IEEE
COMMUNICATIONS
MAGAZINE

March 2016, Vol. 54, No. 3

• Critical Communications and Public Safety Networks

• Radio Communications

• Network Testing

IEEE ComSoc
IEEE Communications Society

A Publication of the IEEE Communications Society

www.comsoc.org

IEEE Communications Magazine

March 2016, Vol. 54, No. 3

- Critical Communications and Public Safety Networks
- Radio Communications
- Network Testing

IEEE
IEEE ComSoc
IEEE Communications Society
A Publication of the IEEE Communications Society
www.comsoc.org

IEEE COMMUNICATIONS MAGAZINE

MARCH 2016, vol. 54, no. 3
www.comsoc.org/commag

4 THE PRESIDENT'S PAGE

6 CONFERENCE CALENDAR

7 GLOBAL COMMUNICATIONS NEWSLETTER

128 ADVERTISERS' INDEX

CRITICAL COMMUNICATIONS AND PUBLIC SAFETY NETWORKS, PART 1: STANDARDS, SPECTRUM POLICY, AND ECONOMICS

GUEST EDITORS: MEHMET ULEMA, ALAN KAPLAN, KEVIN LU, NIRANTH AMOGH, AND BARCIN KOZBE

12 GUEST EDITORIAL

14 TOWARD MOVING PUBLIC SAFETY NETWORKS
Romain Favraud, Apostolos Apostolaras, Navid Nikaein, and Thanasis Korakis

22 PERVASIVE SPECTRUM SHARING FOR PUBLIC SAFETY COMMUNICATIONS
Murat Yuksel, İsmail Güvenç, Walid Saad, and Naim Kapucu

30 NEXT GENERATION PUBLIC SAFETY NETWORKS: A SPECTRUM SHARING APPROACH
Munawwar M. Sohul, Miao Yao, Xiaofu Ma, Eyosias Y. Imana, Vuk Marojevic, and Jeffrey H. Reed

38 INDIRECT RETURNS AND USE OF NPV IN ECONOMIC VIABILITY MODELING OF CRITICAL COMMUNICATIONS NETWORKS
Natalia Boliari

RADIO COMMUNICATIONS: COMPONENTS, SYSTEMS, AND NETWORKS

SERIES EDITORS: AMITABH MISHRA AND TOM ALEXANDER

44 SERIES EDITORIAL

46 LOAD MODULATED ARRAYS: A LOW-COMPLEXITY ANTENNA
Mohammad A. Sedaghat, Vlasis I. Barousis, Ralf R. Müller, and Constantinos B. Papadias

53 A FLEXIBLE 5G FRAME STRUCTURE DESIGN FOR FREQUENCY-DIVISION DUPLEX CASES
Klaus I. Pedersen, Gilberto Berardinelli, Frank Frederiksen, Preben Mogensen, and Agnieszka Szufarska

60 DIGITAL PREDISTORTION FOR MITIGATING SPURIOUS EMISSIONS IN SPECTRALLY AGILE RADIOS
Mahmoud Abdelaziz, Zhu Fu, Lauri Anttila, Alexander M. Wyglinski, and Mikko Valkama

NETWORK TESTING

SERIES EDITORS: YING-DAR LIN, ERICA JOHNSON, AND IRENA ATOV

70 SERIES EDITORIAL

71 UNVEILING NETWORK AND SERVICE PERFORMANCE DEGRADATION IN THE WILD WITH MPLANE
Pedro Casas, Pierdomenico Fiadino, Sarah Wassermann, Stefano Traverso, Alessandro D'Alconzo, Edion Tego, Francesco Matera, and Marco Mellia

80 ACCURATE AND AFFORDABLE PACKET-TRAIN TESTING SYSTEMS FOR MULTI-GIGABIT-PER-SECOND NETWORKS
Mario Ruiz, Javier Ramos, Gustavo Sutter, Jorge E. López de Vergara, Sergio López-Buedo, and Javier Aracil

88 OBSERVING REAL SMARTPHONE APPLICATIONS OVER MULTIPATH TCP
Quentin De Coninck, Matthieu Baerts, Benjamin Hesmans, and Olivier Bonaventure

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: $27 per year print subscription. $16 per year digital subscription. Non-member print subscription: $400. Single copy price is $25.

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*, 3 Park Avenue, 17th Floor, New York, NY 10016.

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 sumitted 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. For further information contact Zoran Zvonar, Associate Editor-in-Chief (zoran.zvonar@mediatek.com). All submissions will be peer reviewed.

Accepted from Open Call

94 **Toward a Net Neutrality Debate that Conforms to the 2010s**
Patrick Maillé, Gwendal Simon, and Bruno Tuffin

100 **Highly Dynamic Spectrum Management within Licensed Shared Access Regulatory Framework**
Aleksei Ponomarenko-Timofeev, Alexander Pyattaev, Sergey Andreev, Yevgeni Koucheryavy, Markus Mueck, and Ingolf Karls

110 **Why to Decouple the Uplink and Downlink in Cellular Networks and How To Do It**
Federico Boccardi, Jeffrey Andrews, Hisham Elshaer, Mischa Dohler, Stefan Parkvall, Petar Popovski, and Sarabjot Singh

118 **The ITU-T's New G.fast Standard Brings DSL into the Gigabit Era**
Vladimir Oksman, Rainer Strobel, Xiang Wang, Dong Wei, Rami Verbin, Richard Goodson, and Massimo Sorbara

Currently Scheduled Topics

Topic	Issue Date	Manuscript Due Date
Next Generation 911	November 2016	March 15, 2016
Integrated Communications, Control, and Computing Technologies for Enabling Autonomous Smart Grid	December 2016	April 1, 2016
New Waveforms and Multiple Access Methods for 5G Networks	November 2016	April 2, 2016
Impact of Next-Generation Mobile Technologies on IoT-Cloud Convergence	January 2017	April 15, 2016
Practical Perspectives on IoT in 5G Networks: From Theory to Industrial Challenges and Business Opportunities	February 2017	May 1, 2016

www.comsoc.org/commag/call-for-papers

TOPICS PLANNED FOR THE APRIL ISSUE

Critical Communications and Public Safety Networks

Integrated Circuits for Communications

FROM THE OPEN CALL QUEUE

Network Function Virtualization in 5G

Adaptive and Cognitive Communication Architecture for Next-generation PPDR Systems

Scalable and Mobile Context Data Retrieval and Distribution for Community Response Heterogeneous Wireless Networks

A Survey on Rapidly Deployable Solutions for Post-Disaster Networks

A Unifying Perspective on Proximity-based Cellular-Assisted Mobile Social Networking

Millimeter-wave Gbps Broadband Evolution towards 5G: Fixed Access and Backhaul

Reducing complexity of Virtual Machine networking

A Scalable Architecture for Handling Control Plane Failures in 5G Heterogeneous Networks

Buffer Sizing in Wireless Networks: Challenges, Solutions, and Opportunities

Multi-Comm-Core Architecture for Terabit/s Wireless

IEEE'S STRATEGY FOR THE YEAR 2030

Harvey Freeman

Celia Desmond

Last month we talked about ComSoc's strategy and goals for the next few years. This month the focus is on IEEE's strategy preparing for the year 2030. The IEEE Board of Directors (BoD) in January 2015 felt that IEEE needed a governance structure that can effectively handle the increasing strategic complexity of a changing and dynamic world. This would require a diverse, efficient, and effective board that represents the members and listens to the voices of members, the technical professional community, and the public, in order to make informed decisions. In addition, the BoD wants to strengthen the role of the member in IEEE governance. Since then there have been many meetings, discussions, proposals, arguments for and against, presentations, and the like, to develop constitution and organization changes that will prepare IEEE for 2030.

I have asked Celia Desmond, currently the Division III (ComSoc) Director and member of the IEEE 2030 Committee, to bring us up to date on the current status of this work. Celia is Lead Project Manager for Echologics, as well as President of World Class Telecommunications, providing telecommunications management training. In IEEE she has held numerous senior management positions, including Project Director for Certification in Wireless Engineering Technology, Director and Secretary of IEEE, IEEE Vice President – Technical Activities, President of the IEEE Communications Society, President of IEEE Canada and its Division Director, and Vice President of the Technology and Engineering Management Society. Celia holds MSc. Engineering and B.Sc. Mathematics & Psychology degrees, Ontario Teaching Certificate, and PMP certification.

There has been a lot of activity at the senior levels of IEEE lately on the 2030 strategy. Two recent endeavors include: (1) a proposed amendment to the IEEE Constitution; and (2) consideration of a proposal to change the structure of the upper levels of governance within the organization. This column provides a brief summary of some of this activity, with a pointer to the source of the most up to date details for those who want more information. The IEEE Board of Directors, at its meeting in November 2015, endorsed the Constitutional Amendment for submission to the members and then issued the following statement to accompany the election materials later this year:

> *The Board of Directors proposes revisions to the IEEE Constitution and recommends each IEEE member vote FOR the amendment. If adopted, these modifications improve the members' voice in governing IEEE and allow future changes to the organizational structure to better respond to the demands of a complex and changing world.*

Specifically the changes:

- Provide members with the possibility of an increased role selecting the Board of Directors, allowing directors to be elected by the full eligible voting membership of IEEE.
- Add language encouraging a diverse Board of Directors.
- Add the IEEE executive director, the most senior IEEE staff executive, as a non-voting member of the Board to participate from inception in setting the strategic direction of IEEE.
- Separate the role of an IEEE delegate from an IEEE director, so that directors need not also be delegates.
- Separate the requirement that corporate officers must also be directors. This will allow corporate officers as currently defined to serve in important leadership positions other than on the Board of Directors.
- Establish a new role for IEEE delegates, who are members of the IEEE Assembly, to recommend and consult with the Board on revisions to IEEE Bylaws.

For most members, this requires some explanation. At this time, when someone is elected to the Board of Directors of IEEE, except for the case of the IEEE President Elect, who then becomes President the following year, that person is elected by a subset of the members of IEEE. Individuals are elected as Delegates to the IEEE Assembly, and by virtue of that election they become Directors on the IEEE Board of Directors.

- Division Directors are associated with one or more Societies/Councils. Each Division Delegate/Director is elected as a Delegate to the IEEE Assembly by the members of the Societies/Councils that the Delegate represents. Division Delegates also are Directors on the IEEE Board of Directors and sit on the Technical Activities Board (TAB) as members. On TAB each Director represents his or her constituency.
- Region Delegates/Directors are elected by the members of their geographic Regions as a Delegate to the IEEE Assembly by the members of the Region that the Delegate represents. Region Delegates also are Directors on the IEEE Board of Directors and sit on the Member and Geographic Activities Board (MGAB) as members. On MGAB each represents his or her constituency.
- The Vice President Technical Activities is elected by all members who are members of any Society/Council.
- The President IEEE Standards and President IEEE-USA are elected by members of their constituencies.
- The Vice President Educational Activities, Vice President Member and Geographic Activities, Vice President Publications Services and Products, Treasurer, and Secretary are elected by the IEEE Assembly, which consists of all 20 Division and Region Delegates mentioned above, plus the President, Past President, and President Elect.

The 23 Delegates, as mentioned above, are Directors on the IEEE Board of Directors, and in this capacity they no longer represent any constituency. They must think and act for the good of IEEE.

The amendment to the Constitution will be sent to all

IEEE members with the 2016 Election ballot. Please exercise your right to vote when you receive your ballot and make the informed decision on these proposed changes that you feel is best for the organization. For more information about the Constitutional Amendment visit:

http://www.ieee.org/about/corporate/election/2016_constitutional_amendment.html

Please do vote on this when you receive your ballot.

The second effort, consideration of a proposal to change the governance structure at the highest levels of IEEE, is a work in progress, and the proposed changes remain under discussion. As feedback comes in, the committee working on this effort takes that feedback into account and makes any required changes to the proposal. The current proposal, its complete history, and all the comments received in building it are available to all members at:

http://www.ieee.org/about/corporate/ieeein2030_archive_m.html

As you will see from the material posted on this site, today the business of the senior levels of IEEE is handled by two bodies: the Board of Directors and the Assembly. The current proposal is to add one additional body to this structure, and distribute the work somewhat differently, with the new Enterprise Board handling the operational management, leaving the Board of Directors free to focus on the strategic directions and issues. In the current proposal all Board of Directors members would be elected by the full membership. Since this is a work in progress, the proposal is evolving over time. It is useful to check the site periodically to understand the progress.

OMBUDSMAN

ComSoc Bylaws Article 3.8.10

"The Ombudsman shall be the first point of contact for reporting a dispute or complaint related to Society activities and/or volunteers. The Ombudsman will investigate, provide direction to the appropriate IEEE resources if necessary, and/or otherwise help settle these disputes at an appropriate level within the Society."

IEEE Communications Society Ombudsman
c/o Executive Director
3 Park Avenue, 17 Floor
New York, NY 10017, USA
ombudsman@comsoc.org

www@comsoc.org "About Us" (bottom of page)

UPDATED ON THE COMMUNICATIONS SOCIETY'S WEB SITE

www.comsoc.org/conferences

2016

MARCH

DRCN 2016 — 12th Int'l. Workshop on Design of Reliable Communication Networks, 14–17 March
Paris, France
https://drcn2016.lip6.fr/

ICBDSC 2016 — 3rd MEC Int'l. Conference on Big Data and Smart City, 15–16 Mar.
Muscat, Oman
http://www.mec.edu.om/conf2016/index.html

OFC 2016 — Optical Fiber Conference, 20–24 Mar.
Anaheim, CA
http://www.ofcconference.org/en-us/home/

IEEE ISPLC 2016 — 2016 IEEE Int'l.Symposium on Power Line Communications and Its Applications, 21–23 Mar.
Bottrop, Germany
http://www.ieee-isplc.org/

IEEE CogSIMA 2016 — IEEE Int'l. Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support, 21–25 Mar.
San Diego, CA
http://www.cogsima2016.org/

WD 2016 — Wireless Days 2016, 23–25 Mar.
Toulouse, France
http://wd2015.sciencesconf.org/

APRIL

IEEE WCNC 2016 — IEEE Wireless Communications and Networking Conference, 3–6 Apr.
Doha, Qatar
http://wcnc2016.ieee-wcnc.org/

IEEE INFOCOM 2016 — IEEE Int'l. Conference on Computer Communications, 10–15 April
San Francisco, CA
http://infocom2016.ieee-infocom.org/

WTS 2016 — Wireless Telecommunications Symposium, 18–20 Apr.
London, U.K.
http://www.cpp.edu/~wtsi/

FRUCT18 2016 — 18th Conference of Open Innovations Association FRUCT and Seminar on Information Security and Protection of Information Technology, 18–22 Apr.
St. Petersburg, Russia
http://fruct.org/cfp

IEEE/IFIP NOMS 2016 — IEEE/IFIP Network Operations and Management Symposium, 25–29 Apr.
Istanbul, Turkey
http://noms2016.ieee-noms.org/

MAY

IEEE CQR 2016 — IEEE Int'l. Communications Quality and Reliability Workshop, 9–12 May
Stevenson, WA
http://www.ieee-cqr.org/

ONDM 2016 — Int'l. Conference on Optical Network Design and Modeling, 9–12 May
Cartagena, Spain
http://ondm2016.upct.es/index.php

IEEE CTW 2016 — IEEE Communication Theory Workshop, 15–18 May
Nafplio, Greece
http://www.ieee-ctw.org/

ICT 2016 — Int'l. Conference on Telecommunications, 16–18 May
Thessaloniki, Greece
http://ict-2016.org/

IEEE ICC 2016 — IEEE International Conference on Communications, 23–27 May
Kuala Lampur, Malaysia
http://icc2016.ieee-icc.org/

JUNE

IEEE BlackSeaCom 2016 — 4th Int'l. Black Sea Conference on Communications and Networking, 6–9 June
Varna, Bulgaria
http://www.ieee-blackseacom.org/

IEEE NETSOFT — IEEE Conference on Network Softwarization, 6–10 June
Seoul, Korea
http://sites.ieee.org/netsoft/

IEEE LANMAN 2016 — 22nd IEEE Workshop on Local & Metropolitan Area Networks, 13–15 June
Rome, Italy
http://www.ieee-lanman.org/

IEEE HPSR 2016 — IEEE 17th Int'l. Conference on High Performance Switching and Routing, 14–17 June
Yokohama, Japan
http://www.ieee-hpsr.org/

IEEE IWQOS — IEEE Int'l. Symposium on Quality and Service, 20–21 June
Beijing, China
http://www.dongliangxie.com/

MED-HOC-NET — Mediterranean Ad Hoc Networking Workshop, 20–22 June
Vilanova I la Geltru, Spain
http://craax.upc.edu/medhocnet2016/

EUCNC 2016 — European Conference on Networks and Communications, 27–30 June
Athens, Greece
http://eucnc.eu/

IEEE ISCC — Int'l. Symposium on Computers and Communications, 26–30 June
Messina, Italy
http://iscc2016.unime.it/

IEEE SECON — 2016 IEEE Int'l. Conference on Sensing, Communication and Networking, 27–30 June
London, U.K.
http://secon2016.ieee-secon.org/

JULY

ICUFN 2016 — Int'l. Conference on Ubiquitous and Future Networks, 5–8 July
Vienna, Austria
http://www.icufn.org/main/

CITS 2016 — Int'l. Conference on Computer, Information and Telecommunication Systems, 6–8 July
Kunming, China
http://atc.udg.edu/CITS2016/

–Communications Society portfolio events appear in bold colored print.
–Communications Society technically co-sponsored conferences appear in black italic print.
–Individuals with information about upcoming conferences, Calls for Papers, meeting announcements, and meeting reports should send this information to: IEEE Communications Society, 3 Park Avenue, 17th Floor, New York, NY 10016; e-mail: p.oneill@comsoc.org; fax: + (212) 705-8996. Items submitted for publication will be included on a space-available basis.

GCN GLOBAL COMMUNICATIONS NEWSLETTER

March 2016
ISSN 2374-1082

DISTINGUISHED LECTURER TOUR

Distinguished Lecturer Tour of Prof. Tilman Wolf in Central America

By Carlos Eugenio Martínez-Cruz, El Salvador Chapter Chair

Prof. Tilman Wolf was invited by the Guatemala, El Salvador, and Panama ComSoc Chapters to conduct a Lecturer Tour in November 2015. All three sections asked for two lectures. The first two lectures were held at the World Heritage Site of Antigua Guatemala. This city that features Spanish Baroque-influenced architecture was the tour's first stop. Tilman Wolf presented two of his best demanded talks: 'Economic Principles for Future Internet Architecture' and 'Attacks and Hardware Defenses for Network Infrastructure.' Both lectures supported the join effort made by Guatemala's IEEE Computer Society and Communications Society.

The second two lectures were given on November 11 in San Salvador, El Salvador's capital. Prof. Wolf was awakened much earlier than 6:00 am, and without having any breakfast he traveled to the place were the lecture was arranged. The first talk started at 7:00 am, on the topic 'Hardware Defenses for Network Infrastructure.' In this presentation, Prof. Wolf discussed his research group's recent work that illustrates the challenges in providing security in Internet routers. He provided an example that showed how vulnerable packet processors can be attacked through the data plane of the network. Using hardware monitors, he showed how his team has developed an effective defense mechanism against such attacks. Finally, it was explained how this work is also applicable to securing general-purpose embedded processing systems.

After finishing his presentation, Prof. Wolf traveled to the University of El Salvador. There the IEEE student branch called for a big meeting. More than 100 students attended the lecture. Again, the talk 'Hardware Defenses for Network Infrastructure' was presented. After this presentation Prof. Wolf interacted with the students. He explained to them how the functionality of routers inside the Internet continues to grow, as does their complex protocol processing operations for content adaptation, security, and network management. Most of the students were sophomores in the electrical engineering department who were attracted by the computer security issues.

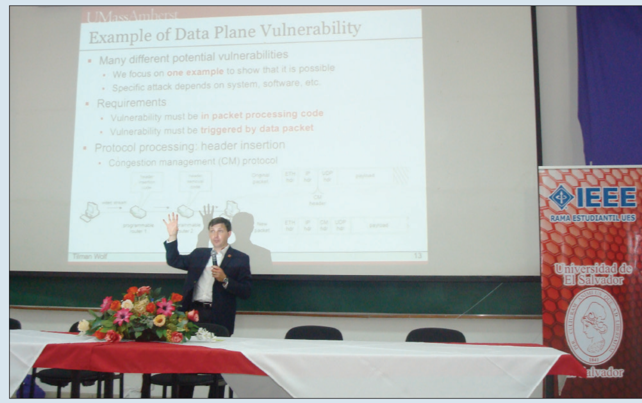

Lecture of Prof. Wolf in San Salvador.


Prof. Wolf suffering the inclemency of the rainy forest weather in El Salvador..

Lecture of Prof. Wolf in San Salvador.

But not everything was work. Central American countries are countries of volcanoes. Our guest was invited to visit El Boqueron, El Salvador's national park, at the top of San Salvador Volcano. There, Dr. Wolf experienced the difference in weather: rain at the top and dry weather at the bottom, where San Salvador is located. For a few minutes it was possible to appreciate the beautiful sight of the San Salvador valley. From above, at a distance, it was possible to see at the same time the crater-lake of Ilopango and a fast-moving tropical rain on the verge of covering the lake.

The final talks were held on November 13 in Panama City at the Hotel Continental. This lecture was supported by the Panama ComSoc chapter made up mostly of professors and students from the Panama University of Technology.

It was a pleasure for the Guatemala, El Salvador, and Panama ComSoc Chapters to host the Distinguished Lecturer tour of Prof. Wolf.

CONFERENCE REPORT

Highlights of the IEEE COMCAS 2015 Conference, Tel-Aviv, Israel

By Shmuel Auster, IEEE COMCAS 2015 General Chair, and Itai Dabran IEEE COMSOC Chapter Chair, Israel

The fifth biennial IEEE International Conference on Microwaves, Communications, Antennas, and Electronic Systems (IEEE COMCAS 2015) took place at the David Intercontinental Hotel in Tel-Aviv on the 2–4 November 2015. The first two days of the conference included the plenary session together with six regular parallel sessions and poster sessions, while the third day was reserved for tutorial and workshop tracks in three parallel sessions.

IEEE COMCAS is organized as a multidisciplinary international conference where scientists, engineers, and students can meet and discuss their common interests. In addition, it holds an exhibition with more than 100 booths that is attended by top internationally recognized scientists and engineers active in fields such as communication, antennas, microwave, and systems engineering. Colleagues who are developing the products and systems of tomorrow particularly find the exhibition helpful to their projects and careers.

This conference is sponsored by the IEEE Communications Society, the IEEE Microwave Theory and Techniques Society, and the IEEE Israel AP/MTT joint chapter. It is technically co-sponsored by additional chapters, societies, and professional organizations including the IEEE's AES (Aerospace and Electronic Systems) and APS (Antennas and Propagation); the Association of Engineers and Architects in Israel; the European Microwave Association (EuMA); and the GAAS (GaAs and other Compounds Association).

This year the organizers significantly increased the number of tutorials and industry-oriented invited sessions, making this conference more relevant than ever before. The emphasis continues to be on applications-oriented research and development, from antennas and devices to systems and software, including GaN technology and applications, biomedical systems and applications, phased array radars, SDR and 5G cellular mobile. The event had strong industry patronage and participation, and the final program comprised 150 technical presentations, talks and tutorials, arranged in 45 sessions over three days with up to six parallel sessions running concurrently for much of the time.

Frank Traut, Keynote speaker from MACOM Technology, USA, talking about "GaN Technology and Applications" at the Opening (Plenary) session.

Oren Hagai, Interlligent CEO and IEEE COMCAS Exhibition Committee Chair (in black jacket) and his staff at Interlligent during the IEEE COMCAS 2015 Exhibition.

Ingmar Kallfass, Invited Speaker from the University of Stuttgart, Stuttgart, Germany on "MMIC Chipset for 300 GHz Indoor Wireless Communication".

At the Posters Session, with Beer Cocktail (sponsored by Carlsberg beer). From left: Solon Jose Spiegel, Shraga Kraus, Shmuel Auster, Claudio Jackobson and Ofer Shaham.

Some of the presented topics at IEEE COMCAS 2015 were: green communication, radar and microwave technologies, device modeling, antennas, networking technologies, defense and UAS applications, communication measurements, implementations and resource allocation, sensors and advanced frequency synthesizers.

The conference also featured a special Women in Engineering session on 'Diversity in High-Tech – What's Working and Why?,' presented by Prof. Orit Hazzan from the Technion, that illustrated how the creation of a culture that enhances diversity benefits the entire STEM community. This invited speaker presentation asserted that it is in the interest of the high-tech world, rather than in the interest of any specific underrepresented group in the community, to enhance diversity in general, and gender diversity in particular.

Hosting such conferences in various locations around the world is an excellent way to encourage IEEE members and societies and the technical community at large to participate. This helps improve recognition and generate growth for the IEEE Communications Society, the IEEE Microwave Theory and Techniques Society, and their Sister Societies. The fifth IEEE COMCAS conference was highly successful, with more than 1500 participants from 38 countries who expressed their very positive feedback and comments.

COMCAS continues to be very successful and will likely be well attended in the future by many technologists who lead innovation and have an impact on the microwave, communications, antenna, solid-state circuits, and electronic systems fields.

Distinguished Lecturer Tour of Tarik Taleb: From Finland to Malaysia

By Fazirulhisyam Hashim, Hafizal Mohamad, and Nur Idora Abdul Razak, IEEE Malaysia ComSoc/VTS Joint Chapter

From left: Sumei Sun, Hamzah Burok, Khairul Akmal Zahri, Tarik Taleb, Tan Ching Seong, Hafizal Mohamad, and Fazirulhisyam Hashim at MCMC, Cyberjaya.

In November 2015, the IEEE Malaysia ComSoc/VTS Joint Chapter hosted an IEEE ComSoc Distinguished Lecturer Tour (DLT) by Professor Tarek Taleb at three different venues, namely MCMC Cyberjaya, UiTM Shah Alam, and the Hilton Kuching. The speaker, Prof. Taleb, is a Professor at the School of Electrical Engineering, Aalto University, Finland. Prior to his current academic position, he worked as senior researcher and 3GPP standards expert at NEC Europe Ltd, Heidelberg, Germany. Before joining NEC and until March 2009, he worked as an assistant professor at the Graduate School of Information Sciences, Tohoku University, Japan, in a lab fully funded by KDDI, the second largest network operator in Japan.

Prof. Taleb arrived in Malaysia on November 22 and was met by Fazirulhisyam Hashim, the current Chair of the IEEE Malaysia ComSoc/VTS Joint Chapter, at Kuala Lumpur International Airport (KLIA), and later checked-in at the Everly Putrajaya Hotel. After finalizing the check-in process, they had a light dinner at the Black Canyon Restaurant in the Alamanda Shopping Mall, which is a three-minute walk from the hotel.

The following day at 10.30 am, Prof. Taleb delivered his first lecture at the Malaysian Communication and Multimedia Commission (MCMC), Cyberjaya. There were approximately 60 attendees, mainly from various research institutes in Malaysia and industry. In the afternoon, Prof. Taleb went to the Faculty of Electrical Engineering, Universiti Teknologi MARA (UiTM) Shah Alam, to deliver his second lecture. There were approximately 20 attendees at this lecture; 90 percent of them were postgraduate students taking telecommunications courses, and the rest were from academia. The audiences were very interested in the topic, which is relevant to their field of studies and research. The program ended with some refreshments, at which time the audience took the opportunity to interact with the speaker.

On November 24 Prof. Taleb took a flight to Kuching, Sarawak to deliver his third lecture, which was scheduled on the following day at the Hilton Kuching, in the Main Ballroom. He was one of the keynote speakers at the 2015 IEEE 12th Malaysia International Conference on Communications (MICC2015). This lecture attracted 100 attendees, mainly from academia and industry.

In his three lectures, Prof. Taleb started with a video clip and introduction of Aalto University and the city of Helsinki. In general, the audiences were exposed to 5G technology and the research work in this area. Prof. Taleb highlighted some issues on the existing telecommunications networks that are

Audience at MCMC Auditorium, Cyberjaya.

hardware-oriented and have a centralized architecture. With the increasing demand for future mobile applications, networks need to be dynamic and flexible. The solution is to migrate to software-defined networking, virtualization, and cloud computing. Prof. Taleb also demonstrated his research work on the topic. In addition, he showcased the feasibility of on-demand creation of cloud-based elastic mobile networks, along with their lifecycle management. The lecture introduced a set of technologies and key architectural elements to realize such a vision, turning end-to-end mobile networking into software engineering.

All three of the lectures were very successful and well attended. We noted that some in the audiences were non-IEEE members. Thus, this event was an ideal platform for us to spread the message of IEEE to postgraduate students and industry. The IEEE Malaysia ComSoc/VTS Joint Chapter would like to thank Prof. Taleb for his lecture, and we hope that we will meet again in the near future. Above all, we would like to extend our special appreciation to IEEE ComSoc for arranging such a wonderful program.

Prof. Taleb during his lecture at MCMC.

Audience at Universiti Teknologi MARA (UiTM) Shah Alam.

CONFERENCE REPORT

IEEE Latin America Network Operation and Management Symposiums

By Fernando Menezes Matos, Augusto Venâncio Neto, Aldri Luiz dos Santos, and José Neuman de Souza, Brazil

The 8th Latin American Network Operations and Management Symposium (LANOMS) 2015 was held in the beautiful and pleasant coastal city of João Pessoa, jointly organized by the Federal University of Paraíba and the Federal University of Rio Grande do Norte. LANOMS is a biannual conference, and the 2015 event was the first time it was located in the Northeast region of Brazil. It was technically co-sponsored by the IEEE Communications Society (ComSoc) and the International Federation for Information Processing (IFIP). The conference has the goal to bring together international representatives of academic areas, civil services, businesses, and other entities in order to create a stimulating environment to present and discuss ideas, solutions, and lessons learned to solve technical and research challenges in network and service management. Moreover, LANOMS also aims to be the main Latin American conference in the network management field, alongside APNOMS (Asia-Pacific Network Operations and Management Symposium) and NOMS (Network Operations and Management Symposium), the main Asian and worldwide conferences in network management, respectively.

This LANOMS edition received 33 paper submissions, of which 12 were accepted as full papers and selected for oral presentation through four technical sessions (Network Virtualization, Security Management, Monitoring and Management Approaches, and Wireless Network Management). Moreover, five papers were selected for short paper presentations, and five papers were selected for poster presentations. An application session was also part of the LANOMS 2015 program, where applications papers, focusing on tools for contemporary and important paradigms, such as cloud computing and software defined networking (SDN), were presented. In addition, two keynotes were scheduled in the conference program, addressing relevant and up-to-date topics. In the first keynote, Prof. Rui Aguiar from the University of Aveiro, Portugal, outlined the management challenges of 5G networks; in the second keynote, Masum Z Hasan, from Cisco Systems (USA),

Local Organizing Committee.

Technical session.

talked about seamless cloud extensions over multiple disparate clouds. This year, LANOMS 2015 paid a small tribute to Prof. Luiz Fernando Gomes Soares, who developed several research efforts on network management and led the development of Ginga-NCL middleware. Ginga-NCL middleware is part of the Brazilian digital TV system, and has been adopted by several countries in Latin America. Finally, a social event was held at a barbecue restaurant, allowing approximately 40 participants to talk about the event, future cooperation, and the like.

The success of this LANOMS edition was largely due to all the people involved in the conference. First, we are most grateful to the authors who submitted their work and the technical program committee members who devoted their time to support the peer-review process. We also need to thank the Federal University of Paraíba (UFPB), the Federal University of Rio Grande do Norte (UFRN), the Federal University of Ceará (UFC), and the Federal University of Paraná (UFPR), the Brazilian Computing Society (SBC), and the National Laboratory of Computer Networks (LARC), for their invaluable support. We also want to thank the organizing committee members and the volunteers who all worked hard on the details and important aspects of the conference. Finally, but not least, we acknowledge the support of our sponsors (CAPES and CGI.br) who helped us cope with the costs of organizing LANOMS 2015, seeking to ensure a high-level event for all participants. We also invite our community to attend the future events.

For further information about LANOMS, please visit the conference portal at http://www. lanoms.org

GCN GLOBAL COMMUNICATIONS NEWSLETTER

STEFANO BREGNI
Editor
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

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

CARLOS MARTINEZ (CARLOS.ELSALVADOR@GMAIL.COM)
EWELL TAN, SINGAPORE (EWELL.TAN@IEEE.ORG)
ALEXEY VINEL (ALEXEY.VINEL@HH.SE)
LISANDRO ZAMBENEDETTI GRANVILLE (GRANVILLE@INF.UFRGS.BR)

IEEE ComSoc IEEE Communications Society

www.comsoc.org/gcn
ISSN 2374-1082

OMBUDSMAN

ComSoc Bylaws Article 3.8.10

"The Ombudsman shall be the first point of contact for reporting a dispute or complaint related to Society activities and/or volunteers. The Ombudsman will investigate, provide direction to the appropriate IEEE resources if necessary, and/or otherwise help settle these disputes at an appropriate level within the Society."

IEEE Communications Society Ombudsman
c/o Executive Director
3 Park Avenue
17 Floor
New York, NY 10017, USA
ombudsman@comsoc.org
www@comsoc.org "About Us" (bottom of page)

CRITICAL COMMUNICATIONS AND PUBLIC SAFETY NETWORKS PART 1: STANDARDS, SPECTRUM POLICY, AND ECONOMICS


Mehmet Ulema


Alan Kaplan


Kevin Lu


Niranth Amogh


Barcin Kozbe

Public safety networks have been utilized by first responder agencies in emergency situations, such as police officers, firefighters, and ambulances. Critical communications networks have been used in various sectors such as construction, transportation, utilities, factories, and mining operations.

Technologies used in public safety networks and critical communications networks today are at a crossroads with next-generation solutions and applications. Many of the existing technologies such as TETRA and P25 have been in use for about 20 years now. They are mature, reliable, and cost effective in supporting mission-critical voice applications. However, they are not designed to support higher bandwidth applications.

Partly due to the ubiquitous availability of commercial broadband applications, and partially due to increasing demand by public safety agencies, the possibilities of broadband data services for public safety networks that are primarily based on Long Term Evolution (LTE) technology are being discussed increasingly in many countries including those in North America and Europe. For example, in 2012 the U.S. government created an authority to establish a national public safety broadband network called FirstNet or First Responders Network. The standards organizations and stakeholders in developing narrowband technologies have indicated that their future strategy is to evolve into LTE-based solutions for public safety systems.

To realize the promises of broadband technologies for critical communications and public safety networks, many obstacles in implementations and deployments need to be overcome. Although early trials and deployments are underway, various challenges in architecture, protocols, operations, economics, and finance are being addressed in industry, academia, and standards organizations. These Feature Topic articles are intended to provide an in depth overview of the field, to analyze the current critical communications and public safety networking landscape, and to provide a comprehensive overview of the key challenges associated with standards, spectrum policy, and economics.

We anticipate that these discussions will encourage further research and development, leading to more advanced solutions. Finally, we hope that the authorship by multiple authors from a broad array of organizations will demonstrate to the magazine readers that the entire industry sees IEEE Communications Society as a consensus-building venue.

The Call for Papers for this Feature Topic generated many interesting and strong submissions, and the paper review process resulted in the acceptance of 11 papers. Therefore, it was decided to spread the accepted papers into two parts. The articles in Part 1 will focus on general topics such as survey, spectrum policies, and economics, whereas the articles in Part 2 will focus on more technical issues and solutions.

The following is a brief introduction to the articles in this issue.

The first article, "Toward Moving Public Safety Networks" co-authored by Romain Favraud, Apostolos Apostolaras, Navid Nikaein, and Thanasis Korakis, provides a good introduction to the field by surveying public safety use cases, current status of the related standards activities within the Third Generation Partnership Project (3GPP), and challenges ahead. The article also discusses backhauling in moving cells enabling dynamic meshing among LTE base stations.

The second article in this series is "Pervasive Spectrum Sharing for Public Safety Communications," co-authored by Murat Yuksel, Ismail Guvenc, Walid Saad, and Naim Kapucu. This article provides a comprehensive discussion on the need for sharing the spectrum, and introduces a new architecture where spectrum sharing is pervasive and necessary to realize a new perspective of increased heterogeneity in the next generation public safety networks.

The next article, "Next Generation Public Safety Networks: A Spectrum Sharing Approach," co-authored by Munawwar M. Sohul, Miao Yao, Xiaofu Ma, Eyosias Y. Imana, Vuk Marojevic, and Jeffrey H. Reed, also focuses on the spectrum aspects of public safety networks with an emphasis on spectrum sharing among public safety agencies and commercial carriers. The article presents different initiatives undertaken by Virginia Tech for deployment in next generation public safety networks. It also presents a vision for a flexible, rapidly deployable, and reconfigurable system to meet the demands of critical communications infrastructure.

The final article in Part 1 of this Feature Topic, "Indirect Returns and Use of NPV in Financial Viability Modeling of Critical Communications Networks" by Natalia Boliari, discusses economics and financial aspects of critical communications and public safety networks. The article focuses on the potential of public safety networks to generate long-term indirect returns in the form of various socioeconomic benefits. The article develops a net present value model that accounts for both direct and indirect benefits.

We hope that you will find these articles as interesting, informative, and challenging as we did. We would like to thank all the authors who submitted their articles to this Feature Topic, and the reviewers, who have given their time generously to provide valuable feedback and comments on the articles and thus to make these issues a reality. We also want to thank the Editor-in-Chief Osman Gebizlioglu and Joe Milizzo of IEEE for their support.

BIOGRAPHIES

MEHMET ULEMA (mehmet.ulema@manhattan.edu) is a professor with Computer Information Systems at Manhattan College, New York. Previously, he was with AT&T Bell Laboratories, Bellcore, Daewoo Telecom, and Hazeltine. He also serves as the Director of Standards Development in ComSoc. He was the TPC Chair of GLOBECOM 2009 and General Co-Chair of NOMS 2016. He is on the Editorial Boards of *IEEE Journal of IoT* and *Springer Journal of Network and Services Management*. He received his Ph.D. from Polytechnic University, Brooklyn, and his B.S. and M.S. from Istanbul Technical University.

ALAN KAPLAN (kaplana@ieee.org) is CTO of Drakontas, developing software for public safety agencies. He was formerly with Panasonic Princeton Research Lab, Clemson University, and Flinders University of South Australia. He holds Ph.D. and M.S. degrees in computer science from the University of Massachusetts Amherst and a B.S. in computer science from Duke University. He is also currently a lecturer in the Department of Computer Science at Princeton University.

KEVIN LU (klu@ieee.org) is an adjunct professor of electrical and computer engineering at Stevens Institute of Technology. He is a member of the IEEE Standards Association (IEEE-SA) Standards Board and is the IEEE-SA contact for the Global Standards Collaboration task force on emergency communications. He was a chief scientist and executive director at Telcordia applied research until 2012. He received his D.Sc. in systems science and mathematics from Washington University in St. Louis.

NIRANTH AMOGH (namogh@huawei.com) is a principal researcher at the Huawei India R&D Center at Bangalore. He is responsible for the wireless networks research within the organization. His research areas include broadband critical communications, M2M/IoT, SDN/NFV, and NGSON. He has filed several patents in his research areas and holds leadership positions in several SDOs in India and globally. In critical communications standardization, he is actively contributing to the 3GPP SA6 (Mission-Critical Applications) WG.

BARCIN KOZBE (kozbe@yahoo.com) is currently a senior consultant at NGen Solutions. Prior to joining NGen Solutions, he was a technical solutions manager at Ericsson Inc. He has been working in the field of computer science, specializing in information technology for telecommunications, for 20 years. He received his M.Sc. degree in computer engineering from Chalmers Technology University in Sweden. His research interests include public safety networks, network management systems, software defined networks, and cloud computing.

Toward Moving Public Safety Networks

Romain Favraud, Apostolos Apostolaras, Navid Nikaein, and Thanasis Korakis

Fourth generation LTE has been selected by U.S. federal and EU authorities to be the technology for public safety networks that would allow first responders to seamlessly communicate between agencies and across geographical locations in tactical and emergency scenarios.

ABSTRACT

Fourth generation LTE has been selected by U.S. federal and EU authorities to be the technology for public safety networks that would allow first responders to seamlessly communicate between agencies and across geographical locations in tactical and emergency scenarios. From Release 11 on, 3GPP has been developing and specifying dedicated nationwide public safety broadband networks that will be scalable, robust, and resilient, and can address the specific communication needs of emergency services. In this realm, the requirements and scenarios for isolated E-UTRAN with no or limited backhaul access to the core network are still in progress. In this article, we survey possible public safety use cases with the induced network topologies, discuss the current status of the 3GPP standards, and highlight future challenges. We further elaborate on the need to support mobile backhauling in moving-cell scenarios and describe two LTE-based solutions to enable dynamic meshing among the base stations.

INTRODUCTION

MOTIVATION

Long Term Evolution (LTE), specified by the Third Generation Partnership Project (3GPP), is becoming the technology reference for fourth generation (4G) cellular networks, as it is increasingly adopted by all major operators all over the world.

LTE is now rising to the challenge of addressing several issues (e.g., cellular networks' capacity crunch, ultra-high bandwidth, ultra-low latency, massive numbers of connections, superfast mobility, diverse spectrum access) that speed up the pace toward 5G. Moreover, LTE is expected to be an important part of the 5G solution for future networks and to play an essential role in advancing public safety (PS) communications. In the United States, LTE has been chosen up as the next appropriate communication technology to support PS, and it is likely to be the same in the European Union soon. Thus, several vendors (e.g., Ericsson, Nokia-Alcatel, Huawei, Cisco, Motorola, Thales) are now starting to propose LTE-based PS solutions, and some of them have been put to real field experimentation.

While existing PS solutions (e.g., Project 25, P25, and terrestrial trunked radio, TETRA) are mature and provide reliable mission-critical voice communications, their designs cannot meet the new requirements and the shift to higher bandwidth applications. In addition, LTE systems were suited to commercial cellular networks in the initial 3GPP releases but not to PS services and the corresponding requirements like reliability, confidentiality, security, and group and device-to-device communications. Therefore, the question raised is whether LTE suffices to be an appropriate solution for PS networks. To address those issues, 3GPP has started to define the new scenarios that LTE will have to face, and has released several studies and specificationson proximity-based services, group and device-to-device communications, mission-critical push-to-talk (MCPTT), and isolated Evolved Universal Terrestrial Radio Access Network (E-UTRAN). These studies define the requirements regarding user equipment (UE) and evolved NodeB (eNB — LTE base station) to provide PS services depending on the E-UTRAN availability and architecture.

In particular, the studies on isolated E-UTRAN target use cases when one or several eNBs have limited or no access to the core network (Evolved Packet Core, EPC) due to a potential disaster, or when there is need to rapidly deploy and use an LTE network outside of the existing infrastructure coverage.

However, 3GPP studies do not define how such isolated eNBs of a single set should communicate together, and leave that to the use of other technologies and vendor-specific solutions.

CONTRIBUTION

In this article, we discuss possible directions and challenges to evolve the LTE network architecture toward 5G in order to support emerging PS scenarios. Starting from the current status of standards on mission-critical communications and focusing on an isolated E-UTRAN case, we delineate two innovative solutions that allow for interconnection of eNBs using LTE, while qualifying the requirements defined by 3GPP for PS scenarios. Such solutions present several advantages when compared to dedicated technologies (e.g. WiFi, proprietary RF links), in that they support network mobility scenarios, and topology split and merge while being cost effective.

The first solution utilizes legacy UEs and evolves them in order to operate as active elements within the network (UE-centric), thus being capable of associating with multiple eNBs and restoring the disrupted links between them.

Romain Favraud is with DCNS and EURECOM; Apostolos Apostolaras and Thanasis Korakis are with the Centre for Research & Technology Hellas (CERTH) and University of Thessaly; Navid Nikaein is with EURECOM.

 0163-6804/16/$25.00 © 2016 IEEE

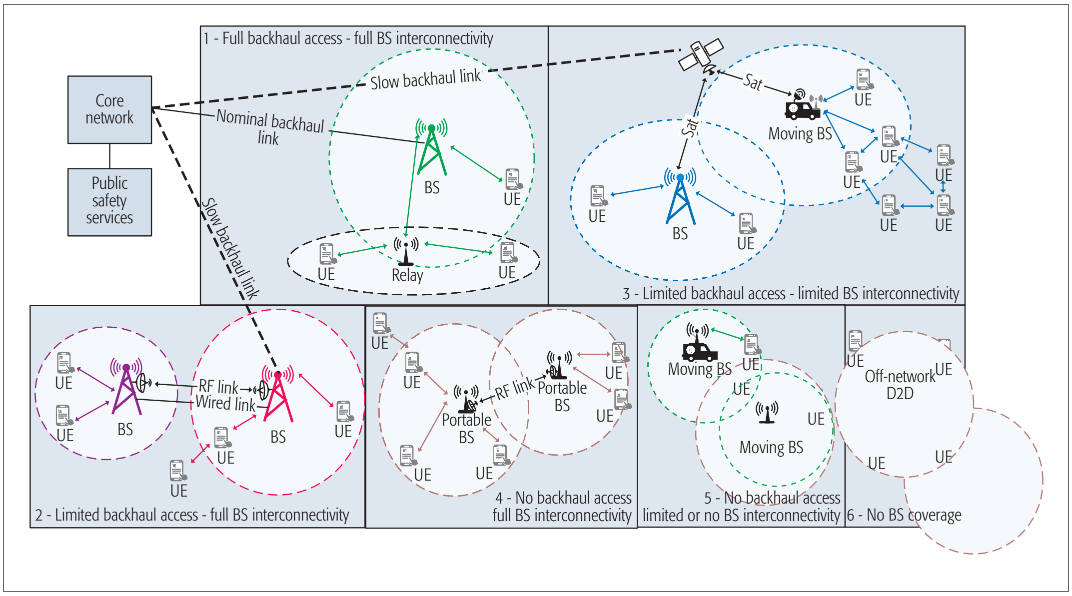

Figure 1. Public safety use cases. *Case 1*: a planned network with fixed BS deployment and backhaul connectivity. *Case 2*: a planned network with fixed BS deployment and limited backhaul connectivity. *Case 3*: a network with fixed BS deployment and moving cells with limited backhaul connectivity assisted by satellite links, proximity services, and device-to-device communications. *Case 4*: no backhaul access in an unplanned network deployment of portable BSs. *Case 5*: moving cells in an unplanned network deployment. *Case 6*: missing BS coverage and proximity services.

In nominal conditions, a nationwide broadband wireless PS network relies on a wired network supporting fixed wireless base stations providing planned coverage and bringing services to mobile entities relying on seamless access to the core network.

The second solution relies on extension of the eNB functionality to allow it to detect and connect directly to neighboring eNBs by encompassing multiple virtual UE protocol stacks (network-centric). These two solutions evolve and restore already existing and potentially disrupted wireless air interfaces such as Uu (eNB-UE radio interface), Un (eNB-relay radio interface), and X2 (inter-eNB logical interface). They create connectivity links among eNBs that can be used to form dynamic mesh networks allowing the size of an isolated E-UTRAN to be extended in fixed and mobile scenarios.

Use Cases and Topologies

Public safety users and first responders encounter a wide range of operational conditions and missions. To effectively address them, they need to rely on sufficient voice and data communication services. While voice services have already been used in tactical communication systems (e.g., TETRA and P25), the absence of a technology that could offer sufficient data rate left the associated services unexploited.

In nominal conditions, a nationwide broadband wireless PS network relies on a wired network supporting fixed wireless base stations (BSs) providing planned coverage and bringing services to mobile entities (e.g., handheld UE or vehicle integrated devices) relying on seamless access to the core network.

A key requirement for the network is that it must be robust, reliable, and not prone to malfunctions and outages. Despite that, it may not survive against unexpected events such as earthquakes, tidal waves, and wildland fires, and may not cover distant lands due to costly deployment.

Figure 1 illustrates six different topologies corresponding to possible use cases that PS users may encounter depending on the operational situation. These six topologies are differentiated based on four criteria:

- Availability of the backhaul link (access to the core network from the BS)
- BS interconnections
- BS mobility
- BS availability (UEs on- or off-network)

In the nominal case (case 1 in Fig. 1), BSs are fixed and benefit from planned coverage as they receive complete services support, and experience full access to the core network and to the remote PS services with no intermissions (e.g., continuous link connectivity with the operation center, monitoring, billing). Therefore, the network can provide nominal access to PS UEs; this case refers to the majority of operations (e.g., law enforcement, emergency services, fire intervention) occurring in covered cities and (sub)urban environments where the network deployment has been previously designed and planned, and services are provided within a large coverage expansion.

In the case of backhaul link failure due to faulty equipment, power outage, or physical damages on the backhaul wires or RF antennas, the core network may not be fully accessible any longer to the fixed BSs (cases 2 and 3). However, depending on either the type and position of failure, or the availability of backup solutions (e.g., satellite backhaul as given in case 3),[1] the BSs may still maintain adequate interconnectivity with each other. Portable BSs (fixed once deployed) can be exploited in order to provide coverage on site, where fixed BSs have not been fully deployed yet or have faulty operation (case 4). In the same way, moving BSs can be utilized in a more dynamic fashion (e.g., for a fight against a fast moving forest wildfire, in vehicular communication on land or at sea [1, 2]) where it is not possible to plan inter-BSs links (case 5). In these cases of portable or moving BSs use,

[1] In such a case, the communication protocol is usually improved by a performance-enhancing proxy (PEP) as specified in Internet Engineering Task Force (IETF) RFC 3135 and RFC 3449.

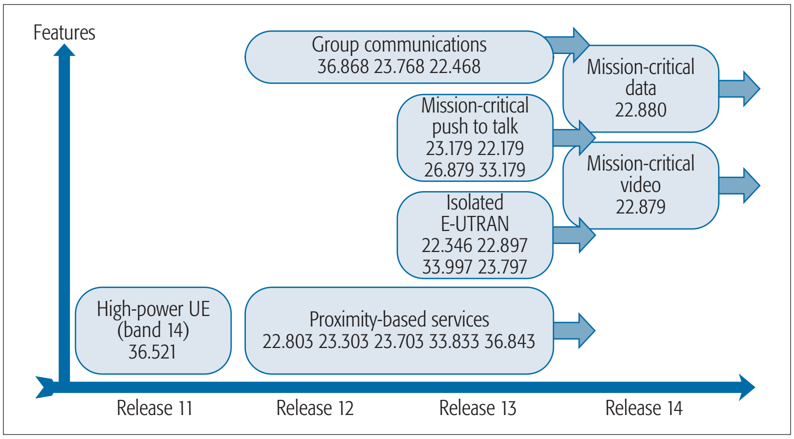

Figure 2. 3GPP PS oriented work items.

it can be hard or impossible to maintain a good connectivity with the macro core network (cases 3, 4 and 5).

Finally, it is likely that due to mobility, users would get out of the coverage service area provided by the BSs (cases 2, 3, and 4), or that in-time service provisioning to users would fail due to intense mobility (case 6). Therefore, due to their own inherent limitations (access to the core network, inter-node connectivity, BSs and UEs mobility), all previous topologies may not be able to provide the same services with a sufficient level of quality to users. For instance, the billing and monitoring services might not be available in some cases. Nevertheless, PS users must be able to use vital services like voice and data group communications in all situations regardless of the network topology dynamic. That is why PS wireless communications cannot rely solely on a planned network of fixed BSs.

Standards Development

The simmering interest of public authorities in LTE for PS use has encouraged 3GPP to tackle this subject and to evolve LTE specifications. Specifically, significant standardization activities have been conducted after the creation of the First Responder Network Authority (FirstNet) in the United States. As illustrated in Fig. 2, the first work dedicated to PS was launched in 3GPP Release 11 with the introduction of high-power devices operating in band 14 (used in the United States and Canada for PS) to extend the possible coverage servicing area. Since then, several work items have been defined in Releases 12 and 13 to study and address the specific requirements of a broadband PS wireless network, not least of which are:

- Guaranteed access: A PS network should be accessible at any time.
- Quality of service (QoS): Guarantee and priority should be ensured for critical calls.
- Reliability: PS networks should provide the services as defined with no interruption when online.
- Resiliency: A PS network should be able to evolve with technology advancements and changes to operational requirements.
- Roaming: UEs should be able to seamlessly use the deployed PS network as well as commercial networks in case of unavailability of the first.
- Spectrum efficiency, capacity, coverage: Spectrum has to be effectively shared to provide the required capacity and coverage.
- Talk around/simplex: Users should be able to communicate even in the case of broadband network unavailability or disruption.

The gaining momentum of LTE networks around the globe has relied on its architecture to provide packet-based network services that are independent of the underlying transport-related technologies. A key characteristic of the LTE architecture is the strong dependence of every deployed eNB on the EPC for all types of services that are provided to the covered UEs. However, this feature prevents UEs from seamless communication service when an eNB is disconnected from the EPC as eNB service to the UEs is interrupted even for local communications. To tackle the aforementioned shortcoming, 3GPP has launched two series of work items: the first one refers to device-to-device communications for enabling proximity-based services (ProSe), and the second one refers to the continuity of service for PS UEs by the radio access network (RAN) and eNBs in the case of backhaul failure for enabling operation on isolated E-UTRAN.

As defined in 3GPP technical specification (TS) 22.346, isolated E-UTRAN aims at the restoration of the service of an eNB or a set of interconnected eNBs without addressing their backhaul connectivity. The goal of isolated E-UTRAN operation for PS (IOPS) is to maintain the maximum level of communications for PS users when eNB connectivity to the EPC is either unavailable (no backhaul) or non-ideal. Isolated E-UTRAN can take place on top of nomadic eNBs (portable BSs, c.f. TS 23.797) deployments or on top of fixed eNBs suffering failures. It should support voice and data communications, MCPTT, ProSe, and group communications for PS UEs under coverage as well as their mobility between BSs of the isolated E-UTRAN, while maintaining appropriate security.

Subsequent to TS 22.346, TS 23.797 provides a solution to the no backhaul IOPS case relying on the availability of a local EPC co-located with the eNB or on the accessibility of the set of eNBs. PS UE(s) should use a dedicated universal subscriber identity module (USIM) application for authentication and use the classical Uu interface to connect to these IOPS networks. If an eNB cannot reach such a local EPC, it must reject UE connection attempts. However, the aforementioned solution does not address issues of scenarios with non-ideal backhaul connectivity. Moreover, requirements for the inter-eNB link connectivity are not specified, even though the operation for a group of interconnected eNBs is defined.

In this article, we advocate the need for novel inter-eNB wireless connectivity as a key for the efficiency of isolated E-UTRAN operation that would allow broadening the network and enhancing the level of cooperation between adjacent nodes, leading to better service provision to the users. We also consider moving cells and meditate on eNB mobility, which is often encountered by (highly) mobile PS entities, in a potential split and merge network.

Future Challenges in Public Safety

Given the wide range of applications, PS communications must be able to provide to a large extent flexibility and resiliency. Being able to adapt under various circumstances and mobility scenarios that are characterized by disrupted communication links (e.g., damaged S1 interface and no EPC network access) and volatile infrastructure operation is of utmost importance. Although there is increasing interest in the development of public safety solutions for isolated E-UTRAN scenarios both by industry and academia, there are still open challenges. Next, we discuss the main ones.

Moving Cells and Network Mobility

In a crisis or tactical scenario, it is vital that field communications can be highly mobile and rapidly deployable to provide network access and coverage on scene. Currently, E-UTRAN is considered fixed, and detection as well as discovery of a network while moving cells are being deployed remains unspecified. When high mobility occurs, the problem becomes network availability as link connections to the EPC servers are dropped. Moreover, due to the limited coverage of moving cells as compared to fixed eNBs [1], enabling inter-cell discovery features for proximity awareness is required as a tool of network intelligence for self-healing. eNBs must be able to search for other eNBs in their proximity either directly or relying on the assistance of enhanced UEs (i.e., UEs with extended capabilities that can interconnect between two eNBs) and eventually synchronize to the most suitable one and re-establish access to the network.

Device-to-Device Discovery and Communications

In the absence of network coverage (off-network case), PS UEs need to discover and communicate with each other by taking partial control of the functionality of the network [3]. UEs should be able to provide network assistance when infrastructure nodes (i.e., eNBs) are missing due to network and/or terminal mobility, or unavailable due to outage and malfunctioning. In such situations, UEs are promoted to assist with time synchronization reference (e.g., based on sidelink power measurement or UEs' own timing), authentication, detection, network discovery, and attachment functions, among others. In addition, UEs may need to request the identity of neighboring UEs (i.e., who is here) belonging to different PS authorities, which calls for over-the-air sensing and self-reconfiguration functionality at the UE side. What is more challenging for PS-UEs is the support of (stored) data relaying from (isolated) neighboring UEs to either other UEs (UE-to-UE relay) or the network (UE-to-network) when they are in coverage.

Programmability and Flexibility

Programmability and flexibility in future PS systems shall allow the rapid establishment of complex and mission-critical services with specific requirements in terms of service quality. A high degree of programmable network components will be able to offer scalable and resilient network deployment on the fly without the need for previous network planning by using network function virtualization and software-defined networking (SDN). Thus, it will result in availability of open network interfaces, virtualization of networking infrastructure, and rapid creation and deployment of network services with a flexible and intelligent control and coordination framework. Such a control and coordination framework is required to manage the entire life cycle of the PS network from configuration and deployment to runtime management and disposal. This is very challenging as it has to optimize the resource allocation across multiple eNBs, manage the topology (especially during the network split and merge), and determine the IP addressing space among the others.

Traffic Steering and Scheduling

The decisions about traffic steering concern control plane actions enabled to form a wireless mesh network. Selecting one or a subset of eNBs to steer the data plane traffic allows users to be connected to the best suited network according to their QoS requirements and the network resources availability. Aimed at overall network optimization, traffic steering techniques can be leveraged to balance the network load, and satisfy carrier and user demands by properly enabling data offloading, interference management, and energy saving policies. Furthermore, the control and data planes should be decoupled as the routing decision and eNB selection are performed at the higher layers while data transfer is operated at the lower layers. Therefore, a novel mechanism to support the BS meshing by giving access to the forwarding table at the lower layers is required. It can be implemented either locally or over the network. In the former case, the forwarding table can be built simply based on the routing table. In the latter case, an SDN approach can be applied to interface between the control and data planes.

> The decisions about traffic steering concern control plane actions enabled to form a wireless mesh network. Selecting one or a subset of eNBs to steer the data plane traffic allows users to be connected to the best suited network according to their QoS requirements and the network resources availability.

Optimization of Performance Metrics to Support Sufficient QoS

A PS network requires provision of sufficient services when a serving eNB currently experiences interruption on backhaul connectivity. Apart from the initiation of isolated E-UTRAN operation, such as exploitation of inter-eNB connectivity links for recovery of system connectivity, a PS network also requires a mechanism to invoke the appropriate complementary resources (e.g., additional bandwidth, alternate communication links, complementary bearers) for self-healing operation and re-establishment of disrupted end-to-end bearers. For more efficient operation of the network, it is important that the same mechanism makes decisions by considering not only the availability of the complementary resources, but also the indicators and the metrics that characterize communication performance (latency, throughput, spectral efficiency, etc.) on the links and priority-level assignment on the Evolved Packet System (EPS) bearers.

Toward Moving Public Safety Networks

In current LTE architectures, eNBs are perceived as the active elements responsible for management and control of the RAN. On the opposite side, UEs are passive clients from the eNB per-

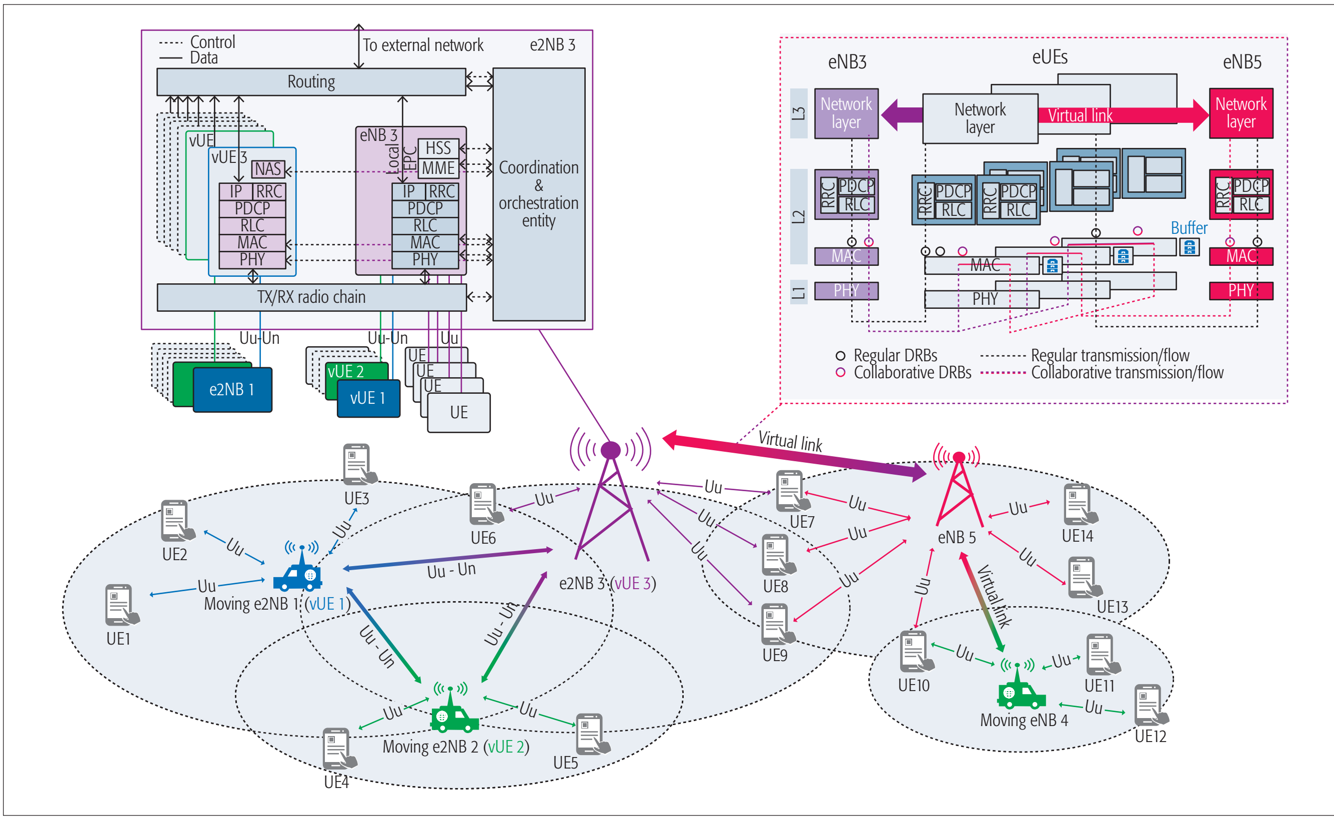

Figure 3. eUE and e2NB architecture for public safety: meshing of isolated or moving eNBs is enabled either (i) by leveraging eUEs as intermediate packet forwarders (UE-centric), thus creating virtual links between eNBs; or (ii) by leveraging e2NBs' functionality of encompassing multiple UEs (network-centric), thus restoring disrupted eNB-eNB communication.

spective, obeying certain rules and complying with the eNB's policies. Thus, the relationship between eNBs and UEs follows the master-slave communication model that is designed to meet the requirements of a fixed network topology. However, network mobility is increasingly gaining interest, and mobile scenarios where portable or moving cells are essentially required for rapidly deployable networks render networking elements with enhanced capabilities more and more attractive. We advocate the need to address those future mobility objectives as a means to meet PS requirements in an isolated E-UTRAN operation. In this direction, the role of legacy eNBs and UEs should be reconsidered within the network.

Following this approach, we delineate two novel solutions that allow inter-eNB link connectivity to be realized and the disrupted air interface to be restored by utilizing:

- Evolved UEs (denoted as eUEs)
- Enhanced eNBs (denoted as e2NBs)

The first refers to a UE-centric network-assisted solution. UEs are assigned enhanced capabilities of associating with multiple eNBs using multiple UE stacks, and thus interconnecting adjacent eNBs. They act as 3GPP UE terminals, maintaining their initial operation, and also as slaves from the eNB perspective. The second concerns a network-centric solution. The eNB stack is extended with several UE stacks, in what we call an e2NB, allowing it to discover and connect to neighboring eNBs, forming a wireless mesh network. A potential but achievable topology is illustrated in Fig. 3, along with a concise depiction of the eUE and e2NB architectures.

Evolved UEs

Evolved UE, like legacy UE, interprets the scheduling information coming from the eNB on the downlink control and signaling channels so as to enable traffic routing and forwarding relying on the allocated physical resource blocks (RBs). Moreover, they report measurements of channel state information (CSI) and buffer status report (BSR) back to the eNB. Furthermore, eUEs have enhanced capabilities of associating with multiple eNBs and thus interconnecting adjacent eNBs [4]. As a consequence, eUEs can also be used to extend the cell servicing area and provide backhaul access to core-isolated eNBs and hence to isolated E-UTRAN scenarios. eUEs can act as intermediate nodes so as to forward the traffic originating from or destined to eNBs. They belong to the control of the RAN of the bridged eNBs.

Enhanced eNB (e2NB)

The e2NB solution relies on the legacy 3GPP eNB and UE functions [2]. The e2NB solution consists of:

- The ability to provide service to mobile UEs and maintain the legacy eNB operation as a standalone node
- The ability to form a wireless mesh network when it is in close proximity to other e2NBs while maintaining service for the mobile entities

The former is achieved by extending the eNB functionality with that of the core network (i.e., mobility management entity, MME, and home

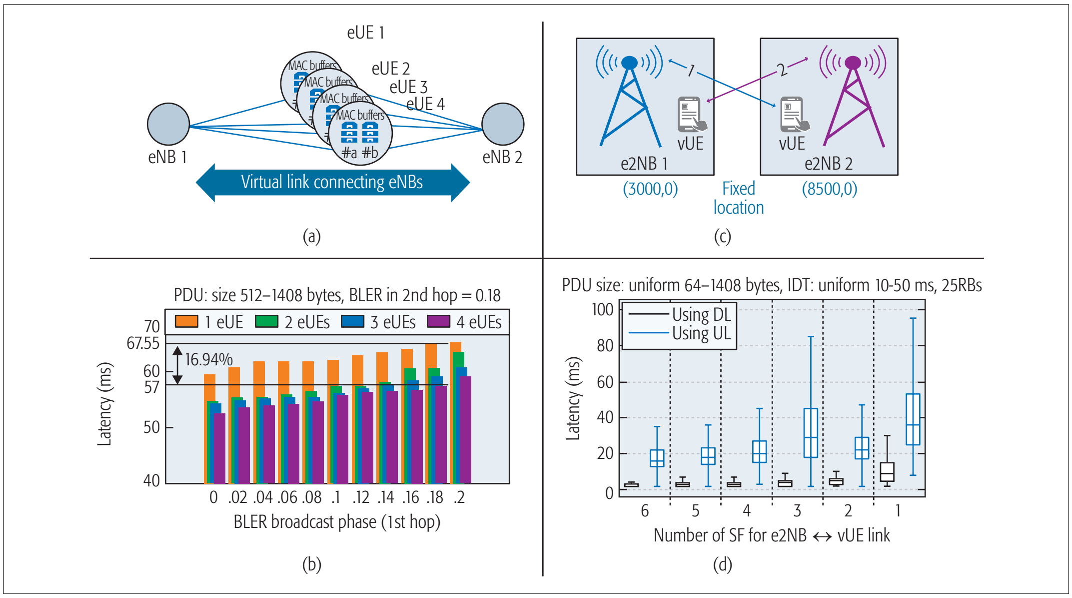

Figure 4. Logical topology for the performance evaluation scenario in OAI: a) four eUEs are leveraged to interconnect two eNBs; b) performance results for latency in the eUEs scenario; c) two e2NBs establish link connectivity using vUEs; d) performance results for latency in e2NB-vUEs scenario.

WiFi solutions are promising if the higher layers and protocols allow for efficient and dynamic meshing, similar to the proposed LTE-based solutions (eUE and e2NB). However, additional dedicated equipment and antennas are needed for WiFi backhauling, thus increasing the cost of BSs.

suscriber server, HSS), which allows it to manage UEs and provide PS services as is proposed by the 3GPP isolated E-UTRAN no-backhaul solution. The latter leverages the Uu and Un interfaces of the 3GPP UE and relay node. An e2NB encompasses multiple virtualized UEs (vUEs), integrating full LTE UE stacks, and one eNB. They share the radio resources and front-end. VUEs are used to discover other e2NBs and can be instantiated on demand to connect to the neighboring eNBs using Uu interface and UE connection procedures before switching to the Un interface. The discovery and on-the-fly connection features allow the e2NB to surpass the classical LTE relay [5] by enabling BS mobility and multiple connections to neighbors, re-establishing inter-eNB connectivity.

Evaluation of Feasibility and the Impact on Latency

In order to evaluate the performance of the above isolated E-UTRAN solutions in a practical and real setting, an implementation prototype of the proposed solutions was tested using the OpenAirInterface platform [6]. Specifically, OpenAirInterface is an open source software implementation of the 4G mobile cellular system that is fully compliant with the 3GPP LTE standards and can be used for real-time indoor/outdoor experimentation and demonstration. After thorough experimentation, results demonstrated the feasibility of the proposed approaches, as these have been presented in [2, 4]. Indicatively, in Fig. 4 we demonstrate two topologies for the isolated E-UTRAN problem where backhaul connectivity is not present. In Fig. 4a, four UEs are leveraged to restore the link connectivity between two eNBs. Performance evaluation results reveal (as shown in Fig. 4b) a significant reduction in latency (up to 16.94 percent), which depends on the number of active cooperating eUEs (up to 4).

In Fig. 4c, two e2NBs enable inter-eNB connectivity utilizing vUE operation. Two vUE-e2NB links are created allowing use of a subset of uplink and downlink subframes (SFs) from one e2NB to the other (six scenarios in Fig. 4d). An important finding that concerns latency performance is that whether using uplink (UL, UE to eNB) or downlink (DL, eNB to UE), the latency improves as the number of available SFs increases. More importantly, DL shows significantly lower latency performance overall as this is not only related to the resource allocation policy but also to the scheduling choice of using the UL or DL path. Thus, flows with different QoS requirements should be mapped on the corresponding link; for instance, low-latency services (e.g., voice calls) should go over DL paths.

Discussion

Some research articles provide insight into solutions when no backhaul is available, providing inter-eNB connectivity relying on WiFi links and including D2D communications that are not yet defined by the ProSe specifications of 3GPP studies [7]. Other technologies are usually used to establish wireless backhaul supporting fixed LTE networks: point-to-point (PTP) RF or free space optics (FSO) links, and point-to-multipoint (PTMP) RF links. In the case of portable BSs, satellite backhaul links are sometimes used. However, we can easily see that these wireless solutions are not adequate for the establishment of a network of BSs enabling voice and data communications in moving cell scenarios.

For instance, Table 1 shows the main differentiating criteria. Despite great performance, PTP and PTMP solutions usually require line-of-sight wireless connectivity with careful network planning, which makes them not applicable to moving cell scenarios. Satellite backhauling, on the other hand, provides the best possible coverage, but may require dedicated tracking antennas and suffers from high cost and high latency (≥ 200 ms) that limit voice and data services [8]. WiFi solutions are promising if the higher layers and protocols allow for efficient and dynamic mesh-

BS back-hauling	PTP/PTMP/FSO	SAT	WiFi	eUEs	e2NBs/vUE
Frequency band	ISM or licensed	Licensed	ISM, possibly licensed	Licensed	Licensed
Link latency	Very low	High	Low–medium	Low–medium	Low
BS mobility support	No	If tracking antenna	If omni-antennas	Yes	Yes
Cost	+++	++++	++	++	+
Topology	Star/mesh	Star	Star/mesh	Mesh	Mesh

Table 1. Main characteristics of base station backhauling solutions.

ing, similar to the proposed LTE-based solutions (eUE and e2NB). However, additional dedicated equipment and antennas are needed for WiFi backhauling, thus increasing the cost of BSs. In addition, commodity WiFi works on industrial, scientific, and medical (ISM) bands, and thus can experience more interference compared to LTE using licensed bands.[2] Studies on commercial networks have shown that WiFi latency is on average a bit higher and has more jitter than that of LTE, although results might differ for PS networks [9]. Moreover, carrier aggregation and full duplex communications are expected to greatly increase LTE global throughput in such mesh topologies, although similar techniques could be used for WiFi.

Some Reflections and Conclusion

Commoditization and virtualization of wireless networks are changing network design principles by bringing IT and cloud computing capabilities in close proximity of network and users. This will facilitate the deployment and management of PS networks by offering a service environment so that adequate (e.g., missing) network functions and applications can be dynamically instantiated for isolated network segments to maintain communication, service, and application as desired [10]. Packet core network functions (e.g. MME, HSS), IP multimedia subsystem (IMS), routing, and topology management are network functions that can be enabled at the BS to restore communication links. Traffic steering, video analytics, content sharing, and localization are examples of network applications that can extend BS functions in order to preserve user service and applications.

In this article, we elaborate on innovative solutions in the context of public safety networks to support efficient isolated E-UTRAN operation. We identify the shortcomings in the state-of-the-art technology, which is currently unable to sufficiently deliver seamless and continuous backhaul connectivity in moving cell scenarios, thus depriving first responders and tactical forces of critical communications. Specifically, we indicate that in the volatile and dynamic environment for public safety communication, the following are needed:

- Evolving UEs as active network elements to restore disrupted air interfaces between bridging eNBs
- Enhancing the role of legacy eNBs to encompass dual protocol stack operation for enabling base station meshing, which is of utmost importance to preserve the integrity of communication

Reviewing the open challenges that pose significant requirements in the field of services provision, we outline the most important and discuss related open research directions.

[2] To solve this problem, certain countries define their own licensed bands for PS WiFi.

Acknowledgment

Research and development leading to these results have received funding from DCNS and from the European Framework Programme under H2020 grant agreement 671639 for the COHERENT project.

References

[1] Y. Sui *et al.*, "Moving Cells: A Promising Solution to Boost Performance for Vehicular Users," *IEEE Commun. Mag.*, vol. 51, no. 6, June 2013, pp. 62–68.
[2] R. Favraud and N. Nikaein, "Wireless Mesh Backhauling for LTE/LTE-A Networks," *MILCOM 2015*, Oct. 2015, pp. 695–700.
[3] R. Liebhart *et al.*, *LTE for Public Safety*, Wiley, 2015.
[4] A. Apostolaras *et al.*, "Evolved User Equipment for Collaborative Wireless Backhauling in Next Generation Cellular Networks," *2015 12th Annual IEEE Int'l. Conf. Sensing, Commun., and Networking,* June 2015, pp. 408–16.
[5] Y. Yuan, *LTE-Advanced Relay Technology and Standardization*, Springer-Verlag, 2013.
[6] N. Nikaein *et al.*, "Demo: Openairinterface: An Open LTE Network on a PC," *Proc. 20th ACM Annual Int'l. Conf. Mobile Computing and Networking*, ser. MobiCom '14, 2014, pp. 305–08; http://doi.acm.org/10.1145/2639108.2641745.
[7] K. Gomez *et al.*, "Enabling Disaster-Resilient 4G Mobile Communication Networks," *IEEE Commun. Mag.*, vol. 52, no. 12, Dec. 2014, pp. 66–73.
[8] M. Casoni *et al.*, "Integration of Satellite and LTE for Disaster Recovery," *IEEE Commun. Mag.*, vol. 53, no. 3, Mar. 2015, pp. 47–53.
[9] J. Huang *et al.*, "A Close Examination of Performance and Power Characteristics of 4G LTE Networks," *Proc. 10th ACM Int'l. Conf. Mobile Systems, Applications, and Services*, ser. MobiSys '12, ACM, 2012, pp. 225–38; http://doi.acm.org/10.1145/2307636.2307658.
[10] Y. C. Hu *et al.*, "Mobile Edge Computing — A Key Technology Towards 5G," ETSI White Paper No. 11, tech. rep., Sept. 2015; http://www.etsi.org/images/files/ETSIWhitePapers/etsi_wp11_mec_a_key_technology_towards_5g.pdf.

Biographies

Romain Favraud (romain.favraud@eurecom.fr) received his M.S. degree in engineering from Grenoble Institute of Technology, France, in 2013 and his M.S. degree in electronic engineering from Polytechnic University of Turin, Italy, in 2013. He has been working since 2014 as a research engineer for DCNS, Toulon, France, and is a Ph.D. candidate at EURECOM, Sophia-Antipolis, France, under the supervision of Navid Nikaein. His current research interests are LTE mesh networks and new architectures for future mobile networks.

Apostolos Apostolaras received his Ph.D. (July 2014) from the Department of Electrical and Computer Engineering, University of Thessaly, Greece (under the supervision of Prof. Leandros Tassiulas) and is currently working as a research scientist. His research interests lie in but are not limited to the definition of methodologies derived from optimization theory and distilled into practical system implementation. He is also interested in experiment-based research in testbeds for the evaluation and improvement of networking system performance.

Navid Nikaein is an assistant professor in the Mobile Communication Department at Eurecom. He received his Ph.D. degree (docteur és sciences) in communication systems from the Swiss Federal Institute of Technology, EPFL, in 2003. He leads a research group focusing on experimental system research related to wireless systems and networking. His contributions are in the areas of wireless access layer techniques and networking protocols, cloud-native and programmable radio networks, real-time prototypes, and scalable emulation and simulation.

Thanasis Korakis received his B.S. and M.S. degrees from the Informatics and Telecommunication Department, University of Athens, in 1994 and 1997, respectively, and his PhD degree in Computer and Communications Engineering from the University of Thessaly, Greece, in 2005, under the supervision of Prof. Leandros Tassiulas. Currently, he is an assistant professor in the Electrical and Computer Engineering Department, University of Thessaly. His research interests are in the field of wireless networks with emphasis on access layer protocols, cooperative networks, directional antennas, quality of service provisioning, and network management.

Critical Communications and Public Safety Networks

Pervasive Spectrum Sharing for Public Safety Communications

Murat Yuksel, İsmail Güvenç, Walid Saad, and Naim Kapucu

Next-generation public safety communication systems must sustain high-speed, ultra-reliable wireless data transmissions. Moving toward this next generation of PSCs warrants a new perspective of increased heterogeneity in emerging wireless architectures and increased multiplexing of wireless spectrum.

Abstract

Next-generation public safety communications (PSC) systems must sustain high-speed, ultra reliable wireless data transmissions. Moving toward this next generation of PSCs warrants a new perspective of *increased heterogeneity* in emerging wireless architectures and *increased multiplexing* of wireless spectrum. To realize this vision, models that incentivize users to *opportunistically share* their spectrum as substrates over possibly multiple hops, and *decentralized and open* techniques that seamlessly exploit these substrates for public safety applications are much needed. The value of such *multihop and multi-technology pervasive spectrum sharing (PSS)* is more pronounced for application scenarios in which the need for spectrum access is vital, and infrastructure-less operation is necessary. This article introduces PSS as a new architecture where sharing is the norm, and outlines its vision, principles, and technical challenges.

Introduction

Public safety communications (PSC) carry critical importance to save lives and property in case of incidents such as fires, terrorist attacks, and natural disasters. The National Broadband Plan (NBP) [1] included the enhancement of the nation's PSC capabilities as one key priority. Three major challenges face our nation's public safety agencies in their use of radio communications [2, 3]:

1. Lack of capacity (radio spectrum allocated for public safety use is highly congested, especially in urban areas and during emergencies)
2. Lack of interoperability (multiple frequency bands, incompatible radio equipment, and a lack of standardization)
3. Lack of functionality (e.g., support for high definition video)

Remarkably, until recently, PSC has been handled via narrowband technologies that fall short on addressing the stringent quality of service (QoS) requirements of critical public safety applications.

To move U.S. PSC capabilities toward the next generation, there is an urgent need for pervasive availability of the spectrum with more open boundaries (Fig. 1). This is particularly critical for scenarios that require little or no infrastructure support and involve disaster response and recovery situations [1]. Emerging wireless standards such as fourth generation (4G) Long Term Evolution (LTE) and relatively shorter-range technologies such as WiFi have the potential to transform the capabilities of next-generation PSC systems. In particular, LTE is emerging as a dominant technology to support PSC, as evidenced by its adoption in the United States, Australia, and other countries. These global decisions triggered the Third Generation Partnership Project (3GPP) standardization group to specify advanced functionalities of 4G LTE technology and its evolution, such as device-to-device (D2D) communications, to support specific requirements of PSC [4].

Several advanced methods have been introduced to increase the efficiency of spectrum sharing, such as auctions [5, 6]. Although these advanced approaches have been successful at increasing the efficiency of spectrum sharing in a confined local neighborhood (a.k.a. one-hop relationships), improving spectrum access and efficiency on a larger horizon, such as within PSC and D2D scenarios, requires a truly interdisciplinary effort solving the technical, economic, and policy problems that are involved.

As recently recognized in the Boston Marathon bombings,[1] in an emergency scenario with limited infrastructure and a large number of users overloading the spectrum, it is of paramount importance to utilize all available substrates such as cellular, WiFi, Bluetooth, and multihop communication capabilities (e.g., via WiFi-Direct[2] and/or LTE-Direct[3]) for efficient usage of the spectrum by victims and first responders.

One promising direction in this regard is the recent introduction of D2D communication over cellular and WiFi bands. Indeed, while resource sharing between wireless devices has been traditionally restricted to short-range technologies such as Bluetooth or Zigbee, enabling D2D over cellular and WiFi presents a high-reward opportunity for realizing a *highly participatory and pervasive sharing* of heterogeneous, multi-purpose wireless spectrum resources, to which we will refer hereinafter as pervasive spectrum sharing (PSS). Providing incentives for such pervasive sharing of a valuable resource involves many techno-economic challenges, such as:

This work is supported in part by the U.S. National Science Foundation under Grants AST-1444077, AST-1443999, AST-1506297, and AST-1443946. The authors thank Thomas Quint, Arvind Merwaday, Meryem Simsek, and Brie Haupt for discussions and support.

[1] "Cellphone Networks Overwhelmed after Blasts in Boston," The Boston Globe, April 16, 2013; https://www.bostonglobe.com/business/2013/04/16/cellphone-networks-overwhelmed-blast-aftermath/wq7A-X6AvnEemM35XTH152K/story.html

[2] WiFi Direct, http://www.wi-fi.org/discover-and-learn/wi-fi-direct

[3] LTE Direct, http://www.qualcomm.com/research/projects/lte-direct

Murat Yuksel is with the University of Nevada, Reno; İsmail Güvenç is with Florida International University; Walid Saad is with Virginia Tech and Kyung Hee University; Naim Kapucu is with the University of Central Florida, Orlando.

0163-6804/16/$25.00 © 2016 IEEE

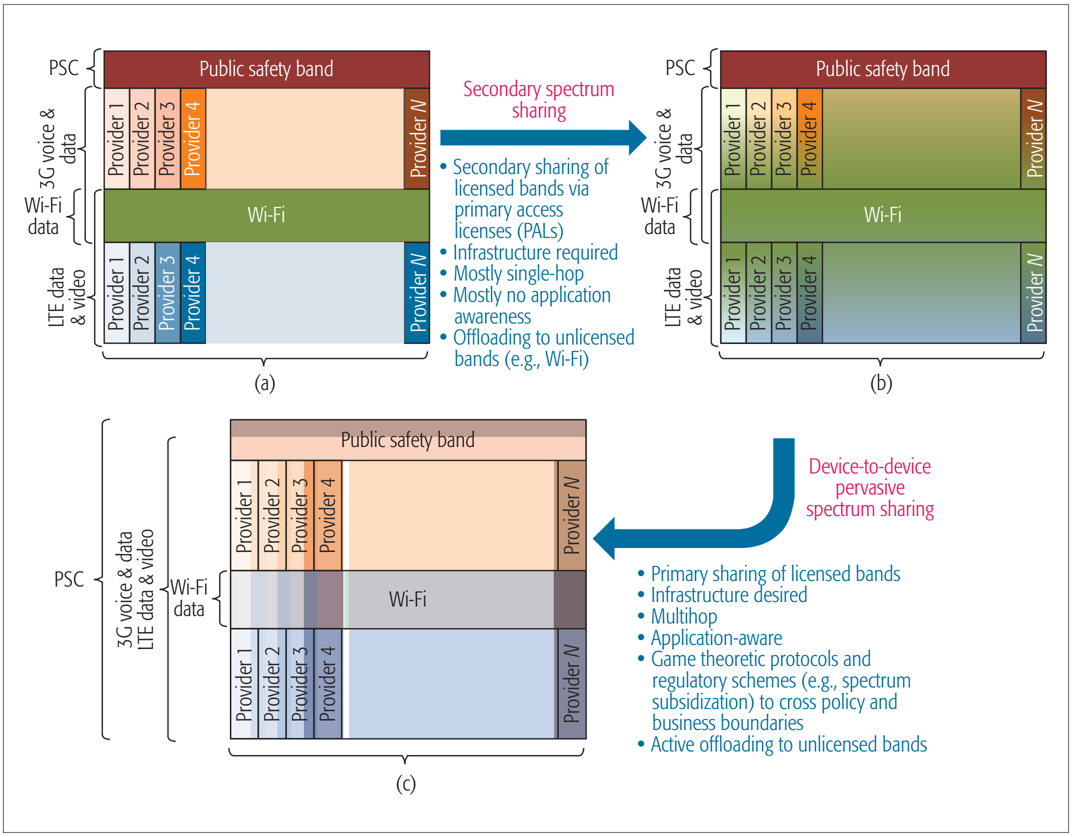

Figure 1. D2D pervasive spectrum sharing. Toward a smoother spectrum usage across licensed, unlicensed, and public safety bands: a) scattered and bordered spectrum with no sharing; b) cross-provider sharing of secondary spectrum when the primary user is idle and offloading 3G and LTE data to WiFi, relieving the spectrum scarcity to some extent; c) cross-provider sharing of primary and secondary spectrum. Coping with critical needs of future services like PSC requires much more active and pervasive sharing at the primary level (e.g., even when the primary user is busy), and across licensed, unlicensed, and restricted bands such as the public safety band.

PSC for threat prevention and emergency response involves swift usage of available resources. Hallmarks of such a PSC situation are high node density, partial availability of heterogeneous network infrastructure, the urge of surviving against attackers, further/cascading emergency events, and a heavily congested spectrum.

- Incentivization of providers and users to cooperate and share their resources over multiple hops
- Policy decisions and regulations to foster more sharing at all levels from regulatory bodies to the device users
- Seamless D2D negotiation and sharing of wireless connectivity and the spectrum
- Formation and design of multiple, coexisting, and interdependent spectrum sharing groups over large and possibly infrastructure-less areas

The following PSC scenarios illustrate the need for D2D-based PSS.

Scenario I: Scarce Capacity — Trying to Reach Infrastructure Nodes. PSC for threat prevention and emergency response involves swift usage of available resources. Hallmarks of such a PSC situation are *high node density*, partial availability of heterogeneous network infrastructure, and the urgency of surviving against *attackers*, *further/cascading emergency events*, and a *heavily congested spectrum*. As seen in Fig. 2a, each device seeks to reach a close infrastructure node (e.g., access point) to communicate with its destination, for example, to report a scene to an official or contact a loved one. However, the problem of composing usable end-to-end paths is complex, and requires fast and seamless settlement of which device is going to use which resource over a highly dynamic topology. It is further complicated as each device has different capabilities and can only use a certain set of spectrum substrates.

Scenario II: Scarce Power — Trying to Reach Public Safety Officials. In more devastating situations, communicating with a public safety official can make the difference between life or death. Consider devices stranded in rubble after an earthquake. Hallmarks of such a situation include infrastructure-less operation, fast discovery, and, most importantly, using device power wisely. The key metric for PSC in such cases is the outage probability or energy efficiency. The number of devices will likely be sparse; thus, capacity will be less of a concern. But the availability of multiple substrates to contact a nearby public safety official is vital, as seen in Fig. 2b. The devices must resolve among each other how to schedule and use heterogeneous substrates for reliable and low-power communication.

Next, we outline the PSS vision and its architectural principles. Then we discuss challenges in realizing PSS's principles in legacy PSC systems and offer various ideas to tackle them. Finally, we conclude the article.

PSS Vision and Architectural Principles

Given the recent saturation of the licensed radio bands, the adoption of new policies and princi-

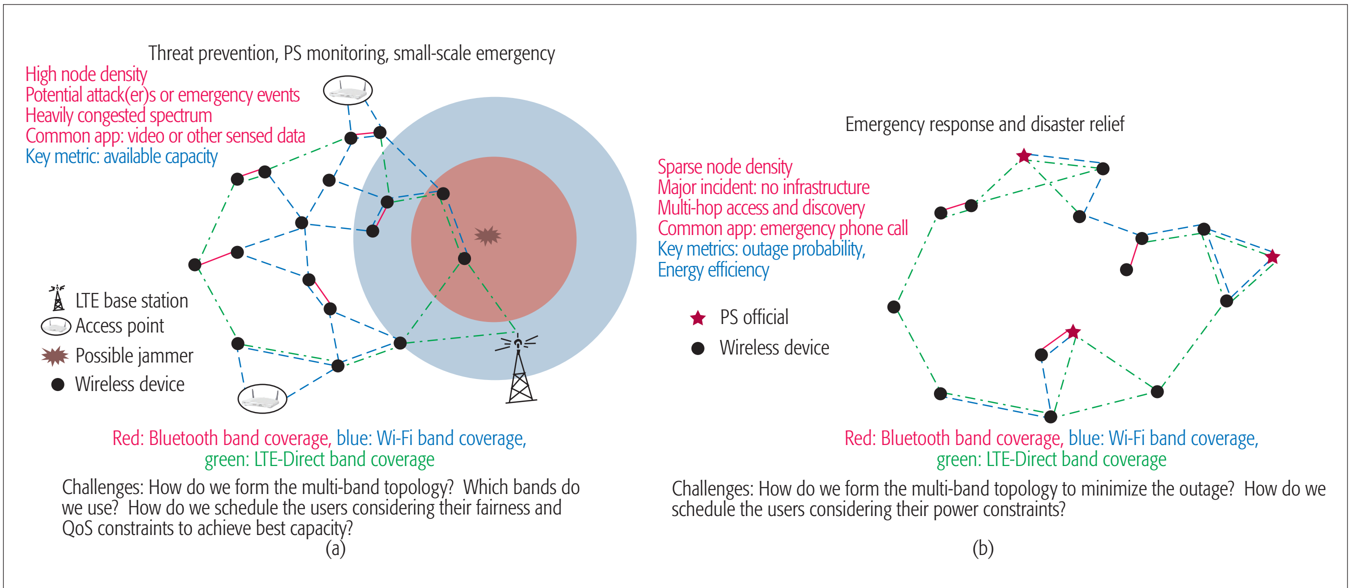

Figure 2. PSC scenarios where D2D PSS will be beneficial: a) scenario I: scarce capacity; b) scenario II: scarce power.

ples is required to address the future needs of PSCs. We envision the following hallmarks of a PSS environment to serve the future PSC applications:

Providers Are Motivated to Share. Currently, spectrum sharing happens at secondary levels, that is, if the primary user is not occupying the channel, another user can use it. Effective utilization of the spectrum is heavily dependent on cross-provider sharing, not just at the secondary level but at the primary one too.

Sharing Is the Norm. Protocol and market designs must be revamped into regimes where sharing is the norm rather than being opportunistic. Wireless protocols and business models sharing spectral resources should retain larger value and generate more revenue.

Government Power Is Wisely Used to Incentivize More Sharing, Particularly for the "Greater Good." Performance-based governmental support and policies should be put in place to foster sharing of the spectrum so that end users can enjoy a better quality service. This is particularly important for PSC applications (e.g., 911 calls and emergency response) that serve the good of the whole society.

These characteristics indicate a spectrum management vision where sharing is pervasive. Attaining such visionary goals requires the following architectural principles for future wireless and PSC systems.

Bottom-Up when Seeking Lower-Level Optimizations. D2D is a great way to discover and exploit spectrum sharing across users and providers. Involving centralized solutions to optimize a local situation may easily become prohibitive due to overhead. Furthermore, PSC scenarios with no or little infrastructure availability force local designs like D2D systems.

Top-Down when Trying to Enforce a Sustainable "Larger Good" Policy. Stakeholders of a multi-owner system like the wireless service provisioning ecosystem rightfully compete for more revenue. Although this competition ensures a healthy market, it can become too aggressive in optimizing the individual benefit at the risk of the larger good. Sharing of a precious resource like spectrum thus requires well designed top-down approaches to policy and regulation. Governments and regulatory agencies must maintain policies for incentivizing operators to share.

Game-Theoretic Designs when Crossing Trust and Administrative Boundaries at Scale. Success of a highly participatory sharing system heavily depends on the incentive (or even urge) of individual device owners. Naturally, most device owners will want to be free riders, unwilling to share their devices' resources. As observed in peer-to-peer systems (e.g., BitTorrent), game-theoretic designs are successful in enforcing sharing at scale. D2D protocols must incorporate simple and effective negotiations to seamlessly form coalitions on the fly so that two conflicting goals can be achieved:
- Fast and efficient wireless downloads/transfers for a device
- Sharing of resources

Challenges Ahead

How to Converge on Best Policies and Regulations

In the current national PSC system, one of the greatest challenges relates to spectrum sharing policies. The U.S. government is seeking to make spectrum more available for mobile use and other services involving wireless broadband technologies [7]. Regulations should allow for growth of wireless and mobile broadband networks to modify and generate new spectrum sharing regulations while also exploring the impact on the effective and efficient utilization of wireless systems. Key assumptions for a model of spectrum sharing include:
- Authority over and responsibility for the PSC system is given to local governments.
- Responsible authorities are limited in the ability to connect and make use of commercial networks for wireless services.
- There are regulations, and spectrum and needed equipment must be dedicated entirely to PSC.

- The principal application is narrowband real-time voice communications [3].

Current practices include authorization by the Federal Communications Commission (FCC) or the President in specific circumstances, and a return to a potential market-based distribution approach. Such a market-based approach was originally developed in the late 1950s with spectrum considered as property. This was efficiently implemented by private users who were considered the best for management purposes since it was assumed that they internalized benefits and costs, and would sell valuable bands to assist the economy. Issues with this exclusive market-based approach soon surfaced regarding license allocation and costs. Although exclusive access eliminated interference, license distribution removed the capability of sharing and limited access. However, tension surfaced between primary and secondary users regarding performance and protections.

According to the President's Council of Advisors on Science and Technology (PCAST), a more constructive management system utilizing allocations and incentives pertaining to spectrum in a market system can use a three-tier interference protection: incumbent, secondary, and general authorized access. The authors in [8] proposed a two-stage pricing combination. The first uses a sound static pricing policy that sets a specific level of commercial traffic. This is followed by an optimal dynamic policy for admission control. The benefits of such a combination include efficient spectrum sharing without requiring additional availability, more stable revenue between commercial networks and users, and an ability to adapt quickly to network conditions [8]. The current management systems of spectrum sharing include the spectrum access system (SAS) and the emergency response interoperability center (ERIC). A SAS allows spectrum allocation between commercial and federal entities, while an ERIC is a committee-based partnership to establish a common technical framework through issues of security, roaming, and priority access. The FCC is responsible for conducting an incentive auction to reallocate spectrum for mobile broadband uses and funded FirstNet,[4] which is the first high-speed nationwide broadband network dedicated to public safety.

Regardless of policy, a pervasive notion of spectrum sharing is dependent on a shift in mindset from the traditional operators [7]. Operators must adapt and utilize cognitive technologies to navigate dynamic spectrum availability. Making decisions regarding spectrum sharing regulations affects a multitude of stakeholders due to band availability along with long-term and short-term needs as well as variations between licensed and unlicensed bands [7]. Acknowledging who is utilizing spectrum is an important aspect for government to be aware of when generating policies for effective spectrum sharing in response to unexpected public safety challenges.

More policy adaptations occurred during the establishment of the Department of Homeland Security (DHS) in 2003, and Title XVIII of the Homeland Security Act of 2002 leading to the establishment of the DHS Office of Emergency Communications. The evolution continued with the National Preparedness Goal promoting *shared responsibility* across all sectors as well as a Quadrennial Homeland Security Review identifying threats with strong implications for national resilience and preparedness. The National Response Framework (NRF) provided a template for agencies to determine appropriate levels for federal involvement regarding domestic incidents. Moreover, this plan supported harmonization and an inter-agency incident management system to handle determined incidents of national significance.

To fund this venture, responsibility was given to the FCC to conduct a two-sided auction for spectrum reallocation and to continue development of their main emergency communications components:

- The 911 call processing and delivery system
- The emergency alert system
- The radio/broadcast or television system

In addition, the NBP [1] was developed to strategize a 10-year implementation plan for a PSC infrastructure. The NBP is a multi-faceted approach to wireless infrastructure through:

- Hardened radio access network infrastructure to enable a higher degree of coverage and resilience
- Priority roaming on commercial networks for additional capacity and increased network resilience
- Mobile technology for coverage during failures or remoteness

The collection of these services influences the broadband ecosystem in four ways:

- Maximizes consumer welfare, investment, and innovation through policies designed for robust competition
- Encourages competitive entry and network upgrades through government influences or controls to ensure management and efficient allocation
- Boosts adoption and utilization and ensures affordability through reform relating to current deployment of universal service mechanisms
- Maximizes benefits for various sectors through policy, standards, incentives, and law reform

Since events like 9/11, emergency and disaster management planning has focused on enhancing and managing collaborations between stakeholders regarding access and operation [9]. In addition, the integration of policies and procedures is challenging and requires a great deal of time. Once strategic plans are in place, policy makers must begin to predict future needs, such as changes due to population and terrain, as gaining access to spectrum and connecting infrastructure will, at some point, compromise public safety objectives. Regardless of the challenges, the development and growth of a national public safety system is not a hopeless cause, as seen through the coordination of response agencies during the Boston Marathon bombings [9].

Once strategic plans are in place, policy makers must begin to predict future needs, such as changes due to population and terrain, as gaining access to spectrum and connecting infrastructure will, at some point, compromise public safety objectives

How to Incentivize Providers

A major impediment for PSS is the providers' tendency to protect the bands they earned with a lot of licensing and operating costs. Adopting new technologies to facilitate D2D spectrum

[4] FirstNet: First Responder Network Authority, http://www.firstnet.gov

Beyond incentivizing providers, there is also a need to incentivize the users themselves to share spectrum resources. In particular, the wide-scale use of D2D communication is of paramount importance in public safety scenarios where it is likely that the infrastructure will be damaged.

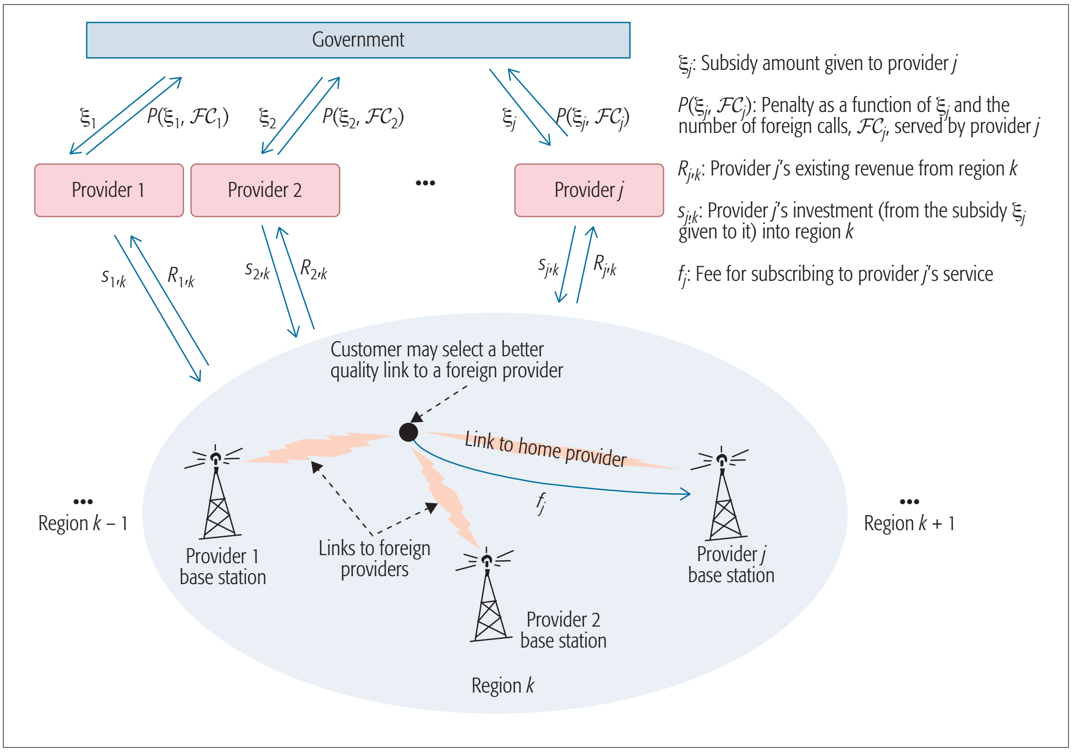

Figure 3. Model for a performance-based spectrum subsidy market.

sharing has become a key policy consideration for spectrum management. Current spectrum usage [10] heavily follows competitive auctions, which can balance standardization trends, Current policies [1] dictate that competitive auctions must remain intact, while simultaneously incorporating new ways to share and manage spectrum usage (e.g., via D2D). In such policies, developing *a governance structure for public safety broadband networks* and *making more spectrum available for PSC* is critical [10] and necessary for enabling heterogeneous spectrum sharing.

