying.

In 3GPP, design targets (overload control, improved coverage, device enhancements for longer battery life and efficient power consumption) for M2M are evaluated using the GSM EDGE radio access network (GERAN) as a baseline technology candidate. Few M2M applications like smart metering are already using cellular networks such as GSM, general packet radio service (GPRS), and narrowband IoT (NB-IoT). Currently, study items in 3GPP are focusing on LTE radio network enhancements for machine-type communication (MTC)[1] [9]. Important design aspects at the access end, and front- and backhaul have been studied and finalized in recent releases of 3GPP.

IEEE standards committees (802 LAN and MAN) explore requirements and challenges of M2M communications in terms of RAN and capacity. For instance, IEEE 802.15.x, with its continuous timely upgrading, serves various applications in personal area networks. Another recently introduced standard, IEEE 802.11ah, supports low-power transmissions with wider coverage for M2M communications. The IEEE 802.16p Task Group (TG) targets elevation of basic WiMAX standards of IEEE 802.16e and IEEE 802.16m for M2M, especially medium access control (MAC) related requirements of channel access, grouping, and device addressing. Some other TGs are also analyzing technologies that have ZigBee as the base standard. However, the main challenge is the lack of efficient backhaul, which limits the network scalability and coverage. The strong 5G backhaul support is expected to overcome this issue and provide optimum support using virtualization techniques.

[1] The term MTC is identical to M2M communication.

Currently, the majority of M2M devices in short range are served by capillary networks (WiFi, Bluetooth, non-3GPP). However, in the future both cellular and capillary networks will provide integrated, seamless, and ubiquitous handover and offloading for these devices. Also, in 5G, conventional mobile broadband (MBB) is expected to be replaced with extreme MBB, and the use cases under consideration include massive, mission-critical, and ultra-reliable MTC. Figure 4 shows an evolution of MTC communication over cellular, wireless, and integrated networks and how 5G is expected to combine all these technologies in a single platform supported by cloud technologies.

Services for Future M2M Communication

The first and foremost service for future M2M is connecting thousands of devices per cell with different QoS requirements and effective group management. Joint resource allocation and concurrent access for massive and M2M devices are well supported in 5G via NFV and cloud-RAN [10]. Control and data channel segregation by adaptive and software-driven core networks provide further improvement in M2M system efficiency. Priority-based ultra-reliable communication for mission-critical services is the design consideration of 5G. Due to UDNs, local caching and effective computation and communication of important data is a very challenging task, which is expected to be catered for in future M2M.

The emerging D2D communication allows devices to exchange data with each other directly based on proximity, hence saving scarce network resources. Devices can also help relay data or information to a nearby sink using this technology. Vehicle-to-vehicle communication is a major user of this technology. Moreover, mission-critical applications in hostile environments can use this technology to locate survivors and assist rescue teams to move in the right directions. Such M2M and IoT applications in the future will be injected with self-healing and fault-tolerant survivability mechanisms currently under development in 5G systems.

Strengths of 5G over Traditional Technologies

Due to compatibility issues and the bottleneck at backhaul networks, legacy technologies are unable to provide optimum services. 5G systems not only present a unified solution by compatible integration of enabling technologies, but also appear to end users as one system with consistent but adaptable features. This adaptability best suits M2M communications with a diverse set of requirements in terms of latency, data rate, and traffic characteristics. For instance, some applications like smart metering require network resources to send data packets (uplink) periodically. Thus, more resources in the uplink are required compared to the massive downlink traffic consumed by another type of application.

Ubiquitous and reliable connectivity and guaranteed QoS for different M2M services are some of the technological advantages of 5G over contemporary systems. Previous communication systems typically have dedicated hardware, whereas 5G, using software-defined networking (SDN) and NFV, will be more software oriented as configuration of network parameters can easily be performed at the application layer using application programming interfaces (APIs).

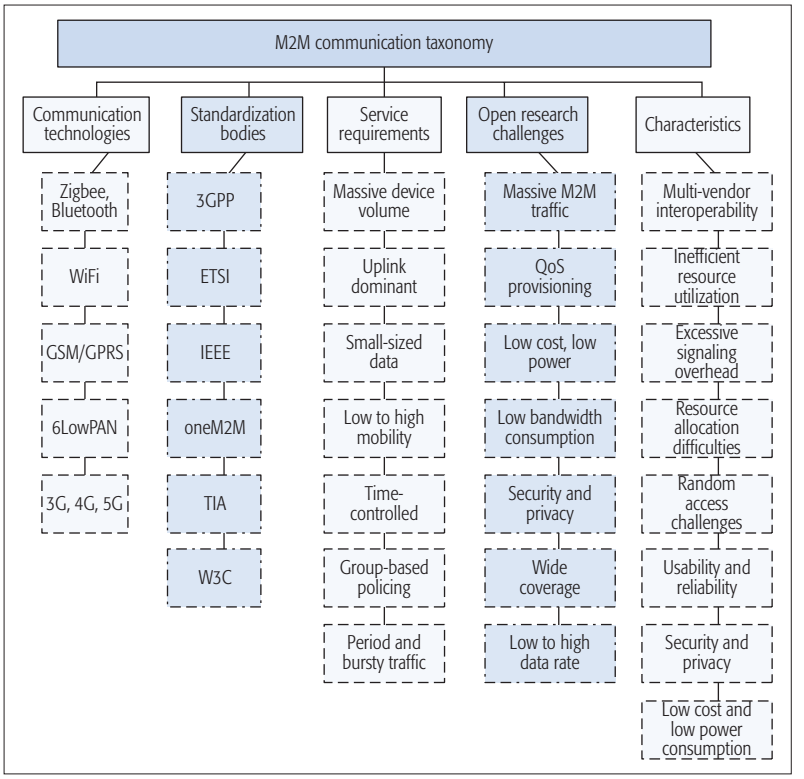

Figure 5. An overview of M2M communication technologies, service requirements, and open research challenges.

Open Research Challenges

Existing MTC related problems require either amendments in the legacy networks or designing a new network architecture. For the upcoming 3GPP Releases 13 and 14, operators and vendors are submitting proposals for a new architecture. Figure 5 illustrates a generic overview of an M2M taxonomy. This section particularly highlights key research challenges of M2M communication in 5G.

Multi-Vendor Interoperability: Multi-vendor interoperability is one of the major requirements for future M2M and IoT applications. Until an interoperability framework is established, M2M/IoT will not be able to acquire the required scale and flexibility it needs to fulfill numerous applications.

Inefficient Radio Spectrum Utilization: In mobile communications, existing standards such as GSM, UMTS, and LTE transport block sizes are essentially optimized for typical VoIP traffic. Since radio spectrum is a valuable asset and scarcely available, allocating one resource block to a single device may significantly degrade its utilization. Therefore, current telecommunication standards should be modified to deal with the future M2M/IoT traffic. Data concatenation, aggregation, and compression can be an important factor as the payload size from an individual node is usually small, hence needing fewer spectrum resources.

Excessive Signaling Overhead: In mobile and wireless communications, M2M nodes share the limited radio spectrum along with traditional cellular traffic in order to exchange their control and

Parameters	Release 8 (cat-4)	Release 8 (cat-1)	Release 12 (cat-0)	Release 13
MTC device bandwidth	20 MHz	20 MHz	20 MHz	1.4 MHz
Peak data rate-downlink	150 Mb/s	10 Mb/s	1 Mb/s	200 kb/s with a TBS of 1000 bits unicast traffic
Peak data rate-uplink	50 Mb/s	5 Mb/s	1 Mb/s	200 kb/s with a TBS of 1000 bits unicast traffic
Duplex mode	Full duplex	Full duplex	Half duplex (opt.)	Half duplex (opt.)
Max. device Tx power	23 dBm	23 dBm	23 dBm	20 dBm
Modem complexity compared to Cat-1	125%	100%	50%	25%

Table 1. A comparison of MTC device power consumption in 3GPP release 8, release 12 with respect to release 13.

data information. However, one of the biggest issues in handling the M2M traffic is the signaling overhead due to the expected and massive rise in number of devices per cell. Thus, it is more likely that the performance of the contemporary networks will be significantly degraded. However, handling interference would be a challenging task. Centralized coordination would be effective in order to limit interference at the cost of increased signaling overhead. However, distributed resource management using cognitive M2M communication could be an effective approach in order to limit interference [11]. In addition, another approach can be used in which the small-sized M2M/IoT data packets can be sent on signaling bearers along with the non-access-stratum (NAS) messages.

Resource Allocation Difficulties: Recently, interference is the key issue of the diversified signaling patterns. Hence, a radio spectrum partitioning strategy is required to overcome the aforementioned problem. This partitioning allocates resources to various devices with numerous constraints on available resources or transmit power. Therefore, the above stated approach can improve the signal-to-interference-plus-noise ratio and energy efficiency of M2M devices.

Random Access Challenges: Since mobile systems are admired for delivering satisfactory services due to long range, increased data rate, high mobility support, and security, devices may access the network simultaneously and thus contend for radio resources [12]. This results in low random access success rate and high network congestion due to the limited number of preambles. It further results in more power consumption, higher packet loss, unpredictable delays, and excessive utilization of radio resources. Additionally, when a huge number of devices try to reconnect, the physical random access channel (PRACH) will be further overloaded. To cater to the above issues, existing solutions include optimized MAC operations, access class barring schemes, separation of radio resources, dynamic allocation of the RACH, code expanded random access (RA), prioritized RA, and backoff adjustment schemes [13]. Moreover, the authors in [14] proposed a scheme that prioritizes control signals based on the combination of control messages such as RA control messages, uplink grants, and downlink assignments. However, exponentially increasing device numbers demand further investigation of techniques that can potentially support M2M/IoT applications.

Usability and Reliability from the Consumer Perspective: Depending on different use case scenarios and requirements, providing software-driven customized network services is a challenging task for future networks. For instance, IoT devices and corresponding data for elder care and fall detection are supposed to operate in automated modes with high precision and reliability. Meanwhile, kids in the same environment may need high speed and bandwidth for video gaming and streaming applications. Segregating user data and providing corresponding desired services to each user from a single cost-effective platform need significant research effort and development.

Security and Privacy: A gigantic increase in IoT and M2M devices connected to the Internet also creates more threats to data security, integrity, confidentiality, and privacy [15]. For instance, devices for medical care may contain personal and critical health information that is transported without any encryption and makes user privacy vulnerable to attack. The conventional authentication, authorization, and accounting (AAA) framework needs optimization for IoT devices to introduce legitimate access to devices and data. Thus, M2M networks demand sophisticated filtering procedures, restricted firewalls, and robust intrusion detection and prevention procedures due to the fact that data may belong to confidential business processes, personal health care, and other mission-critical applications.

Low-Cost and Low-Power Devices: In M2M communications, power consumption and cost are more prominent design factors to make cellular networks feasible and applicable compared to traditional wireless technologies. M2M devices cost less and have lower power consumption compared to VoIP and multimedia services for battery-driven applications like sensors and water meters. The key features in upcoming Release 13 are presented in Table 1. Future research directions include transparent business revenue models, operator return on investment, and production costs that are imminent issues that need thorough analysis.

Conclusion

This article presents an overview of M2M architecture and the technological advancements expected in 5G networks for enhancing mobile M2M communication. Based on this study, we conclude that oneM2M functional architecture is adopted widely for the real deployment of cellular-based M2M applications. Furthermore, 5G networks are particularly designed to meet the unique data rate, reliability, and latency requirements of M2M and IoT. Distinct QoS and shared radio spectrum for M2M and IoT are expected to be efficiently utilized using cloud-assisted and NFV supported systems. Moreover, 5G networks are expected to be one of the major facilitators for cellular-M2M communications through low-cost devices (up to $10), reduced power consumption, and improved security procedures.

Acknowledgments

We especially thank the late Prof. Dr. Carmelita Görg, former head of ComNets, University of Bremen for all her support and guidance. Furthermore, we thank the International Graduate School for Dynamics in Logistics (IGS), doctoral training group of LogDynamics, University of Bremen, Germany, for the financial support of this work. Imran's work is supported by the Deanship of Scientific Research at King Saud University through Research group No. (RG # 1435-051).

References

[1] A. Osseiran *et al.*, "The Foundation of the Mobile and Wireless Communications System for 2020 and Beyond," *Proc. IEEE VTC-Spring Wksp.*, Germany, 2013.
[2] F. Boccardi *et al.*, "Five Disruptive Technology Directions for 5G," *IEEE Commun. Mag.*, vol. 52, no. 2, Feb. 2014, pp. 74–80.
[3] Y. Mehmood *et al.*, "Mobile M2M Communication Architectures, Upcoming Challenges, Applications, and Future Directions," *EURASIP J. Wireless Communications and Networking*, no. 1, 2015.
[4] ETSI, "Machine-to-Machine Communications (M2M), Functional Architecture," tech. rep., ETSI TS 102 690 V2.1.1, Oct. 2013.
[5] M. J. Booysen, S. Zeadally, and G.-J. van Rooyen, "Survey of Media Access Control Protocols for Vehicular Ad Hoc Networks," *IET Commun.*, vol. 5, no. 11, 2011, pp. 1619–31.
[6] 3GPP TSG RAN Meeting, "Narrowband IoT (NB-IoT)," tech. rep., Sept. 2015.
[7] oneM2M, "oneM2M Project Release 1 Specifications." http:www.onem2m.org/technical/published-documents, accessed 1 Jan. 2016.
[8] T. G. I. P. P. Partnership, "5G Vision: The Next Generation of Communication Networks and Services," https://5g-ppp.eu/, accessed 30 Dec. 2015.
[9] H. Shariatmadari *et al.*, "Machine-Type Communications: Current Status and Future Perspectives Toward 5G Systems," *IEEE Commun. Mag.*, vol. 53, no. 9, Sept. 2015, pp. 10–17.
[10] C.-Y. Oh, D. Hwang, and T.-J. Lee, "Joint Access Control and Resource Allocation for Concurrent and Massive Access of M2M Devices," *IEEE Trans. Wireless Commun.*, vol. 14, no. 8, 2015, pp. 4182–92.
[11] A. Aijaz and A. H. Aghvami, "Cognitive Machine-to-Machine Communications for Internet-of-Things: A Protocol Stack Perspective," *IEEE Internet of Things J.*, vol. 2, no. 2, 2015, pp. 103–12.
[12] M. Hasan, E. Hossain, and D. Niyato, "Random Access for Machine-to-Machine Communication in LTE-Advanced Networks: Issues and Approaches," *IEEE Commun. Mag.*, vol. 51, no. 6, June 2013, pp. 86–93.
[13] A. Laya, L. Alonso, and J. Alonso-Zarate, "Is the Random Access Channel of LTE and LTE-A Suitable for M2M Communications? A Survey of Alternatives," *IEEE Commun. Surveys & Tutorials*, vol. 16, no. 1, 2014, pp. 4–16.
[14] T. P. De Andrade, C. A. Astudillo, and N. L. da Fonseca, "Allocation of Control Resources for Machine-to-Machine and Human-to-Human Communications over LTE/LTE-A Networks," *IEEE Internet of Things J.*, vol. 3, no. 3, 2016, pp. 366–77.
[15] M. Security, "M2M Security and Privacy Challenges," http:www.prnewswire.com/news-releases/machine-to-machine-m2m-security-and-privacy-challenges-andopportunities- 272940451.html, accessed 19 Jan. 2016.

Biographies

Yasir Mehmood completed his Master's in electrical (telecommunications) engineering from the Military College of Signals (MCS), National University of Science and Technology (NUST) Islamabad, Pakistan. He is currently a doctoral researcher at the Sustainable Communication Networks (ComNets) research group, University of Bremen, Germany, in the framework of the International Graduate School (IGS) for Dynamic in Logistics (a doctoral training group at the University of Bremen). His major research area includes cellular communications, mobile M2M communications, and cellular Internet of Things (C-IoT).

Noman Haider received his B.S. degree in electronics engineering from Mohammad Ali Jinnah University, Islamabad, Pakistan, in 2011 and his M.S. degree in electrical and electronics engineering from Universiti Teknologi Petronas, Malaysia, in 2014. He is currently a Ph.D. student in the School of Computing and Communications at the University of Technology Sydney, Australia. He has worked on different academic and industry sponsored projects in the field of wireless communications and networking. His research interests include modeling, analysis, and design of heterogeneous cellular networks in 5G.

Muhammad Imran has been an assistant professor in the College of Computer and Information Sciences, King Saud University (KSU) since 2011. He worked as a postdoctoral associate on joint research projects between KSU and the University of Sydney, Australia. He is a visiting scientist ay Iowa State University. His research interests include mobile ad hoc and sensor networks, WBANs, M2M, IoT, SDN, fault-tolerant computing, and security and privacy. He has published a number of research papers in refereed international conferences and journals. His research is financially supported by several grants. Recently, the European Alliance for Innovation (EAI) appointed him as a Co-Editor-in-Chief of *EAI Transactions on Pervasive Health and Technology*. He also serves as an Associate Editor for *IEEE Access*, *IEEE Communications Magazine*, the *Wireless Communication and Mobile Computing Journal* (SCIE, Wiley), the *Ad Hoc & Sensor Wireless Networks Journal* (SCIE), *IET Wireless Sensor Systems*, the *International Journal of Autonomous and Adaptive Communication Systems* (Inderscience), and the *International Journal of Information Technology and Electrical Engineering*. He has served/is serving as a Guest Editor of *IEEE Communications Magazine*, *Computer Networks* (SCIE, Elsevier), *MDPI Sensors* (SCIE), the *International Journal of Distributed Sensor Networks* (SCIE, Hindawi), the *Journal of Internet Technology* (SCIE), and the *International Journal of Autonomous and Adaptive Communications Systems*. He has been involved in more than 50 conferences and workshops in various capacities such as Chair, Co-Chair, and Technical Program Committee member. These include IEEE ICC, GLOBECOM, AINA, LCN, IWCMC, IFIP WWIC, and BWCCA. He has received anumber of awards such as an Asia Pacific Advanced Network fellowship.

Andreas Timm-Giel received his diploma and Dr.-Ing. degrees from the University of Bremen in 1994 and 1999, respectively. From 1994 to 1999, he led a group on mobile and satellite communications at the University of Bremen. From 2000 to 2002, he was the project manager and manager of network operations at MediaMobil GmbH and M2sat Ltd. From December 2002 to November 2009, he was a senior researcher, project manager, and lecturer at ComNets/University of Bremen, where he also led the concerted activity Adaptive Communications at TechnologieZentrum Informatik (Center for Computing Technologies) (TZI). Since November 2009, he has been a professor at Hamburg University of Technology heading the Institute of Communication Networks.

Mohsen Guizani [F] received his B.S. (with distinction) and M.S. degrees in electrical engineering, and his M.S. and Ph.D. degrees in computer engineering from Syracuse University, New York, in 1984, 1986, 1987, and 1990, respectively. He is currently a professor and the ECE Department Chair at the University of Idaho. Previously, he served as the associate vice president of Graduate Studies, Qatar University, Chair of the Computer Science Department, Western Michigan University, and Chair of the Computer Science Department, University of West Florida. He also served in academic positions at the University of Missouri-Kansas City, the University of Colorado-Boulder, Syracuse University, and Kuwait University. His research interests include wireless communications and mobile computing, computer networks, mobile cloud computing, security, and smart grid. He currently serves on the Editorial Boards of several international technical journals, and is the founder and Editor-in-Chief of *Wireless Communications and Mobile Computing* (Wiley). He is the author of nine books and more than 450 publications in refereed journals and conferences. He has guest edited a number of special issues in IEEE journals and magazines. He has also served as a TPC member, Chair, and General Chair of a number of international conferences. He has received teaching awards multiple times from different institutions as well as the Best Research Award from three institutions. He was the Chair of the IEEE Communications Society Wireless Technical Committee and the Chair of the TAOS Technical Committee. He served as a IEEE Computer Society Distinguished Speaker from 2003 to 2005. He is a Senior Member of ACM.

Distinct QoS and shared radio spectrum for M2M and IoTs are expected to be efficiently utilized using cloud-assisted and NFV supported systems. Moreover, 5G networks are expected to be one of the major facilitators for cellular-M2M communications through low cost devices (up to $10), reduced power consumption, and improved security procedures.

ACCEPTED FROM OPEN CALL

Virtualized Cloud Radio Access Network for 5G Transport

Xinbo Wang, Cicek Cavdar, Lin Wang, Massimo Tornatore, Hwan Seok Chung, Han Hyub Lee, Soo Myung Park, and Biswanath Mukherjee

The authors present a new 5G architecture, called V-CRAN, moving toward a cell-less 5G network architecture. They leverage the concept of a V-BS that can be optimally formed by exploiting several enabling technologies such as SDR and CoMP transmission/reception.

ABSTRACT

Current radio access networks (RANs) need to evolve to handle diverse service requirements coming from the growing number of connected devices and increasing data rates for the upcoming 5G era. Incremental improvements on traditional distributed RANs cannot satisfy these requirements, so the novel and disruptive concept of a cloud RAN (CRAN) has been proposed to decouple digital units (DUs) and radio units (RUs) of base stations (BSs), and centralize DUs into a central office, where virtualization and cloud computing technologies are leveraged to move DUs into the cloud. However, separating RUs and DUs requires low-latency and high-bandwidth connectivity links, called "fronthaul," as opposed to traditional backhaul links. Hence, design of the 5G transport network, that is, the part of the network that carries mobile data traffic between BSs and the core network and data centers, is key to meet the new 5G mobile service requirements and effectively transport the fronthaul traffic. Today, consensus is yet to be achieved on how the fronthaul traffic will be transported between RUs and DUs, and how virtualization of network resources will occur from a radio network segment to the centralized baseband processing units. In this article, we present a new 5G architecture, called virtualized cloud radio access network (V-CRAN), moving toward a cell-less 5G network architecture. We leverage the concept of a virtualized BS (V-BS) that can be optimally formed by exploiting several enabling technologies such as software defined radio (SDR) and coordinated multipoint (CoMP) transmission/reception. A V-BS can be formed on a per-cell basis or per-user basis by allocating virtualized resources on demand. For the fronthaul solution, our approach exploits the passive optical network (PON), where a wavelength can be dynamically assigned and shared to form a virtualized passive optical network (VPON). Several use cases of the V-CRAN are presented to show how network architecture evolution can enhance system throughput, energy efficiency, and mobility management.

INTRODUCTION

Fifth generation (5G) networks are envisioned to support 1000× more traffic than today; however, the cost and energy consumption should support affordable services and sustainable growth. 5G radio access network (RAN) design is not only about traffic increase, but also about supporting a large variety of services and devices with unprecedented quality of experience (QoE). Considerable growth of traffic demand and service types poses serious issues of scalability and management for network operators. Even though mobile network capacity can be enhanced by (1) densification, that is, by deploying more radio units and adding another capacity layer with small cells, and (2) usage of larger spectral resources, traditional RANs, based on a distributed architecture, fall short in satisfying these requirements. This calls for a new centralized RAN architecture where processing resources are shared to minimize cost and enable advanced coordination techniques between base stations (BSs). Distributed and centralized RANs and the driver for this evolution are explained below.

DISTRIBUTED RAN

In distributed RAN (DRAN), each BS has two components, co-located in the same cell site: (1) a digital unit (DU) or baseband unit, and (2) a radio unit (RU) or remote radio head. The DU is responsible for baseband processing, while the RU is responsible for transmitting/receiving and digitizing radio signals. The communication channel between each BS and the core network is called backhaul. In DRAN, one way to increase capacity is to densify the network, but cost also increases with additional BSs deployed, as a DU serves its associated RU, and processing resources of a DU cannot be shared among other RUs. Another way to increase capacity is to use advanced technologies such as coordinated multipoint (CoMP) transmission/reception, which can reduce interference and increase throughput. However, this applies stringent delay constraints for control and signaling to guarantee on-time coordination between BSs [1]. But the processing resources of distributed BSs in DRAN are not designed for CoMP, and signaling exchanges undergo long delay (4-15 ms) over backhaul links connecting BSs to the core network [1]. Therefore, DRAN is not a future-proof and scalable solution for 5G.

EVOLUTION OF CLOUD-RAN

Recently, the centralized/cloud-RAN (CRAN) architecture has been proposed [2]. The first generation of CRAN (where C stands for centralized) consists of moving the DU from distributed BSs to a centralized location, where DUs are pooled together, but each remains designated to a specific RU. If DUs

Digital Object Identifier:
10.1109/MCOM.2017.1600866

Xinbo Wang, Lin Wang, and Biswanath Mukherjee are with the University of California, Davis; Cicek Cavdar is with KTH Royal Institute of Technology; Massimo Tornatoreis with Politecnico di Milano; Hwan Seok Chung, Han Hyub Lee, and Soo Myung Park are wtih ETRI Korea.

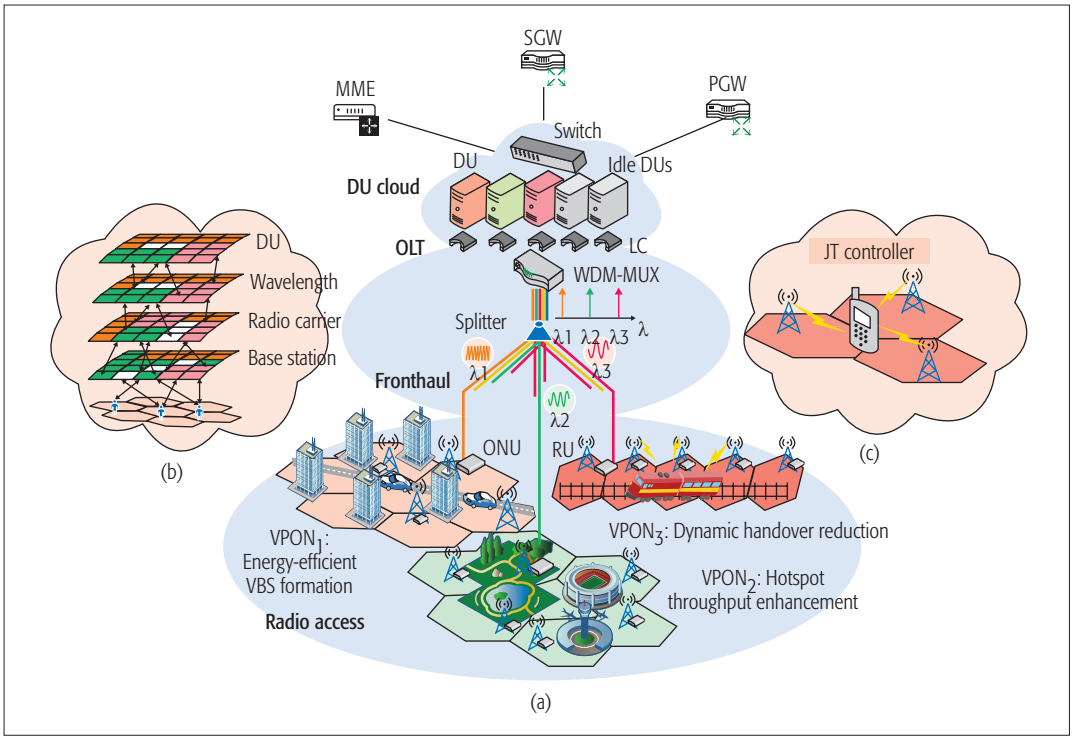

Figure 1. a) Virtualized-CRAN architecture; b) illustration of joint optimization of resources. c) illustration of joint transmission. *Mobility management entity (MME), service gateway (SGW), and PDN gateway (PGW) are functions of mobile core networks. They can be deployed in the mobile core or V-CRAN.

are implemented on general-purpose servers that can be shared by different RUs and serve as a DU cloud, the term "cloud" is justified for CRAN [3].

In CRAN, DUs and RUs need to exchange signals and data, which requires a transport network with high bandwidth, short latency, good jitter performance, and so on. This part of the network is called "fronthaul." Optical channels, over individual fibers or wavelengths, are viable options. However, high-capacity optical channels need to be shared by different BSs adaptively. A fronthaul solution that implements time-wavelength-division multiplexed (T-WDM) channels over a common fiber network is provided by the passive optical network (PON) architecture, known as TWDM-PON [4]. CRAN over TWDM-PON was first proposed in [5] as a promising candidate for implementing CRAN. With upcoming 5G technologies such as millimeter-wave and massive multiple-input multiple-output (MIMO), new trends are developing in terms of different DU and RU splits to relax the bandwidth and latency requirements in fronthaul [6–8].

As novel networking paradigms, such as software-defined networking (SDN) and network functions virtualization (NFV), mature in backbone networks, the trend of virtualization moves toward the access networks, characterizing a virtualized CRAN (V-CRAN). Recently, the concept of a virtualized base station (V-BS) has been proposed to virtualize the computing resources in DU cloud. Pros and cons of different V-BS architectures are compared in [9], while authors in [10] studied the reduction of overall computing resources in DU cloud by consolidating stochastic computational tasks for V-BSs. However, for V-CRAN, virtualization happens not only in DU cloud, but also in other segments, such as fronthaul and radio sites, so efficient management of massive resources in heterogeneous network segments is an important problem. Novel resource allocation solutions are needed to orchestrate computing resources in DU cloud, bandwidth resources in fronthaul, and radio resources in radio sites, to adaptively satisfy per-user demand in V-CRAN.

A Resource Allocation Framework for V-CRAN

In this article, we propose a joint resource allocation framework that can efficiently reconfigure and allocate network resources on a per-cell or per-user basis, by jointly forming a virtualized PON (VPON) and a V-BS.

The first building block is VPON, a virtualized communication channel over a wavelength between many RUs and a DU as in an independent PON. TWDM-PON can provide many such VPONs. A VPON can associate geographically adjacent RUs with the same DU, which can have global information about these RUs and coordinate them with specialized hardware/software for CoMP, and VPON can provide dedicated transport of layer 1 digitized signals for a DU, considering the latency budget of CoMP. Thus, V-CRAN can ease the implementation of CoMP through VPON formation.

A V-BS represents a combination of processing resources in DU-cloud, shared VPON in fronthaul, and a set of RUs. For V-BS, processing resources in a DU are virtualized as functional entities that can be migrated within DU-cloud. We can form a V-BS for an RU, or even for a particular user on demand. When a user is mobile, CoMP provides seamless communication by re-forming dynamic clusters of RUs that can jointly transmit signal to the user. For resource allocation in V-CRAN, V-BS design principles need a cross-layer optimization framework, which assigns resources in an end-to-end manner as shown in Fig. 1b, that is, by allocating DU processing resources, fronthaul

A VPON can associate geographically adjacent RUs with the same DU, which can have global information about these RUs and coordinate them with specialized hardware/software for CoMP, and VPON can provide dedicated transport of layer 1 digitized signals for a DU, considering the latency budget of CoMP.

VPON formation enables not only resource sharing, but also BS coordination. RUs located in an area can be grouped into a VPON and controlled by the same DU. The whole radio-access area can be partitioned into many service areas by formation of VPONs.

transmission resources, and radio resources for each user. This requires that V-BS design jointly consider constraints in different segments of V-CRAN, for example, interference avoidance in the radio network, capacity in fronthaul, and processing capacity in DU cloud.

The contributions of this article are three-fold:

- We present our visions on the physical architecture of V-CRAN.
- We propose a joint resource allocation framework based on VPON and V-BS, and we discuss the design principles.
- We develop a suite of optimization tools to perform joint formation of VPON and V-BS, and we quantitatively evaluate the benefits of this framework for V-CRAN in terms of throughput enhancement, energy saving, and handover reduction.

Note that we benchmark our proposal to the initial CRAN (with only centralization but no virtualization).

Virtualized Cloud Radio Access Network Architecture

Physical Architecture and Implementation Challenge

We illustrate the V-CRAN architecture in Fig. 1a. In DU cloud, commercial servers can be customized to provide real-time baseband processing (and other layer 2/layer 3) functions, and such a server can play the role of a DU [7]. DUs are interconnected by a high-speed layer 2 switch, which exchanges signaling and data among DUs.

TWDM-PON can be used in fronthaul because it can satisfy the stringent delay requirements of the fronthaul segment and provide abundant bandwidth at low cost and energy consumption [4]. TWDM-PON consists of an optical line terminal (OLT) and many optical network units (ONUs). The OLT is collocated with DU cloud, and provides each DU with an optical transceiver and a line card (LC) that can deliver traffic over a wavelength and provide optical-electrical conversion. LCs are connected with a WDM multiplexer/de-multiplexer that can separate traffic on a per-wavelength basis. The passive splitter located remotely from DU cloud can branch the fiber to enlarge the coverage of a TWDM-PON. At the end of a fiber, there is an ONU, which is equipped with a reconfigurable transceiver to tune its transmission wavelength. An ONU is co-located with an RU as an agent of the fronthaul optical network.

The implementation challenges of V-CRAN come from the implementation of virtualization and from joint control/management of virtualized transport, computing, and radio resources. First, network resources must be virtualized and provisioned dynamically, so virtualization techniques used in the IT industry must be tailored to satisfy time-sensitive wireless tasks. Second, there is a trade-off between virtualization gain and implementation complexity, for example, whether to allocate resources on a per-user or per-cell basis.

Virtualized Passive Optical Network Formation

A group of ONUs can share one or multiple wavelengths to form a VPON, that is, a virtualized channel between RUs and a DU. The control software in DU cloud can reconfigure the VPONs by instructing ONUs to retune their operating wavelengths. Since a DU is associated with a VPON via an LC, multiple RUs can be dynamically grouped and controlled by the same DU. Thus, the RU-DU association can be changed in V-CRAN via VPON formation. As wavelength reconfiguration may induce extra delay due to the tuning time and signaling exchange, VPON formation is envisioned to be carried out infrequently (e.g., on a timescale of minutes or hours) to avoid posing an excessive burden on the network control system.

Traffic generated in an RU's coverage area (i.e., inside a "cell" or a sector in a cell) can be grouped into a VPON to share a wavelength's bandwidth. Traffic load in a RAN is dynamic and has a temporal pattern [11] (e.g., alternating busy and idle periods) and a spatial pattern (e.g., with shifted windows of busy hours for urban and residential areas). During busy hours of a day when cells are highly loaded, each cell can be assigned a dedicated VPON. During idle hours, it is desirable to group cells onto fewer VPONs so that idle VPONs and their associated DUs, as well as optical components and interfaces, can be turned off (or put into sleep mode). Besides, DUs and VPON that were assigned to an area under idle hours can be shared by areas currently experiencing busy hours. Thus, the network can be utilized more efficiently, and operational costs are decreased. When traffic load rises and falls in CRAN, VPONs can be re-formed accordingly (e.g., as shown in Fig 1a). VPON1 (orange) can be formed to serve the urban area during idle hours.

Any load consolidation/balancing technique can serve this purpose. A layer 2 switch with enough ports can redirect each frame from any cell to any DU. But when the RAN enlarges and is densified with more cells, the complexity of hardware and control software will grow, and switch latency may negatively impact the performance of BS coordination algorithms such as CoMP. If traffic load from a cell tends to be stable over sustained periods (e.g., minutes), load consolidation/balancing should be done on a per-cell basis in layer 1 instead of a per-frame basis in layer 2. The V-CRAN performs this operation in the optical domain on the fronthaul. This design can reduce DU cloud's complexity and shorten the latency, considering CoMP's rigid latency requirement (typically 1 ms or less) [1].

VPON formation can be formulated as an optimization problem, and solved by integer linear programming (ILP) [12], to minimize usage of wavelengths and DUs. The solution provides wavelength assignment of each ONU, and association of each RU with a DU using the following design principles:

- To associate an RU with a DU, the ONU equipped with the RU must be tuned on the same wavelength of the DU, which is called the wavelength uniformity constraint.
- Total traffic load from all RUs belonging to a VPON must not exceed the capacity of a wavelength or a DU.

Base Station Coordination Assisted by VPON Formation

VPON formation enables not only resource sharing, but also BS coordination. RUs located in an area can be grouped into a VPON and controlled by the same DU. The whole radio access area can be

partitioned into many service areas by formation of VPONs (Fig. 1a). Within each service area, V-CRAN can implement CoMP to improve RAN coverage, bit rate, and throughput. Attractive CoMP techniques providing the highest gains are inter-cell interference cancellation (ICIC) and joint transmission (JT). ICIC reduces interference by allocating different physical resource blocks (RBs) to users in neighboring cells to avoid overlapping of RBs while scheduling users (the spectrum band of an RU is divided into continuous RBs with fixed size). JT is applied in contexts where multiple adjacent BSs cooperate to transmit and receive signals for a user over the same RB so that interference can be converted to useful information. ICIC needs global information about occupation status of RBs in different cells. JT requires not only global information, but also heavy processing resources and short transmission latency to achieve high performance. Also, data and signaling information for scheduling should be duplicated and delivered to all coordinating BSs before the scheduling decision is made [1]. DRAN falls short of satisfying these requirements of JT as discussed earlier. Such delay can be shortened in V-CRAN for two reasons. First, TWDM-PON can provide multiple independent VPONs with fast transport and processing of data between RUs and DU. Second, as shown in Fig. 1c, each DU has global information of all RUs associated with it and is equipped with dedicated hardware/software as a JT controller to provide JT computation.

VPON can also be formed with the objective of maximizing BS coordination by modifying the optimization problem stated above. By solving it, we can split the radio access area into multiple service areas, each served by a cluster of RUs and an associated VPON. A modified optimization problem can also minimize the usage of wavelengths and DUs.

Joint Formation of V-BS and VPON

Depending on the objective, a V-BS formation can be classified as cell-centric or user-centric.