To foster more sharing of the spectrum via large-scale D2D, regulatory power can be introduced. In fact, the NBP [1] recommends the widespread development of the concept of "spectrum subsidy," for example, licensing of the D block for commercial use if *public safety partnerships* are considered by the licensee. Here, we leverage this idea of subsidizing the spectrum to the providers with lower costs in return of "proof of sharing." Thus, providers will be offered discounted bands, potentially at different locations, but will be asked to cover users not subscribed to them to maintain their subsidy incentives from the government (i.e., to sustain spectrum sharing via D2D).

Recent studies suggest significant market and user welfare gains under such subsidization (e.g., data subsidy for offering minimal data plans to users for free) [11]. To understand spectrum subsidization, we introduce a game-theoretic market model with three types of players, as shown in Fig. 3: customers, providers, and the government. Customers are end-user devices spread out to regions who engage in localized spectrum sharing markets, as discussed in the next subsection. Customers are subscribed to a "home" provider. In a quest for better experience, they are given the option to dynamically select another provider's base station if the signal quality may be better.

Providers operate in all regions and receive monetary subsidies, which they use for improving their infrastructure in regions where their service quality is weaker. After a subsidy interval (e.g., a month), a provider j may have to return some or all of its subsidy, ξ_j, back to the government if its "sharing performance" was not good. As a proof of sharing to avoid the penalty, provider j keeps track of the number of foreign customers it served, $\mathcal{FC}_j$. The penalty, $P(\xi_j, \mathcal{FC}_j)$ is a monotonically increasing function of the proof of sharing, $\mathcal{FC}_j$. In such markets, provider j solves the following optimization:

$$\max_{\{s_{jk}\}, f_j} \sum_{k=1}^{K} R_{jk} + (\xi_j - P(\xi_j, \mathcal{FC}_j)) - \sum_{k=1}^{K} s_{jk} \quad (1)$$

where the first term is the total revenue, the second is the leftover subsidy money after the penalty is deducted due to less than required sharing of spectrum, and the last term is the total amount of expenses the provider uses from its subsidy. The provider's controls are the amount of investment it makes into a region k, $\{s_{jk}\}$, and the subscription fee, f_j, that it charges to its customers.

Finally, the government's decision variables are ξ_j and the penalty function $P(.)$. An important feature of this model is that the government motivates the providers to give service to customers who are far away from the base stations of their *home providers*. The government aims to motivate foreign providers by reducing each provider's subsidy if the provider does not service enough "foreign" customers. This penalty motivates the providers to serve foreign customers at the primary level, sometimes instead of their own customers. Hence, the penalty function $P(.)$ makes this subsidization scheme a "performance-based" one, so providers that attained more sharing of their licensed spectrum bands will have to return less of their subsidy to the government.

Our preliminary results show the impact of subsidy on provider revenue for a simple two-provider scenario. We considered two regions with two providers having 10 and 60 users, which corresponds to a weak and a dominant provider, respectively. The government's total subsidy budget is set to $\xi_1 + \xi_2 = 500$. We looked at two scenarios with $\beta = 20$ and $\beta = 40$, where β denotes the average number of calls made by a user via its home provider. Assuming users' calls via providers are proportional to their monetary and hence infrastructural strength [12], we solved Eq. 1 to find the optimal revenue. Figure 4 shows the dependence of the individual and total provider revenue on the government subsidy. Results show that the subsidy monotonically improves the revenue, and the government can control (ξ_1, ξ_2) to facilitate spectrum sharing between providers, while guaranteeing their profit. Our recent work [12] further showed that subsidization will:

- Be more beneficial for smaller providers, allowing them to compete better against large providers
- Motivate providers to invest in regions with weaker coverage

How to Incentivize Users

Beyond incentivizing providers, there is also a need to incentivize the users themselves to share spectrum resources. In particular, the widescale use of D2D communication is of paramount importance in public safety scenarios where it is likely that the infrastructure will be damaged. In such scenarios, the key challenges include:

- Neighbor discovery
- Enabling multiple levels of cooperation between devices ranging from sharing spectrum to performing standard cooperative transmission
- Analyzing how the devices can interact with one another and form D2D groups

For neighbor discovery, traditional D2D typically relies on detecting uplink cellular transmissions. However, in PSC such detection may not be possible due to lack of infrastructure and thus the lack of any uplink transmissions. Here, one can develop new techniques built on some concepts that are routed in ad hoc networks. For example, rendezvous techniques that rely on temporary traffic ad hoc control channels can be used. Alternatively, devices can use historical data from D2D communication or historical encounters to attempt to discover their D2D neighbors.

Due to its promise in proximity services (ProSe) and PSC, D2D device discovery has also received significant interest from 3GPP. In 2012, a new study item was created to study LTE ProSe, and its initial focus was D2D user discovery. In particular, in a typical communication environment, users have to select discovery resources and transmit discovery signals so that they can be identified by other nearby users. To this end, in [13], we compared the performance of three different discovery techniques considering 3GPP-compliant simulations: random, greedy, and centralized discovery resource selection. In the random approach, each user randomly selects a discovery resource to transmit its discovery signal, which may result in collisions if the same resources are used by multiple nearby users. The greedy approach, on the other hand, selects resources with minimal interference levels. Finally, the centralized approach centrally assigns discovery resources to users, assuming that locations of all users are known at the centralized scheduler.

In Fig. 5, we show the performance of the

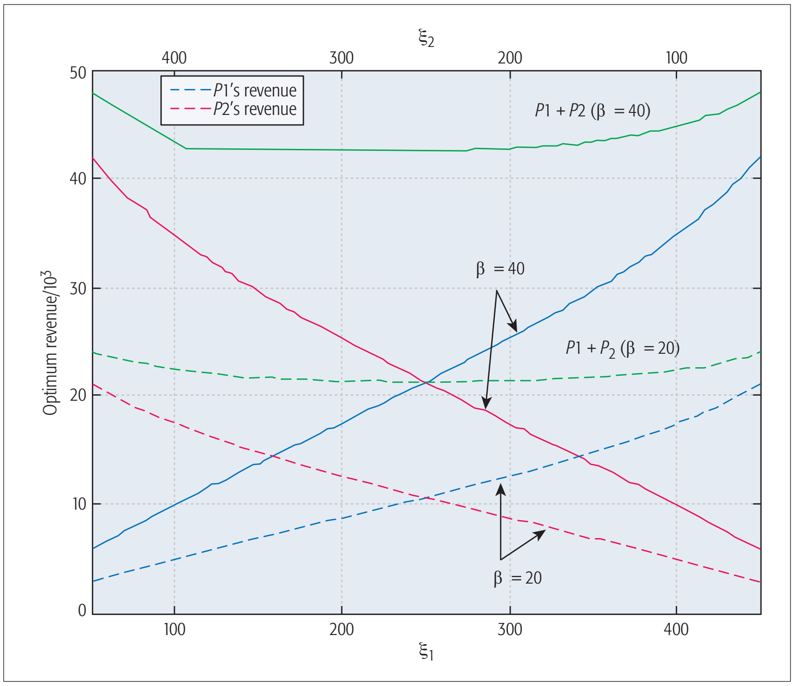

Figure 4. Provider revenue vs. subsidy.

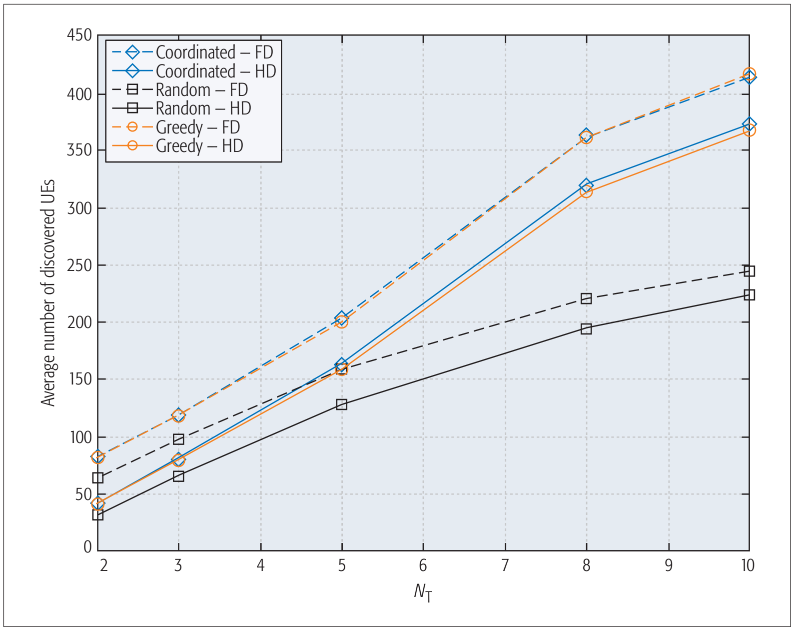

Figure 5. Number of discovered users vs. discovery symbols N_T in 3GPP system-level simulations.

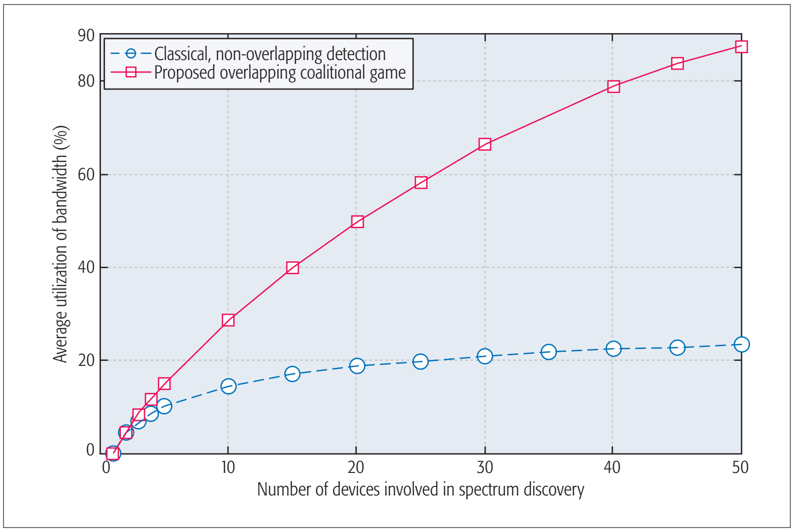

Figure 6. Benefits of overlapping coalition formation for spectrum detection.

D2D discovery algorithms in terms of average number of discovered user equipments (UEs) vs. the number of discovery opportunities, N_T. We consider a macrocell consisting of three sectors with different ISDs for each simulation layout. A hotspot simulation layout that may happen occasionally in public safety scenarios is also considered. Here, based on the simulation assumptions agreed on in 3GPP [13], two-thirds of the UEs are dropped within a circle with radius of 40 m, and the remaining UEs are uniformly distributed. The solid lines in Fig. 5 represent the half-duplex (HD) discovery case and dashed lines the full-duplex (FD) scenario. For each D2D discovery algorithm, it can be observed that the FD case always outperforms the HD case since devices transmitting in the same slot can also be discovered in FD. For small N_T, the algorithms show similar performance. For larger N_T, the random discovery selection method shows the weakest performance due to collisions during random discovery resource selection of UEs. The coordinated discovery resource selection algorithm, which is a centralized algorithm based on path loss estimation using known D2D distances, obtains slightly better performance than the greedy approach, which is a distributed approach based on known received powers.

Once D2D discovery is done, the next step is to effectively share resources and cooperate. Here, cooperation can be at different levels and network layers. For example, the devices can simply cooperate to form a local D2D LAN to share information (e.g., a local D2D LAN between first responders), or they can cooperate to form a spectrum market. Similarly, devices can cooperate to relay each other's data over multiple hops to disseminate certain emergency messages. In some cases, only part of the infrastructure is damaged, so users can cooperate to use spectrum-sensing-like techniques to detect neighbors. However, since only a few neighbors will access the infrastructure, the devices will need to cooperate to improve their detection capabilities in the presence of limited infrastructure access. Clearly, device cooperation for PSC must satisfy key characteristics:

- Lack of infrastructure
- Presence of multiple coexisting and interdependent cooperative groups
- Heterogeneity in terms of resources and node types
- Distributed decision making such that each device can, during an emergency, individually decide on which resource to share and with which neighbors

This, in turn, requires introducing new self-organizing approaches to incentivize users to form cooperative groups, or *coalitions*, to share resources, and cooperate at multiple levels. Here, one suitable framework is that of *coalitional game theory*. Coalitional game theory enables multiple devices to individually weigh in the mutual benefits and costs of cooperation and then decide on whether to cooperate or not. Despite the surge of works on coalitional games for wireless networks, most existing approaches have one limitation: they assume that users can only belong to one coalition. In public safety scenarios, users can share resources with multiple coalitions simultaneously. To this end, one must expand existing models to account for *overlapping coalition formation* cases in which a device can belong to multiple coalitions simultaneously. While the mathematical details of overlapping coalition games are outside the scope of this article (the reader is referred to [14] for one possible approach), it is important to note that such game models can yield a suite of algorithms that can be used by devices to autonomously form cooperative groups, which can include spectrum markets, multihop communication, cooperative sensing, or simply cooperative formation of overlapping D2D LANs.

Using an overlapping coalitional game, one can characterize how a PSC network can autonomously form local D2D communication pairs to share resources and incentivize users to forward each other's packets. Several design questions must be addressed such as how to model mutual benefits and costs that pertain to QoS metrics such as energy, rate, and even neighbor discovery performance; how to handle the interdependence between user-level resource sharing and providers' participation; and how to ensure that cooperation is beneficial not only to individual users, but also to the public safety system as a whole. Moreover, within each formed coalition, devices may be engaged in other optimization or game-theoretic mechanisms. Therefore, one must build multi-level games that include an overlapping coalitional game with underlaid uncooperative or even auction games for resource sharing within each coalition. Last but not least, public safety scenarios may require the deployment of mobile base stations that are integrated in public safety personnel cars or even on unmanned aerial vehicles. The interdependence between such mobile base stations and D2D formation is critical in PSC, as studied in our work in [15].

To illustrate the benefits of overlapping coalition formation, in [14], we adopted this framework to allow neighboring devices to collaboratively detect available spectral bands by sharing spectrum sensing results. These results are then collectively combined within a coalition to get a final decision on whether a certain band is vacant or not. In this model, a device may share its sensing result with multiple coalitions simultaneously. This collaborative detection automatically yields better spectrum usage. Figure 6 shows that additional sharing of spectrum detection results via an overlapping coalitional game formulation can yield significant gains in terms of the percentage of bandwidth (i.e., spectrum) utilization, compared to traditional collaborative approaches with no overlapping coalition.

Clearly, large-scale cooperation between devices of a PSC system is a critical challenge that must be addressed in order to generate a new breed of systems with users who can autonomously form coalitions and cooperate effectively under infrastructure-less scenarios.

Clearly, large-scale cooperation between devices of a PSC system is a critical challenge that must be addressed in order to generate a new breed of systems whose users can autonomously form coalitions and cooperate effectively under infrastructure-less scenarios.

Summary

In this article, we have studied the potential of PSS for PSC applications. We have outlined three hallmarks for successful PSS: providers are motivated to share; sharing is the norm; and government power is wisely used to incentivize more sharing. We have shown how the realization of PSS requires multiple architectural principles to be adopted: bottom-up approach when seeking lower-level oa ptimizations, a top-down approach when trying to enforce sustainable and "larger good" policies, and game-theoretic designs when crossing trust and administrative boundaries at scale. In a nutshell, we have shown that moving toward a new breed of PSC systems requires overcoming key technical and regulatory challenges that include how to converge on best policies and regulations, how to incentivize providers, and how to incentivize users.

References

[1] FCC, "Connecting America: The National Broadband Plan," Mar. 2010; http://download.broadband.gov/plan/national-broadband-plan.pdf.

[2] L. E. Miller, "Wireless Technologies and the SAFECOM SoR for Public Safety Communications," tech. Rep, NIST, Gaithersburg, MA, 2005; http://w3.antd.nist.gov/pubs/WirelessAndSoR060206.pdf.

[3] J. Peha, "Sharing Spectrum through Spectrum Policy Reform and Cognitive Radio," *Proc. IEEE*, vol. 97, no. 4, 2009, pp. 708–19.

[4] X. Lin *et al.*, "An Overview of 3GPP Device-to-Device Proximity Services," *IEEE Commun. Mag.*, vol. 52, no. 4, 2014, pp. 40–48.

[5] Y. Zhang *et al.*, "Auction Approaches for Resource Allocation in Wireless Systems: A Survey," *IEEE Commun. Surveys & Tutorials*, vol. 15, no. 3, July 2013, pp. 1020–41.

[6] G. S. Kasbekar and S. Sarkar, "Spectrum Pricing Games with Spatial Reuse in Cognitive Radio Networks," *IEEE JSAC*, vol. 30, no. 1, 2012, pp. 153–64.

[7] L. Doyle and T. Forde, "A Regulatory Perspective on Cognitive Radio and Spectrum Sharing," *Cognitive Radio and Networking for Heterogeneous Wireless Networks*, ser. Signals and Communication Technology, Springer Int'l. Publishing, 2015,, pp. 257–89.

[8] Q. Wang and T. Brown, "Public Safety and Commercial Spectrum Sharing Via Network Pricing and Admission Control," *IEEE JSAC*, vol. 25, no. 3, 2007, pp. 622–32.

[9] Q. Hu, C. C. Knox, and N. Kapucu, "What Have We Learned Since September 11, 2001? A Network Study of the Boston Marathon BombingsResponse," *Public Admin. Rev.*, vol. 74, no. 6, 2014, pp. 698–712.

[10] L. K. Moore, "Spectrum Policy in the Age of Broadband: Issues for Congress," CRS Report for Congress, Feb. 2010, pp. 1–28; http://fpc.state.gov/documents/organization/138723.pdf

[11] S. M. Yu and S.-L. Kim, "Guaranteeing User Welfare in Network Service: Comparison of Two Subsidy Schemes," *Perf. Eval. Rev.*, vol. 40, no. 2, Sept. 2012, pp. 22–25.

[12] M. Yuksel *et al.*, "Fostering Wireless Spectrum Sharing Via Subsidization," *Proc. IEEE Annual Allerton Conf. Commun., Control, and Comp.*, Urbana-Champaign, IL, Oct. 2013, pp. 1192–99.

[13] M. Simsek *et al.*, "Device-to-Device Discovery Based on 3GPP System Level Simulations," *Proc. IEEE GLOBECOM Wksps.*, 2013, pp. 555–60.

[14] T. Wang *et al.*, "Cooperative Spectrum Sensing in Cognitive Radio Networks: An Overlapping Coalition Formation Game Approach," *IEEE Trans. Commun.*, vol. 62, no. 9, Sept. 2014, pp. 3144–60.

[15] M. Mozaffari *et al.*, "Drone Small Cells in the Clouds: Design, Deployment and Performance Analysis," *GLOBECOM*, San Diego, CA, Dec. 2015.

Biographies

Murat Yuksel [SM'11] (yuksem@unr.edu) received his B.S. degree in computer engineering from Ege University, Izmir, Turkey, in 1996, and his M.S. and Ph.D. degrees in computer science from Rensselaer Polytechnic Institute (RPI), Troy, New York, in 1999 and 2002, respectively. He is an associate professor at the Computer Science and Engineering Department, University of Nevada-Reno (UNR). Prior to UNR, he was with the Electrical and Computer Science Engineering Department, RPI, as a postdoctoral associate and a member of adjunct faculty until 2006. He was a software engineer at Pepperdata (Sunnyvale, California) and a visiting researcher at AT&T Labs and Los Alamos National Laboratory. His research interests are in networked, wireless, and computer systems with a recent focus on big data networking, UAV networks, optical wireless, public safety communications, device-to-device protocols, economics of cyber security and cyber sharing, routing economics, network management, and network architectures. He has been on the Editorial Board of *Computer Networks*, published more than 100 papers in peer-reviewed journals and conferences, and received the IEEE LANMAN 2008 Best Paper Award. He is a Senior Member and Life Member of the ACM, and was a member of Sigma Xi and ASEE.

İsmail Güvenç [SM'10] (iguvenc@fiu.edu) received his Ph.D. degree in electrical engineering from the University of South Florida in 2006, with an outstanding dissertation award. He was with Mitsubishi Electric Research Labs during 2005, and with DOCOMO Innovations Inc. between 2006 and 2012, working as a research engineer. Since August 2012, he has been an assistant professor with Florida International University. His recent research interests include heterogeneous wireless networks and 5G wireless systems. He has published more than 100 conference/journal papers and book chapters, and several standardization contributions. He co-authored/co-edited three books for Cambridge University Press, served as an Editor for *IEEE Communications Letters* (2010–2015) and *IEEE Wireless Communications Letters* (2011–present), and as a Guest Editor for several other journals. He is an inventor/coinventor in 23 U.S. patents, and has another 4 pending U.S. patent applications. He was a recipient of the 2014 Ralph E. Powe Junior Faculty Enhancement Award and the 2015 NSF CAREER Award.

Walid Saad [SM'15] (walids@vt.edu) received his Ph.D degree from the University of Oslo in 2010. Currently, he is an assistant professor and the Steven O. Lane Junior Faculty Fellow at the Department of Electrical and Computer Engineering at Virginia Tech, where he leads the Network Science, Wireless, and Security (NetSciWiS) laboratory within the Wireless@VT research group. His research interests include wireless networks, game theory, cybersecurity, and cyber-physical systems. He was the recipient of the NSF CAREER award in 2013, the AFOSR summer faculty fellowship in 2014, and the Young Investigator Award from the Office of Naval Research (ONR) in 2015. He was the author/co-author of five conference best paper awards at WiOpt 2009, ICIMP 2010, IEEE WCNC 2012, IEEE PIMRC 2015, and IEEE SmartGridComm 2015. He was the recipient of the 2015 Fred W. Ellersick Prize from the IEEE Communications Society. He serves as an Editor for *IEEE Transactions on Wireless Communications*, *IEEE Transactions on Communications*, and *IEEE Transactions on Information Forensics and Security*.

Naim Kapucu (Kapucu@ucf.edu) is a professor of public policy and administration and director of the School of Public Administration at the University of Central Florida (UCF). He is also the founding director of the Center for Public and Nonprofit Management (CPNM) at UCF. His research interests are collaborative governance, emergency and crisis management, decision making in complex environments, social inquiry and public policy, and scholarship in teaching and learning. He teaches network governance, collaborative public management, emergency and crisis management, methodology, and leadership courses. He received his Ph.D. in public and international affairs from the Graduate School of Public and International Affairs at the University of Pittsburgh, Pennsylvania, in 2003. Prior to that, he earned a Master of Public Policy and Management degree from Heinz College's School of Public Policy and Management, Carnegie Mellon University, Pittsburgh, Pennsylvania, in 1997.

Next Generation Public Safety Networks: A Spectrum Sharing Approach

Munawwar M. Sohul, Miao Yao, Xiaofu Ma, Eyosias Y. Imana, Vuk Marojevic, and Jeffrey H. Reed

The authors discuss their vision of a flexible, rapidly deployable, and reconfigurable public safety system to meet imminent and future demands on critical communications infrastructure and discuss the value of spectrum sharing for supporting this long-term vision..

Abstract

Wireless communications play a critical role in national security and disaster relief and are extensively utilized by first responder teams, such as police officers, firefighters, and ambulances. The first responders in emergency situations need to always be connected with one another and with the remote service centers for effective cooperation and coordination. However, the existing public safety (PS) services fail to satisfy the PS user requirements in many emergency scenarios, which usually leads to exceptionally high traffic loads. As a result, increasing the network capacity is one of the primary concerns, and spectrum sharing has been deemed the key solution. With appropriate spectrum sharing partnerships among PS agencies and commercial networks, PS users can access licensed PS spectrum, shared spectrum, as well as commercial networks when the need arises. In this article we present different initiatives undertaken by Wireless @ Virginia Tech toward the next generation PS network. We also discuss our vision of a flexible, rapidly deployable, and reconfigurable PS system to meet imminent and future demands on critical communications infrastructure and discuss the value of spectrum sharing for supporting this long-term vision.

Introduction

The next generation public safety (PS) network, with its mission-critical aspect and the ever increasing demand for rich-content-based applications, requires a paradigm shift not only in sharing and managing the spectrum, but also in designing the radio architecture. Whereas commercial wireless networks are evolving at a high pace to sophisticated standards, such as Long Term Evolution-Advanced (LTE-A), the PS network has not gone through such an evolution mainly due to lack of economic incentives. PS communications still primarily use land mobile radio (LMR) system standards, such as Project 25 (P25) and terrestrial trunked radio (TETRA) [1]. These systems are mature, reliable, and cost effective in supporting mission-critical voice applications. However, they are not designed to support higher-bandwidth applications.

PS involves a very broad range of stakeholders and contexts such as law enforcement, fire response, medical emergencies, natural/man-made disasters, and a number of unexpected situations. An interesting but important aspect of a PS situation is that it begins in a local context, but more often than not ends up as a much more complex process [2]. This makes the problem of coordination more complex and affects the system designs for interoperability across a broad spectrum of actors. There are approximately 14,000 police departments, 3000 sheriff's offices, more than 6000 911 centers, 65+ fusion centers, 1.2 million employees in city, county, state, and Federal law enforcement, and 800,000 working in the private security sector in the United States [5]. Designing and developing an efficient and coordinated system for this diverse workforce is a challenging task, and the effectiveness of such a system is heavily dependent on the accuracy and detailed nature of the information shared within this complex topology.

The coordination among these different types of entities in a PS network requires more than voice communications. Public safety personnel, being used to leveraging commercial wireless technology, desire data, voice, and video communication, and access to the Internet to aid in response to particular emergencies. Public safety networks have a mission-critical aspect that places special requirements on the underlying radio access technologies [3]. The foremost requirement in a PS network is based on priority access to network resources at times of emergency. PS communications also need the ability to communicate in a push-to-talk (PTT) or device-to-device (D2D) direct communication mode. Group calling is another very important feature. All these requirements placed by the desired PS network and ever increasing demand for rich-content-based applications need additional spectrum.

Spectrum allocated for the PS network plays a significant role to ensure proper functioning of the network. The current frequency allocations in the United States (Fig. 1) assign 763–768 MHz and 793–798 MHz for base station and mobile unit use, respectively. The so-called D block would expand this allocation to include 758–763 MHz and 788–793 MHz. In addition, the PS communication requirements are also served with allocations in the 769–775 MHz and 799–805 MHz bands in 12.5 kHz narrowband increments. These latter allocations are primarily used for voice communications. The use of the

The authors are with Wireless@VT; work performed when Eyosias Y. Imana was at Virginia Tech

0163-6804/16/$25.00 © 2016 IEEE

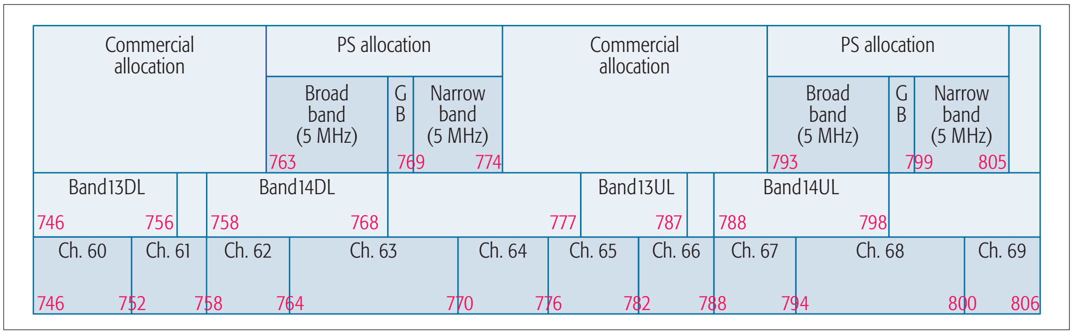

Figure 1. FCC allocated bandwidth in the 700 MHz band [2].

700 MHz spectrum for PS applications is attractive because of the beneficial propagation and penetration characteristics. However, there is practically no room available in the 700 MHz band to extend the spectrum further for future broadband PS applications.

The lack of dedicated wideband spectrum for PS networks significantly throttles the supply–demand equilibrium of the PS ecosystem. There is a huge demand for broadband PS applications from every corner of the community. Spectrum sharing has recently emerged as a potential solution for this. The 3.5 GHz band in the United States is a complementary spectrum band that could be used to supplement efforts to meet the growing capacity demands of wireless communications services. The spectrum sharing approach brings an interesting dimension into the evolution of the PS networks. The First Responder Network Authority (FirstNet), which was created under The Middle Class Tax Relief and Job Creation Act of 2012 [6] as an independent authority within the National Telecommunications and Information Administration (NTIA), aims to provide emergency responders with the first nationwide, high-speed, broadband network dedicated to public safety. In partnership with interested commercial entities, FirstNet is to provide the PS community the first nationwide PS broadband network (NPSBN). These partners will employ any underutilized PS spectrum on an as-needed basis and charge for its use, thereby becoming a true partner in this public/private partnership. Thus, the spectrum sharing approach provides additional spectrum and at the same time economic incentives to implement the NPSBN.

Responding to the PCAST recommendation of sharing up to 500 MHz of federal government RF spectrum with non-government entities [7], the Federal Communications Commission (FCC) issued a Notice of Proposed Rule Making (NPRM) and Order to open up an initial 100 MHz of spectrum (3550–3650 MHz) [4]. This new spectrum-sharing policy of the FCC requires both government and non-government radios to share the newly opened frequency band. Therefore, the new policy requires a paradigm shift not only in sharing and managing the spectrum, but also in designing a suitable radio architecture. Any system design should take into account the possibility of devices operating in multiple frequency bands. Gateway methods (e.g. RF, IP, or application layer conversions) that can bridge among frequencies may prove beneficial [5]. This article examines enabling technologies that can leverage this vision of scalable next generation PS networks.

Wireless @ Virginia Tech Solution: Frequency Translating LTE Repeater

Choosing spectrum sharing as the solution to the spectrum scarcity problem of the PS network comes with various challenges. The spectrum bands identified by the FCC as potential candidates for spectrum sharing are lacking devices, networks, and services that are willing to opportunistically use the shared bands. The PS devices available are not suitable for these bands and hence pose a significant challenge to exploring the suitability of these bands for PS services. In other words, it is extremely expensive to demonstrate the feasibility of spectrum sharing in the 3.5 GHz band because equipment that directly supports the new 3.5 GHz band is not available.

At Wireless@VT we developed a frequency-translating LTE repeater supporting dynamic spectrum access (DSA) techniques. The repeater translates the 700 MHz Long Term Evolution (LTE) band to the 3.5 GHz band and vice versa. Thus, commercial off-the-shelf (COTS) 700 MHz user equipment and femtocells in the market can be used for the feasibility study. This provides a low-cost alternative to demonstrate the feasibility of the new 3.5 GHz band for broadband PS applications.

Example Application Scenario

Our proof-of-concept demonstration analyzes the feasibility of broadband PS applications in the 3.5 GHz band using a frequency translating repeater with DSA capability. The objective was to establish some sort of two-way communications in the 3.5 GHz band between the PSuser equipment (PS-UE) and the base station (eNB). We started with the goal of establishing voice communication or text message exchange. In later phases we extended the communications to support video streaming between PS-UE and eNB. This was motivated by the fact that other than coastal areas, the probability of the primary user being present in the 3.5 GHz band is very low.

As can be seen from the presented scenario (Fig. 2), a search operation is ongoing in a locality. The PS personnel rushes to the location. Each of the PS personnel has a PS-UE that directly connects to the frequency translating repeater. The repeater communicates over the 3.5 GHz

Our proof-of-concept demonstration analyzes the feasibility of broadband PS applications in the 3.5 GHz band using a frequency translating repeater with DSA capability. The objective was to establish some sort of two-way communications in the 3.5 GHz band between the PS user equipment and the base station.

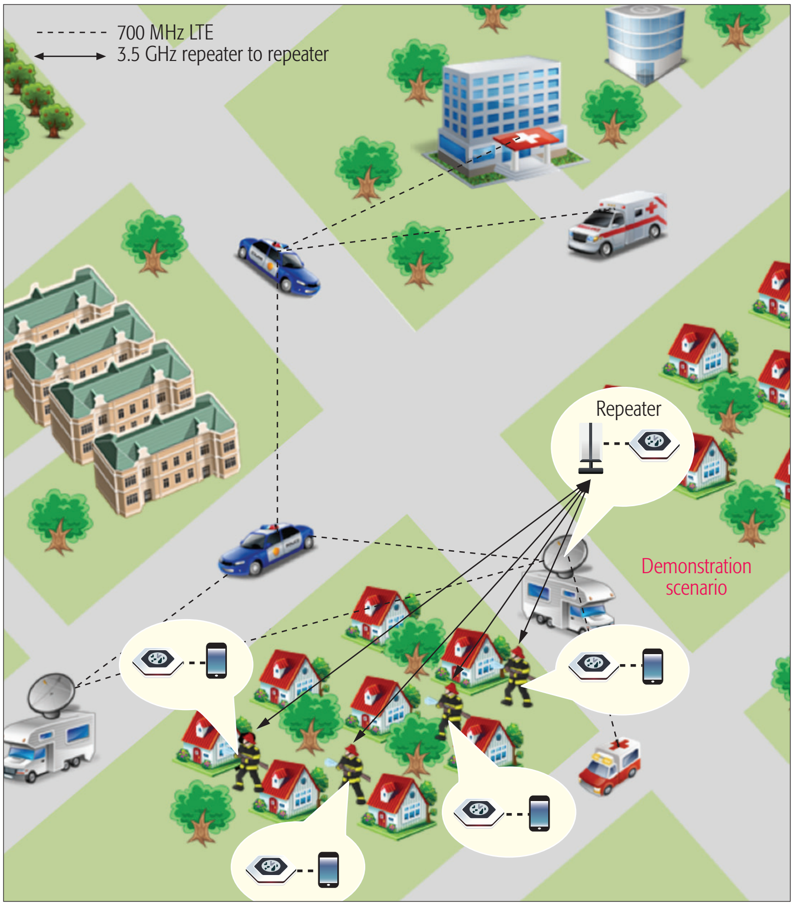

Figure 2. Application scenario for the demonstration.

band to another repeater at the other end of the wireless link. The second repeater is collocated with a femtocell. The femtocell or ad hoc base station can be placed in a vehicle near the scene of an incident. Public safety personnel might want to transmit a live video of the scene to the command post. The camera is connected to the PS-UE and through it to the repeater. The DSA algorithm used for channel monitoring finds an empty channel in the 3.5 GHz band and tunes the repeater accordingly. The repeater then transmits the video to the femtocell-repeater. At the receiving end, the repeater receives the transmission, down converts it to the 700 MHz band, and forwards it to the femtocell.

The 100 MHz bandwidth of the 3.5 GHz band is partitioned into multiple narrowband channels. An external spectrum monitor is introduced in this setup that makes the DSA decisions and communicates them to both repeaters (Fig. 3). If the band is empty, the repeater translates the 700 MHz signal to a 3.5 GHz signal. The second repeater at the other end receives the transmission and translates the signal back to the 700 MHz band. Repeater 2 is connected to the femtocell and completes the communications link between the UE and the eNB. In this way, we can establish a two-way communication in the 3.5 GHz band.

We can add an additional source that will serve as a primary user in the 3.5 GHz band. The intentionally deployed interferer, emulating a primary user, can transmit in one or multiples of these narrowband channels. The spectrum monitor will be able to perform DSA over the entire range of the 3.5 GHz band and provide adecision based on the channel occupancy that tunes the repeater.

Description of the Demonstration Setup

The purpose of the demonstration was to ensure successful operation of the FD-LTE repeaters at both ends of the 3.5 GHz over-the-air link. We used Motorola's LEX-700 [8] as the PS-UE and Rhode & Schwarz's CMW-500 [9] as the LTE eNB. We also emulated pulsating interference and implemented a spectrum monitor to detect the presence and absence of the primary user. In phase 1 of the demonstration, the objective was to observe the impact of a hopping tone interferer on the FD-LTE system in the absence of the external spectrum monitor. The tone interference sweeps the entire bandwidth of the LTE signal in random hops. In phase 2, we introduced the external spectrum monitor and observed the performance improvement. For the interference signal hopping throughout the 3.5 GHz band, we divided the entire band into 10 sub-bands. The interference signal randomly appears in one of these sub-bands at random time intervals. Since this is an FD-LTE link, we assumed four UL-DL pairs: (Ch 0–6), (Ch 1–7), (Ch 2–8), and (Ch 3–9). We used channels 4 and 5 for UL/DL separation.

Figure 4 shows the system setup for the demonstration in an indoor environment. The system consists of four major blocks:
- The PS-UE with the UE-end repeater
- The eNB with the eNB-end repeater
- The primary user emulator
- The spectrum monitor

The PS-UE and eNB have already been discussed. In the remainder of this section we describe the interference source, the UE-end and eNB-end FD-LTE repeaters, and the external spectrum monitor.

The UE-side repeater has two separate RF paths. It attaches to the PS-UE through a band 14 LTE duplexer onboard the repeater. Figure 5 shows the power budget calculation for the UE-side FDD-LTE repeater. In this power budget calculation, we considered the maximum UE transmission power as 24 dBm. Figure 5 also provides information about all components of the UE-side repeater. The eNB-side repeater also has two separate RF paths. Unlike the UE-side repeater, it has two ports to attach to the UL and DL ports of the CMW500. Using a similar calculation, we could also determine the path loss budget of the eNB-side FD-LTE repeater. For this calculation, we considered the transmission power of the eNB to be –35 dBm. The distance between the UE-side repeater and the eNB-side repeater was 1 m.

Experimental Results and Recommendations

Figure 6 shows the impact of interference on the performance of the FD-LTE link. We present the block error rate (BLER) and the throughput over signal-to-interference-plus-ratio (SINR). As expected, the BLER of the FD-LTE link decreases as the power of the interference source decreases (increase in SINR). Also as expected, the throughput of the FD-LTE link increases with increasing SINR. It is interesting to observe that there are flat regions in both curves. The LTE system tries to compensate the interference with channel quality indicator (CQI) adjustment. The flat regions can be attributed to slow CQI adjustments.

We also analyzed the benefits of employing the spectrum monitor in terms of BLER and throughput of the FD-LTE link. As expected and shown by the BLER and throughput curves, employing DSA through the external spectrum monitor improves the FD-LTE link performance. The external spectrum monitor detects the presence or absence of the primary user. Based on this information, it then helps the LTE eNB and PS-UE to switch to an empty channel by retuning the local oscillator (LO) of both repeaters. Another interesting improvement due to spectrum monitoring is the shorter recovery time after interference. The reason behind this is the reduced amount of time the LTE link is interfered by the primary user.

We explored two different approaches for implementing the DSA algorithm: hard termination and soft termination. In the hard termination approach the LTE signal is terminated to sense the PU signal. Sensing without the use of quiet periods would be possible, for example, if the PU bandwidth is much wider than that of the LTE signal. We can then detect the PU operation even without terminating the LTE signal. This approach is termed soft termination. We run two applications with different quality of service (QoS) requirements: YouTube and Skype. YouTube works fine with hard termination DSA. The periodic termination does not have a noticeable impact because video streaming uses buffering. However, the same is not true for Skype.

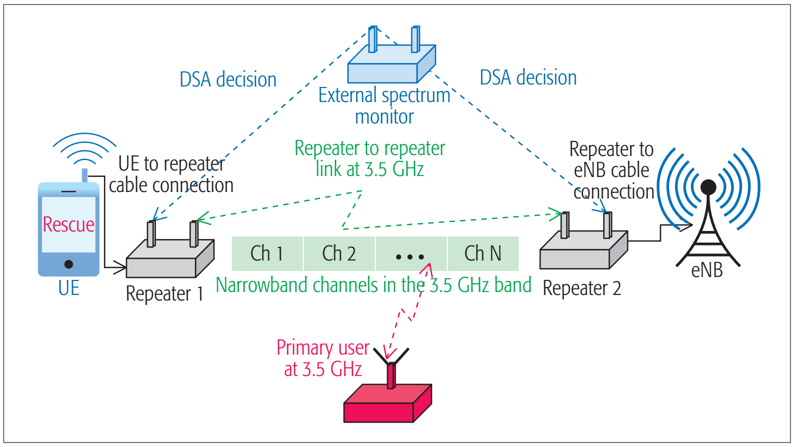

Figure 3. Application scenario of frequency-translating repeater and femtocell.

Hard termination DSA reduces the quality of voice communication. When soft termination is used, however, even Skype can work seamlessly in shared spectrum (Fig. 7). This observation presents an interesting research problem. Based on the QoS requirement of the secondary user's operation and the operational characteristics of the primary user, a matching problem can be formulated with an objective to maximize the sharing experience for both parties.

Next Generation Public Safety Network

Based on our experience of successfully demonstrating the benefits of spectrum sharing for PS networks, we envision the next generation PS networks to be flexible, easy to deploy, reconfigurable, and cognitive. One of the important aspects of the desired PS network is that it has to be able to be deployed fast and adapt to unex-

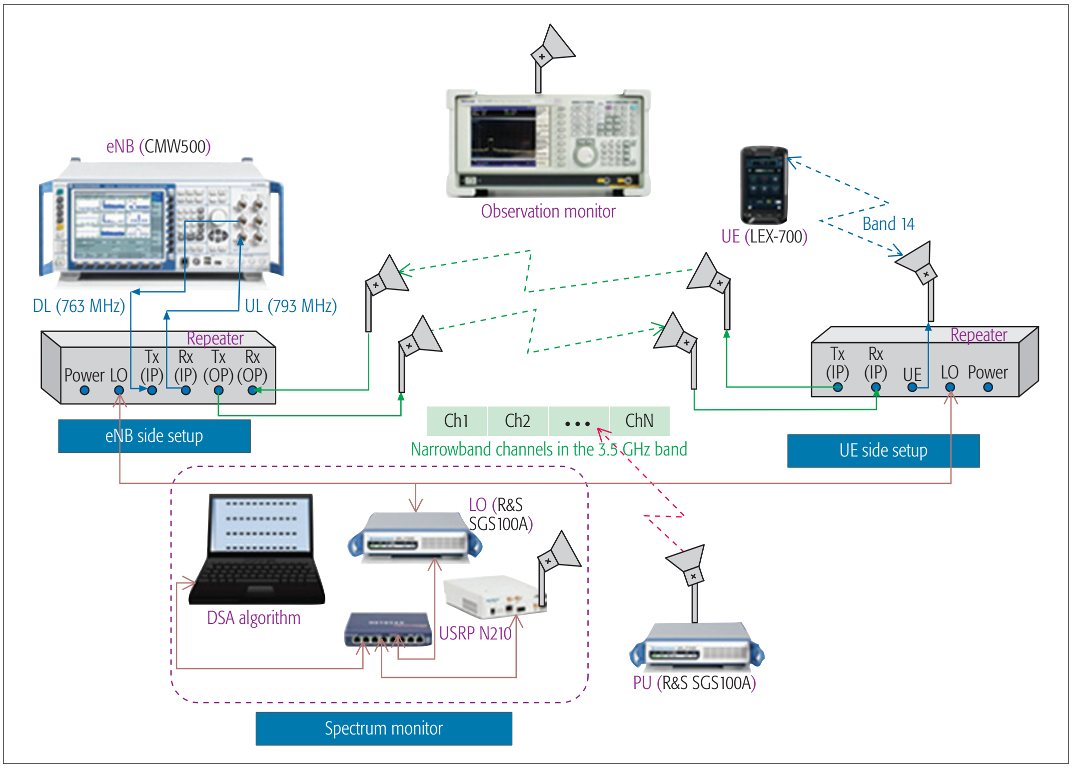

Figure 4. System setup for the FDD-LTE communication with repeaters over the 3.5 GHz band.

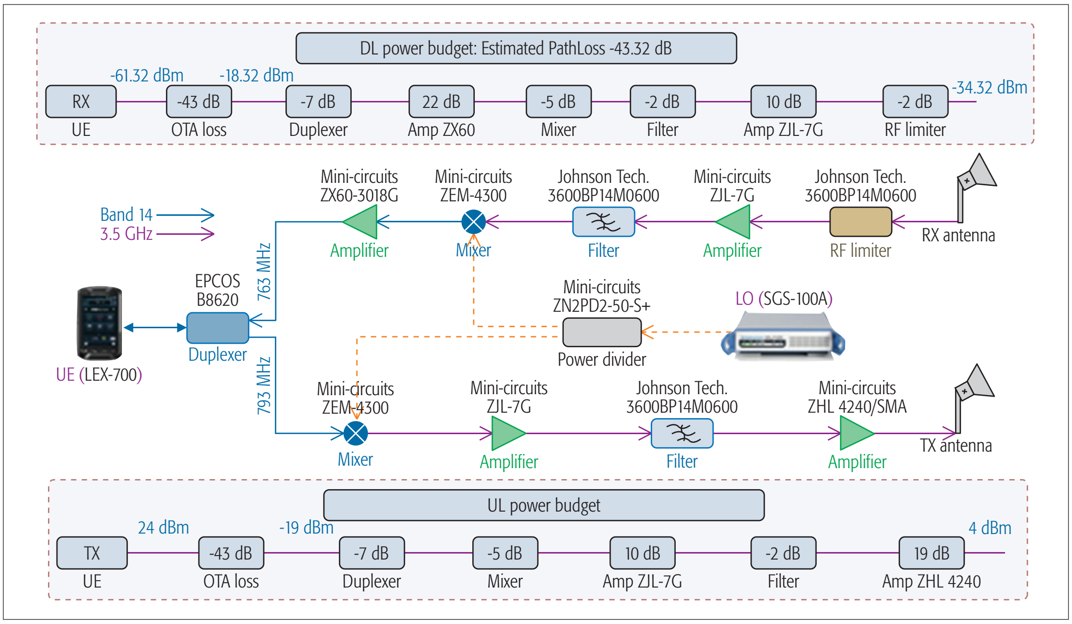

Figure 5. Power budget analysis for the UE-side FDD-LTE repeater.

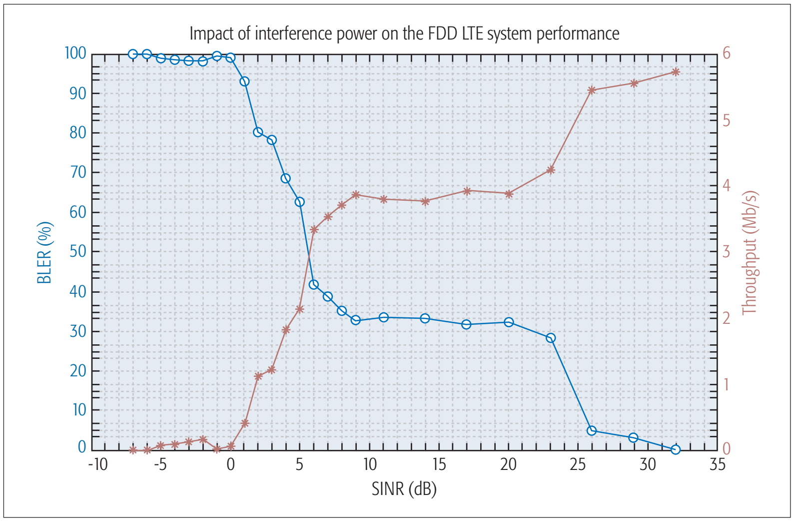

Figure 6. Impact of interference on the FD-LTE link performance.

pected events. We propose a low-cost and sustainable multi-functional mobile station concept to provide multiple wireless services in underserved localities using the new 3.5 GHz band. The system will be set up in a van or truck to ensure the added advantage of mobility. By assigning appropriate priority, the system will be able to address the needs of disaster and non-disaster scenarios, including underserved rural areas.

In order to provide flexibility in deployment and operation, the system architecture is envisioned to support over-the-air (OTA) access to the deployed mobile nodes for reconfiguration and management purposes. It also needs to provide flexibility for data collection and other operating modes. Operational data for any PS situation can be locally stored at the nodes or collected at the control center. The flexibility in the operating mode comes with the system's ability to support both standalone and group operation. The system architecture will be designed to facilitate system reconfiguration capability. The system will use a self-organizing network (SON) concept to achieve flexibility in operating modes.

The proposed system has a centralized architecture with a spectrum access system (SAS) [10] serving as the governing entity (Fig. 8, top). The deployable mobile nodes will serve as the secondary access point (S-AP) or secondary eNB. These S-AP nodes will have the additional capability of spectrum monitoring. The information gathered by these S-APs will be forwarded to the central SAS in the control center. This communication can be done using a commercial fourth generation (4G) data service or direct transmission to the control center. Based on the information received from the spectrum monitors and regulatory database, the SAS will manage the available spectrum opportunities and provide access to the S-APs. The resource scheduling within the secondary cell can be done locally by the S-APs or centrally by the control center.

Over-the-Air Data Collection Capability

The system architecture provides flexibility for storing operational data. The data collection can be done in two ways: If the situation requires analysis of critical, real-time data, the secondary eNBs will use commercial data services to report them back periodically to the control center server. For non-critical data, the system will assume a store and report approach. Each of the mobile nodes will be equipped with a memory device, a local log book that will serve as a temporary storage platform. The mobile nodes will record the experimental results in their memory devices. This will help reduce the reporting overhead. Once the situation is contained and the mobile nodes return to the control center, the log book entries will be automatically transferred into the central result log server of the control center (Fig. 8, bottom). Based on the analysis of the collected data, the PS personnel will be better prepared to address similar situations in future.

Flexible Operating Mode

In order to provide PS personnel freedom during the operation planning phase, the system will support both standalone and group modes of operation. In the standalone mode, the mobile nodes might be used as secondary eNBs interacting with the SAS engine or as spectrum monitors to observe available spectrum opportunities.

The deployable mobile nodes will also be able to work in a group. We plan to use SON concepts, such as self-configuration, coverage and capacity optimization, and mobility robustness optimization, to process the interactions between the mobile nodes and the control center and among the mobile nodes themselves. The SON will take advantage of the deterministic aspect of the deployment. The PS personnel will know the topography, infrastructure, and road/highway coordinates of the area, the number of mobile nodes deployed, the role of each node, and their operating parameters, such as allowed transmission power and communication capabilities. Using this information, the control center and the S-AP mobile node will be able to track and interact with other group members. Member mobile nodes of each group will also know their access point. This will allow the system to have automated and closed-loop optimization capability to adapt to different contexts.

Over-the-Air Reconfiguration Capability

The proposed system also aims to offer flexibility in terms of reconfiguration capability. As modification in hardware is expensive to address different PS situations, the proposed system will focus on providing a reconfigurable software platform. The desired reconfiguration capabilities include the following.

Operation Mode Selection: Depending on the requirement of the PS scenario, the operating mode of the mobile nodes will be set to standalone or group mode.

Role of Individual Nodes: In either of the operating nodes, individual nodes might assume different roles, such as a secondary eNB, external spectrum monitor, or relay.

Technology Selection: The mobile nodes will be equipped with multiple communication technologies, such as cellular and Wi-Fi. Depending on the requirement of the scenario, the appropriate communication technology will be selected.

Repeater Reconfiguration: The target frequency of the repeater can be reconfigured to accommodate the experiment design.

Power Management: The maximum allowable transmission power for individual nodes can be reconfigured before deployment.

Programmable Scheduling Model: The operational command will be able to select from a list of predefined scheduling algorithms. It will also be able to add its own scheduling module to validate different shared spectrum scheduling schemes.

Critical Data List: The control center will be able to modify the critical data list according to the need of the PS scenario in the planning phase of the operation.

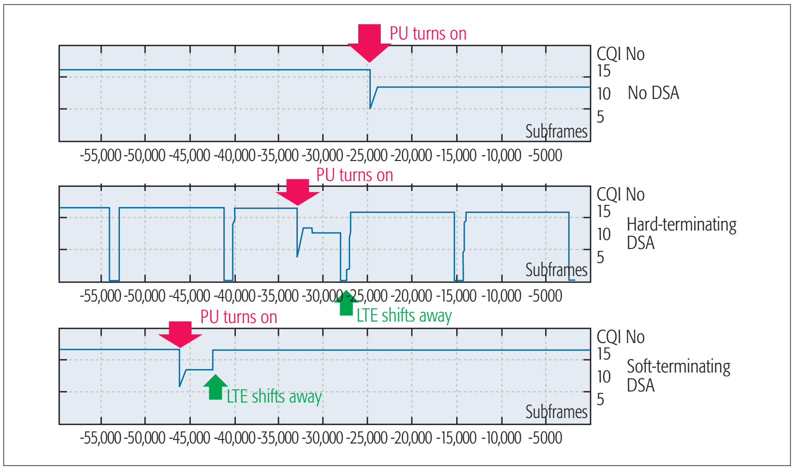

Figure 7. Impact of hard termination and soft termination.

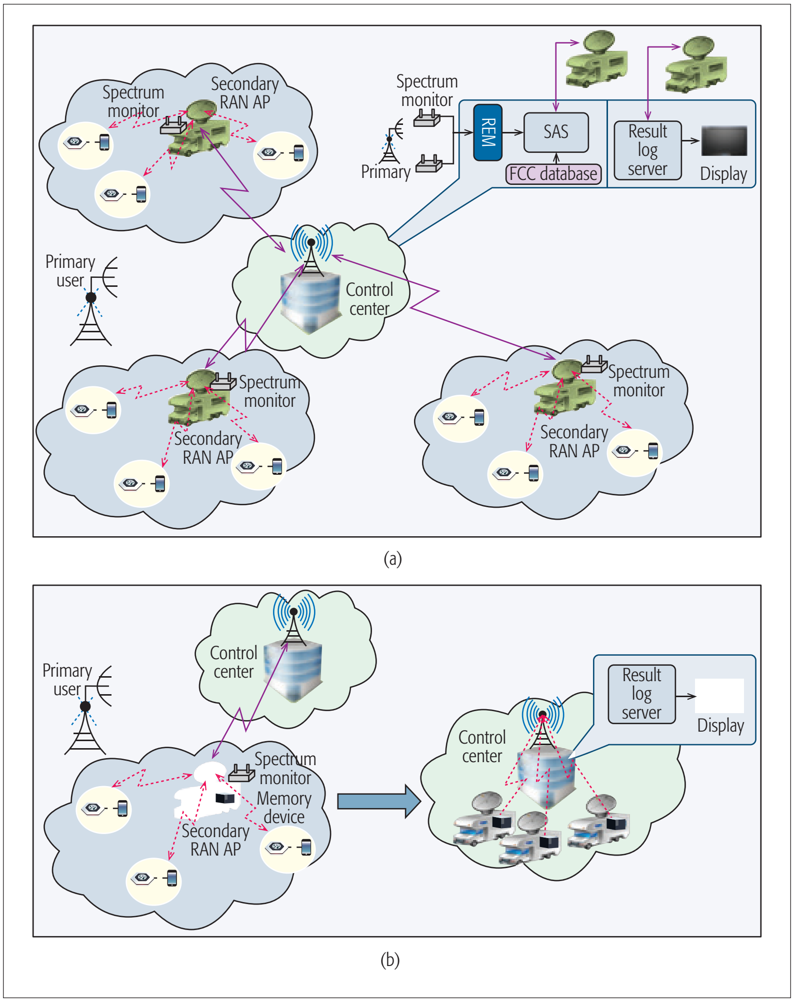

Figure 8. a) System architecture for the proposed testbed for spectrum sharing; b): over-the-air data collection: modile nodes with temporary storage platform.

Conclusion

Wireless communications play a critical role for national security and public safety, and are extensively utilized by first responders. Fortunately, the prevalence of current cellular technologies could facilitate rescue and safety operations during crisis caused by natural events, The successful implementation of spectrum sharing will mark the performance of PS networks and provide the much-needed economic incentive to facilitate the evolution of current systems. This calls for a coordinated effort from the government, industry, and academia.

such as hurricanes, earthquakes, and fires. First responders always need to be connected with one another for cooperation as well as with remote service centers. However, the existing PS services fail to satisfy the PS user requirements in emergency scenarios. As a result, increasing the network capacity is one of the primary concerns, and spectrum sharing was deemed a key solution. With appropriate spectrum sharing agreements, PS users can access licensed spectrum, shared, and, potentially, commercial networks.

This article presents some of the initiatives undertaken at Wireless @ Virginia Tech to help designing the next generation PS networks. We also discuss our vision of a flexible, rapidly deployable, and reconfigurable PS system to meet the service demands when and where needed. The successful implementation of spectrum sharing will mark the performance of PS networks and provide the much-needed economic incentive to facilitate the evolution of current systems. This calls for a coordinated effort from government, industry, and academia to expedite advancements in spectrum sharing R&D that will enable designing and deploying scalable and robust next generation PS networks.

Acknowledgment

The authors acknowledge Mr. Randall Nealy's and Dr. Taeyoung Yang's key contributions to the design and development of the repeater and to project management. This work is supported by the office of the Vice President for Information Technology at Virginia Tech, the Center for Advanced Engineering and Research (CAER), the Institute for Critical Technology and Applied Science (ICTAS), and the National Science Foundation (NSF) under grant number 1343222.

References

[1] T. Doumi *et al.*, "LTE for Public Safety Networks," *IEEE Commun. Mag.*, vol. 51, no. 2, 2013, pp. 106–12.
[2] FCC, "A Broadband Network Cost Model: A Basis for Public Funding Essential to Bringing Nationwide Interoperable Communications to America's First Responders," tech. rep.., Apr. 2010.
[3] U.S. NPSTC Public Safety Communications Report, "Public Safety Communications Assessment 2012–2022, Technology, Operations, and Spectrum Roadmap," June 2012.
[4] FCC NPRM, "In the Matter of Amendment of the Commission's Rules with Regard to Commercial Operations in the 3550–3650 MHz Band," F.C. Commission, Ed., Washington, DC, 2012.
[5] Draft Report and Rec. of the Visiting Committee on Advanced Technology of NIST, "Desirable Properties of a National Public Safety Network," NIST, Ed., 2011.
[6] B. Grossman, "The Middle Class Tax Relief and Job Creation Act of 2012: A Review of Its Scope and Impact on Local Zoning Authorities' Telecommunications Decisions," *Planning & Environmental Law*, vol. 65, no. 6, 2013, pp. 7–8.
[7] "Report to the President: Realizing the Full Potential of Government-Held Spectrum to Spur Economic Growth," PCAST, Ed., White House, Washington, DC, 2012.
[8] P. May, J. Black, and M. R. Fallows, "Portable Communications Device," U.S. Patent No. D711,846. 26 Aug. 2014.
[9] C. Gessner and A. Roessler, *LTE Technology and LTE Test; A Deskside Chat*, Rohde & Schwarz, 2009.
[10] M. M. Sohul *et al.*, "Spectrum Access System for the Citizen Broadband Radio Service," *IEEE Commun. Mag.*, vol. 53, no. 7, 2015, pp. 18–25.

Biographies

Munawwar M. Sohul (mmsohul@vt.edu) received his M.Sc. degree in communication and signal processing from Imperial College, London, United Kingdom, and his B.Sc. degree in electrical and electronics engineering from Bangladesh University of Engineering and Technology, Dhaka. Currently, he is working toward his Ph.D. degree in electrical engineering in the Bradley Department of Electrical and Computer Engineering, Virginia Polytechnic and State University (Virgina Tech), Blacksburg, working with Dr. Jeffrey H. Reed. His research interest in wireless communication and signal processing includes spectrum sharing, opportunistic spectrum access, carrier aggregation, self-organizing networks, and cognitive radios. He was part of Wireless @ Virginia Tech's team that demonstrated the first dynamic-spectrum-access-based LTE operation in 3.5 GHz band.

Miao Yao (miaoyao@vt.edu) received his B.S. degree in electrical engineering from Beijing Jiaotong University in 2007 and M.S. degree in microelectronic and solid state electronics from Tsinghua University in 2010. He is currently working toward his Ph.D. degree at Virginia Tech. His primary research interests are dynamic-spectrum-access-related research issues and energy-efficient resource allocation algorithms in next generation cellular communication systems.

Xiaofu Ma [S'13] (xfma@vt.edu) received his B.S. degree in electronics from Northwest University, Xi'an, China, in 2008, and his M.S. degree in computer science from Tongji University, Shanghai, in 2011, and is currently pursuing his Ph.D. degree in electrical engineering at Virginia Tech. His research interests include network protocol and optimization, cognitive radio networks, dynamic spectrum access, spectrum sharing, chipless RFID, wireless healthcare, and wireless local area networks.

Eyosias Y. Imana (eyoseph@vt.edu) received his B.Sc. in electrical engineering from Bahir Dar University, Ethiopia, in 2007 and M.Sc in electrical engineering from Florida Institute of Technology in 2009. Then he completed his Ph.D. degree in electrical engineering from Virginia Tech in 2014. From January to June 2015, he was postdoctorial associate at Wireless @ Virginia Tech. In summer 2014, he interned at Analog Devices and carried out research on massive MIMO technology. In summer 2015, he joined Qualcomm's modem hardware development team. He has conducted various research projects related to spectrum sharing and cognitive radio. He was part of Wireless @ Virginia Tech's team that demonstrated the first dynamic-spectrum-access-based LTE operation in 3.5 GHz band. He is also involved in research work related to massive MIMO and distributed antenna systems.

Vuk Marojevic (maroje@vt.edu) graduated from the University of Hannover (M.S.), Germany, and Polytechnic University of Catalonia (Ph.D.), Spain, both in electrical engineering. He joined Wireless @ Virginia Tech in 2013. His research interests are in software-defined radio, spectrum sharing, 4G/5G cellular technology, and resource management with application to public safety and mission-critical networks and unmanned aircraft systems.

Jeffrey H. Reed [F] (reedjh@vt.edu) is the Willis G. Worcester Professor in the Bradley Department of Electrical and Computer Engineering, Virginia Tech. He currently serves as founder of Wireless @ Virginia Tech, one of the largest and most comprehensive university wireless research groups in the United States, and is the founding faculty member of the Ted and Karyn Hume Center for National Security and Technology. Since joining Virginia Tech in 1992, he has been PI or co-PI of approximately 100 projects covering areas such as software radio, cognitive radio, ultra-wideband, and channel modeling. He is cofounder of CRT Wireless, a company that is developing cognitive radio techniques for commercial and military systems, and Power Fingerprinting, a company focused on embedded device security. He has served on panels, coordinated numerous workshops and conferences, and served on advisory groups for the Department of Commerce, Department of Defense, State of Virginia, and NSF as well as technical advisory boards for many companies. Recently he served as Associate Editor for *Proceedings of the IEEE* issues on cognitive radio. He is a Fellow of the IEEE for contributions to software radio and communications signal processing and for leadership in engineering education, and is a past recipient of the College of Engineering Award for Excellence in Research. He is the author of three books and over 200 journal and conference papers. He received his B.S., M.S., and Ph.D. degrees from the University of California, Davis.

IEEE 2016 GLOBECOM

GLOBAL COMMUNICATIONS CONFERENCE, EXHIBITION & INDUSTRY FORUM

IEEE — IEEE COMMUNICATIONS SOCIETY

Freedom Through Communications

4-8 December 2016 // Washington, DC, USA

The 2016 IEEE Global Communications Conference (GLOBECOM) will feature a comprehensive technical program including 13 symposia, tutorials, workshops and an industrial program featuring prominent keynote speakers, technology and industry forums and vendor exhibits.

CALL FOR PAPERS

WASHINGTON, DC

#IEEEGC16

– TECHNICAL SYMPOSIA –

We invite you to submit original technical papers in the following areas:

- **Ad Hoc and Sensor Networking**
- **Communication & Information System Security**
- **Cognitive Radio and Networks**
- **Communication QoS, Reliability and Modeling**
- **Communication Theory**
- **Communications Software, Services and Multimedia Applications**
- **Green Communications**
- **Mobile and Wireless Networking**
- **Next-Generation Networking**
- **Optical Networks and Systems**
- **Signal Processing for Communications**
- **Wireless Communications**

Selected Areas in Communications:

- **Access Networks & Systems**
- **Big Data**
- **Cloud Networks**
- **Data Storage**
- **e-Health**
- **Internet of Things**
- **Molecular, Biological and Multi-scale Communications**
- **Power Line Communications**
- **Satellite & Space Communications**
- **Social Networks**

Please address questions regarding the Technical Symposia to TPC Chair Sherman Shen (sshen@uwaterloo.ca) and/or TPC Vice Chairs Maite Brandt-Pearce (mb9q@eservices.virginia.edu) and Sennur Ulukus (ulukus@umd.edu).

Accepted and presented technical and workshop papers will be published in the IEEE GLOBECOM 2016 Conference Proceedings and submitted to IEEE Xplore®. See the website for author requirements of accepted authors.
Full details of submission procedures are available at www.ieee-globecom.org.

IMPORTANT DATES:

Symposia Papers - **1 April 2016** // Tutorial Proposals - **15 March 2016**

– COMMITTEE –

General Chair
Bijan Jabbari, George Mason University

TPC Chair
Sherman Shen, University of Waterloo

TPC Vice Chairs
Maite Brandt-Pearce, University of Virginia
Sennur Ulukus, University of Maryland

Tutorial Co-Chairs
Brian Mark, George Mason University
Yahong Rosa Zheng, Missouri University of Science and Technology

Workshop Co-Chairs
Shuguang Cui, Texas A&M
Stefano Secci, University of Paris 6

Executive Chair
Paul Cotae, University of the District of Columbia

General Vice Chair
Jerry Gibbon, US Department of Commerce

GIMS Advisor
Chi-Ming Chen, AT&T

GITC Advisor
Peter Rost, Nokia Networks

- www.ieee-globecom.org -

Indirect Returns and Use of NPV in Economic Viability Modeling of Critical Communications Networks

Natalia Boliari

This article suggests that broadband critical communications networks can be highly beneficial and desirable to implement, particularly from the public and social points of view.

Abstract

Broadband-technology-based critical communications networks are predicted and known to generate significant direct returns. Less attention is given to their potential to generate long term indirect returns in the form of various socioeconomic benefits. This article focuses on identifying such indirect returns, and suggests that they are included in financial and economic viability analyses for better accuracy in the cost-benefit analyses of broadband critical communications networks. Developing a net present value model that accounts for both direct and indirect benefits, the article suggests that broadband critical communications networks can be highly beneficial and desirable to implement, particularly from the public and social points of view.

Introduction

As critical communications networks are at a crossroads with next generation performance-hungry applications, the use of broadband communications technologies such as Long Term Evolution (LTE) stands out with its superior features. Many developed countries and large organizations have already declared their preference for broadband technology and interest in investing in strategies evolving into broadband as it addresses many of the shortcomings of today's narrowband communications infrastructure for critical communications. It is widely believed that, once deployed, a broadband-technology-based critical communications network will make a significant difference that deeply affects organizations as well as individuals.

The high deployment cost of broadband critical communications networks, however, forces parties to evaluate its economic viability carefully. While economic viability covers wider grounds (including technical, marketing, operational, environmental, and other aspects that are beyond the scope of this study), financial viability implies the commercial viability of a project showing its financial strengths. Economic viability should therefore be considered as an extension of the financial analysis since the concern is for the society or economy as a whole.

In an attempt to quantify financial viability under different scenarios, Hallahan and Peha [4] consider a public-private partnership from a for-profit point of view. Analyzed are the cost of building out and operating the necessary wireless infrastructure and the revenue that could be derived from serving commercial and public safety subscribers (see also [3, 5]). The study computes the revenue of a public-private partnership scenario using the projections based on population covered by the network assuming the revenue is basically the fee collected from commercial and public subscribers. Computed is a net present value (NPV) with a time horizon of 10 years, parallel to the assumed initial licensing period. A discount rate of 8 percent is used, noting its consistency with similar works.

The returns used in the above study are direct returns obtained from subscribers in the form of subscription fees. Broadband-technology-based critical communications also produce significant indirect returns due to tremendous increase in capacity, response time, and multimedia capabilities compared to current critical communications networks based on narrowband technologies such as Project 25 (P25), terrestrial trunked radio (TETRA), and digital mobile radio (DMR). We believe that, in addition to direct returns, identifying indirect returns and including them in the viability analyses will produce more complete results, and enable researchers and decision makers to better understand the overall returns of critical communications networks. Including the indirect returns in economic viability studies will result in more accurate return computations and less bias toward total costs.

The types of indirect returns discussed in this study are mostly socioeconomic returns created by the advances in and superior performance of broadband technology used by upcoming mission-critical networks. They may last many years and create significant benefits to organizations and society. Therefore, this study focuses on indirect returns, and proposes possible ways of predicting their values and including them in financial and economic viability analyses. The study then discusses the popular methods that might be considered to evaluate the viability of broadband critical communications networks.

Indirect Benefits

Indirect benefits are the returns other than subscription fees obtained from moving into broadband-technology-based wide-area mission-critical

The author is with Manhattan College.

0163-6804/16/$25.00

mobile networks. A 2013 study [2] labels them as socioeconomic benefits and evaluates them in the context of the European Union. In this comprehensive work, socioeconomic benefits are assessed in five distinct categories: safety, efficiency from enhanced safety, operational efficiency, use by traffic police, and reduction in mortality by ambulance services. The study reports that for an area covering approximately 60 percent of the European Union, with a population of 300 million, the assessment process estimated an annual consolidated socioeconomic value of around €20.9 billion ($32 billion).

We classified the indirect benefits in four similar categories described as follows:

Reduction in Labor Used in Back, Middle, and Front Office

It is widely believed that as a result of moving to a much more powerful technology, reductions will be realized in police budgets, in frontline police, in the number of back office officers as well as in police stations, front counters, and shared locations [2, p. 38]. Such reduction is also applicable to other first responders and should be considered as actual cash return to the organization or agency paying for the first responders. It is possible to quantify and integrate the amount of this indirect return by simply comparing the numbers of similar size localities with and without broadband critical communications networks and using the relative populations as the correction parameter. This indirect return will not be included in financial viability studies if the wide area networks are entirely private.

Crime Reduction through Interventions

The impact of crime on any society is significant. It is well known that criminal activity facilitates the consumption of illegal goods and services, damages and destroys property, and brings costs to society in the form of higher insurance premiums, extra spending on safety, pain and suffering, and loss of life [1]. Crime affects public and private sectors alike.

Broadband-technology-based critical communications networks will reduce crime by creating additional community policing, by allowing for more effective crime fighting such as the ability to perform visual identity checks, and closed-circuit television (CCTV) evidence capture to help secure convictions (See [2] for a more detailed analysis of benefits).

Any project potentially reducing criminal activity, therefore, is an investment decision. Machin and Marie [6] report that in early 2000, a street crime reduction initiative in the United Kingdom, with incremental policing cost of €24.1 million, was estimated to have delivered a net socio-economic benefit of €107 million to €130 million. In more detail, the average number of robberies recorded during the control period of 1999/2000 and 2000/2001 translates into between 12,751 and 10,846 fewer robberies resulting from the street crime reduction initiative introduction. To obtain the monetary benefits from robbery reduction, they multiply these figures with the average social cost of a recorded personal robbery of £12,094 as estimated by the Home Office. After deducting the £24.1 million average annual cost of the initiative, they find a net social benefit of the policy of between £107 and £130 million. The resulting return was approximately five times the investment. Another study [1] reports that crime's total cost accounts for 7.7 percent of the United Kingdom's and 11.9 percent of the United States' gross domestic product (GDP).

This indirect return obtained from crime reduction is realized by lowering expenditures for police, courts, prisons, and property damage and destruction. It also includes gains such as direct personal earnings and productivity loss due to crime, such as injury; and wider productivity loss in other parts of the economy. The benefits obtained from crime reduction will generate value to society in the form of lower annual spending to maintain higher standards and in the form of direct annual cash flows to any investment project. The benefits created in this category will not modify the viability analysis if the networks are designed as entirely private, but it will still have certain effects, creating rigidities for subscription fees since residents now will have higher real income generated by broadband communications networks. Subscribers will be more willing to accept higher fees once they realize the advantages of living in a low crime area. Residents' real income increases due to crime reduction and makes them more willing to pay for the service causing the reduction.

Mortality Reduction through Reduced Response Time

Broadband-technology-based critical communications networks are recognized as a fundamental component of emergency services [8]. Firefighting, ambulances, and police are considered as services that rely on mission-critical mobile applications (see [9]). The United States has more than 2 million first responders, including about 630,000 police patrol officers, and 300,000 firefighters and other public safety workers, as well as 100,000 federal employees in protective service occupations [13]. Obviously, any efficiency created in this field will have a very significant cost reducing effect.

According to a 2006 study, [11], concerning the U.K., €3.98 billion annual socioeconomic benefits could result from the utilization of mobile broadband to assist ambulance crews. This will result from better informing the crews and saving an additional 1858 out-of-hospital cardiac arrest victims faster, especially within the 8 min target that is critical for Type A life threatening responses. The benefits created in this category will not modify the viability analysis if the networks are designed as entirely private, although one can claim that awareness of mortality reduction created by wide-area networks will change subscribers' price elasticity and make them more accepting of higher subscription fees.

Increase in Real Estate Values and Resulting Property Taxes

According to estimates for the United States, 10 and 25 percent reduction in violent crime could increase housing prices in 8 major cities by 0.83 and 2.1 percent, respectively, in the following year. This would equate to $16 and $41 billion, respectively [12].

Broadband-technology-based critical communications networks will reduce crime by creating additional community policing, by allowing for more effective crime fighting such as the ability to perform visual identity checks, and closed-circuit television evidence capture to help secure convictions.

The method	Formulation
Net present value (NPV)	$NPV = \sum_{t=0}^{N} \frac{CF_t}{(1+R)^t}$
Internal rate of return (IRR)	$NPV = \sum_{t=0}^{N} \frac{CF_t}{(1+IRR)^t} = 0$
Payback period (discounted) (PP)	Payback $= A + (B/C)$
Modified internal rate of return (MIRR)	$MIRR = \left[\left(\frac{1}{CF_0}\right)\sum_{t=1}^{N} CF_t(1+R)^{N-t}\right]^{1/N} - 1$

Table 1. Formulations of popular investment viability methods.

Quantifying this benefit is not difficult and can be estimated by using crime reduction related real estate value appreciation statistics of similar areas, or using the above statistics adjusted for population. This is a direct benefit to home owners and to local governments. It raises property values causing increase in the wealth of owners. Rise in wealth, on the other hand, causes value increase in the locality that may cause increase in creditworthiness, higher propensity to consume, and higher sales tax collection for the local area.

As a function of increase in real estate values, local governments will benefit from proportional increases in real estate taxes resulting from increased real estate values. These are very safe cash flows that need to be discounted with the risk-free rate. It is important to include this indirect return, especially for organizational (or government) partnerships. Any municipality will have to take this indirect benefit into consideration in viability studies. Quantification of this tax increase can easily be achieved by the municipalities since they have the actual tax data of the locality. This return would not affect viability if the networks are entirely private as tax collection would not change the cash flows of the private investor.

The Method

Economic projects are generally evaluated using four alternative popular methods: the net present value (NPV), internal rate of return (IRR), payback period (PP), and modified internal rate of return (MIRR). There are many studies comparing and contrasting those methods in a variety of project evaluation scenarios (e.g., [7, 10]). While we do not intend to provide a comprehensive review of those methods here, we would like to provide a brief description of each method, its basic formulation, and its most widely criticized aspects.

Table 1 shows the basic formulation of each method, showing how they work and what they emphasize. Following the order in Table 1, the NPV method compares a project's initial investment outlay with the present value of expected future cash flows generated by the project. It requires projecting future cash flows as well as a discount rate to compute the present value of those cash flows. If NPV is positive, the project is considered to be viable.

The IRR method computes a discount rate that makes the present value of the above-mentioned cash flows equal to zero. In other words, the IRR calculates the discount rate that makes the NPV equal to zero. It is, in a sense, a break-even rate, implying that if the percentage return (IRR) falls short of this deployment cost of capital (in percentage), the project is not viable and should not be accepted. A popular method in the industry, the IRR unrealistically assumes earned cash flows are reinvested at the project-specific rate and produces unreliable results, especially when the IRR is significantly different than the cost of capital. Wide area networks may have relatively low cost of capital if deployed by governments or with a governmental partnership, causing greater bias due to a wrongly assumed reinvestment rate. In addition to the unrealistic reinvestment rate of the IRR method, it will, as a higher order polynomial, produce confusing multiple solutions, creating another issue that needs to be addressed carefully.

Our third method, the PP, is a rather simple method of computing the number of years and months required to recover the initial investment outlay. It does not tell us to accept or reject the project using a natural benchmark like the ones used in NPV or IRR, but a preferred benchmark may be used for this sort of decision making. It may use the regular cash flows or discounted cash flows. The payback period is not a relevant method in assessing financial or economic viability in this case, since we are not concerned about how long it takes to recover the initial investment. Moreover, since the PP entirely ignores the cash flows generated after the payback point, it is not quite relevant for projects with many years of future cash flows. Finally, as a hybrid method, the MIRR uses the best parts of NPV and IRR, and computes a value that can be used in viability analysis. MIRR computes a rate of return that equates the initial investment outlay to a terminal value that is computed using the cost of capital as the reinvestment rate of the cash flows.

Table 1 summarizes the formulations, and Table 2 summarizes some strengths and weaknesses of the four methods. Throughout the table, CF is the cash flow, t is the time subscript, N is the total number of years, and R is the discount rate. A is the number of years until the positive cumulative cash flow is obtained, B is the amount left to recover the initial investment outlay at point A, and C is the amount of cash flows generated in year $A + 1$.

One of the most common and most widely accepted methods for project evaluation is the NPV method. As noted in the table, the NPV is the method of choice for many economic viability analyses due to its reliability in treating the reinvestment rates properly using the market rate of borrowing as the discount rate and eliminating multiple solutions that may confuse evaluators. If we combine the pros and cons of the most popular viability methods, NPV is the undisputed choice (see, e.g., [10]). NPV requires knowing the initial cost of a project, which may spread over a number of years, and the future cash flows generated by the project, which may be perpetual, limited to a number of years due to a licensing arrangement, or limited to the natural life of the project.

Computing the present cost of an assumed project involves more certainty since the computations are performed at the present time using the present cost figures, as opposed to trying to obtain reliable predicted values of the future benefits. The initial investment outlay of any project is the present time or close to present time phenomenon with no need for a risk-adjusted discount rate to bring the future numbers into the present time (it may require a discount rate when initial investment extends beyond the present year, but will not be difficult to obtain due to the close proximity of those years to present time). Project evaluations lose predictive power as the expected cash flows occur farther from the present time. The potential problems of computing the reliable present value of future cash flows may be summarized in two major categories; they are the difficulty of:

- Projecting the future cash flows with accuracy since accuracy almost always declines with time
- Obtaining the proper risk-adjusted rate needed to discount those cash flows to the present time

Projecting Future Cash Flows

Projecting future cash flows with accuracy is a complicated task. The process is far from being easily predictable due to the difficulty of itemizing the indirect benefits and the complexities surrounding the proper risk-adjusted discount rates on long or even infinite horizons. Moreover, the benefits obtained from critical communications networks are cumulative in nature; especially the indirect ones are expected to evolve in a number of stages. In an attempt to integrate this into the viability studies, this article suggests using a discounted cash-flow model evaluating the cash flows generated by critical communications networks in three broadly defined stages generally used in structuring the interest rates. In a typical setting, debt instruments are classified as short-, medium-, and long-term securities. Short-term has up to a one-year life; intermediate-term traces from 1 to 5–7 years, and involves less certainty; and long-term covers the years after the intermediate stage. The latter is the one with least certainty.

The three-stage scenario is based on the assumption that once a new system is integrated in place of the previous one, it will produce cash flows with changing certainty levels similar to the debt securities of different lives. Furthermore, using the treasury yield curve driven forward rates in discounting produces more realistic discount rates, especially for the periods covered in the intermediate and long-term stages.

Stage I may be called the short-term benefits stage, which starts once the broadband critical communications networks are properly deployed and become functional. Following financial terminology, we can define this stage using an assumed life limited to one full year from the start date. This is the stage that utilizes the immediate safety related benefits of broadband critical communications networks; it is also the stage where residents will benefit from superior

The method	Pros and cons
Internal rate of return (IRR)	IRR assumes project specific IRR as the reinvestment rate creating overestimation. IRR may produce multiple solutions, confusing the user.
Net present value (NPV)	NPV's reinvestment rate is correct, and there is no multiple solution possibility. Well understood by industry and is the choice of academics.
Payback period (PP)	Ignores time value of money and cash flows generated beyond the recovery period, causing gross underestimation of the future cash flows. It has arbitrarily set thresholds.
Modified internal rate of return (MIRR)	Generally reliable and may be used in any viability analysis without issues. It is much more complex than the other procedures. Not well understood by the industry and creates interpretation issues.

Table 2. Pros and cons of the methods.

quality ambulatory and police related interventions and innovations. In this stage, no re-engineering of infrastructure is assumed. The cash flows are based on the collected subscription fees and are very safe. The proposed base discount rate here is the one-year zero rate with further adjustments for other risks, if any.

Stage II (the intermediate term) will be assumed when stage I is completed and may be set for cash flows generated during the four to five years following from the endpoint of stage I. This is the stage that implements efficiency related innovations generating increases in cash flows obtained in stage I. Interventional efficiency and operational productivity will result in re-engineering (based on the new broadband technology) health organizations and the police force, among other things. Subsequent results involve crime reduction and efficient traffic flows with fewer accidents and reduced mortality across the board. Stage II can be viewed as the one where society starts to enjoy the cumulative benefits of broadband-based critical communications networks. Future cash flow estimations generated by this stage may be obtained comparing the cash flows of stage I for a reference period with the cash flows of stage II, as this percentage may be used to simulate future years' values. In other words, an annual growth obtained comparing last year's benefits with this year's will help us predict future years' expected cash flows. Discounting of those cash flows will be performed using the corresponding discount rates obtained from the treasury yield curve's related range similar to common applications in medium-term debt securities; it may be adjusted further, if necessary. These cash flows are also very safe cash flows since they are based on safe subscription fees and indirect returns obtained from crime and mortality reduction rates. However, the section based on crime and mortality reduction may be considered less certain, requiring a higher discount rate that can be explained by the rate difference of the longer-term debt instrument from the one-year zero rate.

Stage III involves the cash flows generated after the completion of stage II. This is the stage that reflects the benefits of broadband-based critical communications networks at the macro level. Such benefits are increases in real estate values that in turn may generate significant income for owners and consequently increase in revenue col-

Since the discount rate used in the NPV method is a firm specific figure, and its value is affected by both the riskiness of the cash flows and the cost of capital of the firm, a project with negative NPV only means that the project is not viable for the firm in question.

lected from property tax. Cash flows generated in this stage, similar to the previous ones, are very stable and safe. While we do not anticipate additional risks in stage III that are not included in stage II, for the sake of longer-term unforeseeable elements of risk, we suggest using a discount rate corresponding to the forward rate assumed for longer-term debt securities.

A Typical NPV Model

A typical NPV model, as explained above, compares the discounted present values of future cash flows of the process with its initial deployment cost. However, a typical NPV model uses only one discount rate applied to all future cash flows and will not have sections with distinct discount rates as suggested in the modified version below.

Our suggested critical communications network viability equation is as follows:

$$NPV = -CF_0 + \sum_{t=1}^{N_1} \frac{CF_{1t}}{\left(1+R_1\right)^t} + \sum_{N_{1+1}}^{N_2} \frac{CF_{2t}}{\left(1+R_2\right)^t} + \sum_{t=N_2+1}^{N} \frac{CF_{3t}}{\left(1+R_3\right)^t}$$

where $N_1 < N_2 < N$ and the first component is the cost component that enters the equation as a negative figure. More specifically, $-CF_0$ is the total deployment cost of the project, computed by the cost evaluators and simplified here as a present year figure; however, it may easily be aggregated using multiple years *e*.

CF_{1t}, CF_{2t}, and CF_{3t} are the cash flows generated during stages I, II, and III, and R_1, R_2, and R_3 are the discount rates used for stages I, II, and III, respectively. Note that for simplicity, we do not include an additional subscript to R values; however, unless the yield curve is flat, each year may require a different discount rate, in line with the forward rates of the years in question.

The Discount Rates and How to Interpret NPV

As briefly explained above, the discount rates used in the NPV equation are connected to the riskiness of the corresponding cash flows of the process as well as the capital cost of the firm undertaking the project.

It is crucially important to understand the function of the *discount rate* in NPV computations. Since the discount rate used in the NPV method is a firm specific figure, and its value is affected by both the riskiness of the cash flows and the capital cost of the firm, a project with negative NPV only means that the project is not viable for the firm in question. In other words, a project with given cash flows may produce a positive NPV for one firm but a negative NPV for another if the capital structures of the firms are different. This would create different cost of capital values to be used as the discount rate. Positive or negative NPV would not imply universal viability to a project, but would show if the project is profitable for a particular firm with its given capital cost. Therefore, firms with higher efficiencies will have lower discount rates, and the projects undertaken by those firms will produce higher NPV values. Certain partnerships may create efficiencies driving the discount rates lower and making the projects more attractive for all partners, while this may not be viable for each of the partners if they work alone.

Due attention should also be paid to the cash flows generated by a project. In general, if the cash flows are traditional steady cash flows for all participants, we can fairly evaluate the alternative viabilities for different firms. However, just like the wide area networks evaluated here, in the case of government involvement, if certain cash flows can only be collected by the government, our approach to an NPV-based viability study will be different. Most indirect returns, such as increases in real estate taxes or reduced expenses due to crime reduction, can only be collected by the government. This creates advantages for the government and therefore incentive to get involved in the project, at least as a partner. Still, this will require a careful analysis of how the direct and indirect returns should be distributed among partners.

The cash flows used here are direct subscription fees and other indirect fees mentioned throughout the article. It is assumed here that if the critical communications networks are deployed by the government, those cash flows are not riskier than similar maturity government debt instruments with no default risk. Therefore, under the assumption of government owned critical communications networks, we are using the discount rates obtained from the treasury yield curve.

If the critical communications network is assumed to be entirely private, company-specific discount rates may be appropriate since a privately owned company's capital cost should be taken into account in assuming the proper discount rate. As an alternative, the treasury zero rates may be adjusted for the added risk coming from private parties using their bond ratings, which is an indication of the credit risk they bear. However, it is also well known that using company-specific rates may create biases against certain types of projects. For example, if a company uses the weighted average capital cost for discounting, this will create unfair rejection of lower-risk projects and acceptance of higher-risk projects, eventually increasing the overall risk level of a firm. Similarly, if a project is undertaken by a partnership composed of many firms with different cost structures, it will be difficult to assume a discount rate that will be used to discount the generated cash flows and one which reflects the characteristics of all partners.

Finally, some firms may use project-specific discount rates, but this is also known to be uneasy to obtain as it requires knowing the project's risk or identifying the competing firms operating only in the industry in question, and to use this information for the risk assessment. In an attempt to overcome these difficulties, we suggest the use of the treasury zero rates to predict the forward rates. Under the assumption of entirely private networks, those values may be adjusted using a firm's bond rating and existing cost of equity and debt relationship. With regard to public projects, they are not necessarily designed in light of profitability as socially desirable projects need not always be financially profitable. This, however, would not change the fact that public projects should still be evaluated as if they are private projects, and the feasibility results must reflect the standard parameters used in regular feasibility analyses. While the accept/reject decisions may not entirely depend on the results of financial viability, the results should reflect the true cost of the project.

Conclusion

A project is considered economically viable when the present value of the project's generated future cash flows exceeds the present cost of the project. Within this context, there are some studies performing NPV analysis to assess the viability of new broadband-technology-based critical communications networks.

This article contributes to the viability studies by highlighting two issues. First, it identifies the generally ignored indirect returns that will create significant cash flows, making the project much more desirable. Second, it suggests that those returns may only be collected by the society or community as a whole, thus making public agencies very good candidates to participate in the deployment of broadband-technology-based critical communications networks.

Overall, the indirect returns create cash flows to publicly owned projects, or public and private partnerships, and much lower cash flows for an entirely private network. This alone is an important conclusion suggesting public agency involvement in the form of subsidies (a fair price for the benefits). The economic implications of such involvement should be evaluated in detail in future studies since a significant part of the indirect benefits are enjoyed by the public regardless of the government's involvement. In case of entirely privately deployed networks, the public will be free riders since they still collect the indirect returns. The private sector should take this into account in viability studies and negotiate for fairness.

The article also highlights the fact that the NPV is not a method showing universal viability of a project, but one that shows if the project is viable for the firm in question with its specific cost of capital. This also suggests forming an efficient partnership to enjoy the lowest possible discount rate that will encourage the deployment of broadband-based critical communications networks.

Our work underlines certain issues in discount rate determinations and highlights the common issues around the discount rate selections as applied to the NPV analysis. It is suggested that if the project is publicly owned, the treasury zero rates should be used to determine the forward discount rates; and in the case of less, partial, or no government involvement, those rates should be adjusted using the involved private parties' bond ratings and other leverage parameters.

References

[1] C. Detotto and E. Otranto, "Does Crime Affect Economic Growth?" *Kyklos*, vol. 63, pp. 330–45. 2010, doi: 10.1111/j.1467-6435.2010.00477.x.

[2] A. Grous, 2013, "Socioeconomic Value of Mission Critical Mobile Applications for Public Safety in the EU: 2x10MHz in 700MHz in 10 European Countries," Centre for Economic Performance, London School of Economics and Political Science, last accessed Oct. 11, 2015; http://www.lse.ac.uk/businessAndConsultancy/LSEEnterprise/pdf/tetraReport.pdf.

[3] R. Hallahan and J. M. Peha, "Quantifying the Costs of a Nationwide Broadband Public Safety Wireless Network, Carnegie Mellon University Research Showcase @ CMU," *36th Telecommun. Policy Research Conf.*, 2008, last accessed Oct. 11, 2015; http://repository.cmu.edu/cgi/viewcontent.cgi?article=1039&context=epp.

[4] R. Hallahan and J. M. Peha, "The Business Case of a Nationwide Wireless Network that Serves Both Public Safety and Commercial Subscribers. Carnegie Mellon University Research Showcase @ CMU," *37th Telecommun. Policy Research Conf.*, 2009, last accessed Oct. 11, 2015; http://repository.cmu.edu/cgi/viewcontent.cgi?article=1038&context=epp

[5] J. P. C. Kleijnen, 1984, "Quantifying the Benefits of Information Systems," *Euro. J. Operational Research*, vol. 15, no. 1, pp. 38–45.

[6] S. Machin and O. Marie, 2011, "Crime and Police Resources: The Street Crime Initiative," *J. Euro. Economic Assn.*, vol. 9, no. 4, pp. 678–701.

[7] Parliamentary Joint Committee on Law Enforcement. "Spectrum for Public Safety Mobile Broadband," Parliament of Australia, 2013, last accessed Oct. 12, 2015; http://www.aph.gov.au/Parliamentary_Business/Committees/Joint/Law_Enforcement.

[8] L. Phalippou, "The Hazards of Using IRR to Measure Performance: The Case of Private Equity," *J. Performance Measurement*, vol. 12, no. 4, 2009, pp. 55–67.

[9] M. Osborne, 2010, "A Resolution to the NPV-IRR Debate?," *Quarterly Rev. Economics and Finance*, vol. 50, no. 2, pp. 234–39.

[10] J. M. Peha, "How America's Fragmented Approach to Public Safety Wastes Spectrum and Funding. Carnegie Mellon University Research Showcase @ CMU," *33rd Telecommun. Policy Research Conf.*, 2005, last accessed Oct. 12, 2015; http://repository.cmu.edu/cgi/viewcontent.cgi?article=1029&context=epp.

[11] L. Price, "Treating the Clock and Not the Patient: Ambulance Response Times and Risk," *Quality and Safety in Health Care*, vol. 15, no. 2, 2006, pp. 127–30, doi: 10.1136/qshc.2005.015651.

[12] R. J. Shapiro and K. A. Hassett, 2012, "The Economic Benefits of Reducing Violent Crime: A Case Study of 8 American Cities," *Centre for American Progress*, last accessed Oct. 11, 2015; .https://cdn.americanprogress.org/wp-content/uploads/issues/2012/06/pdf/violent_crime.pdf.

[13] The White House, "The Benefits of Transitioning to a Nationwide Wireless Broadband Network for Public Safety," 2011, last accessed Oct. 11, 2015; http://www.whitehouse.gov/sites/default/files/uploads/publicsafetyreport.pdf.

Biography

Natalia Boliari (natalia.boliari@manhattan.edu) is an assistant professor of economics at Manhattan College, New York. She holds a Ph.D. degree in public policy with focus on economic development policy, an M.A. degree in public administration with concentration on international development, and a B.S. degree in industrial engineering. Specializing in the interdisciplinary field of new institutional economics, her research interests include the role of institutions and property rights for economic development, the political economy of international trade, impacts of regional trading agreements, and trade in services and finance.

RADIO COMMUNICATIONS: COMPONENTS, SYSTEMS, AND NETWORKS

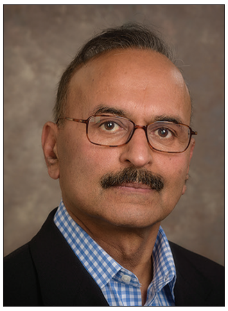

Amitabh Mishra

Tom Alexander

In this issue we focus on air interfaces for next generation digital wireless communications technologies. A new generation of wireless technology has been deployed approximately every 10 years for three straight decades. Each generation has produced significant gains in performance, but also an order of magnitude increase in complexity. With 5G almost upon us, we will need yet another advance in our air interfaces.

A brief recap of the history of digital air interfaces will underline this point. Second generation (2G) digital wireless technologies of the 1990s were simple time-division or code-division multiple access (TDMA or CDMA); easily understood and managed. 3G technologies increased data rate by 10× using wider channels and higher-order modulation formats. 4G (LTE) increased data rates by another order of magnitude via still more complex approaches such as orthogonal frequency-division multiple access (OFDMA) and multiple-input multiple-output (MIMO). The goals for 5G are not yet agreed upon, but it is commonly expected that another 10× capacity improvement will be achieved. This may be driven by extremely wide millimeter-wave channels, massive MIMO, and modulation formats bonding together very large RF channels.

It is worth remembering that the original 3GPP GSM standard offered under 300 kb/s for the whole of an occupied channel! The fact that we are today able to routinely stream high-definition movies in real time to our wireless devices, courtesy of 300 Mb/s LTE-Advanced radios, is a testament to the endless ingenuity of air interface researchers and engineers. To make 5G architectures possible, wireless engineers will need to push the limits in multiple areas. For a start, they have to squeeze more capacity out of the available spectrum, which may entail using RF spectrum in innovative ways (e.g., carrier aggregation). The utilization of the propagation environment must be improved, possibly through massive MIMO. The efficiency of the radio system must be improved by matching user and application needs to air interface characteristics. And we need to do all this without increasing size and power consumption; we have grown used to our thin and light smartphones that run for days without recharging.

This issue of our Series has three articles dealing with timely and relevant air interface topics for 5G. The first, "Load Modulated Arrays: A Low-Complexity Antenna Technology for Massive, Distributed, and Small Cell Wireless Networks," introduces recent work in massive MIMO, which needs to support dozens of RF chains without driving cost and complexity to unsustainable levels. The article compares various approaches to the problem, such as straightforward multi-RF arrays, parasitic arrays, and load-modulated arrays, and then discusses some key implementation aspects of load-modulated arrays as a possible solution. Challenges and issues are discussed to introduce the various engineering trade-offs that must be performed when realizing a massive MIMO front-end.

The second article, "A Flexible 5G Frame Structure Design for Frequency-Division Duplex Cases," looks at a new frame structure for the air interface of 5G networks. Modern wireless RANs must accommodate a diverse set of human and machine user types with an extremely diverse set of service requirements. If these characteristics are not taken into account, most of the air interface technology gains can be lost due to inefficiency of utilization. A key to maintaining utilization is an efficient and flexible PHY/MAC frame structure, which can simplify a lot of the service requirements. The article therefore introduces the usage cases and requirements of proposed 5G advances as they pertain to MAC-level framing.

Our third article deals with reducing spurious emissions in modern digital transmitters. This is a particular issue with advanced radio front-ends: it is difficult to meet stringent requirements on emissions and occupied bandwidth, which conflict with the need for higher-order modulation formats and more efficient power amplifiers. Carrier aggregation merely aggravates the problem, and new approaches are urgently required for 5G. The article "Digital Predistortion for Mitigating Spurious Emissions in Spectrally Agile Radios" surveys techniques for cleaning RF PA emissions without losing efficiency. Sub-band digital predistortion has been proposed as a promising technique, and the authors investigate the complexity and efficiency of sub-band digital predistortion in the context of carrier aggregation and LTE-Advanced.

We will continue to place similar thematic topics before you in future issues, and encourage our readership to submit timely surveys, tutorials, and presentations of ongoing work discussing emerging trends in wireless communications. Of immediate interest are articles in 5G communications and networking, for which we have recently issued a Call for Papers. We particularly thank our reviewers; all are volunteers who have taken time out from their busy professional lives to help filter and improve the papers submitted to us. Without their unstinting efforts, we would be unable to maintain the quality of the Series.

CALL FOR PAPERS
IEEE COMMUNICATIONS MAGAZINE
COMMUNICATIONS STANDARDS SUPPLEMENT

BACKGROUND

Communications standards enable the global marketplace to offer interoperable products and services at affordable cost. Standards development organizations (SDOs) bring together stakeholders to develop consensus standards for use by a global industry. The importance of standards to the work and careers of communications practitioners has motivated the creation of a new publication on standards that meets the needs of a broad range of individuals, including industrial researchers, industry practitioners, business entrepreneurs, marketing managers, compliance/interoperability specialists, social scientists, regulators, intellectual property managers, and end users. This new publication will be incubated as a Communications Standards Supplement in *IEEE Communications Magazine*, which, if successful, will transition into a full-fledged new magazine. It is a platform for presenting and discussing standards-related topics in the areas of communications, networking, and related disciplines. Contributions are also encouraged from relevant disciplines of computer science, information systems, management, business studies, social sciences, economics, engineering, political science, public policy, sociology, and human factors/usability.

SCOPE OF CONTRIBUTIONS

Submissions are solicited on topics related to the areas of communications and networking standards and standardization research, in at least the following topical areas:

Analysis of new topic areas for standardization, either enhancements to existing standards or of a new area. The standards activity may be just starting or nearing completion. For example, current topics of interest include:

- 5G radio access
- Wireless LAN
- SDN
- Ethernet
- Media codecs
- Cloud computing

Tutorials on, analysis of, and comparisons of IEEE and non-IEEE standards. For example, possible topics of interest include:

- Optical transport
- Radio access
- Power line carrier

The relationship between innovation and standardization, including, but not limited to:

- Patent policies, intellectual property rights, and antitrust law
- Examples and case studies of different kinds of innovation processes, analytical models of innovation, and new innovation methods

Technology governance aspects of standards focusing on both the socio-economic impact as well as the policies that guide it. These would include, but are not limited to:

- The national, regional, and global impacts of standards on industry, society, and economies
- The processes and organizations for creation and diffusion of standards, including the roles of organizations such as IEEE and IEEE-SA
- National and international policies and regulation for standards
- Standards and developing countries

The history of standardization, including, but not limited to:

- The cultures of different SDOs
- Standards education and its impact
- Corporate standards strategies
- The impact of open source on standards
- The impact of technology development and convergence on standards

Research-to-standards, including standards-oriented research, standards-related research, research on standards

Compatibility and interoperability, including testing methodologies and certification to standards

Tools and services related to any or all aspects of the standardization life cycle

Proposals are also solicited for Feature Topic issues of the Communications Standards Supplement.