Cell-Centric V-BS Formation

We can allocate "just enough" virtualized resources for each cell to form a cell-centric V-BS, consisting of a shared VPON and virtualized functional entities of a DU, as shown by red dotted lines in Fig. 2 to illustrate a V-BS for Cell2. Two types of functional entities exist in a DU: a cell processor (CP) and a user processor (UP) [13]. A DU can set up a CP for each cell associated with it (Fig. 2). A CP provides part of baseband processing, multiplexes/de-multiplexes user traffic from a cell, and processes cell control messages. A cell can change its serving DU through CP shifting, which first retunes the wavelength of the ONU to change the serving VPON, then sets up a new CP in the destination DU, and finally reclaims the old one in the source DU. A UP provides remaining baseband processing and customized service for each user. A UP can be redirected from one DU to another by UP redirection, in which an overloaded DU redirects some UPs to DUs with extra processing resources through active ports of the internal switch. For example, in Fig. 2, when DU2 cannot accommodate all UPs of Cell2, it can redirect some UPs to DU1. Thus, V-BSs form a cell-less architecture, where there is no dedicated

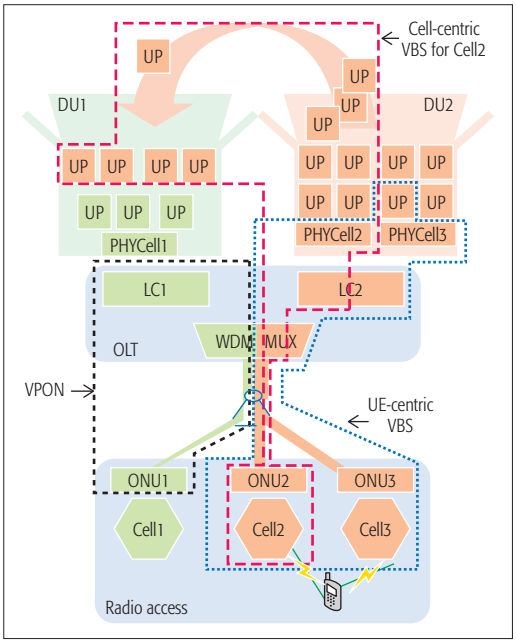

Figure 2. Illustration of VPON and V-BS.

wavelength in fronthaul and functional entities in DU for a specific RU.

The cell-centric V-BS formation problem can be modeled by following the principles described for the optimization problem above with the objective of minimizing usage of wavelengths, CPs and UPs in the DU, and switch ports [12]. Thus, besides the findings, we can also determine for each cell which DU to select to assign a CP; for each user, which DU to select to assign a UP; and number of active switch ports. Hence, we add the following design principles:

- For a cell, a CP must be set up in the DU with which the RU of the cell is associated.
- If a DU is heavily loaded, UPs can be redirected and set up in other DUs, but the source and destination DUs of UP redirection must occupy active ports of the internal switch.

User-Centric V-BS Formation

We can also form a user-centric V-BS for each user, comprising a group of RUs, a shared VPON, and functional entities in a DU, as shown by blue dotted lines in Fig. 2. User-centric V-BS formation assisted by ICIC and JT can enhance the throughput and QoE for users. To implement JT, a set of RUs that can provide good signals for a user needs to be selected from neighboring cells into the *JT-set*. Although throughput of a user can be enhanced by selecting more RUs into the JT-set to form a "bigger" V-BS, it consumes more RBs and increases interference for other users, whose throughput might be jeopardized. ICIC can alleviate this problem by assigning different RBs to users in neighboring cells; but when traffic is heavy, it is hard to assign RBs without conflict due to diminishing available RBs.

A V-BS can be jointly formed with a VPON with the objective of maximizing total throughput of all users in a V-CRAN [14]. By solving the model using constraint programming, we obtain VPON formation and RU-DU association, and which RB is assigned to a user, which RU provides JT service for a user, and which RU sends

VPON can also be formed with the objective of maximizing BS coordination by modifying the optimization problem stated above. By solving it, we can split the radio access area into multiple service areas, each served by a cluster of RUs and an associated VPON. A modified optimization problem can also minimize the usage of wavelengths and DUs.

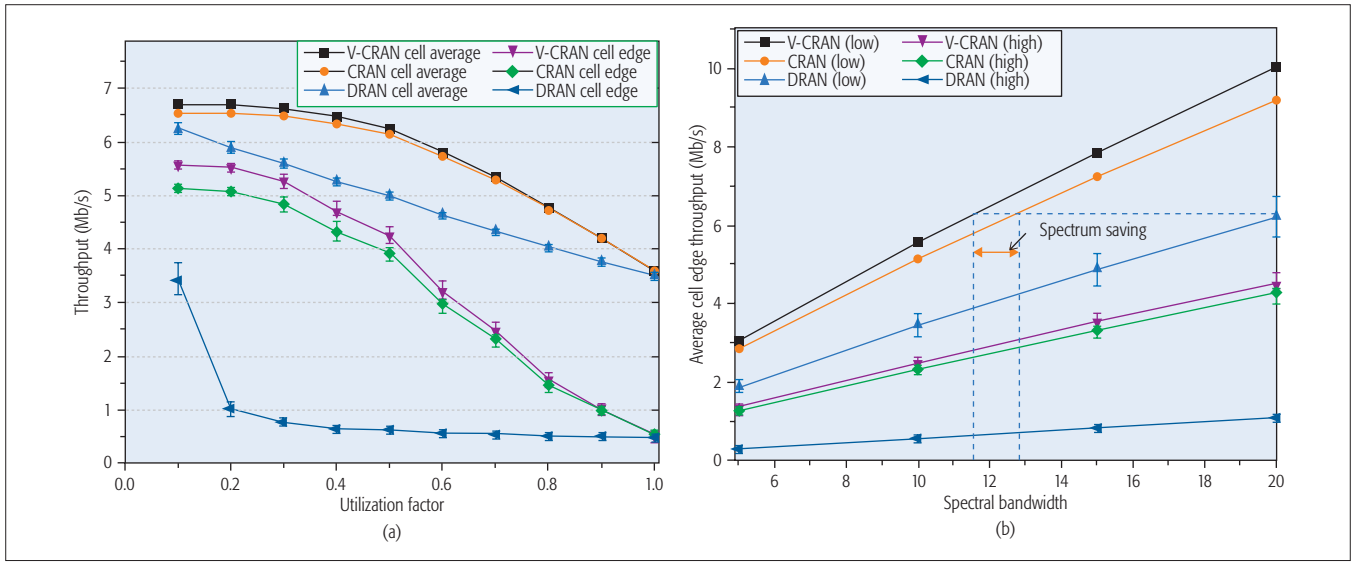

Figure 3. Comparison of throughputs achieved by V-CRAN, CRAN, and DRAN. Simulation experiments are conducted on a 19-cell hexagonal-like cellular network in an urban area (with wrap-around and the Third Generation Partnership Project, 3GPP, urban path loss model). The inter-cell distance is 500 m. We assume a 10 MHz LTE system for each cell. RBs of a cell can be grouped as larger resource block groups (RBGs), conceptually the same as the "resource block" in this article, and assigned to users equally where each user gets a maximum of one RBG. Users are stationary and uniformly distributed in the network. Utilization factor (u-factor) is the ratio of the number of users (equal to the number of occupied RBGs in DRAN scenario) to the total number of RBGs in the network. Bandwidth of a wavelength is 10 Gb/s: a) comparing the average throughput (per user) of three architectures for all users (cell average) and 5 percent worst case users (cell edge); b) comparing the cell edge throughput at low u-factor (0.1) and high u-factor (0.7), respectively. More details can be found in [14].

interfering signal to a user. V-BS formation should follow all design principles described above, especially the wavelength uniformity constraint, which requires RUs in a JT-set to be in the same VPON so that they can be controlled by the same JT controller in the DU. There are additional design principles, as follows:

- All RUs in a JT-set must transmit common data to a user over the same RB, which is called the resource block continuity constraint. This can significantly simplify the transceiver required by the user.
- Inside the DU, the CP of cells in the JT-set and the UP of the user must be set up together so that JT data and control messages can be exchanged and processed fast.
- To avoid providing unfair JT services, the JT-set of a user can only be selected from cells from which the user can receive signals, and must contain at least the host cell, which is the nearest RU to the user.
- RUs that can transmit signals to a user but are not selected in its JT-set should avoid allocating the RB that the user is occupying to reduce interference.
- The number of RBs of an RU allocated to users must not exceed the capacity of the RU's bandwidth resources.
- An RB at an RU cannot be assigned to more than one user at a time.

Moving-User-Centric V-BS Formation

When a user is moving, V-BS can be formed dynamically, surrounding the user, delivering data to the area at the user's arrival. Typically, when a user is going across the cell edge, signal strength from the serving cell diminishes. If the signal strength is lower than a threshold and there is another target cell that can provide a stronger signal, a handover is triggered for the user between serving cell and target cell [15]. When the user is moving within the service area served by a common VPON but traversing a cell edge, a V-BS can be formed to provide a strong enough signal for the user so that handover will not be triggered.

Thus, we add the following design principles for the moving-user-centric V-BS formation problem (besides those described above):

- A V-BS cannot be formed for a mobile user if the wavelength uniformity constraint is violated, when it goes across the boundary of two VPONs. A handover is needed to set up a new UP in the destination DU, migrate information of the user, and reclaim the old UP in the source DU, which we call "inter-VPON" handover.
- A V-BS cannot be formed for a mobile user moving within a VPON if the resource block continuity constraint is violated, even though it is moving within a VPON. A handover is needed to assign a new RB for the user, which we call "intra-VPON" handover.

In V-CRAN, both types of handovers are allowed within DU cloud through internal communication, involving complex signaling exchanges between BSs. Handover latency and failure rate can also be reduced.

Case Studies

Now, we present three use cases of V-CRAN and demonstrate how V-CRAN can outperform the two reference architectures — DRAN and CRAN — with respect to throughput, energy efficiency, and mobility management, and where the enhancement comes from. In DRAN, every DU is co-located with its RU at the cell site. Hence,

a cell needs an independent "housing" facility, a DU remains active all the time, and no BS coordination is deployed at the cell. In traditional CRAN, although DUs are co-located in the DU hotel, there is no sharing of DUs and wavelength; thus, every cell needs an active DU and optical transceiver dedicated to service it, and only limited base station coordination is deployed (ICIC but no JT). We assume simple CRAN architecture as an intermediate state of RAN evolution because it helps us understand where the superiority of V-CRAN comes from.

Increase in System Throughput

In V-CRAN, VPON and V-BS can be jointly formed with the objective of maximizing total throughput of all users, considering the design principles described earlier.

In Fig. 3, we present simulation results to compare throughput performance of V-CRAN, CRAN, and DRAN. In Fig. 3a, we plot average throughputs for cell average and cell edge users, with change of utilization factor ("u-factor") (for definition, see Fig. 3). DRAN achieves lowest throughput and suffers sharp throughput degradation because there is more interference when u-factor becomes larger. CRAN achieves better performance than DRAN (at most 23 and 573 percent for cell average and cell edge users, respectively), because ICIC can reduce interference, especially for low u-factor, where it is easier to avoid overlap between RB-user assignments when RB resources are sufficient. V-CRAN can further enhance the throughput because of JT, with about 25 percent improvement for cell average users compared to DRAN, and throughput enhancement is more significant for cell edge users (almost 645 percent). But when u-factor becomes larger, there is less throughput enhancement for V-CRAN, because a VPON can accommodate fewer cells, and thus fewer JT services can be provided.

In Fig. 3b, we plot cell edge throughput of the three architectures for available bandwidth of each cell ranging from 5 to 20 MHz under two scenarios: low and high u-factors. Confirming Fig. 3a, V-CRAN achieves the highest throughput for various spectral bandwidth availability. Note the bandwidth saved by V-CRAN and CRAN to achieve the same throughput of DRAN with 20 MHz configuration. For low u-factor, V-CRAN and CRAN save around 8.5 MHz (43 percent) and 7.1 MHz (35.5 percent), respectively. For high u-factor, V-CRAN and CRAN with smallest available bandwidth can achieve higher cell edge throughput than DRAN with largest bandwidth. V-CRAN's superiority is more noticeable when u-factor is small, confirming Fig. 3a. Results in Fig. 3b show that the evolution of network architecture can not only bring higher data transmission rate, but also more efficient use of precious spectral bandwidth.

In Table 1, we list the amounts of various resources needed to achieve the cell edge throughput of DRAN with 20 MHz configuration in the three architectures. DRAN consumes as many DUs as number of cells because of co-location of DU and RU, but it does not need the optical transport network. CRAN consumes the same number of DUs and wavelengths as the number of cells, and it can save some spectral bandwidth. V-CRAN consumes much less spectrum, DUs, and wavelength resources. Note that although V-CRAN has slight improvement over C-RAN in terms of throughput, it is much more resource-efficient, thanks to the virtualization and sharing of VPON and the V-BS concept by utilizing 89 percent fewer wavelengths and DUs than C-RAN.

	Spectral bandwidth	Number of DUs	Number of wavelengths
V-CRAN	11.5 MHz	2	2
CRAN	12.9 MHz	19	19
DRAN	20 MHz	19	0

Table 1. Resource consumption of V-CRAN, CRAN, and DRAN to achieve cell edge throughput of DRAN with 20 MHz at low utilization factor.

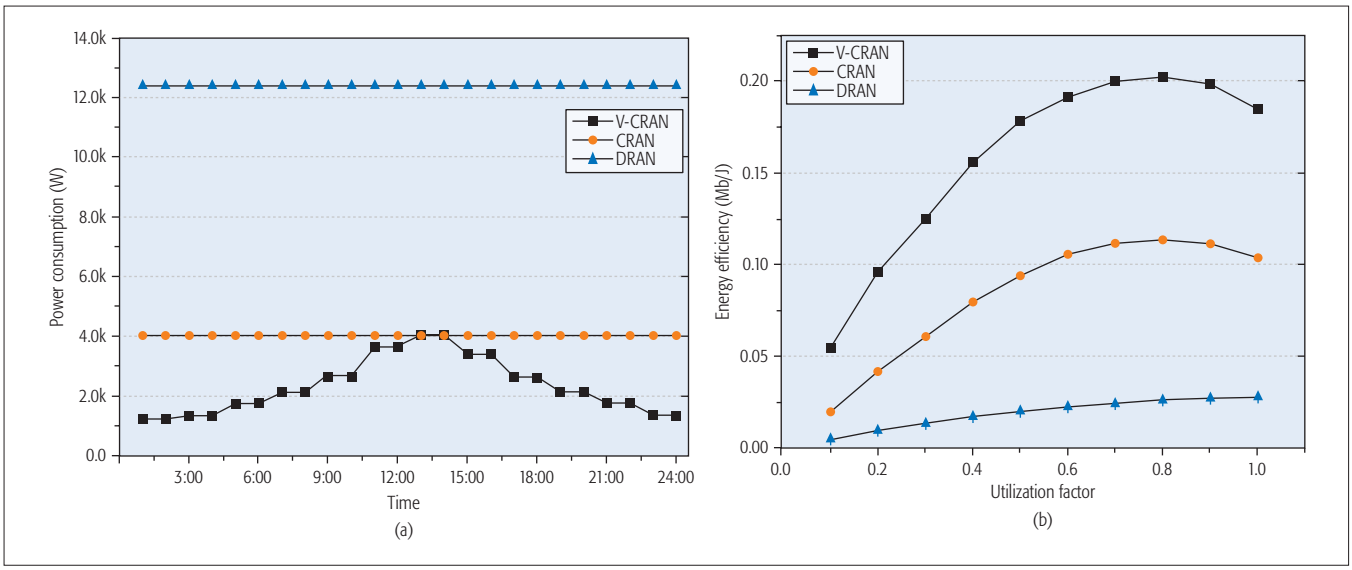

Figure 4. a) Comparing power consumption of V-CRAN, CRAN, and DRAN architectures during a day. The daily 24 hours are slotted into 12 periods, each of 2-hour length. Users have busy hours from 10:00 to 17:00, and traffic loads reach the peak at 13:00-14:00; b) comparing the energy efficiency of three architectures for different load ratios. Note that energy efficiency is the number of bits that can sent by consuming 1 J. More details can be found in [12].

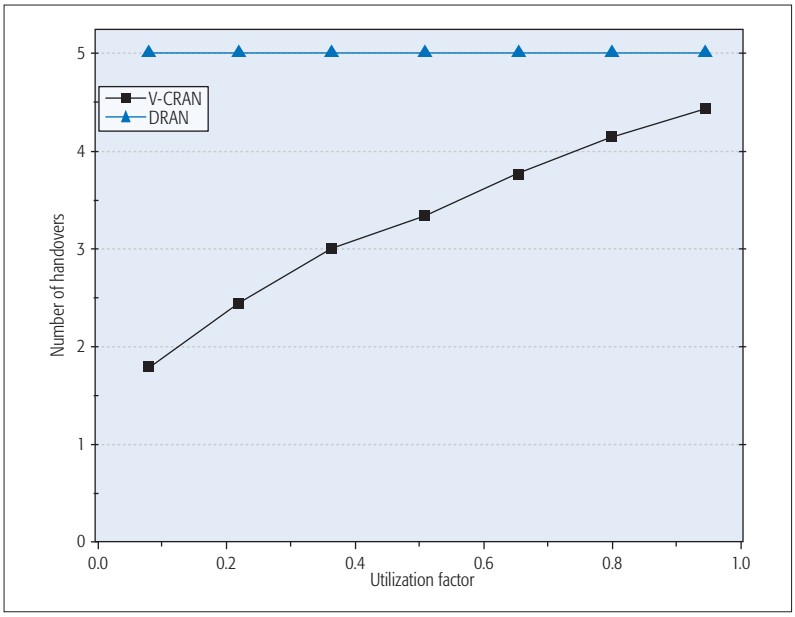

Figure 5. Comparison of average number of handovers suffered by a user for different load ratio when it is traversing the network. Simulation experiments are conducted in a network with one DU cloud connected to 100 cells by a TWDM-PON. A modified random waypoint mobility model is assumed. Every user is connected to the network for a time duration during which it is supposed to keep the communication active, and must traverse five cells before it gets disconnected from the network. Such a user arrives to get connected to the network following a Poisson distribution. 100,000 such users are simulated. This modification is made because we want to fix the number of handovers in DRAN to benchmark the handover reduction achievable for V-CRAN while eliminating the confounding factors (e.g., randomness) induced by the mobility model. Note that we did not plot the performance of CRAN because we assume there is no JT implemented in traditional CRAN architecture, and thus it has the same average number of handovers as DRAN.

Enhancement in Energy Efficiency

V-CRAN can save energy because of two aspects: pooling resources and forming cell-centric V-BSs. First, pooling DUs can save a large amount of energy consumed by housing facilities, which require power to ensure proper operational conditions (e.g., cooling) for DUs, although they do not perform network functions. Although this decoupling of DU and RU needs a transport network, TWDM-PON provides an energy-efficient solution because the optical fiber and splitter are passive devices that consume little energy, and OLTs and ONUs are low-energy-consuming devices.

Second, energy is saved by forming a V-BS for each cell using "just enough" virtualized resources and timely reclaiming of extra resources with variations of traffic demands. Also, we can shut down more DUs proactively by offloading their remaining traffic load to others through UP redirection described earlier.

In Fig. 4a, we plot energy consumptions of the three architectures during a day. On a per-day basis, V-CRAN saves 46.1 percent and 84.1 percent power consumption compared with CRAN and DRAN, respectively. By comparing CRAN and DRAN, we see the savings in power consumption (and enhancements in energy efficiency) due to pooling of resources, because we only have DU pooling but no sharing mechanism for the reference CRAN architecture. V-CRAN can further save power consumption (and further enhance energy efficiency) compared with CRAN because of V-BS formation. By forming V-BS adaptive to the traffic variation during a day, a large amount of power can be saved by shutting down unused network components during idle hours.

In Fig. 4b, we further compare the energy efficiency of three architectures. V-CRAN achieves much higher energy efficiency than other two references because it can enhance the throughput and reduce power consumption. We find that optimal energy efficiency can be achieved when u-factor is around 0.8.

Improvement in Mobility Management

When mobility of users is considered, number of handovers can be minimized by forming a user-centric V-BS for a mobile user.

In Fig. 5, we plot average number of handovers suffered by a user in V-CRAN and DRAN for different u-factors. V-CRAN needs fewer handovers for various u-factors compared to DRAN. Handover reduction is more noticeable when u-factor is lower because of two reasons. First, the number of inter-VPON handovers is less because when traffic load in each cell is small, a single VPON can cover a larger area, and a DU can provide more BS coordination for users moving within the area. Second, intra-VPON handover is also less because when traffic load is low, there are plenty of RBs available in each cell, so there are more opportunities to find the same available RB in cooperating cells for a user. But when u-factor increases, both numbers of inter-VPON and intra-VPON handovers increase, because the VPON is smaller and RBs are fewer.

Conclusion

5G mobile networks need to not only provide higher data rates but also support diverse quality of service requirements coming from emerging mobile services and user equipment. The traditional distributed radio access network is not scalable and cost-efficient for managing the expanding network infrastructure and resources in a flexible manner to adapt to different service requirements. Cloud RAN centralizes the digital unit of a base station to a central office through a high-speed optical transport network, that is, a TWDM-PON. But simply pooling them together cannot fully achieve the gains of CRAN. In this article, we present the virtualized CRAN, which virtualizes network resources, including DUs, BSs, and TWDM-PON between them. A joint framework, a virtualized base station that can be formed for either a cell or a user, was proposed for resource sharing and base station coordination. Within this framework, joint optimization of heterogeneous resources can be achieved in 5G transport networks. We present several use cases of V-CRAN to show how network architecture evolution can enhance system throughput, energy efficiency, and mobility management.

Acknowledgment

This work was supported by an Institute for Information & Communications Technology Promotion (IITP) grant funded by the Korea government (MSIP) (B0132-15-1004, SDN-based wired and wireless converged optical access networking).

References

[1] NGMN, "RAN Evolution Project: CoMP Evaluation and Enhancement," Mar. 2015.
[2] China Mobile Research Institute, "C-RAN: The Road towards Green RAN V3.0," Dec. 2013; http://labs.chinamobile.com/cran/2014/06/16/c-ran-white-paper-3-0/, accessed Jan. 2017.
[3] T. Pfeiffer, "Next Generation Mobile Fronthaul and Midhaul Architectures," *IEEE J. Optical Commun. Networking*, vol. 7, no. 11, Nov. 2015, pp. B38–B45.
[4] ITU-T G.989 Rec. Series, "40-Gigabit-Capable Passive Optical Networks (NG-PON2)," Mar. 2013.
[5] D. Iida *et al.*, "Dynamic TWDM-PON for Mobile Radio Access Networks," *Optics Express*, vol. 21, no. 22, Nov. 2013, pp. 26,209–18.
[6] N. Shibata *et al.*, "Dynamic IQ Data Compression Using Wireless Resource Allocation for Mobile Front-Haul with TDMPON (Invited)," *IEEE J. Optical Commun. Networking*, vol. 7, no. 3, Mar. 2015, pp. A372–78.
[7] NGMN, "Further Study on Critical C-RAN Technologies," Mar. 2015.
[8] C. L. I *et al.*, "Rethink Fronthaul for Soft RAN," *IEEE Commun. Mag.*, vol. 53, no. 9, Sept. 2015, pp. 82–88.
[9] D. Pompili *et al.*, "Elastic Resource Utilization Framework for High Capacity and Energy Efficiency in Cloud RAN," *IEEE Commun. Mag.*, vol. 54, no. 1, Jan. 2016, pp. 26–32.
[10] J. Liu *et al.*, "Statistical Multiplexing Gain Analysis of Heterogeneous Virtual Base Station Pools in Cloud Radio Access Networks," *IEEE Trans. Wireless Commun.*, vol. 15, no. 8, Aug. 2016, pp. 5681–94.
[11] P. Chunyi *et al.*, "Traffic Driven Power Saving in Operational 3G Cellular Networks," *Proc. ACM Mobicom*, U.S., 2011, pp. 121–32.
[12] X. Wang *et al.*, "Energy-Efficient Virtual Base Station Formation in Optical-Access-Enabled Cloud-RAN," *IEEE JSAC*, vol. 34, no. 6, Jan. 2016.
[13] X. Wang *et al.*,"Joint Allocation of Radio and Optical Resources in Virtualized Cloud RAN with CoMP," *Proc. IEEE GLOBECOM*, Washington, DC, 2016,pp. 1–6.
[14] B. Haberland *et al.*, "Radio Base Stations in the Cloud," *Bell Labs Tech. J.*, vol. 18, no. 1, May 2013, pp. 129–52.
[15] LTE; Evolved Universal Terrestrial Radio Access (E-UTRA) and Evolved Universal Terrestrial Radio Access Network (E-UTRAN); Overall Description; Stage 2. 3GPP TS 36.300 v. 10.5.0, Release 10, Nov. 2011.

Biography

Xinbo Wang received his B.Eng. degree in telecommunication engineering from Beijing University of Posts and Telecommunication, China, in 2013 and his Ph.D. degree in computer science from the University of California, Davis, in 2017. During April to September 2016, he was a visiting researcher at KTH Royal Institute of Technology, Sweden, and a visiting research engineer at Orange Labs and Nokia Bell Labs, France. He is currently a research scientist at Facebook Inc. He served/is serving as a TPC member of WCNC 2017 and WCNC 2018. His research interests include cloud radio access networks and optical transport networks for 5G.

Cicek Cavdar is a senior researcher in the Communications Systems Department at KTH Royal Institute of Technology. She has been leading a research group in the Radio Systems Lab composed of eight researchers focusing on design and planning of intelligent network architectures, direct air to ground communications, and IoT connectivity platforms. She finished her Ph.D. studies in computer science at the University of California, Davis in 2008 and at Istanbul Technical University (ITU), Turkey, in 2009. After her Ph.D., she worked as an assistant professor in the Computer Engineering Department, ITU. She served as Publications Chair of IEEE/IFIP Network Operations and ManagementSymposium 2016) and Co-Chair of the IEEE ICC 2017 Green Communications Symposium. She chaired the Green Broadband Communications Workshop and the Panel "Green 5G Mobile and Broadband Access Networks" at IEEE GLOBECOM 2014. She has led several EU EIT Digital projects such as "5GrEEn: Towards Green 5G Mobile Networks" and "ICARO-EU: Seamless DA2GC in Europe."

Lin Wang received her B.Eng. degree in telecommunication engineering from Beijing University of Posts and Telecommunication in 2013 and her M.S. in computer science from the University of Southern California, Los Angeles, in 2015. She is currently pursuing a Ph.D. degree in computer science at the University of California, Davis. Her research interests include next generation Ethernet passive optical networks and data center networking.

Massimo Tornatore [S'03, M'06 SM'13] is an associate professor with the Departmentof Electronics, Information, and Bioengineering, Politecnico di Milano. He also holds an appointment as an adjunct associate professor with the Department of Computer Science, University of California, Davis. He is an author of more than 200 peer-reviewed conference and journal papers. His research interests include performance evaluation, optimization and design of communication networks (with an emphasis on the application of optical networking technologies), cloud computing, and energy-efficient networking. He is a member of the Editorial Board of *Photonic Network Communications*, *Optical Switching and Networking*, and *IEEE Communication Surveys & Tutorials*. He is an active member of the Technical Program Committees of various networking conferences such as IEEE INFOCOM, OFC, ICC, and GLOBECOM. He has been a co-recipient of 10 best-paper awards.

Hwan Seok Chung [SM'12] received his Ph.D. degree in electronics engineering from Korea Advanced Institute of Science and Technology (KAIST) in 2003. In 2003, he was a postdoctoral research associate at KAIST, where he worked on a hybrid CWDM/DWDM system for metro area networks. From 2004 to 2005, he was with KDDI R&D laboratories Inc., Saitama, Japan, engaged in research on wavelength converters and regenerators. Since 2005, he has been with the Electronics and Telecommunication Research Institute (ETRI), Daejeon, Korea, where he is currently a director of the optical access research section. His current research interests include mobile fronthaul, high-speed PON, and modulation format. He has served as a Technical Committee member of OFC, OECC, COIN, ICOCON, MWP, and Photonic West. He was the recipient of the Best Paper Awards at OECC in 2000 and 2003 as well as ETRI in 2011 and 2012.

Han Hyub Lee [M] received his B.S., M.S., and Ph.D. degrees in physics from Chungnam National University, Daejeon, in 1999, 2001, and 2005, respectively. His doctoral research included the application of a Raman fiber amplifier and gain-clamped SOA for WDM systems. From 2006 to 2007, he was a postdoctoral researcher at AT&T Laboratory, Middletown, New Jersey, where he worked on extended WDM/TDM hybrid PONs using a wideband optical amplifier. In 2007, he joined ETRI as a senior researcher and has worked on optical access networks. He has contributed to the development of international standardizations.

Soo Myung Park received his B.S. degree in computer science engineering from the University of Dankuk in 1990, and M.S. and Ph.D. degrees in computer engineering from the University of Konkuk in 1992 and 1999, respectively. Since May 2000, he has been with the Department of Communication and Internet, ETRI. His current research interests include the transport (MPLS-TP packet transport network and optical transport network) SDN area and open-source-based SDN projects including OpenDaylight, ONOS, and so on.

Biswanath Mukherjee [S'82, M'87, F'07] is a Distinguished Professor at the University of California, Davis, where he was Chairman of Computer Science during 1997-2000. He received his B.Tech. (Hons) degree from the Indian Institute of Technology, Kharagpur (1980) and Ph.D. from the University of Washington, Seattle (1987). He was General Co-Chair of IEEE/OSA OFC 2011, Technical Program Co-Chair of OFC '09, and Technical Program Chair of IEEE INFOCOM '96. He co-founded and served during 2007–2010 as the Steering Committee Chair of IEEE ANTS (the leading networking conference in India promoting industry-university interactions), and served as General Co-Chair of ANTS in 2007 and 2008. He is Editor of Springer's Optical Networks Book Series. He has served on eight journal Editorial Boards, most notably *IEEE/ACM Transactions on Networking* and *IEEE Network*. In addition, he has Guest Edited Special Issues of *Proceedings of the IEEE*, the *IEEE/OSA Journal of Lightwave Technology*, *IEEE JSAC*, and *IEEE Communications Magazine*. To date, he has supervised 69 Ph.D.s to completion and currently mentors 15 advisees, mainly Ph.D. students. He was the winner of the 2004 Distinguished Graduate Mentoring Award, the 2009 College of Engineering Outstanding Senior Faculty Award, and the 2016 UC Davis International Community Building Award at UC Davis. He is co-winner of 10 Best Paper Awards, including three from IEEE GLOBECOM Symposia, four from IEEE ANTS, and two from the National Computer Security Conference. He is the author of the graduate-level textbook *Optical WDM Networks* (Springer, January 2006). He served a five-year term on the Board of Directors of IPLocks, a Silicon Valley startup company (acquired by Fortinet). He has served on the Technical Advisory Boards of over a dozen startup companies, including Teknovus (acquired by Broadcom). He was the winner of the IEEE Communications Society's inaugural (2015) Outstanding Technical Achievement Award "for pioneering work on shaping the optical networking area."

The traditional distributed radio access network is not scalable and cost-efficient for managing the expanding network infrastructure and resources in a flexible manner to adapt to different service requirements. Cloud RAN centralizes the digital unit of a base station to a central office through a high-speed optical transport network, that is, a TWDM-PON. But simply pooling them together cannot fully achieve the gains of CRAN.

ACCEPTED FROM OPEN CALL

Efficient Use of Paired Spectrum Bands through TDD Small Cell Deployments

Adrián Agustín, Sandra Lagen, Josep Vidal, Olga Muñoz, Antonio Pascual-Iserte, Zhiheng Guo, and Ronghui Wen

The authors introduce TDD SeNB to operate in the unused resources of an FDD-based system. This proposal alleviates the saturated DL/UL transmission commonly found in FDD-based systems through user offloading toward a TDD system based on SeNBs. In this context, the flexible duplexing concept is analyzed from three points of view: regulation, LTE standardization, and technical solutions.

ABSTRACT

Traditionally, wireless cellular systems have been designed to operate in FDD paired bands that allocate the same amount of spectrum for both DL and UL communications. Such design is very convenient under symmetric DL/UL traffic conditions, as used to be the case when voice transmission was predominant. However, due to the overwhelming advent of data services, which involves large asymmetries between DL and UL, the conventional FDD solution becomes inefficient. In this regard, flexible duplexing concepts aim to derive procedures to improve spectrum utilization by adjusting resources to actual traffic demand. In this work, we review these concepts and propose the introduction of TDD SeNB to operate in the unused resources of an FDD-based system. This proposal alleviates the saturated DL/UL transmission commonly found in FDD-based systems through user offloading toward a TDD system based on SeNBs. In this context, the flexible duplexing concept is analyzed from three points of view: regulation, LTE standardization, and technical solutions.

INTRODUCTION

The most salient feature in the evolution of mobile services is the imposition of data services over voice traffic demand, which requires the redefinition of current wireless cellular networks and communication standards. Second and third generation wireless cellular systems were designed under a symmetric traffic assumption as a result of the predominance of voice traffic. Accordingly, the common technical solution adopted worldwide was the use of paired bands under frequency-division duplexing (FDD). The legacy of this assumption has survived in fourth generation (4G) systems, even though a time-division duplexing (TDD) frame definition was also defined early on.

The new habits of users have produced high asymmetries in data traffic demand, that is, the amount of data transmitted in the downlink (DL) connection is usually much larger than the amount of data in the uplink (UL) transmission [1]. The most conservative measured DL:UL traffic asymmetry ratio across different macro eNBs (MeNBs) is 4:1 [2] due to video downloading and internet browsing. In its turn, the uploading of shared contents in social media is generating the opposite tendency, attaining ratios of 1:4 [3]. Such time/space-varying imbalance of data traffic negatively affects the spectral efficiency of FDD-based systems since its inflexibility translates into an underuse of one band while the other band may be congested. This inefficiency could be reduced by adopting unpaired band technologies based on TDD, in which the use of radio resources dedicated to DL and UL transmissions can be flexibly adapted as a function of the actual traffic demand.

The present work explores how the spectral efficiency of LTE FDD-based systems can be improved under traffic asymmetries by means of the flexible duplexing concept [4]. In particular, the proposed solution assumes an FDD MeNB serving area in which TDD-based small eNBs (SeNBs) are deployed and operate in the unused resources. With the objective of having a clear idea of the benefits of the proposed solution, a simple scenario with one SeNB is considered, as shown in Fig. 1. In this context, the following challenges have to be faced.

Coexistence of Adjacent FDD/TDD Systems: Because of non-ideal transmit filters, adjacent channel interference (ACI) originates from systems operating in adjacent bands. ACI can be managed by either imposing a minimum distance between transmitting nodes [5] or defining a set of guard bands and power spectrum masks [6].

Impact of Different TDD-LTE Frame Pattern Configurations: Conventionally, in TDD mode, all eNBs transmit simultaneously in DL, and, at a different time instant, all UEs transmit in UL. This approach aims to limit the active nodes that generate interference in each case. However, the use of flexible TDD MeNB/SeNBs entails a dynamic decision of their own TDD DL-UL frame pattern that introduces new types of interference in cellular systems; that is, MeNBs/SeNBs can be interfered by other MeNBs/SeNBs. Nevertheless, if this interference is properly managed, significant throughput gains can be obtained in low to medium system loads [5].

Shared Spectrum Access: Interference management is important when eNBs with different maximum transmitting power levels operate in the same radio resources. However, deploying outdoor SeNBs, with a maximum equivalent isotropical radiated power (EIRP) of 30 dBm and a height below 12 m, is a simple solution that

Digital Object Identifier:
10.1109/MCOM.2017.1601173

Adrián Agustín, Sandra Lagen, Josep Vidal, Olga Muñoz, and Antonio Pascual-Iserte are with Universitat Politècnica de Catalunya; Zhiheng Guo and Ronghui Wen are with Huawei Technologies Co., Ltd.

reduces interference and allows large reuse of the spectrum [7].

The proposed solution in this work has also been investigated in [8] where TDD cognitive SeNBs are allowed to exploit the FDD-DL spectrum. In that case, the TDD SeNB listens to the FDD UL signals in order to detect if there are active FDD-UEs (or MUEs in the following) in the neighboring area. If this is not the case, the UEs associated with the TDD SeNB (SUEs in what follows), and of course the SeNB itself, are allowed to transmit in the FDD-DL band. Another implementation of a TDD system in the unused FDD-UL spectrum is proposed also in [9], where the interference between the FDD-UL and TDD systems is avoided thanks to a tight time coordination between FDD and TDD systems.

Introducing the flexible duplex concept by deploying multiple TDD SeNBs in wireless FDD-based communications systems entails dealing with many issues that are not currently found in current cellular wireless systems, in which SeNBs and MeNBs are assumed to operate either in different spectrum bands or, in the case of coexisting in the same band, to use the same duplexing mode. With the objective of elucidating how the proposed approach impacts on system performance, a simplified scenario is analyzed here, where just a single TDD SeNB in the macrocell area adopts the flexible duplex concept (Fig. 1). How the position of the SeNB and its transmitted power impact system performance and which FDD band can be reused (UL or DL) are important aspects that need to be understood. This knowledge will be decisive to determine which of the existing TDD SeNBs deployed in the macrocell area can operate in the underutilized spectrum.

This article examines alternatives and challenges in the implementation of the flexible duplexing concept in LTE. Specifically, the next section details the proposed schemes along with the pros and cons of reusing either the FDD-UL or the FDD-DL bands. The third section addresses the limitations of applying the flexible duplexing concept as a result of the application of current regulation and/or LTE standard constraints. The fourth section reviews future research lines related to multiple flexible duplex SeNBs. The final section underlines some conclusions.