Articles should be submitted to the IEEE *Communications Magazine* submissions site at

http://mc.manuscriptcentral.com/commag-ieee

Select "Standards Supplement" from the drop-down menu of submission options.

Load Modulated Arrays: A Low-Complexity Antenna

Mohammad A. Sedaghat, Vlasis I. Barousis, Ralf R. Müller, and Constantinos B. Papadias

The authors describe promising recent progress in the area of massive antenna array architectures with low front-end hardware complexity. The presented technology enables the design and implementation of antenna arrays with large numbers of elements, while obtaining significant front-end hardware savings as compared to the conventional solutions.

Abstract

This article describes promising recent progress in the area of massive antenna array architectures with low front-end hardware complexity. The presented technology enables the design and implementation of antenna arrays with large numbers of elements, while obtaining significant front-end hardware savings as compared to the conventional solutions. This newly appearing design approach could be used in order to design either massive arrays with complexity that would be prohibitive with the current technology, or smaller arrays that offer high spatial degrees of freedom and are suitable for future small yet powerful cell nodes. RF hardware architectures with a single RF chain are reviewed, compared, and found superior to conventional MIMO implementations in terms of cost, dissipated heat, and physical size. The proposed improvements on the RF side allow the merging of the two dominant cellular technologies of virtual (distributed) and massive (centralized) MIMO into a hybrid approach of antenna arrays that is suitable for both large base stations and small (possibly cooperative) units such as remote radio heads.

Introduction

To satisfy the continuously increasing demand for higher data rates and mobility, and to meet the upcoming users' expectations, industrial partners and operators should be prepared for major changes ahead. The existing cellular network architectures seem to gradually be reaching their performance limits, indicating that the advent of new technologies and the advantageous exploitation of additional resources (e.g., more spectrum or antennas) are necessary.

Recent technological advances propose the concept of *distributed base stations* that are expected to eventually replace the existing bulkier base station (BS) units and play a crucial role in the near future. A key enabler in this new approach is the so-called remote radio head (RRH) [1], which is a compact and lightweight radio module that implements advanced wireless standards such as Long Term Evolution (LTE) and LTE-Advanced (LTE-A). Besides the radio interface, RRHs also carry an appropriate optical interface to allow connections to the fixed backbone network via optical fiber *fronthauling*. Under this new network perspective, intelligence is pushed away from the bulky BSs and aggregated toward the network backhaul where central processing units handle fast coordination/cooperation across a large number of cells. Moreover, today's large BSs are replaced by smaller and lightweight radio nodes that are distributed in space and allow for widely geographically distributed access via radio-over-fiber connections to the central node or BS.

In contrast to this vision, an alternative approach called massive multiple-input multiple-output (massive MIMO) is based on a sparser but *centralized* network infrastructure where a large amount of radio elements and processing are conferred to fewer but larger BSs. The more powerful BSs are able to eliminate intercell interference by focusing the signal power closely around the intended receiver in both the angular and radial domains, thus eliminating the need for additional cell-to-cell coordination and interference balancing/mitigation techniques that would significantly increase the network's burden. Massive MIMO can substantially increase the capacity as a result of the very high multiplexing gain that is potentially available (more details can be found in [2]). Furthermore, a significant improvement of the radiation efficiency can be achieved, mainly resulting from the high beamforming gain and the ability of targeting the transmitted power to a small region in space (i.e., the intended receiver).

On the other hand, the wide adoption of this approach is hampered by several limiting factors. Due to the large number of antenna elements, channel estimation becomes a nontrivial task. It seems to be feasible only if the system is operated in time-division duplex (TDD) mode relying on the sensitive property of channel reciprocity [2]. Pilot contamination, that is, the lack of enough signal dimensions to fit orthogonal pilot sequences for channel estimation [2], still requires further research. Recent promising progress can be found, for example, in [3–5]. From the hardware point of view, the current state of the art of massive MIMO faces two issues: cost and size restrictions. The RF-related costs grow linearly with the number of antennas. Each RF

Mohammad A. Sedaghat and Ralf R. Müller were with the Norwegian University of Science and Technology and are now with Friedrich-Alexander Universität Erlangen-Nürnberg; Vlasis I. Barousis was with Athens Information Technology; Constantinos B. Papadias is with Athens Information Technology.

 0163-6804/16/$25.00 © 2016 IEEE

chain in the current state of the art of massive MIMO transmitters includes a linear power amplifier, a mixer, and a digital-to-analog converter (DAC). Furthermore, to avoid the destructive effect of mutual coupling, antennas are spaced at least half a wavelength apart from each other, which leads to a size problem. Single-RF MIMO transmitters have been proposed recently to alleviate the aforementioned issues. It is shown in this article that single-RF MIMO systems enable large antenna arrays with significantly reduced RF cost. Moreover, single-RF MIMO systems allow for compact arrays without requiring any complicated RF circuitry. Note that single-RF MIMO systems can be utilized in centralized massive MIMO systems as well as distributed MIMO.

In this article, we envision the balanced hybrid combination illustrated in Fig. 1, that is, a highly distributed BS architecture in which the intelligence and all necessary processing are centralized in a main node that is connected to multiple powerful RRHs, each carrying a massive array. Bearing in mind the strict requirements of RRHs, *the circuit front-end hardware complexity of the massive array structures emerges as the key enabler technology that will bring this hybrid approach into reality*. As shown in Fig. 1, each RRH could serve its dedicated area (e.g., a femtocell) and could be equipped with a massive array with low hardware front-end complexity, which could be driven by the central unit through optical links.

The RRHs can be deployed, for example, on a rooftop or tower, or could be mounted on a wall. Assuming this *unified* architecture, our main focus in the following is reviewing the novel front-end circuit architectures that allow significant hardware savings and are required in order to make massive antenna arrays applicable to such architectures. In particular, the next section presents a general antenna architecture in a unified way, as well as two special cases that lead to low hardware complexity front-ends. As a clear step forward, we then focus on a novel front-end architecture that requires only a single carrier feeding and is best suited for massive array implementations.

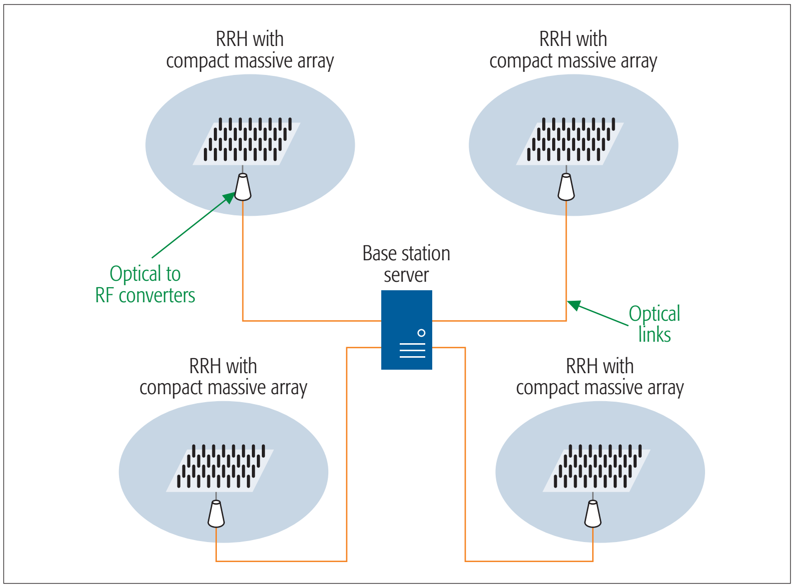

Figure 1. Proposed hybrid architecture. Each RRH can serve a femtocell, and all RRHs are aggregated to a common central unit that holds the intelligence.

A Unified Treatment of Antenna Arrays and Front-End Hardware Complexity Architectures

Multi-RF Arrays

This section provides a unified and qualitative description of the front-end feeding architectures for antenna arrays, with a special focus on those that offer significant hardware savings. The main concept is illustrated in Fig. 2, where the arrays are considered in the transmitting mode. The conventional MIMO implementation (for simplicity hereinafter called "multi-RF array") is shown in Figs. 2a and 2b, where each antenna is connected to its dedicated RF chain. If the antennas are closely spaced (as would typically happen in compact MIMO arrays), mutual coupling should be taken into account. In the presence of mutual coupling, a multi-port matching network could be used in order to decouple the antennas. However, implementing such a multiport matching network for a large array is very complicated. Therefore, in order to avoid the harmful effects of mutual coupling, the antennas are usually spaced about half a wavelength apart. This reduces the mutual coupling effect significantly. In this case, the impedance matrix of the antenna array is approximately diagonal, as shown in Fig. 2b. Note that in massive antenna arrays, antenna decoupling is very complicated; therefore, so far almost all of the proposals for massive MIMO assume that antennas are spaced well apart in order to minimize mutual coupling. The main drawback of the approaches in Figs. 2a and 2b is the cost of the RF chains. For N antennas, N power amplifiers, N mixers, and N DACs are required. On the other hand, we have single-RF MIMO arrays that provide more efficient implementation, as shown in Figs. 2c and 2d. In the following subsections, the single-RF MIMO concept is explained in greater detail.

Parasitic Array Architectures

An attempt to achieve hardware savings is shown in Fig. 2c, which describes the principle of a special antenna structure called *single-RF parasitic array*, in which only a single *active element* is fed by an external RF chain, while the remaining antennas are terminated through tunable analog loads and are known as *parasitics*. Such arrays were initially proposed for low-cost analog beamforming applications [6]. In contrast to the conventional arrays, functionality of parasitic arrays is essentially based on the requirement for strong mutual coupling among all antenna elements. This guarantees that the sole feeding at the active element can induce adequate currents on all of the parasitics, and hence all elements participate in the shaping of the radiation pattern. Multiple-active multiple-passive antenna systems have also been proposed using parasitic antennas (see [7, Ch. 8] for more details). The input currents can be further controlled by changing the

It turns out that the use of highly linear power amplifiers is mandatory in order to avoid possible distortion of the output signal. However, the transistor in this case should be biased in order to conduct during the entire cycle of the input signal, which causes power efficiency degradation.

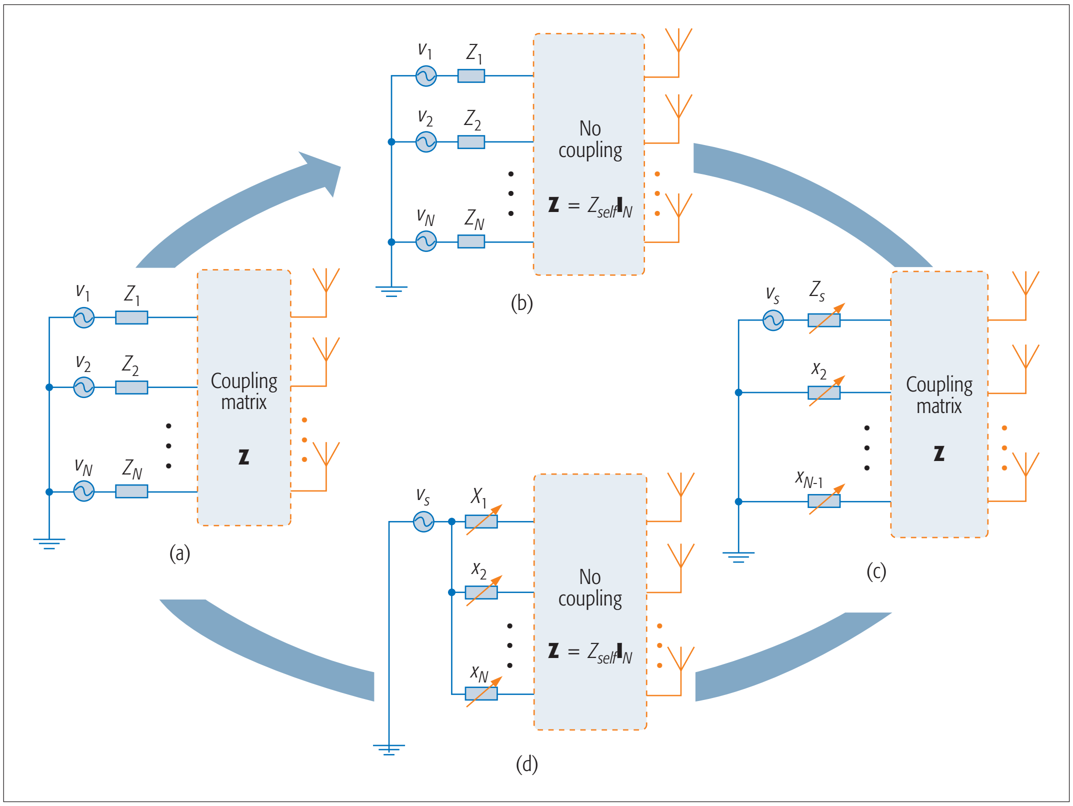

Figure 2. Unified illustration of different front-end circuit feeding architectures for antenna arrays.

effective couplings, which in turn are attained by tuning the analog loads. The analog loads can be implemented using controllable elements such as varactor diodes, or transistors. Apparently, this approach offers significant hardware savings, since the number of active RF chains is reduced to one, and the remaining RF chains are replaced by easy-to-implement tunable analog loads.

Thanks to the tunable analog loads and the consequent beam-shaping capability, parasitic arrays were recently proposed for *single-RF MIMO transmission*. The pioneering work in [8] proposes a smart analog switching technique in order to multiplex two symbols over the air, starting from the simple case of on-off keying (OOK) and continuing to low-order phase shift keying (PSK) modulation formats. This introductory work laid the foundation for further research on single-RF MIMO transmission with parasitic arrays (a detailed overview can be found in [7]). A proof-of-concept experiment for PSK modulation was reported in [9], which validates the functionality of single-RF MIMO transmitters. To support constellations beyond PSK, the authors in [10] propose *an active circuit design* with field effect transistors (FETs), which implements a tunable complex loading for parasitic elements with the real (i.e., resistive) part ranging from negative to positive values. As a clear step forward, the design guidelines and specifications have been described that should be met in order to achieve an arbitrary space-time precoding scheme [11]. *The common baseline in all these works is to design appropriate antenna arrays and loading circuits in order to emulate the MIMO effect in the air by switching to different transmit radiation patterns in consecutive signaling periods.* Since this technology strongly requires mutual coupling among adjacent elements, it is not appropriate for base stations with large inter-element distances.

The Load-Modulated Array Architecture

Inspired by load tuning as a means to control the array beam pattern, Fig. 2d illustrates a newly introduced approach, the main competitive advantage of which is that *it completely eliminates the need for any RF chain*. As observed, all elements are connected to a common carrier signal source through passive and lossless two-port loading networks, each implemented with tunable reactance elements. There is a central power amplifier that amplifies the carrier signal, and there is no need for any mixer and DAC. The purpose of the tunable loads is to adjust the input currents to all radiating elements, *thereby implementing a desired signal constellation in the analog domain.* This alternative solution is known as *load-modulated arrays* (LMAs) [12, 13], reflecting their principle of operation. In LMAs, in contrast to the parasitic array architecture, all antenna elements are connected directly to a common amplifier, and couplings among the elements are not necessary. However, the case of compact and to some extent coupled arrays should not be excluded. In load-modulated arrays, all the processing, such as precoding, coding, mapping, and multi-carrier processing, can be done in the baseband. In fact, the required antennas' currents are calculated in the baseband, and then the states of the tunable elements are

	Advantages	Drawbacks
Multi-RF arrays	• Simple design and implementation. • Digital spectral shaping. • Low distortion.	• The number of RF chains is equal to the number of antennas. • Require amplifiers with very high backoff, leading to low power efficiency. • Require DACs and mixers. • Require a very complicated matching network in the presence of mutual coupling.
Parasitic arrays	• Mutual coupling is permitted (in fact it is required). • Allow for compact designs. • A sole full RF chain is required (consisting of amplifier, DAC, and mixer). • High power efficiency. • Suitable for compact devices with low or moderate number of antennas.	• Not suitable for large arrays. • Analog spectral shaping. • Requiring active loads in the case of arbitrary modulation. • Crossover distortion effects in switching time (which could be eliminated by appropriate filtering).
Load-modulated arrays	• One RF-chain. • One amplifier without any backoff. • No need for DAC and mixer. • Arbitrary modulation with passive loads only.	• Suitable for large arrays. • Analog spectral shaping. • Crossover distortion effects in switching time (which could be eliminated by appropriate filtering).

Table 1. A short list of pros and cons of multi-RF, parasitic, and load-modulated arrays architectures.

changed accordingly. Note that the structure of the transmitter does not change when type of precoding or modulation changes.

In load-modulated single-RF MIMO, a spectral shaping filter is implemented in the analog domain; thus, the sample rate can be reduced to the symbol rate. This simplifies the baseband block a little and adds the complexity of the RF part. Furthermore, since there is no need for the DAC and mixer in load-modulated single-RF MIMO transmitters, the RF part is much cheaper than in the classical MIMO transmitters (multi-RF MIMO).

Table 1 presents a summary of the pros and cons of the above mentioned technologies.

Massive Array Architectures with Low Front-End Hardware Complexity

Front-End Architecture and Challenges

The requirement for power amplifiers with highly linear characteristics and low power dissipation is a crucial problem in wireless networks. Indeed, Conventional MIMO systems often require a high amplifier backoff to serve signals with high peak-to-average-power ratio (PAPR) that might reach the level of 8 to 12 dB [12], for example, in orthogonal frequency-division multiplexing (OFDM) and/or linear precoding. Therefore, it turns out that the use of highly linear power amplifiers is mandatory in order to avoid possible distortion of the output signal. However, the transistor in this case should be biased in order to conduct during the entire cycle of the input signal, which causes power efficiency degradation. From the circuit design point of view, switching power amplifiers, comprising a symmetric topology of complementary blocks (e.g., push-pull) and multiple active elements, can be used to increase the power efficiency. Although the switching mode improves the total power efficiency, this comes at the cost of linearity caused due to, for example, crossover distortion effects produced when the transistors switch over from one to the other. Crossover distortion happens during the time interval in which one of the transistors shuts off and the other turns on (e.g., in class A/B power amplifiers) [14].

Proposed Massive Array Architecture and Indicative Performance Results

The novel architecture in Fig. 2d provides a low-cost massive array implementation with significant hardware saving since there is no need for any DAC and mixer in it; this would in turn bring the envisioned hybrid network topology in Fig. 1 closer to reality. As shown in Fig. 3, in load-modulated MIMO transmitters, all antenna elements are connected to the sole power amplifier through passive two-port lossless loading networks of "T" or "P" topology [13]. The purpose of such loading networks is to adjust the input currents of the radiating elements based on the desired signaling format. Note that the central power amplifier in LMAs amplifies a sinusoid signal with a PAPR of 0 dB. Unlike the parasitic arrays, the advantageous point here is that all antenna elements are connected to a common source, which is a fixed sinusoid carrier signal. This allows us to use a highly efficient nonlinear power amplifier. A circulator protects the power amplifier against the reflected power, which (if there is any) will be consumed onto the resistance [12, 13].

The proposed front-end architecture becomes more favorable in the massive MIMO regime with large numbers of antennas since hardware drawbacks tend to vanish in this regime, as explained in the following parts.

The PAPR Tends Asymptotically to One: In the LMAs shown in Fig. 3, the PAPR of the signal at the output of a circulator is merged to one as the number of antenna elements becomes larger. This is explained by the law of large numbers. In fact, it results from the averaging effect, which kicks in when many antenna elements coexist.

Assuming complex Gaussian input data signals and N antenna elements, the authors in [13] explain that the output power is chi-squared distributed with $2N$ degrees of freedom. The reduction of the PAPR value due to the averaging is drawn in Fig. 4a for different values of clipping probability. As observed, for $N = 100$ elements and 0.1 percent clipping probability, the PAPR is found to be around 1.2 dB. Refer-

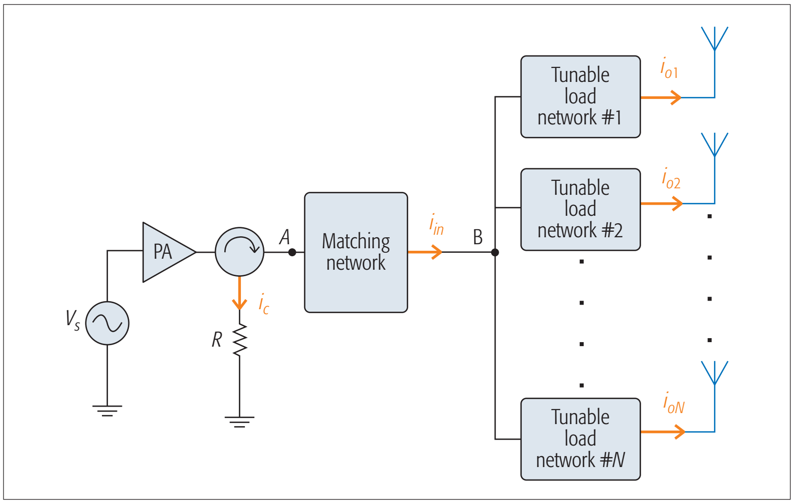

Figure 3. Load-modulated arrays for low-cost massive MIMO design with low hardware complexity and high efficiency.

ring to the discussion above, this definitely indicates that there is no need for high linearity and the use of low-efficient class A/B power amplifiers. Instead, a single and highly efficient class F power amplifier for constant envelop signals can be used that offers significantly higher efficiency of around 80 percent due to its switching mode of operation [14]. Note that the loss in the resistance connected to the circulator should be taken into account in the total power efficiency calculation [12].

Good Matching Conditions Are Achieved: As opposed to conventional MIMO transmitters that are equipped with fixed RF front-end modules and externally varying voltages, Fig. 3 shows that load-modulated arrays are driven by a fixed carrier signal source. Furthermore, the signals on antennas are controlled by tuning the impedance characteristics of the loading networks. Thus, *the input impedance, as seen from the analog source's side, varies with time as a function of the desired currents in consecutive signaling periods.* Thankfully, as the number of elements grows and the massive regime is reached, the load modulation tends to be insensitive to mismatch effects. This is illustrated in Fig. 4b, which shows the probability density function (pdf) of the voltage standing wave ratio (VSWR) vs. the number of antenna elements. It is observed that

$$\lim_{N\to\infty} VSWR = 1$$

with probability 1, which implies that as the number of elements grows, the VSWR converges to 1, and the power amplifier is well matched to the antennas.

Low Signal Distortion: In conventional MIMO systems (multi-RF arrays), the input signals are clipped independently, since every antenna is connected to an individual power amplifier. On the contrary, in load-modulated arrays the input signals are clipped only if the sum power becomes higher than the maximum power of the sole amplifier. The authors in [12] analyze two clipping strategies: *clipping with respect to the minimum mean squared error* (MMSE) and *equal clipping of all signals*. The first approach minimizes the total current distortion, and the second approach clips all signals equally in order to satisfy the power constraint of the sole power amplifier. Clearly, the clipping effect introduces a signal distortion that thankfully weakens as the number of antennas grows. In fact, clipping in single-RF MIMO transmitters is less harmful than in conventional MIMO transmitters. Indicatively, Fig. 4c draws the distortion vs. the number of antennas for fixed values of total power efficiency. It is observed that the two clipping strategies exhibit almost the same behavior, and the distortion becomes negligible as the number of antenna elements becomes large.

Implementation Aspects of the Load Modulators

Tunable loads can be implemented in two ways:
- Soft tuning of some tunable capacitors implemented by varactors or transistors
- Discrete tuning using some RF switches implemented by PIN diodes, Schotkky diodes, or micro-electro-mechanical systems (MEMS)

The first scheme suffers from the nonlinearity of the tunable devices. On the other hand, the second scheme omits the nonlinearity by using hard switching. Discrete tuning is much easier to handle since there is no nonlinear component in the circuit and, except in the short switching time interval, the circuit is linear. Switching on the order of nanoseconds can be done using fast PIN or Schottky diodes [15]. Switching in the discrete tuning case is done by changing the bias voltage of diodes. This can be done by a simple level shifter circuit, and there is no need for a DAC.

The load modulators can be implemented by a combination of "T" or "Π" two-port network topologies [12]. As an extension to this approach, a six-port analog modulator has been demonstrated in [15] that uses Schottky diodes and achieves high switching speed, enhancing the achievable data rate, for example, up to 1.2 Gb/s as reported in [15].

Although the exact computation of the load values that result in a desired signaling format is known (e.g., [11, 12]), a possible deviation from those desired values cannot be excluded in the real world. Reasons for that could be some additional parasitic effects due to the high frequency of operation, or any practical non-idealities caused by the analog elements (e.g., varactor diodes). Inevitably, any loading deviation would affect the output signals and in turn would degrade the performance to some extent. However, the robustness of the proposed architectures has been evaluated and drawn, showing that the proposed architecture performs satisfactorily [10–12].

Challenges and Issues

Utilizing load modulators in MIMO transmitters leads to some challenges and issues. The first challenge is about the trade-off between the number of switches and the loss in each load modulator. The higher the number of switches, the lower the output signal distortion. On the other hand, every switch adds an insertion loss to the circuit. One needs to optimize the number of switches by designing the circuit carefully. Furthermore, designing very low-loss switches is extremely important here since it allows for high resolution signaling. The second challenge is implementation of the analog filters. The analog filters at the output of load modulators add some losses and should be designed carefully in order to mitigate out-of-band radiations created by the switching. Cross-talk between antennas is another challenge here. By changing the state of each switch in one load modulator, the current on the other antennas will also be affected. This effect can be compensated in the baseband by modeling the cross-talk between the branches.

Conclusions and Outlook

Hardware savings in MIMO transceivers have been a crucial and challenging problem. Massive arrays are gaining increasing attention and are expected to be one of the leading technologies in the emerging fifth generation heterogeneous/small cell network architectures. Furthermore, the expressed interest in unlicensed bands in higher microwave and even millimeter-wave regimes will make massive arrays necessary in order to overcome the path loss at such frequencies. Finally, the smaller wavelengths will aid their inclusion into size-restricted access points. This article promotes a newly appearing front-end RF architecture that offers significant hardware savings and is best suited for massive array deployments. The proposed low-complexity RF architecture will be the key enabling technology to pave the way forward toward novel and emerging cellular architectures. The proposed massive array architecture needs only a single-carrier feeding signal, thus eliminating completely the need for a complete RF chain. With discrete load modulators, it even eliminates the need for a DAC. Under this perspective, the input currents to the diverse antenna elements are adjusted by tuning the front-end feeding circuit itself, which in turn is attained by switching some diodes in some load modulators that are attached to the antenna elements. Letting the number of elements increase, it has been shown that a switching power amplifier suffices and allows for significant power savings. Moreover, the mismatch effects as well as the signal distortion tend to attenuate and become negligible by increasing the number of antenna elements.

Overall, this article provides a roadmap for further exploration and aims to trigger more intensive research for future lightweight powerful massive MIMO transceivers that will bring flexible cellular networks with powerful nodes closer to reality. As further research, analyzing the cross-talk effect between the branches in load-modulated arrays is an important next step. Furthermore, finding appropriate filtering methods, such as analog spectral shaping filtering, as well as investigation of load modulation design in the case of closely spaced antennas will be the subject of future work. Finally, building LMAs and tuning them to the demands of emerging RRH, massive, and mmWave systems is of high priority.

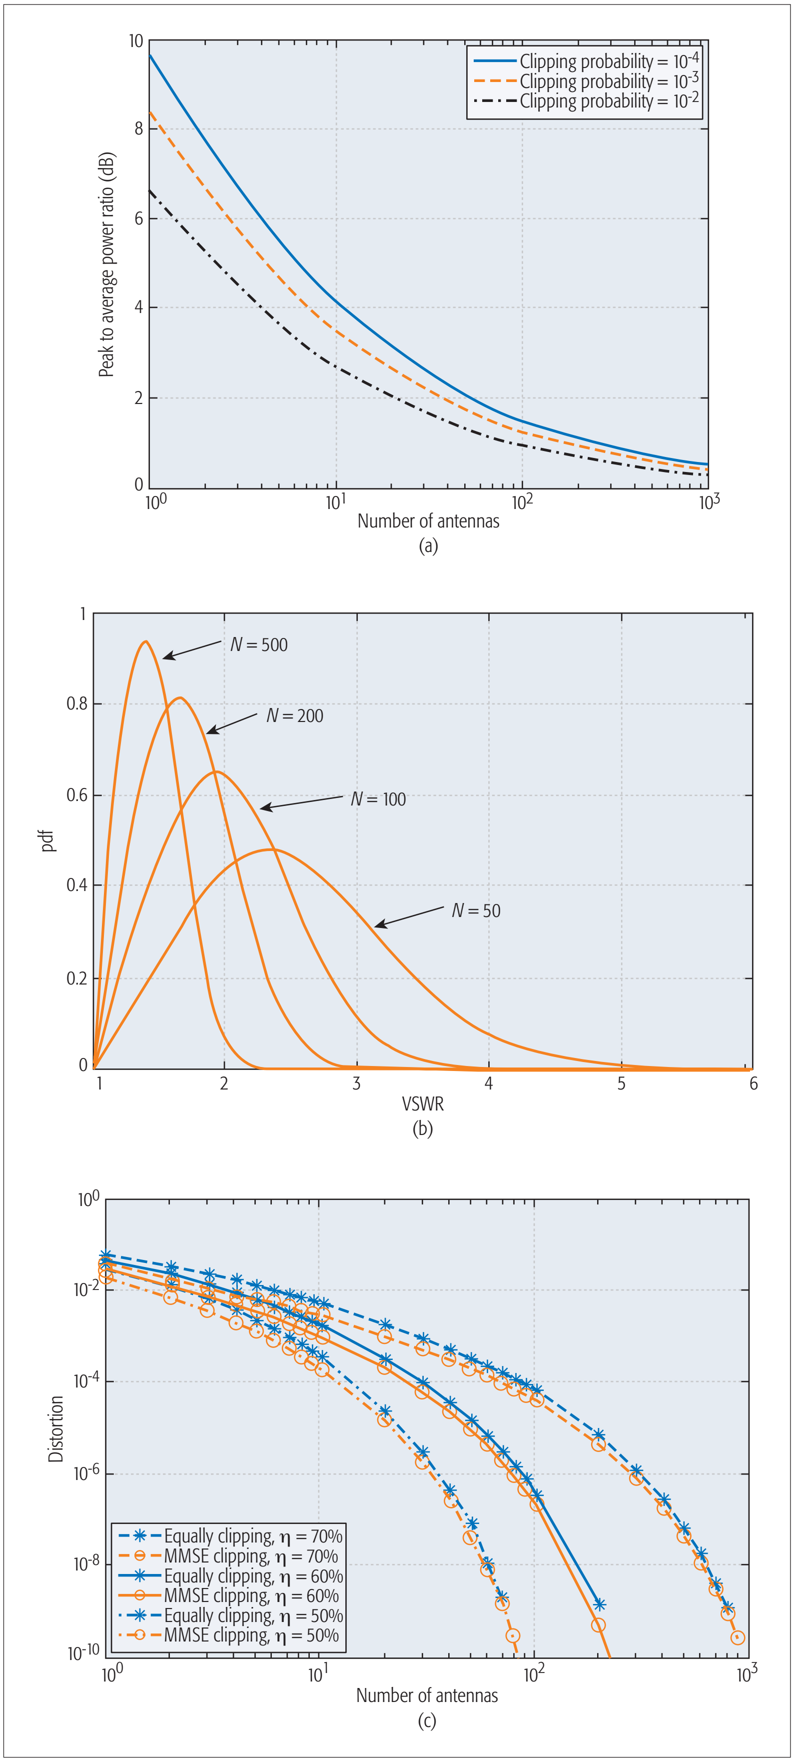

Figure 4. a) PAPR vs. number of antennas for different clipping probabilities; b) VSWR distribution for different numbers of antenna elements; c) signal distortion vs. number of antennas for different clipping methods and total efficiencies.

Acknowledgment

This work has been supported by the EC FP7 project HARP (http://www.fp7-harp.eu/) under grant number 318489.

References

[1] G. Kardaras and C. F. Lanzani, "Advanced Multimode Radio for Wireless and Mobile Broadband Communication," *Proc. IEEE Euro. Wireless Tech. Conf.*, Sept. 2009, pp. 132–35.
[2] E. Larsson *et al*., "Massive MIMO for Next Generation Wireless Systems," *IEEE Commun. Mag.*, vol. 52, no. 2, Feb. 2014, pp. 186–95.
[3] R. R. Müller, L. Cottatellucci, and M. Vehkaper, "Blind Pilot Decontamination," *IEEE J. Sel. Topics Signal Processing*, vol. 8, no. 5, Oct. 2014, pp. 773–86.
[4] H. Yin *et al*., "A Coordinated Approach to Channel Estimation In Large-Scale Multiple-Antenna Systems," *IEEE JSAC*, vol. 31, no. 2, Feb. 2013, pp. 264–73.
[5] H. Yin *et al*., "Pilot Decontamination Using Combined Angular and Amplitude Based Projections in Massive MIMO Systems," *Proc. 16th Int'l. Wksp. Signal Processing Advances in Wireless Commun.*, June/July 2014, Stockholm, Sweden, pp. 216–20.
[6] T. Ohira and K. Gyoda, "Electronically Steerable Passive Array Radiator Antennas for Low-Cost Analog Adaptive Beamforming," *Proc. IEEE Int'l. Conf. Phased Array Sys. and Tech.*, 2000, pp. 101–04.
[7] A. Kalis, A. Kanatas, and C. B. Papadias, Eds., *Parasitic Antenna Arrays for Wireless MIMO Systems*, Springer, 2014.
[8] A. Kalis, A. Kanatas, and C. B. Papadias, "A Novel Approach to MIMO Transmission Using a Single RF Front End," *IEEE JSAC*, vol. 26, no. 6, 2008, pp. 972–80.
[9] O. Alrabadi *et al*., "A Universal Encoding Scheme for MIMO Transmission Using a Single RF-Front End," *IEEE Trans. Wireless Commun.*, vol. 8, no. 10, Oct. 2010, pp. 5133–43.
[10] B. Han *et al*., "A Single RF MIMO Loading Network for High Order Modulation Schemes," *Int'l. J. Antennas Propagation*, vol. 2014, Article ID 879127, 10 pages, 2014; doi:10.1155/2014/879127
[11] V. Barousis and C. B. Papadias, "Arbitrary Precoding with Single-Fed Parasitic Arrays: Closed-Form Expressions and Design Guidelines," *IEEE Wireless Commun. Lett*., vol. 3, no. 2, April 2014, pp. 229–32.
[12] M. A. Sedaghat, R. R. Müller and G. Fischer, "A Novel Single-RF Transmitter for Massive MIMO," *Proc. 18th Int'l. ITG Wksp. Smart Antennas,* Mar. 2014, Erlangen, Germany, pp. 1–8.
[13] R. R. Müller, M. A. Sedaghat, and G. Fischer, "Load Modulated Massive MIMO," *IEEE Global Conf. Signal and Info. Processing*, Dec. 2014, Atlanta, GA, pp. 622–26.
[14] S.C. Cripps, *RF Power Amplifiers for Wireless Communications*, Artech House, 2006.
[15] J. Östh *et al*., "Schottky Diode as High-Speed Variable Impedance Load in Six-Port Modulators," *Proc. IEEE Int'l. Conf. Ultrawideband*, Sept. 2011, Bologna, Italy, pp. 68–71.

Biographies

Mohammad Ali Sedaghat (mohammad.sedaghat@iet.ntnu.no) received his B.Sc degree in electrical engineering from Amirkabir University of Technology, Iran, in 2006 and his M.Sc degree in communication systems from Sharif University of Technology, Iran, in 2008. Currently, he is working toward his Ph.D degree at the Norwegian University of Science and Technology (NTNU), Norway. His current research interests include the general field of wireless communications with emphasis on MIMO systems.

Vlasis I. Barousis received his B.Sc. degree in electronics engineering in 2004 and his M.Sc. in digital communications and networks in 2006. Since 2011 he holds a Ph.D. from the University of Piraeus, Department of Digital Systems, in the field of wireless communications. His current research interests span the area of smart antennas, especially the design of single-RF compact parasitic arrays for MIMO systems.

Ralf R. Müller [S'96, M'03, SM'05] is a professor at Friedrich-Alexander-University Erlangen-Nuremberg (FAU), Germany. He is also an adjunct professor at NTNU. He has held visiting appointments at Princeton University, Institute Eurcom, France, University of Melbourne, Australia, University of Oulu, Finland, National University of Singapore, Kyoto University, Japan, Technical University of Munich, Germany, Babes-Bolyai University, Romania, and FAU.

Constantinos B. Papadias [S'88, M'96, SM'03, F'13] is dean of Athens Information Technology (AIT), Greece, and head of AIT's Broadband Wireless & Sensor Networks (B-WiSE) Research Group. He is also an adjunct professor at Aalborg University in Denmark. He is a graduate of the National Technical University of Athens, Greece (1991), and ENST, France (PhD, 1995). Prior to his current positions he held posts at Eurecom, Stanford, Bell Labs, Columbia University, and Carnegie Mellon University.

A Flexible 5G Frame Structure Design for Frequency-Division Duplex Cases

Klaus I. Pedersen, Gilberto Berardinelli, Frank Frederiksen, Preben Mogensen, and Agnieszka Szufarska

Abstract

A 5G frame structure designed for efficient support of users with highly diverse service requirements is proposed. It includes support for mobile broadband data, mission-critical communication, and massive machine communication. The solution encompasses flexible multiplexing of users on a shared channel with dynamic adjustment of the transmission time interval in coherence with the service requirements per link. This allows optimizing the fundamental tradeoffs between spectral efficiency, latency, and reliability for each link and service flow. The frame structure is based on in-resource physical layer control signaling that follows the corresponding data transmission for each individual user. Comparison against the corresponding LTE design choices shows attractive benefits.

Introduction

Research toward a new fifth generation (5G) air interface is currently ongoing in both academia and industry. This includes defining 5G requirements and identifying candidate techniques to be included in a future system design. Despite the relatively short time of 5G research, the open literature includes an impressive number of 5G related studies; hence, we only provide pointers to some of those in the following. Among others, the METIS project has outlined its 5G vision in [1], the 5GNOW project presented their proposal for asynchronous access and related waveform designs in [2], while the use of more advanced centralized network architectures for 5G was suggested in [3]. Furthermore, small cell optimized design has been identified as being of particular importance to be able to meet the future mobile broadband traffic requirements [4, 5]. There is consensus on the fact that 5G should push the performance limits significantly further toward having virtually zero latency and multi-gigabit-rate end-user experience, and efficient machine-type communication (MTC), depending on the application requirements [6, 7]. For fulfilling such diverse (and sometimes conflicting) requirements, our hypothesis is that a highly flexible and configurable air interface is needed. In that context, the radio frame structure, and especially the methods for multiplexing (mux) of users, are some of the key design choices.

Our focus is therefore on presenting a flexible frame structure capable of fulfilling the challenging 5G requirements for efficient support of a mixture of diverse services. We start by identifying the main requirements and spectrum availability. Although we strive toward having an agnostic solution that is carrier-frequency-independent, we primarily focus on use cases for below 6 GHz for early 5G deployments around 2020. This is motivated by the fact that spectrum regulators will discuss band allocations for mobile communications above 6 GHz no sooner than 2019. The derived flexible frame structure is presented for the frequency-division duplex (FDD) use case applicable for macro-cell deployments. However, several merits of the suggested solution are equally applicable for time-division duplex (TDD) bands. An air interface with orthogonal frequency-division multiple access (OFDMA) is assumed, where users are scheduled on a time-frequency grid of resources [8]. However, the proposed frame structure is also applicable for other candidate waveforms that offer a time-frequency symbol space for a commonly shared channel per cell. The corresponding relationship between physical (PHY) layer control and data channels is outlined, and numerical results are presented. Throughout the article, the Long Term Evolution (LTE) 4G standard [9, 10] is used as our reference for motivating and quantifying the benefits of the new 5G frame structure.

Requirements and Spectrum

Air Interface Requirements

The International Telecommunication Union (ITU) has recently defined challenging requirements for international mobile telecommunications (IMT) in 2020 and beyond [6]. Among others, peak data rates of even up to 20 Gb/s and uniform availability of end-user-experienced data rates of 100 Mb/s to 1 Gb/s are listed. Support for mobile broadband (MBB) requires relatively large bandwidth and frequent transmissions. In addition to offering connectivity for humans, 5G should also be designed for efficient MTC. MTC use cases include massive machine communication (MMC) with a large number of connected

The authors propose a 5G frame structure designed for efficient support of users with highly diverse service requirements. It includes support for mobile broadband data, mission-critical communication, and massive machine communication. The solution encompasses flexible multiplexing of users on a shared channel with dynamic adjustment of the transmission time interval in coherence with the service requirements per link.

Klaus Pedersen and Preben Mogensen are with Nokia Networks and Aalborg University; Gilberto Berardinelli is with Aalborg University; Frank Frederiksen and Agnieszka Szufarska are with Nokia Networks.

0163-6804/16/$25.00 © 2016 IEEE

High contiguous carrier bandwidths of 100–200 MHz are especially relevant for the 3–6 GHz spectrum range, while carrier bandwidths of up to 40–100 MHz for sub-3-GHz FDD deployments are likely adequate in combination with efficient spectrum aggregation techniques.

low-cost devices (e.g., sensors). MMC is characterized by infrequent access, typically transmitting only moderate size payloads with relaxed latency requirements. Devices for MMC are typically associated with requirements for extremely low energy consumption and low cost, meaning that it is desirable to have such devices operate with relatively low radio bandwidth transmission and reception leading to lower transceiver complexity. The second class of MTC use cases is mission-critical communication (MCC). MCC requires lower end-to-end latency and a high degree of reliability to, for example, support vehicular use cases and factory automation processes. In this context, ITU has set a target to achieve a 1 ms over-the-air communication round-trip time (RTT) for a single transmission. This includes transmission of the payload until the corresponding acknowledgment (Ack) is received. Depending on the application, reliability constraints of up to six-sigma (99.99964 percent) are mentioned [1]. For more information on 5G requirements, also see [7].

Designing a system that supports all of the MBB, MMC, and MCC targets is rather challenging, especially since there are fundamental trade-offs in wireless systems between offering high spectral efficiency, low latency, and high reliability [11]. As an example, the performance of MBB can approach the Shannon capacity limit, while there is a cost of reduced spectral efficiency if operating under strict latency and reliability constraints. Our hypothesis is therefore that this calls for a flexible air interface design that allows optimizing each link according to its service requirements. This suggests having a dynamic frame structure that offers the possibility to perform trade-offs between spectral efficiency, energy efficiency, latency, and reliability in coherence with the requirements per link.

Spectrum Availability

Nowadays, the allocated spectrum for mobile communication is all below 6 GHz, and the 2015 World Radiocommunication Conference (WRC) will focus on further sub-6-GHz band allocations. WRC 2019 is expected to also consider band allocations above 6 GHz (e.g., for future 5G deployments). This essentially means that the first commercial 5G deployments around 2020 will need to focus on frequencies below 6 GHz due to regulatory band constraints for mobile communications. As illustrated in Fig. 1, the spectrum below 6 GHz is rather fragmented and composed of a mixture of bands for operating with FDD and TDD, also referred to as paired and unpaired bands, respectively. Depending on the region, potentially up to 2 GHz of spectrum is available below 6 GHz for future mobile radio communication, with nearly equal availability of bands for FDD and TDD deployments. It is especially worth noting that less than half of the potentially available spectrum for mobile communications below 6 GHz is used today. Bands for FDD are primarily available below 3 GHz, although some FDD bands are also available at higher frequencies. The efficient utilization of the spectrum below 6 GHz calls for supporting different carrier bandwidths and flexible spectrum aggregation techniques. In LTE, spectrum

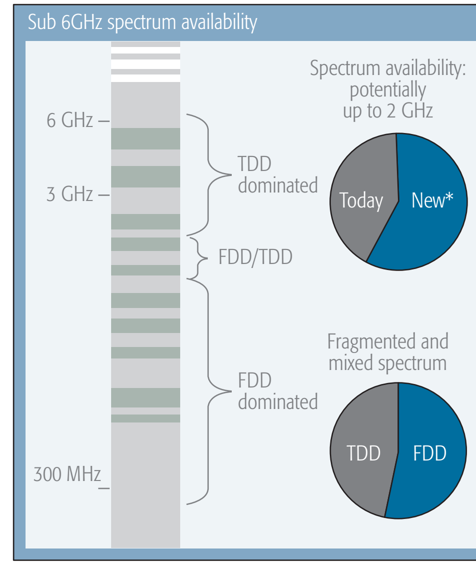

*Potentially new spectrum available for mobile communications, subject to spectrum regulator decisions.

Figure 1. Overview of spectrum availability below 6 GHz.

aggregation is supported in the form of carrier aggregation (i.e., cell aggregation) [12], while one enhancement under study for 5G is support for aggregating fragmented spectrum to form one logical cell. High contiguous carrier bandwidths of 100–200 MHz are especially relevant for the 3–6 GHz spectrum range, while carrier bandwidths of up to 40–100 MHz for the sub-3-GHz FDD deployments are likely adequate in combination with efficient spectrum aggregation techniques.

In this study, our focus is on the design of a flexible spectrum-agnostic frame structure for the spectrum below 6 GHz, with emphasis on solutions for licensed FDD (paired) deployments. The lower FDD bands are especially attractive for providing wide area coverage and outdoor-to-indoor coverage due to the more favorable radio propagation conditions compared to using higher frequency bands. Specifics related to frame design for unlicensed small cell TDD operating below 6 GHz are outside the scope of this study.

Flexible Frame Structure

Time-Frequency Multiplexing of Users

The ability to efficiently adapt and optimize the radio resources for each user in coherence with its service requirements is desirable. Among other things, this requires a highly flexible frame structure. The basic concept is illustrated with the time-frequency grid depicted in Fig. 2a, where a number of users are flexibly multiplexed over the available resources with different transmission time interval (TTI) durations. Each tile refers to the smallest allocation unit of time duration Δt and frequency size Δf. In practice, Δt would equal an integer number of orthogonal frequency-division multiplexing (OFDM) symbols, while Δf corresponds to an integer number of subcarriers. Those values could equal just a few symbols and/or subcarriers. The value of Δt determines the minimum TTI size for scheduling a user,

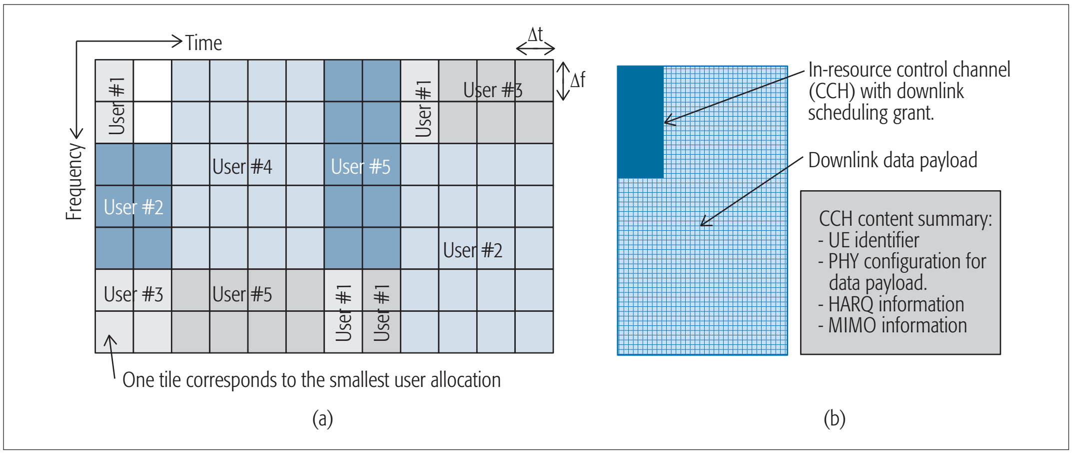

Figure 2. Dynamic time-frequency multiplexing of users and related scheduling grants: a) time-frequency multiplexing of users; b) in-resource control signaling.

LTE features a strict periodic time-division separation of the PHY control and data, by sending the control information in the first set of OFDM symbols. For example, the physical downlink control channel is transmitted over the full system bandwidth in the first OFDM symbols of the TTI having a fixed duration of 1 ms.

as well as the resolution for other TTI scheduling options. Given the most stringent latency requirement of 1 ms for MCC, there is consensus in the research community that a minimum TTI size of no more than 0.2–0.25 ms is needed. The reduced TTI size, combined with stricter network and device processing requirements, allows a sufficient delay budget for sending the payload, receiving and processing it, followed by sending the corresponding Ack. As an example, the 5G small cell concept presented in [4] proposes $\Delta t = 0.25$ ms.

The frame structure allows the TTI size for each scheduling instant of the users to be dynamically adjusted. Thus, some users can be scheduled with a short TTI size of Δt to fulfill the RTT requirement for MCC. However, scheduling all users with this short TTI is not optimal. Using long TTIs allows us to benefit from larger coding gains to approach the Shannon capacity limit, and it also imposes lower control overhead. This comes, however, at the expense of latency increase; in that respect, the usage of longer TTIs is more beneficial for MBB users for which the required data rate may be high, and the latency requirements are less stringent. Setting the TTI size per scheduling grant furthermore offers the possibility to optimize the MBB services using TCP. During the initial data transmission session, the end-user-experienced performance is primarily determined by the RTT due to the slow start TCP procedure (i.e., TCP flow control). Therefore, it would be advantageous to first perform scheduling of the MBB TCP users with short TTIs, followed by longer TTI sizes when reaching steady state operation. In addition to the time-domain scheduling flexibility, the frame structure also allows dynamic frequency-domain scheduling, where users are served on different parts of the carrier bandwidth. This includes scheduling users on non-consecutive frequency blocks to benefit from frequency domain scheduling diversity, as known from LTE. Moreover, scheduling of low-cost MTC devices with reduced bandwidth capabilities on a small portion of the carrier bandwidth is supported as well. The MTC devices served within a narrow bandwidth can be scheduled with a longer TTI size to gain from time diversity, that is, to compensate for the lack of frequency diversity.

In the uplink direction, the flexibility to schedule users with different TTI sizes offers further advantages. While users with moderate path loss toward their serving base station are schedulable on a larger bandwidth with short TTI sizes, coverage-challenged UEs need to be scheduled with longer TTIs on a narrow bandwidth to have a sufficiently high received energy (and power spectral density) at the base station. The latter is, for example, the case for deep indoor users that experience high indoor-to-outdoor penetration loss. Restricting the uplink scheduling to always have short TTI size would therefore have a cost in terms of reduced uplink data coverage.

In-Resource Control Signaling

In-resource physical layer control signaling for sending the scheduling grant pointing to the users' data transmission allocation is proposed. The main idea is to use embedded on-the-fly information to the users on its allocated time-frequency resources, as well as the additional information needed to decode the data. This is referred to as the users scheduling grant sent on a dedicated PHY control channel (CCH). The scheduling grant also contains information such as the allocated time-frequency resources for the users (number of consecutive time symbols per TTI, subcarrier allocation), the modulation and coding scheme (MCS), hybrid automatic repeat request (HARQ) information, and multi-antenna transmission information (e.g., number of spatial streams). The in-resource CCH is mapped at the start of the resource allocation for the user in the first time symbol(s) and over a limited part of the frequency resources, as shown in Fig. 2b. Note that the flexible allocation of in-resource CCH differs significantly from the solutions adopted in the current LTE standard. LTE features a strict periodic time-division separation of the physical layer control and data, by sending the control information in the first set of OFDM symbols [7, 13]. For example, the physical downlink CCH (PDCCH) is transmitted over the full system bandwidth in the first OFDM symbols of the TTI having a fixed duration of 1 ms.

Each uplink data transmission needs an

It is naturally desirable to gain from using beamforming for both PHY CCH and data channels. This is possible due to the in-resource position of the CCH, which allows using beamforming for both the CCH and the corresponding downlink data transmission in case of single stream transmission.

uplink scheduling grant that is sent in the downlink. In that respect, we opt for an uplink grant solution as illustrated in Fig. 3. In Fig. 3a, joint downlink and uplink grants are multiplexed on the same control resources dedicated to a specific user (illustrated by the purple scheduling grant), with the fundamental difference that the downlink part provides information for decoding the associated data block, while the uplink part points to a successive uplink data transmission allocation. If downlink data transmission does not occur for the user, the uplink grant can be transmitted independently, as shown in Fig. 3b (green scheduling), where multiple uplink scheduling grants are stacked in one downlink resource unit; that is, scheduling users #3 and #4 in the uplink, while scheduling users #1 and #2 in the downlink (dark blue scheduling grants).

In the interest of UE complexity and power consumption, it should be possible for the network to configure each UE with a tile pattern for monitoring the downlink CCH scheduling grants. This allows configuring low-cost MTC UEs with low data rate requirements to only monitor the downlink CCH transmissions on a narrow bandwidth of Δf and at a sparse time resolution. On a similar note, UEs with MCC can be configured to monitor for CCH transmissions on a larger bandwidth every Δt to fulfill stricter latency and reliability requirements. Finally, MBB users could be configured to for example, monitor only every nth and mth tile in the time and frequency domain, respectively. The configuration of each UE with a tile pattern for monitoring downlink CCH scheduling grants corresponds to a flexible time-frequency domain discontinuous reception (DRX) mechanism. The time-frequency domain DRX mechanism offers the possibility to control the trade-off between scheduling flexibility and UE power consumption for each link. The DRX configuration of the users is assumed to happen via higher-layer signaling, asynchronously among users, and primarily configured at connection setup. Note that LTE only supports the configuration of time-domain DRX patterns as the CCH carrying the scheduling grants is transmitted on the full carrier bandwidth.

Hybrid Automatic Repeat Request

HARQ is assumed to be an essential technique for 5G. In order to allow a high degree of scheduling flexibility, asynchronous and adaptive HARQ is assumed for both the downlink and uplink. This is in contradiction with LTE, where a synchronous HARQ solution is selected for the uplink, reducing signaling overhead, but also reducing the time-wise scheduling flexibility and implying rigid timing requirements at the same time. It is therefore proposed to have the timing for sending Acks and negative Ack (Nacks) configurable per link, as well as the number of parallel stop-and-wait (SAW) channels. The latter not only offers increased scheduling flexibility, but also more degrees of freedom for network implementations. The latter is of particular relevance for cases where the PHY and medium access control (MAC) hosting the HARQ functionality are separated on different hardware units with inter-communication delays. This is important, for instance, for supporting centralized radio network implementations with different fronthaul latencies.

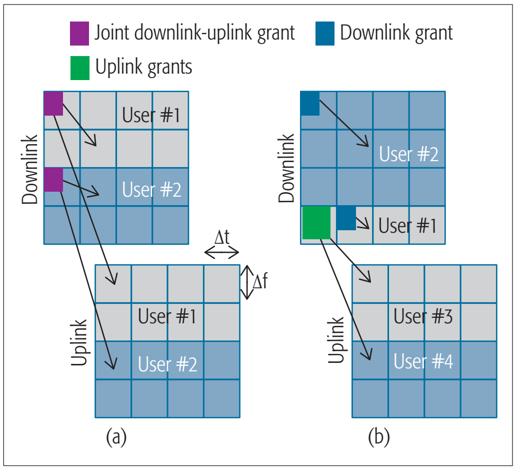

Figure 3. In-resource control channel design for a) joint uplink and downlink scheduling; b) separate uplink and downlink scheduling.

Beamforming and Interference Coordination

Beamforming and massive multi-antenna techniques are important techniques for improving the performance of 5G. In this context, it is naturally desirable to gain from using beamforming for both PHY CCH and data channels. This is possible due to the in-resource position of the CCH, which allows using beamforming for both the CCH and the corresponding downlink data transmission in case of single-stream transmission. The former is a consequence of being able to use the same set of dedicated reference symbols for channel estimation (and coherent demodulation) for those PHY channels. This is a clear advantage over LTE, where the PDCCH is transmitted with open loop transmit diversity mode, due to the time-wise disjoint position of the CCH (PDCCH) and data (physical downlink shared channel — PDSCH). Common reference symbols (CRS) are used for both PDCCH and PDSCH transmissions [9]. In LTE-Advanced (LTE-A) there is support for dedicated reference symbols for the PDSCH demodulation when using Transmission Mode 9, while the PDCCH is still relying on common reference symbols. Additionally, with LTE-A, there is partial support for beamforming on the CCH through the enhanced PDCCH (E-PDCCH), but the initial access configuration would still need to be addressed through the PDCCH. Moreover, resource configuration for the E-PDCCH happens via radio resource management (RRC) signaling to the UE.

Furthermore, inter-cell interference is also expected to be a challenge for the 5G era, calling for both the possibility to use efficient network-based inter-cell interference coordination (ICIC) techniques, as well as receiver-based interference cancellation/suppression schemes. Since the in-resource CCH signaling for the proposed frame structure follows the data allocations, it allows efficient time-frequency domain ICIC for both the CCH and data transmission in the case of synchronized base stations. As an example, if a cell mutes a certain set of its time-frequency domain resources, users scheduled on that set of time-frequency resources in neighboring cells will experience improved signal-to-interference-plus-noise ratio (SINR) for both the CCH and data reception. The same flexibility for ICIC is not possible for LTE due to the strict time division of PDCCH and data in each TTI, where the PDCCH transmission is distributed over the full cell bandwidth [10, 14].

Performance Analysis

Physical Layer Numerology

The PHY design shall naturally be constructed to support the proposed frame structure, offering the necessary symbol space that fits with the requirements for the minimum time-frequency allocation. The current LTE PHY design, with OFDMA waveform and 14 OFDMA symbols per 1 ms TTI [10], does not fit the desire to be able to schedule users with a minimum TTI size of 0.2–0.25 ms. Table 1 summarizes the assumed 5G numerology for further assessment of the proposed frame structure, assuming the traditional cyclic prefix (CP) OFDMA [8], although other waveforms are naturally also considered for a future 5G design. We consider options for CP duration on the order of ~2 and ~4 μs, respectively. The shorter CP of ~2 μs could be sufficient for dense urban macrocell deployments, given the typical values of excess delay spread in such environments (e.g., 1.9 μs for the ITU Urban Macrocell channel model [15]). The longer CP of ~4 μs is closer to the LTE setting, allowing more diverse deployments. The assumed 5G subcarrier spacing corresponds to the LTE subcarrier spacing multiplied by a factor of 16/15 and 32/15, respectively. Hence, the corresponding sample rate can then be synthesized with the same common clock for both LTE and 5G, which is advantageous from an implementation point of view. The larger 5G subcarrier spacing (compared to LTE) offers increased robustness to phase noise. Notice from Table 1 that the symbol capacity in terms of available resource elements (i.e., subcarrier symbols) per 1 ms and 20 MHz carrier bandwidth are identical for 5G and LTE, and hence offers a fair comparison. This is a result of also assuming 90 percent bandwidth efficiency for 5G (as is the case for LTE), meaning that the effective transmission bandwidth is 18 MHz for a 20 MHz carrier configuration. Note that the 5G case with ~2 μs CP results in a lower fast Fourier transform (FFT) size, which is of particular importance for higher carrier bandwidths of, for example, 40 or 100 MHz. As the research on waveform selection and PHY numerology is ongoing, the settings in Table 1 should only be considered as an example used in this study for further assessment of the proposed frame structure.

	5G assumptions		LTE specifications
Subcarrier spacing	32 kHz	16 kHz	15 kHz
Number of subcarriers for 20 MHz bandwidth	560	1120	1200
FFT size for 20 MHz bandwidth	1024	2048	2048
Cyclic prefix	2.0833 μs	4.1667 μs	5.21 μs[1] 4.69 μs
Cyclic prefix overhead	6.25%	6.25%	6.67%
Symbol time (including cyclic prefix)	33.33 μs	66.66 μs	71.87 μs[2] 71.35 μs
TTI size	0.2 ms	0.2 ms	1 ms
OFDM symbols per TTI	6	3	14
Resource elements per 1 ms	16,800	16,800	16,800

[1] The CP equals 5.21 μs for the 1st and 8th symbols, while it equals 4.69 μs for other symbols.
[2] The symbol time equals 71.87 μs for the 1st and 8th symbols, and 71.35 μs for other symbols.

Table 1. 5G PHY numerology (examples only) and corresponding assumptions for LTE (20 MHz carrier).

Control Channel Overhead

As the scheduling grant control channel (CCH) for the proposed 5G frame structure will essentially carry the same information as the LTE scheduling grants on the PDCCH, we

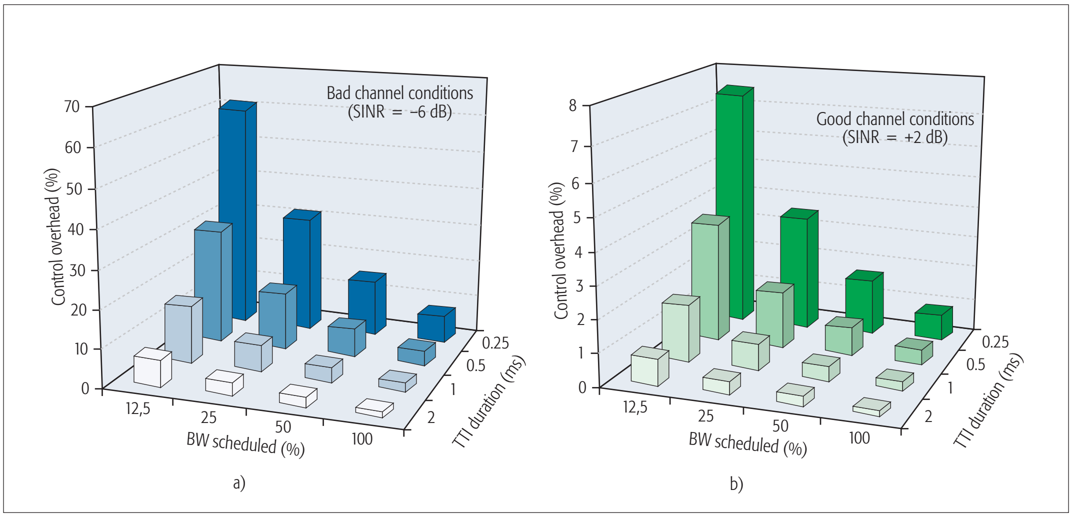

Figure 4. Control channel overhead for different time-frequency scheduling allocations and terminal experienced SINR conditions; a) CCH overhead for bad channel conditions; b) CCH overhead for good channel conditions.

	5G proposal	LTE
TTI size	Variable TTI size. Adjustable per user and per scheduling instant in steps of 0.20–0.25 ms.	Fixed 1 ms TTI size.
RTT	Below 1 ms when scheduling with a short TTI size of 0.20–0.25 ms.	8 ms.
PHY CCH and data channel mux	In-resource control signaling, where CCH and data channel transmissions are aligned, using the same bandwidth.	Strict time mux between PHY control and data. PHY control (PDCCH) is sent as wideband. Control channel blocking can occur.
DRX	Flexible time-frequency configuration of pattern for UE monitoring of downlink CCH scheduling grants.	Flexible time-domain-only configuration of pattern for UE monitoring of downlink control scheduling grants.
HARQ	Asynchronous HARQ for uplink and downlink with configurable number of parallel stop-and-wait channels supporting incremental redundancy.	Synchronous HARQ for uplink and asynchronous HARQ for downlink. Fixed number of parallel stop-and-wait channels supporting incremental redundancy.
UE bandwidth operation	UE can operate on a fraction of the carrier bandwidth – especially attractive for low-cost MTC devices.	UE needs to monitor the full carrier bandwidth up until Rel-12 as the PDCCH is transmitted on the full carrier bandwidth (MTC enhancements coming in Rel-13).
ICIC	Full support for dynamic time-frequency domain ICIC, offering protection for CCH and data channels due to the in-resource control signaling design.	Only time-domain ICIC for PHY control, and time-frequency domain ICIC for PDSCH.
Beamforming	Support for (rank-1) beamforming for both the PHY CCH and data channel due to the in-resource control design.	Open loop transmit diversity for PHY control and beamforming support for PDSCH. E-PDCCH supports beamforming,
Carrier bandwidth	5 MHz, 10 MHz, 20 MHz, 40 MHz, 100 MHz. (potentially with additional options within this range)	1.4 MHz, 3 MHz, 5 MHz, 10 MHz, 15 MHz, 20 MHz.
Spectrum aggregation	Support for aggregation of fragmented spectrum to form one cell, as well as aggregation of cells.	Carrier aggregation, that is, aggregation of individual cells.

Table 2. Summary of proposed 5G characteristics vs. assumptions for LTE.

assume the same basic structure and air interface decoding performance. The PDCCH for a user in good SINR conditions can be sent with quadrature phase shift keying (QPSK) and coding rate 7/10 on a total of 36 resource elements (REs). Such a configuration results in an average reception block error rate (BLER) of less than 1 percent if the post-detection SINR is at or above 2 dB. On the other extreme, the PDCCH could also be sent with QPSK rate 1/11, which in turn would be able to provide the needed 1 percent BLER for users in challenging SINR conditions of down to –6 dB. This requires a total of 288 REs. Notice that for the 3GPP defined macro scenarios, less than 1 percent of the users have a post-detection SINR of –6 dB (assuming standard 2 × 2 single-user open loop transmission diversity). More details on the LTE PDCCH performance can be found in [10, 14]. In addition to the CCH overhead, it is assumed that 10 percent of the REs are used for reference symbols to facilitate channel estimation and coherent demodulation.

Given these assumptions for the required number of REs for the CCH, the relative control overhead for the proposed 5G frame structure is calculated. The relative control overhead is defined as the ratio of used REs for the CCH overhead vs. the total number of used REs for data, control, and reference symbols. The results in Fig. 4 show how the CCH varies depending on the relative scheduling bandwidth and TTI duration. The example in Fig. 4 assumes a 20 MHz carrier bandwidth. It is observed that the CCH overhead scales linearly with the scheduling of users due to the in-resource CCH signaling. For instance, if a low-bandwidth user in poor channel conditions is scheduled with very low latency (short TTI size), the associated overhead of the scheduling equals 61 percent, while scheduling a user in good channel conditions with larger bandwidth and longer TTI size will result in a much lower overhead of less than 1 percent. It should be noted that the CCH overhead values experienced with the proposed 5G frame structure will be between these two extreme values, and will be a result of the traffic in the network and the applied scheduling policy. Thus, trade-offs between CCH overhead and TTI size, or equivalently RTT, are allowed. The fact that the CCH overhead is not hard limited to values of 7, 14, and 21 percent, as in LTE, presents a more flexible solution, where CCH blocking is further reduced; see results on LTE PDCCH blocking in [14] with realistic QoS-aware scheduling.

Key Characteristics vs. LTE

The key characteristics of the proposed 5G frame structure are summarized in Table 2, including comparison against the corresponding design choices for LTE Release 12 (Rel-12). Table 2 shows attractive benefits of the proposed 5G solution, which essentially map to increased flexibility for efficient multiplexing of users with extremely diverse service requirements on the same air interface. Among other characteristics, the proposed design offers shorter RTT when needed, the flexibility to optimize for high throughput at the expense of latency, as well as efficient MTC support for each link.

Conclusion

A flexible 5G FDD frame structure is presented for multiplexing users with highly diverse service requirements and radio conditions. This allows us to optimize the resource allocation on a per link basis. The concept is based on in-resource physical layer control signaling that follows the corresponding data transmission for each individual user. The proposed design offers a short air interface round-trip time when needed, the flexibility to optimize for high throughput at the expense of latency, as well as efficient machine-type communication support. Given these merits, it is suggested to continue the work on such a frame structure. As an example, it remains to be studied how to arrange downlink common channels like system broadcast information, as well as how to most efficiently facilitate multiplexing of uplink control information including HARQ feedback, channel state information, and so on. Similarly, adaptation of the frame structure to TDD cases is another topic of interest.

Acknowledgments

This work has been performed in the framework of the Horizon 2020 project FANTASTIC-5G (ICT-671660), receiving funds from the European Union. The authors would like to acknowledge the contributions of their colleagues in the project, although the views expressed in this contribution are those of the authors and do not necessarily represent the project.

References

[1] A. Osseiran *et al.*, "Scenarios for 5G Mobile and Wireless Communications: The Vision of the METIS Project," *IEEE Commun. Mag.*, vol. 52, no. 5, May 2014, pp. 26–35.

[2] G. Wunder *et al*., "5GNOW: Non-Orthogonal Asynchronous Waveforms for Future Mobile Applications," *IEEE Commun. Mag*., vol. 52, no. 2, Feb. 2014, pp. 97–105.

[3] C. Lin *et al.*, "Towards Green and Soft: A 5G Perspective," *IEEE Commun. Mag.*, Feb. 2014.

[4] P. Mogensen *et al*., "5G Small Cells Optimized Radio Design," *Proc. IEEE GLOBECOM 2013, Int'l. Wksp. Emerging Techniques for LTE-Advanced and Beyond 4G*, Dec. 2013, pp. 111–16.

[5] N. Bhushan *et al*., "Network Densification: The Dominant Theme for Wireless Evolution into 5G," *IEEE Commun. Mag*., vol. 52, no. 2, Feb. 2014, pp. 82–89.

[6] ITU-R SGs, IMT Vision, "Framework and Overall Objectives of the Future Development of IMT for 2020 and Beyond," Feb. 2015.

[7] E. Dahlman *et al*., "5G Wireless Access: Requirements and Realization," *IEEE Commun. Mag*., Commun. Standards Supp., vol. 52, no. 12, Dec. 2014, pp. 42–47.

[8] J. Vihril *et al.*, "On the Waveforms for 5G Mobile Broadband Communications," *Proc. IEEE VTC-Spring*, May 2015, pp. 1–5.

[9] A. Larmo *et al*., "The LTE Link-Layer Design," *IEEE Commun. Mag*., vol. 47, no. 4, Apr. 2009, pp. 52–59.

[10] K. I. Pedersen *et al*., "An Overview of Downlink Radio Resource Management for UTRAN Long-Term Evolution," *IEEE Commun. Mag.*, vol. 47, no. 7, July 2009, pp. 86–93.

[11] B. Soret *et al*., "Fundamental Trade-offs among Reliability, Latency and Throughput in Cellular Networks," *IEEE Proc. GLOBECOM Wksps*., Dec. 2014, pp. 1391–96.

[12] K. I. Pedersen *et al*., "Carrier Aggregation for LTE-Advanced: Functionality and Performance Aspects," *IEEE Commun. Mag*., vol. 49, no. 6, June 2011, pp. 89–95.

[13] G. Boudreau *et al.*, "Interference Coordination and Cancellation for 4G Networks," *IEEE Commun. Mag*., vol. 47, no. 4, Apr. 2009, pp. 74–81.

[14] D. Laselva *et al.*, "On the Impact of Realistic Control Channel Constraints on QoS Provisioning in UTRAN LTE," *Proc. IEEE VTC-Fall*, Sept. 2009, pp. 1–5.

[15] ITU-R, "Guidelines for Evaluation of Radio Interface Technologies for IMT-Advanced," ITU-R M.2135-1, Dec. 2009.

Biographies

Klaus I. Pedersen (klaus.pedersen@nokia.com) received his M.Sc. degree in electrical engineering and Ph.D. degree from Aalborg University, Denmark, in 1996 and 2000, respectively. He is currently a senior wireless network specialist at Nokia Networks, Aalborg, and a part-time professor at Aalborg University in the Wireless Communications Network (WCN) section. He is the author/coauthor of more than 100 peer-reviewed publications on a wide range of topics, as well as an inventor on several patents. His current work is related to 5G air interface design, including radio resource management aspects, and the continuation of LTE and its future development. Is currently part of the EU funded research project FANTASTIC-5G, which focuses on development of a new multi-service 5G air interface for below-6-GHz operation.

Gilberto Berardinelli (gb@es.aau.com) received his first and second level degrees in telecommunication engineering, cum laude, from the University of L'Aquila, Italy, in 2003 and 2005, respectively, and his Ph.D. degree from Aalborg University in 2010. He also received a second level Master's degree in techniques and economics of telecommunications in 2006 from the University of Padova, Italy. In 2006 he worked with the Radio Frequency Engineering Department in Vodafone NV, Padova, Italy. He is currently an associate professor at the Wireless Communication Networks section of Aalborg University, also working in tight cooperation with Nokia Networks. His research interests are mostly focused on physical layer and radio resource management design for 5G systems, thus including multiple access schemes, waveforms, multi-antenna systems, as well as on software defined radio prototyping. He is an author or co-author of more than 60 international publications, including conference proceedings, journal contributions, and book chapters. He is currently part of the EU funded research project FANTASTIC-5G.

Frank Frederiksen (frank.frederiksen@nokia.com) received his M.Sc. degree in electrical engineering from Aalborg University in 1994. He is currently working at Nokia Networks as a wireless networks specialist. During his employment, he has worked on concept creation and research within various radio access technologies, ranging from HSPA, LTE, LTE-A, and 5G, and the main focus has been on the physical layer and lower medium access control operation. He is an author and co-author of several patents.

Preben Mogensen (preben.mogensen@nokia.com) received his M.Sc. and Ph.D. degrees from Aalborg University in 1988 and 1996, respectively. He became a full professor at Aalborg University in 1999, and currently heads the Wireless Communication Networks section in the Department of Electronic Systems. Since 1995, he has also worked part time at Nokia, and was nominated an NSN fellow in 2009. He has supervised more than 35 successfully finalized Ph.D. candidates and has co-authored more than 300 scientific publications His current research is on 5G and IoT.

Agnieszka Szufarska (agnieszka.szufarska@nokia.com) received her M.Sc.E.E. degree in 2004 from the Technical University of Gdansk, Poland. She is currently with Nokia (in Wroclaw, Poland), where she leads a research team. Her background is in radio concept creation and system-level simulations for LTE and 5G. Her research interests include spectrum coexistence of different systems and interference mitigation for heterogeneous deployments with the mixture of small and wide area cells.

It remains to be studied how to arrange downlink common channels like system broadcast information, as well as how to most efficiently facilitate multiplexing of uplink control information such as HARQ feedback, channel state information. Similarly, adaptation of the frame structure to TDD cases is another topic of interest.

Digital Predistortion for Mitigating Spurious Emissions in Spectrally Agile Radios

Mahmoud Abdelaziz, Zhu Fu, Lauri Anttila, Alexander M. Wyglinski, and Mikko Valkama

Recent advances in DPD for non-contiguous transmission scenarios are discussed, with a focus on mitigating the spurious emissions in the concrete example case of non-contiguous dual-carrier transmission. The techniques are compared to more traditional DPD approaches in terms of their computational and hardware complexities, as well as linearization performance.

Abstract

Spectrally non-contiguous transmissions pose serious transceiver design challenges due to the nonlinear PA. When two or more non-contiguous carriers with close proximity are amplified by the same PA, spurious emissions inside or in the vicinity of the transmitter RF band are created. These spurious emissions may violate emission limits or otherwise compromise network coverage and reliability. Lowering the transmit power is a straightforward remedy, but it will reduce throughput, coverage, and power efficiency of the device. To improve linearity without sacrificing performance, several DPD techniques have recently been proposed to target the spurious emissions explicitly. These techniques are designed to minimize the computational and hardware complexity of DPD, thus making them better suited for mobile terminals and other low-cost devices. In this article, these recent advances in DPD for non-contiguous transmission scenarios are discussed, with a focus on mitigating the spurious emissions in the concrete example case of non-contiguous dual-carrier transmission. The techniques are compared to more traditional DPD approaches in terms of their computational and hardware complexities, as well as linearization performance.

Introduction

Third Generation Partnership Project (3GPP) Long Term Evolution-Advanced (LTE-A) and cognitive radio (CR) type developments are recent examples of wireless communication systems that seek to utilize spectrally non-contiguous carriers in order to increase data rates and spectral flexibility. In LTE-A, carrier aggregation (CA) enables bandwidth expansion up to 100 MHz by aggregating five 20 MHz LTE carriers [1]. In the TV white space (TVWS) standard, IEEE 802.22, known as the first CR standard, the aggregation of several carriers is called channel bonding [2]. In this article, we refer to all such transmission schemes simply as CA, and use the attributes contiguous and non-contiguous to separate spectrally contiguous and non-contiguous allocations.

One of the most challenging engineering concerns in the RF part of a transmitter is the power amplifier (PA), particularly due to its nonlinearity. Multicarrier transmissions exhibit very large peak-to-average power ratios (PAPRs), and when they are combined with the nonlinear behavior of the PA, can generate severe unwanted spectral emissions in the adjacent channels or in more distant portions of the spectrum. The levels of these emissions are mandated and enforced by federal regulators, such as the U.S. Federal Communications Commission (FCC; see, e.g., [3]). Recently, it has been demonstrated in LTE-A mobile transmitters with non-contiguous CA or multicluster type transmissions that PA nonlinearities lead to spurious emissions that can violate the given spectrum emission limits [1, 4], or lead to receiver desensitization in frequency-division duplexing (FDD) systems [4]. Within the CR context, the spurious emissions can also interfere with primary user (PU) transmissions or other secondary users (SUs).

One obvious solution to decrease nonlinear distortion is to back off the transmit power. In 3GPP terminology, this is known as maximum power reduction (MPR), and MPR values up to 16 dB are allowed in some CA use cases for mobile terminals [1]. However, this approach yields a significantly lower PA efficiency as well as a substantial reduction in the uplink coverage and throughput. For example, if we assume the well-known COST Hata propagation model for urban areas,[1] along with a base station antenna height of 50 m, a 10 dB reduction in transmit power will result in halving the coverage.[2] Thus, there is a clear need for alternative linearization solutions that do not possess such drastic adverse effects. This article focuses on one linearization solution in particular: digital predistortion (DPD).

Most of the current DPD literature is focused on base stations and other infrastructure nodes where DPD is the de facto solution for linearization. However, only a few published works consider DPD for mobile terminals and other low-cost devices. One of the main reasons is that usually the cost of implementing a DPD solution, in terms of hardware and needed computations, is relatively high. Furthermore, with traditional contiguous transmit spectrum, linearization of mobile transmitters is usually not needed since transmit power levels are generally smaller and emission limits looser relative to base stations or infrastructure nodes. However, this is about

[1] *Digital Mobile Radio Towards Future Generation Systems*, COST 231 Final Report; http://146.193.65.101/cost231/

[2] Coverage is defined here as the distance over which a certain received signal-to-noise ratio is obtained.

Mahmoud Abdelaziz, Lauri Anttila, and Mikko Valkama are with Tampere University of Technology; Alexander M. Wyglinski is with the Wireless Innovation laboratory at Worcester Polytechnic Institute; Zhu Fu was with the Wireless Innovation laboratory at Worcester Polytechnic Institute, and is now with Qualcomm Inc.

to change as non-contiguous CA waveforms are being deployed across a variety of systems.

This article addresses the nonlinearities of terminals that transmit non-contiguous CA waveforms, the resulting spurious emissions, as well as their suppression with DPD. The primary focus is on LTE-A mobile terminals, while we also acknowledge that another potential application area is in TVWS devices. The objectives of this article are generally three-fold:

- To highlight the problem of PA-induced spurious emissions with non-contiguous transmissions, and the potential problems it may create in different systems.
- To demonstrate that back-off is not the only feasible solution to decreasing unwanted emissions, but DPD is a potential solution that does not sacrifice the power efficiency, coverage, or throughput of the mobile user.
- To increase the awareness of the readership of this magazine with respect to recent advances in DPD for non-contiguous transmission schemes. We consider this to be one key element in enhancing the flexibility of radio spectrum usage in emerging radio communication systems.

The rest of this article is organized as follows. The following section gives an overview of the linearity challenges with spectrally non-contiguous transmissions, and presents the relevant emission regulations of 3GPP LTE-A. Recent advances in DPD techniques for spurious emission suppression are then discussed. We then discuss the implementation complexity of the DPD methods. Simulation and RF measurement results are then provided. Finally, several concluding remarks are made.

Linearity Challenges in Employing Non-Contiguous Transmissions

3GPP LTE-Advanced and Unwanted Emissions

In order to meet the requirements of the International Mobile Telecommunications-Advanced (IMT-A) specifications, 3GPP began introducing CA from Release 10 onwards. Release 10 is also the first release where the technology is called LTE-A. For the uplink, Release 10 introduced contiguous CA, allowing only contiguous user equipment (UE) resource allocation per individual carrier. Uplink specifications were again updated with Release 11, which introduced so-called multicluster transmission, which means that non-contiguous resource allocation can be performed within a single uplink (UL) carrier. Release 12 yielded two additional major updates, as non-contiguous intraband CA and interband CA were introduced, where the LTE-A UE can be allocated UL resources across aggregated spectrum consisting of two or more component carriers (CCs). The CCs may possess a bandwidth of any single-carrier bandwidth defined within the LTE specifications, and a maximum of five CCs can be aggregated, with a theoretical maximum aggregated bandwidth equal to 100 MHz.

In general, transmitter unwanted emissions can be divided into three parts:

- Emissions within the carrier bandwidth
- Out-of-band (OOB) emissions
- Spurious emissions.

Out-of-band emissions are unwanted emissions immediately outside the assigned channel bandwidth resulting from the modulated waveform characteristics and nonlinearity of the transmitter, but excluding spurious emissions, and are specified in terms of the spectrum emission mask and adjacent channel leakage ratio (ACLR). Spurious emissions, in turn, are unwanted emissions that are caused by the transmitter nonlinear effects such as harmonic and intermodulation products, parasitic emissions, and frequency conversion products, with OOB emissions excluded. The exact requirements for spectral emissions are specified in [1]. In this article, we are mainly interested in the spurious emission limits. For carrier frequencies above 1 GHz, where most LTE-A bands are located, the spurious emission limit is –30 dBm over a 1 MHz measurement bandwidth. This is a general guideline defined by International Telecommunication Union — Radiocommunication Standardization Sector (ITU-R) [3], which means that it applies to all land mobile services, not just LTE-A.

For carrier frequencies above 1 GHz, where most LTE-A bands are located, the spurious emission limit is –30 dBm over a 1 MHz measurement bandwidth. This is a general guideline defined by ITU-R [3], which means that it applies to all land mobile services, not just LTE-A.

Linearity Challenges with Non-Contiguous Dual-carrier Transmissions

When a non-contiguous dual-carrier signal is applied to the PA input, the PA nonlinearity leads to intermodulation and cross-modulation products affecting many different frequencies, as shown in Fig. 1. Assuming a CC separation of Δf, in addition to the spectral regrowth around the main carriers, intermodulation distortion appears at integer multiples of Δf away from the main CCs. The third-order intermodulation (IM3) sub-bands are centered at Δf from the main CCs, the fifth-order intermodulation (IM5) sub-bands are located at $2\Delta f$ from the main CCs, and so on. These intermodulation distortion (IMD) components may affect the system in two different ways, depending on the CC locations and their separation, as well as on the bandwidth of the transmit filter.

First, the intermodulation spurs created by the nonlinear PA will likely be in the spurious domain, and therefore should obey the more strict spurious emission limits, as opposed to the spectral emission limits that are defined in the vicinity of the carriers [1, 3]. In the context of interband CA, where the CCs are located at different RF bands, the spurs will likely be filtered out by the transmitter RF filter. However, in the case of intraband CA, where the CCs are inside the same RF band, some of the spurious IM sub-bands may be located in-band, and therefore will not be mitigated by the transmit RF filter. Narrower bandwidth allocations per CC are more problematic, since the energy of the spurs is more concentrated, therefore more easily violating the spur limit [4]. To avoid spur limit violations, the mobile device should therefore back off its output power to operate within the linear range of the PA. However, this backoff will lead to a reduction in the UL coverage range, as well as inefficient operation of the PA. In the next section, DPD processing is introduced as a solution to allow a mobile device to meet the spurious emission limit while considerably reducing the amount of required backoff.

Second, in FDD terminals employing

In order to meet linearity requirements without sacrificing performance, digital pre-distortion has been proposed and recognized as a promising PA linearization technique due to its accuracy and cost effectiveness, particularly for high-power high-end base station devices.

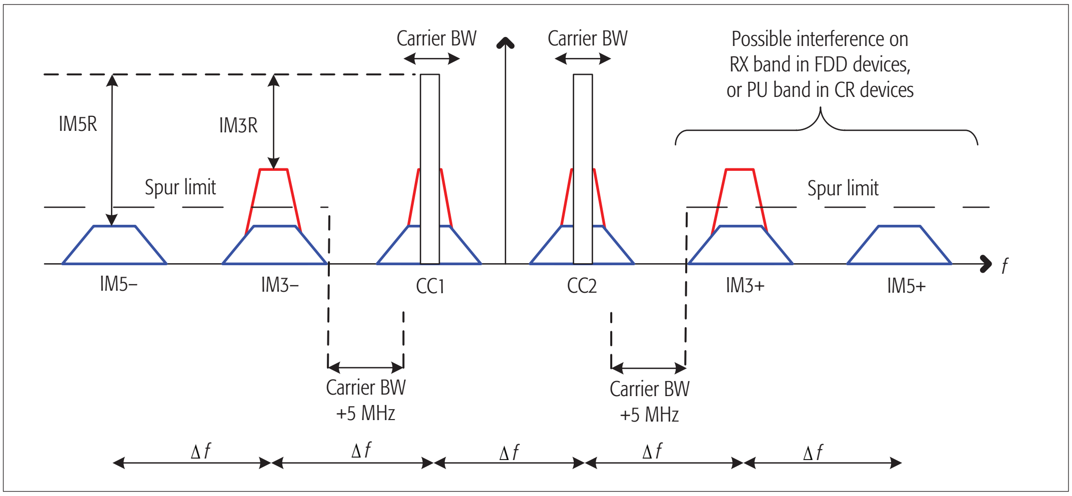

Figure 1. Example transmit spectrum with non-contiguous carrier aggregation under a nonlinear PA. In this illustration, nonlinear distortion products up to order five are shown, including nonlinear distortion around the main carriers as well as at the IM3 and IM5 sub-bands.

non-contiguous intraband transmissions, the spurious components may potentially overlap with the receive band of the device, causing *desensitization of its own receiver* [4]. Carrier aggregation reduces the duplexing distance between TXs and RXs, thereby reaching sufficient isolation between TX and RX using RF duplexers is increasingly difficult.[3] Receiver desensitization can also be alleviated with digital predistortion, as described in the following.

Low-Complexity Sub-Band Digital Predistortion for Non-Contiguous Transmission Schemes

In order to meet linearity requirements without sacrificing performance, digital pre-distortion has been proposed and recognized as a promising PA linearization technique due to its accuracy and cost effectiveness, particularly for high-power high-end base station devices. There have been substantial research efforts over the past 20 years with respect to developing efficient and elaborate DPD techniques for various single-band transmission schemes where linearization for the whole transmit band is essentially pursued. These conventional DPD approaches take as their inputs the full composite transmit band, and we thus refer to these DPD approaches as *full-band DPD*. The tutorial paper in [5] provides an overview of these techniques, while this article only addresses these techniques from the DPD complexity perspective. There have also been a handful of recent works on efficient linearization techniques for multi-band transmitters that employ only a single PA, where it has been assumed that the CCs are separated by enough distance such that the spurious emissions are filtered out by the transmit filter. This approach, referred to as *concurrent linearization*, is treated in a recent overview paper [6]. Such techniques are also out of the scope of this article.

The scope of this article is to introduce efficient low complexity DPD techniques for mitigating the spurious emissions of non-contiguous transmissions, while not concentrating specifically on the linearization of the CCs. This approach is motivated by two main factors. First, the emission limits in the spurious region are stricter than in the spectral regrowth region around the CCs, and are therefore more easily violated in the context of intraband non-contiguous CA. Second, by concentrating linearization efforts on certain spurious emissions only, the processing and hardware complexity of a device can be significantly reduced, thus facilitating the implementation of DPD linearization in lower-cost devices and potentially even mobile terminals. The main target application for these techniques is thus the mobile terminal transmitter, where computational and hardware complexity are critical, although there are no technical limitations to applying these techniques in base stations or infrastructure nodes as well.

There have been some recent studies in the literature that consider the mitigation of the spurious emissions explicitly. In [7, 8], such processing was added to complement a concurrent linearization system, while assuming frequency-flat PA responses within the processed spurious sub-bands. In [7], the DPD parameter estimation was done offline, based on extracting the quasi-memoryless PA parameters using a large-signal network analyzer (LSNA). In [8], a memoryless least-squares fit between the observed IMD at the considered sub-band and certain basis functions was performed. The estimated and regenerated IMD was then injected to the PA input, oppositely phased, such that it cancels the IMD at the PA output. In [9], on the other hand, the IM3 spurious sub-bands were specifically targeted, and CC linearization was not pursued. Furthermore, the DPD parameter identification was based on closed-loop feedback with a decorrelation-based learning rule. A block-adaptive version of this DPD solution was developed and tested in a real-time field programmable gate array (FPGA) implementation in [10], demonstrating a stable response with good linearization performance. All of the above works assume a special DPD structure for the spurious emissions, which we here refer to as

[3] This problem has been acknowledged by 3GPP, for example, in R4-123797, "*UE reference sensitivity requirements with two UL carriers*," Ericsson and ST-Ericsson; http://www.3gpp.org

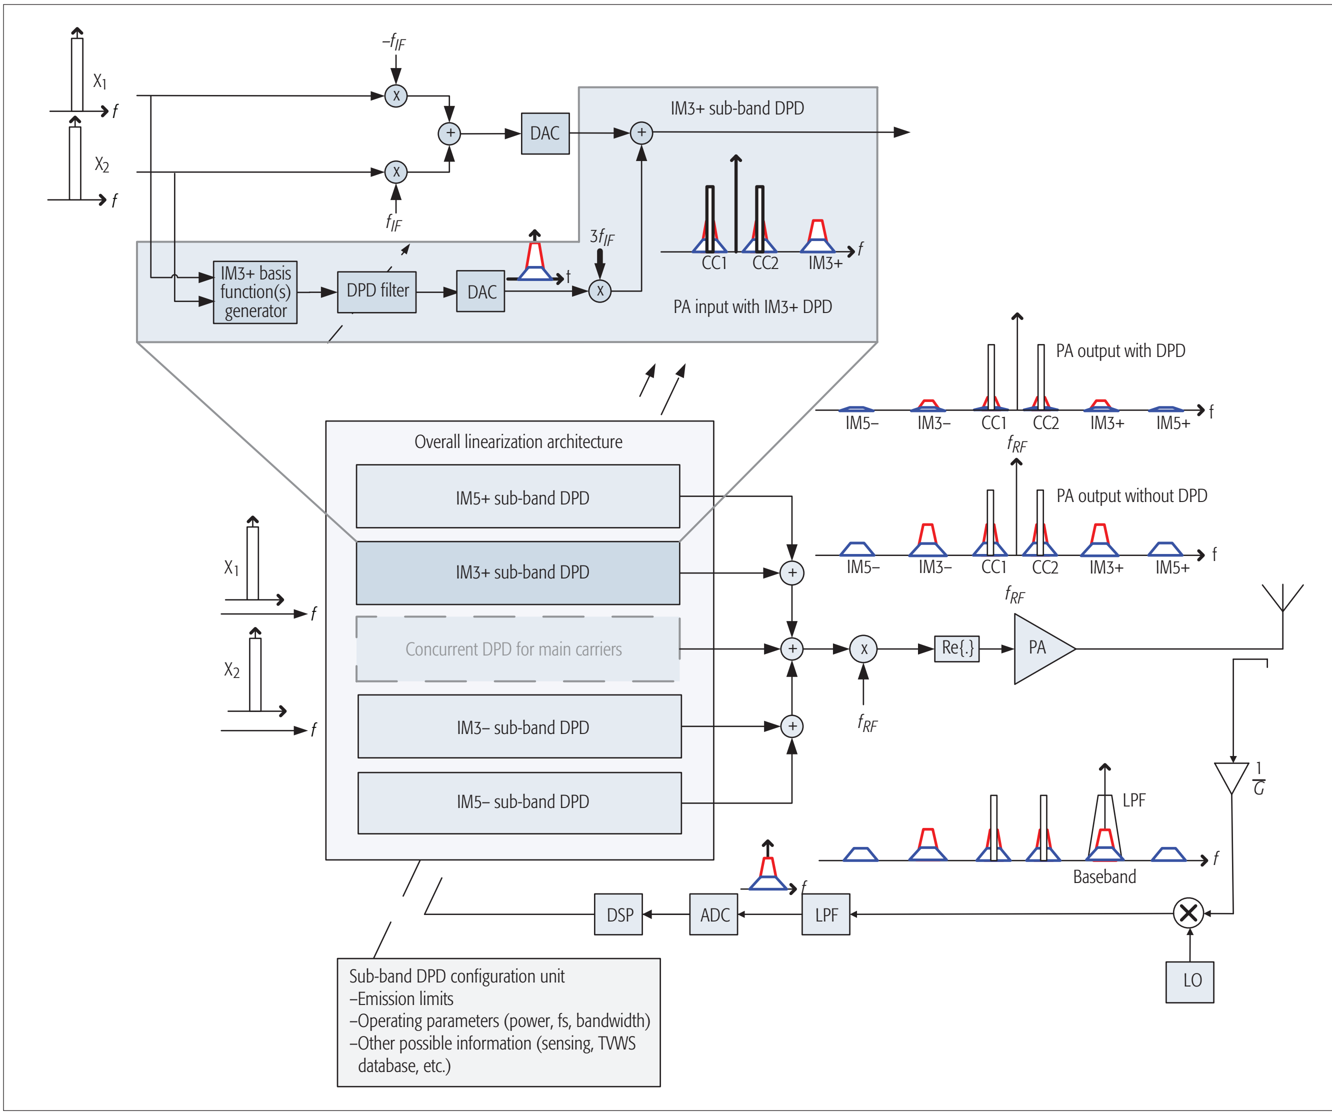

Figure 2. Top: a detailed block diagram showing the spur-injection-based sub-band DPD operating principle for the positive IM3 sub-band; bottom: complete linearization architecture with multiple sub-band DPD stages, which can be flexibly activated based on the prevailing emission limits and spectrum access scenario. The architecture can also be complemented with the concurrent DPD block, shown in gray, if main carrier linearization is pursued.

sub-band DPD. In such processing, the spurious components that we wish to remove from the PA output are modeled explicitly at the digital baseband, and injected into the PA input with proper amplitude and phase such that at the PA output the spurs are suppressed. Such spur-injection-based sub-band DPD is depicted in Fig. 2, which illustrates in more detail the mitigation of the positive IM3 sub-band. The overall linearization architecture with multiple sub-band DPD blocks is illustrated in the same figure. This architecture is also called frequency-selective DPD in [7]. For more technical details, refer to [7–10].

In addition to the sub-band DPD concept, several full-band DPD solutions toward mobile terminals are worth mentioning. These solutions have the capability to tune their linearization efforts to certain sub-bands/frequencies based on the current spectrum allocations and emission limits, to reduce the complexity of the DPD, or both [11–13]. These solutions, however, entail a similar complexity disadvantage compared to sub-band DPD as does any full-band DPD technique. The complexity perspectives are reviewed in more detail in the next section.

Computational and Hardware Complexity Perspectives

General Aspects

The relative computational requirements of the sub-band DPD approach when compared to traditional full-band DPD processing are generally reduced, especially when the LTE-A CCs are significantly separated. In this section, we provide a more thorough analysis of the computing and hardware complexity in these two DPD architectures. As an illustrative example, a full-band DPD with fifth order nonlinearity needs to run at a sample rate that is five times the composite dual-carrier signal bandwidth, which quickly becomes impossible when the carrier separation increases. With the sub-band technique, on the other hand, the minimum sample rate is less than or equal to five times the bandwidth of the wider

	Running complexity			Performance	
	Coeffs	Fs [MSPS]	GFLOPS	EVM [%]	Positive IM3R [dBc]
No DPD	N/A	N/A	N/A	2.23	30.94
Full-band DPD	24	217	43.6	0.25	46.65
Sub-band DPD	6	5	0.35	2.29	59.30

Table 1. Quantitative running complexity and linearization performance comparison of full-band vs. sub-band DPD. Two 1 MHz component carriers separated by 30 MHz. SC-FDMA component carrier waveforms with QPSK data modulation.

CC. To give a numerical example, let us assume a CA scenario with two 1 MHz carriers separated by 30 MHz. The minimum sample rate with a traditional fifth-order full-band DPD would be 155 MHz, whereas with the sub-band DPD, it is only 5 MHz. In addition, the required filter lengths in the full-band DPD are likely to be substantially longer than with sub-band DPD, since the former is predistorting a signal band which is an order of magnitude larger. The overall complexity difference therefore grows when the CC spacing is increased and/or higher order DPDs are considered.

Full-Band and Sub-Band DPD Running Complexities

We shall focus here on the running complexity of the DPD [14], which is most critical for mobile-type devices. It involves the number of computations performed per second when the DPD is operating and the device is transmitting. To quantify the computational complexity differences between the full-band and sub-band DPD architectures, we shall use the number of floating point operations (FLOPs) per sample and the required sample rate in the predistortion path. The running complexity is further divided into two parts: basis function generation, and predistortion using the basis functions [14].

The full-band DPD architecture that we use as a reference is based on the parallel Hammerstein (PH) nonlinear model. The overall running complexity for a seventh order PH model with memory order six, used as the reference solution below, is 201 FLOPs per output sample [14].

In the sub-band architecture, on the other hand, the third and fifth order basis functions read $x_2^*(n)x_1^2(n)$ and $2|x_1(n)|^2\, x_2^*(n)x_1^2(n) + 3|x_2(n)|^2\, x_2^*(n)\, x_1^2(n)$, respectively, when the positive IM3 sub-band is considered.[4] Consequently, the number of FLOPs required for basis function generation in the sub-band DPD is 22 FLOPs per sample per sub-band. In the sub-band architecture, we can assume a smaller filter length due to the significantly smaller bandwidth of the predistorted signal. Additionally, the main carriers are not predistorted. Altogether, in the filtering stage, assuming filter length of 3, 48 FLOPs per sample are required per sub-band. Thus, the total number of FLOPs required by the sub-band DPD in our example is 70 FLOPs per sample per processed sub-band.

The required sample rates in the example scenario of two 1 MHz CCs with 30 MHz carrier spacing are 217 MHz and 5 MHz for the full-band and sub-band DPDs, respectively. Based on these numbers and the above FLOP analysis, the number of FLOPs per second (FLOPS) becomes 43.617 GFLOPS for the full-band DPD, and 0.35 GFLOPS per sub-band for the sub-band DPD. These numbers are also summarized in Table 1.

System Power Efficiency Perspectives

When a DPD is adopted, a less linear but more efficient PA, which can operate near its saturation region, can generally be used. However, the overall power efficiency of the device is only improved if the extra power consumed by the DPD stage is less than the power savings due to increased PA efficiency. Here, we address this aspect from a mobile device perspective.

We consider a practical scenario where the transmit power at the output of a mobile PA is +26 dBm (i.e., 400 mW), stemming from 3GPP LTE-A requirements [1], and assume that the TX duplexer filter and connector insertion losses are 3 dB. Then practical example figures of PA power efficiency, when operating in non-linear or linear modes, are around 30 and 20 percent (or even less), respectively. This means that the power consumed by a nonlinear PA is roughly 1300 mW, while the corresponding linear PA consumes roughly 2000 mW. In other words, adopting a more power-efficient nonlinear PA saves 700 mW of power in this particular example.

On the other hand, in the CA scenario assumed in the previous subsection, the sub-band DPD requires 350 MFLOPS per sub-band. If we focus our linearization on the two IM3 sub-bands, 700 MFLOPS are required for the sub-band DPD processing. As a practical example, the state-of-the-art 28 nm DSP implementation in [15] has a processing power consumption of 23.4 mW at 230 MHz. Assuming 4 FLOPs per clock cycle, this DSP can support 920 MFLOPS, which is clearly sufficient in this example case. Thus, since the DPD already saves 700 mW of power in the PA interface, while consuming only 23.4 mW in the processing, the power budget is clearly in favor of using the sub-band DPD in this example. If the DPD processing is implemented using hardware, an even more power-efficient solution can be realized.