Flexible Use of the Paired Band

In an FDD-based system, a guard band (usually of several megahertz) is required to separate the paired UL and DL bands (usually a few megahertz). Therefore, the current underutilized spectrum cannot accommodate a new FDD-based system. Fortunately, TDD-based systems are not affected by such guard band constraint. Accordingly, we investigate the deployment of a TDD SeNB operating in the unused spectrum of an FDD-based system. The proposed schemes for multiplexing MeNB and SeNB are described below. Since it is important to know how many resources are needed by the MeNB as a function of the traffic demand, we discuss how the resource provisioning could be estimated. Finally, the benefits in terms of user throughput and resource utilization of the proposed flexible duplexing concept are presented.

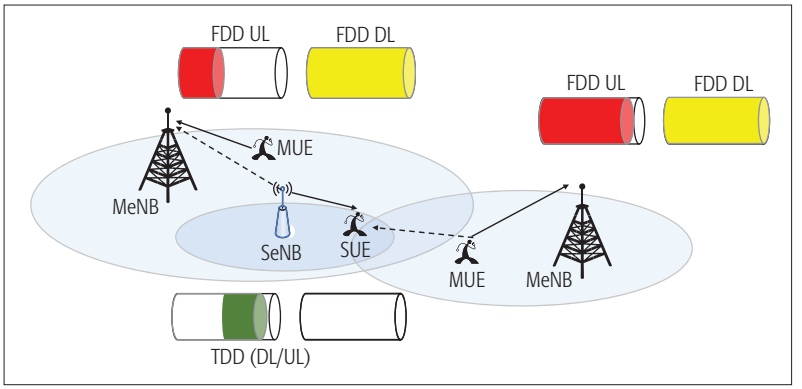

Figure 1. Example of the flexible duplex concept. There is one macrocell area with unused resources in the UL but saturated in the DL. A TDD SeNB is allowed to operate in the unused resources while coexisting with the MeNB.

Flexible Duplexing Methods in LTE

The options to implement the flexible duplexing concept depend on whether it is possible to release part of the licensed FDD spectrum or not, paying special attention to those frequency resources in LTE reserved for control channels. The first option, referred to in the following as *efficient in-band orthogonal use of the licensed spectrum*, subsumes the case where the SeNB works in the band of interest through an orthogonal access manner. The implementations using the orthogonality in time and in frequency are depicted in Figs. 2a and 2b, respectively. Both cases are described in detail below.

The second option, illustrated in Fig. 2c, assumes that the MeNB can make good use of the intra-band component carrier (CC) aggregation LTE feature, adapting its operational bandwidth according to the MeNB traffic demand. MeNB and SeNB work orthogonally in frequency, but, in contrast to the approach in Fig. 2b, the FDD-based system can adapt its reserved resources, and the SeNB can access the channel through frequency multiple access.

Efficient In-Band Orthogonal Use of the Licensed Spectrum: SeNB and MeNB operate in the same band through orthogonal multiple access. The SeNBs might operate at the unused MeNB band in time-multiplexing access (TMA) or frequency-multiplexing access (FMA) (Figs. 2a and 2b) when FDD-UL band is reused. The second approach, depicted in Fig. 2b, requires that the receiver at the MeNB be equipped with additional analog filters to avoid the undesired interference from the sidelobe signal transmission in adjacent bands coming from SeNBs. This is due to the fact that the bandwidth of the receive filter at the MeNB corresponds to the whole operational bandwidth. Therefore, this option comes at the cost of losing flexibility since unused time frequency resources of the FDD band have to be identified beforehand and have to be fixed over time because analog filters are necessarily inflexible.

Reusing the FDD-UL band with the in-band TMA approach demands paying special attention to physical UL control channels (PUCCHs) allocated at the edge of the band (see dark red resources in Fig. 2). Those resources are used by MUEs to transmit acknowledgments at a predetermined

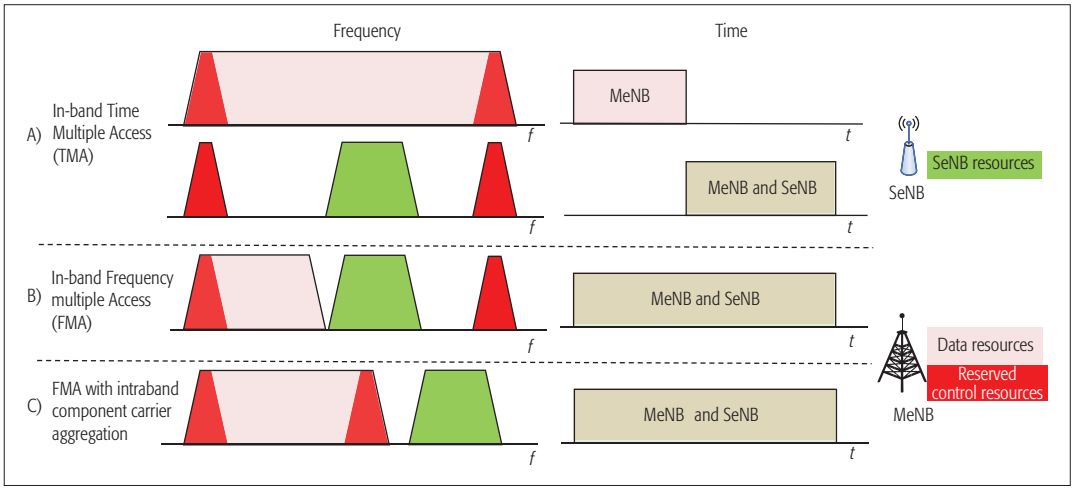

Figure 2. a), b): An overlaid TDD SeNB operates in the unused resources of FDD-UL under time multiple access and frequency multiple access, respectively. In c), the MeNB adjusts its operational LTE FDD-UL bandwidth using CC aggregation, and the TDD SeNB works in the released resources. Dark red resources on the left denote reserved frequencies for FDD control plane communications.

delay after DL transmission. Therefore, SeNB and MeNB should agree on accessing those resources at different time instances, or SeNB and MeNB should be equipped with additional analog filters to preserve the adequate reception of PUCCH at the MeNB. Similarly, the in-band FMA needs transmitting and receiving nodes to be equipped with analog filters (Fig. 2b), and in-band interference becomes ACI.

Reusing the FDD-DL band is more complicated due to the current frame structure defined by LTE. The reserved frequency resources for control plane communications are found in the central part of the FDD-DL band, which complicates the adoption of the in-band FMA approach. Regarding the TMA approach, it is even worse because LTE defines several DL subframes for transmitting system information, which leaves just two subframes (out of ten) for the TDD SeNB, as detailed later.

General system parameters	
MeNBs deployment	One sector of a macrocell area with 1 FDD MeNB. Hexagonal deployment of FDD MeNBs. Two interfering MeNBs operate in FDD with normal usage at Inter-site distance of 500 m
SeNB deployment	1 TDD SeNB at a distance 100 m from MeNB.
UEs deployment	50 UEs uniformly distributed within the macrocell area
Spectrum	Licensed paired FDD: 10 MHz for DL and 10 MHz for UL
Frequency carrier	2.5 GHz
Transmit power	46 dBm (MeNB), 24 dBm (SeNB), 23 dBm (UE)
MeNB antenna system	17 dBi, 3D, sectorized, 2 antennas
SeNB antenna system	5 dBi, 2D, omnidirectional, 2 antennas
UE antenna system	0 dBi, 2D, omnidirectional, 2 antennas
Noise figure	5 dB at MeNB and SeNB, 9 dB at UEs
Noise spectral density	–174 dBm/Hz
Propagation conditions	Pathloss and shadowing as in [5]. Frequency selective fading follows the typical urban model.
Cell association	UEs are associated with the MeNB or SeNB using the reference signal received power (RSRP) combined with a cell range expansion (CRE) bias. CRE is adjusted so as to get approximately 10 UEs associated with an SeNB.
Traffic model	FTP model 3. Packet size 2 Mb.

Table 1. Simulation assumptions.

Efficient Use of the CC-Based Licensed Spectrum: LTE allows reconfiguring bands thanks to the CC aggregation concept [10]. Here, it is assumed that the licensed spectrum is divided in CCs, and there is an entity responsible for the dynamic long-/medium-term resource allocation that selects the number of required CCs by the MeNB as a function of the traffic demand. In this scenario, the unused CCs can be engaged by the TDD SeNB.

Using this approach, the FDD MeNB and the TDD SeNB operate on orthogonal carriers, which are isolated thanks to the respective analog transmit filters (Fig. 2c). The source of interference is ACI between both systems: TDD SeNB (in DL) over the FDD MeNB when the FDD-UL spectrum is reused, and FDD MeNB over the TDD SeNB (in UL) when the FDD-DL spectrum is reused. In all cases, the interference level depends on the probability of line of slight (LOS) between MeNB and SeNB, but this can be reduced by deploying SeNBs at a lower height than MeNBs.

Resource Provisioning

The use of the flexible duplexing concept requires a tight estimation of the amount of spectrum that will not be employed by the FDD MeNB system. A suitable key performance indicator (KPI) that contains this information is the resource utilization (RU). This metric reports the ratio between the total number of resources used by data traffic over the total number of resources available for data traffic. More explicitly, the RU of the kth cell in the dth transmit direction (either DL or UL) (ρ_k^d) can be estimated based on the traffic statistics and the average spectral efficiency as [11]

$$\rho_k^d = \frac{\text{average offered traffic}}{\text{traffic served}} = \frac{\lambda_k^d L_k^d}{x_k^d} \frac{1}{C_k^d}, \quad (1)$$

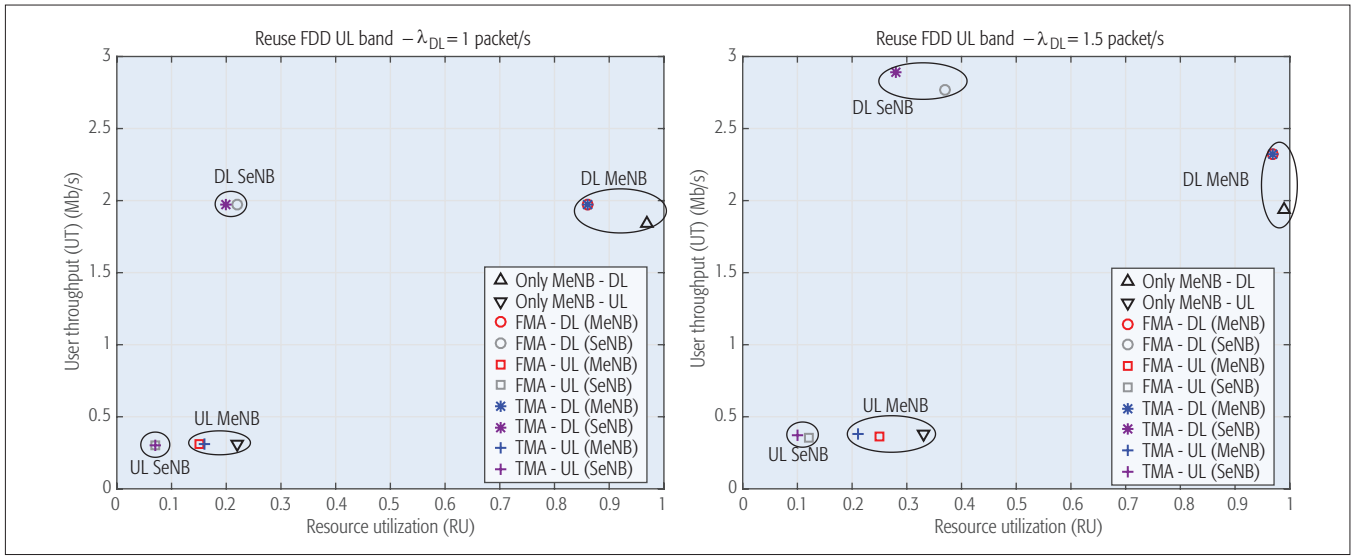

Figure 3. Mean UT (in mb/s) vs. RU for FDD-UL band with d = 100 m: left, λ_{DL} = 1; right, λ_{DL} = 1.5 (packets/s); traffic asymmetry: $\lambda_{DL} = 10 \times \lambda_{UL}$.

where λ_k^d is the mean packet arrival rate (in packets per second), L_k^d denotes the mean packet length (in bits per packet), x_k^d refers to the number of resources, and C_k^d is the average spectral efficiency of the kth cell in the dth transmit direction (in bits per second per resource). Furthermore, the average number of bits in the queue of the kth cell in the dth transmit direction (W_k^d) can be modeled as a function of the RU when packet arrival instants follow a Poisson process. In such a case, for $\rho_k^d \le 1$, we have the following expression [11]:

$$W_k^d = \frac{\rho_k^d}{1-\rho_k^d} \frac{\sigma_k^{2d} + \left(L_k^d\right)^2}{2L_k^d}, \quad (2)$$

where σ_k^{2d} denotes the variance of the packet length. Note that a large RU factor implies a large average queue size.

From Eqs. 1 and 2, one can either estimate how many resources are needed as a function of the traffic statistics or derive a maximum RU for a given target W_k^d, respectively. This information is useful to define the operation bandwidths of MeNB and SeNB from traffic statistics.

Performance Evaluation

In this section, we present some simulation results for the access methods described above. The scenario consists of one FDD MeNB and one TDD SeNB, which operates at the unused portion of the FDD spectrum. Both are separated by 100 m. The complete scenario and simulation assumptions are included in Table 1.

The following techniques are evaluated:

- FMA (in-band FMA and CC-based FMA): FMA with frequency multiplexing between FDD MeNB and TDD SeNB, including ACI
- TMA (in-band TMA): TMA with time multiplexing between MeNB and SeNB, including 1 guard subframe
- Only MeNB: one MeNB operating in the paired spectrum

Data traffic follows a Poisson process with packet arrival rates λ_{DL} and λ_{UL} (in packets per second) for DL and UL transmissions, respectively. Two different traffic asymmetries are considered:

- $\lambda_{DL} = 10 \times \lambda_{UL}$: DL traffic is 10 times larger than UL traffic, leaving unused resources in the FDD-UL band.
- $\lambda_{UL} = 10 \times \lambda_{DL}$: UL traffic is 10 times larger than DL traffic, leaving unused resources in the FDD-DL band.

For each case, two packet arrival rates are evaluated: λ_{DL} = {1, 1.5} packets/s and $\lambda_{DL} = 10 \times \lambda_{UL}$, and λ_{UL} = {1, 1.5} packets/s for $\lambda_{UL} = 10 \times \lambda_{DL}$ (emulating medium and high traffic loads).

By taking into account the values of the packet arrival rates for each traffic asymmetry, the measured RU at the FDD MeNB (according to Eq. 1) is less than 30 percent in the FDD-UL band (under $\lambda_{DL} = 10 \times \lambda_{UL}$ traffic asymmetry) and less than 14 percent in the FDD-DL band (under $\lambda_{UL} = 10 \times \lambda_{DL}$ traffic asymmetry). The TDD SeNB operates in the unused resources in the following way,

FMA: In FDD-UL band reuse, SeNB operates over 7 MHz bandwidth with a TDD duplexing ratio 7:3 (7 DL and 3 UL subframes). In FDD-DL band reuse, the SeNB operates over 8.7 MHz bandwidth with a TDD duplexing ratio 1:9.

TMA: It is just evaluated for FDD-UL band reuse due to the reasons provided earlier. The SeNB operates over the whole spectrum but only 60 percent of the time, with a TDD duplexing ratio 4:2.

Two metrics are considered as performance indicators:

- The RU presented in Eq. 1, measured as the average number of resource blocks needed for the communication over the total number of available resource blocks
- The mean of the user throughput (UT) in megabits per second

The following figures show the mean UT vs. the RU of the MeNB and SeNB for the two considered traffic asymmetries ($\lambda_{DL} = 10 \times \lambda_{UL}$ in Fig. 3 and $\lambda_{UL} = 10 \times \lambda_{DL}$ in Fig. 4). Several conclusions can be drawn.

When there are unused resources in the FDD-UL band and the ones in the FDD-DL band are nearly saturated (Fig. 3):

•An improvement of "UT DL" at both the MeNB and the SeNB is obtained with FMA and

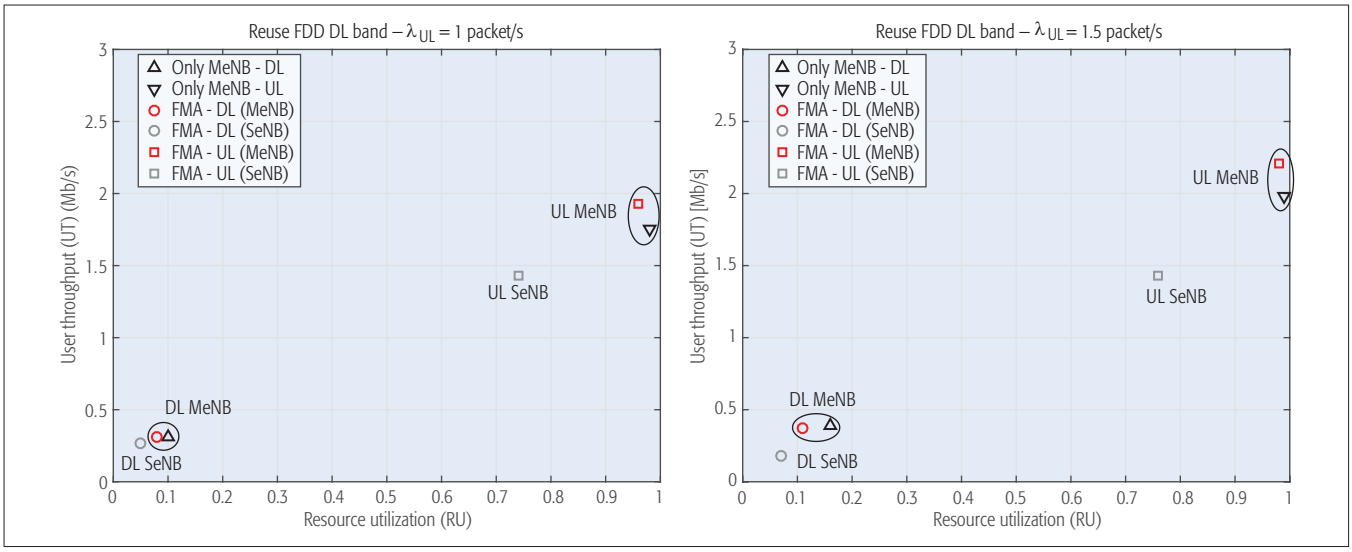

Figure 4. Mean UT (in megabits per second) vs. RU for FDD-DL band reuse with d = 100 m. Left, λ_{UL} = 1; right, λ_{UL} = 1.5 (packets/s); traffic asymmetry: $\lambda_{UL} = 10 \times \lambda_{DL}$.