Flexibility and Feedback Receiver Perspectives

Another clear advantage of the sub-band architecture is that each sub-band DPD block can be switched on or off according to the induced emission levels at these sub-bands, as well as possible band- or area-specific emission limitations. The overall DPD architecture in Fig. 2 has a configuration unit that can activate/deactivate the sub-band DPD blocks according to the input coming from two different sources. The first is the anticipated/measured spurious emissions relative to the emission limits. The second source is the operating system parameters (power, bandwidth, carrier frequencies, etc.) that are used to decide whether DPD is needed in the current system configuration. In TVWS applications, additional information from spectrum sensing and/or a TVWS database can also be employed. In general, this kind of flexibility in the sub-band DPD architecture can lead to substantial power

[4] This is based on an extension of the analysis done in [9], with higher order nonlinearities considered instead of only third order nonlinearity. Here, $x_1(n)$ and $x_2(n)$ denote the digital baseband waveforms of the two CCs.

savings, as the linearization can be tailored and directed only to those frequencies where intermodulation suppression is really needed. Such flexibility is not available in the traditional full-band DPD solutions since, by design, the predistorter tries to linearize the full composite transmit band.

In addition to the complexity reduction and flexibility of the DPD main path, the feedback path complexity used for DPD parameter estimation is also greatly reduced. In order to estimate the parameters of the positive IM3 DPD, as an example, we only need to observe the positive IM3 sub-band at the PA output instead of observing the whole signal band (including the IM sub-bands), which is the case with full-band DPD. This reduction in the observation bandwidth reduces the cost, complexity, and power consumption in the feedback path, thus allowing the use of simpler instrumentation and making the approach suitable for mobile-type devices.

Simulation and RF Measurement Examples

In order to demonstrate and quantify the performance/complexity trade-offs between full-band and sub-band DPD architectures, both Matlab simulation and RF measurement results are presented in this section. We quantify the linearization performance primarily through the obtained suppression of the spurious energy at the considered sub-band, while passband error vector magnitude (EVM) values are also reported. Measurement examples with a commercial LTE-A mobile PA and a small-cell base station PA are also provided as a strong proof of concept of the sub-band DPD architecture.

Simulation Scenarios and Results

Now we illustrate two possible applications of the proposed sub-band DPD solution. The first example shows the suppression of spurious IMD components in order to meet the transmitter spurious emission limits without applying a large backoff. In the second example, the sub-band DPD approach is used to suppress spurious IMD components falling onto the RX band in an FDD transceiver, thus preventing RX desensitization.

Meeting Spurious Emission Limits: In this example, a carrier aggregated LTE-A uplink signal with quadrature phase shift keying (QPSK) subcarrier modulation is applied to a wide-band parallel Hammerstein PA model of nonlinearity order 5, which is based on measurements of a real mobile PA. The carrier separation between the two CCs is 30 MHz. A transmit filter model based on a real measured mobile duplexer is also used in the simulations. The response of the transmit filter is shown in Fig. 3a. The occupied bandwidth within each CC is 1 MHz, and the carriers are allocated such that the positive IM3 sub-band lies inside the transmit filter passband.

The sub-band DPD is tuned to predistort only the positive IM3 sub-band, since the negative one is already filtered by the transmit filter, as shown in Fig. 3a. The employed fifth-order sub-band DPD utilizes two basis functions based on an extended architecture of the one presented in [9], and each basis function is filtered by a three-tap filter.

In the full-band DPD case, implemented for reference, a seventh order parallel Hammerstein DPD based on the indirect learning architecture (ILA) in [5] is used. Altogether four filters, all with six taps, are used in the full-band DPD. Three ILA iterations are used. In the full-band DPD, an inherent 1.5 dB backoff needs to be applied to the transmit path in order to account for the slight increase in PAPR due to the predistortion, something that is not needed in the sub-band DPD. Therefore, for a fair comparison, the transmitter power levels are adjusted such that the output power after the PA is the same for both DPDs.

In Fig. 3a, the Tx power level is +18 dBm, and it can be seen that without applying any DPD processing, the transmitter clearly violates the spurious emission limit of –30 dBm/MHz. After adopting DPD processing, the spur levels are below the limit with both DPD architectures.

Table 1 shows a quantitative comparison of the complexity and performance of the full-band and sub-band DPDs. The complexity is clearly in favor of the sub-band DPD, where the required number of FLOPs per second (FLOPS) is substantially lower than in the full-band DPD, as discussed already. As for the actual linearization performance, both DPDs give good suppression for the IM3 distortion. We quantify the suppression of intermodulation power at the IM3 sub-bands through the power ratios relative to the CC wanted signal power, as shown in Fig. 1 and defined as $IM3R_{dB} = 10 \log_{10}(P_{wanted,\text{cc}}/P_{IM3})$. The IM3R at the positive IM3 sub-band is shown in Table 1 for the CC allocations in Fig. 3a. The sub-band DPD provides better performance in terms of positive IM3R compared to the full-band DPD in this particular scenario. The inband distortion (i.e., EVM) is also measured, where the full-band DPD outperforms the sub-band DPD. This is expected, since the full-band DPD linearizes the whole transmit band, including the main CCs. However, the EVM with the sub-band DPD is only around 2.3 percent, which is clearly enough for modulations at least up to 64-quadrature amplitude modulation (QAM). Additionally, the EVM degradation due to the sub-band DPD, as compared to that without DPD, is only around 0.06 percent.

Figure 3b shows the required MPR for different transmission scenarios relative to maximum transmit power of +23 dBm, assuming duplexer and other connector losses of 3 dB. Both full-band and sub-band DPDs are compared to no DPD, and with the allocation bandwidth per CC varied. A reduction of almost 8 dB in the required MPR can be achieved when the sub-band DPD is used. This directly impacts the uplink coverage of the mobile network, as shown in Fig. 3c, where the percentage of the UL coverage relative to the maximum coverage at +23 dBm is shown with and without DPD, assuming the COST Hata propagation model. As can be seen, coverage extensions on the order of several tens of percentages can be achieved by adopting the sub-band DPD processing.

When a DPD is adopted, a less linear but more efficient PA, which can operate near its saturation region, can generally be used. However, the overall power efficiency of the device is only improved if the extra power consumed by the DPD stage is less than the power savings due to increased PA efficiency.

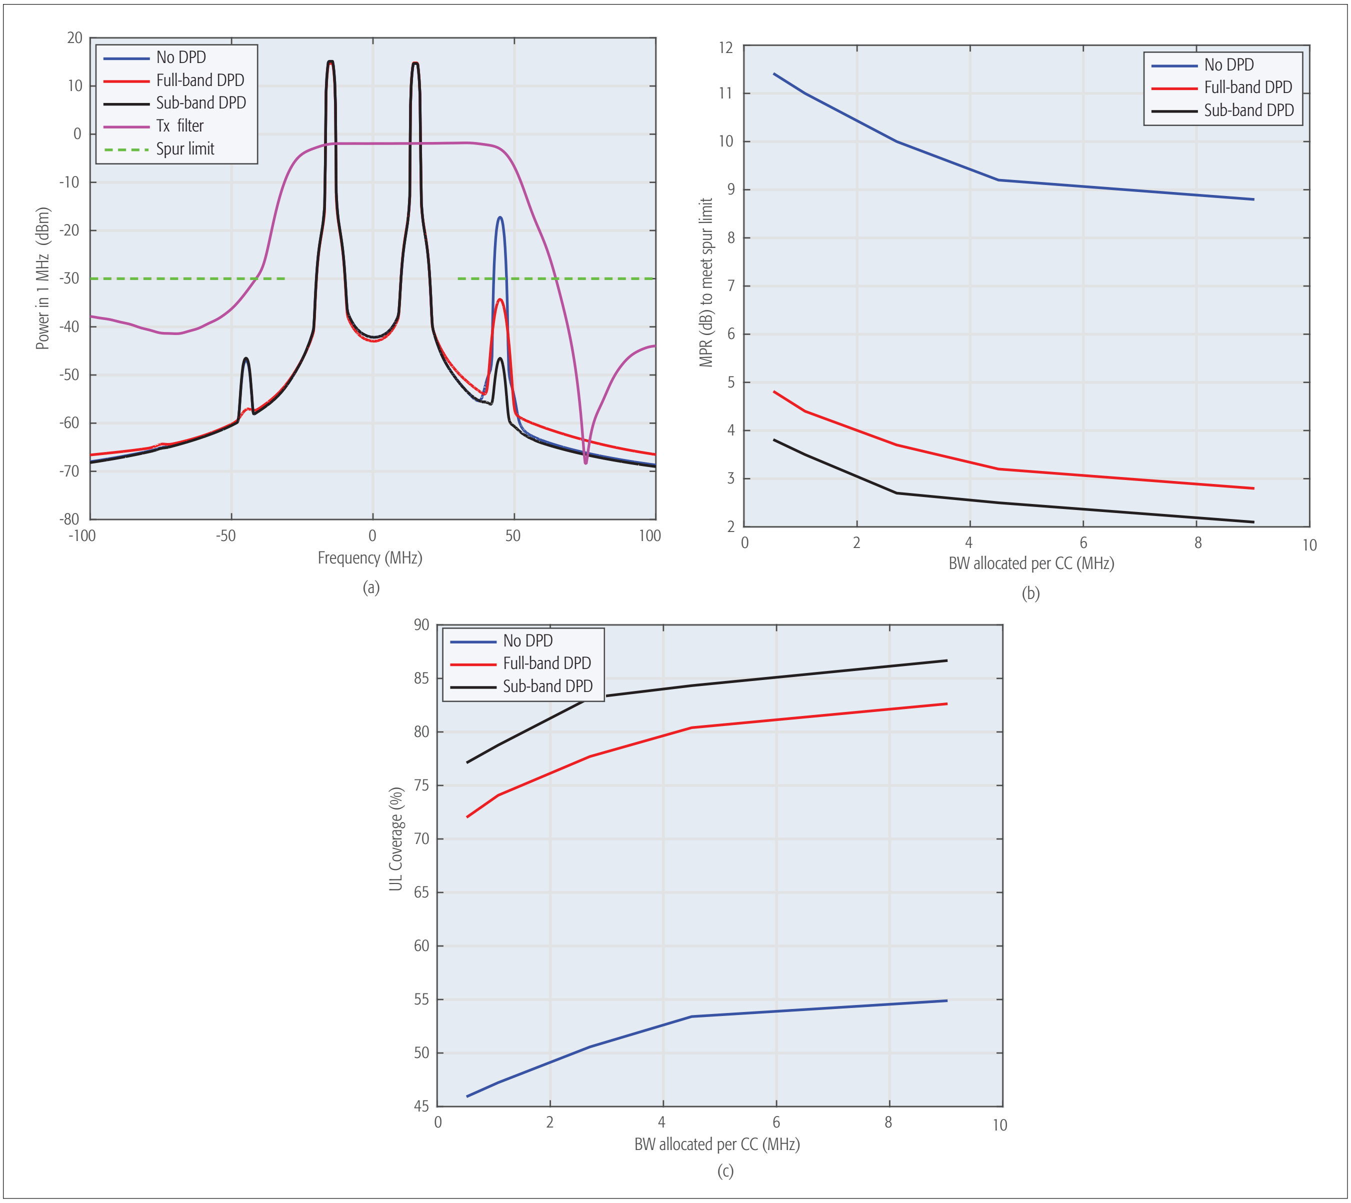

Figure 3. Dual-carrier mobile transmitter baseband equivalent power spectra, required MPRs and UL coverage. The full-band and sub-band DPDs are compared to not using DPD. Full-band DPD is seventh order Parallel Hammerstein, while the sub-band DPD builds on third and fifth order basis functions: a) two 1 MHz CCs are transmitted with Tx power of +18 dBm; b) the required MPRs to meet the TX spurious emission limits are shown, without and with DPD, and with different allocated bandwidths per CC; c) UL network coverage after applying the needed backoff (MPR) to meet the emission limits are shown, in percentage of the full coverage at +23 dBm TX power and linear PA. The COST Hata propagation model is assumed for radio propagation between the mobile and the base station, assuming base station antenna height of 50 m.

Suppressing Spurious Components at RX Band: In FDD transceivers, the TX spurious components may also overlap with the RX band and thus desensitize the RX [4]. The proposed IM3 sub-band DPD solution can be effectively applied to relax this problem. In this example, the transmit signal is again a carrier aggregated LTE-A uplink signal with QPSK subcarrier modulation. Each CC is allocated 50 LTE resource blocks (RBs, 10 MHz CC bandwidth), and the CCs are separated in frequency by 50 MHz. The PA is modeled with a fifth-order Wiener nonlinearity, that is, a linear time-invariant (LTI) filter followed by static nonlinearity.

Following typical commercial duplexer stopband responses, the duplexer filter is modeled to have frequency-selective stopband attenuation of around 50 dB at the RX band. The desired RX signal is an LTE-A downlink orthogonal frequency-division multplexing (OFDM)(A) signal with 50 RBs and QPSK subcarrier modulation. The spurious IMD at the positive IM3 sub-band due to PA nonlinearity now overlaps the RX band, thus interfering the desired RX signal. The reference thermal noise power level at RX input is –104 dBm/10 MHz, RX noise figure is 9 dB, and the reference sensitivity level is –93.5 dBm [1, 4].

Figure 4a shows the spectra of different signal components at the RX band at +20 dBm TX power, with and without using the proposed IM3 sub-band DPD. The spectra are normalized relative to the thermal noise level. It can be observed

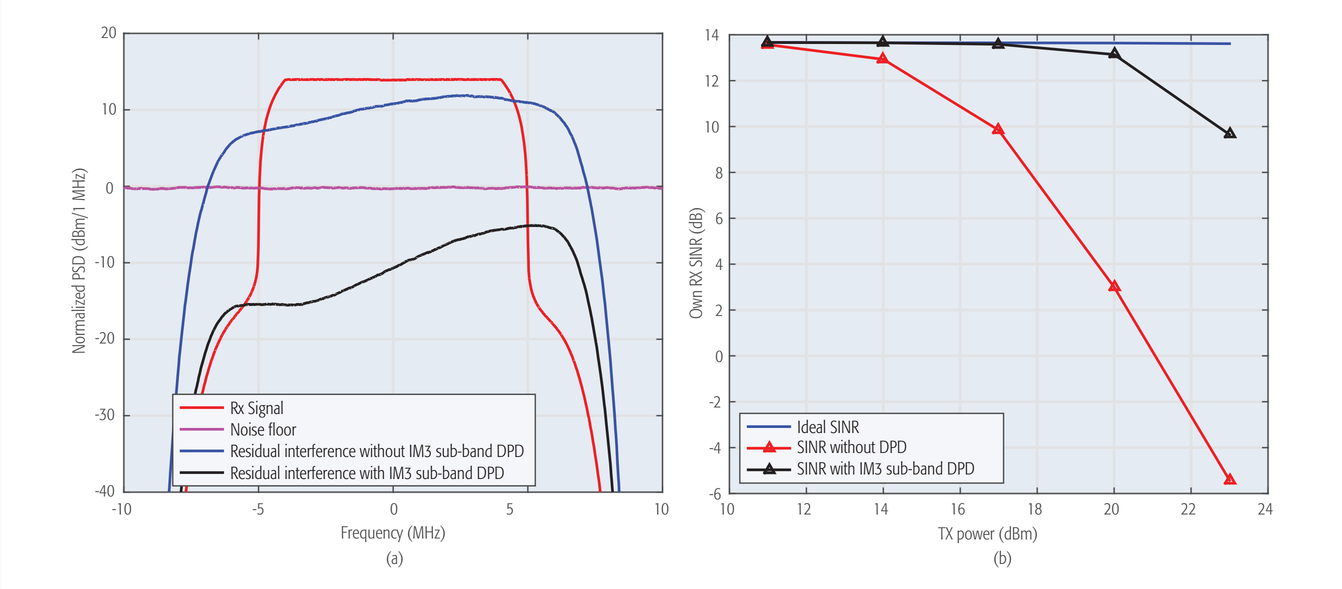

Figure 4. Using DPD to relieve RX desensitization, in the case that the TX IM3 sub-band overlaps with the RX band in an FDD transceiver. A fifth order Wiener PA model is used, with PA gain, IIP3, and 1-dB compression point equal to 20 dB, +17 dBm, and +27 dBm, respectively. The transmit signal is a CA LTE-A uplink signal with QPSK subcarrier modulation, 10 MHz allocation per CC, and 50 MHz CC spacing. A duplexer filter model with frequency selective stop band attenuation of 50 dB at the RX band is assumed. The desired RX signal is a 10 MHz LTE-A downlink OFDM(A) signal with QPSK subcarrier modulation: a) spectra of different signal components at RX band at +20 dBm TX power; b) illustration of RX signal-to-interference-plus-noise ratio (SINR) for different TX powers, with and without sub-band DPD.

that strong spurious IMD interference leaks through the duplexer filter and appears at the RX band, hence corrupting the reception. The proposed sub-band DPD is able to reduce the unwanted emissions and push the interference below the noise floor.

To further quantify the performance, the receiver signal-to-interference-plus-noise ratio (SINR) is evaluated against different TX power levels. The obtained SINR curves with different TX power levels are shown in Fig. 4b. The impact of spurious IMD interference on the receiver performance can now be seen more clearly, indicating that when TX power level increases above +14 dBm, the spurious IM3 starts to deteriorate the RX SINR. The sub-band DPD solution can enhance the SINR at the RX band by up to 15 dB, and thus substantially extend the usable TX power range.

RF Measurement Examples

The sub-band DPD solution is next tested through actual RF measurements using a commercial LTE-A mobile terminal power amplifier designed for LTE UL band 25 (1850–1915 MHz). The RF transceiver used is a National Instruments (NI) PXIe-5645R vector signal transceiver (VST). The digital baseband waveform is first generated locally on the computer, then transferred to the VST to perform I/Q modulation at 1880 MHz. The VST output signal, with –10 dBm power, is then fed to the mobile PA with output power of +20 dBm. The PA output at the considered IM sub-band is observed with the VST through a 40 dB attenuator. The observed and filtered baseband I/Q samples, together with the transmitted samples, are then used by the computer for DPD learning. The DPD learning is performed with a sequence of 50k samples after which the resulting obtained parameters are applied for measuring and quantifying the PA output. The DPD learning algorithm is based on an extension of [9], where now also fifth, seventh, and ninth order nonlinearities are considered at IM3 sub-band in addition to third order nonlinearities. Figure 5a shows the power spectral density at the PA output with and without using sub-band DPD. The signal used is a multicluster LTE-A signal with 1 MHz per cluster and 10 MHz cluster separation. The intermodulation distortion at the considered IM3 sub-band is suppressed by 20 dB with the higher-order sub-band DPD, and giving an additional 10 dB gain compared to the basic third order sub-band DPD.

For further demonstration, another measurement example using an off-the-shelf small cell base station PA[5] is presented in Fig. 5b. In this example, the LTE-A cluster bandwidths are 2 MHz each, with 8 MHz cluster separation. Here, both IM3 and IM5 spurs are separately predistorted using ninth order sub-band DPDs. The distortion at the IM3 and IM5 sub-bands is successfully suppressed by 20 dB and 12 dB, respectively. The power spectrum with simultaneous deployment of IM3 and IM5 sub-band DPDs is also shown in Fig. 5b, demonstrating sufficient spur suppression in this case as well.

Conclusions

In this article, we consider the problem of transmitter power amplifier induced spurious emissions with non-contiguous carrier aggregation, and discuss low-complexity digital predistortion for their mitigation in the mobile terminal transmitter. For non-contiguous CA transmissions, a recently proposed sub-band DPD structure was shown to yield linearization results beyond the conventional full-band DPD solution, while also

[5] The measurement setup for the base station PA example is provided by an online weblab sponsored by National Instruments and Chalmers University of Technology, at http://dpdcompetition.com/access-weblab/

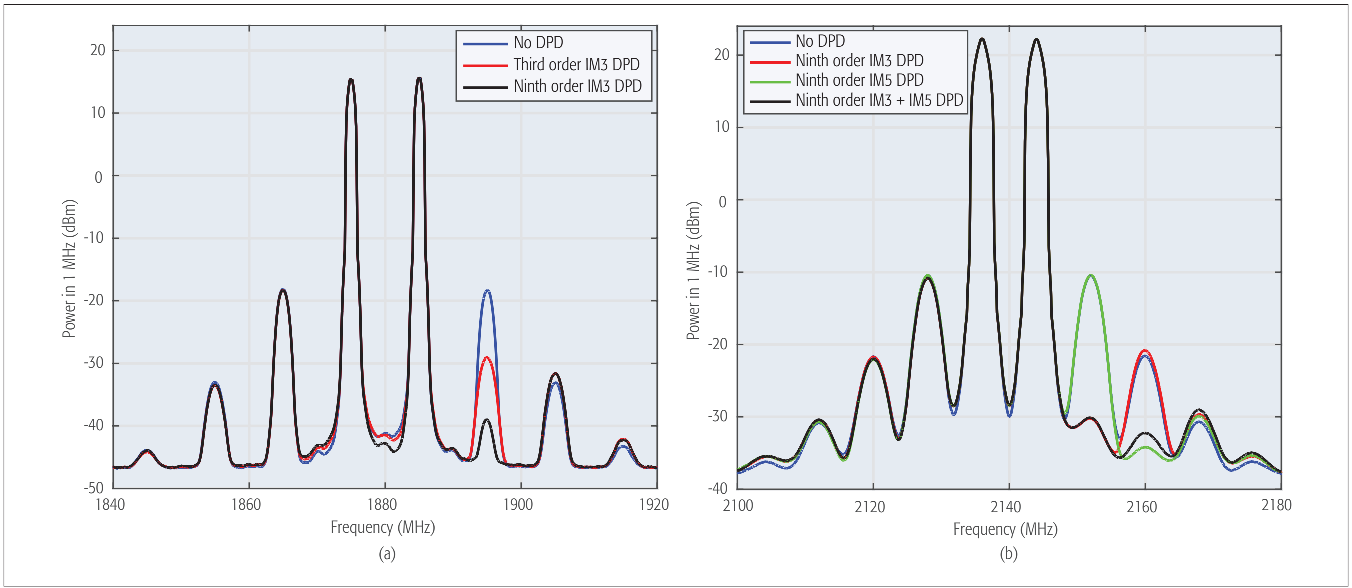

Figure 5. Two RF measurement examples showing the gain from using a sub-band DPD focused on the positive IM3 and/or IM5 sub-bands: a) IM3 spur reduction with both third-order and ninth-order sub-band DPDs are demonstrated, using a real commercial mobile PA operating at +20 dBm. A multicluster TX scenario with 1 MHz allocated per cluster and 10 MHz cluster spacing is used; b) both IM3 and IM5 spurs are predistorted separately as well as simultaneously using ninth order sub-band DPDs. The measured PA is a commercial off-the-shelf small cell base station PA, operating at +25 dBm output power. A multicluster TX scenario with 2 MHz allocated per cluster and 8 MHz cluster spacing is used.

reducing the computational complexity substantially. The example results indicate and demonstrate that with sub-band DPD, the spurious emission limits can be met with much smaller power back-offs. Thus, by employing sub-band DPD, uplink coverage, throughput, and PA efficiency need not be sacrificed in order to meet the emission limits. Sub-band DPD techniques can also be employed to protect the FDD terminal receiver from desensitization, as well as to protect the primary users in cognitive radio networks. This article seeks to spark discussion, raise awareness, and catalyze further research in the field of reduced-complexity DPD for non-contiguous transmissions, in particular for mobile devices.

Acknowledgment

This work was partially supported by the Finnish Funding Agency for Innovation (Tekes) and the National Science Foundation (NSF) under the projects "Cross-Layer Modeling and Design of Energy-Aware Cognitive Radio Networks" and "Future Uncoordinated Small-Cell Networks Using Reconfigurable Antennas." The work was also supported by the Academy of Finland under grants 284694 and 288670.

References

[1] 3GPP TS 36.101 V12.4.0, "LTE Evolved Universal Terrestrial Radio Access (E-UTRA): User Equipment (UE) Radio Transmission and Reception," Rel. 12, June 2014.

[2] IEEE Std 802.22-2011, "Cognitive Wireless RAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Policies and Procedures for Operation in the TV Bands," July 2011.

[3] ITU-R Rec. SM.329-12, "Unwanted Emissions in the Spurious Domain," Sept. 2012.

[4] C. S. Park *et al.*, "Carrier Aggregation for LTE-Advanced: Design Challenges of Terminals," *IEEE Commun. Mag.*, vol. 51, no. 12, Dec. 2013, pp. 76–84.

[5] F. M. Ghannouchi and O. Hammi, "Behavioral Modeling and Predistortion," *IEEE Microwave,* vol. 10, no. 7, Dec. 2009, pp. 52–64.

[6] P. Roblin *et al.*, "Concurrent Linearization: The State of the Art for Modeling and Linearization of Multiband Power Amplifiers," *IEEE Microwave*, vol. 14, no. 7, Nov. 2013, pp. 75–91.

[7] J. Kim *et al.*, "A Generalized Architecture for the Frequency-Selective Digital Predistortion Linearization Technique," *IEEE Trans. Microwave Theory Tech.*, vol. 61, no. 1, Jan. 2013, pp. 596–605.

[8] S. Bassam, M. Helaoui, and F. Ghannouchi, "Channel-Selective Multi-Cell Digital Predistorter for Multi-Carrier Transmitters," *IEEE Trans. Microwave Theory Tech.*, vol. 60, no. 8, Aug. 2012, pp. 2344–52.

[9] M. Abdelaziz *et al.*, "Reduced-Complexity Power Amplifier Linearization for Carrier Aggregation Mobile Transceivers," *Proc. IEEE Int'l. Conf. Acoustics, Speech, and Signal Processing '14*, Florence, Italy, May 2014.

[10] M. Abdelaziz *et al.*, "Sub-Band Digital Predistortion for Noncontiguous Transmissions: Algorithm Development and Real-Time Prototype Implementation," *Proc. 49th Asilomar Conf. Signals, Systems, and Computers*, Pacific, Grove, CA, Nov. 2015.

[11] Z. Fu *et al.*, "Frequency-Selective Digital Predistortion for Unwanted Emission Reduction," *IEEE Trans. Commun.*, vol. 63, no. 1, Jan. 2015, pp. 254–67.

[12] C. Yu *et al.*, "Band-Limited Volterra Series-Based Digital Predistortion for Wideband Rf Power Amplifiers," *IEEE Trans. Microwave Theory Tech.*, vol. 60, no. 12, Dec. 2012, pp. 4198–4208.

[13] R. Braithwaite, "Closed-Loop Digital Predistortion (DPD) Using an Observation Path with Limited Bandwidth," *IEEE Trans. Microwave Theory Tech.*, vol. 63, no. 2, Feb. 2015, pp. 726–36.

[14] A. S. Tehrani *et al.*, "A Comparative Analysis of the Complexity/Accuracy Trade-Off in Power Amplifier Behavioral Models," *IEEE Trans. Microwave Theory Tech.*, vol. 58, no. 6, June 2010, pp. 1510–20.

[15] M. Saint-Laurent *et al.*, "A 28 nm DSP Powered by an On-Chip LDO for High-Performance and Energy-Efficient Mobile Applications," *IEEE J. Solid-State Circuits*, vol. 50, no. 1, Jan. 2015, pp. 81–91.

Biographies

Mahmoud Abdelaziz (mahmoud.abdelaziz@tut.fi) received his B.Sc. degree (with honors) and M.Sc. degree in electronics and electrical communications engineering from Cairo University, Egypt, in 2006 and 2011. He is currently pursuing his doctoral degree at Tampere University of Technology, Finland, where he works as a researcher with the Department of Electronics and Communications. From 2007 to 2012 he has been working as a communication systems and signal processing engineer at Newport Media Inc. (currently acquired by Atmel), as well as other companies in the wireless industry. His research interests include statistical and adaptive signal processing in flexible radio transceivers, wideband digital pre-distortion, and cognitive radio systems.

Zhu Fu is a senior engineer working at Qualcomm, San Jose, California. She received her Bachelor's and Master's degrees from Beijing Institute of Technology, and her Ph.D. degree from Worcester Polytechnic Institute (WPI), Massachusetts, in May 2014, all in electrical engineering. Her research interests include digital pre-distortion, spectrally agile waveforms, and RF front-end characterization.

LAURI ANTTILA [S'06, M'11] received his M.Sc. and D.Sc. (Tech) degrees (with honors) from Tampere University of Technology in 2004 and 2011. He has published over 60 articles in international peer-reviewed journals and conferences, as well as two book chapters. Currently, he is a senior research fellow at the Department of Electronics and Communications Engineering at Tampere University of Technology. His research interests include signal processing for wireless communications, digital front-end signal processing, transmitter and receiver linearization, power-efficient transmitter structures, and radio implementation challenges in 5G cellular radio and full-duplex radio.

ALEXANDER M. WYGLINSKI [S'99, M'05, SM'11] is an associate professor of electrical and computer engineering at WPI and director of the Wireless Innovation Laboratory (WI Lab). He received his B.Eng. and Ph.D. degrees in 1999 and 2005 from McGill University, and his M.Sc.(Eng.) degree from Queen's University in Kingston in 2000, all in electrical engineering. Throughout his academic career, he has published over 35 journal papers, over 75 conference papers, nine book chapters, and two textbooks. His current research activities include wireless communications, cognitive radio, software-defined radio, dynamic spectrum access, spectrum measurement and characterization, electromagnetic security, wireless system optimization and adaptation, and cyber physical systems. He is currently or has been sponsored by organizations such as the Defense Advanced Research Projects Agency (DARPA), the Naval Research Laboratory (NRL), the Office of Naval Research (ONR), the Air Force Research Laboratory (AFRL) - Space Vehicles Directorate, The MathWorks, Toyota InfoTechnology Center U.S.A., and the National Science Foundation. He is a member of Sigma Xi, Eta Kappa Nu, and the ASEE.

MIKKO VALKAMA [S'00, M'01, SM'15] received his M.Sc. and Ph.D. degrees (both with honors) in electrical engineering from Tampere University of Technology in 2000 and 2001, respectively. In 2002, he received the Best Ph.D. Thesis award from the Finnish Academy of Science and Letters for his dissertation, *Advanced I/Q Signal Processing for Wideband Receivers: Models and Algorithms*. In 2003, he worked as a visiting researcher with the Communications Systems and Signal Processing Institute at San Diego State University, California. Currently, he is a full professor and department vice-head at the Department of Electronics and Communications Engineering at Tampere University of Technology. He has been involved in organizing conferences like IEEE SPAWC '07 (Publications Chair) held in Helsinki, Finland. His general research interests include communications signal processing, estimation and detection techniques, signal processing algorithms for software defined flexible radios, cognitive radio, full-duplex radio, radio localization, 5G mobile cellular radio, digital transmission techniques such as different variants of multicarrier modulation methods and OFDM, and radio resource management for ad fhoc and mobile networks.

NETWORK TESTING

Ying-Dar Lin

Erica Johnson

Irena Atov

Going forward, in the updated Call for Papers for future issues, we have expanded the Network Testing Series to Network Testing and Analytics. Network analytics complements testing and modeling activities during not only the design phase of a network or system, but also the implementation and subsequent production and operation phases. To ensure that network and system solutions deliver their required service quality, a range of testing and analytics capabilities and practices are required to model, measure, and evaluate the effectiveness of the solutions over their life cycles. Our newly added topics of interest include network analytics infrastructure, custom analytics solutions for planning, traffic management and overall improvement of network infrastructure and operations, network performance monitoring and diagnostics, predictive network analytics to help detect network and application issues, and analytic tools for troubleshooting and identifying root causes.

We thank the former Editor-in-Chief of *IEEE Communications Magazine*, Dr. Sean Moore, for his support in making the Network Testing and Analytics Series a reality. Also, we are grateful to the current Editor-in-Chief, Dr. Osman Gebizlioglu, for his support and guidance in extending and preparing this Series. As a new Series Editor, Dr. Irena Atov brings new perspectives from communications service providers.

For this issue, still under the title of Network Testing Series, we accepted three articles from 12 submissions. These feature the latest developments in various areas of technology related to network testing, from test equipment for lab testing, measurement and analysis platforms for field operational networks, to smartphone test frameworks for multi-path TCP under heterogeneous cellular and Wi-Fi networks.

Large-scale and distributed Internet measurement platforms are explored in the next article by Pedro Casas *et al.*, "Unveiling Network and Service Performance Degradation in the Wild with mPlane." The authors develop applications and use cases for their network measurement and analysis platform, mPlane, and focus on highlighting and troubleshooting end-customer Internet issues. By deploying mPlane in both fixed-line and cellular operational Internet service provider networks, the authors illustrate how to use this measurement system to automatically detect and diagnose network and service performance issues with very different root causes in different scenarios.

The advantages and use of field programmable gate array (FGPA) network platforms combined with high-level synthesis (HLS) tools for achieving reliable and affordable measurement systems for high-speed networks (10 GB/s and beyond) is investigated by Mario Ruiz *et al.* in their article "Accurate and Affordable Packet-Train Testing Systems for Multi-Gigabit-per-Second Networks." In this first article the authors focus on developing hardware-based packet-train testing solutions and explore the benefits of these compared to software counterparts with respect to accuracy and cost. Two FPGA-based testing solutions are developed for 1 GB/s and 10 GB/s, respectively, and their applicability is illustrated in an operators' production network. The results obtained with hardware development closely match the theoretical values, and overall, the proposed solutions provide consistent capacity and delay measurements, whereas the accuracy of the software solutions is shown to be more limited, particularly at higher (multi-gigabit-per-second) data rates.

The issue concludes with an article related to Multipath Transmission Control Protocol (TCP), a recently standardized TCP extension by the Internet Engineering Task Force, to allow TCP connection to use multiple paths in order to maximize resource usage within the network and increase redundancy. The article is by Quetin De Coninck *et al.*, with the title "Observing Real Smartphone Applications over Multipath TCP." The authors analyze a use case, useful in the context of wireless networks, where smartphones can use both Wi-Fi and a cellular network to provide higher bandwidth. They develop an open source framework for automating test measurements in this environment, and use this framework to conduct a comprehensive analysis on how eight popular smartphone applications (e.g., Facebook, Messenger, Firefox, YouTube) interact with Multipath TCP under different network conditions.

We hope readers will find these articles useful in highlighting recent developments, and as well in motivating their own work. We thank all authors and reviewers for contributing to this Series.

BIOGRAPHIES

YING-DAR LIN [F] (ydlin@cs.nctu.edu.tw) is a Distinguished Professor at National Chiao Tung University in Taiwan. He received his Ph.D. in computer science from the University of California Los Angeles in 1993. He is the director of 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).

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 Erica to 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.

IRENA ATOV [SM] (i.atov@ieee.org) received her Ph.D. from RMIT University, Melbourne, Australia, in 2003. She has worked in academia in both teaching and research roles, consulted for industry, and worked for Telstra in their Network Technology group as program director of Network Analytics and Resilience. Her research has led to development of commercial IT software products in the area of network planning and currently she is co-founder and CEO of InCluesive Technologies, a networking consulting company.

Unveiling Network and Service Performance Degradation in the Wild with mPlane

Pedro Casas, Pierdomenico Fiadino, Sarah Wassermann, Stefano Traverso, Alessandro D'Alconzo, Edion Tego, Francesco Matera, and Marco Mellia

ABSTRACT

Unveiling network and service performance issues in complex and highly decentralized systems such as the Internet is a major challenge. Indeed, the Internet is based on decentralization and diversity. However, its distributed nature leads to operational brittleness and difficulty in identifying the root causes of performance degradation. In such a context, network measurements are a fundamental pillar to shed light on and unveil design and implementation defects. To tackle this fragmentation and visibility problem, we recently conceived mPlane, a distributed measurement platform that runs, collects, and analyzes traffic measurements to study the operation and functioning of the Internet. In this article, we show the potentiality of the mPlane approach to unveil network and service degradation issues in live operational networks, involving both fixed-line and cellular networks. In particular, we combine active and passive measurements to troubleshoot problems in end-customer Internet access connections, or to automatically detect and diagnose anomalies in Internet-scale services (e.g., YouTube) that impact a large number of end users.

INTRODUCTION

Since the early days of the Internet, network measurements have always constituted a pillar in understanding the behavior of the network, especially when something goes wrong. To address this issue we recently conceived mPlane, a large-scale network measurement and analysis framework. mPlane is a distributed measurement architecture to coordinate traffic measurements, with built-in support for iterative measurement and, most of all, automated and advanced analysis. Probes, which perform measurements; Repositories, which store, aggregate, correlate, and analyze them; and a Supervisor, which orchestrates components are the basics of mPlane. Reasoners are intelligent building blocks that extract knowledge, and offer support to network administrators.

The complete mPlane architecture was previously presented in [1], where only some simple examples of applications were discussed. In this article, we briefly recall the mPlane architecture. Then we fully develop use cases that focus on highlighting and troubleshooting on end-customer Internet issues. In particular, we show how to use mPlane to detect and diagnose network and service performance degradation events in operational Internet service provider (ISP) networks, exploiting the richness of the measurements mPlane probes perform, and the analysis capabilities in terms of anomaly detection. We focus on three different case studies:

- Diagnosis of performance degradation in end-customer Internet access connections using hybrid measurements
- Detection and diagnosis of service availability issues in cellular networks
- Detection and diagnosis of quality of experience (QoE)-relevant issues in YouTube

We additionally present a fourth analysis scenario in which we implement a proximity location service to analyze inter-AS paths between selected servers (e.g., YouTube servers, Facebook servers, etc.), based on distributed active measurements through RIPE Atlas.

After briefly summarizing related platforms in the second section, we give an overview of the mPlane architecture. We then provide details on the operational ISPs considered in the case studies and detailed traffic datasets. The results obtained in the application of the deployed mPlane framework are then presented. Finally, we conclude the article.

RELATED PLATFORMS

Many measurement platforms have been proposed in the past, such as iPlane [2], PerfSONAR [3], and RIPE Atlas.[1] Each of them targets specific needs, and focuses mostly on the monitoring of the network layer, relying exclusively on active measurements. For instance, iPlane focuses on network topology discovery via traceroute measurements to build a predictive model of path latency; PerfSONAR and RIPE Atlas offer active measurements from a distributed platform (throughput, ping, traceroute, etc.), to again monitor network paths. They offer limited processing capabilities or algorithms for advanced diagnosis, even if it is possible to run (custom) algorithms on top of collected measurements.

Unveiling network and service performance issues in complex and highly decentralized systems such as the Internet is based on decentralization and diversity. However, its distributed nature leads to operational brittleness and difficulty in indentifying the root causes of performance degradation.

Pedro Casas, Pierdomenico Fiadino, Sarah Wassermann, and Alessandro D'Alconzo are with the Telecommunications Research Center of Vienna; Stefano Traverso and Marco Mellia are with Politecnico di Torino.; Edion Tego and Francesco Matera are with Fondazione Ugo Bordoni

[1] https://atlas.ripe.net/.

The proposed distributed platform is an instance of the more generic mPlane architecture. The mPlane monitoring system is composed of four different entities: probes, repositories, supervisor, and reasoners. They inter-operate thanks to a standard protocol, and are logically organized into three layers.

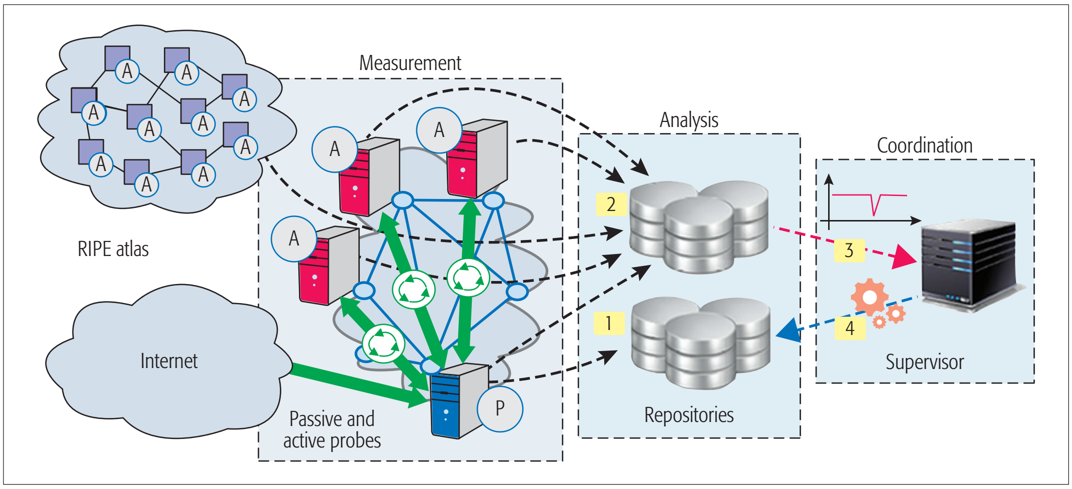

Figure 1. The distributed measurement platform following the mPlane architecture for ISP network troubleshooting and its workflow. Green arrows correspond to both active and passive measurements. Black arrows correspond to measurement data that are exported from probes to repositories. Red arrows correspond to anomaly notification reports. Blue arrows depict the requests made by the Supervisor to trigger deeper data analysis, using, for example, external distributed measurement frameworks such as RIPE Atlas.

For instance, Calyam proposed an anomaly detection system specifically designed to be integrated on top of PerfSONAR measurements [4].

The mPlane platform is instead specifically designed to integrate existing and new measurement probes. Indeed, mPlane already integrates RIPE Atlas, so it is possible to instrument RIPE measurements from an mPlane reasoner. In addition, mPlane offers support for passive and hybrid measurements. For instance, in this article we leverage passive measurements collected by Tstat [5] to compare with active test and to feed the anomaly detection algorithms with rich information about the status of complex platforms, such as the YouTube infrastructure. Note that capturing and processing passive measurements in operational networks such as those analyzed in this article pose a much more complex problem in terms of volume and analysis techniques than relying on pure active measurements.

Besides coordinating probes, mPlane offers storage capabilities, where data can be collected and processed by analytic components, which, for instance, integrate methodologies to identify anomalies. Furthermore, mPlane does not target the monitoring of the network layer only. For instance, we show its ability to monitor content delivery networks, and their direct impact on end-user QoE. In summary, mPlane is the first flexible, open, and intelligent platform to aim at providing monitoring and analysis capabilities for Internet services at large, and not only of pure network layer problems.

Considering anomaly detection algorithms, many proposals are present in the literature — see, for instance, the survey of Chandola [6]. Most of them could easily be integrated into mPlane. The most innovative part offered by mPlane is the ability to design complete workflows, where actions can be triggered as a consequence of alarms, and domain knowledge can easily be integrated. This goes in the same directions as the proposal by Kanuparthy in [7, 8], where the Pythia system is introduced; still, different from mPlane, Pythia focuses mainly on network path troubleshooting, relying exclusively on active measurements.

System Description, Components and Approach

The proposed distributed platform is an instance of the more generic mPlane architecture. The mPlane monitoring system is composed of four different entities: *probes*, *repositories*, *supervisor*, and *reasoners*. They interoperate thanks to a standard protocol, and are logically organized into three layers in Fig. 1.

The Measurement Layer: Consists of probes located at vantage points within the monitored networks, which typically generate large amounts of measurement data. The system supports *active measurements* (e.g., `ping` or `traceroute`, *passive measurements* (e.g., the analysis of traffic flowing on a link), and *hybrid measurements* (e.g., the passive observation of active probing traffic). Measurement campaigns may be triggered on demand, with results returned as soon as the measurement is completed, or be run continuously, with results periodically exported into an mPlane repository to limit the storage utilization at the probe. For this specific instantiation of mPlane, the measurement layer includes the RIPE Atlas active platform using so-called mPlane *proxies* [1].

The Analysis Layer: Consists of *repositories*, which collect and aggregate data generated by the probes. Apart from the storage capacity, the analysis layer is provided with a set of analysis modules that process the data imported from the probes. Such processing may involve filtering, grouping, aggregation of the raw data imported from the probes, or more complex analytics such as anomaly detection. The results are a higher level of aggregation and visibility on the monitored network, and can be directly accessed through a standard queryable SQL-like interface.

The Coordination Layer: Consists of the *supervisor* and the *reasoner*. The first orchestrates probes and repositories. The latter receives output of analysis modules, and may trigger alarms or initiate additional operations in a reactive fashion (e.g., perform additional on-demand measurements to investigate the anomaly).

As indicated in Fig. 1, the measurement layer is decomposed into three different measurement approaches: hybrid measurements combine active and passive measurements, and are used in the analysis of end-customer Internet access connections; passive measurements are instrumented by deploying sniffers at points of presence (PoPs) aggregating a large number of customers; finally, active measurements are performed through both the RIPE Atlas and the IQM platforms.

Measurement Components

In more detail, the passive probe used in all fixed-line scenarios is Tstat [5]. Tstat is an open source packet analyzer capable of monitoring links up to several gigabits per second speed using commodity hardware. It extracts information about both TCP and UDP traffic flows at all the layers of the protocol stack, from simple flow size, or RTT average and standard deviation [9], up to layer 7 data (i.e., application-related). When considering cellular networks, we rely on the METAWIN passive monitoring probe [10], which is capable of handling the complete 3GPP protocol stack.

We use two different active platforms: RIPE Atlas probes and Internet QoS Measurement (IQM) probes.[2] Both probes are capable of RTT and path measurements using `ping` and `traceroute` tools; in addition, the IQM offers more complex measurements, for example, speed-tests and HTTP/HTTPS performance metrics. RIPE Atlas probes are maneuvered through a custom interface[3] that launches measurements and retrieve results automatically.

As a repository, we use DBStream [11]. DBStream is a data stream warehouse tailored for large-scale traffic monitoring applications. It continuously analyzes the measurements obtained by the probes. In particular, multiple instances of an anomaly detection analysis module (ADAM) run in parallel on top of DBStream, flagging anomalous behaviors in different traffic features. We briefly describe ADAM next.

Finally, the mPlane Supervisor is the standard one, provided by the mPlane Reference Implementation (RI). [4] The reasoner is a custom set of python code that, interacting with the supervisor, collects the output of the analysis module, and eventually starts measurement for further investigations.

Anomaly Detection and Diagnosis

ADAM [12] is an mPlane analysis module that detects unusual deviations in the probability distribution of a monitored feature over time. In short, the ADAM base algorithm detects anomalies based on the degree of similarity between the distribution of a feature as currently observed, and a set of distributions describing the normal behavior of the monitored feature (i.e., a baseline). The latter is built through a progressive refinement heuristic, which takes into account the structural characteristics of traffic such as time of day variations, presence of pseudo-cyclic weekly patterns, and long-term variations. The baseline thus evolves over time to adapt to the dynamics of the system. The similarity between distributions is computed on the basis of a symmetric extension to the well-known Kullback-Leibler divergence. When the difference is bigger than an (adaptive) confidence threshold, ADAM raises an alarm. More details on ADAM can be found in [12], where we apply the same base algorithms to detect anomalies in content delivery network (CDN) services.

Multiple instances of ADAM run in parallel, each analyzing multiple traffic features at the same time. We split the monitored features into two groups, referred to as symptomatic and diagnostic features. Symptomatic features are defined such that their abrupt change directly correlates to the presence of abnormal and potentially harmful events. Diagnostic features provide contextual details of the anomalies, pointing to their root causes. In a nutshell, by locating those diagnostic signals that show a change at the same time or same temporal scope as the detected anomaly, one gets a more targeted and specific indication of which features might be causing the anomaly.

Diagnostic features shall provide contextual details of the anomalies, pointing to their root causes. In a nutshell, by locating those diagnostic signals that show a change at the same time or same temporal scope to the detected anomaly, one gets a more targeted and specific indication of which features might be causing the anomaly.

Analysis Workflow

The detection and diagnosis of anomalies runs as a continuous process. The supervisor instructs passive probes to run continuously, measuring and exporting the obtained data to the DBStream repository. When an alarm is triggered, the reasoner can instruct active probes to run further measurements. The measurement and analysis workflow depends on the specific case study, but apart from particular instrumentation details, all the scenarios follow the same basic steps, described in Fig. 1.

Step 1 — Passive Traffic Monitoring: Passive probes are deployed at vantage points. In the case of large-scale service monitoring, probes are deployed at PoPs aggregating a large number of customers, and their real traffic is captured, discarding all privacy-sensitive information.[5] In the case of end-customer access monitoring, Tstat is installed at the same server that instantiates active measurements (a simple FTP server used for speed tests), resulting in the aforementioned hybrid measurement approach.

Step 2 — Active Probing: Active probes scattered in the ISP (five per region) continuously perform speed tests to measure the available bandwidth on the network paths reaching the ISP customers, periodically downloading (uploading) files from (to) the FTP server. Active probes log the achieved application layer throughput, transferring the results to the same DBStream repository.

Step 3 — Detection of Anomalies: Multiple instances of ADAM run on top of DBStream, analyzing a set of features. As soon as an abrupt change is detected in a symptomatic feature, an alarm is raised to the reasoner.

Step 4 — Correlating Multi-Source Measurement Data: When alarms suggesting unexpected performance degradation are detected, the reasoner runs correlation analysis (e.g., factor analysis) to investigate which features show

[2] IQM probes are directly deployed by the ISP.

[3] https://github.com/pierdom/atlas-toolbox.

[4] https://github.com/fp7mplane/protocol-ri

[5] The deployment and the information collected for this has been approved by the ISP security and ethics boards.

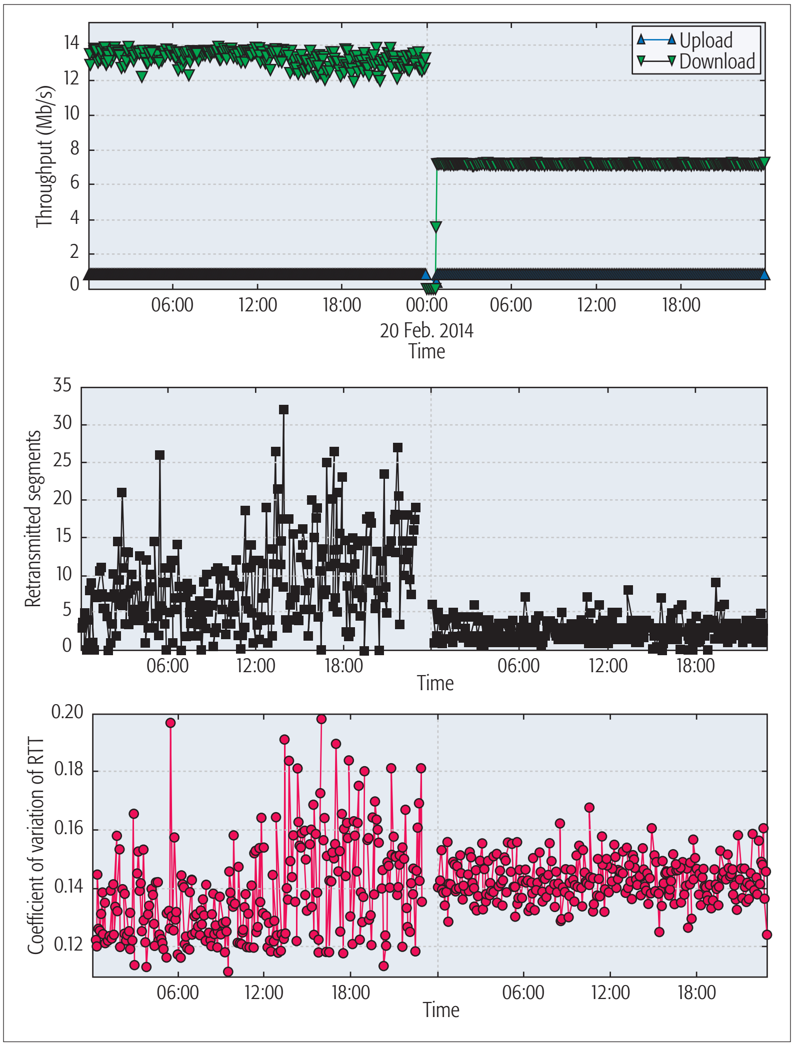

Figure 2. Evolution of time of the throughput measured by one active probe (top), the number of retransmitted segments (center), and the coefficient of variation of the RTT (bottom). U-1 Mb/s/D-16 Mb/s ADSL probe.

a similar abrupt change. Correlated features are then compared against a catalog of known anomaly patterns or signatures, and if a match is found, the most probable cause(s) are reported to the network operator.

Deployment and Datasets

For the sake of end-customer connection monitoring, we deployed active probes in more than 30 locations scattered in the operational country-wide network of the ISP. As regular customers, active probes connect to the network using an asymmetric DSL (ADSL) or fiber to the home (FTTH) access technology. The FTP server is installed in the ISP data center. This ISP offers three configurations for ADSL connections (U-1 Mb/s/D-16 Mb/s, U-1 Mb/s/D-12 Mb/s, and U-0.5Mb/s/D-8Mb/s) and one for FTTH (U-10Mb/s/D-10Mb/s). Each probe periodically runs a speed test by uploading (downloading) files of predefined size, measuring the application layer throughput. Tstat runs on the same server where the FTP is, and logs each TCP flow related by the active tests. We ran the system for three months (February 1–April 30, 2014), observing the time series of application-layer throughput. Speed tests are scheduled every 4 min, resulting in a total dataset of 1.2 million speed test reports produced by the active probes, and as many TCP entries in the log generated by Tstat.[6]

For the analysis of cellular network traffic and service availability, that is, the possibility for a client to obtain a response from servers of given service, we deployed the METAWIN passive probe at the core of a cellular ISP network in the EU. Traffic flows were captured at the well-known Gn interface for two consecutive days in mid-2014. In this case study, we only consider Domain Name System (DNS) traffic. The DNS is the core component of the Internet, providing flexible decoupling of a service's domain name and the hosting IP addresses. Anomalies in Internet-scale services are likely to change the normal DNS usage patterns. For example, users accessing a temporarily unreachable service would generate a new query at every connection retry. For that reason, we extract features derived from the DNS, including frequency of DNS requests, error codes, and so on.

Finally, QoE-based service performance analysis is performed again using Tstat, which passively monitors traffic from a PoP where there are 30,000 residential customers. We focus on the analysis of YouTube traffic, which we extract using Tstat classification modules. The complete dataset corresponds to four weeks of YouTube video traffic flows captured during the second quarter of 2013.

Results and Discussion

In this section we report some of the results obtained with the deployed mPlane framework in the aforementioned case studies. Due to lack of space we do not provide a fully detailed description on the complete diagnosis process for the presented scenarios, but refer the interested reader to https://www.ict-mplane.eu/public/public-deliverables (reports D4.1-4) for further details.

Combining Active and Passive Measurements

For the sake of brevity, we report here two examples of anomalies that could be present in ADSL access links: low SNR in ADSL channels and path congestion.

Low SNR in ADSL Lines: Figure 2 (top) reports the evolution over time of the throughput measured by an active probe accessing the ISP network through a U-1 Mb/s/D-16 Mb/s ADSL interface for two days.[7] Observe that the download throughput curve appears to be noisy during the first day, while after midnight, the ADSL line was re-calibrated to U-1 Mb/s/D-8 Mb/s. After then, speed test measures are much more stable over time. By correlating such output with the statistics provided by Tstat, we notice a fairly large fraction of retransmitted segments during the first day (center), and a constant coefficient of variation of the RTT (bottom). The absence of evident day-night patterns let us exclude that this situation might be due to network congestion, since this typically emerges only during peak periods.

The most probable cause for this anomaly is the occurrence of low signal-to-noise ratio (SNR)

[6] Access lines for experiments are devoted to test only, with no actual customer using those. As such, the testing traffic has minimal interference with customers' traffic that crosses the same path to the FTP server. Similarly, no privacy issues are raised.

[7] The throughput reported in the plots is below the nominal bandwidth since it is measured at the application level.

[8] SNR and BER can be read from SNMP measurements as defined by RFC 2662.

events at the physical link, which can lead to large bit error rate (BER).[8] Losses due to noise cause TCP congestion control to (randomly) slow down the download. The confirmation of this hypothesis is given by the second half of the plots in Fig. 2, when the ADSL modem automatically reduces the downlink capacity to improve the SNR (i.e., negotiating 8 Mb/s instead of 16 Mb/s), thus considerably reducing the packet loss rate and making RTT measurements more stable.

Congestion in the Network: Figure 3 (top) reports the evolution over time of the throughput measured by second active probe (U-1 Mb/s/D-12 Mb/s ADSL). During both days, a clear degradation in throughput is detected, with stable values during the night, that is, when the network is typically lightly loaded. Conversely, available capacity greatly decreases during peak time. This suggests that congestion may appear in the path toward the FTP server. By inspecting the statistics provided by Tstat at the server side, we confirm this intuition. Indeed, notice how the RTT coefficient of variation (bottom) and the rate of retransmitted packets (center) considerably increase during the peak utilization period. From verification with the ISP, such a probe accesses the Internet through a bottlenecked virtual leased line, the available bandwidth of which is out of the control of the operator.

Detecting and Diagnosing Availability Issues

We now present a case study based on the detection and diagnosis of a large-scale anomaly that occurred in the aforementioned cellular network. A significant and anomalous increase in the number of DNS requests is observed between 9:00 and 10:00 of the second day. Conversations with the network operations team revealed that the anomaly caused heavy stress in specific parts of the network. Figure 4a depicts the output of ADAM when applied to the distribution of DNS requests per device, which is defined as the symptomatic feature. ADAM systematically generates anomaly warnings during the one hour duration of the anomaly.

To discover the root causes of the detected anomaly, we define a set of diagnostic features related to the class of problems we target based on expert know-how. In particular, we consider the following set of features: anonymized mobile device identifier (MSID), contacted DNS server IP, radio access technology (RAT), access point name (APN), type allocation code (TAC), DNS requested full qualified domain name (FQDN), device manufacturer, device operating system (OS), and error code of the DNS response (DNS rcode). The first step of the diagnosis consists of identifying which of these features present a significant change in their probability distribution, simultaneous with the alarms generated by running ADAM on the symptomatic feature.

Figures 4b and 4c provide a closer look into the anomaly, comparing the output of ADAM when applied to two of the diagnostic features: the distribution requested FQDNs and the distribution of devices' OS type. Both ADAM outputs also flag anomalies in these two features exactly at the same time as the main anomaly trigger, suggesting that the issue might be due to specific devices (OS) querying for certain services (FQDN).

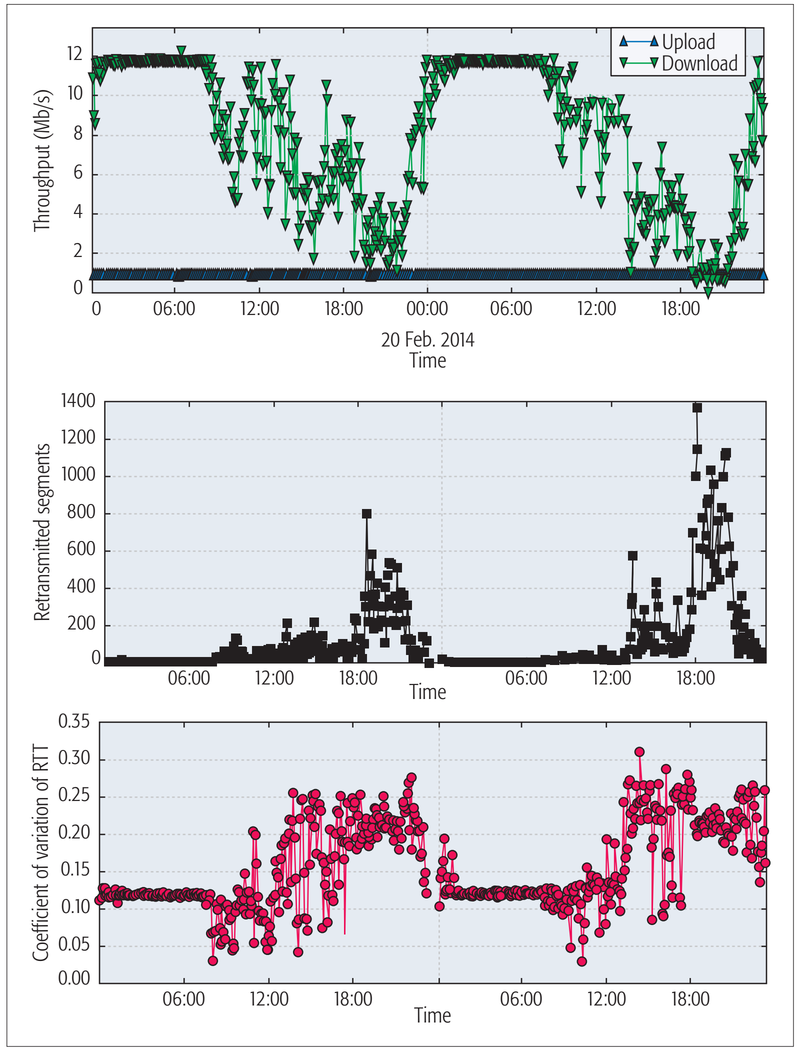

Figure 3. Evolution over time of the throughput measured by one active probe (top), the number of retransmitted segments (center), the coefficient of variation of the RTT (bottom). U-1 Mb/s/D-12 Mb/s ADSL probe.

The next step of the diagnosis is to drill down into each of the dimensions that are highly correlated with the anomaly. This can be achieved, for example, by comparing the heavy hitters before and during the anomaly. For the specific case of the FQDN diagnostic feature, we observed that while some of the top elements present stable behavior (**.facebook.com* and **.google.com*), the FQDNs **.apple.com.akadns.net*, **.push-apple.com.akadns.net*, and *xy-courier.push.apple.com* show a significant increase, pointing to a problem in the availability of the push notification service deployed by Apple.

Detecting and Diagnosing QoE-Relevant Anomalies

The last case study consists of the detection and diagnosis of a major YouTube anomaly impacting the QoE of a large number of customers during several days at peak load times. As the issue was caused by an unexpected cache selection by Google (at least according to our diagnostic analysis), the ISP internal RCA did not identify any problems inside its boundaries. As reported by the ISP operations team, the anomaly occurs on Wednesday, May 8, 2013. We therefore focus on the analysis of the week spanning the anomaly, from Tuesday the 7th to Sunday the 12th. In the following analysis, we generally use 50 percent percentile values instead of averages to filter out outliers.

Figures 5a and b plot the time series of two different symptomatic features related to the YouTube download performance and to the end-user QoE. Figure 5a depicts the median across all YouTube flows of the download flow throughput during the complete week. There is a normal reduction of the throughput on Tuesday at peak load time, between 20 h and 23 h. However, from Wednesday on, this drop is significantly higher, dropping way below a predefined bad QoE threshold of 400 kb/s, as found in [13].

To better monitor the QoE of YouTube videos from flow-level passive measurements, in [14] we introduced a novel QoE-based key performance indicator (KPI) defined as the ratio between the average download throughput (ADT) and the corresponding video bit rate (VBR), β = ADT/VBR. Intuitively, when β is lower than 1, the player buffer becomes gradually empty, ultimately leading to the stalling of the playback, which is the most relevant impact on QoE [13]. In [14] we found that no stallings are observed for β > 1.25. Based on this observation, Fig. 5b actually confirms that the throughput drops, heavily affecting the user experience, as the time series of KPI β falls well into the video stalling region (i.e., β < 1.25).

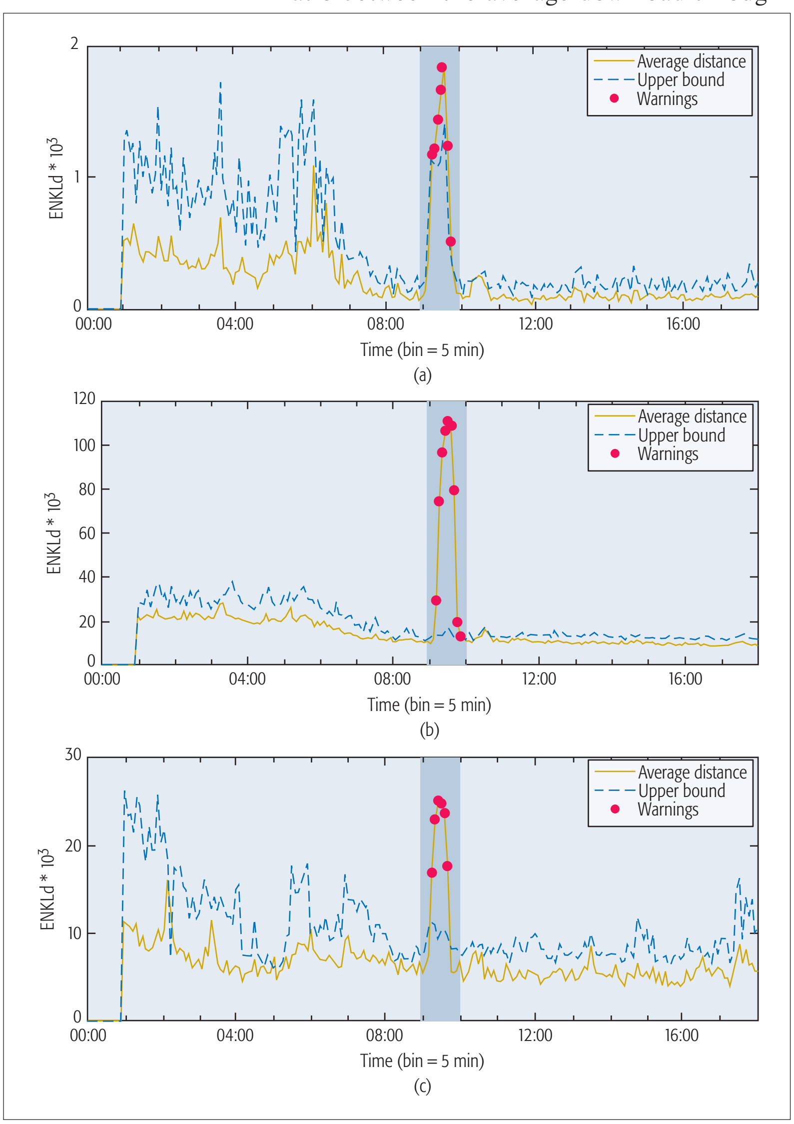

Figure 4. Output of the distribution-based detector for a — the symptomatic features (DNS query count per device) and two diagnostic features (b — FQDN and c — OS type). All the plotted features exhibit distribution changes during the anomalous event: a) ADAM output for the symptomatic feature; b) ADAM output for the diagnostic feature FQDN; c) ADAM output for the diagnostic feature OS type.

To conclude, we report in Figs. 5c and 5d the output of ADAM for two selected features. Figure 5c considers the per /24 YouTube subnetwork served volume as the monitored feature. From Wednesday May 8 onward, ADAM alarms rise from 15:00 to 00:00, which correspond to a different selection of YouTube servers done by Google to serve the monitored customers. Figure 5d reports the ADAM output for the average video flows download rate. In this case, ADAM detects the anomalies only between peak hours (21:00–23:00) from the 8th onward, coherent with the observations drawn from Fig. 5a. Comparing the changes on the served traffic volume distribution against those on the video flows download rate distribution, we observe that the server selection policy used by Google resulted in a QoE degradation only during peak hours on high load days. This suggests that either the selected servers were not correctly dimensioned to handle traffic load peaks, or there is some heavy network congestion at peak time in the paths from the selected Google servers to the customers.

To unveil this kind of Internet paths performance issue within mPlane, we next propose a technique to perform direct `traceroute` measurements in the downlink direction, from the Google servers to the customers (i.e., a *reverse traceroute*). Our technique avoids relying on IP spoofing (normally blocked by many ISPs), as done in previous work [15].

Distributed Active Measurements for Path Analysis

To analyze the performance of server-to-costumers Internet paths in the most general scenario using mPlane, we rely on the RIPE Atlas distributed active measurements framework. We developed DisNETPerf, a distributed Internet paths performance analyzer, to perform direct `traceroute` measurements in the downlink direction. In a nutshell, given a certain source server IP address IP_s, and a destination customer IP address IP_d, DisNETPerf locates the closest RIPE Atlas IP_{DNP} probe to IP_s, and periodically launches `traceroutes` from IP_{DNP} to IP_d, collecting different path performance statistics including RTT per hop, end-to-end RTT, losses, and so on.

Figure 6 (right) depicts the overall idea behind DisNETPerf. DisNETPerf uses a combined topological and geographical-based distance, as probes are located first by AS, Border Gateway Protocol (BGP) routing proximity, and then propagation delay. The selection of IP_{DNP} works as follows: given IP_s, we select all the probes in the same AS (or neighbor ASs if no local probes are found) and launch standard ping measurements toward IP_s. We consider the probe with the smallest minimum RTT as IP_{DNP}.

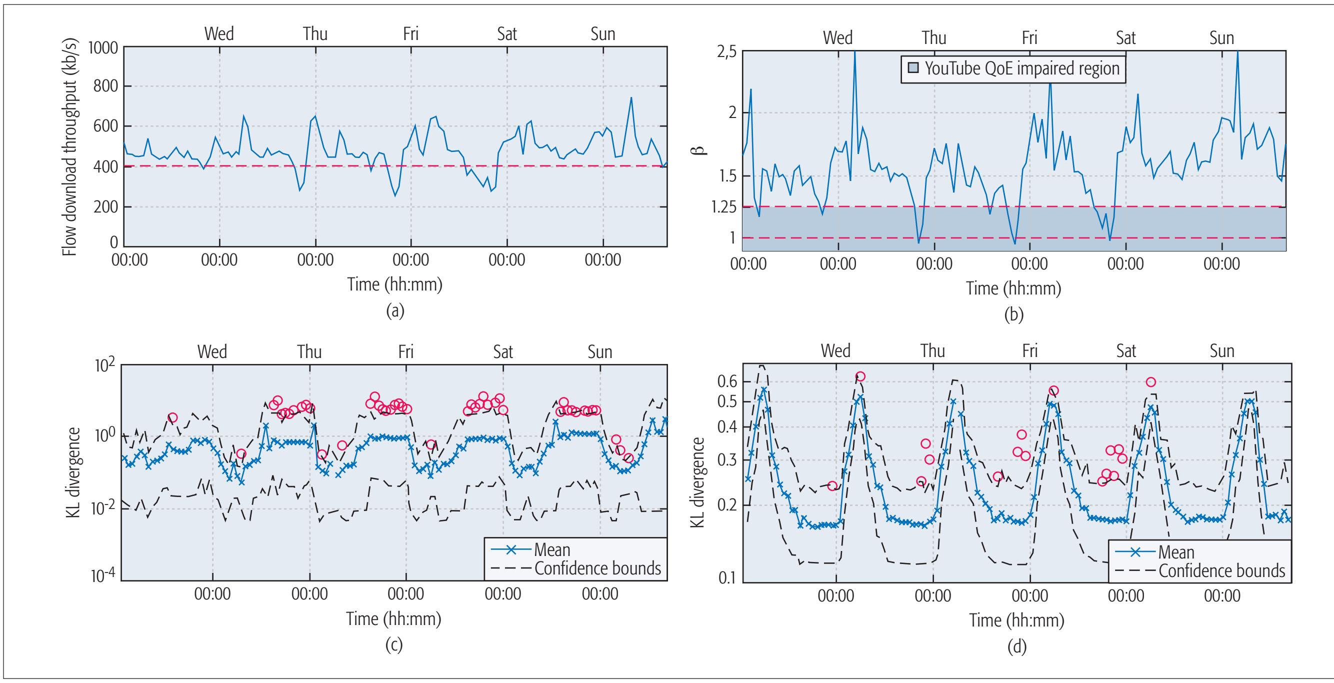

Figure 5. Detecting a QoE-relevant anomaly in a real ISP. There is a clear drop in the download flow throughput from Wednesday till Friday at peak load hours, between 20 h and 23 h. The additional drop in the QoE-based KPI β reveals a significant QoE degradation. The anomalies are flagged by ADAM in the selected symptomatic and diagnostic features: a) median of the flow download throughput per hour for all YouTube flows; b) median of β KPI per hour for all YouTube flows; c) anomalies in traffic volume served by YouTube /24 subnets; d) anomalies in the video flows average download throughput across YouTube users.

We say that the probe selected by DisNETPerf is a good probe w.r.t. IP_s and IP_d if the network path from IP_{DNP} to IP_d is highly similar to the path from IP_s to IP_d. Similar to [16], we define path similarity as the fraction of common links among both paths. Formally, we use the index route similarity (RSIM), defined as

$$RSIM(IP_{DNP}, IP_s, IP_d) = \frac{2 \times common\_links(IP_{DNP}, IP_s, IP_d)}{total\_links(IP_{DNP}, IP_s, IP_d)} \quad (1)$$

where *common_links* refer to the links shared in common by both paths, and *total_links* to the total number of links for both paths. Note that links can be defined at multiple granularities; in particular, for these evaluations we consider links at the AS level, the PoP level, and the router interface (IP) level. IP_s to AS mapping is done through the IP-to-ASN service provided by Team Cymru,[9] whereas IP_s to PoP mapping is achieved through the datasets made available by iPlane [2].

In Fig. 6 (left) we present evaluation results showing the applicability of DisNETPerf in terms of path similarity. The goal of the evaluation is to investigate whether the probe selection approach used by DistNETPerf obtains probes that present the most similar path to the one we want to actually monitor. We use RIPE Atlas probes as source and destination (i.e., IP_s and IP_d) in order to compute the real path (i.e., the ground truth) between source and destination. In the evaluation, we randomly select 100 RIPE Atlas source probes IP_{si}, i = 1, ...,100, and consider a single fixed destination probe IP_d. For each of the sources IP_{si} we run DisNETPerf to locate the 100 closest probes IP_{DNPi}, and obtain both the ground truth path going from IPsi to IP_d and the DisNETPerf path going from IP_{DNPi} to IP_d, and compute the RSIM index $RSIM(IP_{DNPi}, IP_{si}, IP_d)$, $i = 1, ..., 100$. We compute RSIM for AS level, PoP level, and IP level, and plot the CDFs for the three cases. Results are reported for two different groups, the former, in which $RSIM(IP_{DNPi})$ and IP_{si} are located in the same AS (black dotted lines), and the latter, in which $RSIM(IP_{DNPi})$ is located in a neighbor AS (red solid lines). When considering paths at the AS level, there is a significant difference between the groups, and the case of the same AS co-location results in near optimal results. Nevertheless, we observe that about 60 percent of the tests yield an RSIM index > 0.5. Finally, we observe that probes selected by DisNETPerf generally correspond to paths with the highest similarity to the ground truth ones. Indeed, in more than 84 percent of the tests performed, $RSIM(IP_{DNPi})$ results in the highest RSIM index among all the selected candidates.

Concluding Remarks

Unveiling network and service performance issues in complex and highly decentralized systems such as the Internet is a major challenge. mPlane provides a distributed measurement platform, which, among other applications, can be used to shed light on such performance issues. By deploying mPlane in both fixed-line and cellular operational ISP networks, we have shown how to use this powerful and novel framework to automatically detect and diagnose performance issues with very different root causes in different scenarios. Finally, note that all the software tools used in this article are publicly available at the mPlane project website (https://www.ict-mplane.

[9] http://www.team-cymru.org/IP-ASN-mapping.html

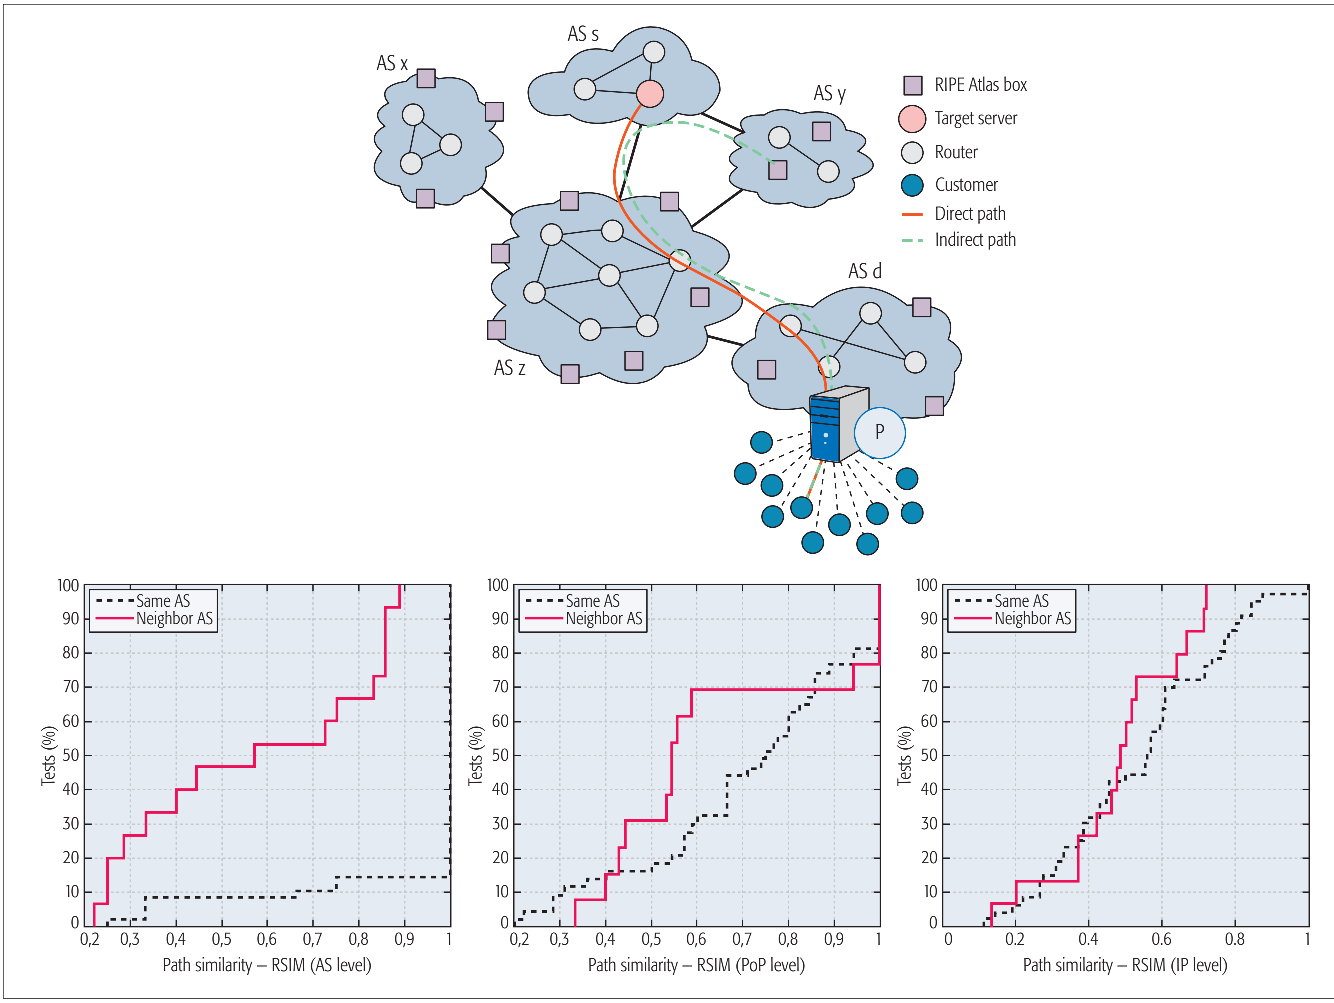

Figure 6. DisNETPerf approach. DisNETPerf achieves almost perfect path similarity at the AS level when IP_{DNP} is located in the same AS of IP_s, and path similarity above 50 percent for most of the cases, also at the PoP and interface level.

eu/). We refer the interested reader to https://www.ict-mplane.eu/public/use-cases for more details on how mPlane is applied to many other relevant use cases.

Acknowledgments

The research leading to these results has received funding from the European Union under the FP7 Grant Agreement no. 318627 (Integrated Project "mPlane"). Special thanks to Eike Kowallik, Stefano Raffaglio, and Andrea Fregosi from Fastweb, Italy, for their contributions to the analysis of anomalies in ADSL lines.

References

[1] B. Trammell *et al.*, "mPlane: an Intelligent Measurement Plane for the Internet," *IEEE Commun. Mag.*, Feature Topic on Monitoring and Troubleshooting Multi-domain Networks using Measurement Federations, vol. 52, no. 5, 2014, pp. 148–56.

[2] H. Madhyastha *et al.*, "iPlane: An Information Plane for Distributed Services," *Proc. 7th USENIX Symp. Op. Sys. Design and Implementation*, 2006, pp. 367–80.

[3] A. Hanemann *et al.*, "PerfSONAR: A Service Oriented Architecture for Multi-Domain Network Monitoring," *Proc. 3rd Int'l. Conf. Service-Oriented Computing*, LNCS 3826, 2005, pp. 241–54.

[4] P. Calyam *et al.*, "OnTimeDetect: Dynamic Network Anomaly Notification in perfSONAR Deployments," *Proc. 2010 IEEE Int'l. Symp. Modeling, Analysis & Simulation of Computer and Telecommun. Sys.*, 2010, pp. 328–37.

[5] A. Finamore *et al.*, "Experiences of Internet Traffic Monitoring with Tstat," *IEEE Network*, vol. 25, no. 3, 2011, pp. 8–14.

[6] V. Chandola *et al.*, "Anomaly Detection: A survey," *ACM Computing Surveys*, vol. 41, no. 3, 2009.

[7] P. Kanuparthy *et al.*, "Pythia: Detection, Localization, and Diagnosis of Performance Problems," *IEEE Commun. Mag.*, vol. 51, no. 11, 2013, pp. 55–62.

[8] P. Kanuparthy and C. Dovrolis, "Pythia: Diagnosing Performance Problems in Wide Area Providers," *Proc. USENIX Annual Technical Conf.*, Philadelphia, PA, 2014, pp. 371–82.

[9] M. Mellia *et al.*, "Passive Analysis of TCP Anomalies," *Computer Networks*, vol. 52, no.14, 2008, pp. 2663–76.

[10] F. Ricciato, "Traffic Monitoring and Analysis for the Optimization of a 3G Network," *IEEE Wireless Commun.*, vol. 13, no. 6, 2006, pp. 42–49.

[11] A. Bär *et al.*, "Large-Scale Network Traffic Monitoring with DBStream, a System for Rolling Big Data Analysis," *Proc. 2nd IEEE Int'l. Conf. Big Data*, 2014, pp. 165–70.

[12] P. Fiadino *et al.*, "On the Detection of Network Traffic Anomalies in Content Delivery Network Services," *Proc. 26th Int'l. Teletraffic Congress*, pp. 1--9, 2014.

[13] P. Casas *et al.*, "YouTube & Facebook Quality of Experience in Mobile Broadband Networks," *Proc. IEEE GLOBECOM Wksp. Quality of Experience for Multimedia Commun.*, 2012, pp. 1269–74.

[14] P. Casas *et al.*, "YouTube in the Move: Understanding the Performance of YouTube in Cellular Networks," *Proc. IFIP Wireless Days*, 2014, pp. 1–6.

[15] E. Katz-Bassett *et al.*, "Reverse Traceroute," *Proc. 7th USENIX Conf. Networked Sys. Design and Implementation*, 2010, pp. 15–28.

[16] N. Hu and P. Steenkiste, "Quantifying Internet End-to-End Route Similarity," *Proc. Passive and Active Measurement Conf.*, 2006.

Biographies

Pedro Casas is a senior researcher at Telecommunications Research Center of Vienna (FTW). He received a Ph.D. in computer science from Telecom Bretagne in 2010, and joined CNRS as a research fellow between 2010 and 2011. His research interests span the analysis of network traffic, QoE modeling and assessment, and data mining in networking. He has co-authored more than 80 papers (45 as first author) in major conferences and journals, and received 8 best paper/best workshop awards in the last seven years.

PIERDOMENICO FIADINO received his B.Sc. and M.Sc. degrees in computer engineering from University of Rome "La Sapienza", and a Ph.D. in electrical engineering from the Institute of Telecommunications of the Technical University of Wien, in 2008, 2010, and 2015 respectively. Since 2010, he is a researcher at FTW. His research interests cover network traffic monitoring and analysis, anomaly detection and diagnosis, machine learning, and data mining.

SARAH WASSERMANN is currently a first year M.Sc. student in computer science at the Université de Liège (ULg), Belgium, where she also earned her B.Sc. degree in 2015. Her main research interests lie in the field of network measurements.

STEFANO TRAVERSO [M'12] has research interests including privacy-preserving systems, network measurements, and content delivery networks. During his Ph.D. and post-doctoral work, he has visited Telefonica I+D research center, Barcelona, Spain;, NEC Laboratories, Heidelberg, Germany; and Alcatel-lucent Bell Labs, Paris, France. He is currently a post-doctoral fellow with the TNG group of Politecnico di Torino.

ALESSANDRO D'ALCONZO received his M.Sc. in electrical engineering and Ph.D. degree from Polytechnic of Bari, Italy, in 2003 and 2007, respectively. Since 2007 he has been a senior researcher at FTW. His research interests cover various topics in the field of telecommunication networks, including traffic monitoring and analysis, network measurements, security and privacy, and detection and diagnosis of network traffic anomalies in content delivery networks.

EDION TEGO received his Bachelor's degree from Polytechnic University of Tirana in 2008 and his Master's degree from the University of Rome "La Sapienza" in 2008, where he is also a Ph.D. candidate. His research interests are in SLA verification, traffic monitoring for anomaly detection and intrusion prevention, and software defined networking.

FRANCESCO MATERA received his Laurea degree in electronics engineering from the University of Rome"La Sapienza" in 1985. Since 1986, he has been a researcher at Fondazione Ugo Bordoni, where he currently works on networks and optical fiber systems. He was the scientific coordinator of the European IST ATLAS "All Optical Terabit per Second Lambda Shifted Transmission" project and has taken part in several other European projects.

MARCO MELLIA [S'08] has research interests in the area of traffic monitoring and analysis, cyber monitoring in general, and big data analytics. He has co-authored over 250 papers published in international journals and presented in leading international conferences. He won the IRTF ANR Prize at IETF-88, and best paper awards at IEEE P2P '12, ACM CoNEXT '13, and IEEE ICDCS '15. He is a member of the Editorial Board of *ACM/IEEE Transactions on Networking*, *IEEE Transactions on Network and Service Management*, and *ACM Computer Communication Review*. He is an associate professor at Politecnico di Torino, Italy.

Accurate and Affordable Packet-Train Testing Systems for Multi-Gigabit-per-Second Networks

Mario Ruiz, Javier Ramos, Gustavo Sutter, Jorge E. López de Vergara, Sergio López-Buedo, and Javier Aracil

Network testing is today more necessary than ever, because networking equipment is over-stressed due to the huge amount of traffic. However, network testing is becoming a complex and expensive task at this traffic load operating point.

ABSTRACT

Communication networks these days face a relentless increase in traffic load. Multi-gigabit-per-second links are becoming widespread, and network devices are under continuous stress, so testing whether they guarantee the specified throughput or delay is a must. Software-based solutions, such as packet-train traffic injection, were adequate for lower speeds, but they have become inaccurate in the current scenario. Hardware-based solutions have proved to be very accurate, but usually at the expense of much higher development and acquisition costs. Fortunately, new affordable FPGA SoC devices, as well as high-level synthesis tools, can very efficiently reduce these costs. In this article we show the advantages of hardware-based solutions in terms of accuracy, comparing the results obtained in an FPGA SoC development platform and in NetFPGA-10G to those of software. Results show that a hardware-based solution is significantly better, especially at 10 Gb/s. By leveraging high-level synthesis and open source platforms, prototypes were quickly developed. Noticeable advantages of our proposal are high accuracy, competitive cost with respect to the software counterpart, which runs in high-end off-the-shelf workstations, and the capability to easily evolve to upcoming 40 Gb/s and 100 Gb/s networks.

INTRODUCTION

Our day-to-day activities are becoming more and more dependent on communication networks: social networks, mobile apps, e-commerce, and so on. Such widespread use of communications has implications at both the access and server sides of the network. At the access side, it drives the development of faster access technologies, such as 10 gigabit passive optical networking (10G-PON) or fourth generation (4G) and beyond mobile networks. At the server side, it causes exponential increases in network traffic being faced by data centers. As a result, network testing is today more necessary than ever, because networking equipment is over-stressed due to the huge amount of traffic. However, network testing is becoming a complex and expensive task at this traffic load operating point.

We note that the main measurement parameters for assessing the quality of Internet access services are, according to [1], upload and download speed, packet loss rate (PLR), delay, and jitter. As link speed grows, measuring these parameters is not only more difficult, but also calls for testing devices that must provide unprecedented accuracy. For example, the transmission of a minimum-size Ethernet frame takes only 67 ns at 10 Gb/s. Apart from the difficulties of measuring at such small timescales, switching time is no longer negligible, making it imperative to include the switching equipment within the test scope as well.

In order to measure these network parameters on a device under test (DUT) such as a network switch or router, or a set of chained DUTs such as a network path, we propose using the packet-train technique. This technique has proven to be effective and highly immune to interference such as cross-traffic load at end-user equipment [2]. Unfortunately, current software tools are severely constrained when it comes to performing this type of measurement at high speeds (e.g., 10 Gb/s), even if they run at kernel level in the operating system.

In this article we show the shortcomings of software-based solutions and how to overcome such limitations using a hardware-based solution. The new field programmable gate array (FPGA) system on chip (SoC), which combines a powerful microcontroller and a programmable logic fabric, can be used to develop accurate and affordable testing devices for high-speed networks. Moreover, we also show how open-source platforms and high-level synthesis (HLS) can efficiently offset the well-known difficulties of FPGA development.

Certainly, the benefits of using FPGAs in multi-gigabit-per-second networks are well known. However, previous works in the network testing area are scarce and mainly focus on replacing configurable traffic generators [3, 4]. Probably the closest proposal to ours is [5], which uses FPGAs to test networking equipment according to RFC 2544 [6]. However, such work considers only 1 Gb/s networks, and no previous state-of-the-art work considers accurate time synchronization mechanisms such as GPS, which is needed to accurately measure one-way delay and jitter in a distributed way. Finally, the benefits of using high-level languages for networking applications are beginning to

The authors are with Universidad Autónoma de Madrid.

 0163-6804/16/$25.00 © 2016 IEEE

be recognized, as shown in [7] where such methodology has been used to implement a 10 Gb/s TCP/IP stack in FPGA.

Hence, we propose the use of novel FPGA SoCs as a means of comprehensively testing network equipment with packet trains. The most remarkable novelties of this work are: first of all, we show that very accurate results can be obtained at both 1 and 10 Gb/s, using two proof-of-concept designs. Second, FPGA SoCs can be used to implement very cost-effective testing appliances, featuring a minimal component footprint and reduced power consumption, which could easily be deployed across the whole network. Moreover, we also quantify how inaccurate software-based solutions can be at multi-gigabit-per-second speeds, unless the corresponding hardware aid comes into play. Finally, we also show the benefits of open source platforms and high-level synthesis in order to reduce FPGA development time and cost, thus making programmable logic competitive with software in terms of design productivity.

The rest of this article is structured as follows. First, we show two network testing use cases where the presented solutions have proven useful. Then the packet-train measurement technique is explained. Next, both the software and hardware approaches are described in detail. A performance evaluation follows, whereby both implementations are compared and discussed. Finally, some conclusions and an outlook of future work are provided.

Use Cases

The range of application of high-speed network testing tools is very diverse. The typical use case focuses on testing the capabilities of network equipment. However, this testing can also be extended to other scenarios, where it is necessary to perform distributed measurements along a network path. Moreover, such testing must be performed continuously in time to monitor the network QoS parameters. Next subsections provide two examples where this type of distributed continuous testing has been applied, apart from the usual testing of network equipment.

Service Level Verification

Nowadays we are witnessing fierce competition between operators to provide more bandwidth in the residential access link for the lowest possible price. In this competition for market share the regulators are playing their role as referees to enforce a given quality of service (QoS) level.

Of particular interest is the case of bandwidth reselling between operators at the access and metro link level. There, the regulator must ensure that the QoS provided by the incumbent operator to the hiring operator meets a QoS level that allows the transmission of interactive multimedia services. Since the number of potential users in the metro network is very large, and the residential bandwidth is growing at a significant pace, we note that the metro network switches, working at multi-gigabit-per-second data rates, must be carefully tested both before deployment and during operation.

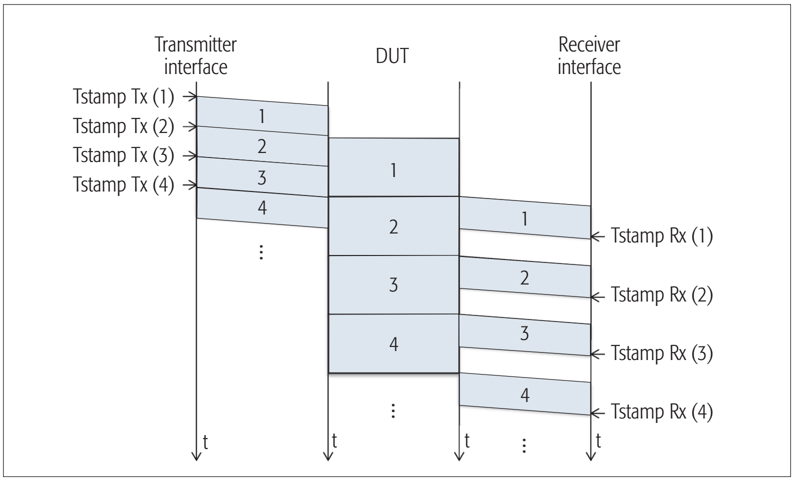

Figure 1. Representation of the packet-train technique.

Measurement of Next-Generation Elastic Optical Network Equipment

Elastic optical networks are being developed to offer the possibility of dynamically changing the signal modulation format and/or the spectrum allocation of optical data links. This dynamic reconfiguration capability paves the way for new operation models, whereby links are no longer statically provisioned, but dynamically adapted to traffic demands. Therefore, the link capacity must be continuously monitored in these elastic optical networks, to check if the underlying network has reconfigured the provided bandwidth or not.

Packet-Pair and Packet-Train Techniques for Network Measurements

Once we have motivated the need for high-speed network testing, we propose to use the packet-train technique, which is an evolution from the previous packet-pair technique. Packet-pair [8] is an active measurement method based on sending multiple packet pairs from a source to a destination endpoint in order to estimate the corresponding QoS parameters. Each pair is composed of equally sized packets, sent back to back at the maximum allowed speed in a link or end-to-end path. At the receiver side, packet dispersion is analyzed to estimate the capacity. As it turns out, packet-pair techniques are prone to both capacity underestimation and overestimation due to interfering traffic, because only two packets are used in the measurements. However, packet-trains [9] provide better accuracy and robustness, simply because more packets are involved in the measurement, and the resulting train is less sensitive to cross-traffic interference than the corresponding pair.

In packet-train techniques, a group of N packets is sent back to back from a sender to a receiver, and the average dispersion of the N packets is used to calculate the capacity, as shown in Fig. 1. Additionally, one-way delay (OWD), jitter, and PLR may also be estimated by including timestamps and sequence numbers on the packets. Increasing the number of packets in the train provides immunity against interfering traffic but also increases the measurement load in the mea-

Traditionally, network measurements have been performed using specialized hardware designed for such a task. In recent years, several software-based solutions that run on top of commercial off-the-shelf systems have also been applied for network measurement and testing tasks.

sured network. We note that this technique is based on flooding the link, and consequently, the measurement time must be kept at a minimum to not interfere with the rest of traffic. Typically, train lengths range from 100 to 1000 packets. Regarding packet sizes, OWD is better measured using minimum-sized packets to reduce the impact of the transmission time on the estimation.

Software-Based Solutions

Traditionally, network measurements have been performed using specialized hardware designed for such a task. In recent years, several software-based solutions that run on top of commercial off-the-shelf (COTS) systems have also been applied for network measurement and testing tasks. The latter provide a cost-effective and flexible solution for the development of network testing probes. For instance, pktgen [10] is a Linux kernel module that enables the generation of traffic with different packet headers and payloads defined by the user (source and destination medium access control, MAC, and IP addresses, UDP ports, etc.) and also with specific statistical features, that is, inter-arrival time and number of flows. Additionally, this kernel module adds a sequence number and the departure timestamp for each packet, which makes this tool suitable for throughput, OWD, jitter, and PLR measurements. The main drawback is that the departure timestamp is taken in the Linux kernel and not in the network interface card (NIC) itself, which adds measurement noise due to the transit time from the Linux kernel to the NIC. In high-speed links (10 Gb/s and beyond), we note that more packets per second must be copied to the kernel, so the measurement noise is more significant. Thus, all traditional software traffic generators cannot exactly mimic the transmission pattern defined by the user, which severely biases the measurement in a high-speed scenario, as stated in [11]. Even if a real-time operating system is used, the interruption timer accuracy is on the order of milliseconds, which is far too coarse for 10 Gb/s networks.

In addition, vanilla network drivers cannot cope with minimum-sized packets at 10 Gb/s rates in neither transmission nor reception, which is essential for testing. Recently, high-speed network engines have been developed [12] to solve this issue. For instance a software traffic generator called *PktGen-DPDK* built on top of Intel's DPDK[1] framework has recently been released. Such a traffic generator is able to transmit either random generated packets or PCAP traces. Although this traffic generator provides 10 Gb/s rates, it cannot add sequence numbers or transmission timestamps to the packets, so it is not able to measure OWD. For this reason, only the Linux pktgen module is considered later in our comparative analysis.

Hardware-Based Solutions

For years, the use of programmable logic devices (more specifically, FPGAs) has democratized hardware design for low-volume users. Nevertheless, the complexity of designing specialized hardware still resides in the FPGA design flow, which is based on demanding hardware description languages (HDLs). To circumvent this issue, several promising high-level synthesis (HLS) tools are appearing. HLS typically uses C/C++ for design entry, instead of the lower-level register transfer level (RTL) descriptions used by HDLs. HLS tools not only improve design productivity, but also bring FPGA technology closer to networking engineers.

In order to develop a hardware-based solution that implements the packet-train technique exploiting HLS design tools, we have worked with two hardware platforms. The first one is used to demonstrate the feasibility of building accurate, affordable, low-power, and portable 1 Gb/s network testing appliances on top of an FPGA SoC device. The second one, provided as a proof of concept for 10 Gb/s networks, is based on the NetFPGA[2] project. For simplicity, we name them HwP1 and HwP2, respectively.

HwP1: ZedBoard

ZedBoard[3] is a low-cost board based on a Xilinx Zynq SoC (XC7Z020-CLG484-1) device that encompasses in a single device a microcontroller based on a dual core ARM Cortex-A9 (referred to as the processing system, PS) and an FPGA (referred to as programmable logic, PL). The board has plenty of input/output connectors, with a standout FPGA mezzanine card (FMC) connector that allows plugging complex peripherals to the system. Furthermore, an operating system such as Linux can be run in the PS, thus enabling complex applications to be developed. Besides, the PL is based on a modern and fast FPGA technology (Xilinx 7-Series), which allows building complex hardware peripherals in the form of hardware modules (also known as IP-Cores, from intellectual property cores). To develop the project, two external Gigabit Ethernet interfaces were used, connected to the PL through the FMC connector.

HwP2: NetFPGA-10G

NetFPGA is an open hardware and software project developed by Stanford and Cambridge Universities in collaboration with Xilinx. It is intended for rapid prototyping of computer network devices. The NetFPGA-10G is based on a Xilinx Virtex-5 FPGA (XC5VTX240TFFG1759-2). It provides four SFP+ interfaces and has an 8X PCI Express Gen 1 interface to the host. Even though it can work as a standalone, it is typically attached to a host PC, as in our testbed. To ease the development process, we leveraged on the current NetFPGA environment, using the Open Source Network Tester (OSNT) project [13] as a development framework.

Hardware Architectures Description

With the aim of speeding up the hardware development cycle and bringing it as close as possible to network application engineers, we have used HLS tools to design the prototypes, as described in [14]. Our tool of choice is Vivado-HLS,[4] which generates synthesizable HDL code from a C/C++ source along with synthesis directives. On the other hand, modules where timing is critical (operations need to be done in an exact number of clock cycles) were implemented using the traditional FPGA design flow based on HDLs (VHDL or Verilog).

[1] http://dpdk.org/browse/apps/pktgen-dpdk/

[2] http://netfpga.org/

[3] http://zedboard.org/

[4] http://www.xilinx.com/products/design-tools/vivado/integration/esl-design.html

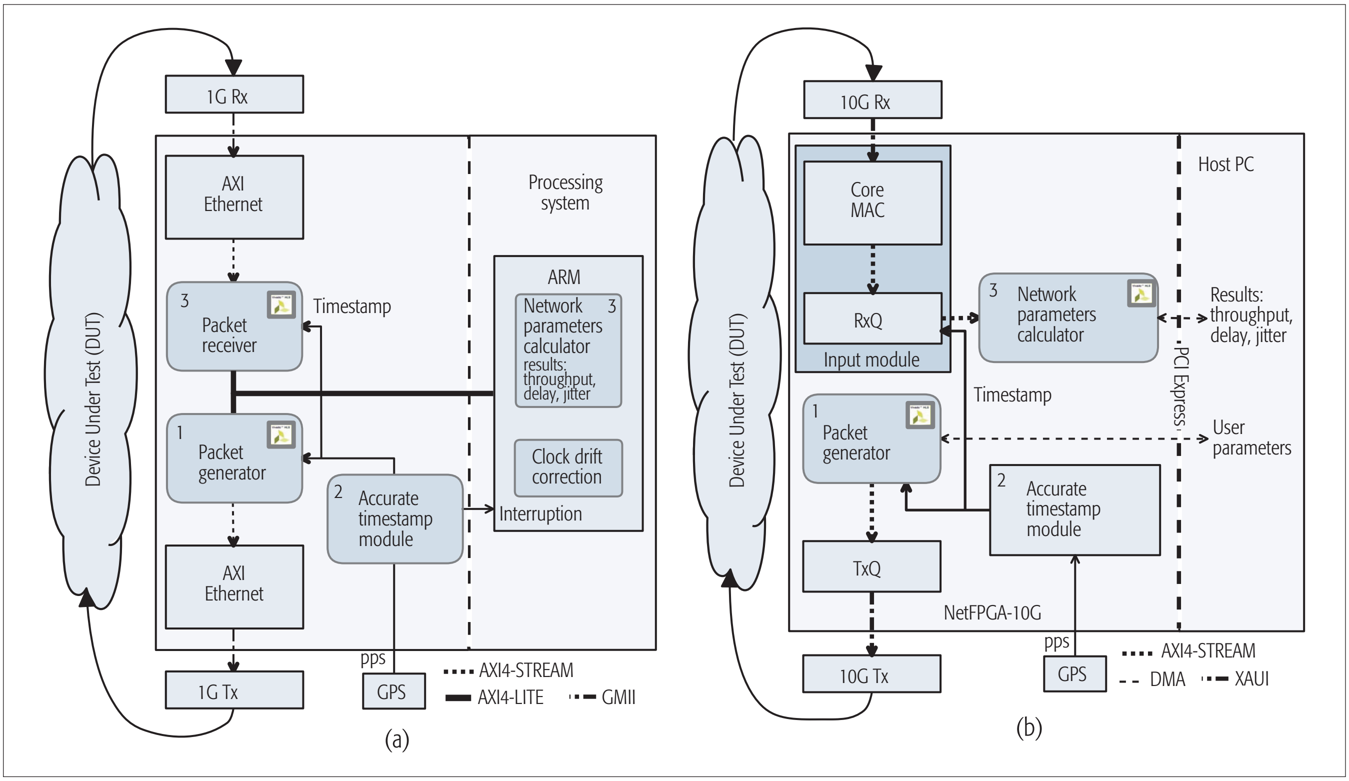

Figure 2. Implemented hardware architectures: a) ZedBoard at 1 Gb/s; b) NetFPGA and OSNT at 10 Gb/s.

The designs communicate with the external physical layer (PHY) chip by means of an AXI4-Stream bus. The difference between both prototypes is in the size of the bus transactions and frequency of operation (32 bits at 100 MHz for HwP1 and 256 bits at 156.25 MHz for HwP2).

In both architectures, depicted in Fig. 2, there are two key IP-Cores with similar behavior in the heart of the system. On one hand, the *packet generator* (1) was developed for both systems with the same HLS code, the only differences being the AXI4-Stream size and frequency of operation. Customizable options include source and destination MAC and IP addresses, UDP ports, packet size, and train length. In addition, each generated packet contains a sequence number, and it is timestamped as close as possible to the PHY chip. On the other hand, the *accurate timestamp module* (2) is in charge of correcting the clock drift. Both implementations use the pulse per second (PPS) signal from a GPS receiver as a reference to compensate the clock drift. At HwP1, this module is split in two parts following a hybrid hardware-software approach: the first part is a variable-rate counter implemented in hardware, which runs at 100 MHz. The second part is an algorithm that runs in the ARM processor. Such an algorithm uses the sum of the previous errors to correct the rate of the counter and sends it back to hardware. Therefore, we can obtain a remarkable timestamp resolution of 10 ns with an extremely low clock drift, thanks to the GPS-based error compensation. At HwP2, we have leveraged on the OSNT project functionality, which implements a direct digital synthesizer (DDS) to correct the clock drift [13]. The design operates at 156.25 MHz, with 6.4 ns resolution.

There are some differences in the architectures mainly because HwP1 has a tightly coupled processor near the FPGA (enabling a hybrid hardware-software approach, as mentioned before), and HwP2 provides a framework with a library of pre-designed modules. In HwP1, following the hybrid hardware-software approach, the *packet receiver* (3), developed in HLS, receives the packets and filters them according to user-defined rules. Then a software program running on the ARM processor computes the *network parameters*. On the other hand, HwP2 uses the infrastructure available at the NetFPGA framework for packet reception and a custom-developed HLS module to compute the network parameters (3).

In summary, our designs are good for sending packet trains with the aim of measuring the required quality parameters (delay, jitter, loss, and throughput). Additionally, both designs use a PPS signal from an external GPS in order to support one-way delay and jitter measurements when testing in a distributed infrastructure. The code of both projects is freely available at GitHub.[5]

In the near future, we expect to port our implementation to a low-cost 10 Gb/s version based on Zynq SoC. Finally, it is worth remarking that both hardware implementations occupy less than 55 percent of most of the available resources in the FPGA. Absolute figures of used resources in both solutions are shown in Table 1.

Performance Evaluation

Evaluation Testbeds

Two different performance evaluation scenarios have been considered for both software and hardware solutions. The first scenario, used for

[5] https://github.com/hpcn-uam/hardwarepackettrain

Feature	SW	HwP1	HwP2
Approximate cost	US$5000	US$1400	US$3500
Power consumption	160 W	8.5 W	120 W
FIFO blocks (FIFO36)	N/A	86 out of 140	11 out of 324
DSP blocks (DSP48E)	N/A	0 out of 220	2 out of 96
Number of slices	N/A	8210 out of 13,300	28,996 out of 37,440
Slice LUTs	N/A	22,224 out of 53,200	79,099 out of 149,760
Slice registers	N/A	24,589 out of 106,400	81,646 out of 149,760
Timestamp resolution	10 μs	10 ns	6.4 ns
Designability	Easy	Moderate	Moderate
Design time	Weeks	Months	Months
Engineer skills needed	Drivers	FPGA	FPGA
Maximum rate supported	10 Gb/s	1 Gb/s	10 Gb/s
Measure throughput 1 Gb/s	Poor	✓	✓
Measure throughput 10 Gb/s	×	N/A	✓
Measure OWD 1 Gb/s	×	✓	✓
Measure OWD 10 Gb/s	×	N/A	✓

Table 1. Features summary of software and hardware prototypes.

calibration, is based on sending measurement packets through an interface and receiving them in another interface of the same testing device in a loopback fashion. The second scenario is based on sending the measurement packets through an interface that is connected to the DUT — in our case, a Cisco Catalyst 2960-S. In the DUT, the measurement traffic is forwarded from one SFP+ port to another SFP+ port that is connected to the traffic receiver interface of the testing device. This scenario addresses the 10 Gb/s case. For the sake of completeness, the performance analysis has been repeated for the 1 Gb/s case, using the same setup but connected to 1 Gb/s DUT ports.

To evaluate the software solution, the pktgen module has been executed on a server running an Ubuntu Linux 14.04 with a 3.16.0 kernel. The server has two Intel Xeon E5-2620 processors with 6 cores each, 32 GB of RAM, and an Intel 82599 10 Gb/s NIC. For all tests the ixgbe vanilla driver has been used along with a *pktgen* 2.75 module.

All the experiments featured packet trains of 100 and 1000 packets with frame sizes of 60, 64, 128, 256, 512, 1024, and 1514 bytes, excluding frame preamble and check sequence. The experiment was repeated 10 times to obtain mean and standard deviation for throughput and OWD. In the case of the *pktgen* module, traffic has been generated using a single transmission queue since packet sequence numbers must be correlative for throughput and delay measurements, and using multiple queues produces packet disorder [15].

Experimental Results

The throughput measurements in the 1 Gb/s scenario using the testing setup with the DUT are shown in Fig. 3a. As can be observed, results obtained with the hardware HwP1 prototype are fairly close to theoretical throughput values, and the standard deviation is very small. In the case of measurements using the software approach, if the train length is 100 packets and the packet sizes are lower than 256 bytes, the results are pretty far from the theoretical values, and present larger deviations. With sizes of 256 bytes and above, we observe that the empirical values approach the theoretical ones. For 1000-packet trains, measurements are more accurate and present less deviation.

Figure 3b shows the throughput measurements for 10 Gb/s. As in the 1 Gb/s case, the results for the hardware systems are very similar to the theoretical values, with low deviation. However, the results for the software systems significantly depart from the theoretical ones, but improve as the packet size and train length increase. Similar results were obtained in the calibration setup, measuring the loopback for 1 and 10 Gb/s scenarios.

Note that for some scenarios, regulatory bodies require that link measurements are performed using minimum-size packets, which implies that software solutions are infeasible. Additionally, thorough network device testing should generate traffic with packet sizes ranging from the minimum size to the MTU.

Regarding OWD measurements, Table 2 shows the estimation results for both loopback and switch setups at 1 Gb/s and 10 Gb/s. As can be observed, software measurements are far from the theoretical values, adding up to 150 μs of error in the worst case scenario — loopback at 1 Gb/s. If accuracy below thousands of microseconds is needed, the hardware solution is the most suitable option. It is also worth noting that the hardware approach not only shows the most accurate results, but also presents extremely low variation (less than 0.01 percent), which makes it well suited for jitter measurements.

Calibration

We note that the measured latency in the hardware development is significantly larger than the theoretical minimum (frame transmission time) in the loopback scenario. This is due to the different elements in the transmission chain. In the 10 Gb/s case, the FPGA features three IP cores that add latency to transmission and reception: the 10G MAC core, 10G attachment unit interface (XAUI) core, and multi-gigabit transceivers (MGTs). In particular, the reception MGT has an elastic buffer in order to use the same clock for transmission and reception, so the design is simplified. Such an elastic buffer will also add uncertainty to the latency measurement. Moreover, we note that this buffer is not the only source of latency; the 10G MAC and XAUI cores can add up to 200 ns (adding transmission and reception latencies). Nevertheless, the main source of latency in the NetFPGA-10G board is the physical medium chip, which performs a conversion from XAUI

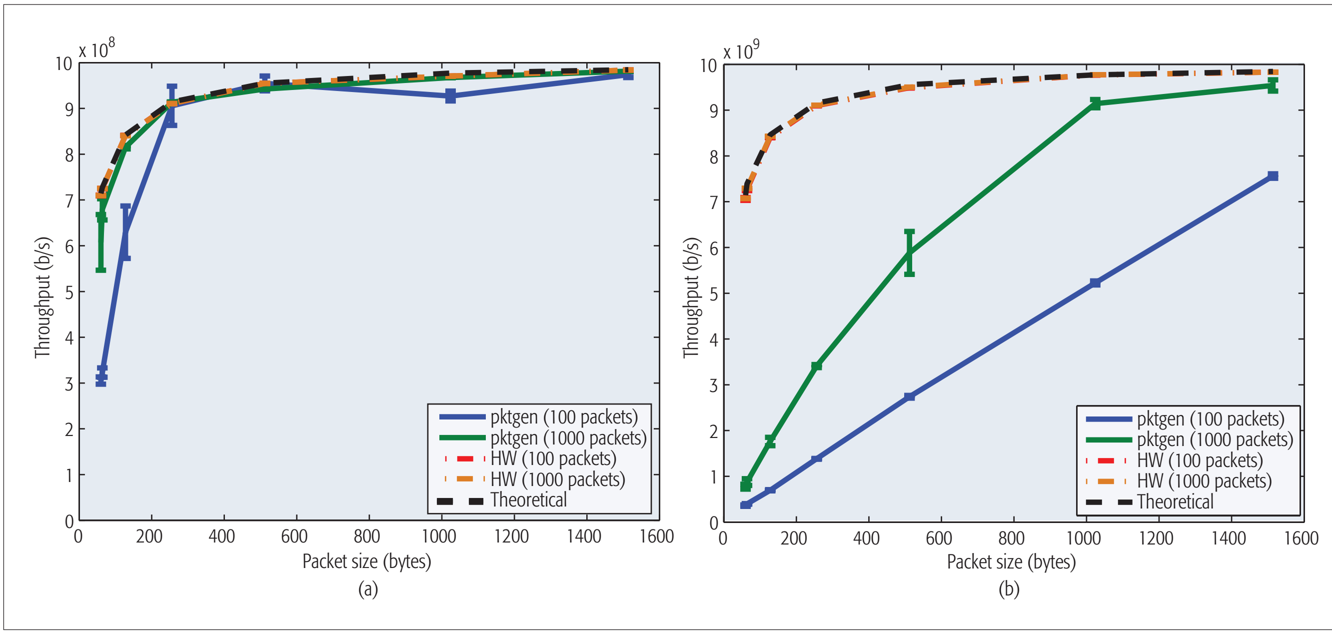

Figure 3. DUT throughput with different packet sizes. Mean and standard deviation estimated in both software and hardware: a) link speed 1 Gb/s; b) link speed 10 Gb/s.

		1 Gb/s				10 Gb/s			
		Loopback		Switch		Loopback		Switch	
Packet size (bytes)	# of packets	OWD *pktgen* (μs)	OWD HwP1 (μs)	OWD *pktgen* (μs)	OWD HwP1 (μs)	OWD *pktgen* (μs)	OWD HwP2 (μs)	OWD *pktgen* (μs)	OWD HwP2 (μs)
60	100	29 ± 5	1.890 ± 0.003	38 ± 14	5.149 ± 0.003	17 ± 6	0.883 ± 0.000	17 ± 8	3.79 ± 0.003
	1000	30 ± 7	1.889 ± 0.000	34 ± 5	5.149 ± 0.001	17 ± 5	0.883 ± 0.000	17 ± 6	3.80 ± 0.002
64	100	33 ± 9	1.941 ± 0.002	34 ± 6	5.236 ± 0.003	17 ± 7	0.883 ± 0.000	18 ± 5	3.81 ± 0.001
	1000	30 ± 5	1.945 ± 0.003	33 ± 4	5.235 ± 0.002	18 ± 6	0.884 ± 0.000	18 ± 6	3.80 ± 0.002
128	100	34 ± 5	2.772 ± 0.001	43 ± 6	6.720 ± 0.005	20 ± 5	0.945 ± 0.000	20 ± 3	3.97 ± 0.007
	1000	37 ± 5	2.773 ± 0.003	44 ± 6	6.728 ± 0.002	28 ± 1	0.946 ± 0.000	29 ± 2	3.97 ± 0.008
256	100	53 ± 8	4.437 ± 0.005	59 ± 8	9.425 ± 0.003	35 ± 8	1.069 ± 0.000	34 ± 7	4.20 ± 0.001
	1000	53 ± 8	4.438 ± 0.007	59 ± 7	9.426 ± 0.002	35 ± 8	1.069 ± 0.000	34 ± 9	4.20 ± 0.000
512	100	83 ± 13	7.763 ± 0.003	92 ± 13	14.793 ± 0.004	59 ± 12	1.317 ± 0.000	58 ± 13	4.65 ± 0.000
	1000	84 ± 11	7.765 ± 0.006	92 ± 11	14.796 ± 0.002	61 ± 11	1.317 ± 0.000	60 ± 10	4.65 ± 0.000
1024	100	122 ± 11	14.423 ± 0.002	134 ± 11	25.556 ± 0.003	90 ± 11	1.812 ± 0.000	89 ± 10	5.56 ± 0.001
	1000	123 ± 3	14.424 ± 0.002	135 ± 4	25.555 ± 0.002	91 ± 4	1.812 ± 0.000	92 ± 8	5.56 ± 0.000
1514	100	170 ± 17	20.798 ± 0.001	172 ± 47	35.854 ± 0.003	130 ± 16	2.322 ± 0.064	129 ± 15	6.48 ± 0.000
	1000	171 ± 6	20.796 ± 0.001	188 ± 9	35.853 ± 0.000	132 ± 6	2.317 ± 0.007	131 ± 9	6.48 ± 0.000

Table 2 . Switch and loopback estimated OWD with different packet sizes and link speeds. Mean and standard deviation.

to the 10 Gb/s serial electrical interface, as well as electronic dispersion compensation (EDC). Such operations add significant latency to the transmission/reception path. Similar considerations should be taken into account for the 1 Gb/s case.

Considering the results of Table 2, we can empirically infer that the aggregate delays (due to the reasons discussed above) linearly depend on the packet size. Therefore, we can use the loopback scenario to extract a calibration function from the delay measurements. To obtain such a function, we have first represented the scatter plot of measured data set (70 points) as shown in Figs. 4a for 1 Gb/s and 4b for 10 Gb/s. As a second step we have fitted the data using a linear regression, as it is the simplest method to calculate such a function. Finally, we have subtracted this function from the theoretical one to obtain the calibration function. This calibration function was later applied to the Catalyst 2960-S switch delay measurements, obtaining comparable results to those reported[6] for a similar device.

[6] http://miercom.com/pdf/reports/20130917.pdf

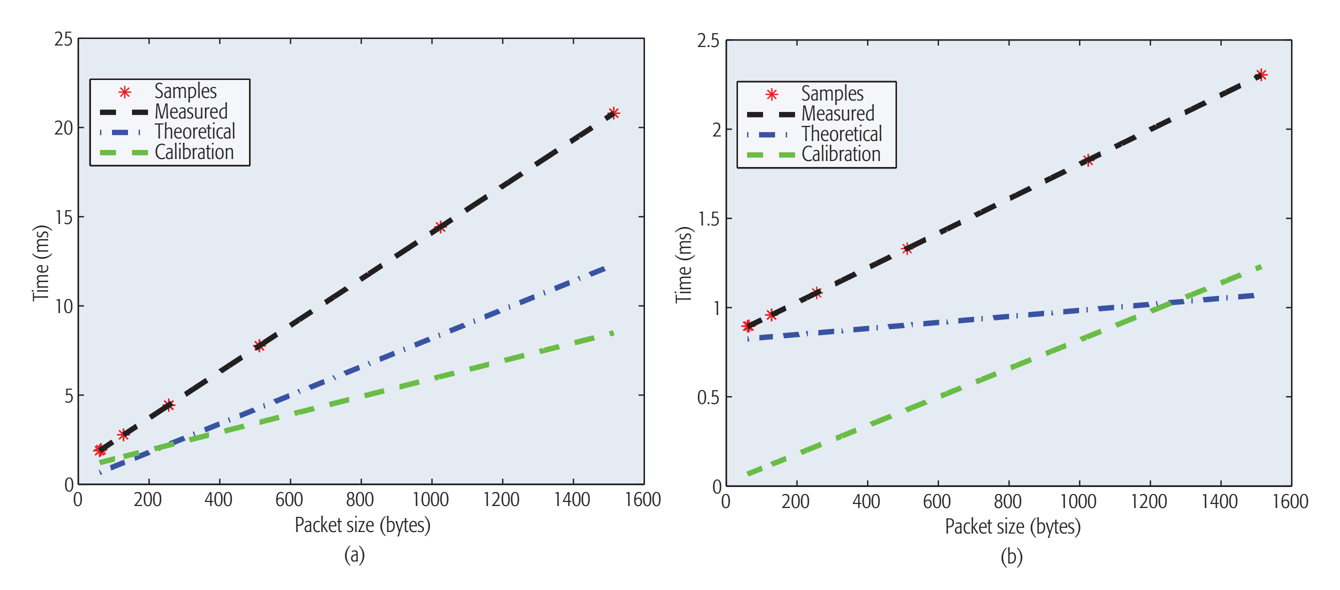

Figure 4. Regression on hardware platforms to calibrate delay measured: a) ZedBoard at 1 Gb/s; b) NetFPGA at 10 Gb/s.

Conclusions and Future Directions

The increasing speed of communication networks poses a serious challenge for testing. In this article we present the advantages of FPGA technology to implement accurate and affordable testing appliances for high-speed networks. Although software-based solutions are undoubtedly the most convenient approach in terms of deployment and development costs, we have shown that the non-determinism of software severely limits the accuracy of such solutions at multi-gigabit-per-second speeds. Moreover, non-determinism arises in the NIC itself and its connection to the system (PCIe, chipset), so using GP-GPU accelerators or multi/many-core architectures does not help to solve such an issue.

The solution proposed in this article is to use new FPGA SoC devices. On one hand, the FPGA fabric can be used to implement an accurate network measuring system. But accurate network measuring is not enough to create a full-fledged network testing appliance. A microprocessor on which to run the testing software and its underlying operating system is needed. This is where the high-performance microcontroller embedded in the FPGA SoC device plays an essential role, providing a hybrid hardware-software approach. On the other hand, FPGA SoCs only require a reduced number of external components in order to create a complete appliance, and its price and power consumption are low enough to consider massive deployment of these appliances throughout the whole network. For example, the price of a Xilinx XC7Z015 device is around US$100, and this device has the capability of implementing a 10 Gb/s Ethernet network port by means of an XAUI interface.

In order to assess the benefits of FPGA measuring systems for high-speed networks, we have developed two proof-of-concept designs. The first one runs at 1 Gb/s and is based on the ZedBoard development platform. The second one runs at 10 Gb/s on top of NetFPGA-10G. We have chosen the packet-train technique to perform the measurements due to its well-known features of accuracy, interference immunity, and low network overhead during testing. With the help of these two proof-of-concept designs, we have shown the advantages of hardware solutions in terms of determinism and accuracy when compared to the software alternative.

Table 1 provides a qualitative and quantitative summary of this hardware vs. software comparison. As made clear in Fig. 3, software solutions only provide good accuracy for throughput measurements at 1 Gb/s and when using large packet sizes. At 10 Gb/s the software accuracy is very poor, even if using kernel-level approaches such as the one evaluated in this article.

In Table 1, it is stated that the development time for the FPGA prototypes has been just months. It is well known that FPGA development is very costly, not uncommonly having development times as long as one year. Another contribution of this article is to present the benefits of open source platforms and high-level synthesis for improving FPGA design productivity. We have shown how high-level synthesis could significantly ease the development of the network parameter calculation core. Additionally, using an open source platform as a starting point, the development time of the 10 Gb/s prototype may be significantly reduced.

Regarding costs and power consumption, the results presented in Table 1 correspond to the prototypes that have been evaluated in this article. For the software solution, a high-end server system with a 10 Gb/s Ethernet card was used. For HwP1, the configured system includes the ZedBoard card, a GPS receiver, and the Gigabit Ethernet FMC card. For HwP2, a NetFPGA-10G card was used (academic price), along with a low-end computer attached to it. It is expected that such costs will be much lower when producing the network testing appliances on a large scale. Finally, the power consumption has been measured when performing the network measurements with packet trains. In HwP2, the standalone NetFPGA-10G consumes less than 30 W.

In the near future, we will see widespread use of 40 and 100 Gb/s networks. Given the results obtained for software solutions at 10 Gb/s, we envision that using dedicated hardware will be a must for testing such networks, especially if we consider the nanosecond accuracy required at those speeds. Fortunately, FPGA SoC devices will be able to implement these testing solutions based on dedicated hardware. At the moment, devices such as Xilinx XC7Z030 are capable of implementing both 10 and 40 Gb/s interfaces at a cost roughly below US$400. For the case of future 100 Gb/s networks, the new generation of Zynq UltraScale+ devices will provide a direct connection to CFP or QSFP28 cages, while at the same time maintaining the low power consumption and moderate price features of the current generation of FPGA SoC devices.

Acknowledgment

This work was partially supported by the Spanish Ministry of Economy and Competitiveness under the project PackTrack (TEC2012-33754) and by the European Union through the Integrated Project (IP) IDEALIST under grant agreement FP7-317999

References

[1] Body of Euro. Regulators for Electronic Commun., "Monitoring Quality of Internet Access Services in the Context of Net Neutrality," tech. rep. BoR (14) 117, Sept. 2014.

[2] J. Ramos *et al.*, "On the Effect of Concurrent Applications in Bandwidth Measurement Speedometers," *Computer Networks*, vol. 55, no. 6, Apr. 2011, pp. 1435–53.

[3] A. Tockhorn, P. Danielis, and D. Timmermann, "A Configurable FPGA-Based Traffic Generator for High-Performance Tests of Packet Processing Systems," *Proc. 6th Int'l. Conf. Internet Monitoring and Protection*, Mar. 2011, pp. 14–19.

[4] T. Groléat *et al*., "Flexible, Extensible, Open-Source and Affordable FPGA-Based Traffic Generator," Proc. First EditionWorkshop on High Performance and Programmable Networking, ser. HPPN '13, Jun. 2013, pp. 23–30.

[5] Y. Wang *et al*., "An FPGA-Based High-Speed Network Performance Measurement for RFC 2544," *EURASIP J. Wireless Commun. and Networking*, vol. 2015, no. 1, June 2015, p. 2.

[6] S. Bradner and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices," IETF RFC 2544 (informational), Mar. 1999.

[7] D. Sidler *et al*., "Scalable 10Gbps TCP/IP Stack Architecture for Reconfigurable Hardware," *Proc. 2015 IEEE 23rd Annual Int'l. Symp. Field-Programmable Custom Computing Machines*, May 2015, pp. 36–43.

[8] C. Dovrolis, P. Ramanathan, and D. Moore, "What do Packet Dispersion Techniques Measure?" *Proc. IEEE INFOCOM '01*, vol. 2, Jan. 2001, pp. 905–14.

[9] A. Johnsson, "On the Comparison of Packet-Pair and Packet-Train Measurements," *Proc. Swedish Nat'l. Computer Networking Wksp.*, Sept. 2003, pp. 241–50.

[10] R. Olsson, "Pktgen the Linux Packet Generator," *Proc. Linux Symp.*, Ottawa, Canada, vol. 2, July 2005, pp. 11–24.

[11] A. Botta, A. Dainotti, and A. Pescape, "Do You Trust Your Software-Based Traffic Generator?" *IEEE Commun. Mag.*, vol. 48, no. 9, Sept. 2010, pp. 158–65.

[12] V. Moreno *et al*., "Commodity Packet Capture Engines: Tutorial, Cookbook and Applicability," *IEEE Commun. Surveys & Tutorials*, vol. 17, no. 3, Aug. 2015, pp. 1364–90.

[13] G. Antichi *et al.*, "OSNT: Open Source Network Tester," *IEEE Network*, vol. 28, no. 5, Sept. 2014, pp. 6–12.

[14] M. Forconesi *et al*., "Bridging the Gap between Hardware and Software Open Source Network Developments," *IEEE Network*, vol. 28, no. 5, Sept. 2014, pp. 13–19.

[15] W. Wu, P. DeMar, and M. Crawford, "Why Can Some Advanced Ethernet NICs Cause Packet Reordering?" *IEEE Commun. Lett*., vol. 15, no. 2, Feb. 2011, pp. 253–55.

Biographies

Mario Ruiz (mario.ruiz@uam.es) is currently a Ph.D. student in the Computer and Telecommunication Engineering program at Universidad Autónoma de Madrid (UAM), Spain. He received his electronic engineering degree from Universidad Nacional de San Juan (UNSJ), Argentina, in 2014. He was a teaching assistant at UNSJ (2012–2014). His research interests include FPGA hardware design for high-speed networks, algorithm acceleration, and development of embedded systems. He is a co-author of a book on analog filters.

Javier Ramos (javier.ramos@uam.es) received his M.Sc. degree in computer science and Ph.D. degree in computer science and telecommunications from UAM in 2008 and 2013, respectively. Currently he works as a post-doctoral researcher at UAM. His research interests are in the analysis of network traffic, quality of service, software defined networks, and network function virtualization.

Gustavo Sutter (gustavo.sutter@uam.es) received his M.Sc. degree in computer science from UNCPBA University in Tandil, Buenos Aires, Argentina, in 1997, and his Ph.D. degree from UAM in 2005. He was an assistant professor at the UNCPBA and is currently an associate professor at UAM. His research interests include algorithms and networking in reconfigurable computing, digital arithmetic, and embedded and high-performance computing. He is the author of three books and more than 100 papers and communications.

Jorge López de Vergara (jorge.lopez_vergara@uam.es) is an associate professor at UAM and a founding partner of Naudit HPCN, a spin-off company devoted to high-performance traffic monitoring and analysis. He received his M.Sc. and Ph.D. degrees in telecommunication engineering from Universidad Politécnica de Madrid in 1998 and 2003, respectively. He performs research on network and service management and monitoring, having co-authored more than 100 scientific papers on this topic.

Sergio López-Buedo (sergio.lopez-buedo@uam.es) is an associate professor at UAM. FPGA technology is his main research interest, especially high-performance reconfigurable computing and communication applications. He has been a visiting researcher at the University of British Columbia (2005), George Washington University (2006–2007), and the University of Cambridge (2015). He holds 3 patents and has more than 50 publications, and is also a co-founder of the spin-off company Naudit HPCN.

Javier Aracil (javier.aracil@uam.es) received his M.Sc. and Ph.D. degrees (Honors) from Universidad Politécnica de Madrid in 1993 and 1995, both in telecommunications engineering, and his Licenciatura (five-year degree) in mathematics from UNED in 2009. In 1995 he was awarded a Fulbright scholarship to pursue post-doctoral research at the University of California, Berkeley. He is currently a full professor at UAM and a co-founder of the spin-off company Naudit HPCN.

In the near future, we will see widespread use of 40 and 100 Gb/s networks. Given the results obtained for software solutions at 10 Gb/s, we envision that using dedicated hardware will be a must for testing such networks, especially if we consider the nanosecond accuracy required at those speeds.

Observing Real Smartphone Applications over Multipath TCP

Quentin De Coninck, Matthieu Baerts, Benjamin Hesmans, and Olivier Bonaventure

A large fraction of smartphones have both cellular and WiFi interfaces. Despite this, smartphones rarely use them simultaneously because most of their data traffic is controlled by TCP, which can only use one interface at a time. Multipath TCP is a recently standardized TCP extension that solves this problem.

ABSTRACT

A large fraction of smartphones have both cellular and WiFi interfaces. Despite this, smartphones rarely use them simultaneously because most of their data traffic is controlled by TCP, which can only use one interface at a time. Multipath TCP is a recently standardized TCP extension that solves this problem. Smartphone vendors have started to deploy Multipath TCP, but its performance with real smartphone applications has not been studied in detail yet. To fill this gap, we port Multipath TCP on Android smartphones, and propose a framework to analyze the interactions between real network-heavy applications and this new protocol. We use eight popular Android applications and analyze their usage of the WiFi and cellular networks (especially 4G/LTE).

INTRODUCTION

Smartphones are the most popular mobile multihomed devices. Many users expect that their smartphones will be able to seamlessly use all available WiFi and cellular networks. Unfortunately, reality tells us that seamless coexistence between cellular and WiFi is not as simple as users would expect despite the huge investments in both cellular and WiFi networks by large network operators.

Several cellular/WiFi coexistence technologies have been proposed in recent years [1]. Some of them have been deployed. Recently, Multipath TCP [2] has received a lot of attention when it was selected by Apple to support its voice recognition (Siri) application. Siri leverages Multipath TCP to send voice samples over both WiFi and cellular interfaces to cope with various failure scenarios. As of this writing, Siri is the only deployed smartphone application that explicitly uses Multipath TCP. But there is no public information about the benefits of using Multipath TCP with it. In July 2015, Korea Telecom announced at Internet Engineering Task Force (IETF) 93 that they use Multipath TCP on Samsung Galaxy S6 smartphones to provide their users higher bandwidth.

Multipath TCP is a TCP extension that allows sending data from one end-to-end connection over different paths. On a smartphone, Multipath TCP allows applications to simultaneously send and receive data over both WiFi and cellular interfaces. It achieves this objective by establishing one TCP connection, called subflow in [2], over each interface. Once the subflows are established, data can be sent over any of the subflows thanks to the Multipath TCP scheduler. Researchers have analyzed the performance of Multipath TCP in such hybrid networks [3–6]. Their measurements show that Multipath TCP can indeed provide benefits by pooling network resources or enabling seamless handovers. However, these analyses were performed with bulk transfers between laptops and servers. As of this writing, no detailed analysis of the performance of real smartphone applications with Multipath TCP has been published.

We fill this gap in this article by presenting two main contributions that improve our understanding of the interactions between smartphone applications and the protocol stack. After a brief overview of Multipath TCP, we first propose a measurement methodology that automates user actions on Android smartphone applications. These actions trigger the creation of real connections. We then analyze how eight popular smartphone applications interact with Multipath TCP under different network conditions with both WiFi and cellular networks. Our measurements indicate that Multipath TCP works well with existing smartphone applications. Finally, we summarize the key lessons learned from this analysis.

MULTIPATH TCP AND RELATED WORK

Multipath TCP is a recent TCP extension that enables the transmission of the data belonging to one connection over different paths or interfaces [2]. A Multipath TCP connection is a logical association that provides a bytestream service. Compared to other multipath transport layer solutions such as SCTP, Multipath TCP can be deployed on TCP-compatible networks. To request the utilization of Multipath TCP, the smartphone adds the `MP_CAPABLE` option in the `SYN` segment sent over its `default` interface (e.g., WiFi). This option contains some flags and a key [2]. If the server supports Multipath TCP, it includes its key in the `MP_CAPABLE` option sent in the `SYN+ACK`. According to the Multipath TCP terminology, this TCP connection is called the initial subflow [2]. The smartphone can use it to exchange data over the WiFi interface. If the smartphone also wants to send data

Quentin De Coninck, Benjamin Hesmans, and Olivier Bonaventure are with Universite Catholique de Louvain; Matthieu Baerts is with Tessares.

0163-6804/16/$25.00 © 2016 IEEE

for this connection over its cellular interface, it sends a new SYN segment with the MP_JOIN option over this interface. This option contains a token derived from the key announced by the server in the MP_CAPABLE option. This token identifies the Multipath TCP connection on the server side. The server replies with a SYN+ACK containing the MP_JOIN option, and the second subflow is established. Multipath TCP sends data over any of the available subflows. Two levels of sequence numbers are used by Multipath TCP: the regular TCP sequence number and the data sequence number (DSN). The DSN is associated with the bytestream. When data is sent over a subflow, its DSN is mapped to the regular sequence numbers with the DSS option that also contains DSN acknowledgments. When losses occur, Multipath TCP can retransmit data over a different subflow. To achieve that, the device sends on another subflow a packet with the same DSN containing data. This operation is called a reinjection [7]. Although at the subflow level it looks like a new packet, a reinjection can be detected by looking at its DSN to see if it was previously sent on another subflow.

The operation of a Multipath TCP implementation depends on several algorithms that are not standardized by the IETF. First, the *path manager* defines the strategy used to create and delete subflows. Second, the *packet scheduler* [8] selects, among the active subflows that have an open congestion window, the subflow that will be used to send the data.

Various researchers have analyzed the performance of Multipath TCP through measurements. Raiciu *et al.* [9] discuss how Multipath TCP can be used to support mobile devices and provide early measurement results. Chen *et al*. [4] analyze the performance of Multipath TCP in WiFi/cellular networks by using bulk transfer applications running on laptops. Deng *et al*. [6] compare the performance of single-path TCP over WiFi and LTE networks with Multipath TCP on multi-homed devices by using active measurements and replaying HTTP traffic observed on mobile applications. They show that Multipath TCP provides benefits for long flows but not for short ones, for which the selection of the interface for the initial subflow is important from a performance viewpoint.

Multipath TCP on Android Smartphones

Several backports of the Multipath TCP kernel on Android smartphones were released in recent years. However, these ports were often based on old versions of the Multipath TCP kernel. For this work, we rely on a backport of the latest version, 0.89v5, of the Multipath TCP Linux kernel[1] on a Nexus 5 running Android 4.4.4. It should be noted that the Linux kernel used on such Android devices is tweaked to use only one interface at a time. We disable this function and configure the kernel to be able to simultaneously use two interfaces. The Multipath TCP kernel controls the utilization of the available interfaces thanks to a path manager. We use the Full Mesh path manager, which creates a subflow over all network interfaces for each established TCP connection. To spread packets over the available paths, we use the default round-trip time (RTT)-based scheduler [8], which sends packets over the available path with the lowest RTT.

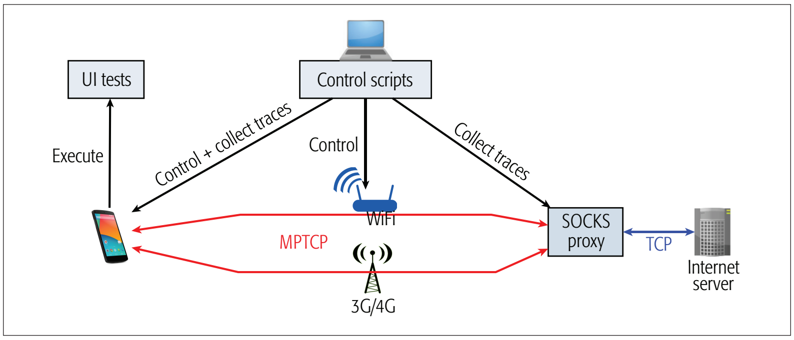

Figure 1. High-level view of the test framework.

Most popular smartphone applications use TCP to interact with servers managed by the application developers. As of this writing, it has not been possible to convince them to install Multipath TCP on their servers. To overcome this issue, we configure the smartphone to use a Multipath-TCP-capable SOCKS proxy server for all its connections, as shown in Fig. 1. This is exactly the same setup as the one launched commercially in Korea in June 2015. Each (Multipath) TCP connection initiated by the smartphone is thus redirected to, and terminated at, the proxy server. The proxy server then establishes a regular TCP connection to the server. Thanks to this setup, the smartphone can use Multipath TCP over the cellular and WiFi interfaces while interacting with legacy servers via the proxy.

The SOCKS server itself uses ShadowSocks and is configured to use the minimum encryption scheme to reduce the overhead. The other settings are set to the recommended values.[2] On the smartphone, we use the standard Android ShadowSocks client.

Automating Measurements

In order to collect a large number of measurements, we developed a test framework that automates the interactions with these applications.[3] A high-level overview is shown in Fig. 1. On this basis, we identify two main tasks: *controlling devices* and *mimicking user interaction*.

The devices are controlled by Python and shell scripts (3100 lines split into different modules). Our controller checks the availability of the smartphones and wireless networks, collects packet traces, and modifies settings such as the protocol (either TCP or Multipath TCP) and interfaces (WiFi, cellular, or both) used by the smartphone.

It was designed to be reusable, modular using parameters, and to cope with unexpected situations caused by the unreliability of this kind of device.

User interactions are simulated through application user interface (UI) tests to produce each high-level scenario. Each of the eight selected applications has its own UI test. These UI tests are implemented by using the MonkeyRunner Android UI testing tool. Each unit test is implemented as a new program, and all of them use a shared Utils class. Thanks to this class, our framework allows a scenario to be built with less than a few hundred lines of code.

[1] http://www.multipath-tcp.org

[2] http://shadowsocks.org/en/config/advanced.html

[3] Results are reproducible, instructions are publicly available: http://github.com/MPTCP-smartphone-thesis/uitests

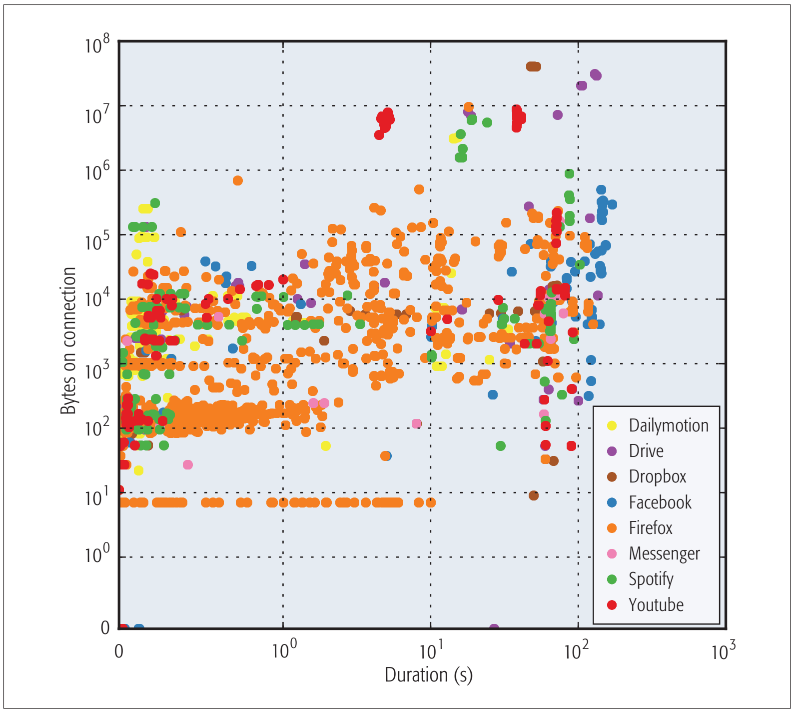

Figure 2. Duration and data transferred by the smartphone applications.

Each test was designed to resist different unusual situations, such as failure of the smartphone, failure of one of the wireless networks, or an unexpected reaction of the application. The measurements presented in this article were performed with the versions of the applications released on November 15, 2014. To avoid network optimization and have repeatable measures, cached files are deleted when launching our tests.

All the tests described in this article were performed during the night to reduce interference with other users on the networks. The WiFi network was provided by a controlled router with an `802.11n` interface on the 5 GHz frequency band with a bit rate of 65 to 72 Mb/s. The router was connected with a 100 Mb/s link to the university network. We ensured that no other WiFi network was emitting in this frequency band in the building. The cellular network is a commercial one, and we configured either 3G or 4G on the smartphone. The test scenarios were run in a random order each day to limit correlation of the results with the time at which they were launched.

Test Scenarios

Now, we provide an overview of the scenarios used to generate network traffic. Our test scenarios can be split into two categories: upload-intensive scenarios and download-intensive scenarios. Each test takes less than 120 s.

Upload-Intensive: We first consider two interactive applications: Facebook and Messenger. With the Facebook application, our test first updates the news feed, then writes a new status, takes and shares a new photo with a description, and finally performs a new check out status. With Messenger, it sends a text message, then puts an emoticon, and finally sends a new photo. Then we consider two cloud storage applications: Dropbox and Google Drive. For both, we create a fresh file containing 20 MB of purely random data and upload it.

Download-Intensive: First, we use Firefox to browse the main page of the top 12 Alexa web sites with an empty cache. Our second application is Spotify. This is a music delivery application. The test plays new music (shuffle play feature) for 75 s. Finally, we consider two popular video streaming applications: Dailymotion and Youtube. For both applications, we play three different videos in the same order and watch them for 25 s. Those videos are available in HD, and we fetch the best possible quality even when using cellular networks.

We used these applications on the testbed shown in Fig. 1. This setup allows us to capture all the packets sent by both the smartphone and the SOCKS server. We captured more than 110,000 connections over about 1400 different tests conducted in February and March 2015 carrying more than 15 GB of data. The entire dataset is publicly available.[4]

Measurements

We use our test framework to analyze the interactions between smartphone applications and the network under various conditions. We first observe our applications over regular TCP, then we study how they behave over Multipath TCP. We use `tstat` [10] and `mptcptrace` [11] to extract information from packet traces.

Single-Path Measurements

The selected applications interact in different ways with the underlying transport protocol. An important factor that influences the performance of TCP is the lifetime of the connections and the number of bytes that are exchanged. To study this factor, we analyze the TCP connections established by our studied applications. Figure 2 shows that they create different types of TCP connections. Each point on this figure represents one captured TCP connection. The x-axis (in logarithmic scale) is the connection duration in seconds, while the y-axis is the number of bytes exchanged on the connection. Firefox is clearly the application that uses the largest number of connections (63.9 percent of all connections), which is not surprising given that our Firefox scenario contacts the 12 top Alexa web sites. Unsurprisingly, streaming and cloud storage scenarios with Dropbox (31.75 percent), Youtube (29.7 percent), Drive (19.9 percent), Dailymotion (9.6 percent), and Spotify (5 percent) are the applications that exchange the largest volume in bytes. On the other hand, our Facebook scenario generates long TCP connections that do not exchange too many bytes.

Some of the connections that we observe are caused by the utilization of a SOCKS proxy. There are hundreds of connections that last up to tens of seconds but only transfer seven bytes of data.

After investigation, Firefox preventively opens new TCP connections but sometimes never uses them. The seven exchanged bytes correspond to the command sent by the SOCKS client. This command contains the IPv4 address and destination port used by the SOCKS proxy to establish the regular TCP connection to the remote servers. Most of the short connections that only transfer about 100 bytes are DNS requests that are sent over TCP by the SOCKS client.

[4] http://multipath-tcp.org/data/IEEEComMag16

Our connections can be categorized into three types:

- Short connections carrying a relatively small amount of data
- Long connections carrying most of the data
- Long-lived connections carrying a small amount of data

In our tests, 74 percent of the connections last less than 1 s. Among the connections that last more than 1 s, 32 percent carry more than 10 kB and represent 98.6 percent of the overall volume. Finally, the remaining 68 percent of the connections that last more than 1 s exchange less than 10 kB of data. This tends to match many measurement studies which identified that most TCP connections are short, and most of the traffic is carried by a small fraction of all TCP connections [12].

The RTT is one of the key factors that influence the performance of TCP connections. We used `tstat` to compute the average RTT for each of the captured TCP connections. Figure 3 provides the CDF of the RTT measures among all the TCP connections used in the upstream (data sent by the smartphone) and downstream directions. The 4G network exhibits an RTT in upstream with a median of 42.6 ms and a mean of 50 ms. In the downstream direction, the median RTT increases up to 38.1 ms. On the WiFi network, 60 percent of the connections have an RTT shorter than 15.4 ms. Unsurprisingly, there is some bufferbloat on the 3G network, mainly in the upstream direction, but the bufferbloat remains reasonable compared to other networks [13].

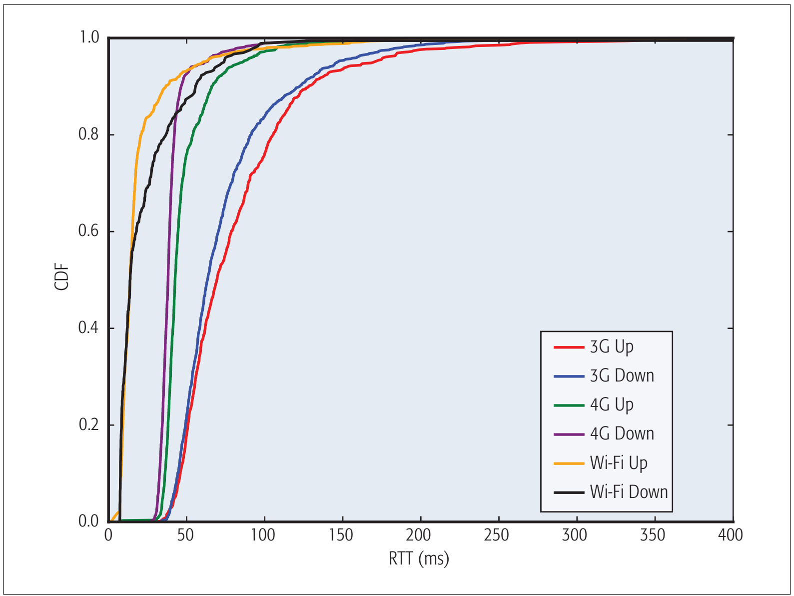

Figure 3. Average round-trip time of the TCP connections over WiFi, 3G, and 4G networks.

Multipath Measurements

The previous section showed that our measurement scenarios cover different utilizations of TCP. We now enable Multipath TCP on our smartphone and perform the same measurements to understand how our eight applications interact with Multipath TCP. The first, but important, point to be noted is that we did not observe any incompatibility between the applications and Multipath TCP.

Multipath TCP can be used in different modes [3] on smartphones. For our measurements, we focus on a configuration where Multipath TCP tries to pool the resources of the cellular and WiFi interfaces simultaneously since the handover and backup performance has already been studied in [3].

When a 4G and a WiFi interface are pooled together, it is interesting to analyze which fraction of the traffic is sent over which interface. With the Multipath TCP implementation in the Linux kernel, this fraction depends on the interactions between the congestion control scheme, the packet scheduler, the underlying networks, and the application.

We first consider Multipath TCP connections using WiFi and 4G interfaces, with WiFi set as the `default` interface. In Fig. 4a, each point corresponds to one Multipath TCP connection, and the *x* axis indicates the number of bytes transfered by this connection from the smartphone to servers. Although we observe connections using both WiFi and cellular interfaces, Fig. 4b shows that 96 percent of the connections only use the WiFi interface. However, Fig. 4c indicates that those connections are small since they carry only 16.3 percent of all the data bytes contained in the considered connections.

Several factors explain why Multipath TCP does not use the cellular network for these short connections. The first factor is the configured `default` route. When an application initiates a connection, Multipath TCP sends the `SYN` over the interface with the `default` route, in our case the WiFi interface. This is the standard configuration of Android smartphones that prefer the WiFi interface when it is active. If the Multipath TCP connection is short and only transfers a few kilobytes or less, most of the data fits inside the initial congestion window and can be sent over the WiFi interface, while the second subflow is established over the cellular interface. 71 percent of the connections sending only on the WiFi interface are in this case. Furthermore, the RTT over the WiFi interface is shorter than over the cellular interface. This implies that most of the time, as long as the congestion window is open over the WiFi interface, Multipath TCP's RTT-based scheduler [8] prefers to send packets over the WiFi interface. Indeed, 84 percent of the connections with both subflows established have a smaller average RTT on WiFi than on 4G.

Those factors explain why data on the short connections are exchanged only over the WiFi interface. We experimentally verified this by performing the same set of measurements with the `default` route pointed to the 4G interface. Figure 5 shows that with this configuration most short connections still exclusively use the 4G network (label 1, Fig. 5), but this concerns only 65 percent of all connections. It seems that even if cellular is the `default` interface, many connections still mainly use WiFi, even for connections exchanging less than 1 kB. This occurs for connections that do not push data as fast as possible. If the connection lasts more than two RTTs, Multipath TCP has enough time to establish the second subflow. The packet scheduler will then select the subflow with the lowest RTT — 88 percent of connections using both subflows have a WiFi subflow with a lower average RTT than the cellular one.

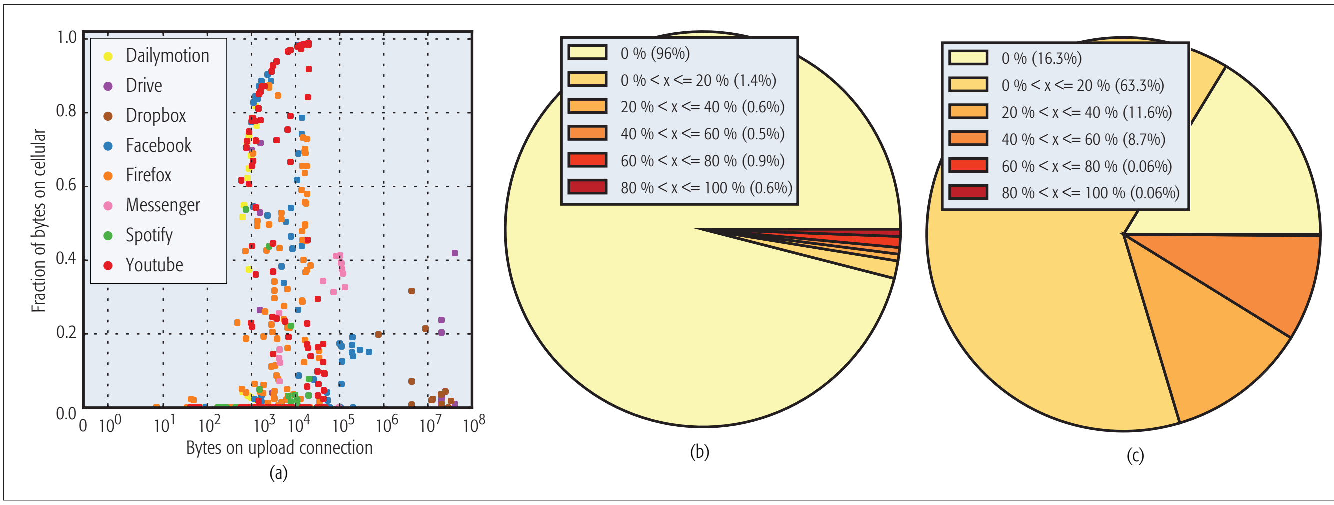

Figure 4. When the default route points to the WiFi interface, Multipath TCP mainly uses this interface for short connections: a) fraction of data bytes sent on a connection depending on its data size; b) connections classified by the percentage of data sent on the cellular interface; c) data bytes classified by the percentage of data sent on the cellular interface on the connection to which it belongs.

This explains the bottom of Fig. 5 (annotated as 2): a group of Firefox connections that transfer less than 10 kB nearly exclusively use the WiFi interface. A closer look at the packet trace reveals that these connections are part of the connection pool managed by Firefox. This behavior does not happen with other applications. When Firefox creates a connection in this pool, the initial handshake and the SOCKS command to our SOCKS server are sent. These packets are exchanged over the cellular interface, and Firefox does not immediately send data over the established connection. This leaves enough time for Multipath TCP to create the subflow over the WiFi interface and measure its RTT. When Firefox starts transmitting data over such a connection, the RTT-based scheduler used by Multipath TCP prefers the WiFi subflow, and no data (except the initial SOCKS command) is sent over the cellular subflow.

When the applications push more data over the Multipath TCP connection, the distribution of the traffic between the cellular and WiFi interfaces also depends on the evolution of the congestion windows over the two subflows. If the application pushes data at a low rate, the packet scheduler will send it over the lowest-RTT interface (WiFi in this case). However, this distribution can be fragile. If one packet is lost, the congestion window is reduced, and the next data might be sent over the other interface. If the application pushes data at a higher rate, the congestion window over the lowest-RTT interface is not large enough, and the packet scheduler will send data over the second subflow.

In some cases, data transfered by Multipath TCP on one flow may be retransmitted again on another flow. This phenomenon is called reinjection [7] and might limit the performance of Multipath TCP in some circumstances [14]. We used `mptcptrace` to compute the reinjections over all observed Multipath TCP connections. In our experiments (WiFi and 4G), reinjections in the upstream direction were rare (less than 0.5 percent of all connections include a reinjection) and short (no more than 5 kB are reinjected on a connection). Looking at the proxy traces in the downstream direction, reinjections are observed on only 2 percent of all connections, and the largest observed reinjection is 30 kB on a 5 MB connection. This overhead is thus low.

An important benefit of the resource pooling capabilities of Multipath TCP is its ability to adapt to various networking conditions. When a smartphone moves, the performance of the WiFi and cellular interfaces often vary. Previous work with bulk transfer applications has shown that Multipath TCP can adapt to heterogeneous networks having different bandwidths and delays [15]. Our measurement framework also allows exploration of the performance of smartphone applications under various network conditions. As an illustration, we analyze the packet traces collected when the smartphone is uploading a file with Dropbox. We first consider a WiFi access point attached to a digital subscriber line (DSL) router having 1 Mb/s of upstream bandwidth and 15 Mb/s of downstream bandwidth. When the smartphone is attached to both this WiFi access point and the 4G network, it sends on average 91 percent of the data over the 4G network. This is expected because although the WiFi has better RTT, the congestion window of this path is quickly full and empties slowly. In that case, the Multipath TCP scheduler selects the next available subflow with the lowest RTT — here the cellular interface. Since the cellular network offers a larger bandwidth, Multipath TCP can take advantage of it and thus avoids being trapped in a low-performance network for big connections.

As a second test case, we consider our standard WiFi access (around 70 Mb/s in both streams) and the 4G network with bandwidth limited down to a few hundred kilobits per second. This is the shaping enforced by our cellular network once we reach the monthly traffic volume quota. In this case, 98.8 percent of the bytes are sent over the WiFi interface.

Conclusion

Multipath TCP is a new TCP extension that has strong potential on smartphones, as shown by its recent adoption by Apple and Korea Telecom. By enabling TCP connections to exchange data over cellular and WiFi interfaces, it brings new possibilities to improve the user experience. Apple's deployment focused on a single use case, and little is currently known about the interactions between real smartphone applications and Multipath TCP. In this article, we have proposed and implemented a measurement testing framework that enables researchers to conduct reproducible experiments with traffic generated by real applications.

We have used our measurement framework to study the interactions between eight very different smartphone applications covering several smartphone use cases and the latest version of the Multipath TCP implementation in the Linux kernel. Several lessons have already been learned from a first analysis of the packet traces that were captured. First, all the studied applications work without any modification with Multipath TCP. This confirms that Multipath TCP is compatible with existing applications. Second, for the short connections often used by the studied applications, Multipath TCP uses the `default` route to forward the data for most connections. As suggested in [6], we confirm that the selection of this default route is thus an important decision on the smartphone. Third, for long connections, Multipath TCP enables the applications to pool the bandwidth on the cellular and WiFi interfaces, and maintains good performance when one of them has bandwidth restrictions. This is important for the user's experience given that smartphones often associate with wireless networks by relying on metrics like signal-to-noise ratio.

We expect that our framework and the collected packet traces will be beneficial to Multipath TCP researchers and implementers by enabling them to study how improvements to the implementation would affect real applications in a reproducible manner. Moreover, this framework could also be used to measure the energy consumption impact of Multipath TCP on mobile devices like smartphones.

Acknowledgment

This work was partially supported by the EC within the FP7 Trilogy2 project. We would like to thank Gregory Detal and Sebastien Barre for the port of the latest Multipath TCP Linux kernel on the Nexus 5, and Patrick Delcoigne and his team for the cellular measurements.

References

[1] F. Rebecchi *et al.*, "Data Offloading Techniques In Cellular Networks: A Survey," *IEEE Commun. Surveys & Tutorials*, vol. 17, no. 2, 2nd qtr. 2015, pp. 580–603.

[2] A. Ford *et al.*, "TCP Extensions for Multipath Operation with Multiple Addresses," IETF RFC 6824, Jan. 2013.

[3] C. Paasch *et al.*, "Exploring Mobile/WiFi Handover with Multipath TCP," *Proc. ACM SIGCOMM CellNet Wksp.*, 2012, pp. 31–36.

[4] Y.-C. Chen *et al.*, "A Measurement-Based Study of Multipath TCP Performance over Wireless Networks," *Proc. 2013 Conf. Internet Measurement*, ACM, 2013, pp. 455–68.

[5] Y.-S. Lim *et al.*, "How Green is Multipath TCP for Mobile Devices?" *Proc. 4th Wksp. All Things Cellular: Operations, Applications, & Challenges*, ACM, 2014, pp. 3–8.

[6] S. Deng *et al.*, "Wifi, LTE, or Both? Measuring Multi-Homed Wireless Internet Performance," *Proc. 2014 Conf. Internet Measurement*, ACM, 2014, pp. 181–94.

[7] C. Raiciu *et al.*, "How Hard Can it Be? Designing and Implementing a Deployable Multipath TCP," *Proc. 9th USENIX Conf. Networked Sys. Design and Implementation*, 2012, pp. 29–42.

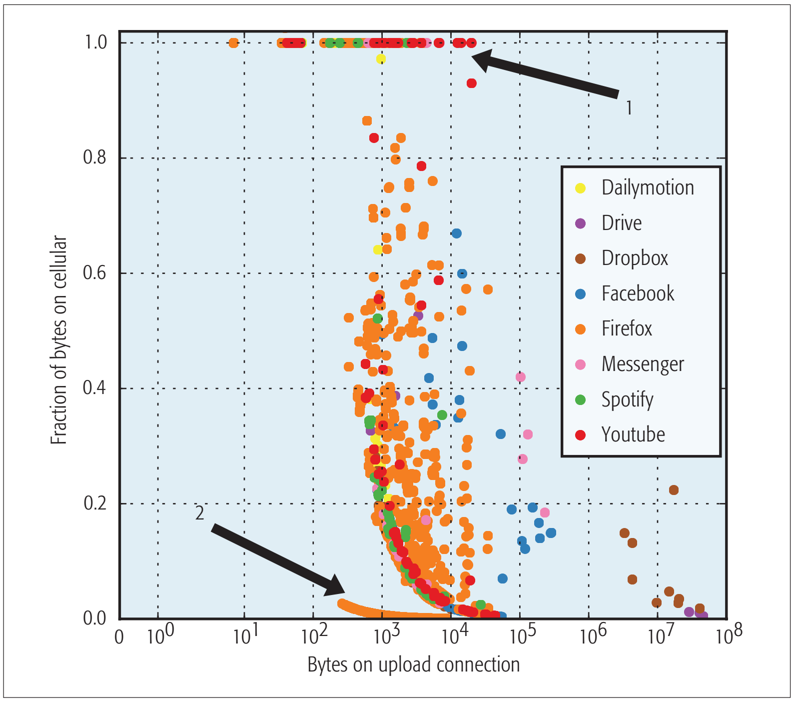

Figure 5. When the default route points to the cellular interface, many connections are tried by the WiFi interface.

[8] C. Paasch *et al.*, "Experimental Evaluation of Multipath TCP Schedulers," *Proc. 2014 ACM SIGCOMM Wksp. Capacity Sharing*, 2014, pp. 27–32.

[9] C. Raiciu *et al.*, "Opportunistic Mobility with Multipath TCP," *Proc. ACM MobiArch 2011*, 2011.

[10] M. Mellia, A. Carpani, and R. L. Cigno, "Tstat: TCP statistic and analysis Tool," *Quality of Service in Multiservice IP Networks*, Springer, 2003, pp. 145–57.

[11] B. Hesmans and O. Bonaventure, "Tracing Multipath TCP Connections," *Proc. 2014 ACM Conf. SIGCOMM*, 2014, pp. 361–62.

[12] A. Gember, A. Anand, and A. Akella, "A Comparative Study of Handheld and Non-Handheld Traffic in Campus Wi-Fi Networks," *Passive and Active Measurement*, Springer, 2011, pp. 173–83.

[13] S. Ferlin-Oliveira, T. Dreibholz, and O. Alay, "Tackling the Challenge of Bufferbloat in Multi-Path Transport over Heterogeneous Wireless Networks," *Proc. 2014 22nd IEEE Int'l. Symp. QoS*, 2014, IEEE, pp. 123–28.

[14] Y.-S. Lim *et al.*, "Cross-Layer Path Management in Multi-Path Transport Protocol for Mobile Devices," *Proc. 2014 IEEE INFOCOM*, 2014, pp. 1815–23.

[15] C. Paasch, R. Khalili, and O. Bonaventure, "On the Benefits of Applying Experimental Design to Improve Multipath TCP," *Proc. 9th ACM Conf. Emerging Networking Experiments and Technologies*, 2013, pp. 393–98.

Biographies

QUENTIN DE CONINCK (quentin.deconinck@uclouvain.be) received his B.Eng. and M.Eng. degrees in computer engineering from Université Catholique de Louvain, Belgium, in 2013 and 2015, respectively. He is now pursuing a Ph.D. at the same institution. His research interests include applications of mobility in computer networks and low-level system architecture.

MATTHIEU BAERTS (matthieu.baerts@tessares.net) received his M. Sc. degree in computer science in 2015 from Université Catholique de Louvain. He is now working with Tessares s.a., which is currently developing hybrid Internet access bonding solutions using Multipath TCP. He is also involved in free software communities, hoping to contribute to technologies available for everyone.

BENJAMIN HESMANS (benjamin.hesmans@uclouvain.be) is a third-year Ph.D. student at Université Catholique de Louvain. His main interest is Multipath TCP performance understanding and its improvements. He has written open source tools such as MPTCPTrace, which provides detailed graphical and statistical information on MPTCP connections based on packet traces.

OLIVIER BONAVENTURE (olivier.bonaventure@uclouvain.be) is a professor at Université Catholique de Louvain, where he leads the IP Networking Lab (INL) (http://inl.info.ucl.ac.be). His research focuses on improving Internet protocols. His recent work includes various improvements to routing protocols, Multipath TCP, and segment routing. He currently serves as an Editor for *SIGCOMM Computer Communication Review*.

Accepted from Open Call

Toward a Net Neutrality Debate that Conforms to the 2010s

Patrick Maillé, Gwendal Simon, and Bruno Tuffin

Network neutrality has been a topic of discussion for the past 25 years, with current legislation/regulation in the United States and Europe targeting the ISPs or "common carriers." But the reality of the Internet in the 2010s is that various actors contribute to the delivery of data, with sometimes contradictory objectives.

Abstract

Network neutrality has been topic of discussion for the past 25 years, with current legislation/regulation in the United States and Europe targeting the ISPs or "common carriers." But the reality of the Internet in the 2010s is that various actors contribute to the delivery of data, with sometimes contradictory objectives. In this article, we highlight the fact that neutrality principles can be bypassed in many ways without violating the rules currently evoked in the debate; for example, via CDNs, which deliver content on behalf of content providers for a fee, or via search engines, which can hinder competition and innovation by affecting the visibility and accessibility of content. We therefore call for an extension of the net neutrality debate to all the actors involved in the Internet delivery chain. We particularly challenge the definition of net neutrality as it is generally discussed. Our goal is to initiate a relevant debate on net neutrality in an increasingly complex Internet ecosystem, and to provide examples of possible neutrality rules for different levels of the delivery chain, this level separation being inspired by the OSI layer model.

Introduction

Net neutrality is "the principle that Internet service providers should enable access to all content and applications equally, regardless of the source, without favoring or blocking particular online services or websites" (*Oxford Dictionary*). This universality principle for the Internet has been present for a long time in the United States. In the 1996 Telecommunications Act [1], the Federal Communications Commission (FCC) declared Internet service providers (ISPs) to be "common carriers" providing a public service. In February 2015, it voted to regulate broadband Internet service as a public utility. The same principle was adopted by the European Parliament in 2014 with a vote that restricts the ability of ISPs to charge specific service providers (SPs) [2]. This principle is often translated as guaranteeing that all packets are treated equally at each intermediary step.

The network neutrality debate started in the 1990s, amid growing concerns about the business models of network operators, which have to deal with several challenging trends, including the increase of traffic volume, the growing traffic asymmetry between operators, and the fast decrease of transit prices. Recently, spectacular disputes between major Internet actors[1] as well as vehement reactions from user associations and governments[2] have drawn attention to that debate. Arguments and discussions on possible regulation of the Internet resulted in refined definitions of network neutrality [4–6].

In this article, we show that the current focus on ISP behavior is too restrictive in the era of *cloud*-based content delivery. Indeed, the regulatory bodies in the United States and Europe — those most active with regard to neutrality — focus on the ISP/internetwork level, but the value chain in the Internet is not restricted to only ISPs and content providers (CPs). In particular, web portals (especially search engines) and content delivery networks (CDNs) have become key components in the delivery chain, while they surprisingly remain absent from the debate. Despite some recurring criticism [7], search engines have not been forced to conform to any universal policy yet; similarly, to the best of our knowledge, the only mention of CDNs in official net neutrality reports is from the Norwegian regulator, according to which "*the ordinary use of CDN servers is not a breach of net neutrality*" [8]. However, we show later that the presence of a CDN, or biases in search listings, can be exploited against the fundamental principles of neutrality.

More generally, we think the net neutrality debate should be extended to all the actors involved in the Internet delivery chain. Our goal is not to discuss the validity of neutrality proponents and opponents, but rather to take a first step toward a global framework that would be more appropriate for the definition of public regulating rules (if any) in the era of the cloud and more generally of *information-centric networks*. We present our vision of such a framework below. To foster scientific activities, we highlight some topics where the rigorous analysis of supposedly "neutral" network scientists could be especially appreciated by decision makers.

Net neutrality is a highly politicized topic. We have tried to be politically neutral and scientifically objective in this article;[3] however, we realize that some readers may interpret some of our statements as being politicized. If so, we assure you that any perceived politicization is unintentional.

[1] For example, Netflix witnessed a 25 percent reduction in traffic rates for Comcast users before a deal was struck [3].

[2] See, for example, Barack Obama's declaration on neutrality in November 2014.

[3] The authors receive funding from Orange (a French ISP) for their research activities, but the content presented in this article is solely the responsibility of the authors.

Patrick Maillé and Gwendal Simon are with Telecom Bretagne (Institut Mines-Telecom); Bruno Tuffin is with Inria.

 0163-6804/16/$25.00 © 2016 IEEE

The Internet in the 2010s

Actors in Today's Internet Delivery Chain

The Internet model usually considered when discussing the relevance of net neutrality is a chain of three actors: users, ISPs, and content/service providers.

ISPs appear as the centerpiece of the debate, being seen as the main intermediary in the delivery chain. But many "newcomers" are now in the picture [9], as illustrated in Fig. 1. In addition to the three aforementioned actors, *transit networks* often act as intermediaries when users and contents are hosted by different ISPs; *CDNs* play a key role by storing content closer to users, thereby reducing transit costs and improving performance for that content;[4] *device builders* may introduce biases through the features of their products (possibly colluding with some other actors); and *search engines* (seen as service providers) directly affect the accessibility (visibility) of content. *Regulatory bodies* therefore face the delicate task of defending fairness and universality principles in this complex ecosystem, where ever changing technical and business conditions prevent (or considerably complicate) analysis and comparison.

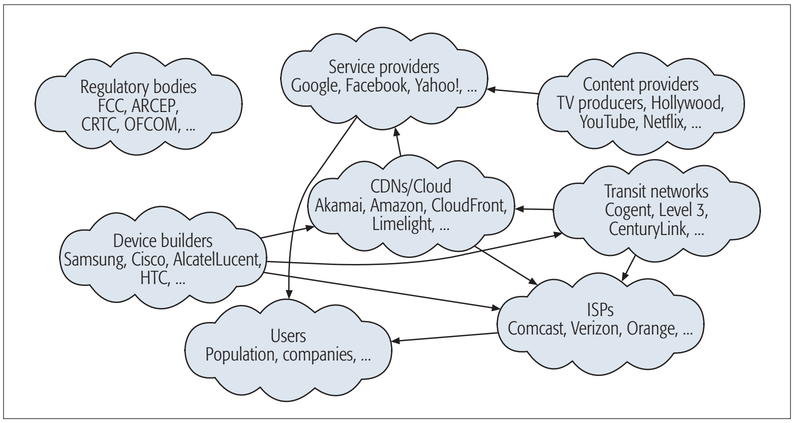

Figure 1. The main categories of actors in the Internet ecosystem; arrows represent a provider-customer (seller-buyer) relationship.

An Intricate Delivery Chain

The success of a given CP partly depends on its interactions with other actors. For example, its *visibility* is strongly affected by the behavior of search engines, which leaves space for biases that we point out later. Another example, the *engagement of end users* for a CP is impacted by the delivery performance [10], which is usually under the responsibility of cloud providers and CDNs. Here again, non-neutral behaviors from those actors impact the CP.

The actors that have direct economic relationships negotiate service level agreements (SLAs). In the case of content delivery, the intricate relationships between actors yield *chains* of SLAs, which are based on network-oriented metrics such as throughput and ratio of packet losses. But inferring the quality of experience (QoE) from a chain of network metrics, potentially cascading, is a difficult task. Most CPs struggle both to assess the QoE of their customers and to identify the failing actor when the quality of delivery is below expectations.

This complex chain of content delivery and the difficulties in assessing the performance of the involved actors potentially leave some room for intermediaries to favor some CPs against others, thereby violating the principles of neutrality. We show later that a packet-level interpretation of neutrality can unfortunately not prevent intermediaries from biasing the competition among CPs.

Existing Protection Tools

Net neutrality opponents often argue that existing legal frameworks already protect fair competition, and thus no additional regulation is needed. This is the case in the United States with the *antitrust law* and the associated Federal Trade Commission (FTC). Such antitrust arsenals do not, however, address some of the points that are frequently raised by net neutrality proponents.

First, the antitrust framework acts reactively and on a per-case basis, although neutrality regulation looks for proactive solutions. With the antitrust framework, an actor that reckons to suffer economically from the (non-neutral) behavior of another actor should file a complaint with the FTC to obtain redress. The framework requires lawsuits and judgments once a plaintiff has built a case. On the other hand, neutrality regulation aims to anticipate the possible problems and prevent them by specifying allowed and forbidden behaviors.

Second, focusing on competition corresponds to a partial view — since it is limited to market aspects — which ignores some other principles raised by net neutrality proponents, such as freedom of speech and innovation fostering (the latter involving equal access to resources regardless of the monetary capacity of CPs).

In our opinion, neutrality rules are essentially *complementary* to antitrust laws, and setting boundaries between the application field of antitrust laws and that of neutrality rules should be part of the debate as well.

Examples of Neutrality Breaches

As previously stated, the supply chain between content and users includes multiple intermediary actors, which act in a free market with reciprocal engagements based on SLAs. We show in the following that these intermediaries can generate significant biases in the network and create a breach in what represents network neutrality for most people. We focus on CDNs and search engines, which are two key actors in the supply chain.

Competition Biases Due to CDNs

First, we consider the impact of a profit-driven CDN provider by developing the two scenarios depicted in Fig. 2. Both scenarios model the economic reality of the traditional interactions between ISPs, CDNs, and CPs.

A CDN serves a large population of end users by means of edge servers, which are usually located near the point of peering (PoP) between ISPs. When a CDN receives a request, it has two handling options: either fetch the requested content from the origin server, or fetch the requested content from the edge server. Option 1 incurs transit costs (which can be low if the CDN also owns a transit network, but large otherwise), while option 2 leads to improved user QoE,[5] and no transit costs are incurred. Thus, in terms of cost and quality, option 2 should always be

[4] The use of CDNs can be seen as against neutrality principles, since they offer an improved delivery for a fee, something new arrivants in the market may not be able to afford. We do not address this question here.

Note that we focus on static content in these models for the sake of clarity. But the CDN models for interactive services, such as live streaming and cloud gaming, are not essentially different: the edge server should still reserve some hardware resources (mostly computing power here) to serve a subset of users.

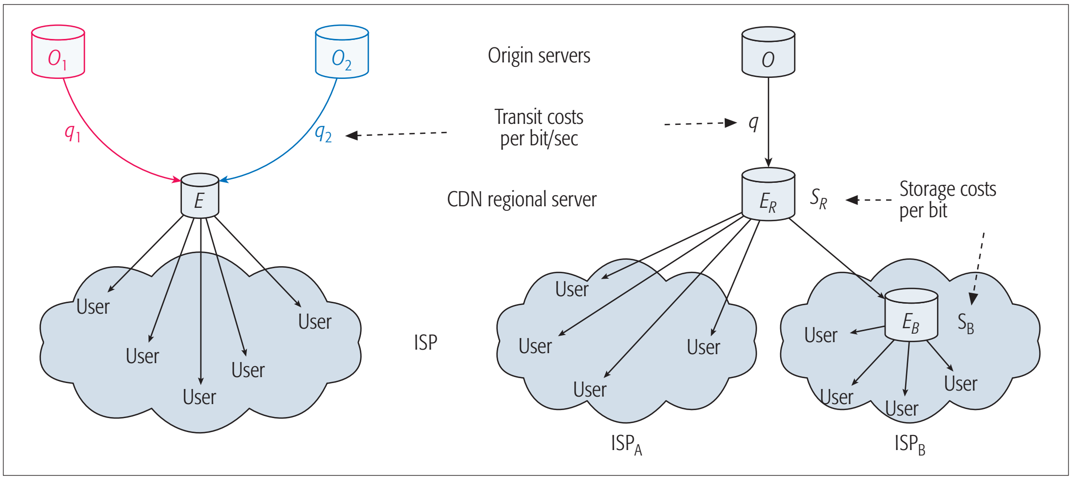

Figure 2. Two cases of neutrality breaches due to the presence of a CDN. On the left, two content providers, each with its own origin server, compete for the storage resources of a regional CDN edge server [11]. On the right, two ISPs are served by a shared CDN regional server, but one ISP (here ISP_B) installs its own edge server.

chosen; however, storage/cache and processing resources at E are finite. Hence, some requests must be handled by option 1 (fetch from origin server). A rational CDN will maximize profits by optimizing the balance between options 1 and 2, playing on caching decisions at E (or possibly on the number/capacity of edge servers).

Note that we focus on static content in these models for the sake of clarity. But the CDN models for interactive services, such as live streaming and cloud gaming, are not essentially different: the edge server should still reserve some hardware resources (mostly computing power here) to serve a subset of users. For example, in live rate-adaptive video streaming systems, the edge server should both get multiple video representations from the origin server and manage the multiple concurrent requests for the videos.

Biasing the Competition among Content Providers: Let us consider the scenario shown on the left of Fig. 2. That scenario is extensively discussed in [11]. To reach customers in a given ISP, two competing CPs subscribe to the CDN, which has to share the privileged resources in the edge server E among the CPs. The scientific literature comprehensively studies the performance of content replacement strategies in dynamic scenarios, but rarely integrates the economic parameters, here:

- The *transit costs* q_1 and q_2, which can differ among CPs due to different paths to reach the ISP of interest
- The *charging policy* of the CDN, with prices per volume possibly differing depending on whether users are served from an edge or an origin server

A profit-maximizing CDN stores in the edge server the content that yields the largest revenues. If transit costs q_1 and q_2 are equal, and if the CDN charges both CPs the same way, it should store the most popular content. However, if the transit costs of one CP (say CP_1) are higher than the other, the interest of the CDN in storing CP_1 content in the edge server exceeds that of storing CP_2 content of comparable popularity, which both benefits CP_1 customers and harms CP_2 customers. In other words, *the internal costs of a CDN can distort competition*.

The situation gets even more problematic when the CDN charges CPs differently. In [11] we study a case based on real popularity measures, CP_2 being (on average) approximately five times more popular than CP_1. Figure 3 displays the "satisfaction" (a normalized customer QoE measure) of both CPs, according to the ratio of content from CP_1 in the edge server.

The "standard scenario" situation corresponds to the "optimal" (for the CDN) sharing of the edge server when both CPs are charged identically. With CP_2 being more popular, approximately five-sixths of the edge server is filled with CP_2 content, and the average satisfaction of CP_2 customers exceeds that of CP_1 customers. Hence, the CDN favors the dominant player, although both CPs pay the same price. Note, however, that the CDN strategy is the most efficient in terms of traffic reduction and overall satisfaction of end users.

Now, assume that one of the CPs pays more for its requests fulfilled by the edge server. Two of the vertical lines correspond to two extremes where one CP pays 10 times the basic price (see [11] for more details). When CP_2 pays more (typically to reinforce its dominant position), the satisfaction gain of CP_2 is small, but competitor satisfaction drops more significantly (from 0.94 to 0.89), possibly leading some CP_1 customers to churn. Alternatively, when CP_1 pays more (typically to increase its audience with better quality), the impact is less spectacular. To summarize, *the dominant player can leverage the CDN to harm its competitors, although challengers cannot.*

Biasing the Competition among ISPs: Consider now the scenario depicted on the right of Fig. 2, where a CDN serves two ISPs. Initially, users of both ISPs (labelled A and B) get data from either the origin server O or a regional "shared" edge server E_R. The size of the cache in E_R is chosen by the CDN according to the transit cost q from O to E_R (based on traffic) and the storage cost s_R at E_R (based on volume).

Now assume one ISP, say ISP_B, installs an edge

[5] Note that QoE is mainly affected for services requiring high bandwidth or low latency.

server E_B within its network[6] with a storage cost s_B, in order to improve the QoE for its users through an agreement with the CDN. The profit-maximizing CDN now has four ways to store content: in none of the edge servers (so only in the origin server), in E_R only, in E_B only, or in both E_R and E_B.

The appearance of edge server E_B can create conflict, in particular when content previously stored in E_R gets stored only in E_B (and no longer in E_R): the QoE of ISP_A users is then degraded. Such a scenario is more likely to occur when the storage cost of E_B is significantly lower than E_R, or when ISP_B traffic is significantly greater than ISP_A traffic. In that case, the content stored in the regional "shared" edge server E_R is the "second" most popular in ISP_B, that is, the most popular among the content not stored in E_B. Thus, *agreements between CDNs and a given ISP can degrade the performance for users in* another *ISP*.

Figure 4 illustrates the impact of edge server E_B on the overall satisfaction (still a normalized QoE measure) for users in both networks, using real movie request data over two years obtained from a leading French video on demand (VoD) service and including the request origin networks. We set ISP_B as the network with the highest traffic, and attribute to ISP_A the average traffic from all other networks. The *y*-axis shows the ratio between the satisfaction with and without the edge server E_B; the *x*-axis corresponds to different storage costs s_B.

We observe that the installation of an edge server in ISP_B has an (expected) positive impact on the satisfaction of ISP_B users, but also an (undesirable) negative impact on the satisfaction of ISP_A users. Observe that the latter exceeds the former for most of the storage cost prices on E_B: a powerful network operator does not need to implement any aggressive pricing policy to degrade the performance of competitors.

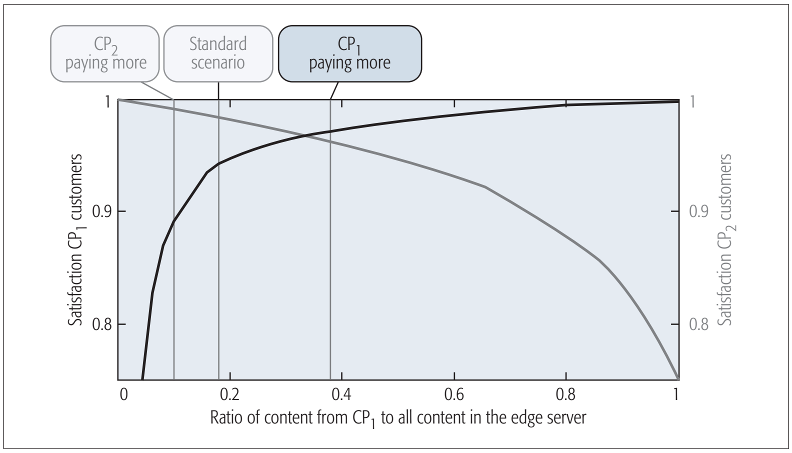

Figure 3. Satisfaction for both CP_1 and CP_2 customers according to the ratio of the edge server filled with content from CP_1.

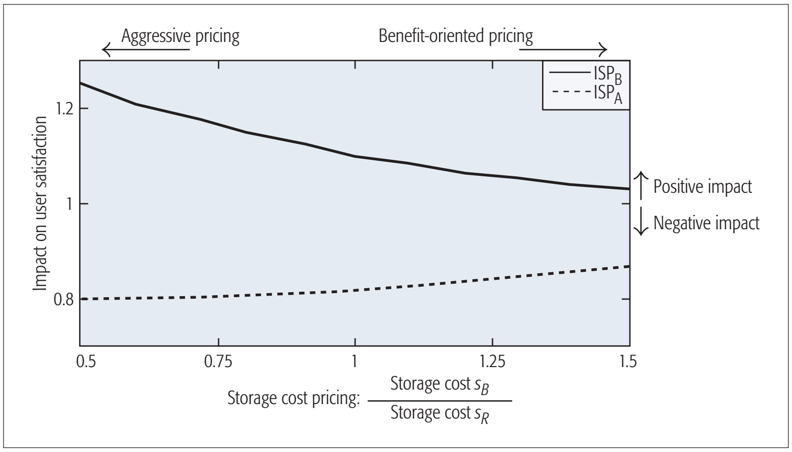

Figure 4. Impact on user satisfaction of the creation of an edge server in ISPB, computed as the ratio

$$\frac{\text{satisfaction with cache } E_B}{\text{satisfaction without } E_B}.$$

Search Engines Biasing the Access to Content

Search engines are the preferred way to discover and access content: according to ComScore, in 2014 about 21 billion searches were treated each month in the United States alone. Search engines usually present their suggestions corresponding to keywords in two categories:

- *Sponsored results*, a list of advertisements related to the search and clearly defined as such, from which the engines make money when the links are clicked
- *Organic results*, a ranked list of links, believed to associate the most relevant items to the requested keywords since promotional results are assumed to be within sponsored links

To increase revenues, search engines can be tempted to include in the organic results some results that are not among the most relevant but can generate (direct and indirect) short-term revenues. A typical case is that of a content provider (called CP_1) vertically integrated with (i.e., owned by) the search engine. The engine is economically incentivized to rank that content higher [12]. Such temptation has raised a *search neutrality* debate parallel to the one on network neutrality [9]. We do not target here any specific search engine nor claim one is biased: our goal is to show through a model that such biases can be motivated, and to investigate their impact on fairness among CPs.

The model designed in [12] analyzes the temptation to return "non-neutral" organic results. This model includes a long-term effect: end users who are unsatisfied with the relevance of the organic results may stop using the search engine. The search engine objective is then to maximize its long-term revenue, which corresponds to the trade-off between the short-term revenue per visit on CP_1 and the longer-term number of visits on the search engine. The model helps to understand revenue-maximizing ranking strategies and to anticipate the impact of regulatory interventions for various scenarios. The principle is illustrated in Fig. 5.

Table 1 illustrates, for some arbitrarily chosen numerical values, the differences in visit rates to CPs due to the revenue-optimal ranking, the resulting average relevance of the engine's output, and the revenues it can generate, for different values of a parameter β, which represents the average revenue directly generated per visit to the search engine (via clicks on sponsored links, their assumed single source of revenue); the smaller β, the larger the bias because of larger

[6] In most cases, the ISP owns the edge server, which it rents to the CDN.

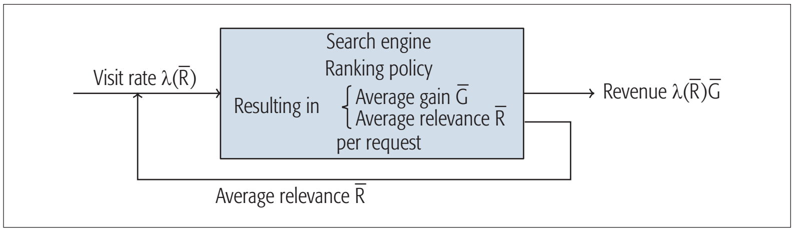

Figure 5. Search engine whose ranking policy produces an average relevance of results and an average gain. The number of visits (i.e., popularity of the engine) depends on the average relevance.

(relative) incentives to create revenue through integrated CPs. In this scenario, the impact on the average relevance remains small (around 10 percent), but the impact on the visibility and revenues of CPs is substantial, possibly threatening the survivability of non-integrated CPs.

Toward a Global Framework

Net neutrality proponents often claim that every packet should be treated equally, which misses important aspects and does not prevent unfair situations. Indeed, that proposal focuses on the regulation of the third layer of the OSI model (the so-called network layer), although the end users get the results of the policies of all intermediaries in the seven layers. As shown in the previous section, intermediaries can introduce significant biases while fulfilling this "neutral" condition at the third layer.

We therefore call for an extension of the debate to any class of actors having an influence in the way end users consume content, including CDNs, as well as the providers of data centers, and search engines.

The need is then for a definition of neutrality that would be applicable to all actors, the enforcement of which would prevent the behaviors pointed out previously. Such a definition would need to go beyond the packet level, and may involve a classification of all actors into categories, for each of which specific principles should be followed. The regulation would then be specific to each level; below are a few examples:

- At the ISP level, no differentiation among packets according to the origin, destination, service, or price paid (the common suggestion)
- At the CDN level, caching of only the most popular content, independent of the provider, the type of content, or any monetary aspect
- At the search engine level, ranking results based only on relevance to the query (and possibly the user), independent of the result page owner, the type of query, or possible payment

Our purpose in the remainder of this section is not to put forward a clear framework for a new neutrality definition, ideally robust to network evolution, but rather to launch a discussion in the community toward that goal. To start that discussion, we hereafter pose a few questions, and provide limited and probably flawed suggestions, which will (hopefully) ignite the debate.

Where Should Neutrality Be Applied?

The whole topology of the Internet should be abstracted to capture all the relevant actors that may introduce inappropriate biases. This effort should not be restricted to the physical infrastructure, but should also consider the actors that contribute to the consumption of content by any means. Typically, search engines are not captured by the current framework targeting the packet level and ISPs. This effort should also consider the new actors that are expected to implement and integrate their network management services into the network devices compatible with network functions virtualization (NFV) standards. More generally, all actors that transform the content and are part of the delivery should be considered.

A definition of neutrality should study behaviors (actor actions) *at each actor level*: this would imply defining levels for the different types of actions on content in a way comparable to the OSI reference model. For instance, we would need to characterize (in terms of level) the services provided by CDNs and by search engines. Players at a given level/layer should then be neutral in terms of protocols applied at this level by avoiding any kind of differentiation through the compliance of level-specific rules, extending the current rules defined for ISPs at the network/transport layer.

Note that unfortunately, it is not easy to identify the level at which an actor plays. For example, in France, Skype is considered by the regulator as a virtual network operator, which thus competes with ISPs, and not as a service (or application) provider.

On What Notion Should Neutrality Focus?

What is the goal of neutrality; what do we want to preserve? The current focus is on equal treatment: no differentiation would be allowed whatever the source, application, or service; hence, there is some induced notion of *fairness*.

Among the underlying principles is the universality of access. This is a baseline (as it only focuses on blocking), but we need stronger constraints in a neutrality definition. Slowing down some applications, or favoring others, is generally refused by neutrality proponents and does not conform to the packet-based equal treatment proposal. On the other hand, the supposed equal treatment at the packet level can be criticized since most Internet applications use TCP at the transport layer, which leads to unfair qualities of service offered since the TCP throughput varies with the user/provider distance (it is indeed inversely proportional to the square root of the round-trip time). This characteristic of TCP can harm innovation in regions with bad network quality, even with perfectly equal packet treatment.

Defining the "no differentiation" principle when talking about CDNs, for instance, seems subtle, because CDNs precisely differentiate services when selecting what content to cache. The notion of neutrality could be in terms of avoiding money-based differentiation: fairness considerations would help decide what to cache, but there can be several potential interpretations:

- Should any packet have the same probability of being cached? We do not see any reason to do that.
- Should it rather be the most popular content that is cached (which is usually assumed natural to limit the load on the network)? But could we not reason differently, saying that all content providers should be equally cached, or in the end offer the same QoE (something not done for the traditional neutrality principle though)?

It seems reasonable to set a "neutrality rule" requiring CDNs to cache the most popular content, *independent* of the source *identity*. But the focus could also be on reducing the overall network load by favoring the caching of distant content (which contributes more to the load by using more links) in addition to considering popularity.

In order to (partially) solve these questions, we suggest to define fairness, like neutrality, at the layer the actor plays (related to the question "where?"). This type of definition could also encompass search engines and ISPs. The basic idea is to have "equal treatment," but one could also look at other notions, such as social welfare, user welfare, or even service welfare (note, however, that such efficiency measures are in general incompatible with fairness, an argument often used by non-neutrality defendants). The goal is to determine whether a behavior harms the desired goal at the layer where it is operating.

When Should It Be Controlled?

Should we ensure that any newcomer is given a particular treatment? Regarding the current strategies of CDNs, CPs in their initial stages are disadvantaged against incumbents. However, favoring these companies would be in disagreement with the current neutrality rules of equal treatment. It is therefore a slippery slope, opening the door to "mission creep" situations where new types of biases would defeat the initial goal of achieving neutrality.

With the current focus of the debate, note that this question is not really a problem, but it is if we try to have a broader view of the Internet network.

Potential Next Steps

To conclude this article, we would like to present a few issues to be tackled by the community.

Taking into Account Economics in Future Internet Studies

In general, the question of actor profit is rarely central in the literature related to network and service management. Computer science and electrical engineering scientists generally aim to maximize the efficiency of their proposal regarding technical objectives and not economic ones. Typically, to our knowledge the literature related to content-centric networking (CCN) does not deal with the economics of actors, although in our opinion, the revamping of the Internet cannot be seriously studied without taking economic factors into account.

Evaluating the Impact of Possible Regulatory Policies

The role of a regulating agency is to recommend policies that guarantee widely accepted principles. As previously mentioned, the stated principles are not (yet?) well defined when it comes to neutrality. It is up to the scientific community to provide rigorous studies, as unbiased as possible, about potential policies, their impact on the considered actor level, and also the possible impact on the overall delivery chain.

Defining and Studying New Fairness Models

Our brief discussion above highlights the lack of a better definition of fairness and in particular the subtle equilibrium between the guarantee of fair competition among actors pursuing similar objectives (the notion of actor level) and the preservation of the motivation for investment toward satisfying end users better.

	Relevance	CP_1 revenue	Other CP revenue	CP_1 visit rate	Other CP visit rate
Neutral	0.635	0.028	0.0283	0.057	0.057
Non-neutral (for $\beta = 1$)	0.618 (–3%)	0.066 (+136%)	0.0243 (–14%)	0.112 (+96%)	0.049 (–14%)
Non-neutral (for $\beta = .5$)	0.592 (–7%)	0.084 (+200%)	0.0215 (–24%)	0.140 (+146%)	0.043 (–25%)
Non-neutral (for $\beta = .25$)	0.568 (–11%)	0.093 (+232%)	0.0193 (–32%)	0.158 (+177%)	0.039 (–32%)

Table 1. Impacts of ranking when the search engine owns CP_1.

References

[1] FCC, Telecommunications Act, http://transition.fcc.gov/telecom.html, 1996.

[2] European Commission, "Proposal for a Regulation of the European Parliament and of the Council Laying Down Measures Concerning the European Single Market for Electronic Communications and to Achieve a Connected Continent, and Amending Directives 2002/20/ec, 2002/21/ec and 2002/22/ec and Regulations (ec) no 1211/2009 and (eu) no 531/2012," http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=COM:2013:0627:FIN:EN:PDF, 2013.

[3] J. Brodkin, "Netflix Performance on Verizon and Comcast Has Been Dropping for Months," *Ars Technica*, Feb 2014.

[4] T. M. Lenard and R. J. May, Eds., *Net Neutrality or Net Neutering: Should Broadband Internet Services Be Regulated*, Springer, 2006.

[5] A. Odlyzko, "Network Neutrality, Search Neutrality, and the Never-Ending Conflict Between Efficiency and Fairness in Markets," *Rev. Network Economics*, vol. 8, no. 1, pp. 1–21, 2009.

[6] T. Wu, "Network Neutrality, Broadband Discrimination," *J. Telecommun. and High Tech.*, 2003.

[7] F. A. Pasquale and O. Bracha, "Federal Search Commission? Access, Fairness and Accountability in the Law of Search," *Cornell Law Rev.*, Sept. 2008.

[8] Norwegian Communications Authority, "Norwegian Report on Content Delivery Networks (CDNs)," NKOM report, May 2012.

[9] P. Maillé and B. Tuffin, *Telecommunication Network Economics: From Theory to Applications*, Cambridge Univ. Press, 2014.

[10] S. Shunmuga Krishnan and Ramesh K. Sitaraman, "Video Stream Quality Impacts Viewer Behavior: Inferring Causality Using Quasi-Experimental Designs," *IEEE/ACM Trans. Networking*, vol. 21, no. 6, 2013, pp. 2001–14.

[11] P. Maillé *et al*., "How Neutral is a CDN: An Economic Approach," *Proc. 10th IEEE Int'l. Conf. Network and Service Management*, Nov. 2014.

[12] P. L'Ecuyer *et al*., "Revenue-Maximizing Rankings for Online Platforms with Quality-Sensitive Consumers," tech. rep., Inria, Feb. 2014; https://hal.inria.fr/hal-00953790.

Biographies

Patrick Maillé (patrick.maille@telecom-bretagne.eu) graduated from Ecole Polytechnique and Telecom ParisTech in 2000 and 2002, respectively. Since then, he has been with the Networks, Security, and Multimedia Department of Telecom Bretagne, obtaining a Ph.D. in 2005. His research interests are in all economic aspects of telecommunication networks, from pricing schemes at the user level to auctions for spectrum and regulatory issues (net neutrality, search neutrality).

Gwendal Simon (gwendal.simon@telecom-bretagne.eu) graduated from University Rennes 1 (France). During his Ph.D. work (obtained in 2004), he worked at Orange Research Labs. From 2004 to 2006 he was a researcher at Orange Labs. Since 2006, he has been an associate professor at Telecom Bretagne. He was a visiting researcher at the University of Waterloo from September 2011 to September 2012. His research interests include multimedia delivery systems (video and gaming) and network management.

Bruno Tuffin (bruno.tuffin@inria.fr) received his Ph.D. degree in applied mathematics from the University of Rennes 1, France, in 1997. He is currently a senior researcher at INRIA (the French National Institute for Research in Computer Science and Control). His research interests include developing Monte Carlo and quasi-Monte Carlo simulation techniques for the performance evaluation of telecommunication systems and telecommunication-related economical models.

Highly Dynamic Spectrum Management within Licensed Shared Access Regulatory Framework

Aleksei Ponomarenko-Timofeev, Alexander Pyattaev, Sergey Andreev, Yevgeni Koucheryavy, Markus Mueck, and Ingolf Karls

Historical fragmentation in spectrum access models accentuates the need for novel concepts that allow for efficient sharing of already available but underutilized spectrum. The emerging LSA regulatory framework is expected to enable more advanced spectrum sharing between a limited number of users while guaranteeing their much needed interference protection.

ABSTRACT

Historical fragmentation in spectrum access models accentuates the need for novel concepts that allow for efficient sharing of already available but underutilized spectrum. The emerging LSA regulatory framework is expected to enable more advanced spectrum sharing between a limited number of users while guaranteeing their much needed interference protection. However, the ultimate benefits of LSA may in practice be constrained by space-time availability of the LSA bands. Hence, more dynamic LSA spectrum management is required to leverage such real-time variability and sustain reliability when, for example, the original spectrum user suddenly revokes the previously granted frequency bands as they are required again. In this article, we maintain the vision of highly dynamic LSA architecture and rigorously study its future potential, from reviewing market opportunities and discussing available technology implementations to conducting performance evaluation of LSA dynamics and outlining the standardization landscape. Our investigations are based on a comprehensive system-level evaluation framework, which has been specifically designed to assess highly dynamic LSA deployments.

INTRODUCTION AND BACKGROUND

CURRENT SPECTRUM ACCESS AND MANAGEMENT MODELS

Over the years, radio spectrum has become a critical resource for numerous purposes: from economic and social to cultural and scientific. However, its management has largely remained unchanged in the course of the past three decades due to the underlying complexity of the process and insufficient maturity of radio technology. Along these lines, various distinct approaches to spectrum management have historically taken shape.

"Command-and-Control" Spectrum Management: This age-old antiquated paradigm executes static spectrum allocation. Accordingly, a regulatory body assigns a frequency band to a particular entity while imposing strict constraints on such use. Naturally, this approach led to barriers in spectrum access, bringing along difficulties in meeting the increasing demand for wireless spectrum-based services. In addition, the corresponding assignment of frequency bands never relied on market mechanisms, hence resulting in very low economic profits. Conventionally, spectrum ownership rights have been granted as the result of so-called "beauty contests" and required considerable lobbying of regulation authorities.

Exclusive Use of Spectrum: This model is centered around a long-term (15 to 30 years) spectrum band license awarded to utilize a particular band. Correspondingly, the resulting use is subject to certain well defined rules, such as maximum power levels and geographical coverage. Exclusive licenses empower their respective owners (e.g., cellular network operators) with unrestricted interference management capabilities, thus enabling quality of service (QoS) guarantees, but at the same time impose high market entry barriers (i.e., billions of euros). As opposed to the legacy "beauty contests," assignment was transformed by sales of spectrum: most regulators have now adopted market-centric approaches (e.g., auctions) to redistribute frequency allocations.

Shared Use of Primary Licensed Spectrum: In this concept, the frequency bands of a licensed owner (the primary user) are shared by a non-license holder (the secondary user). Importantly, access by the secondary user may sometimes occur without notifying the primary user and requires the respective protection of the latter, such that the intended operation of primary communication is not deteriorated. In this regard, there has been a recent surge in software-defined radio technologies, cognitive and adaptive radio networks, as well as reconfigurable networking to enable the intended dynamic spectrum access (e.g., in TV white space). However, the fundamental limitation of this form of access is that it is unclear how the secondary user may deliver reliable QoS guarantees over such shared spectrum.

Shared Use of Unlicensed Spectrum: When a spectrum band is allowed for "open access," no entity can claim its exclusive use, and the target spectrum should be made fairly accessible to everyone. An example of such spectrum usage is the industrial, scientific, and medical (ISM) bands, where multiple potential users (e.g.,

Aleksei Ponomarenko-Timofeev, Alexander Pyattaev, Sergey Andreev, and Yevgeni Koucheryavy are with Tampere University of Technology; Markus Mueck and Ingolf Karls are with Intel Mobile Communications.