TMA as compared to the case of having only the MeNB active since more DL traffic can be served. The relative gains are up to 7 percent in terms of "UT DL" both at MeNB and SeNB for λ_{DL} = 1 packets/s. Nevertheless, when the traffic increases up to λ_{DL} = 1.5 packets/s, "UT DL" relative gains significantly improve: 20 percent at MeNB and 49 percent at SeNB (Fig. 3 right).

•For medium traffic loads (λ_{DL} = 1 packets/s in Fig. 3, left), thanks to the reuse of the licensed bandwidth for UL with FMA and TMA, the MeNB DL reduces its RU compared to the case of having only the MeNB active. Thus, the MeNB is decongested.

•The ACI (SeNB to MeNB and SUE to MeNB) in FMA imposes a lower transmission rate, so more resources are needed ("RU DL SeNB," "RU UL MeNB," and "RU UL SeNB"). For example, the "RU DL SeNB" is 0.37 with FMA and 0.28 with TMA for λ_{DL} = 1.5. The UT becomes lightly degraded with FMA at high traffic loads due to the activity of the SeNB. It generates ACI toward the MeNB, so the transmission rate of MUEs is lowered, but active longer, and negatively impacting the SeNB when it is in UL ("UT UL" in Fig. 3, right). This effect is avoided with TMA. "UT UL" at MeNB and SeNB is 0.36 and 0.35 Mb/s with FMA, respectively, while it takes values of 0.38 and 0.37 Mb/s at MeNB and SeNB, respectively, under TMA.

On the other hand, when the underutilized resources are in the FDD-DL band, results are affected by two main impairments:

DL interference from neighboring MeNBs (co-channel external interference): This can significantly degrade the system performance since the MeNB interferes with the SeNB transmissions in both DL and UL

ACI at SeNB: This is important because transmitters in the neighbouring band are MeNBs. From Fig. 4 we infer that:

- An improvement of the "UT UL MeNB" is obtained compared to the case of having only the MeNB active, as more UL traffic can be served. The relative gains are: 10 percent for λ_{DL} = 1 packets/s and 12 percent for λ_{DL} = 1.5 packets/s.
- For the two packet arrival rates, the "UT UL SeNB" is significantly degraded with respect to the "UT UL MeNB." The relative losses are of 19 percent for λ_{DL} = 1 packets/s and of 28 percent for λ_{DL} = 1.5 packets/s. This is due to co-channel external interference from neighboring MeNBs, which significantly impacts on "UT UL SeNB." In addition, the effect of co-channel external interference also degrades "UT DL SeNB."
- There is nearly no impact on "UT DL MeNB" when reusing the FDD-DL band and on "UT UL MeNB" when reusing the FDD-UL band, and similar values are obtained as in the case of having only the MeNB, since all the traffic that arrives at the system is being served. According to Eq. 2, low resource utilization denotes a small average queue size in the system.

Current Limitations

In spite of the promising benefits shown by the flexible duplexing concept, its implementation in the short term must face several challenges from the point of view of regulation and LTE standardization.

Regulation

Radio spectrum regulators define which type of transmissions are allowed on different parts of the spectrum. In general, the FDD-UL spectrum can be employed by mobile stations or end users, but not by base stations. In this regard, a survey to different regulators revealed that at least in the United States, the flexible duplex concept is allowed in the band 1719–1755 MHz, [12]. On the other hand, in Europe (ECC PT1) and Japan (ARIB), flexible use of UL and DL for FDD bands is not allowed. Nevertheless, the use of SeNBs with a transmit power equivalent to the maximum allowed in the UL by regulation would satisfy all technical requirements. In this regard, ECC PT1 is open to receiving new results about the benefits of the flexible use of the licensed spectrum bandwidth.

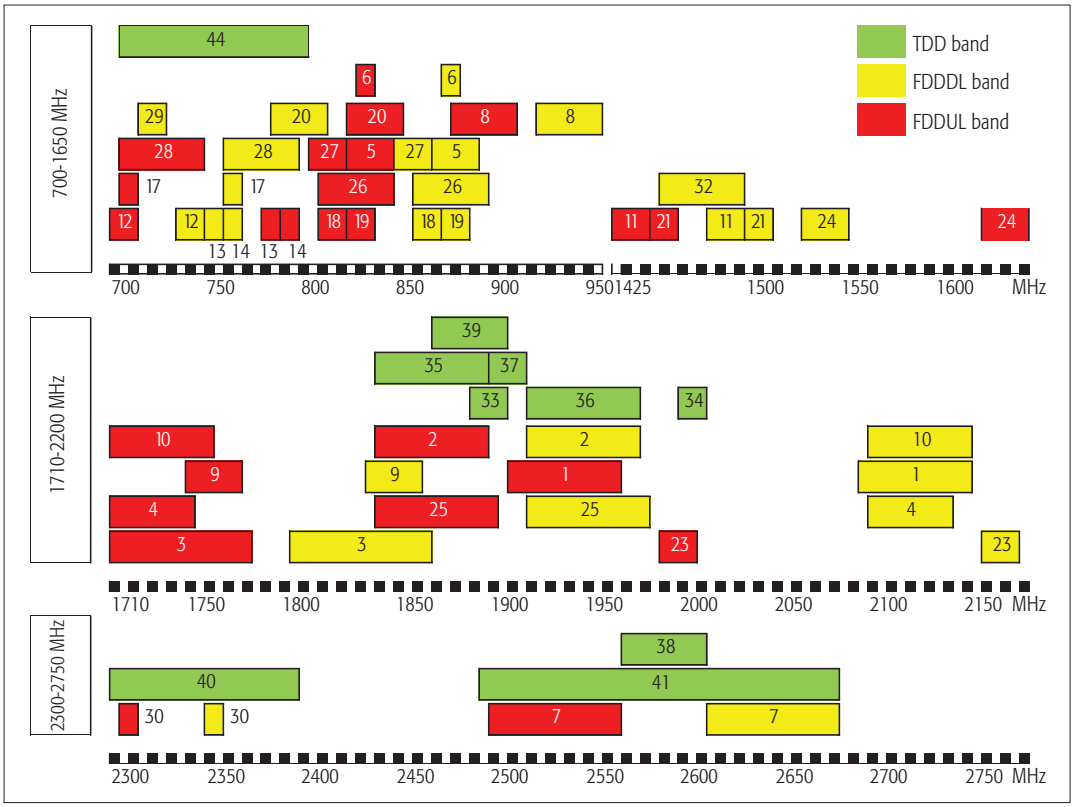

Figure 5. E-UTRA FDD (red and yellow) and TDD (green) operating bands. Proposed flexible duplexing methods would be standard-compliant in those bands where TDD and FDD overlap.

The flexible duplexing concept might be considered in Europe and Hong Kong because band 7 is actually in use, so that TDD SeNBs operating in band 41 could reuse the underutilized spectrum. The UEs in the area should just measure the control channels of MeNB and SeNB, and decide their association to one of them.

LTE Compliance

The following aspects limit the definition of LTE standard-compliant procedures for the implementation of the flexible duplexing concept.

Operating Band Definition: LTE defines a set of operating bands along with its use: FDD (bands 1–32) or TDD (bands 33–44) [10]. From the comparison of UL-DL FDD bands and TDD bands (Fig. 5), we can observe that, with the current definition, not all FDD bands could be reused by TDD systems. However, the flexible duplexing concept might be considered in Europe and Hong Kong because band 7 is actually in use, so TDD SeNBs operating in band 41 could reuse the underutilized spectrum. The UEs in the area should just measure the control channels of MeNB and SeNB, and decide their association with one of them.

PUCCH in the FDD UL Band: The current design of LTE FDD systems places PUCCH in the resource blocks located at the edge of the band [13]. Those resources are devoted to transmit system information of UE. This system constraint limits the flexibility for seizing the unused time-frequency resources in the in-band TMA approach as described above.

Frame Structure in FDD DL Band: Even in situations of low DL traffic, the information necessary to operate the system (to be more specific, synchronization signals and system information and paging) needs to be transmitted by the FDD-DL cell so that terminals can detect and connect to a cell. In FDD, the subframes where such information is provided are subframes 0, 4, 5, and 9 within an LTE frame composed of 10 subframes. Therefore, these subframes must be used by the FDD cell and cannot be used by a TDD cell. Similarly, the TDD SeNB also needs to transmit the information necessary to operate the system. In TDD, such information is provided through subframes 0, 1 (special subframe), 5, and 6 (special subframe) within a frame of 10 subframes. It turns out that only two subframes could be reused for data transmission with the in-band TMA approach. On the other hand, the SeNB should be deployed in a narrower bandwidth placed at one side of the band for the in-band FMA approach because all synchronization signals of FDD-DL occupy the central part of the band.

Carrier Aggregation: Currently, the 3GPP standard does not allow tackling situations where the traffic asymmetry is higher in the FDD-UL band because, by definition, the CCs in the UL are smaller than in the DL [13]. This feature does not allow extending the concept explained earlier to reuse the underutilized FDD spectrum in the FDD-DL band.

Future Challenges and Research

Future research encompasses three different lines:

- The management of the operation of multiple SeNBs in the unused spectrum for the multiplexing methods explained above
- The study of opportunistic multiple access (OMA) methods
- The design of advanced interference cancellation receivers for OMA

One of these research lines presents important technical challenges that require further investigation for the deployment of the flexible duplex concept.

Multiple SeNBs: Resources among SeNBs

The benefits of the flexible duplexing concept can be enlarged when combined with the deployment of multiple SeNBs and the application of interference management techniques, but its actual implementability is also tied to the limitations imposed by regulations and standards.

should be distributed by taking into account the actual traffic demand per SeNB. The RU presented earlier, in addition to allowing the identification of the required spectrum, is a useful metric to derive resource provisioning schemes in a multi-cell scenario (e.g., multiple TDD SeNBs that exploit the spectrum released by the FDD MeNB). Efficient resource provisioning should distribute resources among cells in a balanced way while trying to avoid very different occupancies. In this sense, a meaningful optimization criterion is the minimization of the maximum RU among cells so that resources are fairly distributed and more resources are given to those cells with larger traffic loads and/or those cells experiencing greater delays. For example, long-term graph coloring-based resource provisioning schemes are presented in [15] with the objective of optimizing the RU factors of multiple TDD SeNBs when either orthogonal access is assumed or reuse of resources among non-interfering SeNBs is considered.

OMA in OFDM-Based Systems: Allowing an SeNB to work in part of the spectrum of the MeNB in an opportunistically way (i.e., non-orthogonal) demands tight synchronization (timing offset adjustments) so as to maintain the orthogonality of orthogonal frequency-division multiplexing (OFDM) carriers at SeNB and MeNB. For example, in the heterogeneous scenario described in [14], which consists of one MeNB and one SeNB both sharing band and duplexing, it was shown that UEs have to advance their UL transmissions (pre-compensation) not only by taking into account the propagation delay with the SeNB but also by considering the propagation delay with the MeNB and the received frame boundary. The cyclic prefix in OFDM systems combats this issue, in addition to maintaining the orthogonality among subcarriers. However, in the flexible duplexing concept, SeNBs work in TDD while MeNB is FDD-UL, which means that synchronization is more challenging because the SeNB DL transmission should be pre-compensated by taking into account the neighboring FDD MeNBs and the TDD SeNB working in UL.

Advanced Interference Canceller in OMA: If signal time offsets are properly pre-compensated to avoid time misalignments with an FDD MeNB, OMA can be used by the TDD SeNBs. The interference received by the MeNB from the transmitting SeNBs might be very high due to line of sight condition in the MeNB-SeNB link. A successive interference canceller (SIC) would alleviate the effect of interference, but the performance of SIC will depend on the received interference level (distance between MeNB and SeNB) and the bit rate selected by the SeNB (DL transmissions to a SUE).

Conclusions

The flexible duplexing concept allows for improvements in the system efficiency of pair-based systems by using TDD SeNBs. When allowing a TDD SeNB to operate in the underutilized FDD-UL band, we have observed that the DL user throughput is improved and the congestion of the MeNB is reduced. The effect of ACI for the FMA schemes becomes relevant when:

- The SeNB is close to the MeNB.
- The SeNB transmits with high power.
- The activity of UEs and the SeNB is high.

On the other hand, when a TDD SeNB reuses the underutilized FDD-DL band, the potential gains are reduced because of the external interference coming from neighboring MeNBs that are transmitting DL signals. The benefits of the flexible duplexing concept can be enlarged when combined with the deployment of multiple SeNBs and the application of interference management techniques, but its actual implementability is also tied to the limitations imposed by regulations and standards.

Acknowledgments

This work has been supported by Huawei Technologies Co., Ltd. and by the "Ministerio de Economía, Industria y Competitividad" of the Spanish Government and European Regional Development Fund (ERDF) TEC2013-41315-R DISNET and "Agencia Estatal de Investigación" and ERDF TEC2016-77148-C2-1-R 5G&B-RUNNER-UPC

References

[1] DIGITALEUROPE: Call for Timely Harmonisation of the 1452–1492 MHz and 2300–2400 MHz Bands to Support Delivery of the EU Radio Spectrum Policy Programme Objectives, Brussels, 21 Feb. 2012; http://digitaleurop.org.
[2] Qualcomm: Wireless Broadband Future and Challenges, Samena Telecommunications Council Convergence to Casablanca, Oct. 27, 2010.
[3] 3GPP RP-140062, "Motivation of New SI proposal: Evolving LTE with Flexible Duplex for Traffic Adaptation," Mar. 2014; http://www.3gpp.org.
[4] W. Lei, Z. Mingyu, and W. Rongui, "Evolving LTE with Flexible Duplex," *Proc. IEEE GLOBECOM Wksps.)*, Dec. 2013.
[5] 3GPP TR 36.828, "Further Enhancements to LTE Time Division Duplex for Downlink-Uplink Interference Management and Traffic Adaptation," Release 11, June 2012; http://www.3gpp.org.
[6] ECC Report 119, "Coexistence between Mobile Systems in the Frequency Band at the FDD/TDD Boundary", June 2008; http://www.erodocdb.dk.
[7] "Real Wireless, Low-Power Shared Access to Spectrum for Mobile Broadband," Ofcom project MC/073, Mar. 2011; http://ofcom.org.uk.
[8] R. Berangi *et al.*, "TDD Cognitive Radio Femtocell Network (CRFN)," *Proc. IEEE Int'l. Symp. Personal, Indoor Mobile Radio Commun.*, Sept. 2011.
[9] 3GPP. R1-134295, "FDD-TDD CA/Dual Connectivity Solution Exploiting Traffic Asymmetry in Duplex-Neutral Bands," Oct. 2013; http://www.3gpp.org.
[10] 3GPP TS 36.101, "User Equipment Radio Transmission and Reception," Release 13, July 2015;. http://www.3gpp.org.
[11] Kumar, D. Manjunath, and J. Kuri, *Communication Networking: An Analytical Approach*, Morgan Kaufmann, Elsevier, 2004.
[12] 3GPP TR 36.882, "Study on Regulatory Aspects for Flexible Duplex for E-UTRAN," Release 13, June 2015.
[13] E. Dahlman, S. Parkvall, and J. Sköld, *4G LTE/LTE-Advanced for Mobile Broadband*, Academic Press, 2011.
[14] N. Himayat *et al.*, "Synchronization Uplink Transmissions from Femto AMS," IEEE C802.16m-09/3075r2, Jan. 2010; http://ieee802.org/16/tgm
[15] S. Lagen *et al.*, "Long-Term Provisioning of Radio Resources Based on their Utilization in Dense OFDMA Networks," *Proc. IEEE Int'l. Symp. Personal, Indoor Mobile Radio Commun.*, Valencia, Spain, Sept. 2016.

Biographies

Adrián Agustín received his M.S. degrees in telecommunication (00) and electronic engineering (2002) and his Ph.D. degree (2008) from Universitat Politcnica de Catalunya (UPC), Barcelona, Spain. After graduation, he joined Indra-Espacio, and he joined the Signal Theory and Communications Department at UPC in 2002, becoming a research associate in 2008. He has contributed to 3GPP RAN1 standardization and participated in several European Commission funded projects. His research

interests include wireless interference management, IoT, and backhaul management.

Sandra Lagen received her Telecommunication Engineering degree (2011), M.S. degree (2013), and Ph.D. degree (2016) from UPC. From 2012 to 2017 she was a research assistant in the Signal Processing and Communications group at UPC. Since May 2017 she has been a researcher in the Mobile Networks Department at Centre Tecnolgic de Telecomunicacions de Catalunya (CTTC). Her research interests include signal processing, wireless communications, MIMO, and optimization theory.

Josep Vidal is a professor at UPC. He has coordinated five collaborative EC-funded projects in different areas of MIMO relay communications, self-organization, cooperative transmission, and heterogeneous networks, areas where he has authored +170 journal and conference papers. He served as an Associate Editor of *IEEE Transactions on Signal Processing* and is a member of the IEEE ComSoc Signal Processing for Communications and Electronics Technical Committee

Olga Muñoz received her M.S. degree (1993) and Ph.D. degree (1998) in electrical engineering from UPC. Since 2001, she has been an associate professor at UPC. She has participated in European Commission funded projects such as FIREWORKS, ROCKET, FREEDOM, TROPIC, and TUCAN3G on the topics of cooperative communications, relaying systems, coordinated radio resource allocation, and heterogeneous backhauling. She has published over 50 papers in books, international conferences, and journals in the areas of signal processing and communications

Antonio Pascual-Iserte received his Electrical Engineering (2000) and Ph.D. (2005) degrees from UPC. After graduation he was with Retevision R&D, working in the DVB-T and T-DAB networks in Spain. In 2001 he joined the Department of Signal Theory and Communications (UPC), where he is now an associate professor. His research focuses on optimization applied to communications. He has worked in several research projects and published papers in conferences and journals.

Zhiheng Guo received his Bachelor's degrees in telecommunication (2002) and Ph.D. degree 20(07) from Beijing University of Posts and Telecommunications, China. After graduation he joined Ericsson Research, becoming a senior research engineer in 2011, and he joined Huawei in 2014. He has participated in many 3GPP RAN1 meetings for telecommunication technology standards. His research interests focus on physical layers and include flexible duplex, LTE-NR coexistence, multi-user superposition, integrated access and backhaul, and so on.

Ronghui Wen received her Bachelor's degree in telecommunication (2004) and Ph.D. degree (2010) from Harbin Institute Technology University, China. After graduation she joined Huawei in 2010. Her research interests focus on physical layers and include flexible duplex, URLLC, eIoT, and more.

ACCEPTED FROM OPEN CALL

Routing in FRET-Based Nanonetworks

Pawel Kulakowski, Kamil Solarczyk, and Krzysztof Wojcik

The authors focus on nanocommunications via FRET, which was found to be a technique with a very high signal propagation speed, and discuss how to route signals through nanonetworks. They introduce five new routing mechanisms, based on biological properties of specific molecules, and experimentally validate one of these mechanisms.

ABSTRACT

Nanocommunications, understood as communications between nanoscale devices, is commonly regarded as a technology essential for cooperation of large groups of nanomachines and thus crucial for development of the whole area of nanotechnology. While solutions for point-to-point nanocommunications have already been proposed, larger networks cannot function properly without routing. In this article we focus on nanocommunications via FRET, which was found to be a technique with a very high signal propagation speed, and discuss how to route signals through nanonetworks. We introduce five new routing mechanisms, based on biological properties of specific molecules. We experimentally validate one of these mechanisms. Finally, we analyze open issues showing the technical challenges for signal transmission and routing in FRET-based nanocommunications.

INTRODUCTION

With enormous growth and progress in the whole area of nanotechnology, the demand for communication between nanomachines has arisen naturally. Nanomachines, called also nanodevices or nanonodes, may be of both biological and artificial origin. Biological ones (e.g., proteins or whole cells) are building blocks of living organisms. Currently, scientists are working on artificial nanomachines, constructing structures like molecular switches, ratchets, and motors based on their biological counterparts [1]. The application field for nanomachines is extremely wide, extending from environment monitoring, industrial manufacturing, and building labs-on-a-chip to an enormous number of applications in medicine, like drug delivery, diagnostics, tissue regeneration, and surgical operations. The limited size of nanomachines, however, restrict their functions and capabilities, so the ability to perform more complex actions relies on cooperation in larger groups, that is, nanonetworks. In this sense, nanocommunications will play a similar role for nanotechnology as telecommunications currently plays for electronics: it will enable nanomachines to work together. Nanodevices will communicate each other in order to:

- Exchange and forward gathered information
- Coordinate their joint actions
- Interface with other biological and artificial systems

This is why efficient communication between nanodevices is a crucial challenge to be met in order to develop future nanonetworks and the whole area of nanotechnology.

The dimensions of nanoscale devices are comparable to single molecules, so electromagnetic communication based on miniaturized transceivers and antennas is hard to be directly applied. Instead, numerous biologically inspired communication mechanisms have been proposed. Calcium signaling, molecular and catalytic nanomotors, pheromones propagation, and information transfer using bacteria as carriers have already been studied, but these mechanisms are slow and characterized by large propagation delays: the encoded data travels at a speed of several dozens of micrometers per second at maximum [2]. Compared to these techniques, a mechanism based on the phenomenon of Förster resonance energy transfer (FRET) is more promising, especially bearing in mind its low propagation delay [3–6]. FRET is a process in which a molecule, known as a donor, is able to non-radiatively (i.e., without releasing a photon) transfer its energy to another molecule, called an acceptor. For this energy transfer to occur, the donor must be in an excited state, the acceptor has to be located in close proximity to the donor, and finally, the donor and acceptor molecules must be spectrally matched, that is, the donor emission and acceptor absorption spectra (the frequency ranges of emitted/absorbed EM spectrum) should overlap (Fig. 1a). The delay of the energy transfer is usually no more than 20 ns. The donor can be excited in various ways, for example, by photon absorption, another FRET process, or a chemical reaction (bioluminescence). The latter process, called bioluminescence resonance energy transfer (BRET), may be suitable for nanocommunications, as it does not require any external energy source (e.g., light) for donor excitation.

The FRET and BRET processes are very distance-dependent: their efficiency decreases with the sixth power of the donor-acceptor separation. For each two types of molecules, the, so-called Förster distance can be experimentally measured; it is a separation where the respective FRET efficiency is equal to 50 percent. Förster distances for typical donor-acceptor pairs range from 3 to 9 nm. This effectively places a limit on FRET transmission distances to about 10 nm. Recent works [7, 8] have shown that the FRET efficiency and transmission range can be additionally increased when using multiple donors at the transmitter side and multiple acceptors at the receiver side of the nanocommunication channel. Such a technique is called multiple-input multiple-output

Digital Object Identifier:
10.1109/MCOM.2017.1600760

Pawel Kulakowski is with the AGH University of Science and Technology; Kamil Solarczyk is with Jagiellonian University; Krzysztof Wojcik is with Jagiellonian University Medical College.

0163-6804/17/$25.00 © 2017 IEEE

(MIMO)-FRET, following the idea of multiantenna systems well known in wireless communication. The most common molecules that can be used as FRET donors and acceptors are known as fluorophores. Their absorption and emission spectra lie in the visible light range (i.e., from 380 to 760 nm). It should be noted, however, that FRET as a phenomenon is not limited to this range of electromagnetic spectrum, but may also occur for shorter and longer wavelengths. When thinking about interactions of future nanodevices, one may envisage fluorophores performing functions of nanoantennas attached to some larger nanostructures, for example, proteins like antibodies, myosins, kynesins, dyneins, and adenosine triphosphate (ATP) synthases, working as nanomachines (Fig. 1b).

The content of the rest of this article is as follows. The next section introduces five new routing mechanisms available for FRET-based nanonetworks. Then, in order to prove its feasibility, the first of these mechanisms is validated experimentally, and the results of the laboratory experiment are given. Later, the most important open issues in nanorouting are discussed, and finally, the article is concluded.

Routing Mechanisms

While point-to-point communication via FRET, with its parameters like bit error rate, channel capacity, throughput, and signal propagation delay, was analyzed and measured in previous papers [5–8], the main motivation for this article is to go a step further and propose possible routing mechanisms in FRET-based nanonetworks. When thinking about fully operational nanonetworks, there must be not only point-to-point links, but also multihop connections. Nanonodes should be able to forward signals; moreover, the nodes should be able to make routing decisions, that is, decide where to forward the signals. Assuming a nanonetwork communicates among its nodes via FRET, we propose five new techniques of signal routing. They are:

- Proteins with multiple different fluorescent dyes
- Photoswitchable fluorophores
- Quenchers
- Proteins with changeable shape
- ATP synthases

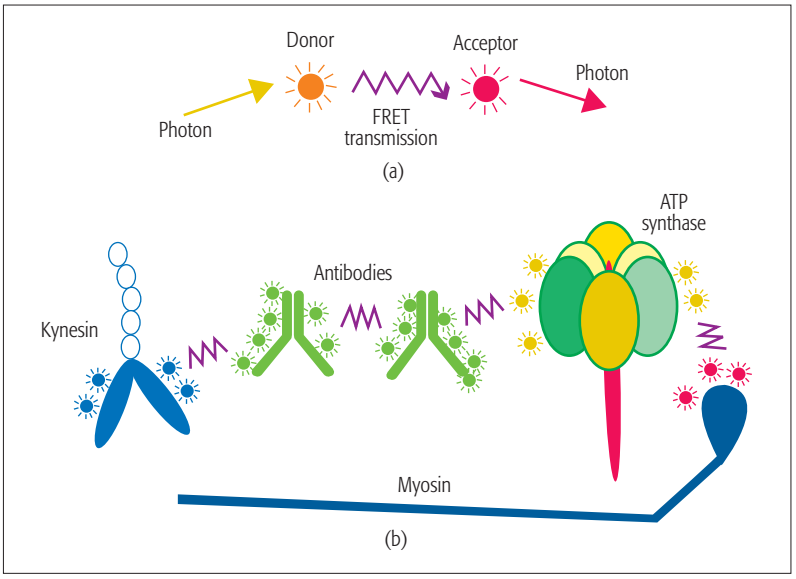

Figure 1. a) The FRET process after excitation of the donor molecule by an external photon. The excitation energy is passed non-radiatively to the acceptor molecule and then can be released as another photon; b) examples of proteins that may work as nanomachines: antibodies, kynesin, ATP synthase, and myosin molecules. Each of them has some fluorophores (marked as small circles with short rays) attached to it. The fluorophores serve as nanoantennas transmitting and receiving signals via FRET.

They are described in the five following subsections and compared parametrically in Table 1. Until now, to our best knowledge there is still no research on these techniques in the area of nanocommunications. There are, however, numerous papers already published in life sciences where it was shown that these five mechanisms could provide suitable means for nanorouting. In the next section, we also report our laboratory experiments proving the efficiency of the first of the proposed routing techniques.

Proteins with a Few Different Fluorescent Dyes

Attaching fluorescent dyes (fluoro (e.g., proteins) is common in biology experiments in order to localize these proteins in their environment. For the purpose of nanocommunications, the fluorophores can serve as nanoantennas, which are able to communicate with each other via FRET.

Nanorouter type	Routing mechanism	Number of links outgoing from the nanorouter	What switches the nanorouter	Switching time
Proteins with multiple different fluorescent dyes	Not needed	2–4	–	–
Photoswitchable fluorophores	Change of absorption and emission spectra	2	Light	1–60 seconds
Quenchers	Fluorescence quenching	2–6	Temperature, pH	1–10 seconds
Proteins with changeable shape	Protein conformation change	2	Adding a ligand	10–100 milliseconds
ATP synthases	ATP synthase rotation	2–10	Voltage, ATP	2–50 milliseconds

Table 1. Routing techniques for FRET-based nanonetworks.

Each fluorophore is characterized by its emission spectrum. The FRET signals emitted by this fluorophore can be received only by neighboring fluorophores with absorption spectra matching this emission spectrum.

Many types of fluorescent dyes can be attached to a single protein (nanomachine) in various ways, usually utilizing active chemical groups present in the dye or in the protein. We can regulate not only the type of attached particles, but also their number, limited by the protein dimension. When the attached molecules have different absorption and emission spectra, they can receive and transmit signals of different frequencies, thus performing as nanorouters. This scenario resembles a situation where a wireless device is equipped with two or more antennas working in different frequency bands. Such a nanorouter is different from other solutions presented below in the sense that it does not need to be switched: it passes upcoming signals depending on their wavelength (Fig. 2a). The known fluorescent dyes that can be used to construct such a nanorouter are fluorophores from the Alexa Fluor (AF), DyLight, and Atto families.

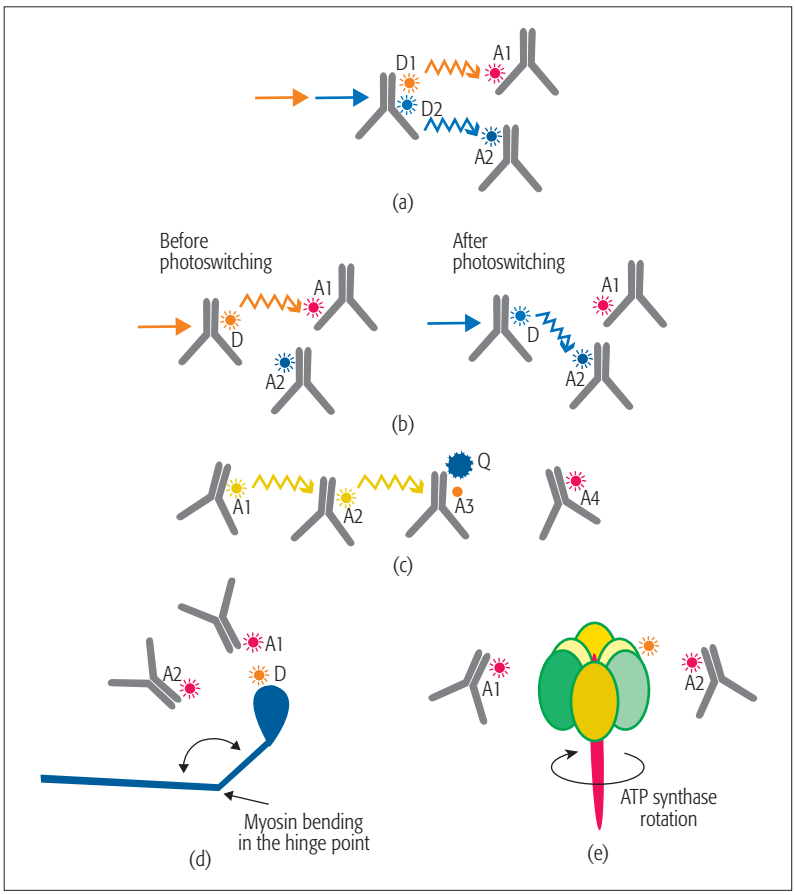

Figure 2. Routing mechanisms for FRET-based nanonetworks: a) a transmitting nanomachine has two fluorophores working as nanoantennas. As the fluorophores have different absorption and emission spectra, signals can be transmitted via D1 to A1 or via D2 to A2; b) fluorophore D may be photoswitched and then its absorption and emission spectra change. Before photoswitching, it sends signals to A1; after photoswitching, it communicates with A2; c) FRET signal may be passed in multiple hops from A1 to A3, but an active quencher Q blocks its further propagation. After deactivation of the quencher, the signal might also reach A4; d) myosin bends, and fluorophore D moves away from A1 and closer to A2. Transmission D-A1 is interrupted, and transmission D-A2 may be initiated; e) ATP synthase rotates along its own axis. Fluorophore D attached to the ATP synthase may periodically communicate with A1 and A2.

Photoswitchable Fluorophores

Photoswitchable fluorophores, also called photoswitches, have been known in life sciences for years, but have recently gained significant attention as they allowed the development of various super-resolution imaging techniques [9]. What is important for nanorouting purposes is that these fluorophores can be reversibly switched between two states by external electromagnetic impulses of a specific power and frequency. In each state, a reversible photoswitch has different absorption and emission characteristics, that is, it is able to receive and transmit signals, also via FRET, in different frequency bands. A typical shift of the fluorophore emission or absorption spectrum is about 100 nm [10]. Therefore, photoswitchable fluorophores can be thought of as tunable antennas that, depending on external control, are able to get FRET signals from and forward them to the chosen neighboring nanonodes. Certain connections in nanonetworks can be open and shut in this way (Fig. 2b). The known reversible photoswitchable fluorophores are, among others, Cy5 molecules, and numerous AF and Atto dyes.

Quenchers

A quencher works when in close proximity to a fluorescent molecule called a reporter; the quencher suppresses fluorescence of the reporter and in turn blocks all its signals. The reporter can resume its transmissions when the quencher is removed or deactivated. This latter effect is usually realized via changes in the environment of the fluorophore (temperature, pH). For communication applications, quenchers may be used in order to temporarily block FRET signals propagating in undesired directions, thus blocking the chosen links between nanonodes (Fig. 2c). Multiple quenchers are commercially available (e.g., TAMRA and Dabcyl), but the most suitable are so-called dark quenchers, such as Black Hole Quenchers, which emit neither photons nor FRET signals.

Proteins with Changeable Shape

Interaction of two or more proteins with each other or a protein with a small molecule (e.g., an ion) may lead to the binding of these entities. This binding sometimes causes a conformational change of the protein, that is, the change of its shape. Some of the conformations are especially attractive for nanorouting purposes, as they change the separations between the fluorophores attached to the protein and other fluorophores located on neighboring nanomachines. The FRET efficiency decreases with the sixth power of the separation between the fluorophores, so via control of this separation, one may in fact open or shut down chosen nanolinks. A good example is myosin, which is a protein that changes its shape (i.e., bends its part) after binding a Ca^{2+} ion binding,. As a result, fluorophores attached to myosin may be put sufficiently close to other fluorophores attached to another nanomachine, just in order to perform a successful FRET transmission (Fig. 2d). In this example, the Ca^{2+} ion may be imagined

as a trigger that activates the switch, enabling a specific nanolink. It should be emphasized that the conformational change of myosin shape is reversible after the Ca^{2+} ions are removed.

ATP Synthase

Finally, a very promising protein is an ATP synthase (a motor enzyme creating ATP), which is able to rotate around its own axis (Fig. 2e). This effect can be used to periodically communicate with some other nanomachines located nearby, either broadcasting the same signal to all the neighbors or dividing the signal among them in a time-division multiple access (TDMA) manner. The rotation of the ATP synthase may be initiated by providing electrical voltage or adding ATP.

To sum up, all these routing mechanisms, based on the properties of specific molecules, enable the control of signals in nanonetworks in numerous different manners, mimicking solutions well known from telecommunication networks. The signals can be routed to the chosen receivers, broadcasted or de-multiplexed, and certain links may be switched on and off. All these techniques may be used in the same nanonetwork, offering a great variety of solutions, depending on the particular need.

Experimental Studies on Nanorouting

In order to present not only theoretical analysis, but also practical experience on nanorouting, we experimentally validated the first of the proposed routing mechanisms (proteins with a few different fluorescent dyes). Communication over nanometer distances occurring between devices of nanometer dimensions cannot be observed in classical telecommunication laboratories. Therefore, we performed experiments in a biophysical lab using a confocal microscope equipped with a fluorescence lifetime imaging microscopy (FLIM) module, which enabled measurement efficiency of the FRET process between nanonodes. First, we constructed a network composed of three nanonodes, as in Fig. 2a. The nodes were proteins: Immunoglobulin G with fluorophores attached working as FRET nanoantennas. In particular, we had one nanonode (rabbit anti-mouse IgG 610-451-C46 produced by Rockland) operating as a nanotransmitter/nanorouter with two types of fluorophores attached to it: Atto 425[1] and AF 680. These two types of fluorophores are characterized by significantly different emission spectra and thus, if excited, may emit FRET signals that can be reached by different receivers. Apart from the nanorouter, we had two other nodes working as nanoreceivers: one with AF 488 (goat anti-rabbit IgG ab150077 produced by Abcam) and the second one with AF 750 (goat anti-rabbit IgG A-21039 produced by Sigma). In this nanonetwork the emission spectra of Atto 425 matches the absorption spectra of AF 488, and similarly, the emission spectra of AF 680 matches the absorption spectra of AF 750 (Fig. 3). This means that we had two nanolinks created: the first one between Atto 425 and AF 488 (from the nanorouter to the first nanoreceiver), and the second one between AF680 and AF750 (from the nanorouter to the second nanoreceiver).

As we could not exactly control the movements of the proteins in the laboratory samples, the nanonodes (Immunoglobulin G molecules) were fixed to a large chain of DNA sequence and other proteins in order to keep the transmitting and receiving nanoantennas (fluorophores) about 815 nm from each other (Fig. 4). The fixing was done only for measurement purposes; in a real scenario, the nanonodes could float freely in their environment (e.g., in a human cell).

The laboratory experiments were performed on fixed HeLa 21-4 cells (a common cell line used in laboratory research). In the nucleus of each cell there were about $3 \cdot 10^6$ nanonetworks (the nanorouter and both nanoreceivers) attached to the DNA sequence. We analyzed the data from 57 nuclei, so we had over 10^8 nanonetworks examined, which was a statistically credible number. The cells containing the nanonetworks were exposed to short pulses of laser light. The wavelength of light utilized to excite the fluorophores attached to the nanotransmitter, Atto 425 and AF 680, were 405 and 640 nm, respectively. During each laser pulse, only a few fluorophores (donors) were excited in the whole sample. It means that, despite having multiple fluorophores at the nanorouter and at the nanoreceivers (Fig. 4), we investigated a scenario of a single donor and multiple acceptors, which was like a single-input multi-output (SIMO) channel. In real applications, all the donor molecules may be excited at once using full MIMO communication, which can additionally increase the FRET efficiency compared to the values shown in this article [7, 8].