0163-6804/16/$25.00 © 2016 IEEE

medical and sensor devices, microwave ovens, cordless phones, WiFi networks) may access the spectrum without external regulation. While such unregulated access significantly lowers market entry barriers, it also produces uncontrolled wireless interference and consequently makes it extremely hard to meet the desired QoS guarantees. In addition, a multitude of spectrum sharers may lead to a situation in which none of the users achieve their expected benefit. This is a very likely course of development today, given the increasing popularity of WiFi and the corresponding emphasis of network operators on different forms of WiFi offloading.

Transformation of the Global Wireless Landscape

To overcome the long-standing effects of fragmentation in spectrum access models, there is a pressing demand for novel frameworks allowing for efficient sharing of already available but underutilized spectrum. The need for this change is becoming increasingly urgent as the pressure on the radio spectrum is steadily building, largely due to the unprecedented explosion in wireless traffic. Indeed, recent forecasts by Cisco predict the growth in mobile data demand at a rate of nearly 60 percent over the next 5 years, which brings along a 10-fold overall increase. In this regard, past traffic growth predictions look overly optimistic in that they heavily underestimate the mobile data acceleration [1].

As data from mobile and wireless devices is expected to soon exceed traffic from wired equipment, the fourth generation (4G) networks of today face the risk of "capacity crunch." To this end, the forthcoming 5G technologies offer a range of decisive improvements in cell capacity [2]. However, more efficient use of existing spectrum will not solely be sufficient to achieve the needed factors of 1000- to 10,000-fold improvement. These targets by 2020 are impossible without the availability of additional frequency resources, which will be required for a range of spectrum-hungry technologies from conventional mobile-to-infrastructure links to complementary device-to-device and multihop communication, as well as wireless front- and backhauling.

Regrettably, given that the traditional approach of repurposing spectrum is reaching its limits (especially on bands below 6 GHz), it is unlikely that more contiguous and broader microwave frequencies will be made available any time soon. At the same time, whereas radio spectrum may be saturated during peak hours and/or in crowded locations, there is presently an extreme variability of load across time and space. Hence, this dynamics may be exploited to manage spectrum more efficiently, especially given the fact that average traffic grows at a much slower rate than busy hour traffic. Currently, though, there are no feasible options to manage spectrum on such small-scale spatio-temporal granularity, which calls for new approaches to spectrum policy and allocation methods.

In light of the above, it appears that the shared use of spectrum becomes unavoidable even for those who have conventionally enjoyed exclusive access rights [3]. However, the existing forms of spectrum sharing (in primary licensed or unlicensed spectrum, see above) do not offer

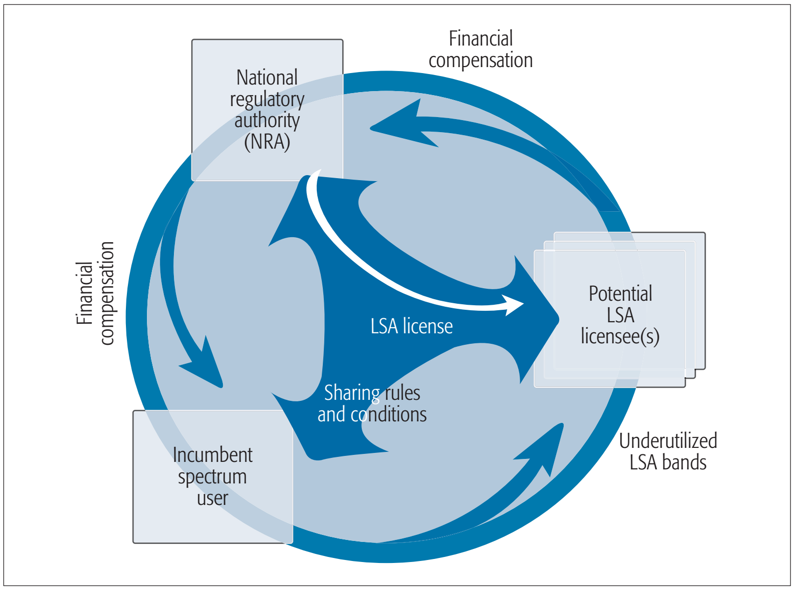

Figure 1. LSA regulatory framework: key stakeholders.

much needed interference protection, thus resulting in insufficient reliability, QoS guarantees, and predictability of operation. By contrast, the emerging licensed shared access (LSA) regulatory concept (Fig. 1) allows for more advanced spectrum sharing between a limited number of entities with carefully defined usage rights, combining the benefits of command-and-control spectrum management with a flexible and innovative market-friendly approach.

Broadly, LSA enables *authorized* spectrum sharing by allowing at least two users, the *incumbent* (i.e., the current holder of spectrum rights) and the LSA *licensee* (i.e., the temporary user of spectrum) to access the same frequency bands in a licensed predetermined manner following a well defined mutual agreement [4]. In other words, LSA guarantees that the incumbent retains spectrum access rights anytime, anywhere, and the LSA licensee(s) will refrain from using this spectrum when needed by the incumbent (or at least will not disrupt the incumbent's operation).

Under the LSA's "individual licensing regime," sharing agreements need to guarantee high predictability in terms of spectrum access for all the involved parties:

- For the incumbent(s), LSA leverages additional economic benefits from underutilized spectrum without imposing any significant operational restrictions on its expected use.
- For the national regulator, LSA harmonizes spectrum usage, opening a path to its optimization via controlled sharing as an alternative to permanent segmentation.
- For the licensee(s), LSA delivers additional frequencies at more affordable costs together with predictable QoS guarantees due to coordinated interference.

However, the licensee's benefits from LSA may in practice be constrained by space and time availability of the LSA bands. As long as LSA usage remains static, it should suffice that a dedicated exclusion zone or time is created to protect the incumbent's use of spectrum. On the other

LSA enables authorized spectrum sharing by allowing at least two users, the incumbent (i.e., the current holder of spectrum rights) and the LSA licensee (i.e., the temporary user of spectrum), to access the same frequency bands in a licensed predetermined manner following a well defined mutual agreement.

hand, in case of dynamic geographic/temporal LSA sharing, on-demand authorization of the LSA licensee(s) is required as a consequence of real-time restrictions imposed by the incumbent. While such dynamic LSA systems are more complex to build and maintain, they also unlock higher potential performance benefits. In what follows, we concentrate on highly dynamic LSA operation allowing a licensee's spectrum access over a particular frequency, time, and location. To this end, we offer our vision of the required functionality for the LSA architecture to support such dynamics. In addition, we summarize our recently completed system-level study of LSA performance with a dedicated set of tools that we contribute to make conclusions across a wide range of LSA-centric use cases and scenarios.

LSA System Architecture and Implementation

Use Cases and Market Opportunities

To ensure pragmatic and efficient LSA operation providing the desired spectrum access flexibility and harmonization, it is crucial to identify viable use cases and scenarios of its application.

Mature Operator Markets: First and foremost, already today LSA may benefit mobile network operators (MNOs) with mature 3G markets, but lacking 4G coverage and capacity benefits due to the lengthy spectrum refarming process. In such markets, where players are typically reluctant to alter their existing business strategies, LSA may change the rules of the competition by allowing smaller MNOs to quickly augment their capacity and coverage.

Smaller and Virtual Operators: Going further, the larger dominating MNOs owning exclusive spectrum licenses may be challenged by smaller MNOs, which in the past had restricted business opportunities due to very little exclusive spectrum. However, dominating MNOs can also strengthen their market positions by acquiring extra LSA bands. In addition, non-MNO players, such as virtual network operators, may proliferate in the market, thus reshaping the existing business ecosystem.

Mobile Broadband Services: LSA may also support forward-looking governmental plans to increase adoption of public services over mobile broadband. Indeed, as predicted by many sources, the novel types and higher numbers of wireless services are very likely to mushroom around 2020. Supported by LSA, the emerging 5G trends may include mobile ultra high-definition holography and multimedia-based immersion, large-scale augmented and virtual reality, big data processing, as well as public safety and disaster relief.

Rural and Machine-Type Markets: Furthermore, LSA holds significant promise for markets with large rural populations, as well as the machine-to-machine, wearable, and Internet-of-Things (IoT) markets. To this end, LSA may help leverage the available secondary spectrum in areas with low population densities.

Ultra-Dense Heterogeneous Networks: Finally, LSA also has the potential to aid the deployment of ultra-dense networks based on multi-radio small cells. Overall, the latest analysis of frequency requirements by the International Telecommunication Union Radiocommunication Sector (ITU-R) indicates significant bandwidth demand across today's heterogeneous network deployments. Consequently, over 1000 MHz of new spectrum is currently required, and more efficient spectrum utilization frameworks, such as LSA, are an important building block to enable certain well defined scenarios.

While a separate LSA business case may be difficult to identify, it can be foreseen that LSA will become one of the potential dynamic spectrum access modes together with exclusive access, co-primary shared access, authorized shared access, unlicensed access, and, perhaps, other options in future 5G systems. Hence, as exclusive access will continue to remain the preferred method of spectrum usage by 5G-grade MNOs, we believe that LSA will be increasingly employed as a complementary approach in conjunction with other spectrum access alternatives, such as unlicensed WiFi in 2.x and 5.x GHz bands, TV bands below 800 MHz, unlicensed cellular access in 1800 MHz, and so on.

Naturally, LSA principles are based on voluntariness, where the regulator is not expected to force incumbent(s) to accept sharing. Instead, driven by their economic benefits the incumbent(s) are stimulated to provide the LSA licensee(s) with access to a part of their spectrum at certain locations and times. In addition, rules must be defined allowing the incumbent to revoke such granted spectrum should it be required again (or if the licensee is causing harmful interference to the incumbent), and the respective mechanisms are considered below.

Prospective LSA System Architecture

As follows from the above, the envisioned LSA ecosystem assumes an intricate interplay between the national regulatory authority (NRA), the incumbent(s), including both governmental and commercial entities, and the potential LSA licensee(s). To define a simple and easy-to-deploy sharing framework, as well as determine appropriate rights of its use, these stakeholders need to engage in intensive bi- and trilateral dialogs [5]. This should allow cellular operators to leverage additional spectrum on a secondary basis, with exclusive and guaranteed access over certain time, frequency, and geographical area. To this end, the prospective LSA system architecture [6] features the LSA repository, the LSA controller, and the mobile wireless communication network operations, administration, and management (OA&M) entity (Fig. 2).

LSA Repository: This is essentially a database that may include a variety of information on both the incumbent and the licensee(s). In particular, it needs to store up-to-date space, time, and frequency information on the incumbent's spectrum utilization. Accordingly, the repository is primarily responsible for delivering information on spectrum availability and associated conditions, but may also add safety margins and even deliberate distortions to such data — the incumbent may not be willing to disclose precise information due to its sensitive nature. The management of the

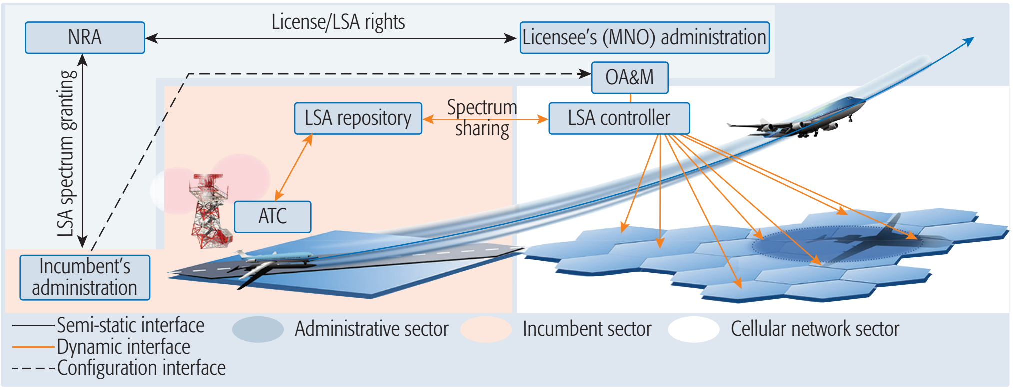

Figure 2. Envisioned LSA architecture and motivating scenario.

LSA repository may be performed by the NRA or the incumbent directly, or can be delegated to a trusted third party.

LSA Controller: This generally manages access to the spectrum made available to the LSA licensee based on sharing rules and information on the incumbent's use provided by the LSA repository. There is typically a direct link between the LSA controller and at least one LSA repository, which allows secure and reliable information transfer and requires a standardized communication interface. Correspondingly, after the incumbent's spectrum use information from the repository has been combined with the sharing rules built on the current LSA usage rights, the controller evaluates the LSA spectrum availability and provides the respective grant to the LSA licensee.

OA&M Entity of the LSA Licensee's Mobile Network: This performs the actual management of the LSA spectrum by issuing the radio resource management (RRM) commands based on the information received from the LSA controller. These RRM commands, after they have been delivered to the MNOs' base stations, enable user equipment (UE) to either transmit on the LSA spectrum or hand over to another frequency band subject to LSA spectrum availability, QoS requirements, or data plan preferences. In addition, OA&M can help the associated base stations with channel and/or transmit power level selection.

The considered LSA system design enables efficient transition from relatively static to significantly more dynamic LSA operation. Indeed, for static incumbents bound to a particular location and time (e.g., a military base or a TV studio), the resulting interference could be controlled by simple pre-planned exclusion methods. However, in case of a dynamic incumbent (e.g., a radar system or a broadcasting service provider), a significantly more capable low-latency interface between the LSA repository and the LSA controller should become available. It needs to allow near-real-time coordination between the incumbent(s) and the LSA licensee(s), as well as timely revocation of the LSA frequency bands by the incumbent in case of emergency or excessive interference from the licensee. We thus continue with reviewing the available implementation options for the dynamic LSA system.

Technology Aspects and LSA Implementation Design

In the first place, LSA needs simple mechanisms allowing the users of a licensee MNO to efficiently enter and vacate the LSA spectrum. For example, after a radio access network (RAN) begins advertising the availability of the LSA band, the idle-mode UE may follow the standard reselection procedures to move to the LSA frequencies. However, such a decision is user-centric in nature — it may cause lengthy delays and uncertainty in intended LSA operation and hence may not be preferred by the MNOs. An alternative network-centric solution is to directly handover the connected-mode UE to a certain component carrier within the LSA frequencies. Unfortunately, in presently deployed cellular networks, such as Third Generation Partnership Project (3GPP) Long Term Evolution (LTE) Releases 8 and 9, the UE may only use one component carrier at a time, which naturally limits the possibility to employ both primary licensed band and LSA band for increased reliability.

Starting with Release 10, LTE technology defines a carrier aggregation (CA) mechanism that essentially enables the utilization of several component carriers simultaneously. Given that CA also remains under the full control of the network, it is more efficient and robust since the UE does not in fact change its underlying operating band. More importantly, CA provides means to implement LSA today, without significant modification of existing MNO deployments. The downside of CA, however, is the need for higher signaling overhead, and research is currently underway to improve CA operation for LSA. In addition, it is expected that practical LSA deployments would require a range of new dedicated mechanisms taking into account the radio technology used for an incumbent's transmissions, such as RRM, interference mitigation, load balancing, and traffic steering schemes, together with respective network planning modifications.

Along these lines, a crucial underlying LSA mechanism is the possibility of the incumbent to revoke the spectrum band while the LSA license is still effective, which may be required for the reasons discussed previously. To do so, the incumbent needs to inform the LSA repository of the change in its spectrum availability by sending what is known as an "evacuation request"

It is expected that practical LSA deployments would require a range of new dedicated mechanisms taking into account the radio technology used for an incumbent's transmissions, such as RRM, interference mitigation, load balancing, and traffic steering schemes, together with respective network planning modifications.

via a dedicated interface. Importantly, to enable LSA spectrum allocation/revocation in dynamic on-demand fashion, the LSA controller needs to have a direct low-latency and high-reliability interface with the corresponding control entity in the MNO's core network, which is, in turn, connected to, for example, the serving gateway or mobile management entity via the S1 interface. This should allow for a more dynamic response to any changes in licensing and offer better predictability in a licensee's spectrum usage.

In what follows, we study a characteristic LSA use case, where the incumbent that owns a spectrum license over a large geographical area requires its frequency resource only occasionally for small and localized portions. We also assume a reasonable cellular network presence in the same area, and the respective MNO has established a direct high-speed interface that enables the incumbent to constrain the interference generated to it by the cellular network explicitly, without additional lengthy negotiations.

Description	Value
Airplane parameters	
Airplane takeoff speed	65 m/s
Interference threshold (I_0)	–85 dBm over 10 MHz
Airplane ascent/glide slope	7 degrees
Airplane acceleration	5 m/s^2
Air-ground propagation model	Free space
Observation period	60 s
Cellular parameters	
Cell radius (R)	288 m
Operator's licensed band radio network plan	1 × 3 × 3
LSA band radio network plan	1 × 1 × 1
Cellular scheduling policy	Proportional Fair
LTE power control parameters (for licensed)	$\alpha = 1$, $SINR_{tgt} = 20$ dB
LTE power control parameters (for LSA)	$\alpha = 1$, $SINR_{tgt} = 5$ dB
Maximum BS transmit power	35 dBm (directional)
Antenna leakage	–35 dB
Propagation model	ITU urban micro
Carrier frequency	2.1 GHz
BS antenna height	15 m
BS antenna sidelobe isolation	20 dB
Shadow fading standard isolation	3 dB
LSA protective margin (K)	10 dB
UE Parameters	
Traffic pattern	Full buffer (saturation)
Maximum UE transmit power	23 dBm (isotropic radiator)
Antenna height	1.5 m

Table 1. Simulation scenario parameters.

In our example scenario, an airport leases its telemetry spectrum to a mobile network (Fig. 2), which uses this spectrum exclusively until an airplane needs to be tracked by air traffic control (ATC). When it happens, the ATC instructs the MNO to restrict its interference around the position of the airplane to allow the telemetry collection. This is feasible since the location of the airplane is known by ATC. In the end, it is up to the MNO to decide how to implement the imposed interference constraints, giving the MNO an opportunity to smoothly transition its UEs from the LSA band to the primary licensed band whenever necessary. The scale of such interference management may range from small transmit power adjustments to full "shutdown" of the LSA bands, and has a number of associated challenges:

- An exact location of the recall source (e.g., an airplane) should be known, as well as the corresponding radio propagation model to guarantee efficient isolation.
- The control interface has to operate with adequate dynamics to keep up with fast-moving objects (e.g., high-speed trains and airplanes) to avoid excessive reservations.
- The control interface must be sufficiently reliable to not affect the operational reliability of the incumbent(s).

In what follows, we construct a realistic LSA scenario to exemplify the operation of such highly dynamic system, and provide numerical insights into its expected performance. Note that this usage model, while does not intend to highlight all of the LSA features, is representative and may be adopted today on vast geographical areas, thus constituting a viable business case.

Modeling Highly Dynamic LSA Operation

Characteristic LSA Scenario

Our motivating scenario for the dynamic LSA operation is demonstrated in Fig. 2. We first note that the majority of today's airports have rather small airfields (e.g., http://www.transtats.bts.gov/airports.asp) and may not even have a tower. For such small airports, it could be relatively expensive to have dedicated radio resources for ATC functioning — these have to be controlled carefully in order to operate. Furthermore, it is not sufficient to only reserve spectrum resources around the airport premises; they must also be available in larger surrounding areas, wherever the airplanes remain relatively close to the ground. Indeed, today's airplanes require time and space to take off and land; hence, the resulting exclusion zones end up being vast, up to 25 km in radius.

In our practical small airport scenario, the airplanes do not arrive/depart every minute. Realistically, we expect an airplane every 10–20 min, or sometimes even less often than that. Moreover, since there are not too many airplanes in the air, they cannot receive signal/interference from the entire exclusion zone, only from a smaller part of it. As a result, the majority of the exclusion zone could often be underutilized, and the corresponding spectrum may thus be available to share with an MNO. Naturally, the cellular network/users would then need to adjust the transmit power based on where the airplanes are in real time to guarantee the required radio channel quality for ATC operation. We are primarily interested in the detailed performance analysis of such a system revealing the degrees of adequate interference control measures by the operator.

To facilitate the corresponding evaluation, several clarifying assumptions have to be adopted. A single runway is focused on, with airplane arrivals and departures separated by at least 5-min time intervals. As a consequence, there are never two airplanes in the exclusion zone in our model. Furthermore, all of the airplanes follow the same ascent profile, and only employ telemetry at lower altitudes. The frequency bands in use by the telemetry are shared between the ATC system and the cellular system with LSA. Given that the telemetry transmission is bidirectional, the airplane receiver must be protected from the interference produced by the licensee MNO.

To this end, the co-located cellular network operates in cooperation with the ATC and attempts to utilize the shared band in the exclusion zone whenever this does not cause excessive interference to the ATC system. Importantly, the base stations (BSs) have directional antennas with downtilt, providing at least 20 dB isolation between their radiation and the airplane in the air. Hence, we investigate the more interesting case when the LSA band is employed for uplink UE communication to augment the existing primary licensed band in use by the operator. Based on the above considerations, we construct an evaluation scenario in what follows.

Proposed System-Level Evaluation Methodology

Our below evaluation (Table 1) concentrates on the transient period when, for example, the airplane takes off from the runway and travels through the cellular network in the immediate vicinity of the airport. A similar but reverse pattern would be observed during landing, and we do not evaluate that situation here to avoid redundancy.

In more detail, our characteristic LSA scenario operates as follows. A grid of MNO cells is laid out next to a simulated airstrip with a size of 5 × 5 cells. Following the respective 3GPP recommendations, cellular users are uniformly located in the grid at an average density of 10 UEs/cell. Furthermore, at time $t = 0$ the airplane is launched from the airstrip, and simulation runs until the airplane leaves the network and reaches its cruising altitude. Meanwhile, the UEs transmit their saturated data in uplink causing interference on the airplanes. As discussed previously, cellular BSs are assumed to have near perfect isolation that prevents them from interfering with the airplane systems in downlink.

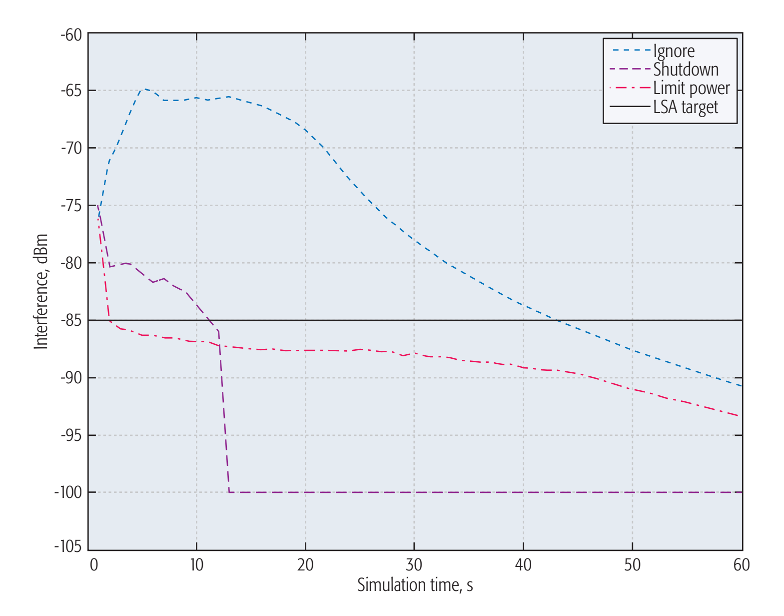

Figure 3. Interference analysis of LSA operation.

Our system-level evaluation is performed for three alternative policies of operation:

1. IGNORE policy: The airplane travels through the network receiving all possible interference from it. This is a benchmark policy and corresponds to what would happen if no coordination is introduced between the LSA incumbent and the licensee.
2. SHUTDOWN policy: All the BSs whose UEs have a chance to cause interference on LSA bands are "powered off" (in practice this may correspond to a variety of measures to disassociate users and stop transmission). This solution may seem to be the most straightforward, but has unexpected side-effects, as we show below.
3. LIMIT POWER policy: All the BSs are forced to reduce the corresponding UE's uplink power whenever instructed by the ATC to meet the interference constraints. This provides a more flexible and efficient control solution based on a heuristic approach to only limit transmit power when necessary.

Similar policies of operation may be defined for cellular downlink, except that it would be the BS transmit power that is controlled, not the UE power. For policies 2 and 3, the cellular network controller has to first learn which cells should be adjusted in response to the airplane presence. Since the exact details of the propagation environment between the cellular entities and the airplane are not known precisely, we could assume the worst case, which corresponds to free space propagation between any transmitter on the ground and the airplane. This may, in fact, be an accurate model since the atmospheric absorption at 2–3 GHz is minimal, and one can observe line-of-sight communication to the airplane from nearly anywhere on the ground, especially around airports.

For the LIMIT POWER policy, we attempt to directly bring the interference below what

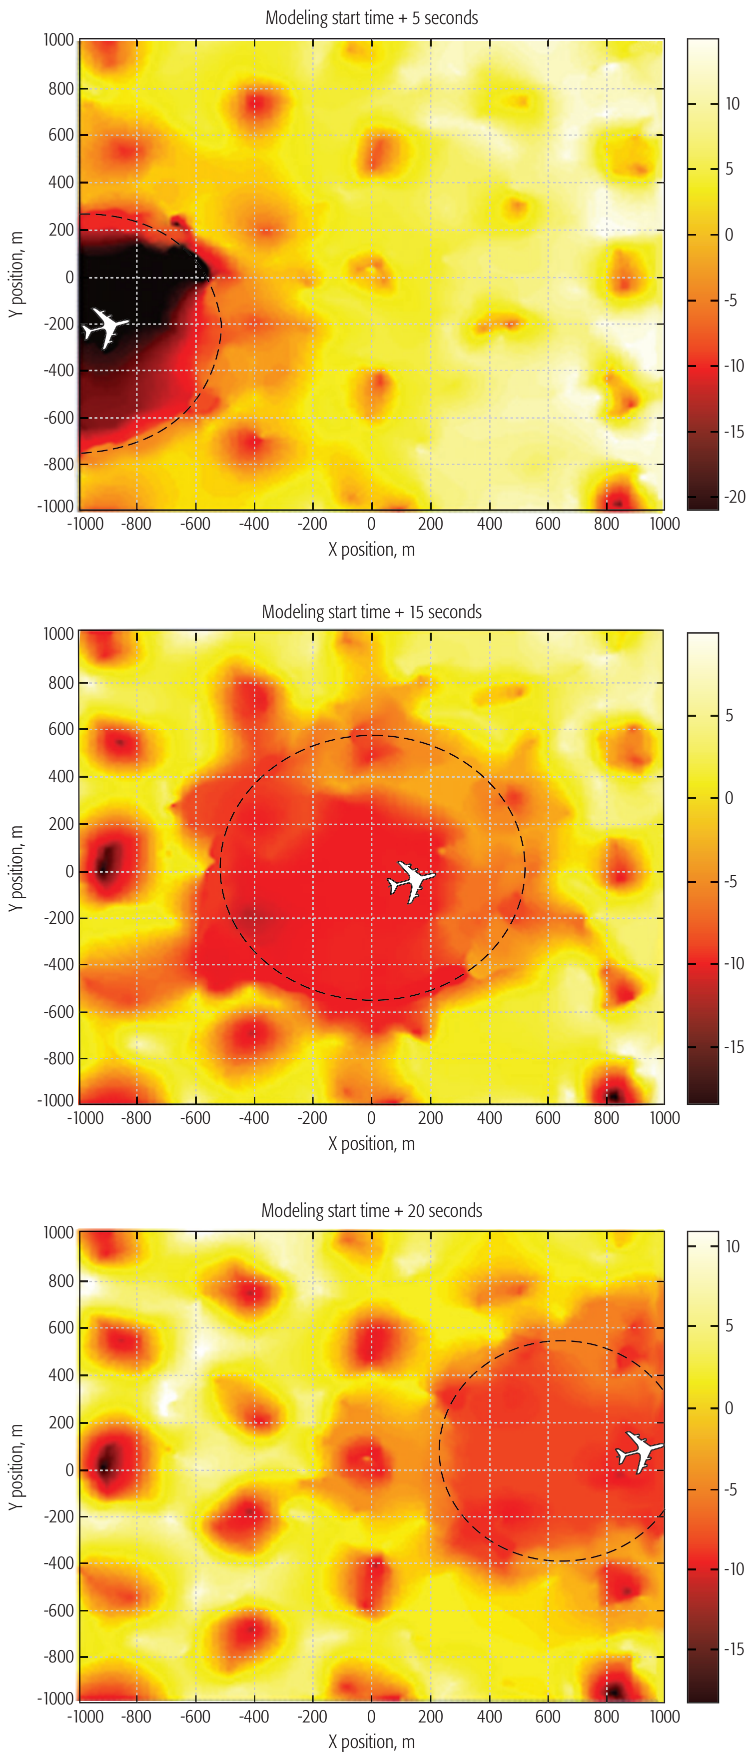

Figure 4. Performance evaluation: an airplane moves across a cellular network.

is required by the ATC. In our evaluations, we assumed a reasonable respective limit of I_0 to be –90 dBm/10 MHz, which is just over the realistic noise floor. In addition, since we have more than a single transmitter interfering with the airplane, we need to introduce a protective margin. For this scenario, as our analysis shows, a sufficient margin is $K = 10$ dB, so all the interference estimates are increased by K. For the SHUTDOWN policy, we need to simply "turn off" any BS and its associated UEs as commanded by ATC. For the IGNORE policy, nothing extra related to power control needs to be done.

The constructed simulation environment is based on our WINTERsim system-level framework[1] and allows us to implement the above interference control mechanisms explicitly. We also model the airplane as a mobile object, which is simulated live with 25 LTE cells for the time it takes the airplane to clear the network coverage area. During this entire time, the incumbent reports airplane position to the network controller, which in turn implements the interference mitigation policies on the BSs. The BSs employ the conventional LTE power control logic to enforce a particular policy on their UEs over the consecutive radio frames. The entire system is compliant with effective LTE specifications and has been calibrated against the reference 3GPP scenarios in our past publications [7].

Performance Results and Their Interpretation

In what follows, our primary focus is on the performance analysis of highly dynamic LSA spectrum sharing and the respective measures to protect the reliable operation of the incumbent's systems. To this end, Fig. 3 demonstrates the levels of interference that the incumbent receives from the cellular users for all the considered control policies: IGNORE, SHUTDOWN, and LIMIT POWER. While the IGNORE policy results in severe interference (as it does not reduce power), both SHUTDOWN and LIMIT POWER satisfy the interference requirements. It is important to note, however, that SHUTDOWN does not result in as prompt a network reaction as one would expect. This is because the shutdown of a cell makes all of the associated UEs join other nearby cells using close-to-maximum power, and as a result does not have the desired effect on the interference, even with excessive shutdown thresholds. By contrast, LIMIT POWER keeps the interference within limits as well as allows users to still transmit data, even with lower allocated power.

Furthermore, we observe how the network responds to the airplane's mobility (Fig. 4). This is best visible for LIMIT POWER policy, as we can then monitor the UE's uplink power as the airplane moves across the network. On the heatmap plots, we clearly see that the airplane casts its radio "shadow," thus causing the surrounding users to decrease their transmit power. The approximate bounds of such a shadow are shown with dashed lines, and we see that it lags behind the airplane as it accelerates.

[1] http://winter-group.net/downloads/

The snapshots in Fig. 4 are not regular in time since the airplane is gaining speed. What is crucial to note here is that we never need to reduce the UE transmit power below –10 dBm to meet the interference constraints of the airplane. Hence, cell-center users can still continue utilizing the LSA bands as usual, sometimes even enjoying higher QoS (in the absence of cell-edge users). To make it happen, the network needs to employ a type of Proportional-Fair scheduler that would allocate most of the resources to the cell-center users, since they are now the only ones with a reasonable signal-to-interference-plus-noise ratio (SINR). In case of SHUTDOWN policy, the radio shadow would actually cause complete power-off of all the affected LSA cells and UEs, resulting in dramatic loss of capacity. Let us further investigate how much performance could be realistically gained by preferring LIMIT POWER over SHUTDOWN.

One of the network's key performance indicators is how much energy has been radiated by the LSA cells altogether while the incumbent's airplane arrives/departs. This important metric can then be translated into throughput, subject to a particular scheduling policy, and Fig. 5 indicates how the three control policies compare against each other in this respect. While IGNORE is not practical, it gives a good idea of how much energy could have been radiated, should there be no airplanes. The SHUTDOWN and LIMIT POWER policies are practical, and their main difference is in the cost of added control interface complexity. For SHUTDOWN, it is a binary command, while LIMIT POWER requires regular updates of power thresholds on each BS.

In practical terms, LIMIT POWER, while not approaching IGNORE levels, still enables considerable additional power that can be collectively radiated from the UEs in the immediate vicinity of the airport. As transmit power directly translates into throughput, more radiated energy generally allows for higher data rates. In the case of LIMIT POWER, we observe that when the airplane is in the middle of the deployment, it affects most of the cells (Fig. 4). Therefore, the radiated power is minimal at around 15 s in Fig. 5.

It is also important to note that Fig. 5 reveals a step-wise behavior for LIMIT POWER. This is tightly connected with how the control mechanism operates. Essentially, the incumbent updates the reference position of the airplane every second in our scenario, and the cellular system responds by changing the power control settings accordingly. These changes in turn are reflected in the interference levels as measured by the airplane and presented in Fig. 3. In reality, the smoothness of the resulting steps will depend on the capacity/latency of the LSA control interface between the incumbent and the licensee.

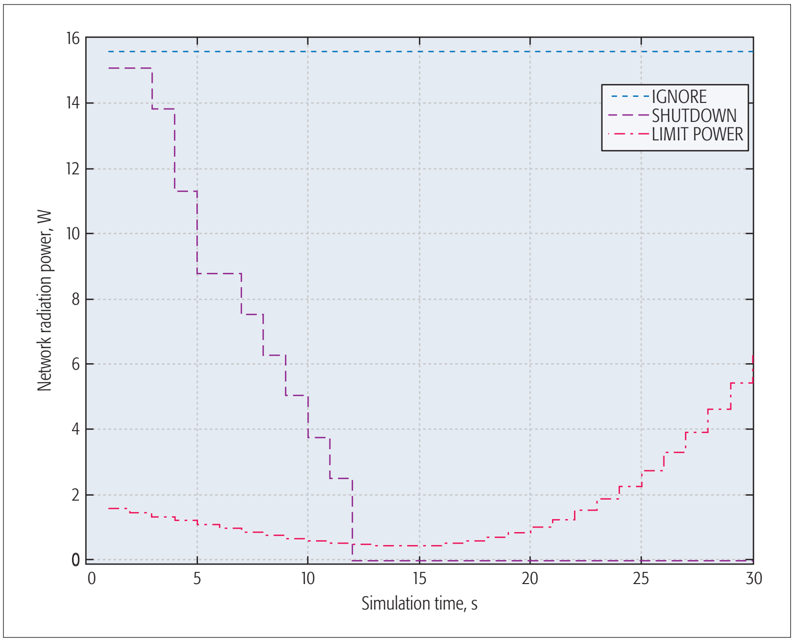

Figure 5. Comparing alternative modes of LSA operation.

Current Status and Future Evolution of LSA

Regulation and Standardization Update

In Europe, LSA is receiving considerable political support from the European Commission (EC), which is the executive body of the European Union (EU). In particular, EC has asked its high-level advisory Radio Spectrum Policy Group (RSPG) to provide their opinion on LSA (RSPG 13-538). As a result, the EC has issued a standardization mandate on mobile/fixed communication networks (MFCN) [8] in 2.3–2.4 GHz bands; the anticipated follow-up EC directives and decisions on LSA would be legally binding for the 28 member states of EU.

Currently, European adoption of voluntary spectrum sharing with LSA is facilitated by Conférence Européenne des Postes et des Télécommunications (CEPT), primarily for MFCN [9]. The corresponding CEPT Working Group on Frequency Management (WG FM) has established two project teams (PT52 and PT53) to ensure that the LSA framework is ready to be introduced to the market from a regulatory perspective. While the former addresses more specific implementation measures of LSA, the latter outlines its general aspects, including possible sharing arrangements and band-specific conditions (if not dealt with by a specific project team) for the implementation of the LSA that could be used as guidelines for CEPT administrations.

Along these lines, PT52 has finalized their Decision (14)02 on harmonized technical and regulatory conditions for the use of 2.3–2.4 GHz bands for MFCN and delivered it to the European Communication Committee (ECC) of CEPT. Their other work included ECC Recommendation (14)04 on cross-border coordination for MFCN (and with other systems) in these bands. Presently, PT52 is developing another ECC Recommendation to provide guidance to administrations in implementing a sharing framework between MFCN and PMSE (Programme Making and Special Events) [10], as well as preparing their response [11] to the EC Mandate on 2.3–2.4 GHz bands. In turn, PT53 delivered their ECC Report 205 on LSA, which was published in 2014.

From another end, the European Telecommunications Standards Institute (ETSI), the main player in European standardization, has

Understanding that LSA may eventually alter the rules of competition, most MNOs have generally adopted the careful neutral behavior expecting when its potential benefits and associated limitations will become more clear. Seeking to resolve their concerns, this article sheds light on the expected LSA operation in a characteristic highly dynamic scenario.

been tasked by EC Mandate M/512 to enable the deployment and operation of CRS under the LSA regime. Correspondingly, ETSI standardization has already issued technical specifications [12, 13] outlining system requirements and thus develops LSA system architecture [6]. This contributes to the overall picture produced by efficient interaction between ETSI standardization, CEPT regulation activities, and alignment with political objectives across Europe. Furthermore, several new LSA-related Work Item proposals have been submitted to 3GPP LTE Release 13, and how 3GPP will address these proposals is currently under discussion.

Complementing European efforts behind licensed spectrum sharing technology, similar approaches are emerging in other geographical regions. In the United States, the Federal Communications Commission (FCC) has issued their Notice of Proposed Rulemaking (NPRM) [14] targeting the use of small cells in 3.5 GHz band based on a scheme that is closely related to LSA. Accordingly, the spectrum is proposed to be managed by a spectrum access system (SAS) incorporating a dynamic database and, potentially, other interference mitigation techniques. While LSA foresees only two tiers of services (incumbents and licensees), the FCC introduces the possibility of three tiers: incumbent access, priority access, and general authorized access (GAA).

In addition to regulation and standardization LSA activities, the respective research efforts are decisively gaining momentum [15]. For instance, the leading community conferences IEEE DYSPAN (Dynamic Spectrum Access Networks) and IEEE CROWNCOM (Cognitive Radio Oriented Wireless Networks) have recently reported an impressive number of research papers on LSA, ranging from spectrum occupancy measurements in 2.3–2.4 GHz bands to LSA trial demonstrations in a live LTE network. This extensive research is supported by several visible project consortia, such as METIS-2, CORE+, CoMoRa, COHERENT, and some other 5G-PPP activities. All of the above creates fruitful soil for prompt development and adoption of the LSA ecosystem, as well as subsequent efficient deployments.

Summary on LSA Performance Promise

While we expect additional global spectrum to be allocated for future mobile services as the result of the World Radiocommunication Conference 2015 (WRC-15), it is very likely that such new frequency bands will also encompass other legacy primary services. Here, LSA may help deliver considerable benefits to the MNOs when employed in conjunction with their primary allocation under exclusive licenses. However, to ensure that LSA will not conflict with exclusive spectrum usage models, it should be based on effective market demand — LSA should come as a complementary solution for accessing spectrum on particular bands, rather than a replacement to conventional exclusive access.

In the future, LSA might also add to the beginning separation between the MNOs and the actual spectrum owners, with the latter offering more capacity to the former together with the associated spectrum availability guarantees. Correspondingly, extremely flexible and adaptable LSA implementations will be required, and our work may serve as an important building block to make it happen. Understanding that LSA may eventually alter the rules of competition, most MNOs have generally adopted careful neutral behavior in expectation of when its potential benefits and associated limitations will become more clear. Seeking to resolve their concerns, this article sheds light on the expected LSA operation in a characteristic highly dynamic scenario.

Acknowledgment

The work of Sergey Andreev is supported with a Postdoctoral Researcher grant by the Academy of Finland as well as with a Jorma Ollila grant by Nokia Foundation.

References

[1] M. Mueck *et al.*, ed., *Novel Spectrum Usage Paradigms for 5G*, 2014.
[2] J. Andrews *et al.*, "What Will 5G Be?," *IEEE JSAC*, vol. 32, 2014, pp. 1065–82.
[3] K. Chatzikokolakis *et al.*, "Toward Spectrum Sharing: Opportunities and Technical Enablers," *IEEE Commun. Mag.*, vol. 53, 2015, pp. 26–33.
[4] K. Buckwitz, J. Engelberg, and G. Rausch, "Licensed Shared Access (LSA) – Regulatory Background and View of Administrations," *Proc. CROWNCOM* (invited paper), 2014.
[5] P. Ahokangas *et al*., "Business Models for Mobile Network Operators in Licensed Shared Access (LSA)," *Proc. DYSPAN*, 2014.
[6] ETSI TS 103 235, "System Architecture and High Level Procedures for Operation of Licensed Shared Access (LSA) in the 2300 MHz–2400 MHz Band."
[7] S. Andreev *et al*., "Cellular Traffic Offloading onto Network-Assisted Device-to-Device Connections," *IEEE Commun. Mag*., vol. 52, 2014, pp. 20–31.
[8] EC Mandate on MFCN for 2.3–2.4 GHz. RSC. Mar 2014.
[9] M. Gundlach *et al.*, "Recent Advances on LSA in Standardization, Regulation, Research and Architecture Design," *Proc. Cognitive Cellular Systems*, 2014.
[10] ECC Recommendation to Provide Guidance to Administrations in Implementing a Sharing Framework Between MFCN and PMSE within 2300–2400 MHz, CEPT WG FM PT52.
[11] Response to the EC Mandate on 2300–2400 MHz, CEPT WG FM PT52.
[12] ETSI TS 103 113, "Mobile Broadband Services in the 2300 MHz–2400 MHz Band under Licensed Shard Access Regime", v. 1.1.1, 2013-07.
[13] ETSI TS 103 154: "System Requirements for Operation of Mobile Broadband Systems in the 2300 MHz–2400 MHz Band under Licensed Shared Access (LSA) Regime," v. 1.1.1, 2014-10.
[14] FCC, Enabling Innovative Small Cell Use In 3.5 GHz Band NPRM & Order, Dec. 12, 2012.
[15] M. Matinmikko *et al*., "Cognitive Radio Trial Environment: First Live Authorized Shared Access-Based Spectrum-Sharing Demonstration," *IEEE Vehic. Tech. Mag.*, vol. 8, 2013, pp. 30–37.

Biographies

Aleksei Ponomarenko-Timofeev (aleksei.ponomarenko-timofeev@tut.fi) is a researcher in the Department of Electronics and Communications Engineering at Tampere University of Technology, Finland. He received his B.Sc. degree from St. Petersburg State University of Aerospace Instrumentation, Russia. He has co-authored a number of publications on spectrum access techniques, the LSA framework, as well as H2H and M2M communications.

Alexander Pyattaev (alexander.pyattaev@tut.fi) is a Ph.D. candidate in the Department of Electronics and Communications Engineering at Tampere University of Technology. He received his B.Sc. degree from St. Petersburg State University of Telecommunications and his M.Sc. degree from TUT. He has publications on a variety of networking-related topics in internationally recognized venues, as well as several technology patents. His primary research interest lies in the area of future wireless networks: shared spectrum access, smart RAT selection, and flexible adaptive topologies.

Sergey Andreev (sergey.andreev@tut.fi) is a senior research scientist in the Department of Electronics and Communications Engineering at Tampere University of Technology. He received his Specialist degree (2006) and Cand. Sc. degree (2009) from St. Petersburg State University of Aerospace Instrumentation, as well as his Ph.D. degree (2012) from Tampere University of Technology. He has (co-)authored more than 90 published research works on wireless communications, energy efficiency, heterogeneous networking, cooperative communications, and machine-to-machine applications.

Yevgeni Koucheryavy (yk@cs.tut.fi) is a full professor and lab director at the Department of Electronics and Communications Engineering of Tampere University of Technology. He received his Ph.D. degree (2004) from the same university. He is the author of numerous publications in the field of advanced wired and wireless networking and communications. His current research interests include various aspects of heterogeneous wireless communication networks and systems, the Internet of Things and its standardization, as well as nanocommunications. He is an Associate Technical Editor of *IEEE Communications Magazine* and Editor of *IEEE Communications Surveys & Tutorials*.

MARKUS MUECK (markus.dominik.mueck@intel.com) oversees Intel's technology development, standardization, and partnerships in the field of spectrum sharing. In this capacity, he has contributed to standardization and regulatory efforts on various topics including spectrum sharing within numerous industry standards bodies, including ETSI, 3GPP, IEEE, the Wireless Innovation Forum, and CEPT. He is an adjunct professor of engineering at Macquarie University, Sydney, Australia, and acts as ETSI Board Member supported by INTEL and as general Chairman of ETSI RRS Technical Body (Software Radio and Cognitive Radio Standardization). He has earned engineering degrees from the University of Stuttgart, Germany, and the Ecole Nationale Supérieure des Télécommunications (ENST), Paris, France, as well as a doctorate degree from ENST in communications.

INGOLF KARLS (ingolf.karls@intel.com), Intel Deutschland GmbH, is program manager of the Next Generation and Standards (NGS) team in Intel's Communication and Devices Group (iCDG). He has accomplished many projects in mobile communication at Siemens AG, Infineon Technologies AG, and Intel for more than 20 years. He has managed the operation of 14 national and European public funded research and development projects in European FP7. He has contributed to national and international regulation and standardization bodies like 3GPP, BITKOM, DLNA, ETSI, IEEE, OMA, and OMTP. He holds several patents in the area of wireless communication. His current research and development interests are the evolution of LTE and new innovative air interface solutions for 5G wireless technologies and standards, and their translation into products.

Why to Decouple the Uplink and Downlink in Cellular Networks and How To Do It

Federico Boccardi, Jeffrey Andrews, Hisham Elshaer, Mischa Dohler, Stefan Parkvall, Petar Popovski, and Sarabjot Singh

ABSTRACT

Ever since the inception of mobile telephony, the downlink and uplink of cellular networks have been *coupled*, that is, mobile terminals have been constrained to associate with the same base station in both the downlink and uplink directions. New trends in network densification and mobile data usage increase the drawbacks of this constraint, and suggest that it should be revisited. In this article we identify and explain five key arguments in favor of downlink/uplink decoupling based on a blend of theoretical, experimental, and architectural insights. We then overview the changes needed in current LTE-A mobile systems to enable this decoupling, and then look ahead to fifth generation cellular standards. We demonstrate that decoupling can lead to significant gains in network throughput, outage, and power consumption at a much lower cost compared to other solutions that provide comparable or lower gains.

Ever since the inception of mobile telephony, the downlink and uplink of cellular networks have been coupled, that is, mobile terminals have been constrained to associate with the same base station in both the downlink and uplink directions. New trends in network densification and mobile data usage increase the drawbacks of this constraint, and suggest that it should be revisited.

INTRODUCTION AND BACKGROUND

From the first to the fourth generation (1G–4G) of mobile networks, the downlink (DL) and uplink (UL) of a given communication session have been coupled: the mobile user equipment (UE) must associate with the same base station (BS) in both the DL and UL. Historically, this was a nearly optimal approach, since the strongest BS-UE connection was the same in both directions. However, this conventional approach has recently come under scrutiny [1] given the possible gains that can be achieved by decoupling the association in the context of a dense heterogeneous cellular network, wherein different BSs can have highly variable transmit powers and deployment topologies.

The arguments in favor of the coupled status quo are several. From a pure network design perspective, the logical, transport, and physical channels are easier to design and operate; this pertains particularly to the synchronization of acknowledgments, call admission and handover procedures, DL/UL radio resource management, and power control, among others. Decoupling the links also requires strong synchronization and data connectivity (e.g., via fiber) between the BSs. From a deployment and topology perspective, until just a few years ago cellular systems were designed and deployed under the assumption of a homogeneous network with macrocells all transmitting with about the same power. From a traffic point of view, the load in both directions was approximately the same in voice-centric 2G and early 3G systems. Moreover, 3.5G (e.g., high-speed packet access, HSPA) and 4G systems are dominated by downlink traffic, justifying the use of DL-centric association procedures rather than UL or decoupled ones.

The emergence of heterogeneous networks (HetNets) [2], where small cells at higher carrier frequencies and/or smaller transmit powers are deployed within the coverage area of macrocells, calls for revisiting the coupled association approach. Range extension has been included in 4G to add a bias in the cell association to offload more traffic from macro to small cells. Data and control plane separation was introduced in [3]: the control information is sent by high-power nodes at lower frequencies, whereas the payload data is conveyed by low-power nodes at possibly higher frequencies. However, both range extension and data/control plane separation are based on a coupled DL/UL association, where DL and UL are associated with the same BS.

The motivation for downlink/uplink decoupling (DUDe) emerges from a holistic view of the two-way (DL/UL) traffic and the association procedure of a UE, rather than adopting a coupled association a priori and then separately optimizing DL and UL transmissions. Since a coupled association is a particular sub-case of a decoupled one, a well designed association policy based on DUDe can in principle outperform a coupled association. But by how much? And at what cost?

More specifically, the main questions this article attempts to answer are:
- What recent trends in cellular network deployment and applications make the gains from DUDe more relevant now than in the past?
- What are the key benefits of a decoupled association in terms of throughput gain, reliability, and power conservation? What are the challenges? How can these gains be realized in current (e.g. Long Term Evolution-Advanced, LTE-A) and future 5G cellular networks?

Federico Boccardi is with Ofcom; however, the views expressed here are his own and do not reflect those of his employer. Jeffrey Andrews is with the University of Texas at Austin; Hisham Elshaer is with Vodafone Group R&D and King's College London; Mischa Dohler are with King's College London; Stefan Parkvall is with Ericsson; Petar Popovski is with Aalborg University; Sarabjot Singh is with Intel.

 0163-6804/16/$25.00 © 2016 IEEE

- How disruptive will these changes be to the network architecture? Are the gains large enough to be worth the trouble?

We note that research developments on DUDe are recent and limited to a few contributions. The interest in decoupling the downlink and uplink was indicated in [4–6], and further explored in a few subsequent contributions. In particular, [7–9] studied the throughput and signal-to-interference-plus-noise ratio (SINR) gains from a theoretical perspective, while [1] assessed DUDe via detailed industry-standard simulations.

We begin the discussion in this article with the five key arguments in favor of DUDe, and provide evidence for the corresponding gains from very recent theoretical analysis and simulation-based experiments. Then we move on to discuss what changes will need to be made in the current and future cellular standards, and explain why, in our view, such changes are quite manageable. DUDe opens up many new interesting research questions as well, which we identify throughout the article.

Five Reasons to Decouple the Downlink and Uplink

We now articulate the five principal arguments in favor of DUDe. Our arguments are supported by a combination of recent theoretical and system-level simulation results by the present authors and others. In particular, the theoretical results are mostly sourced from recent work [7], in which we perform a comprehensive SINR and rate analysis with DUDe in a multi-tier cellular network with spatially random UEs and BSs. The UEs employ fractional UL power control, and small-cell biasing is used to achieve cell range expansion: both very similar to LTE. The results are mathematical and thus transparent, albeit in some cases based on idealized models to allow tractability. We refer to this approach as the *analytical model*.

The simulation results and parameters follow largely from [1], and utilize an existing LTE HetNet deployment in conjunction with a high-resolution 3D ray tracing channel model that takes into account clutter, terrain, and building data. This ensures a highly realistic and accurate propagation model. The BS types and locations are based on a small cell test network in the London area and consist of five macrocells covering a 1 km^2 area with a dense small cell deployment embedded in the area. The UE distribution is based on live traffic measurements, and the UEs use the same UL fractional power control as in the analytical model. We assume that the DL association is based on the DL reference signal received power (RSRP). We refer to this approach as the simulation model.

As we see below, these two distinct approaches to modeling and analyzing DUDe are quite unified in terms of the conclusions they offer. Table 1 contains the cellular network notation and simulation parameters. We also use the same parameters for numerical evaluation of mathematically derived results using the analytical model.

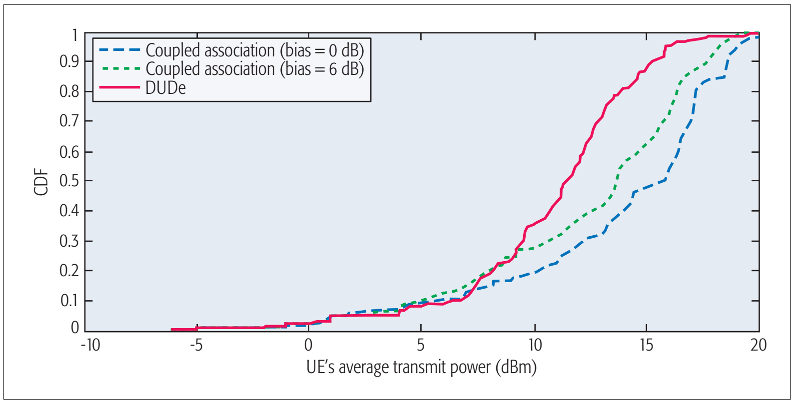

Figure 1. CDF of the UE's UL transmit power via the simulation model. Cell edge users (right side of figure) require higher transmit powers and thus achieve larger power reduction from DUDe.

Increased Uplink SNR and Reduced Transmit Power

In a typical HetNet scenario the DL coverage area of a macrocell is much larger than that of a low-power BS; indeed, this is why they are often called *small cells*. The coverage area disparity is primarily attributable to the differences in DL transmit powers, but is also due to the BS heights and antenna gains. In contrast, in the UL all transmitters have roughly the same maximum transmit power. Therefore, a device that is associated with a macrocell in the DL might instead wish to be associated with a small cell in the UL to take advantage of the reduced path loss [8]. The positive effects are twofold. For UEs that are transmitting at the maximum power, a connection to a closer BS provides a higher signal-to-noise ratio (SNR). Moreover, for a fixed target SNR, the reduced path loss alternatively allows transmit power reduction via power control.

In Fig. 1, we observe the decrease in transmit power via DUDe by comparing three cases via the simulation model. The first case is the baseline with a coupled DL/UL association and no small cell bias. The need for bias arises in a HetNet scenario where, due to load balancing, the UEs are steered toward being associated with a small cell if their received power is lower (up to the bias value in dB) than that received from the macrocell BS. The second case is still coupled, but the small cells have a 6 dB bias. We note that 6 dB has been shown in [10] to be a reasonable value for the bias. The third case is for DUDe. DUDe yields 2.3 dB at 50 percent and 3 dB at 95 percent cumulative distributed function (CDF) relative to the coupled association with 6 dB bias.

Improved Interference Conditions

DUDe also decreases the UL interference due to multiple complementary effects.

First, and as an obvious consequence of the transmit power reduction demonstrated in the previous section, the UL interference generated to other BSs is correspondingly reduced by about 2–3 dB. This is quite significant, especially for the low SINR UEs in the UL, since at low SINR in a dense network, decreasing the interference by 3 dB implies an approximate doubling of data rate.

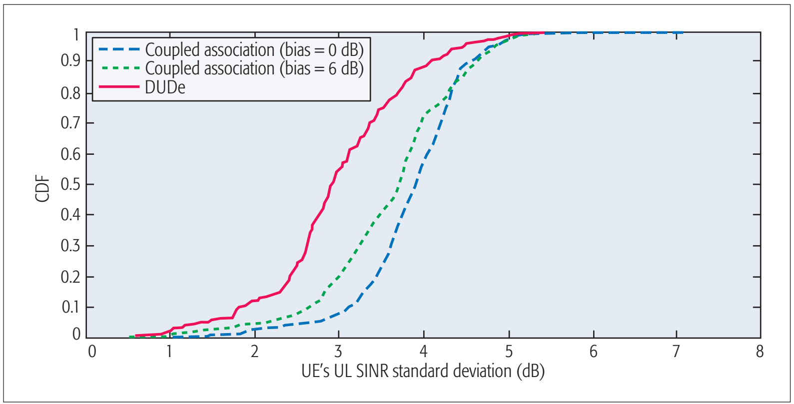

Figure 2. CDF of the UE's SINR standard deviation over time. DUDe reduces the variations and improves performance.

Second, DUDe provides the ability to independently select the association that minimizes interference at both the UE and the BS. Uplink interference in a given spectral band is an aggregation of many different UEs' transmissions in different cells, as received by a given BS, say BS0. The interference generated by each of these UEs depends on its location relative to its own desired BS, the amount of power control, its distance to BS0, and the UL precoding weights. In contrast, the DL interference at a given UE depends on the BSs' transmit power, the DL beamforming weights, and the distance to the different BSs. On top of this, the nearly independent scheduling and loading in the DL and UL causes further randomness in the interference. For all these reasons, average interference levels can be quite different in the DL and UL resources. Therefore, a decoupled association that allows the UE/network to seek out the best interference environment in the two links independently can be expected to substantially outperform a coupled association, which must "split the difference."

Third, DUDe will also prove a boon for device-to-device (D2D) communication, which, as of Third Generation Partnership Project Release 12 (3GPP Rel. 12), will take place in the UL bands. By lowering the UL transmit power and generating less interference, DUDe will create a more benign environment for D2D receivers and thus allow more D2D transmissions to take place.

Finally, in addition to reducing the amount of average interference, DUDe also allows a reduction of the UL SINR variance, as shown in Fig. 2 (obtained via the simulation model), which translates into more efficient and effective UL schedulers and performance gains [11]. Specifically, with respect to a coupled association with a 6 dB bias, the decoupling yields a reduction of 1 dB on average, which is about 25 percent at 50 percent CDF.

Improved Uplink Data Rate

Unsurprisingly, increasing the desired received power and decreasing the interference leads to higher SINR, and hence a higher spectral efficiency and data rate. However, there are additional factors that can complicate the effect of DUDe on the UL rate.

For example, consider an LTE HetNet with small cell range expansion and biasing. On aver-

Parameter	UEs	Macrocells	Small cells
Max transmit power	20 dBm	46 dBm	30 dBm (unless otherwise specified)
Antenna system (simulation)	1 Tx and 1 Rx Antenna gain = 0 dBi	2 Tx and 2 Rx Antenna gain = 17.8 dBi	2 Tx and 2 Rx Antenna gain = 4 dBi
Antenna system (analysis)	The analysis considers 1 Tx and 1 Rx isotropic antenna system for UEs, macrocells, and small cells.		
Downlink bias	N/A	0 dB	Varies from 0 to 8 dB
Spatial distribution (analytical)	Uniform	Poisson point process	Poisson point process
Spatial distribution (simulation)	Hotspot distribution based on realistic traffic measurements	Based on the Vodafone LTE network deployment.	Based on the Vodafone LTE small cell test network deployment.
Spatial density	330 per km^2	5 per km^2	4 small cells per macrocell
Channel model (analytical)	Rayleigh small-scale fading, standard path loss with exponent = 3.5 Lognormal shadowing with standard deviation = 8 dB		
Channel model (simulation)	3D ray-tracing propagation model		
Power control	Uplink fractional path loss compensation		
Operating frequency	2.6 GHz co-channel FDD deployment		
Bandwidth	20 MHz (100 frequency blocks)		
Scheduler	Equipartition of resources among UEs (analysis); proportional fair (simulation)		

Table 1. Cellular network notation and parameter values.

	DUDe vs. picocells (bias = 0)		DUDe vs. picocells (bias = 6 dB)		DUDe vs. femtocells (bias = 0)		DUDe vs. femtocells (bias = 8 dB)	
	Analysis	Simulation	Analysis	Simulation	Analysis	Simulation	Analysis	Simulation
5th %-ile (cell edge) rate gain	115%	90%	50%	30%	270%	260%	140%	95%
50th %-ile (median) rate gain	95%	150%	30%	60%	260%	230%	120%	180%

Table 2. Summary of predicted uplink rate gains averaged over all UEs in the network, as a result of DUDe. Picocells have transmit power of 30 dBm; femtocells, 20 dBm. We note that DUDe also outperforms the baseline when downlink biasing is used.

age, the optimal DL bias is in the neighborhood of 5–10 dB as noted before, although with blanking or interference avoidance, up to 18–20 dB may be used in certain scenarios [2, 10]. DL biasing leads to a better association in both directions even with coupled association, since by expanding the DL small cell coverage region, more UEs associate with the nearby small cells in the UL as well, which is also the main point of DUDe.

Nevertheless, we still observe very substantial rate gains for DUDe even when compared to biased coupled associations. Detailed breakdowns of these rate gains in various configurations are given in [1, 7], with our findings summarized in Table 2. Here, picocells have transmit power of 30 dBm and femtocells 20 dBm. The gains result mainly from the improved channel quality and also from the biasing as discussed above, which gives cell edge (5th percentile) and median (50th percentile) UEs access to more resources, which results in a higher UL rate. It is quite encouraging that two very different models and approaches to evaluating the rate gains both result in the conclusion that gains in the range of 100–200 percent are within reach, although the gains do erode somewhat with biasing since the baseline improves. Finally, we note that a recent paper based on optimization theory with a different model also finds significant gains from DUDe [9].

Different Load Balancing in the Uplink and the Downlink

The load that a given BS has in the UL may be different from the load that the same BS has in the DL. This implies that it is not optimal to have the same set of UEs connected to the same BS in both the UL and DL, so at least some of the UEs should use decoupled access.

Additionally, DUDe allows pushing more UEs to underutilized small cells in the UL only since it is not limited by interference as is the case in the DL. In Fig. 3 we show that this results in a better distribution of the UEs among macro and small cells, which in turn allows for more efficient resource utilization and higher UL rates. We note that DUDe outperforms the baseline for both unbiased and biased association.

Low Deployment Costs with RAN Centralization

Implementing a decoupled cell association in a real network requires excellent connectivity and modest cooperation between different BSs. As we discuss in the subsequent section, the main requirement DUDe imposes is a low-latency connection between the DL and UL BSs to allow fast exchange of control messages,

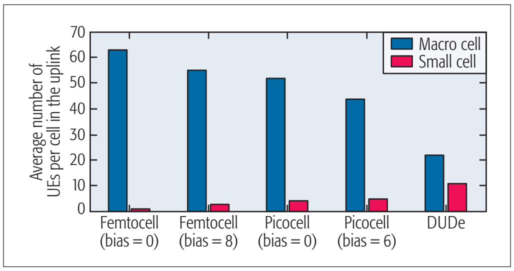

Figure 3. The femtocell, picocell, and DUDe cases. We note that the total number of UEs in the system is kept constant across the different cases (Table 1).

like hybrid automatic repeat request (HARQ) messages. We emphasize that unlike the most sophisticated forms of cooperative multipoint (CoMP), like joint processing, where a high throughput backhaul connection between BSs is required to allow rapid data exchange, DUDe does not impose a tight requirement on the backhaul capacity. Put another way, DUDe allows gains similar to uplink joint processing (about 100 percent edge and average throughput gain, as just seen), but with lower deployment costs. Compared to using multiple-input multiple-output (MIMO) or new spectrum to increase the throughput, the cost comparison is even more favorable to DUDe.

The ongoing trend toward using partial or full radio access network (RAN) centralization in deployments where a high-speed backhaul is available will be an enabler for DL and UL decoupling, as signaling will be routed to a central processing unit with low-latency connections. In particular, partial centralization refers to those local deployments (e.g., indoor) where the transmission points serving the same local area are all connected to the same baseband processing central unit. Full centralization, often referred as cloud-RAN, extends this approach to larger areas, where a large number of RF units are connected to the same baseband processing central unit.

Given this already ongoing trend toward more centralized RAN architectures, which are underpinned by low-latency connectivity between BSs, the incremental cost of DUDe appears negligible in such scenarios.

DUDe in LTE-A: Enabling Architectures

DUDe can, depending on the deployment scenario and backhaul properties, already be supported by the existing LTE/LTE-A specifications. Illustrated in Fig. 4, three specific embodiments are discussed below.

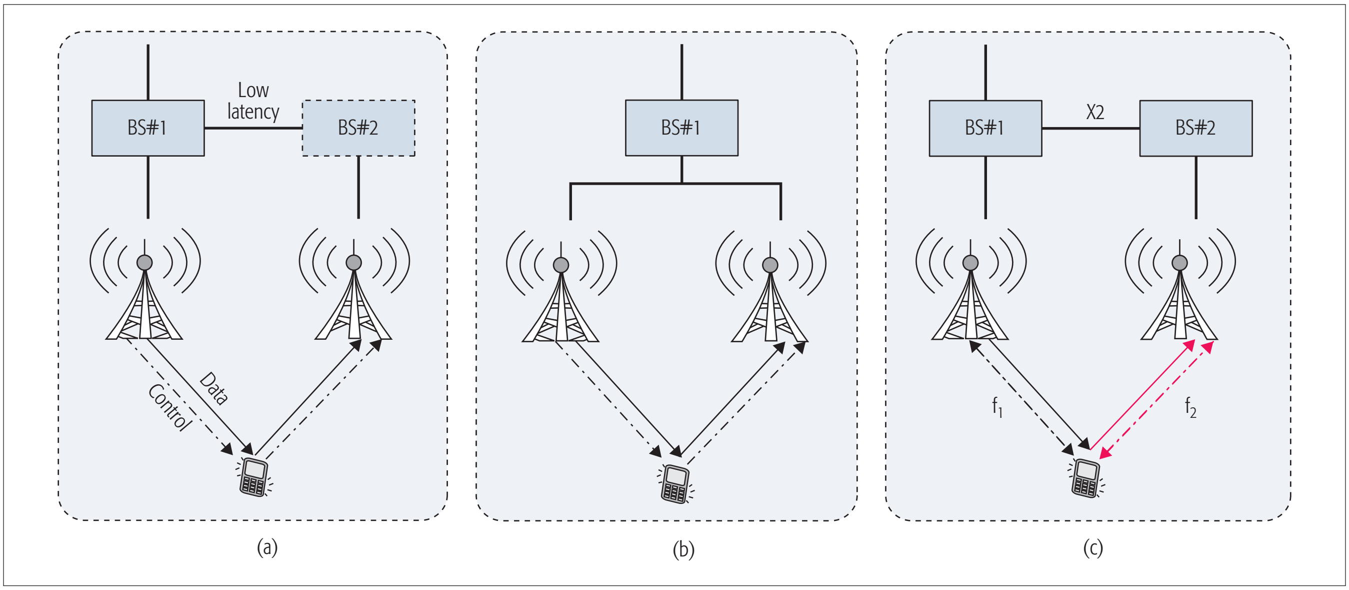

Figure 4. The three discussed embodiments of DUDe are: a) centralized processing unit; b) shared cell-ID; c) the dual connectivity option.

Centralized Processing

As mentioned before, in a deployment scenario with multiple radio units with a different cell-ID connected to a centralized node (as in the case of a centralized RAN, C-RAN), DUDe is possible in LTE-A without additional standardization support (Fig. 4a). The BS used for DL transmission to a specific UE is selected using conventional means, typically based on DL signal strength measurements. Uplink transmissions are received by one, or if macro diversity is desirable, multiple radio units as the specifications do not mandate the reception node. Either UL decoding could be performed at the radio unit (or at the set of radio units), or sampled analog data could be forwarded to the centralized unit via a common public radio interface (CPRI) for further processing.

Uplink-related control signaling (including, e.g., HARQ and power control commands) needs to be transmitted from the DL node. In the same way, DL-related control signaling from the terminal needs to be received by the UL node and forwarded to the DL node over the infrastructure.

Shared Cell-ID

An interesting extension of the approach described above is the so-called shared cell-ID approach [6] (Fig. 4b), where radio units all belong to the *same* cell (i.e., have the same cell-ID). Here, channel state information (CSI) enhancements and quasi-co-location mechanisms introduced in Release 11 as part of the CoMP work are used to rapidly, independently, and, from a terminal perspective, transparently switch transmission and reception points for a given terminal. This is a step away from the traditional cell-oriented paradigm toward viewing the antenna points as resources to be used in the best possible way to maximize performance. Furthermore, node association and mobility are handled via proprietary (non-standardized) solutions, transparent to the mobile terminal, providing better mobility robustness in dense networks compared to methods relying on UE-centric measurements.

Although conceptually straightforward, both centralized processing and shared-ID approaches require a fairly low-latency backhaul to meet the timing requirements (e.g., to send HARQ messages). In a practical LTE-A rollout, the deployment is thus limited to remote radio units connected to a centralized baseband processing node.

Dual Connectivity

While the two solutions described above require a very low-latency backhaul, usually achieved via connecting radio units to the same central unit, DUDe can also be implemented with a less ideal backhaul. *Dual connectivity*, an extension first introduced in Release 12, allows for a terminal to be simultaneously connected to two cells and can be used for DUDe (Fig. 4c). We note that in Release 12, DUDe using dual connectivity is limited to inter-frequency deployments, that is, to deployments where the two cells transmit over different frequency bands; nevertheless, later releases may add support for intra-frequency band deployments. The two cells operate separately, handling their own scheduling and control signaling (e.g., HARQ message), thereby significantly relaxing the backhaul requirements compared to the centralized baseband approach and enabling the standardized X2 interface to be used for inter-BS communication. This solution has advantages and disadvantages. On one hand, a low-latency backhaul connection for the signaling is not needed. On the other hand, mobility must be handled using standardized mechanisms, and the possibilities for proprietary optimization are limited.

DUDe in 5G and Beyond

The next few years will see intense research and development on 5G. The ITU is starting their work on requirements under "IMT-2020," and in 3GPP initial activities on 5G standardization began at the end of 2015 with the overall goal of a large-scale trial around 2018 and commercial operation in 2020. Although any discussion of 5G is by definition speculative, there is an emerging consensus on the data rate requirements and likely key technical features of 5G, including extreme BS densification, massive MIMO, the introduction of millimeter-wave bands, and possibly a "cell-less" architecture [5, 12].

With this view of 5G, in this section we discuss whether 5G (and beyond) standards should include other features to *natively* support DUDe. In other words, we discuss whether a design that is optimized for DUDe from its inception, rather than amended *a posteriori*, could lead to even higher gains.

Major Architectural Changes?

An important question is whether a simple evolution of today's 3GPP architecture design discussed above would be able to efficiently support DUDe in emerging heterogeneous 5G deployments. In the previous section, we discussed how the LTE-A architecture already supports a DUDe implementation when different BSs are connected via fiber to the same radio unit. For the case of different base stations not connected to the same radio unit, we discussed how support for DUDe in 4G is limited to different frequencies. Any future 5G releases in 3GPP should thus simply allow for same-frequency dual connectivity, which, despite having implications on resource and interference management, is not considered to be a major upgrade.

A further tweak is needed to ensure proper encryption of all data and control channels, particularly when communication via the X2 interface is used between BSs. While each eNB can support tens of IP security (IPsec) tunnels, the management of security via IPsec is so cumbersome that operators tend to deploy only a few IPsec gateways (GWs) per country. Indeed, LTE has seen most IPsec GWs deployed close to the serving GW (S-GW), which means that traffic logically going via the X2 is actually routed via the S-GW; this incurs a delay that renders DUDe inefficient. While LTE-A enjoys some more IPsec GWs to deploy closer to the mobile edge, future 5G designs ought to improve security mechanisms and implementations, allowing the encryption of X2 traffic with lower latency.

Furthermore, some integration work is needed with emerging paradigms that have proven useful for current coupled systems. First, as mentioned before, the integration of DUDe with the decoupled control/data plane and licensed assisted access (LAA) will require some architecture modifications. Second, self-organizing networking (SON) paradigms will be instrumental in coordinating, in a non-conflicting manner [13], the increased degrees of freedom in the system.

Given the above discussion, however, we conclude that native support of DUDe does not require major design changes in 5G from an architectural perspective.

DUDe and Hyper-Densification

The importance of decoupled selection of the DL/UL access points may grow significantly in the coming years, as 5G will feature hyper-dense deployments in order to meet the high rate demands in crowded spots. One could argue that at extremely high densities of cells, DUDe will lead to lower gains since nearly all the devices will be associated with the nearest small cell in UL and DL. However, this will only be true if we assume that all the small cells will have the same power, traffic, and deployment characteristics. This is an unrealistic assumption, since future cellular deployments will be characterized by a mixture of user deployed and operator deployed cells, with different power levels, using frequencies ranging from below 1 GHz to tens of gigahertz, providing services for very different types of traffic and natively supporting D2D communications. DL and UL traffic flows will be routed via a mixture of licensed and unlicensed carriers, requiring different allocation criteria.[1] Therefore, we expect that DUDe gains in future deployments will be even higher with respect to the ones presented above, especially if we consider the generalized version in which a UE is associated with multiple points and selects the DL or UL direction dynamically, as part of a scheduling and optimization process.

From a broader perspective, we believe decoupled access necessarily shifts the focus of algorithmic solutions and optimizations toward models that consider two-way traffic from each UE. This is part of a larger trend in wireless network optimization that encompasses full-duplex communication, two-way relaying, and dynamic time-division duplex (TDD).

We believe decoupled access necessarily shifts the focus of algorithmic solutions and optimizations towards models that consider two-way traffic from each UE. This is part of a larger trend in wireless network optimization that encompasses full-duplex communication, two-way relaying, and dynamic TDD.

TDD, FDD, or a New Way of Duplexing?

DUDe can work with both FDD and TDD, with different implications from a system-level perspective and from a spectrum-related perspective.

TDD allows much more flexibility in trading DL and UL resources compared to frequency-division duplex (FDD). With decoupling, as we have seen, fewer UL resources are needed to achieve the same UL rate vs. the coupled case, and those resources could be reassigned to the DL via dynamic TDD, which is in line with the two-way network optimization discussed above. Traditionally, another benefit of TDD is the possibility of estimating the DL channel via UL reference signals. This is particularly important for channels with large dimensionality, such as with massive MIMO. Unfortunately, when DUDe is used, DL and UL transmissions originate and terminate at different locations, respectively, breaking the channel reciprocity. Much of the existing spectrum is paired FDD spectrum, so for both of these reasons massive MIMO may need to be supported without channel reciprocity.

In the medium/long term, DUDe together with different emerging technology trends could require rethinking the traditional FDD/TDD dichotomy. DUDe, hyper-densification, and the use of higher frequencies and highly directional antenna arrays could enable duplexing approaches over the spatial domain. For example, the same band could be used for two different devices, one receiving in DL from a BS, and the other one transmitting in UL to another BS. Assuming an effective DL/UL spatially coordinated scheduling mechanism, this could allow full-du-

[1] Recently, 3GPP finalized the work on LAA, where licensed and unlicensed carriers are aggregated. LAA uses licensed spectrum for control-related transmissions while sending data over both licensed and license-exempt carriers.

Interestingly, major changes to the radio access and core networking technologies are not needed. DUDe can be considered an innovative approach that affects the fundamentals of cellular networks and thus opens up many opportunities for research and design.

plex-like performance leveraging on the spatial domain [14]. In addition, once analog/digital interference cancellation mechanisms become truly operational to support full temporal duplex, the DUDe concept is also beneficial since the generalized decoupling would allow the support of a DL and not necessarily the same UL user in the same spectral band.

DUDe with Millimeter-Wave Frequencies

Above, we discussed why DUDe could make channel estimation via channel reciprocity in TDD more difficult. This effect could be even more pronounced at millimeter-wave (mmWave) frequencies, where the large number of antenna elements used for beamforming would be enabled by channel reciprocity.

However, there are other factors that point to DUDe as an important enabler for mmWave. For example, recent studies on electromagnetic field exposure [15] show that to be compliant with applicable exposure limits at frequencies above 6 GHz, the maximum transmit power in the UL might have to be several dB below the power levels used for current cellular technologies. Since the transmit power has an important impact on UL coverage (in particular for sounding over a non-precoded channel), we believe a pragmatic approach would be to allocate UL over a lower frequency with a better link budget. That is, while in the rest of this article we discussed associating a UE with a macrocell in the DL and with a small cell in the UL, for mmWave the opposite strategy might prove fruitful: associating the UE to the mmWave small cell in the DL and to a sub-6-GHz macrocell in the UL.

Conclusions

In traditional cellular networks, it is practically an axiom that the uplink connection is always associated with the same base station that has been selected for downlink reception. In this article we revisit this axiom and introduce the features of downlink/uplink decoupling (DUDe), a new architectural paradigm where downlink and uplink are not constrained to be associated with the same BS. This is becoming especially relevant in the wake of the densification expected in future cellular networks, where each terminal has multiple access points in proximity. We have identified five key arguments that demonstrate the usefulness of DUDe, based on a blend of theoretical, experimental, and architectural insights. We have shown how DUDe can lead to significant gains in network throughput, outage, and power consumption at a much lower cost compared to other solutions that provide comparable or lower gains. We have discussed the changes needed in the existing LTE-A systems in order to enable DUDe-based operation. We have then presented arguments why DUDe should natively be considered as a part of the future 5G systems. Interestingly, major changes to the radio access and core networking technologies are not needed. DUDe can be considered an innovative approach that affects the fundamentals of cellular networks, and thus opens up many opportunities for research and design.

References

[1] H. Elshaer *et al*., "Downlink and Uplink Decoupling: A Disruptive Architectural Design for 5G Networks," *Proc. IEEE GLOBECOM*, Dec. 2014, pp. 1798–18803.

[2] A. Damnjanovic *et al.*, "A Survey on 3GPP Heterogeneous Networks," *IEEE Wireless Commun.*, vol. 18, no. 3, June 2011, pp. 10–21.

[3] Y. Kishiyama *et al*., "Future Steps of LTE-A: Evolution toward Integration of Local Area and Wide Area Systems," *IEEE Wireless Commun*., vol. 20, no. 1, Feb. 2013, pp. 12–18.

[4] J. G. Andrews, "Seven Ways that HetNets Are a Cellular Paradigm Shift," *IEEE Commun. Mag.*, vol. 51, no. 3, Mar. 2013, pp. 136–44.

[5] F. Boccardi *et al*., "Five Disruptive Technology Directions for 5G," *IEEE Commun. Mag.*, vol. 52, no. 2, Feb. 2014, pp. 74–80.

[6] S. Parkvall *et al*., "Heterogeneous Network Deployments in LTE," *Ericsson Rev.*, Feb. 2011.

[7] S. Singh, X. Zhang, and J. G. Andrews, "Joint Rate and SINR Coverage Analysis for Decoupled Uplink-Downlink Biased Cell Associations in HetNets," *IEEE Trans. Wireless Commun*., vol. 14, no. 10, Oct. 2015, pp. 5360–73.

[8] K. Smiljkovikj, P. Popovski, and L. Gavrilovska, "Analysis of the Decoupled Access for Downlink and Uplink in Wireless Heterogeneous Networks," *IEEE Wireless Commun. Lett.*, vol. 4, no. 2, Jan. 2015, pp. 173–76.

[9] H. Boostanimehr and V. Bhargava, "Joint Downlink and Uplink Aware Cell Association in HetNets with QoS Provisioning," *IEEE Trans. Wireless Commun*., in press.

[10] S. Singh and J. G. Andrews, "Joint Resource Partitioning and Offloading in Heterogeneous Cellular Networks," *IEEE Trans. Wireless Commun*., vol. 13, no. 2, Feb. 2014, pp. 888–901.

[11] N. Abu-Ali *et al.*, "Uplink Scheduling in LTE and LTE-Advanced: Tutorial, Survey and Evaluation Framework," *IEEE Commun. Surveys & Tutorials*, vol. 16, no. 3, Aug. 2014, pp. 1239–65.

[12] J. G. Andrews *et al*., "What Will 5G Be?," *IEEE JSAC*, vol. 32, no. 6, June 2014, pp. 1065–82.

[13] H. Y. Lateef *et al*., "LTE-Advanced Self-Organizing Network Conflicts and Coordination Algorithms," *IEEE Wireless Commun*., in press.

[14] H. Thomsen, *et al*., "CoMPflex: CoMP for In-Band Wireless Full Duplex," submitted to *IEEE Wireless Commun. Lett.*, July 2015; http://arxiv.org/abs/1507.08574.

[15] D. Colombi, B. Thors, and C. Tornevik, "Implications of EMF Exposure Limits on Output Power Levels for 5G Devices Above 6 GHz,," *IEEE Antennas Wireless Propagation Lett*., vol. 14, Feb. 2015, pp. 1247–49.

Biographies

Federico Boccardi [SM] is a principal technology analyst at Ofcom. He received his M.Sc and Ph.D. degrees in telecommunication engineering from the University of Padova, Italy, in 2002 and 2007, respectively, and his postgraduate diploma in strategy and innovation from the Oxford Sad Business School in 2014. Before joining Ofcom, he was with Bell Labs (Alcatel-Lucent) from 2006 to 2013 and with Vodafone R&D in 2014. He has published more than 60 peer-reviewed international research papers and issued patents. He is an Associate Editor for *IEEE Transactions on Wireless Communications* and *IEEE Transactions on Cognitive Communications and Networking*. His current interests fall at the intersection between technology and strategy.

Jeffrey G. Andrews [S'98, M'02, SM'06, F'13] received his B.S. in engineering from Harvey Mudd College in 1995, and his M.S. and Ph.D. in electrical engineering from Stanford University. He is the Cullen Trust Endowed Professor (#1) of Electrical and Communications Engineering at the University of Texas at Austin and Editor-in-Chief of *IEEE Transactions on Wireless Communications*. He developed code-division multiple access systems at Qualcomm from 1995 to 1997, and has consulted for entities including Apple, Samsung, Verizon, AT&T, the WiMAX Forum, Intel, Microsoft, Clearwire, Sprint, and NASA. He is an elected member of the Board of Governors of the IEEE Information Theory Society, has been a co-author of 11 best paper award recipients, including the 2014 IEEE Stephen O. Rice Prize and the 2014 IEEE Leonard G. Abraham Award, and recently received the 2015 ASEE/Hewlett Packard Terman Award.

Hisham Elshaer [S'13] received his B.Sc. from Ain-Shams University in Cairo, Egypt, in 2008. He received his M.Sc. in wireless systems from the Royal Institute of Technology (KTH), Sweden, in 2012. He was with Ericsson R&D in Stockholm as a Master thesis student in 2011. He worked as a consultant at Tieto Sweden AB from 2012 to 2013. Since July 2013 he has been with Vodafone R&D in the United Kingdom as an early stage researcher under the EU Crossfire project where he is working on developing new 5G technologies. He is also studying toward a Ph.D. at King's College London. He won the best paper award at IEEE GLOBECOM 2014 for his work on downlink and uplink decoupling.

Mischa Dohler [S'01, M'03, SM'07, F'14] is a full professor in wireless communications at King's College London, head of the Centre for Telecommunications Research, co-founder and member of the Board of Directors of the smart city pioneer Worldsensing, Distinguished Lecturer of the IEEE, and Editor-in-Chief of *Wiley Transactions on Emerging Telecommunications Technologies* and *EAI Transactions on the Internet of Things*. He is a frequent keynote, panel, and tutorial speaker. He has pioneered several research fields, contributed to numerous wireless broadband, IoT/M2M, and cyber security standards, holds 12 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. He has talked at TEDx, and had coverage on TV and radio.

STEFAN PARKVALL [SM] is currently a principal researcher at Ericsson Research working in research on 5G and future radio access. He is one of the key persons in the development of HSPA, LTE, and LTE-Advanced radio access, and has been deeply involved in 3GPP standardization for many years. He served as an IEEE Distinguished Lecturer (2011–2012) and is a co-author of the popular books *3G Evolution HSPA and LTE for Mobile Broadband*, *HSPA Evolution The Fundamentals for Mobile Broadband*, and *4G LTE/LTE-Advanced for Mobile Broadband*. He has numerous patents in the area of mobile communication. In 2005, he received the Ericsson "Inventor of the Year" award, in 2009 the Swedish government's Major Technical Award for his contributions to the success of HSPA, and in 2014 he and colleagues at Ericsson were nominated for the European Inventor Award, the most prestigious inventor award in Europe, for their contributions to LTE. He received his Ph.D. degree in electrical engineering from KTH in 1996. His previous positions include assistant professor in communication theory at KTH and a visiting researcher at the University of California, San Diego.

PETAR POPOVSKI [S'97, A'98, M'04, SM'10, F'16] received his Dipl.-Ing. (1997) and Magister Ing. (2000) in communication engineering from Sts. Cyril and Methodius University, Skopje, Macedonia, and his Ph.D. from Aalborg University, Denmark (2004). He is currently a professor at Aalborg University. He is an Editor for *IEEE Transactions on Communications* and has served in the past as an Editor for *IEEE Communications Letters*, the *IEEE JSAC* Cognitive Radio Series, and *IEEE Transactions on Wireless Communications*. He is a Steering Committee member for the *IEEE Internet of Things Journal*. His research interests are in communication theory, wireless communications, and networking.

SARABJOT SINGH [S'09, M'15] is a research scientist with Intel Corp., Santa Clara, California. He received his B.Tech. in electronics and communication engineering from the Indian Institute of Technology Guwahati, in 2010, and his M.S.E and Ph.D. in electrical engineering from the University of Texas at Austin in 2013 and 2014, respectively. He has held industrial positions at Nokia in Berkeley, California; Alcatel-Lucent Bell Labs in Crawford Hill, New Jersey; and Qualcomm Inc. in San Diego, California. He is interested in the modeling, design, and analysis of heterogeneous wireless networks with focus on novel MAC and PHY technologies. He was the recipient of the President of India Gold Medal in 2010 and the ICC Best Paper Award in 2013.

The ITU-T's New G.fast Standard Brings DSL into the Gigabit Era

Vladimir Oksman, Rainer Strobel, Xiang Wang, Dong Wei, Rami Verbin, Richard Goodson, and Massimo Sorbara

The standardized G.fast transmission method and advanced crosstalk cancellation techniques are presented by the authors, with specific performance projections and measurement results achieved during the first demonstrations and trials, showing bit rates of 500 Mb/s over 250 m and available reach up to 400 m.

ABSTRACT

This article explores the recently issued ITU-T Recommendations specifying "G.fast" (G.9701 [1] and G.9700 [2]) that bring user bit rates up to 1 Gb/s over twisted pairs from the distribution point to customer premises. The overview and some key research challenges of G.fast are discussed in [3]. The standardized G.fast transmission method and advanced crosstalk cancellation techniques are presented here with specific performance projections and measurement results achieved during the first demonstrations and trials, showing bit rates of 500 Mb/s over 250 m and available reach up to 400 m. A description of standardized tools for dynamic performance maintenance, resource allocation, and power saving enhancing G.fast applications concludes this article.

INTRODUCTION

Modern life depends on the Internet, and thus the demand for high-speed Internet access is rapidly growing. Digital subscriber line (DSL) technology, accordingly, keeps up with both customer demand and the progress in competing access technologies, such as DOCSIS, WiMAX/ Long Term Evolution (LTE), and gigabit passive optical networking (G-PON). In 2010, the International Telecommunication Union Telecommunication Standardization Sector (ITU-T) developed Recommendation G.993.5 [4], which set a 100 Mb/s benchmark in DSL services [5]. The new G.fast Recommendations [1, 2] specify 1 Gb/s access over copper.

To reach such high bit rates, G.fast uses only the last leg of the existing copper access network and in-premises wiring. These wires are usually unshielded, non-conditioned twisted pairs, flat pairs, or quads (four twisted wires) and known for very strong crosstalk, especially inside quads [3]. Reaching high bit rates over such low-quality copper is a difficult task that requires a substantially new approach. The main challenges encountered by engineers and the potential technical choices are discussed in [3], while this article describes the adopted technical solutions.

The G.fast-based access network uses the fiber-to-the-distribution-point (FTTdp) architecture [6], which comprises a distribution point unit (DPU) connected to the central office (CO) by fiber (PON or point-to-point fiber). DPUs are installed close to the customer premises, typically in mini-cabinets mounted in basements of multi-dwelling units, on electrical poles, in curb boxes, or in manholes [3, 6], and connected to customer premises equipment (CPE) via copper pairs. A DPU typically serves 4–20 lines, but bigger DPUs are expected in the future. A single-line DPU may serve as a fiber-to-the-home (FTTH) copper extension. DPUs can be powered locally, remotely, or by subscribers from the customer premises using reverse power feeding (RPF) [7]; the latter is very convenient for small DPUs. The achievable bit rate over a particular line depends on its length and wire type. The maximum reach is 400 m, but the majority of installations are expected to be within 100 m.

An example of a typical G.fast installation using RPF is shown in Fig. 1. The G.fast transceivers in the DPU (FTU-O) and in the CPE (FTU-R) are connected via a copper pair; the FTU-R resides in the network termination unit (NTU). The broadband services delivered to the DPU via a PON feeder are conveyed to the NTU and further distributed to various broadband applications via a high-speed in-premises network (e.g., WiFi). The RPF power sourcing equipment (PSE) in the NTU generates sufficient power to supply the associated FTU-O and common functions of the DPU via the copper pair. The DPU power supply unit (PSU) gathers the power sourced by PSEs of all active lines through corresponding power extractors (PEs). In other installations, the PSE may be separate from the NTU and also feed the NTU. The PSE can work during power outages using a backup battery. More RPF details can be found in [3, 7].

Analog phones are connected through adapters because RPF uses DC: the Foreign Exchange Office (FXO) adapter receives plain old telephone service (POTS) signaling derived from voice over IP (VoIP) service by the analog telephone adapter (ATA) and generates alternative signaling capable of running over in-premises wiring with RPF; the Foreign Exchange Subscriber (FXS) adapter further recovers the original POTS signaling. Both the FXO and FXS

Vladimir Oksman and Rainer Strobel are with Lantiq, an Intel Company; Xiang Wang and Dong Wei are with Huawei Technologies; Rami Verbin is with Sckipio Technologies; Richard Goodson is with AADTRAN Inc.; Massimo Sorbara is with Qualcomm.

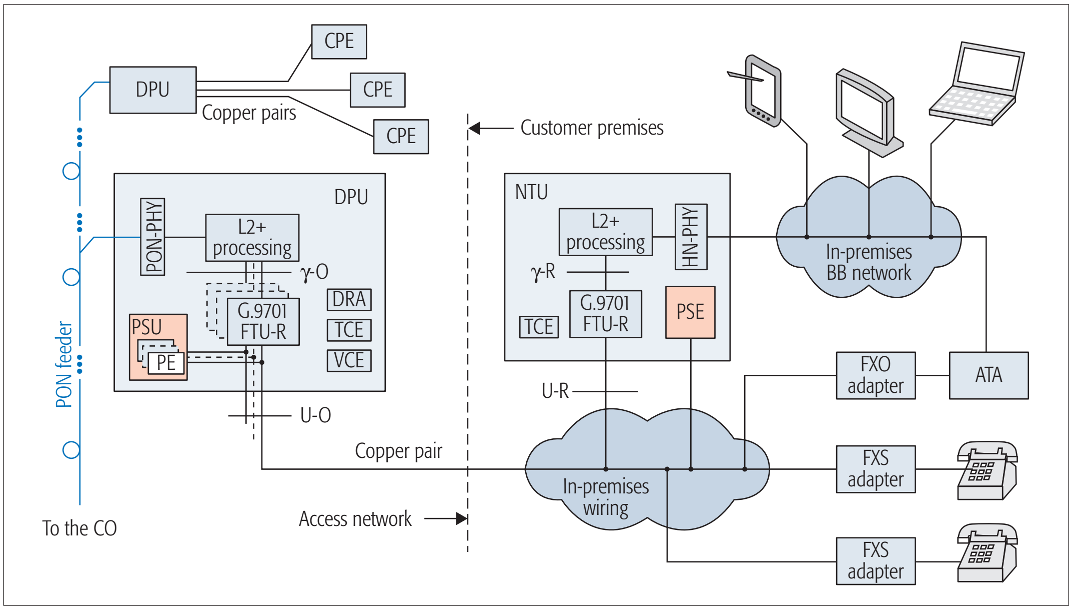

Figure 1. Example of G.fast deployment using RPF and derived POTS.

G.fast is a state-of-the-art copper access technology that provides fiber-grade transmission speed over existing copper, minimizes energy consumption, reduces maintenance cost, and provides great robustness and flexibility for the customers.

are fed by the PSE. To avoid G.fast performance loss, no other devices should be connected to the in-premises wiring.

In other installations, the ATA may reside in the NTU, while derived voice service can be distributed throughout the premises via cordless phone technology or by using smartphone connection to in-premises WiFi; no particular option for voice distribution is implied by G.fast.

G.9701 uses the frequency spectrum from 2.2 to 106 MHz with full crosstalk cancellation between the lines sourced by a DPU: near-end crosstalk (NEXT) is avoided by using synchronized time-division duplexing (STDD), and far-end crosstalk (FEXT) is cancelled using vectoring. No alien crosstalk cancellation is defined. Other G.fast innovations include dynamic allocation of resources between DPU sourced lines, efficient energy-saving techniques, and dynamic performance maintenance. Pair bonding is defined to allow multiplication of customers' bit rate.

G.fast customer installations vary by signal attenuation and noise, and may be influenced by other technologies. For reliable self-installation provisioning, which reduces the operator's cost, G.fast defines a flexible and robust transmission protocol, online reconfiguration, and dynamic adaptation of bit rate, all maintained via robust management channels. Zero-touch management further reduces the operator's cost by avoiding truck rolls for future equipment upgrades and adding new subscribers.

G.fast is a state-of-the-art copper access technology that offers fiber-grade transmission speed over existing copper, minimizes energy consumption, reduces maintenance cost, and provides great robustness and flexibility for customers.

Fundamentals of G.9701 Technology

Transmission Method

G.9701 specifies the functionality of G.fast transceivers (FTU-O and FTU-R) that establish a high-speed transmission path between γ-O and γ-R reference points (Fig. 1). The user's data packets from upper layers (L2+) are mapped into data transmission units (DTUs) that are conveyed transparently over the line. Reed-Solomon forward error correction [8] improves noise immunity: each DTU is assembled from multiple Reed-Solomon codewords, and DTU bytes are interleaved. The number of codewords and their size are configured to fit the throughput of the line. Noise is further mitigated by retransmission of DTUs received in error; the number of retransmissions for a DTU is limited by the latency bound. For retransmission and latency control, each DTU contains a sequence number and associated timestamp.

Discrete multi-tone (DMT) modulation [8] is used for passing DTUs and management data over the line. The advantages of DMT are well known, especially its capability to operate on lines with multiple bridged taps such as in-premises wiring. The specified tone spacing of 51.75 kHz is 12 times that of very-high rate DSL 2 (VDSL2) [8] because G.fast loops are much shorter. Thus, 2048-tone DMT is sufficient to cover the current G.fast frequency spectrum, simplifying the design. Each DMT symbol is cyclically extended using both prefix and suffix. The prefix mitigates inter-symbol interference and is configurable to address a wide range of loop lengths. The suffix is applied for transmit spectrum shaping and overlaps with the following symbol to improve efficiency; suffix size always fits the size of the windowing [8]. The default cyclic extension yields a symbol duration of 20.83 s. Up to 14 bits can be loaded per tone. To further increase bit rates, future versions of G.fast may extend the frequency spectrum to 211.968 MHz.

The advantages of the G.fast duplexing scheme (STDD) compared to FDD are described in [3]. With STDD, the upstream and downstream sets of DMT tones can be selected independently, making STDD flexible in the frequency domain. A particular selection depends on the deployment scenario, and involves channel characteristics and spectrum compatibility issues (see below).

To minimize the performance loss, downstream transmit PSDs are optimized across all the lines with precoder updates, including transmit power reduction of tones causing high crosstalk.

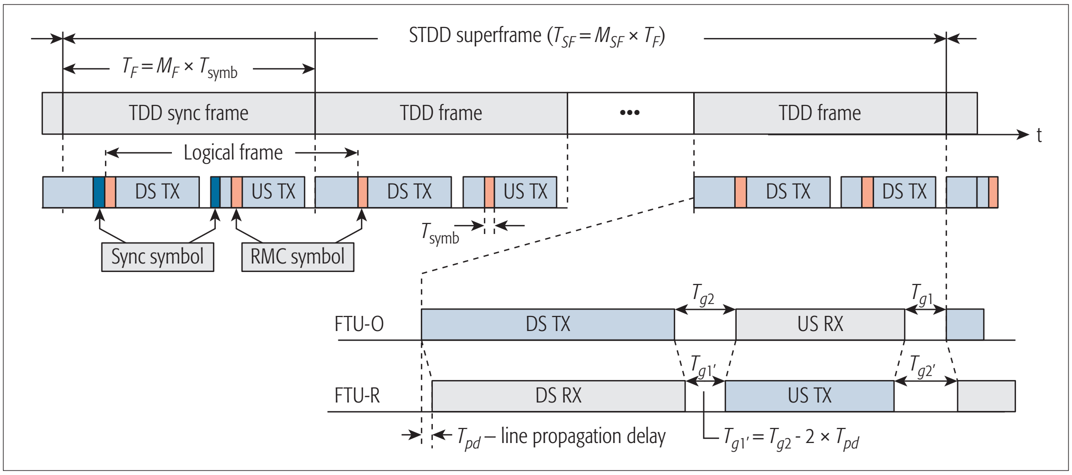

Figure 2. G.fast transmission format.

The G.fast transmission format comprises superframes, each composed of M_{SF} TDD frames (Fig. 2). Each TDD frame contains M_F symbol periods (T_{symb}). One set of contiguous symbol periods is assigned for downstream transmission and another one for upstream transmission. The sum of guard times between upstream and downstream transmissions ($T_{g_1} + T_{g_2}$) is one symbol period.

Superframes follow each other with no gaps; their boundaries are identified by downstream sync symbols. Both downstream and upstream sync symbols reside in a TDD sync frame, and carry probe sequences used for channel estimation and other purposes (described below).

The maximum duration of a TDD frame is bounded by the propagation delay limit to $M_F =$ 36 symbol periods. A setting of $M_F = 23$ reduces round-trip delay. A superframe contains 8 and 12 TDD frames, respectively, so its duration is always about 6 ms, which allows the superframe period to be used as a time base for initialization and management procedures.

The transmission path is maintained by the embedded operations channel (eoc) and robust management channel (RMC). The eoc is multiplexed into DTUs; it has a flexible bit rate that can support high-volume management data, but its robustness is about the same as user data. The RMC, in contrast, is defined to carry short messages and is much more robust due to high-redundancy Reed-Solomon coding, tone selection, and conservative bit loading. Multiple repetitions are applied for critical RMC commands.

One RMC symbol per direction is sent in each TDD frame (Fig. 2). It carries the RMC on dedicated tones and DTU bytes on other tones. The positions of RMC symbols in a TDD frame and the sets of RMC tones are configured at initialization. Symbol positions from one to the next RMC symbol of the same direction represent a logical frame. The RMC carries acknowledgments of received DTUs, supporting DTU retransmission, and conveys management commands facilitating logical frame configuration, online reconfiguration (OLR), and transitions into and out of low-power states.

With STDD, the time positions of superframes, TDD frames, sync symbols, and RMC symbols are aligned across all lines sourced by a DPU (vectored group). The alignment is by symbol boundaries, and only a small deviation is tolerable to avoid NEXT and facilitate FEXT cancellation, discontinuous operation, and fast reconfiguration.

FEXT Cancellation

FEXT cancellation is imperative for reaching high bit rates. Similar to G.993.5, G.fast performs FEXT cancellation at the DPU: the downstream transmit signals are precoded by adding FEXT pre-compensation signals, and a post-processor subtracts FEXT components from the received upstream signal [3, 5]. The vectoring control entity (VCE) at the DPU performs channel estimation, and computes precoder and post-processor matrices for all connected lines. The particular methods of channel estimation, matrix computation, and FEXT cancellation are vendor discretionary. For downstream channel estimation, the VCE may assign the same or different precoder matrices for sync symbols and data symbols (including the use of non-precoded sync symbols).

Like G.993.5, G.fast uses linear precoding. However, the FEXT behavior in G.fast is fundamentally different, especially for quad-twisted cables, due to G.fast's much wider frequency spectrum. Figure 3 shows a typical G.fast FEXT channel: the in-quad crosstalk can be stronger than the direct channel at high frequencies. This complicates precoding because the added pre-compensation signals can substantially increase the transmit PSD. Thus, with a given PSD limit, pre-compensation signals associated with a line generating high crosstalk can suppress the power of the direct signal in other lines of the vectored group, causing substantial performance loss [3].

To minimize this performance loss, downstream transmit PSDs are optimized across all the lines with precoder updates, including transmit power reduction of tones causing high crosstalk. A transmitter-initiated gain adjustment (TIGA) is used to accommodate the change of precoder gain in the peer FTU-R receiver. Figure 3 shows that PSD optimization substantially improves the achievable signal-to-noise ratio (SNR). It may even exceed the single-line SNR due to

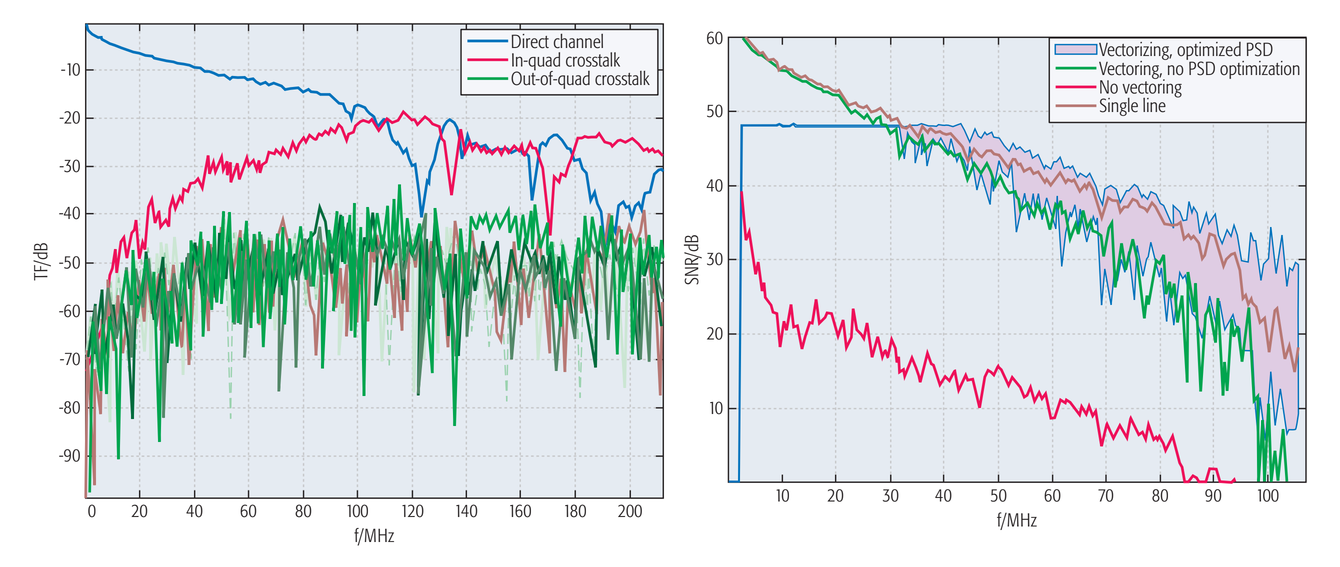

Figure 3. Left: direct and crosstalk channels of PE 0.5 mm quad cable (Germany); right: achievable downstream SNR.

additional direct signal propagation via FEXT channels. The latter also has a negative effect: turning off a line in a vectored group requires a precoder update and likely changes performance of the other lines.

Channel Estimation

Channel estimation is necessary to compute the FEXT channel matrix, and to derive precoder and post-processor matrices. Similar to G.993.5, probe sequences are carried by upstream and downstream sync symbols (Fig. 2). The values of the received probe signals are delivered to the VCE, which computes the channel matrix. The VCE assigns appropriate probe sequences (orthogonal, like Walsh-Hadamard sequences, or pseudo-orthogonal) to the lines of the vectored group using its own channel estimation strategy. Probe sequences are repeated periodically, allowing efficient averaging for noise mitigation. Unlike G.993.5, probe sequences may include 0-elements (no transmission of a sync symbol). By using 0-elements, the vectored group can be virtually divided into sub-groups; this reduces the aggregate FEXT inside each sub-group and speeds up channel estimation.

The assigned probe sequences are communicated to the FTU-R at initialization and can be updated via eoc during showtime (the state when user data is communicated). The VCE may request of the FTU-R to report DFT-output samples that represent the received probe signal in the frequency domain, or error samples that represent a normalized error vector between this received probe signal and the associated reference constellation point. The FTU-R uses the communicated downstream probe sequence to identify this constellation point, since the error may be comparable or even exceed the received signal due to strong FEXT.

The FTU-R report (vectoring feedback) is sent to the VCE via the eoc during showtime and via the special operations channel (SOC) during initialization. If available upstream capacity is insufficient, tone interpolation and decimation of the feedback in time and frequency are used, spreading the transmission over several probe sequence cycles.

Operation of a Vectored Group

Vectored group operation comprises three phases: tracking, joining, and leaving. During tracking, all lines of the group are in showtime: lines neither join nor leave the group, and each line tracks channel variations caused mainly by temperature changes and discontinuous operation. The latter may require frequent precoder updates facilitated by OLR procedures (TIGA, SRA, see below).

In the joining phase, new lines are added to the vectored group. During initialization of new lines, channel matrices, precoders, and transmit PSDs of new and showtime lines are jointly optimized, improving overall performance. This also requires multiple OLR procedures and high-volume transfers of vectoring feedback in showtime lines; both are supported by the eoc.

Lines can leave the group in an orderly or disorderly manner. Orderly leaving first terminates transmission in both directions, then updates the channel matrices of the remaining lines, allowing the FTU-R to then safely disconnect. A disorderly-disconnected FTU-R (e.g., unplugging) usually disturbs other lines substantially due to associated changes in their direct channel and increase of residual FEXT, which is obviously undesirable. Fast rate adaptation (FRA) and retransmission help mitigate error bursts until channel matrices are updated; the performance is restored by associated OLR procedures.

Details of the joining procedure are shown in Fig. 4. After a G.994.1 handshake [9], during which the two sides exchange capabilities, agree on a common operational mode, and set necessary parameters to facilitate STDD, the FTU-O starts training by transmitting superframes containing only sync symbols modulated by a probe sequence (during O-VECTOR-1). This allows the VCE to learn and cancel downstream

During the channel analysis and exchange phase, FTUs establish their desired showtime settings, such as bit loading, DTU size, and RMC tone sets. After CA&E, lines transition into showtime. The expected joining time of a single line is significantly less than in VDSL2 due to the shorter probe sequence cycle.

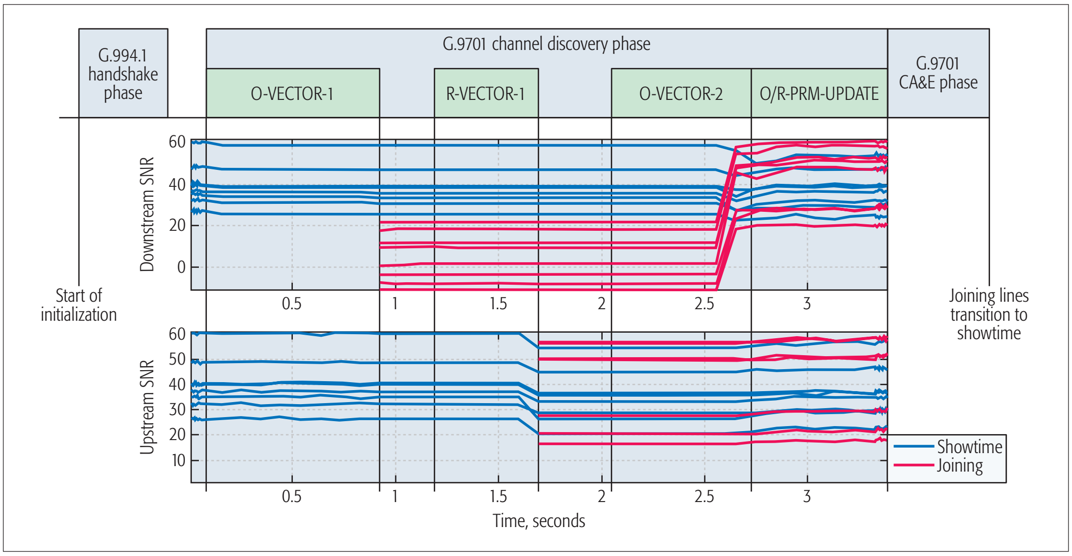

Figure 4. Joining timeline and examples of precoder and post-processor conversion in a 16-line DPU for tone #1160 (60 MHz).

crosstalk from joining lines into showtime lines without disturbing showtime lines. After the precoders of showtime lines are updated and the downstream crosstalk from joining lines is cancelled, the FTU-O turns on the downstream SOC and sends to the FTU-R the necessary upstream initialization data. Since crosstalk between joining lines is not cancelled, data transmitted over the SOC is scrambled using a unique scrambling seed in each joining line to avoid reception from a non-peer FTU-O. Furthermore, to improve robustness, the data is transmitted using repetitions and is modulated by an ID-sequence, which is orthogonal relative to the ID-sequences of other joining lines.

During R-VECTOR 1, the FTU-R transmits sync symbols modulated by a probe sequence, and the VCE estimates the upstream channel. After upstream crosstalk between all joining and showtime lines is mutually cancelled, a high-speed upstream SOC is established to convey vectoring feedback for precoder training (during O-VECTOR-2). After O-VECTOR-2, the downstream crosstalk between all joining and showtime lines is also cancelled.

During PRM-UPDATE, both FTU-O and FTU-R optimize their transmit PSDs in conditions when crosstalk is cancelled. One goal of optimization is to reduce the transmit power on tones with extra SNR margin. Another is to suppress tones generating very high crosstalk (to avoid performance loss in showtime lines). Other criteria, such as total power reduction, may also be applied [10].

During the channel analysis and exchange (CA&E) phase, FTUs establish their desired showtime settings, such as bit loading, DTU size, and RMC tone sets. After CA&E, lines transition into showtime. The expected joining time of a single line is significantly less than in VDSL2 due to the shorter probe sequence cycle.

Power Saving

Low power consumption is vital for G.fast, driven by limited heat dissipation, remote/reverse powering of the DPU, and battery-fed operation of the CPE. Power-saving mechanisms are discontinuous operation (DO) and low-power states.

Discontinuous Operation: DO scales a transceiver's power consumption with actual data throughput by transmitting only when data is available. During the remaining (quiet) symbol slots, essential analog and digital processing may be turned off, bringing substantial power savings.

In a vectored group with high crosstalk, turning off one line may change the direct channel of other lines, causing performance degradation. Therefore, strict coordination of turning slots quiet is applied across all vectored lines. Specifically, each logical frame is divided into a normal operation interval (NOI, the first TTR slots) and a discontinuous operation interval (DOI, the remaining slots). During NOI, no quiet slots are allowed, while in DOI the first TA slots are quiet, and the remaining slots may be quiet if no user data is available or active otherwise. In some NOI and DOI slots with no data available, an FTU may transmit only pre-compensation signals (idle slots). The DO parameters TTR, TA, and total number of active slots (TBUDGET, Fig. 5) are determined by the DRA and VCE based on DPU dynamic resource allocation (see below). They are configured per logical frame and coordinated across the vectored group via the RMC. Bit loadings and transmit PSDs during NOI and DOI are updated independently using OLR procedures. Figure 5 shows an example of downstream DO in a four-line DPU. The NOI includes five slots (TTR = 5), and values of TA are set so that during DOI only one line is transmitting at a time, facilitating crosstalk avoidance; no transmission in other lines and reduced vector processing saves power.

Power Saving States: During low-power states, FTUs save power by transmitting only sync symbols and RMC symbols in a few assigned TDD frames; all other slots are quiet. Furthermore, the transmit PSD and number of active tones in

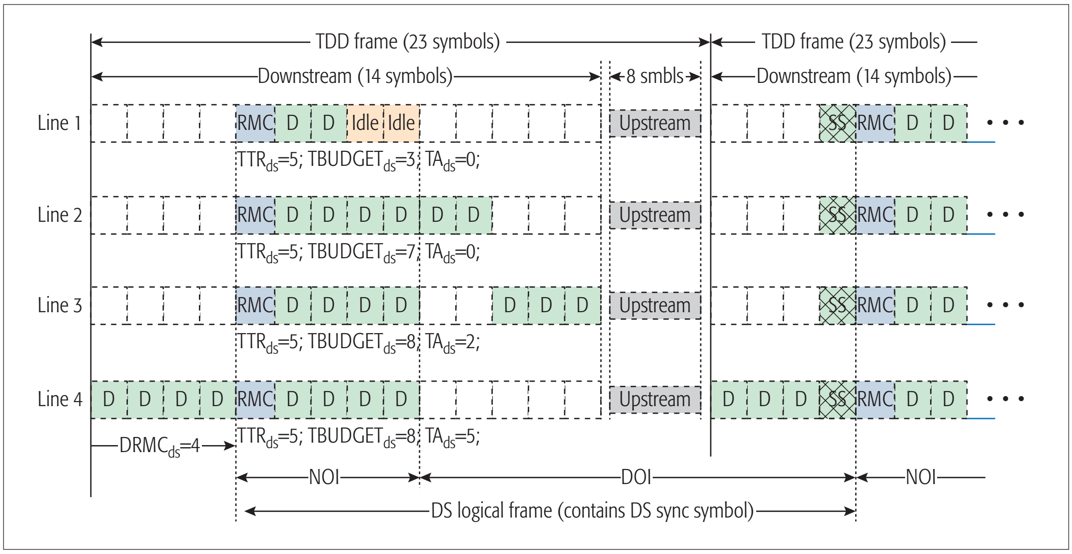

Figure 5. Example of downstream DO in a 4-line group (23-symbol TDD frame).

Keep-alive is very slow, but provides fast recovery of services. An NTU in L2.2 also detects incoming phone calls and temporarily transitions into L2.1 to support them. When the power is restored, the line moves into L2.1 (to support phone calls) and is ready to restart broadband service.

RMC symbols may be reduced. Two low-power states called L2.1 and L2.2 are defined. L2.1 efficiently saves power when broadband services are off, while continuing VoIP services. When broadband service restarts, the line transitions back to normal in less than 1 s.

In case of a power outage at a customer premises equipped with battery backup, L2.2 maintains only keep-alive traffic (no broadband or VoIP services). Keep-alive is very slow, but provides fast recovery of services. An NTU in L2.2 also detects incoming phone calls and temporarily transitions into L2.1 to support them. When the power is restored, the line moves into L2.1 (to support phone calls) and is ready to restart broadband service.

Transitions between states are triggered from the upper layers of the DPU, and by the DRA that monitors broadband traffic and CPE battery status. They are facilitated by eoc/RMC commands, similar to OLR transitions.

Dynamic Performance Maintenance

G.fast maintains services under varying channel conditions, overcoming unpredictable changes in the channel response and noise without dropping the link. Re-acquisition of the channel per transmission frame, as in IEEE 802.11n, is infeasible for G.fast due to its relatively long channel estimation time. Instead, G.fast uses a combination of two OLR types: seamless rate adaptation (SRA) and fast rate adaptation (FRA). SRA is accurate but slow and facilitates steady-state performance optimization, while FRA is coarse and fast, and keeps the link stable under sudden deep drops in SNR. In addition, bit swapping is used to permanently adjust the SNR margin.

For SRA, similar to VDSL2, the FTU receiver computes the optimum bit loading per tone and communicates it to the peer transmitter via the eoc. The bit loading update at the peer FTUs is synchronized to the start of a particular superframe by an RMC command. For robustness, this command is repeated multiple times and contains a count-down to the targeted superframe, which allows the receiving FTU to identify it even in very harsh noise conditions.

For FRA, the G.fast spectrum is divided into up to eight contiguous sub-bands. An FRA command determines a coarse bit loading trim per sub-band. A trim-down avoids line drop upon unexpected substantial SNR loss. The trim request is generated by the receiver and communicated to the peer transmitter via the RMC; the transmitter quickly activates a bit loading update by sending an RMC synchronization command.

Since the RMC is much more robust than the data channel, it remains functional to recover the line from a temporary loss of data connectivity. As a receiver senses critical degradation of the channel, it requests via the RMC to trim down the bit loading in the affected sub-bands. After both the transmitter and receiver synchronously lower their bit loading according to the new channel conditions, data and eoc connectivity is recovered, and vectoring feedback is restored. Once vectoring coefficients are updated, the receiver optimizes bit loading by a trim-up FRA and/or SRA. This way, a stable link is maintained under harsh temporary conditions without sacrificing the steady-state performance by using extra SNR margin. The FRA timing is configured to avoid reacting to impulse noise, which is handled by retransmission.

The TIGA procedure is introduced into G.fast to facilitate updates of the precoder upon joining, leaving, tracking, and DO events, which in a high-crosstalk environment usually requires a change in the downstream transmit PSD and bit loading (see above). A TIGA eoc command is sent by the FTU-O to convey the necessary changes (bit loading and complex gain) to the receiver. The following RMC command provides a synchronous update of the precoder and parameters of the FTU-O transmitter and FTU-R receiver across the vectored group.

Spectral Compatibility and Coexistence with Current Deployments

G.fast spectral compatibility is determined by ITU-T G.9700 [2], which defines the PSD mask, transmit power limit, and a variety of spectrum

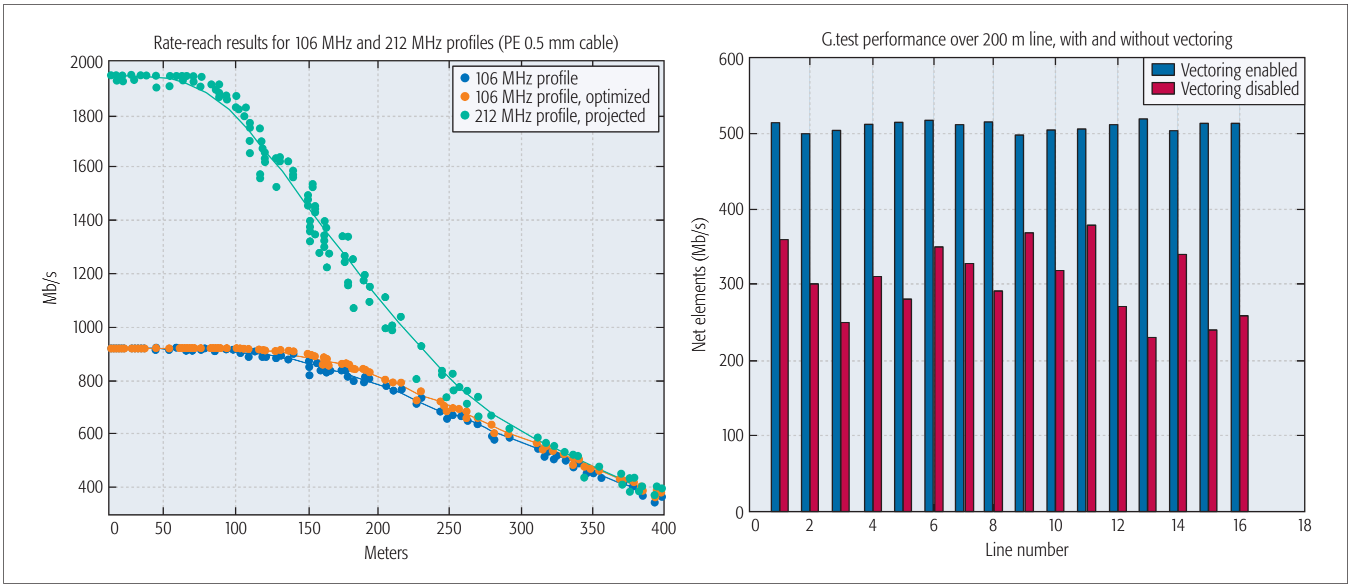

Figure 6. Simulated rate-reach curves (left) and actual measurement results (right) for 0.5 mm cables, 4 dBm TX power.

management tools established to reduce RFI egress into other DSL and radio services. The G.fast transmit power limit is 4 dBm. The PSD mask is –65 dBm/Hz below 30 MHz, drops to –73 dBm/Hz at 30 MHz, and further slopes down to –79 dBm/Hz at 212 MHz. G.fast suffers from RFI ingress, especially from FM radio, and crosstalk generated by VDSL2 and in-premises networks using power line technologies (PLT) [11].

Compatibility with DSL Deployments

Taking into account the expected migration from DSL to G.fast services and unbundling, it is highly desirable that G.fast be spectrally compatible with asymmetrical DSL (ADSL)/ADSL2plus and with VDSL2 deployed from an exchange or a cabinet.

The most practical and reliable way to maintain spectral compatibility between DSL and G.fast is by spectral separation. G.fast is by design compatible with ADSL/ADSL2plus since the lowest frequency of G.fast is 2.2 MHz. For compatibility with VDSL2, the start frequency of G.fast is set above the VDSL2 spectrum; the latter depends on the used VDSL2 profile. Use of spectral overlap between VDSL2 and G.fast is studied in [12].

Compatibility with Broadcast, Amateur Radio, and PLT

For compatibility with broadcast radio, G.fast transmits substantially reduced PSD on frequencies above 30 MHz. If no transmission is allowed inside international amateur-radio bands or the FM-radio band (87.5 MHz — 108 MHz), these frequencies are notched out from the G.fast spectrum.

In-premises PLT networks may impact G.fast performance due to crosstalk between in-premises electrical wiring and phone wiring. This crosstalk is difficult to predict, although recently studied statistical models [11, 13] show that in 90 percent of the cases crosstalk attenuation is 60 dB or even less on frequencies above 30 MHz. ITU-T G.9977 defines a mechanism to reduce this crosstalk by adjusting transmission parameters of G.fast and PLT network nodes via an arbitration device controlled by the operator.

Performance of G.fast

Theoretical Evaluation and Measurement Results

Capacity evaluations for G.fast 106 MHz and 212 MHz profiles over a sample of PE-0.5 mm cable are shown in Fig. 6 (left). The simulation shows bit rates averaged over five 24-pair cable binders, each with line lengths uniformly distributed between 10 m and 400 m, using a start frequency of 2.2 MHz (dots show performance on different pairs). The "optimized" option reflects PSD optimization described above, which is also used for 212 MHz performance projection with linear precoding.

The wideband transmit power and PSD limits meet G.9700. Flat spectrum background noise with a PSD of –140 dBm/Hz is applied to model the effect of QLN and receiver noise factor. For the 106 MHz profile, the aggregate (upstream plus downstream) bit rate of 500 Mb/s is achieved for lines up to 340 m. A similar distance is expected for 1 Gb/s service with 2-line bonding. For the 212 MHz profile, a 1 Gb/s aggregate bit rate can be reached for loops up to 220 m.

The measurement results in Fig. 6 (right) for a 16-pair group of another 0.5 mm 200 m cable with G.fast using the frequency range 23–106 MHz also shows the clear advantage of vectoring; a 500 Mb/s bit rate is supported on all tested lines.

Management

Management Interfaces

Operators manage G.fast through the management information base (MIB) established at the DPU and in some cases at the NTU. At installation, the operator sets system configuration parameters, and during showtime the operator reads out performance and test parameters, reported events, and collected statistics using management objects defined in [14]. Operators can access the DPU MIB remotely via the network management system (NMS). If the DPU is unpowered, a persistent management agent (PMA) acts as a proxy for the NMS. Relevant FTU-R management data are retrieved via the eoc. The NMS can access the NTU MIB, if established, using TR-069 [15].

Diagnostic and Performance Prediction

Means for line diagnostics include collection of line attenuation (HLOG), quiet-line noise (QLN), and signal attenuation (SATN) for both upstream and downstream. Unlike VDSL2, these parameters can be monitored during showtime and, together with reports on crosstalk coupling and SNR margin, provide a detailed picture of the current line status.

A method to measure the downstream HLOG and QLN is based on reporting of the discrete Fourier transform (DFT) samples of the received signal for specific elements of a probe sequence. For instance, setting a certain element to 0 in all lines allows measurement of the QLN. Furthermore, setting a particular element to 1 in one line and to 0 in all other lines allows measurement of the HLOG. The vectoring feedback (from the FTU) can be configured to report DFT samples for selected probe sequence elements, while error samples are reported on other elements. Thus, by adding a few elements to a probe sequence, HLOG and QLN can be monitored without interrupting the service and channel estimation.

Dynamic Resource Allocation

Dynamic resource allocation (DRA) controls the number of transmission slots in each logical frame for all lines of the DPU as a function of traffic loading, environmental conditions, power state, and battery status. By using DRA, service level agreements (SLAs) may be met using minimized DPU power consumption.

The main inputs to the DRA function are:

- Downstream and upstream traffic load indicators for every line in the DPU. Currently, these are simply per-traffic-class buffer occupancies, much like in G-PON. Using these indicators, the DRA adjusts the number of allowed transmission slots in a logical frame as a function of traffic load, thereby scaling the power consumption with traffic. The DRA can also switch a particular line to a low-power state if no broadband traffic is required.
- Environmental conditions, such as DPU housing temperature. This keeps DPU power dissipation within acceptable limits.
- Battery status indicator. The DRA can switch a particular line to a low-power state if the line is battery-fed as a result of power outage.

Based on these inputs, the DRA determines DO parameters (2.5.1) for each line. The parameter TBUDGET is coordinated across all DPU lines to control power dissipation and to address vectoring processing constraints.

Conclusions

G.fast brings DSL technology to a new level, comparable to the FTTH grade of service. It allows operators to offer their customers multiple broadband services, of both constant bit rate and variable bit rate, with total aggregated bit rate up to 1 Gb/s and low propagation delay. Such services include multi-channel HD video, high-quality audio and voice, and modern multi-user interactive gaming. G.fast assumes operation over low-grade copper drop cables and in-premises wiring. Reverse power feeding resolves the issue of cabinet powering, and readiness for customer self-install substantially improves cost effectiveness, simplifies system management, and brings convenience to the customer.

G.fast brings DSL technology to a new level, comparable to FTTH grade of service. It allows operators to offer their customers multiple broadband services, of both constant bit rate and variable bit rate, with total aggregated bit rate up to 1 Gb/s and low propagation delay.

Acknowledgments

The authors wish to acknowledge the great role in project development of Frank Van der Putten (Alcatel-Lucent), Chair, Les Brown (Huawei), Editor, and colleagues Les Humphrey (British Telecom), who led us with his vision of the project, Ian Horsley (British Telecom), Angus Carrick (Swisscom), Hubert Mariotte (France Telecom), Marko Loeffelholz (Deutsche Telecom), Miguel Peeters (Broadcom), Danny Van Bruyssel (Alcatel-Lucent), Ken Kerpez (ASSIA), and Chris Chang (Metanoia). We also thank Tom Starr (AT&T), who initiated development of G.fast in ITU-T, for helpful comments and guidance in this work.

References

[1] ITU-T Rec. G.9701-2014, "Fast Access to Subscriber Terminals (G.fast) Physical Layer Specification."

[2] ITU-T Rec. G.9700-2014, "Fast Access to Subscriber Terminals (G.fast) Power Spectral Density Specification."

[3] M. Timmers, C. Nuzman, and J. Maes, "G.fast: Evolving the Copper Access Network," *IEEE Commun. Mag*., Aug. 2013, pp, 74–79.

[4] ITU-T Rec. G.993.5-2010, "Self-FEXT Cancellation (Vectoring) for Use with VDSL2 Transceivers."

[5] V. Oksman *et al.*, "The ITU-Ts New G.vector Standard Proliferates 100 Mb/s DSL," *IEEE Commun. Mag*., Oct. 2010, pp 140–48.

[6] Broadband Forum, "Architecture and Requirements for Fiber to the Distribution Point," tech. rep. TR-301.

[7] ETSI, "European Requirements for Reverse Powering of Remote Access Equipment," ETSI tech. spec. TS 101 548, V1.2.1, Nov. 2014.

[8] T. Starr *et al.*, *DSL Advances*, Prentice Hall, 2003.

[9] ITU-T Rec. G.994.1-2012, "Handshake Procedures for Digital Subscriber Line Transceivers."

[10] R. Strobel, A. Barthelme, and W. Utschick, "Zero-Forcing and MMSE Precoding for G.fast," *Proc. IEEE GLOBECOM*, 2015.

[11] S. Galli *et al.*, "The Impact of PLC-to-DSL Interference on VDSL2, Vectored VDSL2, and G.fast," *Proc. IEEE ISPLC 2015*.

[12] R. Strobel and W. Utschick: "Coexistence of G.fast and VDSL in FTTDP and FTTC Deployments," *Proc. EUSIPCO*, 2015.

[13] S. Galli *et al*., "Statistical Modeling of PLC-to-DSL Interference," *Proc. IEEE ISPLC 2015*.

[14] ITU-T Rec. G.997.2-2015, "Physical Layer Management for G.fast Transceivers."

[15] Broadband Forum, "CPE WAN Management Protocol," tech. rep. TR-69.

Biographies

Vladimir Oksman (vladimir.oksman@intel.com) received his M.S. and Ph.D. degrees from the Leningrad Radio and Telecommunications College (USSR) in 1976 and 1987, respectively; che urrently directs DSL standardization technical marketing at Intel Inc. He actively participated in development and standardization of emerging DSL technologies starting in 1991 as a researcher, project manager, and technical advisor. Since 2011 his main focus is concept engineering and standardization of G.fast.

Rainer Strobel (rainer.strobel@intel.com) works for Concept Engineering at Lantiq, now Intel, since 2011. He received his Dipl.-Ing.(FH) degree from the University of Applied Sciences Augsburg in 2009 and his M.Sc. degree in electrical engineering from Technische Univsersität München (TUM) in 2011.

He is currently working toward a Ph.D. degree at Intel in cooperation with the Associate Institute for Signal Processing, TUM. His research interests include optimization of wired MIMO communication systems, copper-fiber hybrid networks, and wireline channel modeling.

XIANG WANG (Wangxiang@huawei.com) got his Ph.D degree in computational mathematics from Nanjing University in 2009. He is interested in numerical optimization, linear algebra application, and communications. His work at Huawei has been focused on DSL technology research since 2010. He also represented Huawei at the ITU-T Q4/SG15 meeting on DSL and G.fast.

DONG WEI (weidong@huawei.com) received his Ph.D. degree in electrical engineering from the University of Texas at Austin in 1998. He worked as an assistant professor at Drexel University, Philadelphia, Pennsylvania, and as a principal MTS at SBC Labs, with research interests in signal processing for communications and advanced access technologies. Since 2006, he has been a senior expert at the U.S. R&D Center of Huawei Technologies, working on advanced DSL technologies, actively contributing to various international SDOs.

RAMI VERBIN (rami@sckipio.com) is a CTO and co-founder of Sckipio Technologies, developing G.fast solutions. Prior to Sckipio he was CTO of CopperGate and Sigma Designs. Since 1994 he has been developing wire-line access solutions and leading research in various communication fields. He holds an M.Sc. in electrical engineering and an M.B.A..

RICHARD GOODSON (msorbara@adtran.com) has over 30 years of advanced communication technology experience, from spread spectrum radios to xDSL and PON. He holds a B.S.E.E. from the University of Alabama and an M.S.E.E. from the University of Florida. He participates in the Broadband Forum, ITU-T, FSAN, and ATIS standards organizations. He has been at ADTRAN since 1995 where he is currently director of industry standards and technology analysis in the CTO Office.

MASSIMO SORBARA (msorbara@qca.qualcomm.com) is senior director of technical standards at Ikanos Communications and now at Qualcomm in Red Bank, New Jersey. He has worked on the standardization of DSL systems since 1990, holding various leadership positions in ITU-T, Broadband Forum, and ATIS. His research interests are in signal processing for advanced network access technologies. He received a B.S.E.E. from Manhattan College (1978) and an M.S.E.E. from the University of Santa Clara (1982).

Call for Papers
IEEE Communications Magazine
Next Generation 911

Background

In most parts of the world where a phone is used, users depend on the SOS or emergency service (ES). In North America, this service is accessible by dialing 911 and in most parts of Europe by dialing 112. Since its introduction in the 1950s, the service has adapted to technology changes: from wireline phones with fixed locations to supporting mobile phones. However, the technology is improving faster than the ES infrastructure can keep up. Faster networks coupled with heterogeneous access methods pose a challenge to the evolution of the ES. As if this is not enough, the richer choices available to a user today for communication — video, text messaging, social networking portals like Facebook and Google+, instant messaging, web-based calling, and over-the-top voice over IP applications — make the task of providing ES uniformly ever more difficult. The United States and Europe have decided to approach ES through a clean slate approach. The redesign of the ES in the United States is known as Next Generation 911 (NG911) and in Europe as Next Generation 112 (NG112). The core of NG ES is the Emergency Services IP Network (ESInet). ESInet uses Session Initiation Protocol (SIP) to deliver voice, video, text, and data calls reliably and uniformly to the ES network.

Authors from industry and academia are invited to submit papers for this Feature Topic of *IEEE Communications Magazine* on next generation 911. The Feature Topic scope includes, but is not limited to, the following topics of interest:

- Overview of NG911 implementation efforts
- Technical challenges in implementing NG911 systems
- Status of current implementation of NG911
- Status of NG911 deployment
- Issues in multi-modal NG911 devices
- Impact of social media on NG911 systems
- Societal impacts of the presence (or absence) of such services
- Over-the-top applications and NG911
- Security and privacy of NG911
- Public policy and funding issues with NG911
- Regulatory environment for NG911
- The role of standards (IETF, ITU-T) in NG911

Submissions

Articles should be tutorial in nature, with the intended audience being all members of the communications technology community. They should be written in a style comprehensible to readers outside the specialty of the article. Mathematical equations should not be used (in justified cases up to three simple equations are allowed). Articles should not exceed 4500 words. Figures and tables should be limited to a combined total of six. The number of archivable references is not to exceed 15. Complete guidelines for manuscript preparation can be found via the link http://www.comsoc.org/commag/paper-submission-guidelines. Please send a PDF (preferred) or MS-Word formatted paper via Manuscript Central (http://commag-ieee.manuscriptcentral.com). Register or log in, and go to the Author Center. Follow the instructions there. Select "November 2016 / Next-Generation 911".

Important Dates

- Manuscript Submission Deadline: March 15, 2016
- Decision Notification : June 30, 2016
- Final Manuscript Due Date: August 31, 2016
- FT Publication Date: November 2016

Guest Editors

Vijay K. Gurbani	Salvatore Loreto	Ravi Subrahmanyan	Carol Davids
Bell Laboratories, Alcatel-Lucent, USA	Ericsson, Sweden	Butterfly Network Inc., USA	Illinois Inst. of Tech., USA
vkg@bell-labs.com	salvatore.loreto@ericsson.com	ravi.subrahmanyan@ieee.org	davids@iit.edu

COMPANY .... PAGE

Keysight .... Cover 2

National Instruments .... Cover 4

Rohde & Schwarz GmbH 3

Tutorial .... Cover 3

Wiley 11

CURRENTLY SCHEDULED TOPICS

TOPIC	ISSUE DATE	MANUSCRIPT DUE DATE
NEXT GENERATION 911	NOVEMBER 2016	MARCH 15, 2016
INTEGRATED COMMUNICATIONS, CONTROL, AND COMPUTING TECHNOLOGIES FOR ENABLING AUTONOMOUS SMART GRID	DECEMBER 2016	APRIL 1, 2016
NEW WAVEFORMS AND MULTIPLE ACCESS METHODS FOR 5G NETWORKS	NOVEMBER 2016	APRIL 2, 2016
IMPACT OF NEXT-GENERATION MOBILE TECHNOLOGIES ON IOT-CLOUD CONVERGENCE	JANUARY 2017	APRIL 15, 2016
PRACTICAL PERSPECTIVES ON IOT IN 5G NETWORKS: FROM THEORY TO INDUSTRIAL CHALLENGES AND BUSINESS OPPORTUNITIES	FEBRUARY 2017	MAY 1, 2016

www.comsoc.org/commag/call-for-papers

TOPICS PLANNED FOR THE APRIL ISSUE

CRITICAL COMMUNICATIONS AND PUBLIC SAFETY NETWORKS

INTEGRATED CIRCUITS FOR COMMUNICATIONS

FROM THE OPEN CALL QUEUE

NETWORK FUNCTION VIRTUALIZATION IN 5G

ADAPTIVE AND COGNITIVE COMMUNICATION ARCHITECTURE FOR NEXT-GENERATION PPDR SYSTEMS

SCALABLE AND MOBILE CONTEXT DATA RETRIEVAL AND DISTRIBUTION FOR COMMUNITY RESPONSE HETEROGENEOUS WIRELESS NETWORKS

A SURVEY ON RAPIDLY DEPLOYABLE SOLUTIONS FOR POST-DISASTER NETWORKS

A UNIFYING PERSPECTIVE ON PROXIMITY-BASED CELLULAR-ASSISTED MOBILE SOCIAL NETWORKING

MILLIMETER-WAVE GBPS BROADBAND EVOLUTION TOWARDS 5G: FIXED ACCESS AND BACKHAUL

REDUCING COMPLEXITY OF VIRTUAL MACHINE NETWORKING

A SCALABLE ARCHITECTURE FOR HANDLING CONTROL PLANE FAILURES IN 5G HETEROGENEOUS NETWORKS

BUFFER SIZING IN WIRELESS NETWORKS: CHALLENGES, SOLUTIONS, AND OPPORTUNITIES

MULTI-COMM-CORE ARCHITECTURE FOR TERABIT/S WIRELESS

There's a Better Way to Design 5G Wireless Systems

It starts with an integrated design approach for rapid prototyping.

In the race to 5G, researchers need to accelerate design cycles. With the LabVIEW Communications System Design Suite and NI software defined radio hardware, they can build 5G prototypes fast to decrease the time between an idea and a real-world application.

>> See how at ni.com/5g

800 891 8841

NATIONAL INSTRUMENTS™

Communications Standards

A Supplement to IEEE Communications Magazine

MARCH 2016

www.comsoc.org

Deploying SDN and NFV at the Speed of Innovation

End-to-End Service Management Framework for Anything-as-a-Service

An Edge Operating System Enabling Anything-as-a-Service

Service Description in the NFV Revolution: Challenges and a Way Forward

A Two-Layered Data Model Approach for Network Services

Patents and Standardization, Part 1: A Tutorial on Patents

Standards for Multi-Stream and Multi-Device Media Synchronization

IEEE Standard 802.19.1 for TV White Space Coexistence

3GPP Device-to-Device Communications for Beyond 4G Cellular Networks

LTE-Advanced in 3GPP Rel-13/14: An Evolution Toward 5G

Semantics for Anything-as-a-Service

IEEE ComSoc IEEE Communications Society

A Publication of the IEEE Communications Society

COMMUNICATIONS STANDARDS

A Supplement to IEEE Communications Magazine

MARCH 2016

2 **EDITOR'S NOTE: STANDARDIZING INFORMATION MODELLING**
GLENN PARSONS, EDITOR-IN-CHIEF

3 **COMMENTARY: ABOUT ICAP: THE IEEE STANDARDS ASSOCIATION CONFORMITY ASSESSMENT PROGRAM**
BRUCE KRAEMER, PRESIDENT, IEEE STANDARDS ASSOCIATION

4 **STANDARDS NEWS**

7 **STANDARDS EDUCATION**

ACCEPTED FROM OPEN CALL

10 **PATENTS AND STANDARDIZATION, PART 1: A TUTORIAL ON PATENTS**
KRISTA S. JACOBSEN

16 **STANDARDS FOR MULTI-STREAM AND MULTI-DEVICE MEDIA SYNCHRONIZATION**
M. OSKAR VAN DEVENTER, HANS STOKKING, MATT HAMMOND, JEAN LE FEUVRE, AND PABLO CESAR

22 **IEEE STANDARD 802.19.1 FOR TV WHITE SPACE COEXISTENCE**
STANISLAV FILIN, TUNCER BAYKAS, HIROSHI HARADA, FUMIHIDE KOJIMA, AND HIROYUKI YANO

28 **3GPP DEVICE-TO-DEVICE COMMUNICATIONS FOR BEYOND 4G CELLULAR NETWORKS**
SHAO-YU LIEN, CHUN-CHE CHIEN, FAN-MIN TSENG, AND TIEN-CHEN HO

36 **LTE-ADVANCED IN 3GPP REL-13/14: AN EVOLUTION TOWARD 5G**
JUHO LEE, YOUNSUN KIM, YONGJUN KWAK, JIANZHONG ZHANG, ARIS PAPASAKELLARIOU, THOMAS NOVLAN, CHENGJUN SUN, AND YINGYANG LI

SEMANTICS FOR ANYTHING-AS-A-SERVICE

44 **GUEST EDITORIAL**
NIGEL DAVIS, KAM LAM, YUJI TOCHIO, AND SCOTT MANSFIELD

46 **DEPLOYING SDN AND NFV AT THE SPEED OF INNOVATION: TOWARD A NEW BOND BETWEEN STANDARDS DEVELOPMENT ORGANIZATIONS, INDUSTRY FORA, AND OPEN-SOURCE SOFTWARE PROJECTS**
BRAM NAUDTS, WOUTER TAVERNIER, SOFIE VERBRUGGE, DIDIER COLLE, AND MARIO PICKAVET

54 **E2EUBERIM: END-TO-END SERVICE MANAGEMENT FRAMEWORK FOR ANYTHING-AS-A-SERVICE**
KATARZYNA WAC, MARK CUMMINGS, AND JAYANTA DEY

62 **AN EDGE OPERATING SYSTEM ENABLING ANYTHING-AS-A-SERVICE**
ANTONIO MANZALINI AND NOEL CRESPI

68 **SERVICE DESCRIPTION IN THE NFV REVOLUTION:TRENDS, CHALLENGES AND A WAY FORWARD**
JOKIN GARAY, JON MATIAS, JUANJO UNZILLA, AND EDUARDO JACOB

76 **A TWO-LAYERED DATA MODEL APPROACH FOR NETWORK SERVICES**
CARL MOBERG AND STEFAN VALLIN

Standardizing Information Modelling

Glenn Parsons

Information is the link between business services and ICT (information and communications technologies). Information modelling captures business requirements so they can be communicated in such a way that ICT applications are developed to support services. Information should be stable and organized in a way that enables service deployment flexibility when it comes to process and organization. Information modelling enables a common language within a business or an organization, and between businesses. Often, several different concepts commonly in use mean the same thing. Also, similarly named concepts can have different meanings and definitions. Standardized information modelling is used to align these concepts and their definitions. As we evolve toward a truly networked society, the demand will grow for market driven standardization of modelling and the deployment of standardized business services.

The importance of standards to the work and careers of communications practitioners is the basis of this publication. It is a platform for presenting and discussing standards related topics in the areas of communications, networking, research, and related disciplines.

This issue of the Communications Standards Supplement contains a number of open call papers as well as a feature topic on information modelling. Readers will notice the ongoing Commentary section with a recurring view from the IEEE-SA President. This time we are happy to also include a new IEEE Standards Education section to create awareness and understanding of the importance of standardization and the critical role standards play in industry and society. The Standards News section offers the current status of standards work in various SDOs relevant to service management and information modelling, as well as pointers to SDO material. I trust that the reader will find these informative and illustrative of the fundamental role standards play in the communications networking ecosystem.

The first article in the Open Call section, authored by Krista Jacobsen, is actually the first in a new series on patents and standardization. Jacobsen provides an initial tutorial that explains the patenting process in the United States. Patent Cooperation Treaty (PCT) applications are introduced, which are commonly used to file a single patent application that can then make its way into the national offices of most countries. The article provides an overview of the policies underpinning patent law, the rights accompanying a patent, and the requirements and procedures to obtain a patent. In future editions of the supplement, Jacobsen will contribute articles on: the obligation to disclose patents to standards setting organizations (SSO) and patent offices; the meaning of statements promising to license any such patents on "fair, reasonable, and non-discriminatory" (FRAND) terms; and an overview of antitrust law and the tension between SSOs and antitrust law.

Media synchronization is receiving renewed attention, with ecosystems of connected devices enabling novel media consumption paradigms. The article by van Deventer *et al.* provides an overview of recently published standards, given the commercial interest in media synchronization that has spawned a new wave of international standards for media synchronization.

The article by Filin *et al.* gives an overview of IEEE Std. 802.19.1 for "TV White Space Coexistence Methods". With the understanding of the need to provide coexistence solutions for different cognitive radio systems operating in white space frequency bands, this standard specifies radio technology independent methods for coexistence among dissimilar or independently operated TV band networks.

3GPP Release 12, including its special form of "broadcasting" communications for device-to-device (D2D) communications, has been adopted as the next generation public safety network. The article by Lien *et al.* provides a comprehensive overview of D2D operations in Release 12 to show that this was an unprecedented technology in cellular networks.

Following on with the 3GPP standards activities, the article by Lee *et al.* introduces a set of key technologies expected for 3GPP Release 13 and 14, as part of the continued evolution of LTE-Advanced and as a bridge from 4G and 5G. As the fourth generation (4G) LTE-Advanced network becomes a commercial success, technologies for beyond 4G and 5G are actively studied, from a research perspective as well as a standardization perspectives.

Following the open call articles, Scott Mansfield and his editorial team provide an informative feature topic, "Semantics for Anything-as-a-Service", that shows the connection of end-to-end service management with information and data modelling. Mansfield *et al.* will introduce the feature topic papers in more detail.

Looking to the remainder of 2016, the Communications Standards Supplement will continue quarterly publication, and each issue will be "anchored" around a topic of current market relevance to drive focus. The next issue will contain a feature topic based on the recent ITU Kaleidoscope, "Trust in the Information Society", presenting insight into the means of building information infrastructures deserving our trust. Proposals for future feature topics are welcome as we look forward to the Communications Standards Supplement evolving into a stand-alone magazine in 2017.

Biography

Glenn Parsons [SM] (glenn.parsons@ericsson.com) is an internationally known expert in mobile backhaul and Ethernet technology. He is a standards advisor with Ericsson Canada, where he coordinates standards strategy and policy for Ericsson, including network architecture for LTE mobile backhaul. Previously, he has held positions in development, product management and standards architecture in the ICT industry. Over the past number of years, he has held several management and editor positions in various standards activities including IETF, IEEE, and ITU-T. He has been an active participant in the IEEE-SA Board of Governors, Standards Board and its Committees since 2004. He is currently involved with mobile backhaul standardization in MEF, IEEE and ITU-T and is chair of IEEE 802.1. He is a Technical Editor for IEEE Communications Magazine and has been co-editor of several IEEE Communications Society Magazine feature topics. He graduated in 1992 with a B.Eng. degree in electrical engineering from Memorial University of Newfoundland.

About ICAP: The IEEE Standards Association (IEEE-SA) Conformity Assessment Program

By Bruce Kraemer, President, IEEE Standards Association

Standards Tradition

IEEE has historically been the source of many highly respected and widely used standards. Some of these have become pervasive in our daily lives. However, the products or services that use those standards have not, historically, involved the IEEE in either product verification or product branding. Until recently, those tasks have been delegated to other organizations. History is changing.

New Offering

The IEEE Standards Association now operates an initiative known as the IEEE-SA Conformity Assessment Program (ICAP). This program extends the standards ecosystem by providing a means of producing an early conformance test and trusted products, and thereby serves to accelerate market adoption of conforming products.

ICAP initiatives draw together the same subject matter experts from industry who wrote the IEEE standard to develop test suites associated with that standard and establish compliance test programs that provide the producers and consumers with validated product certifications.

The Benefits of Conformity Assessment

Conformity assessment provides confidence and assurance to the market ecosystem that a product is compliant with industry standards and functions as designed. Markets significantly prefer products that have demonstrated compliance and interoperability. For these reasons, conformity assessment has become a critically important aspect of conducting business in today's global marketplace.

Conformity assessment provides benefits to diverse groups in the supply and demand chain. This includes the consumer, manufacturers, service providers, value-added resellers, and businesses.

• Consumers benefit from conformity assessment as it provides confidence and assurance that a product or service they purchase meets standards of quality [1].

• For manufacturers, conformity assessment minimizes the need to undergo multiple service providers' internal test programs.

• For service providers, conformity assessment readily identifies products that conform and are compliant with the industry standards. For service providers, conformity assessment readily identifies products that conform and are compliant withthe industry standards.

Certified products ensure that key functionality is implemented, while providing interoperability across multiple vendors' solutions.

Examples

The IEEE conformity assessment program has successfully been applied. One conformance certification is used by the Telecom and Power and Energy[2] Industry for IEEE Standard 1588, which is an IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems, developed by the IEEE Instrumentation and Measurement Society. The Telecom Industry has used the IEEE certification to demonstrate compliance with IEEE Standard 1588, thus providing improved networking ability in factory automation, test and measurement, and other telecommunications applications that require very close time synchronization. The 1588 Standard enables precise synchronization of clocks in measurement and control systems implemented with technologies such as network communication, local computing, and distributed objects, as does the IEEE Standard C37.118.1, IEEE Standard for Synchrophasor Measurements for Power Systems. Each of these programs began issuing certifications in 2015 [3].

More information on IEEE Standard 1588-2008 can be found at:

https://standards.ieee.org/findstds/standard/1588-2008.html

The IEEE Power and Energy Society's Power System Relaying Committee has been working on synchrophasor standards for 20 years. Phase measurement units (PMUs) and synchrophasor data provide grid operators with a more precise and higher resolution image of the grid in real-time. Such time sensitive information can help in numerous ways to alert the grid operators to grid stress early on, potentially avoiding power outages and maintaining power quality.

The U.S.-Canada investigation into the 2003 Northeast blackout recommended that synchrophasor data systems be installed immediately across North America for this purpose, hypothesizing that had such a system been in operation the August 14, 2003 blackout preconditions could have been identified, understood, and mitigated without the subsequent grid collapse.

Country-wide or multi-national power systems cover thousands of miles and thousands of pieces of equipment. Keeping such a large distributed system operational 100 percent of the year is a challenge. The Power and Energy industry is gaining unprecedented value from having phasor measurement units undergoing testing and certification prior to deployment to help ensure smooth functioning of the power grid. The IEEE certification process is being used to improve quality control, interchangeability of parts, reliability, development of plug and play devices, and to help ensure global acceptance of products and services for these industries.

A more thorough explanation of Synchrophasor Measurements can be found at

http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=7052413

One description of the benefits of Synchrophasors is described in

https://www.smartgrid.gov/files/Synchrophasor_Report_08_09_2013_DOE_2_version_0.pdf

More to Come

ICAP has additional program committees in the process of completing test suites across a range of technology areas, including communications, power, and consumer electronics. The IEEE-SA supports these volunteer committees facilitating test suite development activities, interaction with laboratory partners for piloting and executing testing per the test suites, and providing an infrastructure for certification operations. The program also supports interoperability test events, relevant workshops and webinars of interest to the conformity assessment ecosystem, and serves as a centralized conformity assessment resource for the IEEE community.

The Importance of Conformity Assessment

Consumers, manufactures, and service providers have well-deserved expectations relating to product reliability, efficiency, and interoperability. Conformity assessment reflects the development of well-defined processes and procedures, which are designed to evaluate and confirm features and functionality defined by industry standards. Conformity assessment has spurred innovation and has been a key factor in enabling products from a variety of producers to freely move across territorial boundaries with the assurance of operating safely, reliably, and within certain measurable performance criteria.

More information about ICAP can be found at:

http://standards.ieee.org/about/icap/index.html

ONF use UML Information Models to Generate Yang

By Nigel Davis, Ciena Fellow and Co-Chair, ONF Information Modeling Project Team

As part of the work to define the SDN controller, several ONF teams have focused on modeling network resources using UML. The tool chosen for this is the Eclipse platform-based open source UML tool, Papyrus. The ONF Core Information Model, which is at version 1.1 (released November 2015), provides a recursive model of network resources that can be applied anywhere from the most abstract level of view through to the device view. The model is freely available via the ONF library (https://www.opennetworking.org/sdn-resources/technical-library and https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/ONF-CIM_Core_Model_base_document_1.1.pdf). The model is derived from previous work by other bodies, notably ITU-T and the TM Forum, and the earlier 1.0 version of the ONF Core IM has been made available via ITU-T in G.7711. Other bodies are exploring/using the ONF model as a basis for their network modeling activities.

While important as an enabler of consistency and coherence, the ONF Core IM alone is not sufficient. So the ONF teams have also focused on interface design. The approach in ONF is to take the Core IM and use an emerging formal process of "pruning and refactoring" the model to form user centric interface views, where pruning is narrowing of the definitions so as to match the needs of the view application, and refactoring is a controlled compacting of the model to remove redundant elements and structure. Both pruning and refactoring maintain the semantics. The user centric views are presented as UML Information Models with clear traceability back to the Core Information Model. The ONF teams have been working on the definition of the transport API as well as the definition of the model for intent-based interaction (which focuses on constraint based outcome oriented requests not cluttered by technology detail).

Again, a user centric view in an Information Model alone is not sufficient; the key is code. To that end, the ONF teams have developed UML to Yang tooling that takes the Papyrus UML output for a view (decorated with the appropriate stereotypes enhancing the UML model definitions) and generates Yang schema that can then be fed into existing Yang tool chains. The UML to Yang tooling is open source and is available via the ONF supported Eagle project (https://github.com/OpenNetworkingFoundation/EAGLE-Open-Model-Profile-and-Tools). The ONF teams are also considering tooling that produces other interface schema forms (especially for constraint based interfaces) and are currently working toward tooling support for the pruning and refactoring process.

Incidentally, in the ONF modeling work, the term Data Model is not used due to concerns over ambiguity of distinction between IM and DM.

Multi-SDO NFV Information Modelling Workshop

Michael Brenner (ClearPath Networks), ETSI NFV ISG Vice-Chair

ETSI's NFV Industry Specification Group (ETSI NFV) convened a ground-breaking industry workshop to study how information modelling approaches could be aligned across the industry to simplify automation for network operators. The workshop brought together the leading Standards Development Organizations (SDOs) and Open Source communities hosted by CableLabs at their Louisville, Colorado (USA) location.

This was a significant and unique event because it is the first time the key standards organizations and open source communities have met together with a common purpose to accelerate alignment of their activities in relation to NFV. Participating organizations included 3GPP, ATIS, Broadband Forum, DMTF, ETSI NFV, IETF, ITU-T SG 15, MEF, OASIS/TOSCA, Open Cloud Connect, ONF, OpenDaylight, OPNFV, and the TM-Forum. Organizations that did not participate in this workshop are welcome to get involved in this collaboration by contacting ETSI.

Different information models and data models are being used among SDOs and open source communities resulting in fragmentation and complexity for implementation leading to increased cost and delaying time to market. Alignment of information models brings clarity of definition and drives consistent open APIs to enable integration across the entire ecosystem, including SDN and NFV.

As a result of this workshop there is increased understanding of the challenges and opportunities in the development and adoption of various modelling approaches. There was positive feedback from many delegates on the high value of the workshop and the increased awareness of efforts in peer organizations, and individual commitment by key experts will significantly boost prospects for industry alignment going forward.

The workshop was co-chaired by ETSI NFV Vice-Chair Michael Brenner (ClearPath Networks) and Klaus Martiny (Deutsche Telekom), ETSI NFV Network Operator Council Vice-Chair. A collaboration plan was agreed on to achieve meaningful progress by the end of 2016. The first feedback from participating organizations will be in March 2016. The participating organizations will independently progress their work, mindful of the collaboration milestones, and regular conference calls will take place to monitor progress, and when necessary another workshop will be convened to promote alignment.

Don Clarke (CableLabs), Chair of the ETSI NFV Network Operator Council, said: "It is vitally important that standards development organizations foster closer collaboration as the industry moves toward a new era of software-based networking. This workshop surprised all of us by the quality of inputs from an unprecedented number of participating SDOs, and the enthusiasm of delegates to maintain momentum in follow-ups bodes well for good progress over the coming weeks and months."

Report from ITU-T Study Group 15 Question 14

Hing-Kam Lam, ITU-T Study Group 15 Question 14 Rapporteur

ITU-T Study Group 15 Question 14 is responsible for the management/control requirements and information and data models of telecommunications transport equipment. The current work program includes the definition of a generic protocol-neutral information model for transport resource (G.7711/Y.1702 (08/15)), protocol-neutral information models for Carrier Grade Ethernet (G.8052), MPLS-TP (G.8152), and OTN (G.874.1), and the management requirements documents for those transport technologies.

Effective information modeling requires the use of UML tools. Over the past couple of years Q14 has been working with industry partners such as the ONF and MEF to develop a set of guidelines for the use of UML in information model development for telecommunications equipment management. G.7711 has an annex with guidelines for the use of UML. Q14 has standardized on the Eclipse Papyrus as the UML tool of choice. Working with the ONF, a set of guidelines that support the use of Papyrus has been created (ONF TR-515 Papyrus Guidelines 1.1). Utilizing the same tool set and guidelines documents allows the ITU-T to work closely and remain aligned with industry partners such as the ONF, MEF, and TM Forum.

Information models are an integral part of a model-driven lifecycle. As the network ecosystem becomes more agile, orchestrated, software-defined, and virtualized, the creation of standardized APIs needs to become faster as well. Q14 is evaluating mechanisms to make the traditional standards development process faster through the use of tooling that allows for information models to be translated into protocol-specific data models (such as YANG or JSON Schema). The ability to drive API development and documentation directly from the Eclipse environment (using a combination of Papyrus and GenDoc) is increasing the efficiency of standards development.

WHAT'S NEW AT THE MEF?

SCOTT MANSFIELD, ERICSSON, MEMBER OF THE MEF BOARD

The MEF Forum (MEF) continues to build on "The Third Network" messaging, which provides an open, secure, on-demand, orchestrated framework for connectivity services. In addition to the Lifecycle Service Orchestration (LSO) reference architecture work, the MEF is organizing engagements with the Open Source community to validate and iteratively enhance the LSO definition.

The LSO reference architecture identifies the actors in the ecosystem and the interfaces that need to be defined between the actors. This allows for the development of open application programming interfaces (APIs) that can be leveraged to enable end-to-end connectivity services.

The MEF is changing its way of working to become a software-aware standards development organization. In order to move from standards that are used to build static connectivity, a more agile approach that allows iterative development of standards along with open running code is needed. The framework requires the ability to prove interoperability and ensure that the services adhere to the specifications. New certification programs and hackathons provide the mechanisms necessary to exercise the framework. Coordination and collaboration with industry partners (including other Standards Setting Organizations, Industry Fora, and Open Source projects) are leveraged to address the complete end-to-end connectivity service ecosystem. Establishing and executing on new open-source coordination initiatives such as OpenCE will provide practical experience with exercising the new agile standards development paradigm.

Exploring different strategies of engagement will continue throughout 2016 for the MEF. In particular, a memorandum of understanding between the MEF and the ITU-T established a high-level framework of cooperation targeting the advancement of Carrier Ethernet and future agile, assured, and orchestrated services, with emphasis on trust/assurance, conformance testing, 5G non radio area, LSO, SDN, NFV and Open Source implementation.

TM FORUM UPDATE ON WORK RELATED TO INFORMATION MODELING

KENNETH DILBECK, TM FORUM

The Information Framework (previously known as the SID) is where the TM Forum focuses its efforts on information modeling. The Information Framework provides a reference model and common vocabulary for all the information required to implement the business processes necessary for a digital service provider. It reduces complexity in service and system integration, development, and design by providing an off the shelf information model that can be quickly adopted by all parties.

For background, the Information Framework currently is composed of approximately 2000 entities, more than 500 added in the last few years in diverse areas such as metrics, charging and billing, payments, umbrella information model, and most recently virtualized network functions. The Information Framework is grouped into eight domains: marketing sales, product, customer, service, resource, engaged party, enterprise and common. The Information Framework is much more than a resource model, and provides the end-to-end associations and relationships necessary to support a complete business.

The Umbrella Information Model (UIM), added a little over a year ago, is the result of a joint work effort between NGMN, 3GPP, and the TM Forum to create a set of umbrella entities that would allow the models from 3GPP and the TM Forum to be integrated in a much easier and predictable way. This work has been adopted by both 3GPP and the TM Forum, and has been an integral part of a number of TM Forum Catalyst projects (specialized POCs) that have proven that the UIM is very useful in managing hybrid networks.

The UIM is also key to a work effort currently being considered by MEF, ONF, and the TM Forum to coordinate modeling related to network resources. MEF and ONF have used the Information Framework as a launching point for their detailed work, so it is believed it would be very beneficial to the industry to coordinate the network resource modeling to eliminate any overlap and maximize the effort of a limited pool of subject matter experts.

The Information Framework has also been the basis for the APIs developed by the TM Forum. This places the APIs into a much broader end-to-end solution context and increases their reusability and interoperability. This is true for the legacy MTMN, MTOSI, OSS/J APIs, and has recently been extended to a set of REST APIs that have been developed over the last two years. Currently there are a number of APIs available for public download at http://projects.tmforum.org/wiki/display/API/TM+Forum+Ecosystem+API+-Portal.

The APIs are available for trouble ticket, product catalog, customer management, product inventory, product ordering, billing management, party management, SLA management, usage management, and performance management. There are a number of APIs currently under development and more in the pipeline addressing areas such as virtualization, security, onboarding, and activation.

Looking to the future, we will be increasing our ability to support NFV/SDN and virtualization, IoT, and specific initiative such as SmartCity, SmartClimate, and IoE.

ETSI NFV ISG: THE BUSINESS TRANSFORMATION OF NETWORK OPERATIONS THROUGH VIRTUALIZATION OF NETWORK FUNCTIONS

STEVEN WRIGHT, AT&T, CHAIR ETSI NFV ISG

The NFV ISG vision of an open ecosystem for NFV enables rapid innovation in networks and services for vendors, network operators, and service providers. The Network Functions Virtualization concept (NFV) was first introduced in a joint-carrier white paper published in October 2012.

Innovation in end-to-end services is enabled by software-based deployment and operationalization of virtualized network functions (VNFs) on independently deployed and operated NFV infrastructure based on high volume industry standard servers.

Network operator business objectives driving the need for operational transformation through virtualization include:

- Rapid service innovation through software-based deployment and operationalization of network functions and end-to-end services.
- Improved operational efficiencies

resulting from common automation and operating procedures.

•Reduced power usage achieved by migrating workloads and powering down unused hardware.

•Standardized and open interfaces between network functions and their management entities so that such decoupled network elements can be provided by different players.

•Greater flexibility in assigning VNFs to hardware.

•Improved capital efficiencies compared with dedicated hardware implementations.

The ETSI NFV mission is to facilitate the industry transformation and development of an open, interoperable ecosystem. As the focal point for the NFV ecosystem, the forum develops and maintains core NFV documentation, as well as collaborative relationships with other specialist SDOs and industry alliances, including open source communities. The ETSI NFV documentation and related open source implementations enable open innovation in the design of VNFs and end-to-end network services composed from them.

Since the initial meeting in January 2013, the ETSI NFV membership has grown to approximately 300 companies (including 38 network operators) and has delivered key NFV specifications covering requirements, use cases, architectural framework, terminology, management & orchestration, security, performance, and reliability. It has also sponsored multi-vendor proofs of concepts to encourage interoperability and growth of an open ecosystem around virtualized network functions. Additional specifications are being developed for publication during 2016.

CALL FOR PAPERS
IEEE COMMUNICATIONS MAGAZINE
NEXT GENERATION 911

BACKGROUND

In most parts of the world where a phone is used, users depend on the SOS, or Emergency service (ES). In North America, this service is accessible by dialing 911 and in most parts of Europe by dialing 112. Since its introduction in the 1950s, the service has adapted to technology changes: from wire line phones with fixed locations to supporting mobile phones. However, the technology is improving faster than the ES infrastructure can keep up. Faster networks coupled with the heterogeneous access methods pose a challenge to the evolution of the ES. As if this is not enough, the richer choices available to a user today for communication --- video, text messaging, social networking portals like Facebook and Google+, instant messaging, web-based calling, and over-the-top Voice over IP applications – make the task of providing ES uniformly ever more difficult. The US and Europe have decided to approach ES through a clean-slate approach. The redesign of the ES in US is known as Next Generation 911 (NG911) and in Europe as Next Generation 112 (NG112). The core of NG ES is Emergency Services IP Network (ESInet). ESInet uses the Session Initiation Protocol (SIP) to deliver voice, video, text and data calls reliably and uniformly to the ES network.

Authors from industry and academia are invited to submit papers for this FT (Feature Topic) of IEEE Communications Magazine on next generation 911. The FT scope includes, but is not limited to the following topics of interest:

- Overview of NG911 implementation efforts
- Technical challenges in implementing NG911 systems
- Status of current implementation of NG911
- Status of NG911 deployment
- Issues in multi-modal NG911 devices
- Impact of social media on NG911 systems
- Societal impacts of the presence (or absence) of such services
- Over-the-top applications and NG911
- Security and privacy of NG911
- Public policy and funding issues with NG911
- Regulatory environment for NG911
- The role of standards (IETF, ITU-T) in NG911

SUBMISSIONS

Articles should be tutorial in nature, with the intended audience being all members of the communications technology communities. They should be written in a style comprehensible to readers outside the specialty of the article. Mathematical equations should not be used (in justified cases up to three simple equations are allowed). Articles should not exceed 4500 words. Figures and tables should be limited to a combined total of six. The number of archivable references is not to exceed 15. Complete guidelines for manuscript preparation can be found via the link http://www.comsoc.org/commag/paper-submission-guidelines. Please send a PDF (preferred) or MS-Word formatted paper via Manuscript Central (http://commag-ieee.manuscriptcentral.com). Register or log in, and go to the Author Center. Follow the instructions there. Select "November 2016 / Next-Generation 911".

IMPORTANT DATES

- Manuscript Submission Deadline: March 15, 2016
- Decision Notification : June 30, 2016
- Final Manuscript Due Date: August 31, 2016
- FT Publication Date: November 2016

GUEST EDITORS

Vijay K. Gurbani	Salvatore Loreto	Ravi Subrahmanyan	Carol Davids
Bell Laboratories, Alcatel-Lucent, USA	Ericsson, Sweden	Butterfly Network Inc., USA	Illinois Inst. of Technol., USA
vkg@bell-labs.com	salvatore.loreto@ericsson.com	ravi.subrahmanyan@ieee.org	davids@iit.edu

Innovation & Competition: Succeeding through Global Standards: A New Massive Open Online Course Delivered on IEEEx.org

By Yatin Trivedi, Member, IEEE Standards Education Committee

We know that many products and technologies that influence and transform the way we live, work, and communicate rely on the development of technical standards. Standards fuel compatibility and interoperability, reduce costs and risk, simplify product development, enable innovation, hasten time-to-market for new products, support commerce, and can determine a company's global competitiveness. Standards play a vital role in helping consumers understand and compare competing products, and give confidence to investors.

In March 2016, a new six-week massive open online course (MOOC) entitled "Innovation & Competition: Succeeding through Global Standards," will begin at the IEEE Standards University. This new course offers a practitioner's view of standards and is geared to graduate students and emerging professionals in the fields of engineering, technology, computing, business, economics and law, particularly those working, or planning to work, in all facets of product planning, development, launch, and support.

The course will enable participants to better contribute to their organizations and advance their careers. University faculty in related areas will also benefit from taking the course and be able to add value to students through their teaching. By completing this MOOC, students can expect to come away with an understanding of:

- The different types of standards.
- How standards impact trade and innovation.
- How standards evolve over time.
- Why companies participate in standards development.
- How standards are changing to meet emerging needs.
- How related activities can be integrated with other organizational functions.
- Strategic implications.
- How standards can be applied to product design and planning.

Consideration will also be given to related conformity assessment, regulation, and intellectual property management.

Although educators may recognize the importance of standards and desire to incorporate standards topics into their course curriculum, they may lack access to industry expertise in this arena. The information provided through this course can be a vehicle for professors to invite industry standards skills into their classrooms.

The IEEE Standards Education Committee has delivered successful standards education workshops for several years, with very positive feedback. This course will significantly increase the impact of the IEEE's standards education efforts to a global audience. Students will be able to attend the course at no cost or receive a certificate of completion for a small fee.

Funding for this course was provided by the IEEE Foundation Fund (FF) and the IEEE Life Members Fund (LMF). This is a joint project of IEEE Educational Activities and the IEEE Standards Association. The MOOC is part of the overall new IEEE Standards University.

Visit IEEEx.org to watch a video introduction about the course and to enroll.

The IEEE Standards University e-Magazine

This online publication is sponsored by the IEEE Standards Education Committee, a joint committee of the IEEE Educational Activities Board and the IEEE Standards Association. Serving the community of students, educators, practitioners, developers, and standards users, we are building a community of standards education for the benefit of humanity.

IEEE's Standards Education Program is committed to:

- Promoting the importance of standards in meeting technical, economic, environmental, and societal challenges.
- Disseminating learning materials on the application of standards in the design and development aspects of educational programs.
- Actively promoting the integration of standards into academic programs.
- Providing short courses about standards needed in the design and development phases of professional practice.

The March edition (http://www.standardsuniversity.org/e-magazine/) contains a number of articles on trust and cyber security.

Abstracts

The Evolving Internet of Things (IoT) Requires New Approaches by Colleges and Universities in Educating All Students in Cyber Security
By William Butler, William "Vic" Maconachy, Helen G. Barker

Abstract

The Internet of Things (IoT) is quickly evolving, with 2016 predicted to be a critical year for the growth of connected devices, the sheer volume of data captured, and the types of 'things' connected (human-to-machine, machine-to-machine). The machine-to-machine category is growing at an exponential rate, along with the data captured and communicated to databases for categorization and analysis. Colleges and universities are challenged to educate students across the curriculum to understand the underlying technologies and the application of this technology to solve everyday problems. Students must learn that the IoT surrounds them and is already a critical aspect of their daily lives. Business, engineering, computer science, and cyber security departments across the country must plan to address student awareness through revamped departmental curricula and interdisciplinary opportunities across departments. It is this generation of future workers who will be tasked to solve the issue of security within the IoT. This article does not advocate a new degree, but rather a comprehensive interdisciplinary systems approach.

IoT Security Standards—Paving the Way for Customer Confidence
By Alan Grau

Abstract

With the opening of the Consumer Electronics Show in Las Vegas, the IoT has moved beyond the initial hype phase and even past the phase of early deployments. In July of 2014, HP Labs did a study of 10 popular IoT devices and found that the security was shockingly bad. The researchers studied 10 devices, looking at the end-to-end security capabilities of these devices, including privacy protection, authorization, encryption, user interface protection, and code security. They found that 70% of the devices had at least one *major* vulnerability! By the time they completed their study,

the researchers identified more than 250 vulnerabilities, an average of 25 security vulnerabilities per device. Security was clearly an afterthought—or worse—for these devices. An average consumer, or even a security savvy consumer, has little ability to know which brand of IoT device has better security or even any security. An OEM may claim "built-in security", but that phrase alone means little. This article discusses the role and importance of security standards, challenges in creating security standards, and what's emerging in the area of standards for IoT security.

Internet of Things Requirements and Protocols
By Kim Rowe

Abstract

More and more protocols are being added for the Internet of Things (IoT) as large vendors address the deficiencies of their products. These higher level IoT protocols are suitable for a broad range of applications. Knowing the correct protocol or set of protocols for a given application which cover communication, security, management, and scalability, is the first design consideration. After this, the best implementation of each of the protocols must be understood. From this understanding, the designer can select the optimal implementation of each protocol for the system and then from these, select the best set of protocol implementations for the system. For the IoT protocol space, standards are not yet converged for particular applications, and the market ultimately decides which of these standards are most relevant. This article examines the range of protocols available and the specific requirements that drive the features of these protocols, and considers the implementation requirements to build a complete system.

Full Articles

Protecting Against Cyber Threats

By Sangeeta Kodukula
Security Consulting Systems Engineer, Cisco Systems

In today's digital age we cannot argue with the idea that the evolution of technology has contributed to many conveniences in our day-to-day life. Online banking, the ability to access corporate data from smart phones, and e-commerce, are just a few examples of how the digital evolution has changed the landscape of how we operate on a daily basis. While these conveniences have contributed to simplifying our daily operations, have we taken a moment to think about the possible repercussions of exposing sensitive data on the Internet? Similar to how the landscape of technology has changed, so has the evolution of the cyber threat.

Today's hackers are more sophisticated than ever, and can make more than $1 million a year! Methodologies such as leveraging exploit kits to deliver malware, ransomware, Distributed Denial of Service (DDoS) attacks, and phishing attacks, are just a few examples of how these hackers infringe on their victims. Industry and the public sector are just as vulnerable as consumers who post sensitive data on the Internet. As a result, regulations such as PCI (Payment Card Industry) compliance, HIPAA (Health Insurance Portability and Accountability), and FIPS (Federal Information Processing Standard), are all examples of guidelines companies in these various verticals must adhere to in order to protect sensitive data from being compromised. These regulatory compliances consist of a comprehensive checklist of requirements companies must align with in order to pass their compliance audits. These standards evolve as threats evolve. These compliances require industry to invest in IT security in order to protect and reassure their customers' sensitive data cannot be accessed from an unauthorized source. IT departments must invest in various layers of security by implementing Next-Gen Firewalls, Next-Gen Intrusion Prevention System (IPS) solutions, encryption technologies, and VPN technologies, to name a few solutions. Why should companies care to invest in these technologies? The cost of a potential breach can cost a company millions of dollars, and it can bruise their overall brand reputation and cause their stock price to go down.

From a consumer standpoint we can take necessary precautions in order to minimize our risk as well. Measures that can be taken include creating strong passwords, changing passwords frequently, not openly distributing bank account details, and checking credit card statements thoroughly. These measures will help prevent one from becoming the next victim of cyber-crime.

Further reading:

PCI Compliance Standards: https://www.pcisecuritystandards.org/

HIPAA Compliance Standards: http://www.hhs.gov/hipaa/for-professionals/security/laws-regulations/

FIPS Compliance Standards: http://csrc.nist.gov/publications/fips/fips140-2/fips1402.pdf

Security and IoT in IEEE Standards

By Cherry Tom
Emerging Technologies Intelligence Manager, IEEE Standards Association

Security elements have been included in numerous IEEE standards and standards projects over many years. If one searches the IEEE standards status report[1] by entering "security," and views the project scope, purpose, and/or abstract, multiple references to security can be seen. These standards and standards projects cover topics as diverse as vehicle communications, smart grid technologies, personal health devices, networking, mobile devices, and storage devices. All these and more could conceivably be part of the Internet of Things (IoT). A number of these standards were developed before the term "Internet of Things" became widely used.

IoT Architecture

IEEE has a specific initiative (one of IEEE's important, multi-disciplinary, cross-platform initiatives) for the Internet of Things (IoT). The IEEE IoT website includes a link for educational resources such as webinars, other videos, and podcasts. The link to the IEEE-SA IoT website is for standards and related information. In particular, the project IEEE P2413, Standard for an Architectural Framework for the Internet of Things (IoT), has a subworking group focused on Quadruple Trust, i.e. "Protection, Security, Privacy and Safety". This involves a holistic end-to-end approach, including development of a threat model for IoT [2]. This considers the various vertical applications for IoT and documentation of architecture needs to address the threat model. The participants in IEEE P2413 include representatives from major corporations involved in IoT from regions around the world and provide expertise in all aspects of IoT, including security and compliance. To involve startup companies, IEEE-SA hosts a number of events where the companies can present their projects for evaluation as well as learn about the IEEE's activities in IoT.

The following are examples of IEEE standards and projects related to security and IoT.

Cryptography

- The IEEE 1363 series of standards for public key cryptography, beginning with IEEE 1363-2000, IEEE Standard Specifications for Public-Key Cryptography, and including IEEE 1363a-2004, IEEE 1-2008, IEEE 1363.2-2008, and IEEE 1363.3-2013, is developed by the 1363 WG.
- The IEEE 1619 series of standards for encryption in storage media, beginning with IEEE 1619-2007, IEEE Standard for Cryptographic Protection of Data on Block-Oriented Storage Devices, and continuing with IEEE 1619.1-2007, and IEEE 1619.2-2010, is developed by SIS-WG, Security in Storage Working Group.

Devices and Sensors/Actuators

•Within the IEEE 1451/21450 [3]/21451 series of standards for transducers for sensors and actuators, including IEEE 21451-1-2010, IEEE 21451-2-2010, IEEE 21451-4-2010, and IEEE 21451-7-2011, a new project, IEEE P24151-1-4, Standard for a Smart Transducer Interface for Sensors, Actuators, and Devices – eXtensible Messaging and Presence Protocol (XMPP) for Networked Device Communication, being developed by the XMPPI–XMPP Interface Working Group, specifically addresses issues of security, scalability, and interoperability in session initiation and protocol transport.

•IEEE 2410-2015, IEEE Standard for Biometric Open Protocol, provides "Identity assertion, role gathering, multilevel access control, assurance, and auditing" [4], and was developed by the BOP – Biometrics Open Protocol working group.

•A new project approved in 2015, IEEE P1912, Standard for Privacy and Security Architecture for Consumer Wireless Devices, being developed by the P1912 WG, will describe a common communication architecture and approaches for end user security, including device discovery/recognition, user authentication, and user control of tracking items/people and sharing of information.

•IEEE 2600-2008, IEEE Standard for Information Technology: Hardcopy Device and System Security, covers printers, copiers, and multifunction devices. It defines security requirements such as authentication, authorization, privacy, integrity, device management, physical security, and information security.

Networking for IoT

•IEEE 802.1X-2010, IEEE Standard for Local and metropolitan area networks–Port-Based Network Access Control, covering common architecture, functional elements, and protocols for mutual authentication and secure communication between the clients of ports attached to the same LAN, and its amendment, IEEE 802.1Xbx-2014, were developed by the 1 – Higher Layer LAN Protocols Working Group.

•IEEE 802.1AE-2006, IEEE Standard for Local and Metropolitan Area Networks: Media Access Control (MAC) Security, specifies "how all or part of a network can be secured transparently to peer protocol entities that use the MAC Service provided by IEEE 802 LANs to communicate"[5]. Its amendment, IEEE 802.1AEbw-2013, expands its security capabilities. These were developed by the 1 – Higher Layer LAN Protocols Working Group.

•IEEE 802.1AR-2009, Standard for Local and metropolitan area networks – Secure Device Identity, enables the secure association of locally significant device identities with manufacturer provisioned identities for use in provisioning and authentication protocols, and was developed by the 1 – Higher Layer LAN Protocols Working Group.

•The latest editions of IEEE 11-2012, IEEE Standard for Information technology–Telecommunications and information exchange between systems Local and metropolitan area networks–Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, developed by the WG802.11 – Wireless LAN Working Group, and IEEE 802.15.4-2015, IEEE Standard for Local and metropolitan area networks–Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs), developed by the WG802.15 – Wireless Personal Area Network (WPAN) Working Group, include extensive sections on security.

•IEEE project 15.9, IEEE Draft Recommended Practice for Transport of Key Management Protocol (KMP) Datagrams, developed by the WG802.15 – Wireless Personal Area Network (WPAN) Working Group, provides guidelines for support of key management in IEEE 802.15.4.

•IEEE 802.21a-2012, IEEE Standard for Local and Metropolitan Area Networks: Media Independent Handover Services – Amendment for Security Extensions to Media Independent Handover Services and Protocol, was developed by the 21 – Media Independent Handoff Working Group.

•The IEEE 1888 series, beginning with IEEE 1888-2014, IEEE Standard for Ubiquitous Green Community Control Network Protocol, and including IEEE 1888.1-2013 and IEEE 1888.2-2014, has a specific standard for security, IEEE 1888.3-2013, IEEE Standard for Ubiquitous Green Community Control Network: Security, and was developed by UGCCNET-SEC/P1888.3 WG – Ubiquitous Green Community Control Network: Security Working Group/ UGCCNET-SEC/P1888.3. It includes security requirements, architecture, authentication, authorization, and security procedures and protocols.

Infrastructure Systems

(Note: intranets may incorporate IoT while not necessarily connected to the public Internet.)

•IEEE 692-2013, IEEE Standard for Criteria for Security Systems for Nuclear Power Generating Stations, developed by WG 3.2 – Security Systems Working Group, addresses security system equipment for "detection, assessment, surveillance, access control, communication, and data acquisition".

•The numerous IEEE smart grid systems standards [6] include a number focused on security, e.g. IEEE C37.240-2014 – IEEE Standard Cybersecurity Requirements for Substation Automation, Protection, and Control Systems, developed by 240 WG – PC37.240 Cyber Security Standard, and IEEE 1686-2013 – IEEE Standard for Intelligent Electronic Devices Cyber Security Capabilities, developed by WGC1 – Substations Working Group C1.

Other Considerations

It should be noted that IEEE P2413 includes in its definition for properties of the "thing" in the Internet of Things, virtual properties such as might be derived from big data analysis. The IEEE Big Data Initiative includes standards development as a key focus area. Privacy and security remains a concern for Big Data.

While not official IEEE standards, the documents "Building Code for Medical Device Software Security" and "Avoiding the Top 10 Software Security Design Flaws" provide guidance for software designers, including those involved in software for IoT. They were developed as part of the IEEE Cybersecurity Initiative.

In addition to IEEE, other organizations are also involved in standards for IoT and security. Another article, "IoT Interoperability Requires Security", includes along with IEEE, descriptions of the work in several of these organizations.

References

[1] http://standards.ieee.org/develop/project/status.html
[2] http://grouper.ieee.org/groups/2413/Intro-to-IEEE-P2413.pdf
[3] http://standards.ieee.org/findstds/standard/21450-2010.html
[4] http://standards.ieee.org/findstds/standard/2410-2015.html
[5] http://standards.ieee.org/findstds/standard/802.1AE-2006.html
[6] http://smartgrid.ieee.org/resources/standards

Patents and Standardization, Part 1: A Tutorial on Patents

In an ideal world, the standards promulgated by standards-development organizations (SDOs), such as the IEEE, ITU, and ETSI, would always result from an impartial comparison of proposed solutions made in an environment "free of distortion, bias, or manipulation by private interests." In the real world, however, standards tend to be a result of compromises and, in some cases, hard-fought battles among parties, often competitors, who have diverse private interests.

Krista S. Jacobsen

Abstract

In an ideal world, the standards promulgated by standards-development organizations (SDOs), such as the IEEE, ITU, and ETSI, would always result from an impartial comparison of proposed solutions made in an environment "free of distortion, bias, or manipulation by private interests."[1] In the real world, however, standards tend to be a result of compromises and, in some cases, hard-fought battles among parties, often competitors, who have diverse private interests. Considerations beyond technical merit are generally in play. For better or worse, one of those considerations is often the understanding that intellectual property may protect a candidate solution. Of the four types of intellectual property — patents, trademarks, copyrights, and trade secrets — the one that looms largest in the standards-development context is patents.

This article provides a foundational overview of patents and patent systems with a specific focus on the patent laws and procedures of the United States. Where feasible, the article also identifies similarities and differences in the patent laws and procedures of other countries.

Policy Considerations for Patent Systems

There is perhaps no more compelling testimonial to the importance of patents than the fact that nearly every country grants patent rights. Patent systems reflect governments' determinations that the public benefits when inventors disclose their innovations publicly instead of keeping them secret. Thus, patent systems execute a bargain made between governments and inventors: in exchange for an inventor publicly disclosing how to make and use a new, useful, and nonobvious invention, the government grants the inventor a right to prevent others from exploiting that invention for a specified period of time, called the *patent term*, or at least the right to collect money from unauthorized parties who use the invention during the patent term.

Patent systems encourage innovation not only by offering a period of exclusivity to inventors in exchange for disclosure of their inventions, but also by encouraging others to innovate. The United States Supreme Court stated in *Aronson v. Quick Point Pencil Company* that "patent law … promotes disclosure of inventions to stimulate further innovation and to permit the public to practice the invention once the patent expires."[2]

As explained further below, "the invention" is defined by the claims of a patent. During the patent term, subject to some important caveats, the public is free to use all information disclosed in the patent except for the claimed invention.[3] After the patent expires, the claimed invention becomes part of the public domain.

What Is a Patent?

A patent is a government-issued, legally-enforceable document that defines an invention and grants the inventor the right to restrict others from exploiting that invention for a limited period of time. Most countries have enacted their own patent laws, which tend to be similar, but not identical, to the laws of other countries. For now, there is no international patent, and applicants must obtain separate patents in each country in which they wish to protect their inventions.[4] The patents granted by individual countries are valid and enforceable only in the country of grant. In other words, a United States patent has no effect in Canada, Germany, or China, and, conversely, a Canadian, German, or Chinese patent has no effect in the United States.

All countries disallow certain types of innovations or discoveries from being patented, regardless of how important they may be. In the United States, for example, abstract ideas, laws of nature, and natural phenomena are not patentable.[5] The overarch-

[1] Amicus Curiae Brief of the American Antitrust Institute in Support of Neither Party, *In re Rambus Inc.*, 2006 WL 2330117 (F.T.C. 2006) (No. 9302).

[2] 440 U.S. 257, 262 (1979); see also *Kewanee Oil Co. v. Bicron Corp.*, 416 U.S. 470, 496 (1974) ("The decision of Congress to adopt a patent system was based on the idea that there will be much more innovation if discoveries are disclosed and patented than there will be when everyone works in secret.").

[3] As explained below, there may be other applications or patents claiming other portions of the disclosure, and the claims of those applications or patents may further restrict the scope of information available for public use.

[4] Europe is in the process of establishing a "European patent with unitary effect," also known simply as the "unitary patent," which will enable a party to obtain a single patent that will be valid in multiple European countries.

[5] The state of the law of subject matter eligibility in the United States is murky at present, and attorneys, examiners, and courts charged with applying the Supreme Court's holdings in recent subject matter eligibility cases are struggling to define how to recognize claims for ineligible subject matter. For example, claims for graphical user interfaces, software, and methods of doing business may be abstract ideas.

The author is with Jacobsen IP Law.

 0163-6804/16/$25.00 © 2016 IEEE

ing policy concern is that granting a patent for an abstract idea (e.g., the creation of a contractual relationship), a law of nature (e.g., gravity), or a natural phenomenon (e.g., a mixture of bacteria that does not have significantly different characteristics from what occurs in nature) would remove "the basic tools of scientific and technological work" from the public domain.[6] The application of a law of nature in a new and useful way may be patentable in the United States, but as the Supreme Court recently stated, a claim "must limit its reach to a particular, inventive application" of the law of nature.[7]

Other countries also forbid the patenting of certain subject matter. For example, Canada does not issue patents for scientific principles, abstract theorems, methods of medical treatment or surgery, higher life forms, or forms of energy (e.g., electromagnetic and acoustic signals, regions of the electromagnetic spectrum, electric currents, and explosions) [1]. The European Patent Cooperation Treaty forbids the patenting of discoveries, scientific theories, mathematical methods, aesthetic creations, presentations of information, playing games, doing business, computer programs, and schemes, rules, and methods for performing mental acts [2]. Consequently, what is patentable subject matter in one country may not be patentable subject matter in another country.

The United States grants three different types of patents for three different types of innovations. Plant patents protect distinct and new varieties of plants that are invented or discovered and then asexually reproduced [3]. Design patents protect ornamental designs embodied in or applied to a manufactured article [4]. Utility patents, which are the most common type of patent, and also the type most likely to be implicated in SDO settings, protect processes, machines, manufactured products, or compositions of matter [5]. Utility patents generally have a term of 20 years from the earliest effective filing date.[8] Other countries grant patents that are similar to the utility patent of the United States.

It may surprise readers to learn that a patent does not confer to the patent holder any affirmative right — that is, a right to do something. For example, a patent does not give the patent holder a right to make a product incorporating the claimed invention. One reason is that by making a product incorporating the claimed invention, the patent holder might also be incorporating another party's innovation, which itself might be protected by a patent. Instead, a patent gives its owner the right to prevent others from making, using, selling, offering to sell, importing, or exporting the claimed invention during the patent term [6, 7]. Thus, at least in theory, a patent confers to its owner a right to exclude others from doing something they might otherwise do.[9] Litigation is often required to determine whether a patent owner can, in fact, prevent another party from pursuing an allegedly-infringing activity.

Like other assets, patents may be transferred from one owner to another. Thus, independent inventors who do not have the resources to take advantage of their patents may sell them to another party who might be in a better position to exploit those patents, such as by licensing the patents to others.[10] Likewise, a company whose employees have invented a new and useful technology may obtain patents to protect that technology and sell those patents to another entity.[11]

Like other assets, patents may be transferred from one owner to another. Thus, independent inventors who do not have the resources to take advantage of their patents may sell them to another party who might be in a better position to exploit those patents, such as by licensing the patents to others.

Patent Applications

The process to obtain a patent on an invention begins with the filing of a patent application in a patent office, such as the United States Patent and Trademark Office (USPTO), the European Patent Office, or the Japan Patent Office. Several countries, including the United States, require inventions made within their borders to be filed in that country's patent office before filing in any other country.[12] One reason for this requirement is so that the government may screen applications for any information that might jeopardize national security. In the United States, if the application passes the screening, the USPTO grants a foreign filing license that authorizes the applicant to file the application in foreign countries.

Assuming a patent applicant has complied with any foreign filing license requirements, the appli-

[6] *Gottschalk v. Benson*, 409 U. S. 63, 67 (1972) ("Phenomena of nature, though just discovered, mental processes, and abstract intellectual concepts are not patentable, as they are the basic tools of scientific and technological work.").

[7] *Mayo Collaborative v. Prometheus Labs.*, 132 S. Ct. 1289, 1302 (2012).

[8] The word "generally" is appropriate because a patent granted on an application filed before June 8, 1995 has a term of 17 years from its issue date. The significance of the effective filing date is discussed below.

[9] Although the right to exclude others from exploiting a patented invention is central to United States patent law, the Supreme Court's 2006 decision in *eBay Inc. v. MercExchange L.L.C.* weakened a patent owner's right to exclude. Before *eBay*, a court would automatically issue a permanent injunction barring an infringer from further infringement of a patent found to be valid and infringed. But in *eBay*, the Supreme Court held that a patent owner is entitled to a permanent injunction only if certain criteria are met. For example, unless the patent holder shows that money damages (i.e., payments to the patent holder from the infringer) are inadequate to compensate the patent holder for the infringement, the court will not grant an injunction to prevent an infringer from further infringement during the patent term. The patent holder is still entitled to receive money damages as compensation for the infringement, but the patent holder cannot stop the infringement altogether. Patent holders may still try to enjoin the importation of infringing products, however, by filing a complaint at the International Trade Commission (ITC).

[10] Parties whose businesses focus on acquiring and licensing patents, rather than developing and manufacturing products, are often pejoratively, and in some cases unfairly, referred to as "patent trolls." A less inflammatory term is "non-practicing entity."

[11] Most companies require employees to sign employment agreements obligating them to assign, to the companies, any inventions conceived in the scope of their employment.

A patent applicant files a nonprovisional application to start the examination process in the USPTO, which is conducted by a patent examiner. If a claimed invention meets all requirements for patentability, the examiner allows the application to issue as a patent.

cant may file a patent application directly in the patent office of each country in which the applicant seeks a patent. For example, an inventor may file a patent application in the USPTO to seek a United States patent, a separate application in the State Intellectual Property Office of the People's Republic of China to seek a Chinese patent, and yet another application in the Canadian Intellectual Property Office to seek a Canadian patent. Because each patent office has its own procedures and fees, and most countries only allow licensed practitioners to file patent applications, the cost and complexity of filing separate patent applications in multiple countries may be high.

As an alternative to filing separate patent applications directly in the patent offices of all of the countries in which an applicant seeks patents, an international treaty called the Patent Cooperation Treaty (PCT) enables applicants to file a single patent application in a designated patent office and later, from this single application, elect to seek separate patents in some or all of the 148 "contracting states" that are signatories to the PCT, including the United States and most other industrialized nations. When the PCT applicant later selects some or all of the 148 contracting states, the PCT application enters the so-called "national phase" in each of those contracting states, at which point it becomes a separate, stand-alone application in each of the selected contracting states. Each contracting state's patent office then determines whether to grant a patent based on the national-phase application.[13] The PCT therefore allows applicants to file and pay fees for a single "placeholder" application and postpone, typically for up to 30 months from the priority date of the application,[14] the decision of exactly where to seek separate patents on the basis of that single application.

In the United States, in addition to applications entering the national phase from PCT applications, there are two general classes of patent applications: nonprovisional and provisional. Other countries' patent applications are equivalent to the nonprovisional application in the United States, but may be subject to content or structure requirements that differ from those in the United States.

Nonprovisional Applications

A patent applicant files a nonprovisional application to start the examination process in the USPTO, which is conducted by a patent examiner. If a claimed invention meets all requirements for patentability, the examiner allows the application to issue as a patent. The interaction between patent applicants and examiners, called *prosecution*, is discussed in a section below.

In the United States, a nonprovisional application must include "a written description of the invention, and of the manner and process of making and using it, in such full, clear, concise, and exact terms as to enable any person skilled in the art to which it pertains, or with which it is most nearly connected, to make and use the same" [8]. Other countries have similar requirements.[15] Most nonprovisional applications include both a written description of the invention and one or more drawings to illustrate the invention, such as scale drawings, block diagrams, or flowcharts.

The legal scope of the invention for which the patent applicant seeks protection is defined by one or more *claims*. In the United States, the claims must particularly point out and distinctly claim the subject matter that the inventor regards as the invention [9]. Other countries have similar requirements.[16]

That the claims, and only the claims, define the scope of the invention cannot be overstated. Although the body of a patent application may disclose a host of information about the invention, including, for example, background or other information previously known and available to the public, it is the claims that establish the scope of the eventual patent holder's right to exclude. Moreover, the claims define what will *not* be within the eventual patent owner's right to exclude, because, subject to some caveats, the public may use information that is disclosed but not claimed in a patent.[17] Thus, although the claims of a patent application may be challenging for non-lawyers to read and parse, the claims are the most important part of a patent application.

The USPTO recognizes several types of non-

[12] See, e.g., 35 U.S.C. § 184(a) ("Except when authorized by a license obtained from the Commissioner of Patents a person shall not file or cause or authorize to be filed in any foreign country prior to six months after filing in the United States an application for patent … in respect of an invention made in this country… ."); Patent Law of the People's Republic of China, Article 20 ("Any unit or individual that intends to apply for patent in a foreign country for an invention or utility model accomplished in China shall submit the matter to the patent administration department under the State Council for confidentiality examination. …").

[13] One exception is in Europe, where the European Patent Office assesses patentability for the member states of the European Patent Organisation.

[14] The priority date of an application is explained below.

[15] See, e.g., European Patent Convention, Article 83 ("The European patent application shall disclose the invention in a manner sufficiently clear and complete for it to be carried out by a person skilled in the art."); Patent Act of Canada, § 27(3) ("The specification of an invention must … set out clearly the various steps in a process, or the method of constructing, making, compounding or using a machine, manufacture or composition of matter, in such full, clear, concise and exact terms as to enable any person skilled in the art or science to which it pertains, or with which it is most closely connected, to make, construct, compound or use it… ."); Australian Patents Act, § 40(2) ("A complete specification must: (a) disclose the invention in a manner which is clear enough and complete enough for the invention to be performed by a person skilled in the relevant art… .").

[16] *See*, e.g., European Patent Convention, Article 84 ("The claims shall define the matter for which protection is sought."); Patent Act of Canada, § 27(4) ("The specification must end with a claim or claims defining distinctly and in explicit terms the subject-matter of the invention for which an exclusive privilege or property is claimed."); Australian Patents Act, § 40(2) ("A complete specification must: … end with a claim or claims defining the invention… .").

provisional applications. In addition to original nonprovisional application filings, there are continuing applications, which claim the benefit of the filing date of an original nonprovisional application referred to as the parent application. Continuing applications include continuation, divisional, and continuation-in-part applications.

A continuation application includes the parent application's written description and drawings, but different claims. The claims of a continuation application may be directed to the same invention as claimed in the parent but using different language, or the continuation application may claim a separate invention disclosed by, but not claimed in, the parent application. For example, if the parent application discloses both an automobile engine and a steering mechanism but only claims the engine, a continuation application might claim the steering mechanism.

An applicant typically files a divisional application in response to an examiner's determination that the claims filed in a parent application are directed to multiple distinct inventions (i.e., they lack "unity of invention"). A patent may only claim a single invention, and the examiner might find, for example, that claims directed to an apparatus are for a different invention than claims directed to a method of making or using that apparatus. Consequently, the examiner may require the applicant to choose which one of the multiple distinct inventions to pursue in the application. The applicant may then pursue the claims of any non-elected inventions in divisional applications. Like a continuation application, a divisional application includes the parent application's written description and drawings, but different claims.

A continuation-in-part (CIP) application includes some or all of the disclosure of the parent application as well as new disclosure in the form of additional written description and, typically, drawings. Material in the CIP application that also appears in the parent application is entitled to the priority date of the parent application, whereas the priority date of the new material is the filing date of the CIP application. The priority date of an invention is important because, as discussed below, it establishes a cut-off date for information that could otherwise prevent the applicant from receiving a patent for the invention.

A continuing application must be filed in the USPTO before a patent issues from the parent application. As long as a parent application is pending, the patent applicant in the United States may file continuation, divisional, and CIP applications claiming the benefit of the parent application. The terms of patents issuing from those continuing applications are determined based on the effective filing date of the earliest-filed related nonprovisional application.

Provisional Applications

Instead of filing a nonprovisional application, an applicant may file a provisional application in the USPTO to establish a priority date for a purported invention without having to fulfill the formal requirements to file a nonprovisional application. The USPTO does not examine provisional applications, and therefore the provisional filing fee is substantially less than the fees due when filing a nonprovisional application. Moreover, there is no required format for a provisional application, and provisional applications do not need to include claims. Therefore, the time required to prepare a provisional application may be less than the time required to prepare a nonprovisional application.

It is important to recognize, however, that a provisional application establishes a priority date for a later-claimed invention only if the provisional application provides a disclosure that enables a person having ordinary skill in the art to make and use the later-claimed invention without undue experimentation. Many attorneys recommend drafting at least a few claims during the preparation of the provisional application to ensure that the provisional application establishes a priority date for the later-claimed invention. Consequently, the time required to prepare a high-quality provisional application may not be significantly less than the time required to prepare a nonprovisional application.

All provisional applications expire one year after filing. To preserve the priority date of material disclosed in a provisional application, within twelve months of the provisional application's filing date the applicant must either convert the provisional application to a nonprovisional application or, more commonly, file a nonprovisional application claiming the benefit of the provisional application.[18]

Patentability

A patent applicant is entitled to a patent in a selected country only if the invention meets that country's requirements for patentability. Although countries vary in their specific wording of the requirements for patentability, the common thread is that a purported invention is only patentable if it is new, has some constructive use, and is nonobvious over what is already known from the perspective of a person having ordinary skill in the relevant art.[19] The hypothetical person having ordinary skill in the art is presumed to be cognizant of the entire universe of *prior art* in existence on the *priority date* of the claimed invention [10].

The priority date of the claimed invention is the filing date of the earliest-filed related application that discloses the claimed subject matter. If an application does not claim the benefit of an earlier-filed application (i.e., it does not claim

In addition to original nonprovisional application filings, there are continuing applications, which claim the benefit of the filing date of an original nonprovisional application referred to as the parent application. Continuing applications include continuation, divisional, and continuation-in-part applications.

[17] As explained below, the applicant may amend the claims during prosecution, and therefore the claims originally filed may differ from the claims in the later-issued patent. Thus, relying on the claims of a published application to determine what information disclosed in the application is dedicated to the public may be risky. Also, as stated above and explained below, the applicant may file continuing applications while a parent application is pending, and these continuing applications may claim portions of the disclosure not claimed in the parent application.

[18] Converting a provisional application to a nonprovisional application adversely affects the term of any patent that eventually issues from the application. Thus, most patent attorneys recommend filing a nonprovisional application claiming the benefit of the provisional application.

Because a patent gives its owner the right to exclude others from taking advantage of the invention during the patent's term, only those innovations that are new, useful, and nonobvious relative to the prior art in existence as of their priority dates are worthy of a patent grant.

the benefit of a provisional or parent application), the priority date is the application's filing date. If, however, the application claims the benefit of an earlier-filed application that discloses the claimed subject matter, the priority date is the priority date of that earlier-filed application (i.e., the filing date of the provisional application or the priority date of the parent application). One can see now why the quality of a provisional application's disclosure matters: a provisional application that does not adequately disclose the invention later claimed in a nonprovisional application claiming the benefit of the provisional application does not establish a priority date that is earlier than the filing date of the nonprovisional application.

The prior art against which the patentability of a claimed invention is assessed includes the patents and printed publications in existence as of the priority date of the claimed invention. Therefore, an earlier priority date is beneficial because it reduces the universe of prior art against which the patentability of the purported invention is evaluated.