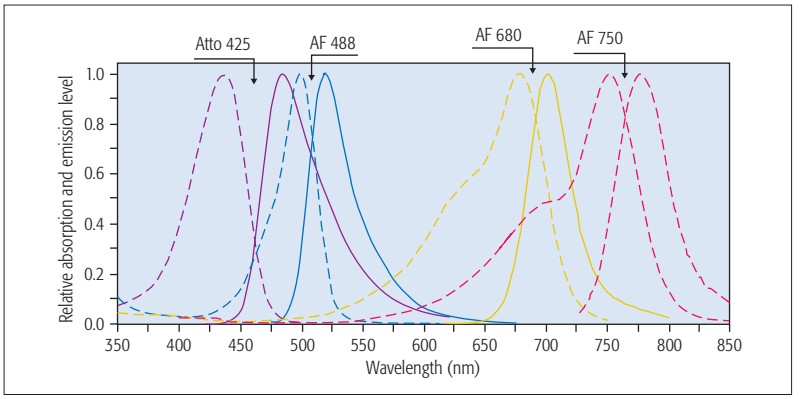

Figure 3. Absorption (dashed lines) and emission (solid lines) spectra of the fluorophores used in the experiment [11].

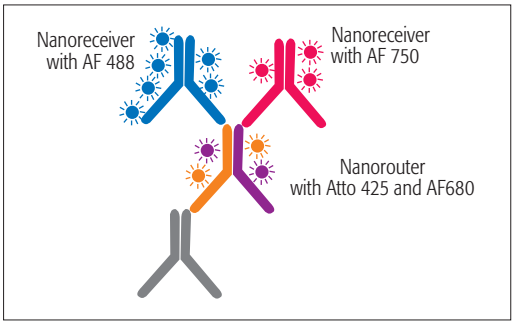

Figure 4. The experimental scenario: a chain of antibodies where one of them acts as a nanorouter with Atto 425 and AF 680 dyes (violet and orange ones), and two others are nano receivers with AF 488 (blue) fluorophores and AF 750 (red) ones.

[1] The numbers of fluorophores, 425, 488, 680, and 750 indicate the wavelengths of their maximum absorption, given in nanometers. Their emission spectra are usually shifted by a few dozen nanometers into higher wavelengths.

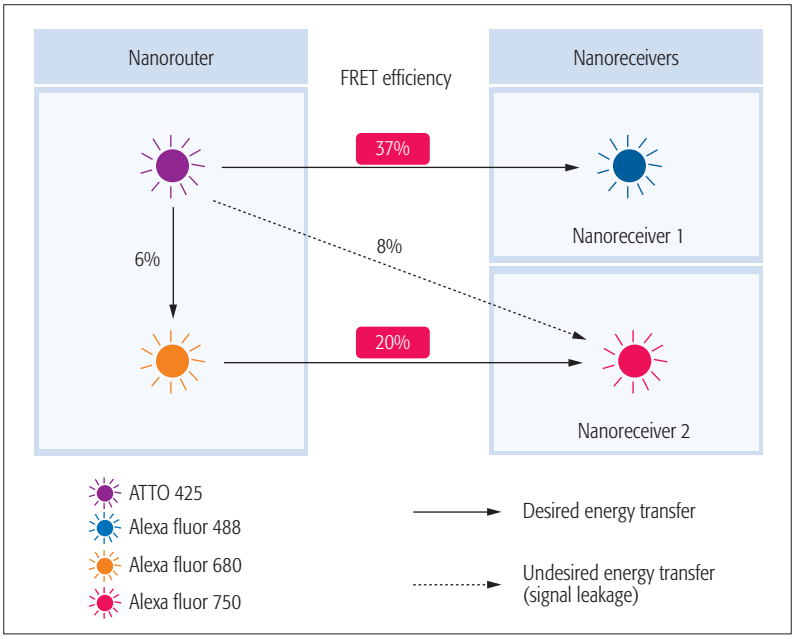

Figure 5. The FRET efficiency values characterizing the signal transmission between the nanorouter and the nanoreceivers.

The main purpose of the experiment was to measure how much of the FRET signal could be passed from the nanorouter to the chosen nanoreceiver. When a donor in an excited state is close to an acceptor, and the emission and absorption spectra of these molecules are matched, energy may be passed from the donor to the acceptor via FRET. In a laboratory sample, FRET can be observed indirectly as a decrease of donor fluorescence intensity (or an increase of acceptor fluorescence intensity). Thus, measurements of fluorescence intensity can give information regarding the efficiency of the FRET transfer. An alternative way to obtain the value of FRET efficiency is to monitor the fluorescence lifetime of the donor. Fluorescence lifetime is defined as the mean time the donor spends in its excited state before returning to the ground state. Since FRET represents an additional way for the donor to depopulate its excited state, the lifetime of the donor decreases in the presence of an acceptor. Comparing the lifetime of the donor (nanorouter) in the absence and presence of the acceptor (nanoreceiver) allows calculation of the FRET efficiency [12]. As the latter approach is less prone to photobleaching effects than intensity-based methods, we have chosen it for our experiments.

The measurements were conducted using a Leica TCS SP5 II SMD confocal microscope (Leica Microsystems GmbH) integrated with FCS/FLIM module from PicoQuant GmbH. The FLIM module enabled combining fluorescence lifetime measurements to optical microscopy. The lifetime measurements were performed in time correlated single photon counting mode. The nanonetworks located in the nuclei of HeLa cells were exposed to a pulse laser exciting in turns the molecules of Atto 425 and AF 680. The detection of photons emitted by these molecules was observed with single photon avalanche diode detectors. The laser repetition rate was 40 MHz, while the laser power was adjusted to obtain a photon counting rate of 200–300 kCounts/s. The acquisition time for each field of view (image) was set to 1 minute. The analysis of the results was done with SymPhoTime II software using tail fitting of the lifetime functions. Lifetime decays of Atto 425 and AF 680 were fitted with a two-exponential function, that is, a sum of two exponential functions with different amplitudes and lifetimes. For the calculation of FRET efficiency, the average lifetime was used. The goodness of fit was estimated based on the weighted residuals and the chi squared value.

The results of the performed experiments are summarized in Fig. 5. The FRET values in the pairs Atto 425 → AF 488 and AF 680 → AF 750 show how efficient the transmission between the nanorouter and the nanoreceivers might be. During the measurements, due to the limited sensitivity of the detectors, it was not possible to excite multiple nanorouter donors at once. Thus, in practice, we have measured the FRET efficiencies for the SIMO case, which is much less effective than the full MIMO-FRET. This is why the FRET values are rather low for telecommunication purposes. Fortunately, using MIMO-FRET and suitable coding techniques with enough redundancy, FRET-based nanocommunications can guarantee bit error rate at the level of 10^{-6} [8]. The results also raise an issue of undesired signal leakages: when the donor Atto 425 at the nanorouter is excited, 37 percent of the signal is transferred to the AF 488 at the first nanoreceiver, but at the same time 8 percent passes to AF 750 at the second nanoreceiver. The same problem occurred during additional control measurements where we observed another signal leakage between the donors at the nanorouter (Atto 425 and AF 680) at the level of 6 percent. Both leakages are caused by the partial and undesired overlap of the Atto 425 emission spectrum and AF 680/750 absorption spectra (Fig. 3).

Open Issues

While the article proposes five new routing techniques for nanonetworks and reports the experimental validation for one of them, it is clear that this research area is at the beginning of its road, and many questions remain unanswered. Below we present the most important open issues.

Routed Signal Leakages

As reported in the previous section, the separation between the links outgoing from the nanorouter is far from perfect: it may happen that the signal from the nanorouter is absorbed at the wrong nanoreceiver. The reason is that the receiving molecules are characterized by broad absorption spectra and may absorb signals emitted at wavelengths even 300 nm shorter than their maximum absorption wavelength (again Fig. 3). The most straightforward solution of this problem is to increase the wavelength separation between the pairs of fluorophores used for communication purposes. Because of the fluorescence spectroscopy requirements, the fluorophores currently available in the market are working in the wavelength range of 300–900 nm, which is in general the visible and partially infrared light range. The molecules with emission/absorption spectra at higher and lower wavelengths would, at least partially, solve this problem. Another possibility would be to use fluorophores with very narrow emission spectra, for example, BODIPY dyes.

Switching Time

While the propagation speed for FRET signals is quite high, the routing itself may not be so fast. The routing mechanisms proposed in this article require, in some cases, change in the shape of a molecule (e.g., myosin) or in its properties (photoswitches, quenchers). The time of this change (i.e., the nanorouter switching time) may even reach several dozens of seconds (e.g., in the case of fluorophores photoswitching) [13], which is too long for most telecommunication purposes. There is a clear need to shorten the switching time. It may be done by finding more sensitive molecules that could act as nanorouters or carefully choosing the switching conditions (pH of the environment, radiation intensity, frequency, etc).

Transmitter Excitation

In the current FRET experiments, the transmitters (donors) are excited by an external laser impulse. While it may be the case of some nanonetworking applications, in most of them it should be the nanomachine itself that initiates the communication. In such a situation, BRET may be applied instead of FRET. In BRET, the excitation energy comes not from an external source, but from a chemical reaction taking place near the donors. Thus, when a nanomachine is going to send a signal, it induces the chemical reaction, which, in turn, excites the donors. The manner in which the nanomachine initiates this reaction still remains an open issue.

Signal Storage

The FRET phenomenon enables communication between the fluorophores, but these molecules cannot hold the received signal. Before releasing its energy (emitting a photon, via FRET or other processes) and returning to the ground state, a molecule spends no longer than a few dozen nanoseconds in the excited state. This creates a problem for the considered nanonetworks, as the nanonodes should have some buffers where the signals could be stored until the nodes decide to resend them further. Currently, there is no clear solution for creating these buffers. One option could be using phosphorescent molecules that are able to store energy even nine orders of magnitude longer than fluorophores [14]. Phosphorescence is a phenomenon similar to fluorescence, but its excitation state lifetime can reach minutes, as releasing energy by a phosphorescent molecule is related to a transition, which is, according to quantum mechanics, forbidden. Storing signals in phosphorescent molecules could be even more important when using nanorouters with long switching times, as mentioned two subsections earlier.

Nanomachine Movements

There is also a general issue of nanomachine movement control. While it is a little out of the scope of this article, it should be noted that relative nanomachine positions and their separation are crucial for the efficiency of FRET-based communication as it depends on sixth power of the transmitter-receiver (donor-acceptor) distance. Thus, manipulating the nanomachines, for example, putting them within closer range for the time of communication, will result in much higher efficiency of information transfer. It can be done using proteins with changeable shapes, as indicated in the section about routing mechanisms, but there are some other options possible, like bridging of free monomers of IgE or JAK receptors [15].

Conclusion

The phenomenon of FRET seems to be a very promising solution for nanocommunications, with a very high propagation speed compared to other techniques proposed so far. While point-to-point transmission with FRET has already been investigated, both theoretically and experimentally, this article expands the topic by analyzing possibilities of routing FRET signals through nanonetworks. The new proposed routing techniques rely on physical and biological properties of specific molecule types, mainly proteins and fluorophores. With current development of biotechnology, these molecules may easily be manipulated and put together in large static or mobile structures. The routing techniques enable maintaining end-to-end communication in whole nanonetworks, which is a crucial step to have future tiny nanodevices working together on complex endeavors.

The new proposed routing techniques rely on physical and biological properties of specific molecule types, mainly proteins and fluorophores. With current development of biotechnology, these molecules may be easily manipulated and put together in large, static or mobile, structures.

Acknowledgment

The authors would like to thank P. Cholda, A. Jajszczyk, A. Lason, and R. Wojcik for reviewing the manuscript and their constructive comments. The work was performed under contract 11.11.230.018 and is also a result of the research project No. DEC-2013/11/N/NZ6/02003 financed by the National Science Center. The confocal microscope was purchased through an EU structural funds grant BMZ no. POIG.02.01.00-12-064/08.

References

[1] S. Erbas-Cakmak *et al.*, "Artificial Molecular Machines," *Chemical Reviews 2015*, vol. 115 no. 18, pp. 10,081–10,206.
[2] T. Nakano, A. W. Eckford, and T. Haraguchi, *Molecular Communications*, Cambridge Univ. Press, 2013.
[3] T. Förster, "Zwischenmolekulare energiewanderung und fluoreszenz," *Annalen der Physik*, vol. 437, nos. 1/2, 1948, pp. 55–75.
[4] L. Parcerisa and I. F. Akyildiz, "Molecular Communication Options for Long Range Nanonetworks," *Computer Networks*, vol. 53, 2009, pp. 2753–66.
[5] M. Kuscu and O. B. Akan, "A Physical Channel Model and Analysis for Nanoscale Communications with Förster Resonance Energy Transfer (FRET)," *IEEE Trans. Nanotechnology*, vol. 11, no. 1, Jan. 2012, pp. 200–07.
[6] M. Kuscu and O. B. Akan, "Multi-Step FRET-Based Long-Range Nanoscale Communication Channel," *IEEE JSAC*, vol. 31, no. 12, Dec. 2013, pp. 715–25.
[7] K. Wojcik, K. Solarczyk, and P. Kulakowski, "Measurements on MIMO-FRET Nanonetworks Based on Alexa Fluor Dyes," *IEEE Trans. Nanotechnology*, vol. 14, no. 3, May 2015, pp. 531–39.
[8] K. Solarczyk, K. Wojcik, and P. Kulakowski, "Nanocommunication via FRET with DyLight Dyes Using Multiple Donors and Acceptors," *IEEE Trans. NanoBioscience*, vol. 15, no. 3, Apr. 2016, pp. 275–83.
[9] E. Betzig *et al.*, "Imaging Intracellular Fluorescent Proteins at Nanometer Resolution," *Science*, vol. 313, 2006, pp. 1642–45.
[10] S. van de Linde *et al.*, "Direct Stochastic Optical Reconstruction Microscopy with Standard Fluorescent Probes," *Nature Protocols* 6, 2011, pp. 991–1009.
[11] Chroma Technology Corp., Spectra Viewer; https://www.chroma.com/spectra-viewer, accessed 26 July 2016.
[12] J. R. Lakowicz, *Principles of Fluorescence Spectroscopy*, 3rd ed., Springer, 2006.
[13] S. van de Linde *et al.*, "Photoinduced formation of reversible dye radicals and their impact on super-resolution imaging," *Photochem. Photobiol. Sci.*, 10, 2011, pp. 499–506.

[14] D. Wasserberg, S. C. J. Meskers, and R. A. J. Janssen, "Phosphorescent Resonant Energy Transfer between Iridium Complexes," *J. Physical Chemistry A*, vol. 111, no. 8, 2007, pp. 1381–88.
[15] B. Alberts, J. Wilson, and T. Hunt, *Molecular Biology of the Cell*, Garland Science, 2008.

Biographies

Pawel Kulakowski (kulakowski@kt.agh.edu.pl) received a Ph.D. in telecommunications from AGH University of Science and Technology, Krakow, Poland, in 2007. Currently he is working there as an assistant professor. He also worked for a few years in Spain, as a visiting postdoctoral researcher and professor at the Technical University of Cartagena, the University of Girona, the University of Castilla-La Mancha, and the University of Seville. He has co-authored about 30 scientific papers, in journals, conferences, and as technical reports. He has been involved in numerous research projects, especially European COST Actions: COST2100, IC1004, and CA15104 IRACON, focusing on topics of wireless sensor networks, indoor localization, and wireless communications in general. His current research interests include molecular communications and nanonetworks. He was recognized with several scientific distinctions, including three awards for his conference papers and a scholarship for young outstanding researchers.

Kamil Solarczyk (kj.solarczyk@uj.edu.pl) received M.Sc. and Ph.D. degrees in biophysics from Jagiellonian University, Krakow, Poland, in 2010 and 2016, respectively. He is a postdoctoral researcher in the Department of Cell Biophysics, Faculty of Biochemistry, Biophysics and Biotechnology, Jagiellonian University. His research interests include DNA repair processes, chromatin architecture, and nanoscale communications.

Krzysztof Wojcik (krzysztof.wojcik@uj.edu.pl) received his M.Sc. and Ph.D. degrees in biophysics from Jagiellonian University in 2003 and 2015, respectively, and an M.D. from Jagiellonian University Medical College (JUMC), Krakow, Poland, in 2007. He was an assistant in the Division of Cell Biophysics Faculty of Biochemistry, Biophysics and Biotechnology, Jagiellonian University (2007-2014). He is an assistant at the Allergy and Immunology Clinic in II Chair of Internal Medicine JUMC. His research interests include confocal microscopy techniques and their applications in autoantibodies research, as well as the use of fluorescent probes in nanocommunications.

NATIONAL INSTRUMENTS

INNOVATE FASTER

WITH FIELD-DEPLOYED 5G PROOF-OF-CONCEPT SYSTEMS

In the race to design next-generation wireless technologies, research teams must rely on platforms and tools that accelerate productivity. And whether you're working in the lab or deploying solutions for field trial test, NI software defined radio hardware and LabVIEW Communications software can help you innovate faster and build 5G proof-of-concept systems to demonstrate new technologies first.

Accelerate your innovation at ni.com/5g.

IEEE Collabratec™

Bright Minds. Bright Ideas.

Introducing IEEE Collabratec™

The premier networking and collaboration site for technology professionals around the world.

IEEE Collabratec is a new, integrated online community where IEEE members, researchers, authors, and technology professionals with similar fields of interest can **network** and **collaborate**, as well as **create** and manage content.

Featuring a suite of powerful online networking and collaboration tools, IEEE Collabratec allows you to connect according to geographic location, technical interests, or career pursuits.

You can also create and share a professional identity that showcases key accomplishments and participate in groups focused around mutual interests, actively learning from and contributing to knowledgeable communities. All in one place!

Network.
Collaborate.
Create.

Learn about IEEE Collabratec at
ieeecollabratec.org

IEEE

Fraunhofer
Heinrich Hertz Institute

FRAUNHOFER INSTITUTE FOR TELECOMMUNICATIONS, HEINRICH HERTZ INSTITUTE, HHI

INDOOR Li-Fi

Make your Light Smarter

Fraunhofer HHI presents the next generation Gigabit Visible Light Communication (VLC) modules for wireless Internet access via light. Outstanding features are the smaller form factor, lower energy consumption, enhanced coverage and multi-user access. A standard Ethernet interface allows easy network integration. The new modules are immediately available for industrial prototyping and field tests.

Facts

- Use of standard high-power LEDs
- No interference with existing Wi-Fi networks
- Multi-user access possible
- Peak data rate 1 Gbps
- Small form factor

also available in other colours

Photonic Networks and Systems

Fraunhofer Heinrich Hertz Institute
Einsteinufer 37 | 10587 Berlin
Germany

products-pn@hhi.fraunhofer.de
www.hhi.fraunhofer.de/vlc