An inventor's own disclosures made before the priority date of a claimed invention may be prior art. In many countries, public disclosure of an invention prior to filing a patent application constitutes prior art to the later-filed application claiming that invention. Therefore, if an inventor presents a conference paper disclosing the invention before filing a patent application, that conference paper may bar a patent from issuing from the later-filed application. Although the United States gives inventors a one-year grace period during which to file a patent application after publicly disclosing an invention, the contours of the grace period changed in 2013, and some uncertainty surrounds the grace period's scope. For this reason, and because public disclosure may bar patentability in other countries, most patent attorneys recommend that inventors always file a patent application before publicly disclosing their inventions.

Patent Prosecution

The interaction between patent applicants and patent examiners is called prosecution. Prosecution is conducted in writing. If the examiner determines that the claimed invention does not meet all requirements for patentability, the examiner sends the applicant a document, called an office action in the United States, setting forth the reasons why the examiner contends that the claimed invention is unpatentable. Reasons examiners provide for rejecting claims may include that the claimed invention is disclosed in the prior art, the claimed invention would have been obvious to a person having ordinary skill in the art as of the priority date, the claims are directed to unpatentable subject matter, or the claims (e.g., in a continuation application) are not adequately supported by the written description and drawings.

In response to an office action, the applicant, generally represented by a licensed patent practitioner, may submit a written response in which the applicant may amend the claims to address the examiner's rejections (e.g., to distinguish the claimed invention from the prior art), or, if the applicant disagrees with the examiner's contentions, explain why, in the applicant's view, the examiner is incorrect. A patent office and applicant may have several rounds of communications before one of two possible outcomes results: either the patent office grants a patent, or the applicant abandons the application. In some cases, the process of prosecuting an application to either grant or abandonment may take years, particularly if the applicant and examiner cannot reach agreement, in which case the applicant may appeal the examiner's decision.

Conclusion

Patent systems are designed to reward those who invest the time and resources to invent and are willing to disclose their innovations to the public. Because a patent gives its owner the right to exclude others from taking advantage of the invention during the patent's term, only those innovations that are new, useful, and nonobvious relative to the prior art in existence as of their priority dates are worthy of a patent grant. This article provided an overview not only of the policies underpinning patent law and the rights accompanying a patent, but also the requirements and procedures to obtain a patent, which include preparation and prosecution of a patent application in a patent office.

References

[1] *Patent Office Manual of Patent Office Practice (Canada), 1998 Edition*, Dec. 2015 update, § 12.05, pp. 12-6–12-8.
[2] European Patent Convention, 15th ed., Sept. 2013, article 52(2), p. 110.
[3] United States Code, Title 35, § 161 (2012).
[4] United States Code, Title 35, § 171 (2012).
[5] United States Code, Title 35, § 101 (2012).
[6] United States Code, Title 35, § 271(a) (2012).
[7] TRIPS Agreement on Trade-Related Aspects of Intellectual Property Rights, article 28, Apr. 15, 1994.
[8] United States Code, Title 35, § 112(a) (2012).
[9] United States Code, Title 35, § 112(b) (2012).
[10] *Manual of Patent Examining Procedure (United States), 9th Ed.*, Rev. 07.2015, Nov. 2015, § 2141.03, pp. 2100-140–2100-141.

Biography

Krista S. Jacobsen [M] (krista@jacobseniplaw.com) is a patent attorney based in Campbell, California, whose solo practice includes expert witness services, patent litigation support, patent preparation and prosecution services, and intellectual property counseling services. She is also a Lecturer in Law at Santa Clara University School of Law. She received a B.S. degree in electrical engineering from the University of Denver in 1991, M.S. and Ph.D. degrees in electrical engineering from Stanford University in 1993 and 1996, respectively, and a J.D. degree from Santa Clara University School of Law in 2009. From 1996 to 2006 she worked as an engineer in the digital subscriber line (DSL) industry. At various times between 1994 and 2006, she was an active participant in meetings of IEEE 802.14, IEEE 802.3ah, ITU-T Study Group 15/Question 4, ETSI TM6, and ANSI T1E1.4. She is the co-editor of two books on DSL technology, a named inventor on eleven United States patents, and the author or co-author of multiple standards contributions, conference papers, magazine articles, technical journal papers, and law journal papers.

[19] See, e.g., 35 U.S.C. § 101 (in the United States, "[w]hoever invents or discovers any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof, may obtain a patent therefor… ."); European Patent Convention, Article 52(1) ("European patents shall be granted for any inventions, in all fields of technology, provided that they are new, involve an inventive step and are susceptible of industrial application."); TRIPS Article 27 ("[P]atents shall be available for any inventions, whether products or processes, in all fields of technology, provided that they are new, involve an inventive step and are capable of industrial application.").

Standards for Multi-Stream and Multi-Device Media Synchronization

Given the commercial interest in media synchronization and disadvantages of proprietary technologies, consumer-equipment manufacturers, broadcasters, and telecom and cable operators have started developing a new wave of international standards for media synchronization. This article provides an overview of recently published standards from the most relevant bodies: IETF, ETSI, MPEG, DVB, HbbTV and W3C.

M. Oskar van Deventer, Hans Stokking, Matt Hammond, Jean Le Feuvre, and Pablo Cesar

Communications Standards

Abstract

Media synchronization is getting renewed attention with ecosystems of connected devices enabling novel media consumption paradigms. Social TV, hybrid TV, and companion screens are examples that are enabling people to consume multiple media streams at multiple devices together. These novel use cases require media synchronization, as unfortunately there are substantial delay differences between the various delivery routes for television and streaming media. Broadcasters have started using proprietary solutions for over-the-top media synchronization, such as media fingerprinting or media watermarking technologies. Given the commercial interest in media synchronization and the disadvantages of proprietary technologies, consumer-equipment manufacturers, broadcasters, and telecom and cable operators have started developing a new wave of international standards for media synchronization. This article provides an overview of recently published standards from the most relevant bodies: IETF, ETSI, MPEG, DVB, HbbTV, and W3C.

Introduction

Media synchronization is relevant whenever two or more associated media streams are played out together. The classic example is synchronization of audio and video for a television broadcast in order to achieve lip synchronization. More recent examples are social television (TV), hybrid TV and companion screen. Social TV, a.k.a. "watching apart together", has multiple users watching the same TV broadcast while communicating with each other by voice, chat, or other social media. Hybrid TV converges multiple media streams from different channels (broadcast, Internet) into one single TV stream (e.g. broadcast video with subtitles or alternative audio received via the Internet). The companion screen provides user interaction or media on tablet devices associated with a television broadcast (e.g. a play-along quiz, alternative audio, or alternative camera views).

M. Oskar van Deventer and Hans Stokking are with the Netherlands Organisation for Applied Scientific Research TNO.

Matt Hammond is with BBC R&D.

Jean Le Feuvre is with Telecom ParisTech.

Pablo S. Cesar is with The National Research Institute for Mathematics and Computer Science CWI.

Requirements on synchronicity differ per use case. Social TV is the least demanding case. If there is no audio crosstalk, users won't notice delay differences of less than a second, and often they do not even notice a four-second difference [1]. Hybrid TV is the most strict case, since lip-sync requires audio and video to be synchronized within 40 milliseconds [2]. The companion-screen case would be between those two extremes, adding the challenge of achieving synchronization between two separate devices where communication latency must be compensated for.

Even the least demanding requirement cannot be met by today's media delivery technologies. There can be up to six seconds difference in delivery of a single broadcast channel in a single country via different providers [3]. Transcoding buffers are a major contribution to those delay differences. Transmission delays are also significant. For example a single satellite hop introduces over a quarter of a second delay due to the non-infinite speed of light. Internet delivery using content delivery networks (CDN) is by far the slowest delivery technology. It can easily take 30 seconds to perform all required delivery steps, from transcoding and segmenting to segment buffering at the media player client. A recent test showed a 72 seconds delay between a UK broadcaster's origination of a television channel and its delivery via Internet outside the UK [3].

Broadcasters have started using over-the-top media-synchronization technologies based on audio fingerprinting or audio watermarking for offering synchronized companion-screen content, as these technologies are relatively easy to deploy even in the absence of standards. However, they fail when the audio level is low or if there is background sound in the viewing environment, and considerable confusion may occur when a clip from a program is being reused in another program. Any system must make a compromise between factors such as recognition speed, robustness, and perceptibility of any changes to the audio or ability to discriminate across a volume of audio material [4]. Both fingerprinting and watermarking poorly handle user interactions, such as pause, seek, rewind, and fast forward. Another concern is cost. There is the cost for changing the workflow to get the watermark into the audio of a broadcast before the encoder. Licensing fees for commercial solutions typically scale with the number of channels, the amount of content on those services being watermarked and/or the amount of activity from client applications (e.g. searches of an audio fingerprinting database). Finally, the lack of standards may result in high vendor switching costs.

Media synchronization has re-emerged as an active field of standardization during the last few years [5], several of which have been published in 2014. This article provides a comprehensive, but approachable, overview of recently published standards for media synchronization from the most relevant standardization bodies: IETF, ETSI, MPEG, DVB, HbbTV, and W3C.

Table 1 provides an overview of the standards for media synchronization discussed in this article. The following sections will provide further information about each of the standards, detail-

Body	Standard(s)	Year	Purpose	Timeline(s)	Correlation	Wall clock	Coordination
IETF	RFC 7272 [6]	2014	Social TV	RTP timestamp	N/A	Any (e.g. NTP)	RTCP extension
ETSI	TS 183 063 [7]	2011	IPTV, Social TV	RTP timestamp	RTCP extension	Same as IETF	RTCP extension
MPEG	TEMI [8]	2014	Timeline	MPEG TEMI	N/A	NTP / PTP	N/A
DVB	TS 103 286-2 [9]	2014	Companion screen	Various	CSS-MRS protocol	CSS-WC protocol	CSS-TS and CSS-TE protocols
HbbTV	TS 102 796 v1.3.1 [10]	2015	Hybrid TV, companion screen	Same as DVB	API	Same as DVB	DVB CSS-TS, API
W3C	SMIL [11]	2008	Multimedia presentations	Interactive	Semantic	N/A	SMIL scheduler

Table 1. Standards for media synchronization.

ing their main features, followed by a discussion about the current status of the standardization effort, and offering some insights on future developments.

Synchronizing RTP Streams: IETF RFC 7272

The IETF has standardized the use of the Real Time Control Protocol (RTCP) for the generic synchronization of Real Time Protocol (RTP) streams at different devices, e.g. social TV. RFC 7272 [6], published in 2014, defines the interaction between a synchronization client (SC) and a media synchronization application server (MSAS). The SC reports synchronization status information to the MSAS. The MSAS receives status information from multiple receivers in a single synchronization group. It can thus calculate a common playout time for all receivers. Based on this, the MSAS sends synchronization settings instructions to the SCs. A common clock that is synchronized across all receivers is assumed. No particular synchronization method is mandated, although several are suggested, including the Network Time Protocol (NTP).

Synchronization status information is defined in an RTCP extended report block, and contains, among others, the following information:

- A synchronization group identifier that distinguishes different groups of SCs that synchronize RTP streams.
- The packet-received RTP timestamp, identifying the RTP packet for which the timing is reported.
- The NTP time that a packet arrived at the input of the device.
- Optionally, the time that the contents of the packet were presented to the user.

The synchronization settings instructions are sent in a similar RTCP extended report. The main difference is that it contains a 64-bit packet-presented timestamp, allowing accurate control of presentation times. This is, for example, required in use cases such as audio beam forming or video wall displays.

Further, RFC 7272 specifies a Session Description Protocol (SDP) parameter that enables RTP entities to advertise their capability in various session control protocols. Furthermore, IETF has specified SDP parameters for negotiation of clock synchronization capabilities in RFC 7273.

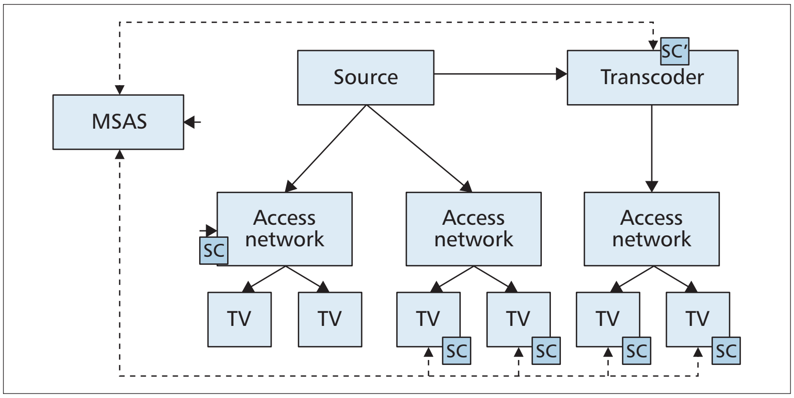

Figure 1. ETSI TS 183 063 functional entities for IDMS.

Social TV: ETSI TS 183 063

ETSI TISPAN specifies inter-device media synchronization as part of their Internet Protocol Television (IPTV) release 3 specifications, first published in 2011 [7]. The main use is social TV, in which various users can enjoy remote shared TV experiences, and which is spoiled by major path-delay differences between systems. The ETSI specification uses and expands upon RFC 7272.

A difference from RFC 7272 is that in the ETSI specification, the SC can be included in a TV, but it can also be part of an access network (see Fig. 1). This enables synchronization of large groups of legacy TVs by buffering at one node in the network. This requires that the delays from this node to those receivers are similar for all receivers, and that the SC compensates for those delays. Also, the MSAS can be a separate entity in the network, but it can also be contained in a TV. This enables a peer-to-peer style of synchronization.

To support this separation between function and element, ETSI specifies the session setup procedures. ETSI uses the Session Initiation Protocol (SIP) protocol for setting up media sessions, and uses the SDP protocol for specifying the synchronization part. In the SDP description exchanged between the TV and the IPTV system, the address of the MSAS is indicated. This can be a function in the network, or a function in another terminal.

An extension by ETSI to RFC 7272 is to support synchronization of the same content in different formats. In such a case, the RTP time-

The backbone of the media broadcasting industry is the MPEG-2 transport stream format, used throughout the broadcast chain. The MPEG-2 TS format specifies how video and audio or other media packets from one or several TV channels, called programs in MPEG-TS terminology, are scheduled within a continuous stream of bytes.

stamps reported by receivers are for different timelines. To support this situation, ETSI has defined a Synchronization Client *prime* (SC′), to be contained in a stream modifying entities such as a transcoder (see also Fig. 1). The SC' reports the mapping of timestamps in the incoming stream with those in the outgoing stream. This allows for conversion between those timestamps by the MSAS, thereby supporting SCs receiving original streams and SCs receiving modified streams that may be located in different networks.

External Media Items for Broadcast: MPEG TEMI

The backbone of the media broadcasting industry is the MPEG-2 transport stream (TS) format, used throughout the broadcast chain. The MPEG-2 TS format specifies how video and audio or other media packets from one or several TV channels, called programs in MPEG-TS terminology, are scheduled within a continuous stream of bytes. One challenge of the standard is to ensure that the target receiver behaves properly and consumes input data in time, avoiding buffer overflow, and to allow building systems with a limited amount of memory. This is achieved by a complex set of constraints on each bit in the byte flow, governed by the program clock reference (PCR) signal sent in the byte stream for each program. Decoding and presentation timestamps reference this clock. This forms an intrinsic timeline for the program, and is typically lost during transcoding or transmuxing of the source. To overcome this drawback, the MPEG group has defined an extension to MPEG-2 TS allowing for the carriage of extrinsic media clocks, along with other features, under the name 'TEMI' (Timing and External Media Information).

Earlier standards define such a mechanism using a dedicated elementary stream in the multiplex for the transport of an extrinsic clock. However, the method, while elegant, can be quite costly; at least one TS packet (188 bytes) has to be used to send the extrinsic clock of a program. When the clock has to be sent with each video frame for frame-accurate synchronization, this adds up to 75 kb/s for a 50 Hz video signal, more than the bandwidth used by some audio streams. TEMI addresses this issue by defining standard signaling, which is not affected by changes to PCR and that can be inserted before the start of a video or audio frame in the same TS packet. This design makes it possible to reduce the above signaling to 4 kb/s.

TEMI provides different ways of defining the extrinsic timeline:

- Sending a media time/media timescale pair, which can be compared with the presentation time of other media. Typically, this is compared with the composition time of an ISOBMFF track, or with the current time in an MPD period of an MPEG-DASH session.
- Sending an NTP or PTP (Precision Time Protocol) timestamp that matches timestamps associated with packets of the other media, for example interpolated NTP time of an RTP packet.
- Sending a time code of the media frame to be matched with a time code embedded in the other media, for example embedded as a track in ISOBMFF files or as an extension header in RTP packets.

TEMI also provides tools to signal the uniform resource location (URL) of one or several additional content items to be played synchronously with the broadcast, along with their MIME (Multi-Purpose Internet Mail Extensions) types. These items are assigned a timeline identifier, associated with each TEMI timing information. This allows sending URLs of associated services at a much lower frequency than timing information. Finally, TEMI provides a way of announcing when additional media content will become active by sending countdown signals for a given timeline identifier. TEMI content may also be marked as splicing points, indicating that the previous non-splicing content will resume at the end of the splice. This helps receivers to optimize their resources.

Figure 2 shows how TEMI may be used to signal and synchronize 3D or 4K enhancements (*URL#1*) and other additional contents such as subtitles or alternate audio (*URL#2*) to an existing broadcast signal, and signal upcoming splicing content (*URL#3*), typically for ad insertion purposes.

As a further continuation of TEMI activity, MPEG is investigating unified signaling of the different timelines defined in its various system layers (MPEG-4, ISO Base Media File Format, MPEG-DASH), along with specifying the missing tools enabling hybrid delivery of media content, such as signaling of coding dependencies between different containers.

Companion Devices: DVB CSS

The group Digital Video Broadcast – Companion Screens and streams (DVB-CSS) has developed a standard [9] to synchronize a media stream on a companion device with a media stream on a television set. The DVB-CSS architecture (Fig. 3) has one TV device and one or more companion screen applications (CSAs) running on companion screen devices that are connected via a home network, typically WiFi. Both TV and CSA independently receive media streams from the broadcaster (not shown). The presentation of the media streams is synchronized by using a set of new protocols and a new material resolution service.

A typical synchronization scenario is as follows. The user tunes their TV to a broadcast service. The TV receives the service, which includes a broadcast stream and metadata for media synchronization. The user pairs their companion device with the TV and starts a CSA. The TV provides the CSA with content identification and other information (CSS-CII protocol), which includes the service endpoints for the other protocols. The CSA queries the material resolution server (CSS-MRS protocol) and obtains material information that describes the structure of the broadcast (composition of materials and sub-materials such as programs, sections within programs, and adverts). It also describes the relationship between this structure and timelines. DVB-CSS supports several types of timelines,

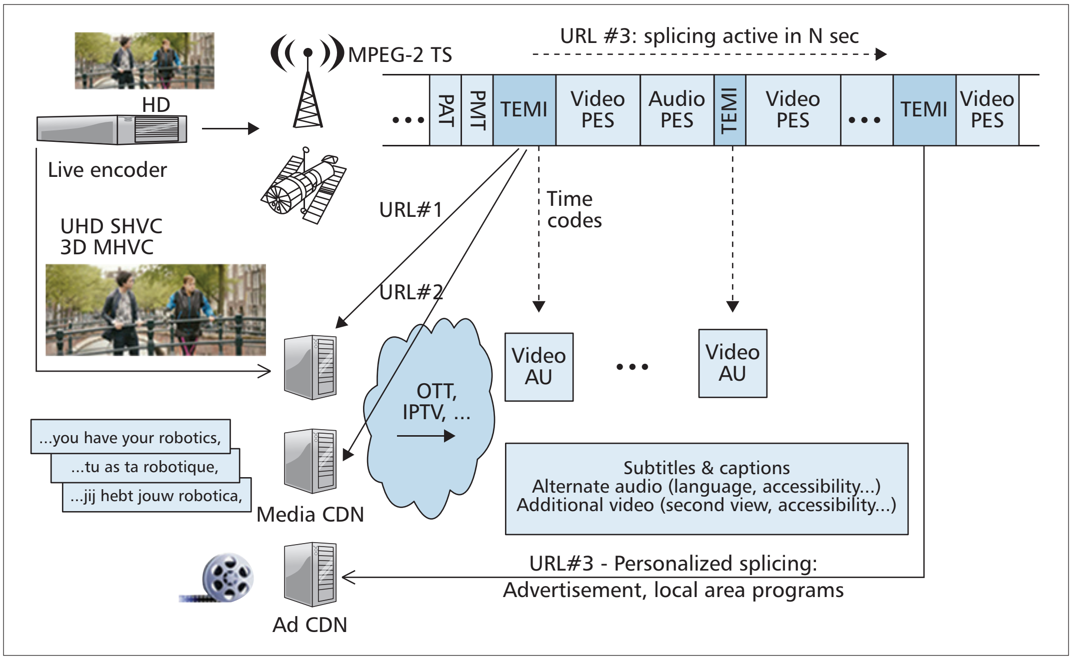

Figure 2. TEMI usage in hybrid broadcast.

HbbTV is an industry forum that specifies a HTML+JavaScript application programming interface (API) for browser-based applications on TVs, launched in 22 countries as of March 2015. The new HbbTV 2.0 specification [10] includes features for media synchronization, which are a profile of DVB CSS.

including MPEG transport stream presentation timestamp, ISOBMFF composition time and time relative to the start of a period in an MPEG DASH presentation. It also supports the use of TEMI as a timeline, but not as a means of specifying the external media stream to be played by the companion screen.

This combination of information from the TV and a server enables the CSA to determine which streams it should present and how its timeline correlates to that of the media being presented by the TV. However, the CSA manages its own behavior and is not directly controlled by the TV.

In parallel, the CSA synchronizes its wall clock with the TV (CSS-WC protocol). When the user starts a selected media stream on its companion screen, the CSA synchronizes the stream's timeline with the timeline of the broadcast stream on the TV (CSS-TS protocol described below). The CSA can also subscribe to trigger events (CSS-TE protocol) that are received by the TV from the broadcaster as part of the signaling within the broadcast stream.

The Wall Clock Synchronization protocol (CSS-WC) is a request-response UDP-based protocol that enables the client (CSA) to estimate a clock at a server (the TV) and measure and compensate for network round-trip delay. The protocol design is similar to the client/server mode of NTP but significantly simplified. Although many devices implement NTP to set their system-wide clocks, a CSA running on a device cannot always check if an NTP client process is functioning or query the accuracy of clock synchronization. Media synchronization also does not require the shared clock to be with reference to absolute real world time and can therefore avoid complexities such as leap-seconds. Frame accurate media synchronization requires accuracy of the order of milliseconds, and the chances of achieving this are improved if the protocol operates directly between the TV and CSA instead of via a hierarchy of intermediate servers on more distant network segments.

The Timeline Synchronization protocol (CSS-TS) is a websocket-based protocol that carries the timing information needed for coordination between the CSA and TV. Messages conveyed by this protocol describe the relationship between wall clock time and timeline position. This enables the CSA to accurately estimate the current TV timeline position despite possible network transmission delays. Timeline positions reported by the TV are expected to take account of any delays between the point at which it samples the timeline position in its media pipeline and the display of the media. Similarly, if a set-top box and an HDMI-connected display is used, then the STB is expected to make a best-effort to compensate for the play-out delay of the display. HDMI signaling may be used for this purpose.

Hybrid TV: HbbTV 2.0

HbbTV is an industry forum that specifies a HTML+JavaScript application programming interface (API) for browser-based applications on TVs, launched in 22 countries as of March 2015. The new HbbTV 2.0 specification [10] includes features for media synchronization, which are a profile of DVB CSS. In an HbbTV 2.0 TV, media synchronization is possible between the TV and a CSA or another HbbTV 2.0 TV acting in the role of a CSA. It is implemented as a profile of the DVB-CSS specification, and it is only activated when an interactive application running on the TV explicitly requests it. The DVB-defined CSS-CII, CSS-WC, and CSS-TS protocols are used in HbbTV 2.0, but HbbTV 2.0 TVs are not required to implement CSS-TE.

Media synchronization functionality between TV and CSA is available for most media types that the TV can be playing, including both broadcast and streamed broadband content. Required support for media synchronization between mul-

Synchronized multimedia integration language and nested context language are the most relevant examples of declarative and structured rich media formats. SMIL has been standardized by the World Wide Web Consortium, while NCL is the multimedia presentation standard for IPTV selected by ITU (ITU H.761).

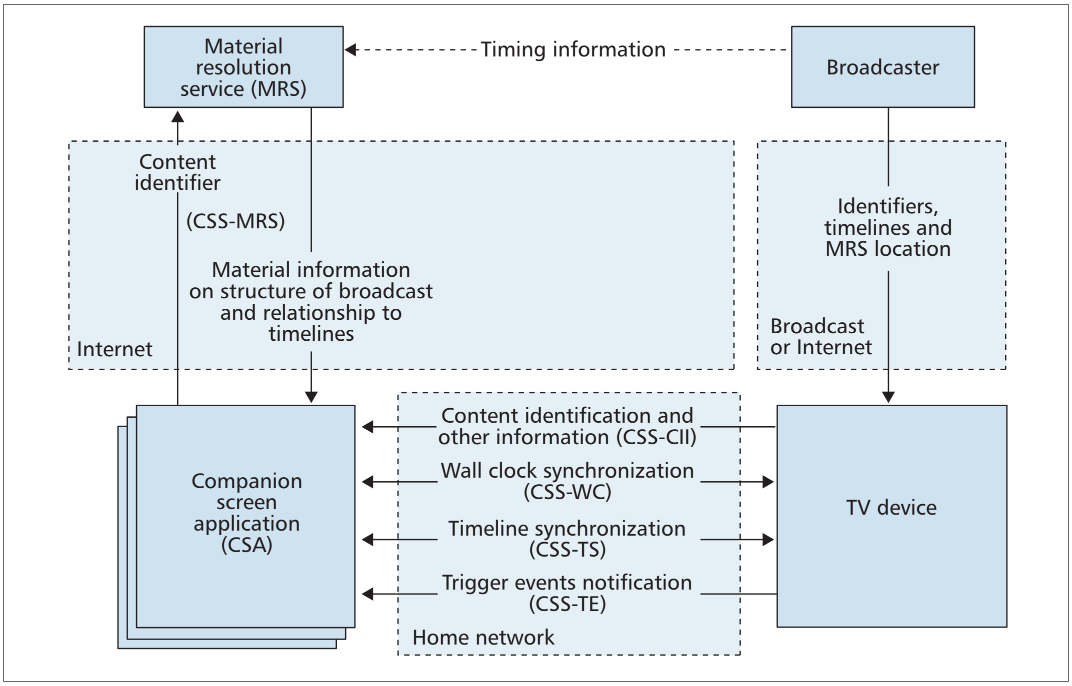

Figure 3. DVB-CSS architecture [9].

tiple streams within the TV is limited to combinations where one stream (possibly together with a synchronizing subtitle data stream) is delivered by broadcast and another stream by broadband.

HbbTV 2.0 specifies a single API that can be used for both single-TV multi-stream synchronization and inter-device synchronization. For the latter, an HbbTV terminal can act as both "master" and "slave," enabling streams on two TVs to be synchronized with each other. The API controls the life cycle of the MediaSynchoniser JavaScript object. This object is initialized by the API and populated with media objects, corresponding to to-be-synchronized media streams. The API also has methods to enable and disable the inter-device protocols explained above.

A media synchronization buffer is optional in HbbTV 2.0. Even without a buffer in the TV, media synchronization may be possible. The broadcaster can preload media streams for the CSA (or a slave HbbTV terminal) on a CDN. The broadcaster could editorially delay the broadcast stream, although this is not typically done for live streams. If any of the media streams is MPEG DASH (HbbTV only supports this type of standards-based adaptive streaming), then it is mandatory to buffer the stream on a CDN. If the TV has a media-sync buffer, then it will be at least 30 MByte large. This is sufficient to reliably buffer at least 10 seconds of encoded high definition television (HDTV) content.

Multimedia Presentations: W3C SMIL and ITU NCL

Synchronized multimedia integration language (SMIL) and nested context language (NCL) are the most relevant examples of declarative and structured rich media formats. SMIL has been standardized by the World Wide Web Consortium (W3C) [11], while NCL is the multimedia presentation standard for IPTV selected by ITU (ITU H.761) [12]. They are both XML-based integration formats, and as such they do not directly define media objects. Instead, they define the temporal and spatial relationships between the different media objects, enabling media synchronization of distributed objects across heterogeneous devices. They both sit on top of other low-level transmission and delivery standards, which are in charge of executing the low-level synchronization primitives.

The core part of these standards is the scheduler. The scheduler is in charge of constructing a time graph of the presentation, based on the duration of the media items and on the temporal synchronization between them. Based on the time graph, media items composing the presentation become active or inactive at specific moments in time. The scheduler is dynamic, allowing for the description of adaptable presentations based on events (from the user, from the network, from third-party entities like a broadcaster). They can be used for a variety of use cases from social television applications (including secondary screen support) to video conferencing services, to late-binding mashup videos [13].

NCL and SMIL have a strict separation between the document's content and structure, and it provides non-invasive control of presentation timing, linking, and layout. In NCL, authors can declaratively describe the temporal behavior of a multimedia presentation using connectors and links. SMIL acts as a container format in which spatial, temporal, linking, and interactive primitives can be used to position, schedule, and control a wide assortment of multimedia presentations. Both languages also allow for some form of procedural control. The biggest difference between the two languages is that while SMIL provides high-level constructs defining a restricted set of temporal relationships, NCL allows an author to create a set of custom relationships from a toolkit of language primitives as objects.

Both languages incorporate the recurrent aspects from a multimedia presentation [14]:

- Media items: defining what to render (video, images, text, and sometimes 3D objects). For selectivity purposes, the document model might also provide mechanisms for rendering one of multiple alternative assets.
- Style: defining how to render media, including multimedia styling options and digital effects such as zooming within an image.
- Spatial composition: defining where to render media in order to provide a meaningful and aesthetically attractive presentation.
- Temporal composition: defining when to render media including the start time and duration of media items, and also synchronization constraints between the items.
- User interaction: defining how to influence the presentation.

Conclusion

Current standardization efforts target specific use cases, such as social TV, hybrid broadcast/broadband services, and companion screens, which require media synchronization. This article provides an overview of them, highlighting the current industry push for such new services, both at the IP media stream level (IETF RTCP, ETSI TISPAN) and the MPEG-2 transport stream level (DVB CSS, MPEG TEMI). It also includes more fundamental standards (W3C SMIL and ITU NCL) that can serve as models for future more general synchronization primitives.

Successful standardization efforts are key for industrial partners. Vendors in HbbTV have already committed to implement at least the mandatory aspects of HbbTV 2.0 (including the profile of DVB-CSS) in their new TV products, with the expectation of seeing compliant products by 2017. Meanwhile broadcasters, including the BBC, are already exploring the services that media synchronization will enable [15].

References

[1] D. Geerts *et al.*, "Are We in Sync? Synchronization Requirements for Watching Online Video Together," *ACM Conf. Human Factors in Computing Systems, Vancouver (ACM-SIGCHI)*, May 2011.

[2] ITU-R BT.1359-1, "Relative Timing of Sound and Vision for Broadcasting," 1998.

[3] W.J. Kooij *et al.*, "Playout Delay of TV Signals: Measurement System Design, Validation and Results," *ACM Int'l. Conf. Interactive Experiences for TV and Online Video (ACM-TVX)*, June 2014, pp. 23–30.

[4] BBC R&D white paper 101, "EBU Tests of Commercial Audio Watermarking Systems," Dec. 2004.

[5] M. Montagud *et al.*, "Inter-Destination Multimedia Synchronization: Schemes, Use Cases and Standardization," *Multimedia Systems*, vol. 18, no. 6, 2012, pp. 459–82.

[6] IETF Inter-Destination Media Synchronization (IDMS) Using the RTP Control Protocol (RTCP), IETF RFC 7272, June 2014.

[7] ETSI IMS-based IPTV Stage 3 Specification, ETSI TS 183 063, Mar. 2011.

[8] MPEG TEMI, ISO/IEC 13818-1:201x/FDAM 2 Delivery of Timeline for External Data, Dec. 2014.

[9] DVB Companion Screens and Streams; Part 2: Content Identification and Media Synchronisation, DVB Document A167-2, July 2014, also published as ETSI TS 103 286-2.

[10] HbbTV Hybrid Broadcast Broadband TV version 2, Feb. 2015, available at hbbtv.org, to be published as ETSI TS 102 796 v1.3.1.

[11] W3C Synchronized Multimedia Integration Language, W3C SMIL 3.0, Dec. 2008.

[12] ITU-T H.761, "Nested Context Language (NCL) and Ginga-NCL for IPTV Services," 2009.

[13] R.L. Guimaraes *et al.*, "Creating Personalized Memories from Social Events: Community-Based Support for Multi-Camera Recordings of School Concerts," *ACM Multimedia Conf. (ACM MM 2011)*, Scottsdale (AZ), USA, Nov. 28–Dec. 1, 2011, pp. 303–12.

[14] J. Jansen, P. Cesar, and D. Bulterman, "A Model for Editing Operations on Active, Temporal Multimedia Documents," *ACM Symp. Document Engineering (ACM DocEng 2010)*, Manchester, UK, Sept. 21–24, 2010, pp. 87–96.

[15] S. Jolly and M. Evans, BBC R&D white paper 242 (Feb. 2013): "Improving the Experience of Media in the Connected Home with a New Approach to Inter-Device Communication."

Biographies

M. Oskar van Deventer (oskar.vandeventer@tno.nl) received M.Sc. and Ph.D. degrees in electrical engineering from the Eindhoven Technical University in 1987 and 1994, respectively. He became a Member (M) of IEEE in 2009. He has worked at KPN Research since 1987 and at TNO since 2003, where he is a senior scientist media networking. His current focus is on the realization of international standards for media networking, including DVB, HbbTV, ETSI, IETF and OIPF. He is the author of one book, and (co)author of more than 100 publications, 60 patent applications, and 600 standardization contributions.

Hans Stokking received the M.Sc. degree in systems engineering, policy analysis & management from the Delft Technical University in 1999. He has worked at KPN Research since 1999 and at TNO since 2003, where he is a senior scientist media networking. His work focuses on multi-domain networking and delivery of IPTV services, with a special research interest in media synchronization. He is the organizer of the Media Synchronization Workshop series and TPC member for various conferences.

Matt Hammond received an M.Eng. degree in computer systems and software engineering from the University of York in the United Kingdom in 2001. From 2001 to 2011 he worked as a research engineer in the Research and Development Department of the BBC in London in the United Kingdom, and as a senior research engineer since 2011. His current work focus is on companion device interaction with televisions and broadcast services. He is an active contributor to the DVB Companion Screens and Streams specification and to HbbTV in the areas of synchronization and other companion screen functionality.

Jean Le Feuvre received his M.Sc. degree in telecommunications in 1999 from TELECOM Bretagne. He is the co-founder of Avipix, an MPEG-4 consulting and software provider company. He has been involved in MPEG standardization since 2000. He joined TELECOM ParisTech in 2005 as associate professor in the Signal Processing and Image Department. He is an expert in multimedia delivery and interactive video systems. He is the project leader and maintainer of GPAC, a multimedia framework based on standards technologies (MPEG, IETF, W3C, and Web3D). He is the editor of several ISO/IEC standards, and the co-author of several books, more than 30 publications and patents applications, and 300 standardization contributions.

Pablo Cesar received his Ph.D. from Helsinki University of Technology in 2005. He leads the Distributed and Interactive Systems Group at Centrum Wiskunde & Informatica (CWI). He has (co)authored more than 70 articles (conference papers and journal articles) about multimedia systems and infrastructures, social media sharing, interactive media, multimedia content modeling, and user interaction. He is involved in standardization activities (e.g. W3C, MPEG, ITU) and has been active in a number of European projects. He is co-editor of the book *Social Interactive Television: Immersive Shared Experiences and Perspectives*, and has given tutorials about multimedia systems in prestigious conferences such as ACM Multimedia and the WWW conference.

Successful standardization efforts are key for industrial partners. Vendors in HbbTV have already committed to implement at least the mandatory aspects of HbbTV 2.0 (including the profile of DVB-CSS) in their new TV products, with the expectation of seeing compliant products by 2017.

IEEE Standard 802.19.1 for TV White Space Coexistence

Understanding the need to provide coexistence solutions for different cognitive radio systems operating in white space frequency bands, in December 2009 the IEEE 802 Executive Committee started project 802.19.1 to develop a standard for "TV White Space Coexistence Methods". The authors present the key concepts of the IEEE 802.19.1 coexistence system, its architecture, and protocol. They also provide an example of the IEEE 802.19.1 coexistence system implementation and operation.

Stanislav Filin, Tuncer Baykas, Hiroshi Harada, Fumihide Kojima, and Hiroyuki Yano

COMMUNICATIONS STANDARDS

Abstract

Understanding of the need to provide coexistence solutions for different cognitive radio systems operating in white space frequency bands, in December 2009 the IEEE 802 Executive Committee started project 802.19.1 to develop a standard for "TV White Space Coexistence Methods". This standard, published in June 2014, specifies radio technology independent methods for coexistence among dissimilar or independently operated TV band networks. In this article we present the key concepts of the IEEE 802.19.1 coexistence system, its architecture, and protocol. We also provide an example of the IEEE 802.19.1 coexistence system implementation and operation.

Introduction

In radio communications, spectrum is a very limited resource. Several measurements and studies have indicated that often many licensed frequency bands are under-utilized [1]. In particular, TV bands have been shown to have unused spectrum, and later measurements have only confirmed such a result [2]. This provides an opportunity for radio systems to identify and use temporally unused spectrum, thus improving spectrum utilization. To enable such an opportunity, two requirements must be satisfied:

1. Radio systems must have the capability to identify and use the temporally unused parts of the spectrum.
2. Radio regulations must be in place to allow radio systems to operate in the temporally unused parts of licensed frequency bands.

The first requirement can be satisfied by the current state of radio communication technology, in particular by using cognitive radio systems [3]. The second requirement is currently satisfied in some countries. For example, the FCC has published several documents allowing secondary user access to TV white space (WS) for fixed and portable devices [4, 5]. In the U.K., Ofcom has published several related documents [6, 7], and has recently approved license-exempt access to TV WS based on its framework [8]. In Japan, MIC has published rules for secondary operation in TV WS [9]. Related studies are ongoing in Singapore, CEPT, and ITU-R.

Frequency bands in which radio regulations allow cognitive radio systems to operate in temporally unused parts of these frequency bands are typically called "WS frequency bands." Here, the term "white space" refers to the temporally unused parts of the frequency bands. Radio systems to which these frequency bands are assigned are called "primary radio systems" or "primary users." Cognitive radio systems that operate in WS of these frequency bands are called "secondary radio systems" or "secondary users." Using U.S. radio regulations [4, 5] as an example, TV broadcast systems and wireless microphones are the primary users, while cognitive radio systems are the secondary users. In this example, WS frequency bands are limited to some channels in TV frequency bands.

In addition to regular functionality of radio systems operating in licensed bands (delivery of user traffic), cognitive radio systems operating in WS frequency bands shall have two additional functions: primary user protection, and coexistence with other secondary users. The first function is typically mandatory for regulatory compliance, since a cognitive radio system is operating as a secondary user. This means that it is allowed to operate only in temporally unused parts of WS frequency bands.

The second function, i.e., coexistence with other cognitive radio systems, is not required by radio regulations. The reason to have this function is as follows. WS is not exclusively assigned to one particular cognitive radio system. Any cognitive radio system that satisfies radio regulations for primary user protection can use WS. Consequently, more than one cognitive radio system can select the same WS for its operation. In such a case, several cognitive radio systems operating in the same WS may cause interference to each other, which may lead to performance degradation or even the inability to continue operation. The mechanisms that allow cognitive radio systems to avoid such interference situations are called "coexistence mechanisms."

The coexistence mechanisms can be put into one of two categories:

- Mechanisms of coexistence between similar cognitive radio systems.
- Mechanisms of coexistence between different cognitive radio systems.

The first category is also called "self-coexistence mechanisms." By similar cognitive radio systems we mean cognitive radio systems operated according to the same radio communication standard. Examples of radio communication standards for radio systems capable of operating in WS frequency bands are IEEE 802.22 [10] and IEEE 802.11af [11]. Typically, self-coexistence mechanisms are incorporated into a radio communication standard, and thus are used by cognitive radio systems operating according to this standard. A self-coexistence mechanism defined in the IEEE standard 802.22 is based on time scheduling. Base stations that need to coexist with each other share one or several superframes for transmission. Within the shared frames, each base station uses only specific frames for transmis-

Stanislav Filin, Fumihide Kojima, and Hiroyuki Yano are with National Institute of Information and Communications Technology.

Tuncer Baykas is with Istanbul Medipol University.

Hiroshi Harada is with Kyoto University.

 0163-6804/16/$25.00 © 2016 IEEE

sion. The IEEE 802.11af self-coexistence mechanism is based on the well known CSMA protocol.

Self-coexistence mechanisms provide various means to solve the coexistence problems for cognitive radio systems operating according to the same radio communication standard. However, such mechanisms cannot be used to solve the coexistence problem for different cognitive radio systems operating according to different radio communication standards. For example, an IEEE 802.11af station cannot join a frame schedule used by IEEE 802.22 base stations. Also, an IEEE 802.22 base station is not capable of supporting the IEEE 802.11af CSMA protocol.

Understanding the need to provide coexistence solutions for different cognitive radio systems operating in WS frequency bands, in December 2009 the IEEE 802 Executive Committee initiated project P802.19.1 to develop a standard for "TV White Space Coexistence Methods" [12]. This standard was published in June 2014. It specifies radio technology independent methods for coexistence among dissimilar or independently operated TV band networks [13], and it defines services and mechanisms to enable coexistence between different cognitive radio systems operating in the TV WS frequency bands.

We will use the term "WS radio system" to refer to a cognitive radio system operating in TV WS frequency bands in the remainder of the article. Also, we use the term "coexistence system" to refer to a management system that is based on the IEEE 802.19.1 standard and that is aimed at providing coexistence services between different WS radio systems operating as secondary users.

In this article we present an overview of the IEEE 802.19.1 standard. The article is organized as follows. We describe the IEEE 802.19.1 system architecture, and key elements of the IEEE 802.19.1 coexistence protocol. We give an example of the IEEE 802.19.1 coexistence system implementation. Finally, we conclude the article.

System Architecture

The IEEE 802.19.1 system architecture, as shown in Fig. 1, is designed to perform three key tasks required to solve coexistence problems between different WS radio systems:

- Discovery of WS radio systems that need to coexist with each other.
- Changing operating parameters of these WS radio systems to improve their performance.
- Providing a unified interface between different types of WS radio systems and a coexistence system.

Neighbor Discovery

The first task of a coexistence system is to discover WS radio systems that need to coexist with each other. To solve the first task, a logical entity called a coexistence discovery and information server (CDIS) is defined. Its key function is to support discovery of the neighbor WS radio systems. Two WS radio systems are neighbors if they are likely to cause one-way or mutual harmful interference to one another if they operate on the same frequency channel.

Neighbor relationships between WS radio systems depend on multiple parameters. The most obvious parameters are the locations of the WS radio systems and their radiated power, which in turn depends on transmission power, antenna height, antenna pattern, and propagation loss. The next set of parameters characterizes the receivers of a potential victim, for example, receiver sensitivity, noise floor, or minimum required SNR. Also, different radio access technologies may have different ranges of operating SINR values, and may have different levels of robustness to interference and to different types of interference. Finally, signal propagation is different in different frequency bands. As a result, two WS radio systems may be neighbors on one frequency channel and may be not neighbors on another channel.

There are many different ways to take into account all these factors. As a result, different neighbor discovery algorithms with different performance and complexity can be developed and implemented. Some examples of such algorithms are given in the IEEE standard 802.19.1.

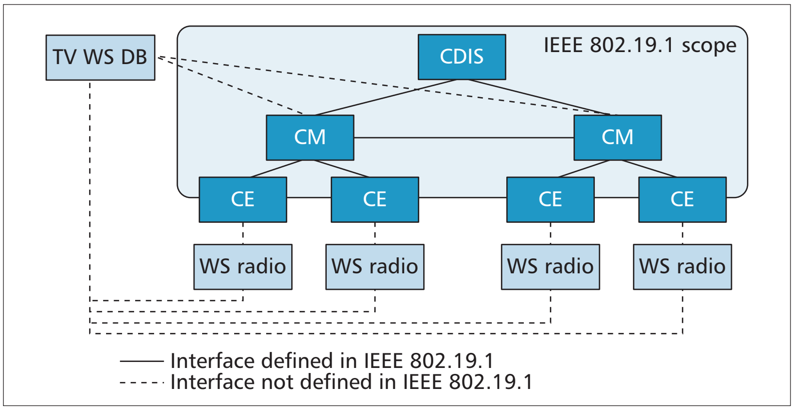

Figure 1. IEEE 802.19.1 system architecture.

Coexistence Management

The second task of a coexistence system is to continuously update operating parameters of WS radio systems in a way that improves their performance. The IEEE standard 802.19.1 provides two coexistence services to solve this task, namely, information service and management service. Within the information service, the coexistence system provides neighbor discovery information to a WS radio system, and the WS radio system autonomously updates its operating parameters. Within the management service, the coexistence system manages the operating parameters of a WS radio system. To provide the management service, a logical entity called a coexistence manager (CM) is defined.

A CM collects registration information from WS radio systems it serves, neighbor discovery information from a CDIS, and available frequency bands information from a TV white space database. The CM evaluates this information and makes coexistence decisions on operating parameters of the WS radio systems it serves. Then the CM requests corresponding reconfiguration of the WS radio systems it serves.

A CM can only control operating parameters of the WS radio systems that are subscribed to its management service and are served by this CM. Neighbor WS radio systems may be served by different CMs. In this case, a CM can negotiate with a neighbor CM to modify operating parameters of the neighbor WS radio systems served by the neighbor CM.

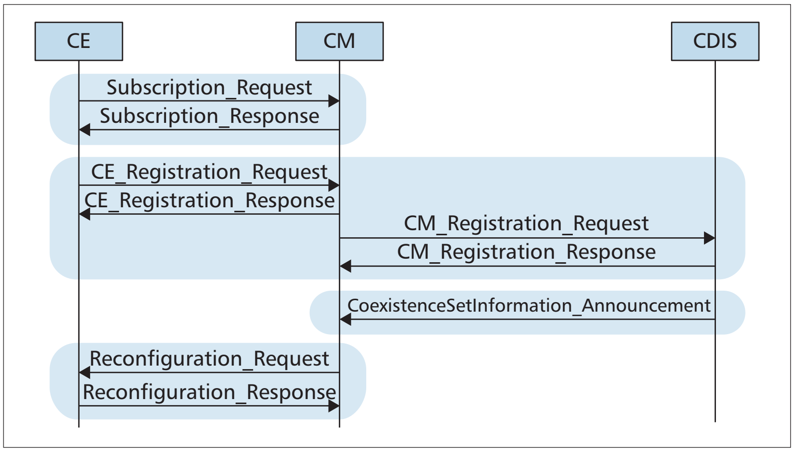

Figure 2. Basic procedures of the IEEE 802.19.1 coexistence protocol.

The IEEE 802.19.1 standard is designed to adapt only high-level operating parameters of a WS radio system. For example, a coexistence system can adapt such high-level parameters as the operating channel, the large-scale transmission schedule, and the transmission power. It does not aim to modify such low-level operating parameters as inter-frame resource allocation, modulation and coding scheme, and the portion of transmission power allocated to individual users.

Similar to neighbor discovery, there can be multiple algorithms for making coexistence decisions on a WS radio system's operating parameters or for performing negotiation between different CMs. Examples of such algorithms are provided in the IEEE 802.19.1 standard.

Unified Interface

Once a coexistence system is deployed and in operation, it is intended to serve various types of white space radio systems. Correspondingly, there is a need to have a unified interface between different types of WS radio systems and a coexistence system. This task is solved by defining a logical entity called a coexistence enabler (CE).

A CE provides a unified interface between a WS radio system and the IEEE 802.19.1 coexistence system. As shown in Fig. 1, the interface between a CE and a WS radio system is outside of the scope of the IEEE 802.19.1 standard. In fact, the service access point is defined in an abstract manner in the standard, while exact implementation is left up to manufacturers. Such an approach is very beneficial, because it does not require any changes in already published and future standards in order to use coexistence services provided by the IEEE 802.19.1 coexistence system. A CE will serve as a translator of WS radio system specific messages to messages exchanged between a CE and a CM.

Coexistence Protocol

The key elements of the IEEE 802.19.1 coexistence protocol include the following procedures:
- Subscription.
- Registration and registration update.
- Providing coexistence set information.
- Reconfiguration.
- Obtaining channel classification information.
- Sharing coexistence set information.
- Coexistence set element reconfiguration.

The subscription procedure is the first procedure that a WS radio system needs to perform in order to join the IEEE 802.19.1 coexistence system. This procedure is used to confirm the WS radio system subscription level, i.e., information service or management service. After the subscription, a WS radio system provides its registration information to the coexistence system using the registration procedure.

The registration information includes general information about the WS radio system such as:
- Network ID.
- Network technology (e.g., IEEE 802.11af, IEEE 802.22).
- Network type (e.g., fixed, mode 1).

The registration information also contains capabilities and operating parameters such as:
- Available frequency bands (from TV WS database).
- Operating frequency bands.

Also, the registration information contains such important elements as discovery information, which includes the following parameters:
- Location information.
- Transmission power related information (e.g., maximum transmission power).
- Receiver information (e.g., sensitivity, noise figure).

The discovery information is forwarded from the CM to the CDIS, where the discovery of potential neighbors is performed. Potential neighbors are WS radio systems that are likely to interfere with each other if they operate on the same channel.

The result of the discovery process is the coexistence set information. For each WS radio system a coexistence set is a set of its potential neighbors. This information is announced from CDIS to CM using the providing coexistence set information procedure.

Based on the information received from the CE and the information received from the CDIS, the CM decides on the best configuration of the WS radio systems that it serves. The IEEE 802.19.1 standard does not mandate specific coexistence decision algorithms to be used by the CM. After such decisions have been made, the CM generates reconfiguration requests and sends them to the WS radio systems that need to be reconfigured using the reconfiguration procedure. The four procedures mentioned above create a basic set of the coexistence protocols. They are illustrated in Fig. 2.

In a typical deployment there will be multiple CMs serving different WS radio systems. When WS radio systems served by different CMs are potential neighbors, there is a need to exchange information between CMs, and there may be a need to ask a neighbor CM to reconfigure its WS radio system.

The sharing coexistence set information procedure is used to exchange the minimum required information needed between different CMs. This procedure is a result of the tradeoff between enabling efficient coexistence decision making and privacy. Different CMs do not disclose private information about WS radio systems they serve, e.g., location information. However, they do provide information necessary for efficient decisions, such as operating channels of the WS radio systems.

In some cases, a CM cannot create an acceptable configuration by reconfiguring only the WS radio systems it serves, while such a configuration could be created by minor reconfiguration of WS radio systems served by a neighbor CM. This is due to the uneven usage of spectrum by different WS radio systems. To enable such reconfiguration, the coexistence set element reconfiguration procedure is used. A reconfiguration request from one CM to another CM is not mandatory, and will be accepted only if it does not result in the performance degradation of the operating WS radio systems.

Implementation Example

An implementation example of the IEEE 802.19.1 TV WS coexistence system is shown in Fig. 3. More details can be found in [14].

At the time of this experiment only IEEE 802.11af systems were available. However, they were implemented in two options: with 5 MHz bandwidth and with 10 MHz bandwidth. Both systems position central frequency in the middle of a TV channel. Correspondingly, a 5 MHz system needs one TV channel for operation, while a 10 MHz system needs three TV channels for operation. These systems cannot leverage the CSMA self-coexistence mechanism of IEEE 802.11af due to their different bandwidths. The experiment shows how IEEE 802.19.1 coexistence system can be used for their coexistence.

The TV WS database has been implemented in a server that is located remotely and can be accessed via the Internet. CDIS and CMs are implemented in general purpose PCs that are located locally. One CDIS and two CMs have been deployed.

Several IEEE 802.11af networks have been deployed. Each network has one access point (AP) and one station (STA). Each AP is coupled with a CE. Three IEEE 802.11af 5 MHz networks are denoted as WS radio systems 1, 2, and 3. They are served by CM 1. Two 10 MHz 802.11af networks are denoted as WS radio systems 4 and 5. They are served by CM 2. All five networks are subscribed to a management service. Also, all five networks are located in such a way that they are all neighbors to each other.

Available channels have been assigned as follows. WS radio system 1 has two available channels: 728–734 MHz and 746–752 MHz. WS radio system 2 has two available channels: 734–740 MHz and 746–752 MHz. WS radio system 3 has one available channel: 746–752 MHz. WS radio systems 4 and 5 have three available channels: 728–734 MHz, 734–740 MHz, and 740–746 MHz.

In total, five WS radio systems have four available channels. Each of three 5-MHz WS radio systems needs one channel for its operation, while each of two 10-MHz WS radio systems needs three channels for its operation (because the 10-MHz WS radio system central frequency is positioned in the middle of a TV channel).

According to the evaluation scenario, WS radio systems 1, 2, 3, 4, and 5 join the IEEE 802.19.1 coexistence system and start operation one by one. When a WS radio system comprised of an access point and a station starts operation, the operating frequency is selected randomly from the available channels.

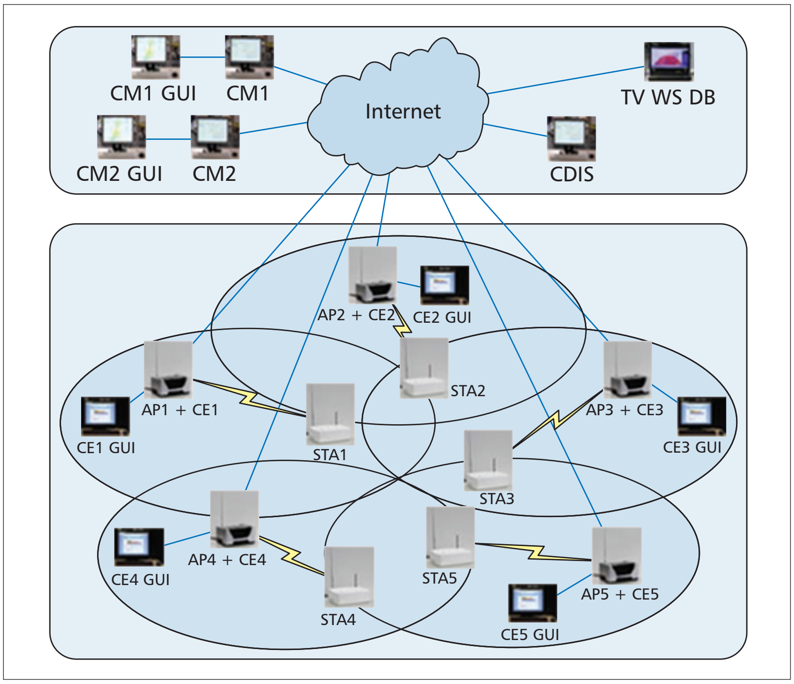

Figure 3. Implementation example of the IEEE 802.19.1 TV WS coexistence system.

Figure 4 shows the resulting configuration after WS radio systems 1, 2, and 3 have started operation, as viewed in a CM GUI. Red bold arrows between WS radio systems show neighbor relationships for a pair of WS radio systems. According to Fig. 4, all WS radio systems are neighbors. Colored circles show coverage areas. Different colors correspond to different operating frequencies, with the legend shown in the right upper part of the CM GUI.

WS radio system 1 has selected the 728–734 MHz channel for operation. WS radio system 2 has selected the 734–740 MHz channel for operation. WS radio system 3 has selected the 746–752 MHz channel for operation.

After WS radio systems 1, 2, and 3 have selected their operating channels, there are no available channels for WS radio systems 4 and 5. Recall that WS radio systems 4 and 5 use 10 MHz of bandwidth and need three consecutive channels for their operation.

To resolve this situation, CM 2 obtains information about WS radio systems 1, 2, and 3 from CM 1, and request CM 1 to reconfigure WS radio systems 1 and 2 to allow WS radio systems 4 and 5 to operate. Figure 5 shows the resulting configuration after all WS radio systems have started operation, as viewed in a CM GUI. Note that for 10 MHz systems, coverage areas are split into three sectors to illustrate three TV channels used by these systems.

Now WS radio systems 1, 2, and 3 use the 746–752 MHz channel for their operation, and WS radio systems 4 and 5 use the 728–734 MHz, 734–740 MHz, and 740–746 MHz channels for their operation. Fig. 5 shows that all five WS radio systems are operating. WS radio systems 1, 2, and 3 are of the same type and can share the same frequency bands using the self-coexistence mechanisms of the IEEE 802.11af. Similarly, WS radio systems 4 and 5 are of the same type (but different from white space radio systems 1, 2, and 3) and can share the same frequency bands using self-coexistence mechanisms of the IEEE 802.11af.

In this evaluation scenario, the IEEE 802.19.1 coexistence system creates such a configuration of the operating frequencies of the WS radio systems that the number of the operating systems is maximized. In particular, the number of operating WS radio systems has increased from three 5-MHz systems to five radio systems comprising three 5-MHz systems and two 10-MHz systems.

Conclusions

The IEEE 802.19.1-2014 standard defines radio transmission technology independent methods for coexistence of different or independently operated cognitive radio systems in TV WS frequency bands. This article presented the key concepts of the IEEE 802.19.1 coexistence system, its architecture, and protocol. Also, an example of an IEEE 802.19.1 coexistence system implementation and operation was shown.

When development of the IEEE 802.19.1 standard began, only a few WS radio system standards were available, including IEEE 802.22 and IEEE 802.11af. Currently, other standardization bodies are developing radio systems for operation in TV WS. For example, the IEEE 802.15.4m standard has been published [15]. Also, there are trials of WS radio systems based on proprietary technologies. The IEEE 802.19.1 coexistence system is expected to become an enabler for the seamless operation of these radio systems.

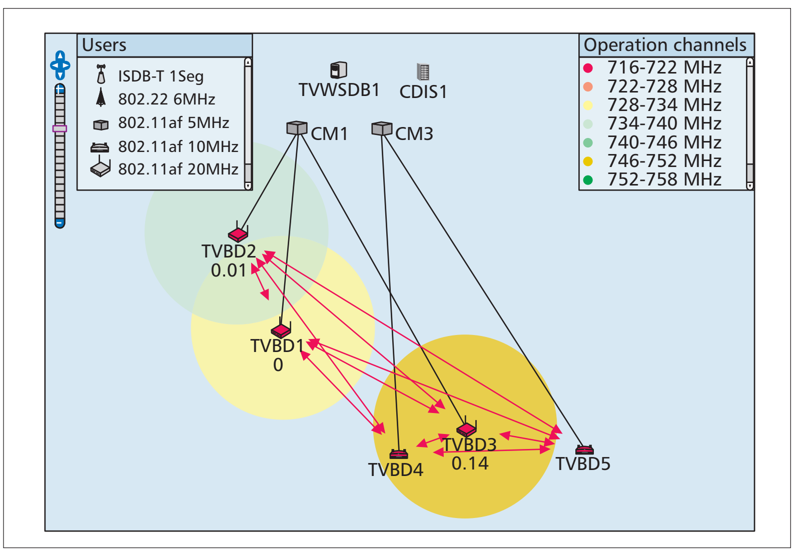

Figure 4. Configuration of operating channels after WS radio systems 1, 2, and 3 have started operation.

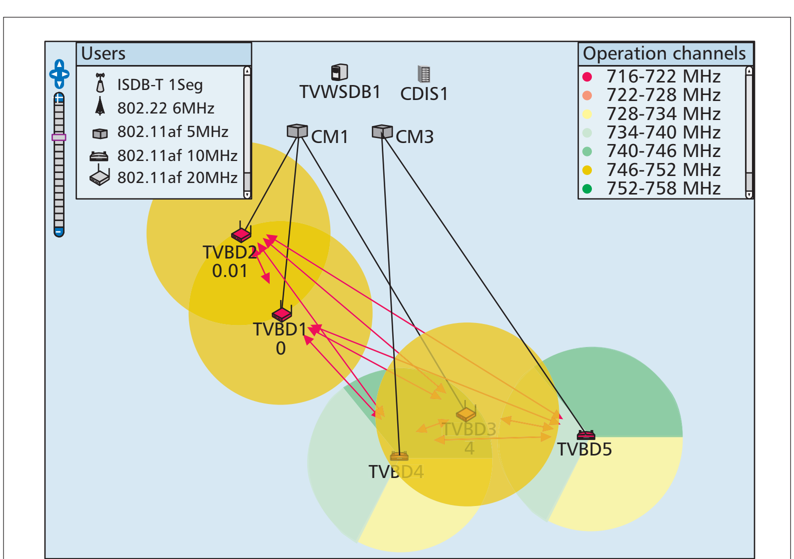

Figure 5. Configuration of operating frequencies after all WS radio systems have started operation.

References

[1] G. Staple and K. Werbach, "The End of Spectrum Scarcity," *IEEE Spectrum*, vol. 41, no. 3, 2004, pp. 48–52.

[2] J. van de Beek *et al.*, "TV White Space in Europe," *IEEE Trans. Mobile Computing*, vol. 11, no. 2, 2012, pp. 178–88.

[3] S. Filin *et al.*, "International Standardization of Cognitive Radio Systems," *IEEE Commun. Mag.*, vol. 49, no. 3, 2011, pp. 82–89.

[4] FCC Second Memorandum, Opinion and Order, "In the Matter of Unlicensed Operation in the TV Broadcast Bands: Additional Spectrum for Unlicensed Devices Below 900 MHz and in the 3 GHz Band," Sept. 2010.

[5] FCC third Memorandum, Opinion and Order, "In the Matter of Unlicensed Operation in the TV Broadcast Bands: Additional Spectrum for Unlicensed Devices Below 900 MHz and in the 3 GHz Band," Apr. 2012.

[6] Ofcom, "TV White Spaces — A Consultation on White Space Device Requirements," Nov. 2012.

[7] Ofcom, "TV White Spaces — Approach to Coexistence," Sept. 2013.

[8] Ofcom, "Implementing TV White Spaces," Feb. 2015.

[9] IEEE 802.11-12/677r0, H. Harada, "Status Report on Usage of TV White Space in Japan," May 2012.

[10] IEEE Standard 802.22-2011, "Cognitive Wireless RAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Policies and Procedures for Operation in the TV Bands," July 2011.

[11] IEEE Standard 802.11af-2013, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 5: Television White Spaces (TVWS) Operation," Dec. 2013.

[12] T. Baykas *et al.*, "Developing a Standard for TV White Space Coexistence: Technical Challenges and Solution Approaches," *IEEE Wireless Commun.*, vol. 19, no. 1, 2012, pp. 10-22.

[13] IEEE Standard 802.19.1-2014, "TV White Space Coexistence Methods," May 2014.

[14] IEEE 802.19-12/138r0, S. Filin and H. Harada, "Feasibility Study of the IEEE 802.19.1 TVWS Coexistence Protocol," July 2012.

[15] IEEE Standard 802.15.4m-2014, "Low-Rate Wireless Personal Area Networks (LR-WPANs) Amendment 6: TV White Space between 54 MHz and 862 MHz Physical Layer," Mar. 2014.

Biographies

Stanislav Filin [SM] (sfilin@nict.go.jp) works as a senior researcher in NICT. He is chair of the IEEE 1900.7 WG on WS Radio. He served as a chair and as a technical editor of the IEEE 802.19.1 TG on WS Coexistence. In 2014 he received an IEEE SA SB award for contribution to the development of the IEEE standard 802.19.1.

Tuncer Baykas [SM] (tbaykas@ieee.org) works as an assistant professor at Istanbul Medipol University. He was the first chair of IEEE 802.19.1 TG. He served as co-editor and secretary for 802.15 TG3c and contributed to many TVWS related projects, including 802.22, 802.11af and 1900.7. His work is supported by the Tubitak Bidep 2232 program.

Hiroshi Harada [M] (harada@ieee.org) is a professor in the Digital Communication Laboratory, Department of Communications and Computer Engineering, Graduate School of Informatics, Kyoto University. He is also executive research director of NICT. He joined the Communications Research Laboratory, Ministry of Posts and Communications, in 1995 (currently NICT). Since 1995 he has researched SDR, cognitive radio, and broadband wireless access systems and wireless smart utility networks. He has served on the Board of Directors of the Wi-SUN alliance and DSA alliance, as a chair of the IEEE DySPAN Standards Committee, and as a vice chair of IEEE 802.15.4g, 802.15.4m, and 1900.4.

Fumihide Kojima [M] (f-kojima@nict.go.jp) received a D.E. degree in electrical communications engineering from Osaka University, Japan in 1999. Currently he is Director of the Smart Wireless Laboratory, Wireless Network Research Institute of the National Institute of Information and Communications Technology (NICT). His current research interests include intelligent MAC protocols for smart utility networks.

Hiroyuki Yano (yano@nict.go.jp) is Director General of the Wireless Network Research Institute at the National Institute of Information and Communications Technology (NICT) in Japan. Dr. Yano worked for Human Communication, Human Interface, Wireless Technologies and Strategic Planning of R&D in NICT. His present research interest is in the area of wireless network systems, including both terrestrial and satellite communications.

IEEE 2016 GLOBECOM

GLOBAL COMMUNICATIONS CONFERENCE, EXHIBITION & INDUSTRY FORUM

IEEE COMMUNICATIONS SOCIETY

Freedom Through Communications

4-8 December 2016 // Washington, DC, USA

The 2016 IEEE Global Communications Conference (GLOBECOM) will feature a comprehensive technical program including 13 symposia, tutorials, workshops and an industrial program featuring prominent keynote speakers, technology and industry forums and vendor exhibits.

CALL FOR PAPERS

WASHINGTON, DC

#IEEEGC16

– TECHNICAL SYMPOSIA –

We invite you to submit original technical papers in the following areas:

- **Ad Hoc and Sensor Networking**
- **Communication & Information System Security**
- **Cognitive Radio and Networks**
- **Communication QoS, Reliability and Modeling**
- **Communication Theory**
- **Communications Software, Services and Multimedia Applications**
- **Green Communications**
- **Mobile and Wireless Networking**
- **Next-Generation Networking**
- **Optical Networks and Systems**
- **Signal Processing for Communications**
- **Wireless Communications**

Selected Areas in Communications:

- **Access Networks & Systems**
- **Big Data**
- **Cloud Networks**
- **Data Storage**
- **e-Health**
- **Internet of Things**
- **Molecular, Biological and Multi-scale Communications**
- **Power Line Communications**
- **Satellite & Space Communications**
- **Social Networks**

Please address questions regarding the Technical Symposia to TPC Chair Sherman Shen (sshen@uwaterloo.ca) and/or TPC Vice Chairs Maite Brandt-Pearce (mb9q@eservices.virginia.edu) and Sennur Ulukus (ulukus@umd.edu).

Accepted and presented technical and workshop papers will be published in the IEEE GLOBECOM 2016 Conference Proceedings and submitted to IEEE Xplore®. See the website for author requirements of accepted authors.
Full details of submission procedures are available at www.ieee-globecom.org.

IMPORTANT DATES:

Symposia Papers - **1 April 2016** // Tutorial Proposals - **15 March 2016**

– COMMITTEE –

General Chair
Bijan Jabbari, George Mason University

TPC Chair
Sherman Shen, University of Waterloo

TPC Vice Chairs
Maite Brandt-Pearce, University of Virginia
Sennur Ulukus, University of Maryland

Tutorial Co-Chairs
Brian Mark, George Mason University
Yahong Rosa Zheng, Missouri University of Science and Technology

Workshop Co-Chairs
Shuguang Cui, Texas A&M
Stefano Secci, University of Paris 6

Executive Chair
Paul Cotae, University of the District of Columbia

General Vice Chair
Jerry Gibbon, US Department of Commerce

GIMS Advisor
Chi-Ming Chen, AT&T

GITC Advisor
Peter Rost, Nokia Networks

- www.ieee-globecom.org -

3GPP Device-to-Device Communications for Beyond 4G Cellular Networks

Avoiding any signal relay via BSs, direct data exchanges among users known as device-to-device (D2D) communications have consequently emerged as the most crucial innovation and the mandatory technology for a public safety network. However, due to the tight schedule for standardization, a special form of "broadcasting" communications is adopted by D2D communications in Release 12.

Shao-Yu Lien, Chun-Che Chien, Fan-Min Tseng, and Tien-Chen Ho

Abstract

Since 2012, 3GPP Release 12 (the beyond fourth generation (B4G) network) has been adopted as the next generation public safety network to be deployed in 2015. In addition to conventional high data rate transmissions in LTE-A, a public safety network is required to further support communications for users out of base station (BS) coverage. Avoiding any signal relay via BSs, direct data exchanges among users, known as device-to-device (D2D) communications, have consequently emerged as the most crucial innovation and the mandatory technology for a public safety network. However, due to the tight schedule for standardization, a special form of "broadcasting" communications has been adopted by D2D communications in Release 12. To fully comprehend such an unprecedented technology in cellular networks, this article provides a comprehensive overview of D2D operations in Release 12. From the system architecture to the radio interface, the provided insights boost the knowledge to practice D2D communications over cellular networks.

COMMUNICATIONS STANDARDS

Introduction

Cellular networks have been deployed for several decades to provide ubiquitous wireless and mobile services to individuals. In past decades, communications among millions of mobile devices has been provided solely relying on connecting to widely deployed BSs. Although this state-of-the-art paradigm provides strong mobility management to enable seamless wireless services, it may not be efficient in terms of spectral utilization, energy consumption, and transmission latency. This inefficiency becomes particularly severe when data exchanges occur between two mobile devices located close to each other. In recent studies, it has been shown that direct data exchanges among two nearby mobile devices (without data relay via BSs) may provide better performance in terms of spectrum, energy, and transmission latency [1–4]. Consequently, these technical merits have driven the development of D2D communications [5].

Technologies for D2D communications have been studied for several years, and successful examples include Infrared Data Association (IrDA), Bluetooth, and WiFi Direct. Operating on high frequency and unlicensed bands, existing D2D technologies enjoy a wider bandwidth and a lower cost of deployment. However, they also suffer from low spectrum efficiency, low power efficiency, short communication distances, and vulnerability to interference. Therefore, developing D2D technology compatible with widely deployed cellular networks turns out to be an effective solution to largely enhance the capability of D2D communications. With potential technical merits, however, D2D communications in cellular networks did not attract much attention in standardization work. This stagnation persisted until 2012, when the Federal Communications Commission (FCC) of the United States endorsed 3GPP Release 12 (also known as the B4G network) as the next generation nationwide public safety network. When public safety events (such as rescue, conflagration, riots, accidents, etc.) occur, the LTE/LTE-A shall be the first network (FirstNet) responsible for providing communication services [6]. In addition to the conventional communication requirements of high data rates, low latency, priority control, and reliability, a public safety network must further support urgent communications when BSs are paralyzed by natural disasters or malicious attacks. For this additional requirement, D2D and group communications have been included as two key features in 3GPP Release 12 [7–10]. However, because of the tight schedule in the standards development, D2D communications are designed to be a special mode of "broadcasting" communications to simplify the design flow. In other words, for public safety, D2D communications are a sort of "open-loop" communications in Layer 1 and "groupcast" (one-to-many) communications in Layer 2. Specifically, a D2D receiver does not provide any feedback message (including channel state information, CSI, and acknowledgements) to a D2D transmitter.

The evolution of cellular network radio access solutions, from Global System for Mobility (GSM), Universal Mobile Telecommunications System (UMTS), LTE, to LTE-A, can be divided into two categories of communications: "closed-loop" and "open-loop". "Closed-loop" implies that feedback messages are provided from a receiver to a transmitter. For example, the UMTS system adopts code division multiple access (CDMA), by which a receiver periodically informs the transmitter about the received power level. The transmitter is thus able to adjust transmission power (such an inner-loop power control is performed 1500 times per second). On the other hand, LTE/LTE-A systems adopt orthogonal frequency division multiple access (OFDMA), by which a receiver continuously estimates the present CSI, and informs the transmitter about the optimum modulation and coding scheme (MCS) for the subsequent data transmission. These closed-loop operations boost Layer 1 and Layer 2 to adapt to dynamic channel variation, which is the key enabler to enhance

Shao-Yu Lien is with National Formosa University.

Chun-Che Chien is with the Institute for Information Industry.

Fan-Min Tseng is with the HTC Corporation.

Tien-Chen Ho is with Overseas Chinese University.

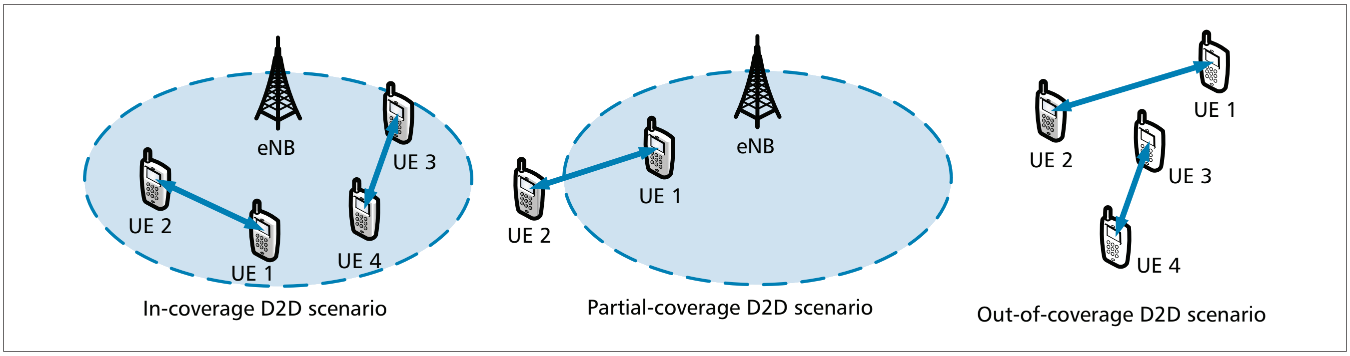

Figure 1. Release 12 D2D communications are divided into three scenarios: a) in-coverage; b) out-of-coverage; and c) partial-coverage.

transmission data rates. On the other hand, GSM and Multimedia Broadcast Multicast Services (MBMS) in 3GPP Release 6 to Release 8 adopt open-loop communications, while only low data rate voice and unreliable multimedia services are supported. Although open-loop may support multimedia transmissions for public safety, the potential and performance to support high data rate, low latency, and reliable communications still remain open for further investigations. Release 12 D2D communications thus open a new paradigm, which may fundamentally inspire and change future designs of cellular networks.

To fully comprehend the practice of such a novel technology, in this article we provide a comprehensive survey of the standardization of 3GPP Release 12 D2D communications. This survey begins top-down from the system architecture and resource/mobility management in the air interface to physical layer procedures, to provide key principles and the design philosophy of Release 12 D2D communications.

System Architecture of D2D Communications

Scenarios and Reference Architectures of D2D Communications

In addition to public safety, the use cases of D2D communications in the future can be in general commercial/social applications, traffic offloading from networks, and integration of existing infrastructure services. To potentially support all these purposes, a general concept of proximity services (ProSe) is used to construct the system architecture of D2D communications. However, Release 12 ProSe only emphasizes public safety communications, which will support urgent D2D communications when:
- All BSs are available.
- A part of BSs are available, while other BSs are paralyzed.
- All BSs are paralyzed.

To support these three cases, Release 12 D2D communications can be divided into the following three scenarios, as shown in Fig. 1. In LTE-A, a BS is referred as an eNB, and we also align with this terminology in this article.
- **In-coverage.** This scenario indicates that the considered user equipment (UE) is within coverage of the eNB(s).
- **Out-of-coverage.** This scenario indicates that the considered UE is out of the eNB's coverage.
- **Partial-coverage.** This scenario indicates that some UEs are within coverage, while other UEs are out of the eNB's coverage.

To support the above scenarios, the reference architecture of ProSe is illustrated in Fig. 2a. In addition to the conventional Uu interface between an eNB and a UE, and the S1 interface between an eNB and the evolved packet core (EPC), the following interfaces are included.

PC5: PC5 is the interface for D2D broadcasting communications. The details of PC5 will be discussed later.

PC1: To provide diverse wireless services based on D2D communications, various ProSe applications (ProSe APPs) can be installed in a UE. These ProSe APPs can exchange data with the remote ProSe APP server via PC1.

PC2: The ProSe functionalities supported by the EPC are referred as ProSe Function. PC2 is the interface defined between the ProSe APP server and ProSe Function. This interface can be used for ProSe Function to update application data for its ProSe database.

PC3: To identify the existence of other UEs, a UE may proceed to D2D discovery. Thus, PC3 is the interface between a UE and ProSe Function used for the configurations of D2D discovery and communications. If two UEs register to different public land mobile networks (PLMNs) or even belong to different radio access networks (RANs) such as LTE-A and WiFi, D2D discovery may rely on ProSe Function in the EPC via PC3.

PC4: As ProSe Function defines the ProSe functionalities supported by the EPC, PC4 specifies the interaction between ProSe Function and the EPC. It can be used to authorize ProSe services for session/mobility management.

PC6: PC6 is the interface between multiple ProSe Functions in different PLMNs. It can be used for D2D discovery between UEs subscribed to different PLMNs.

PC7: PC7 is the interface between ProSe Functions in the visiting PLMN (VPLMN) and the home PLMN (HPLMN). It can be used for HPLMN to handle the ProSe service authorization.

PC8: PC8 is the interface between a roaming UE and HPLMN ProSe Function. It can be used for ProSe Function in the HPLMN to configure D2D communications of UEs.

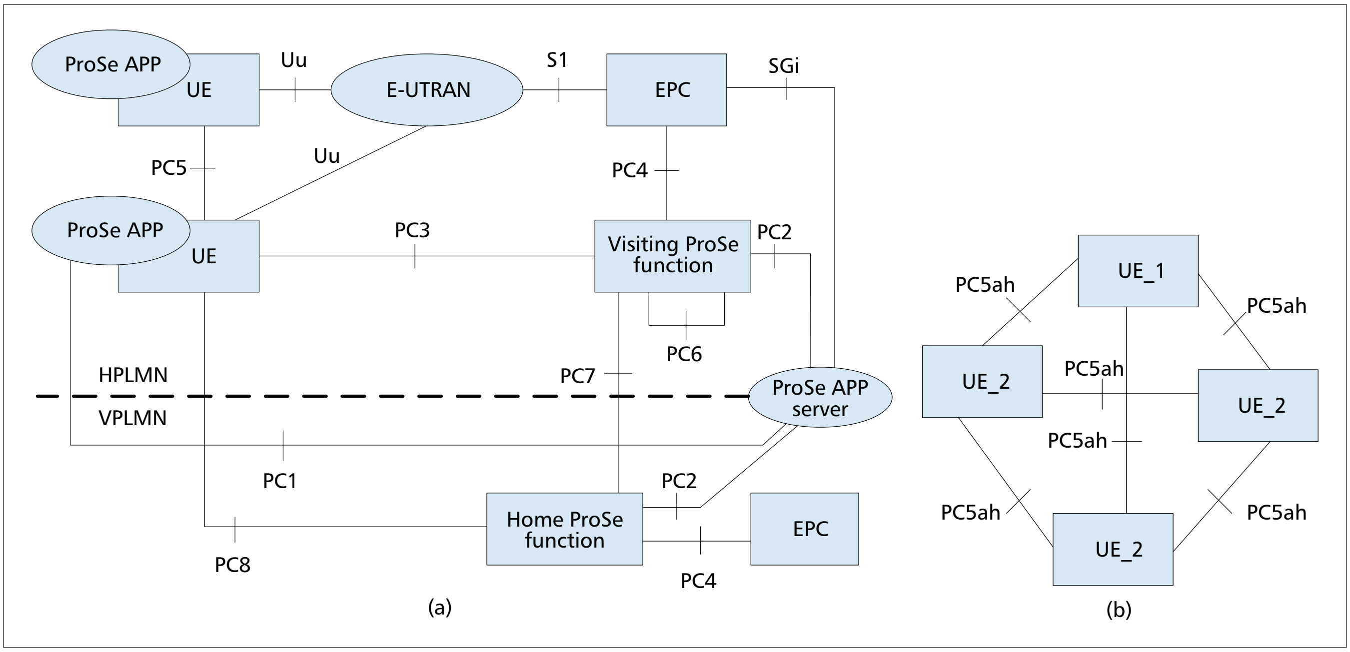

Figure 2. a) Reference architecture of ProSe with roaming; b) reference architecture of ProSe with one-to-many communications in ad hoc mode.

SGi: SGi is the interface used for APP data and the application level control information exchanges.

For UEs out of eNBs' coverage, each UE is able to exchange data directly with other UEs if signal strength is acceptable, thus forming an ad hoc network among UEs, as shown in Fig. 2b. Instead of PC5 for all radio layers, PC5ah is a lower-layer interface (only Layer 1 and Layer 2) for the ad hoc D2D network.

D2D Discovery

A complete D2D discovery procedure involves operations in different layers, and the main steps are illustrated in Fig. 3 and Fig. 4. Suppose that a user (called Jasmin) owning UE-A wishes to enjoy D2D wireless services (as shown in Fig. 3). The APP of the UE-A sends a request to the third party APP server to obtain the friend list of Jasmin. In the application layer, each friend of Jasmin is associated with a friend-identity (Friend-ID), which may be an email address, a phone number, or a nickname in the APP. After obtaining the friend list, the APP of UE-A then sends a request via 3GPP layers to the ProSe server to obtain an expression code for each Friend-ID. Through the prior negotiation, the ProSe server may already have the mapping parameters between each Friend-ID and the corresponding expression code. Thus, the APP is able to obtain the expression code via the reply from the ProSe server. Next, the APP of the UE-A is able to perform two modes of discovery to either announce its existence or find a particular friend. For these two discovery modes, there are two corresponding discovery messages:

- Mode A ("I am here"). This message contains the expression code of Jasmin.
- Mode B ("Who is there/Are you there"). This message contains the expression code of a wanted friend.

If the APP of UE-A chooses to announce its existence, the expression code of Jasmin is forwarded to the 3GPP layers of UE-A, as shown in Fig. 3. The 3GPP layers then broadcast a Mode A discovery message containing the expression code of Jasmin. When a friend of Jasmin (e.g., UE-B) receives a Mode A message, UE-B knows the existence of UE-A. The 3GPP layers of UE-B then send this knowledge to its APP. On the other hand, if the APP of UE-A chooses to find a particular friend (e.g., Alice owning UE-B), the APP of UE-A forwards the corresponding expression code of Alice to the 3GPP layers of UE-A, as shown in Fig. 4. The 3GPP layers of UE-A then broadcast a Mode B discovery message containing the expression code of Alice. When UE-B receives the Mode B message, UE-B replies with a Mode A discovery message to UE-A to inform UE-A of its existence. After identifying each other, the two UEs are ready to perform D2D communications.

Please note that following the above steps, two UEs know the existence of each other in the application layer. However, this does not imply that two UEs are able to identify each other in Layer 1. Since the content of a discovery message is only meaningful for the application layer, in Layer 1 a transmitter does not know the existence of the receiver (and thus the existence of the link). Such a paradigm is very different from the conventional Uu interface, where the discovery procedure between a UE and an eNB is initiated from Layer 1.

Air Interface of D2D Communications

Mobility/Resource Management

In Release 12 D2D communications, the paging procedure is not supported, especially for UEs out of coverage. As a result, mobility management is very limited, which only focuses on radio resource management. In the air interface of D2D communications, there are two modes of resource allocations:

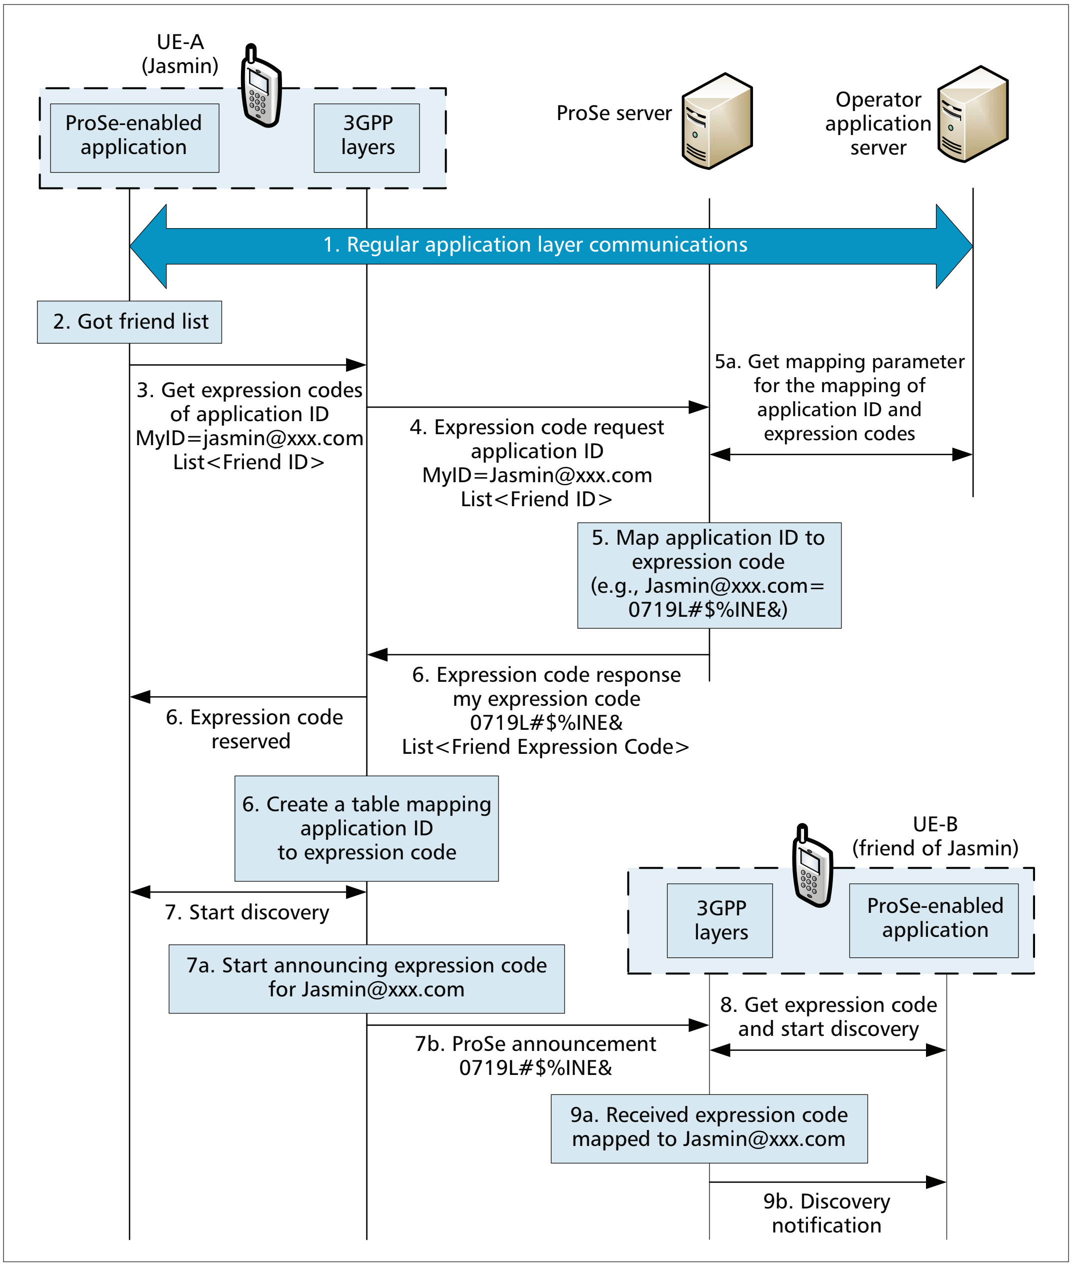

Figure 3. D2D discovery procedure using the Mode A discovery message.

In Release 12 D2D communications, the paging procedure is not supported, especially for UEs out of coverage. As a result, the mobility management is very limited, which only focuses on the radio resource management. In the air interface of D2D communications, there are two modes of resource allocations.

- **Mode 1.** An eNB schedules radio resources for a UE to transmit direct data and control information for D2D communications.
- **Mode 2.** A UE on its own "randomly" selects resources from the pre-configured resource pools to transmit direct data and control information for D2D communications.

Mode 1 resource allocation is supported by in-coverage UEs, while Mode 2 resource allocation can be supported by both in-coverage and out-of-coverage UEs. For Mode 2 resource allocation, if a UE is within eNBs' coverage, then the resource pools for this UE may be configured by eNBs. When a transmitter UE determines the radio resources for D2D data transmissions (either scheduled by the eNB or randomly selected by the UE), the transmitter UE has to inform the receiver UE about its radio resource selection. This information is referred to as sidelink control information (SCI), which will be detailed later. To provide reliability in open-loop communications, transmission repetition is used for both control and data transmissions. Each transmission of SCI is repeated twice, while the number and resources for the transmission repetition of each data packet (known as the time domain resource pattern of transmission (T-RPT)) are indicated by SCI. In addition, SCI also contains the modulation and coding scheme (MCS) for each data transmission, and the receiver group ID of each data transmission. Radio resource allocation for SCI transmissions also follows the Mode 1 and Mode 2 allocations. In other words, before data transmissions, a transmitter UE needs to first determine the radio resources for SCI. However, these radio resources are unknown by a receiver UE. Consequently, a receiver UE needs to perform non-coherent detection of these radio resources. By successfully detecting SCI, a receiver UE is able to obtain

The dynamic MCS is not for link adaptation. Instead, since multiple video and audio coding/decoding standards can be supported by D2D communications, the transmitter UE is able to properly arrange radio resources for various multimedia transmissions through the dynamic MCS.

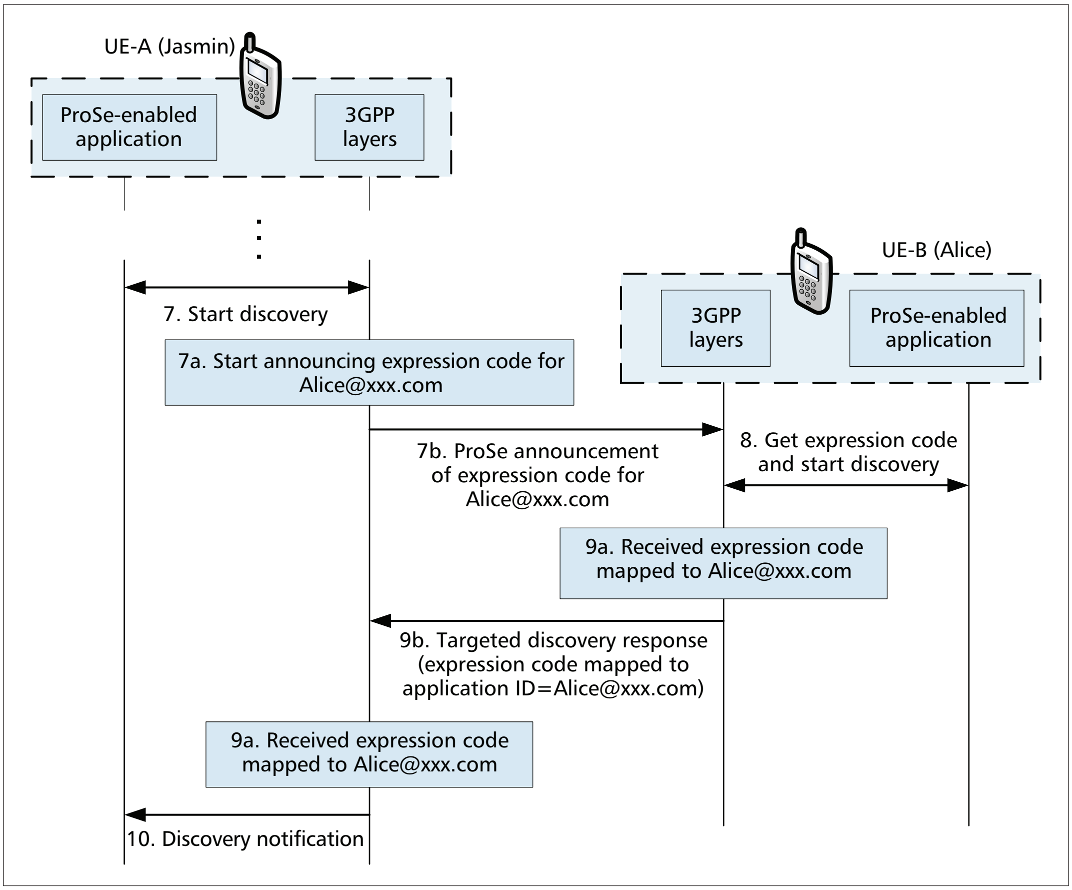

Figure 4. D2D discovery procedure using the Mode B discovery message.

the "locations" of radio resources for data transmissions.

Please note that although closed-loop link adaptation schemes such as adaptive modulation and coding, hybrid automatic repeat request (H-ARQ), and beam-forming, are precluded by Release 12 D2D communications, it is possible for a transmitter UE to dynamically change the MCS. The dynamic MCS is not for link adaptation. Instead, since multiple video and audio coding/decoding standards can be supported by D2D communications, the transmitter UE is able to properly arrange radio resources for various multimedia transmissions through the dynamic MCS.

D2D Physical Signals and Channels

In Layer 1, the physical uplink share channel (PUSCH) in the conventional Uu (eNB-UE) link is reused for D2D broadcasting communications. The reason of adopting the uplink channel is two-fold. First, in the PUSCH, transmitters are UEs and the receiver is an eNB. When a UE is uploading data to an eNB, other UEs are able to perform D2D communications concurrently if interference from the D2D link to the eNB is limited. In contrast, if the physical downlink share channel (PDSCH) is reused for D2D communications, a common transmitter is an eNB. In this case, if two UEs perform D2D communications concurrently with the downlink transmission, the eNB becomes a strong interference source to significantly degrade the performance of D2D. As a result, operating D2D communications at the PUSCH obtains higher spectrum efficiency. Second, single carrier frequency division multiple access (SC-FDMA) is adopted for the PUSCH in LTE/LTE-A, while orthogonal frequency division multiple access (OFDMA) is adopted for the PDSCH. Considering that the peak to average power ratio (PAPR) in SC-FDMA is smaller than that in OFDMA, operating D2D communications at the PUSCH also obtains higher energy efficiency. To support D2D communications, the following physical signals and channels are additionally provided along with PUSCH.

- **Sidelink Synchronization Signals and Physical Sidelink Broadcast Channel (PSBCH).** Since a synchronous operation may provide better performance in terms of interference management, spectrum utilization, and power consumption as compared with an asynchronous operation, it is expected that all eNBs and all UEs will follow the same timing reference. For this purpose, an eNB or a UE can transmit its timing reference for other UEs to synchronize to. An eNB or a UE transmitting its timing reference is referred as a **D2D synchronization source**. As a result, there can be multiple collocated D2D synchronization sources. The transmitted signals carrying the timing reference are thus referred to as sidelink synchronization signals. Therefore, an eNB is certainly a D2D synchronization source, and the corresponding sidelink synchronization signals is the Release 8 primary and secondary synchronization signals. In addition to the sidelink synchronization signals, a D2D synchronization source also transmits the PSBCH, which carries information to support D2D synchronization, such as the frame number for D2D communications, system bandwidth, the ID of the synchronization source, the type of synchronization source (i.e., the timing reference is derived by a UE or an eNB), time division duplex (TDD) or frequency division duplex (FDD) configuration, and the stratum level (this will be detailed in the following section).

•**Discovery Signal and Physical Sidelink Discovery Channel (PSDCH).** As mentioned in Fig. 3 and Fig. 4, the higher-layer discovery procedure is initiated by the APP generating the discovery message. In Layer 1, D2D discovery is achieved using Layer 1 discovery signal exchange via the PSDCH. For this purpose, there are two types of discovery procedures supposed by Layer 1.

- **Type 1:** Radio resources for discovery signal transmissions are allocated on a non-UE-specific basis. That is, the allocated radio resources can be used by all UEs to transmit their discovery signals.
- **Type 2:** Radio resources for discovery signal transmissions are allocated on a per-UE-specific basis. That is, the allocated radio resources are only utilized by a specific UE to transmit the discovery signal.
  ***Type 2A:** Radio resources are allocated at each discovery signal transmission time instant.
  ***Type 2B:** Radio resources are semi-persistently allocated for discovery signal transmissions.

The Type 1 discovery procedure can be supported by out-of-coverage UEs. In this case, the radio resource allocation can be pre-configured by the operator. On the other hand, the Type 2 discovery procedure can be supported by in-coverage UEs, by which radio resources are scheduled by an eNB.

Reference Signal. To assist the receiver UE to demodulate data, the transmitter UE can transmit the demodulation reference signal (DMRS) for the receiver UE to perform channel estimation. In LTE/LTE-A, the DMRS does not provide out-of-band channel information of the channel, thus the sounding reference signal (SRS) is used to estimate CSI outside the allocated radio resources. However, since the radio resource pools for D2D communications are scheduled by an eNB or pre-configured, frequency selective scheduling outside these radio resource pools is not allowed. Therefore, the SRS is avoided by D2D communications.

Physical Sidelink Control Channel (PSCCH) and Physical Sidelink Share Channel (PSSCH). Similar to the PUCCH and the PUSCH used for uplink control and data transmissions, the PSCCH and the PSSCH are used for Layer 1 control and data transmissions in D2D communications. On the PSCCH, a D2D transmitter can transmit sidelink control information (SCI) format 0, which carries the MCS used for D2D communications, the frequency hopping flag to indicate whether to activate frequency hopping in D2D communications, the RB assignment and hopping resource allocation to indicate the RB allocation, T-RPT as aforementioned, and the timing advance (TA) to indicate the timing difference between the downlink timing and the uplink timing.

It is possible that the above signals and channels may occupy the same radio resources. In this case, the sidelink synchronization signals and the PSBCH have the highest priority.

Physical Procedures

The major target of Release 12 D2D is to support public safety communications. Therefore, in addition to the communications between an eNB and a UE that shall be supported when all eNBs function normally, the communications among UEs shall be supported as well when some or all eNBs are paralyzed. To support this target, two physical layer procedures are crucial for partial-coverage and out-of-coverage UEs.

> The major target of Release 12 D2D is to support public safety communications. Therefore, in addition to the communications between an eNB and a UE that shall be supported when all eNBs function normally, the communications among UEs shall be supported as well when some or all eNBs are paralyzed.

D2D Synchronization

The D2D synchronization procedure is the first physical layer procedure as a UE powers on. In D2D communications, each UE needs to determine:

- Whether to be a D2D synchronization source.
- How to derive the timing reference if it decides not to be a synchronization source.

If an eNB is available, then all UEs can derive the timing reference directly from an eNB. However, for partial-coverage and out-of-coverage cases, some UEs may not receive the synchronization signals from an eNB. In these two cases, the D2D synchronization procedure is required to achieve the following two goals.

- In the partial-coverage case, it is expected that the timing references of all UEs (even not in the coverage of an eNB) align with the timing reference of the eNB.
- In the out-of-coverage case, the timing references of all UEs shall align with each other.

To achieve these goals, multi-hop synchronization is supported by D2D. That is, a timing reference generated from a UE or an eNB can be forwarded to multiple UEs, and these UEs can subsequently forward the received timing reference to other UEs. The D2D synchronization procedure includes the following principles, as illustrated in Fig. 5.

•When a UE powers on, it detects the synchronization signals from an eNB. If the synchronization signals from an eNB can be detected, this in-coverage UE synchronizes to the eNB with a similar procedure of the cell search in Release 8.

•In spite of the successful synchronization with an eNB, an in-coverage UE shall continuously detect the D2D synchronization signals and the PSBCH transmitted from other UEs. As previously mentioned, the PSBCH carries the ID of the synchronization source, the type of the synchronization source, and the stratum level (the hop-count of the D2D synchronization signals). If an in-coverage UE detects the D2D synchronization signals and the PSBCH indicating that the synchronization source type is a UE, this implies that there must be an out-of-coverage UE generating its own timing reference. In this case, the in-coverage UE transmits the D2D synchroniza-

D2D communications in Release 12 adopt open-loop transmissions, which is very different from the conventional paradigms in UMTS and LTE/LTE-A adopting closed-loop communications. Nevertheless, Release 12 is just the beginning stage for the development of D2D communications..

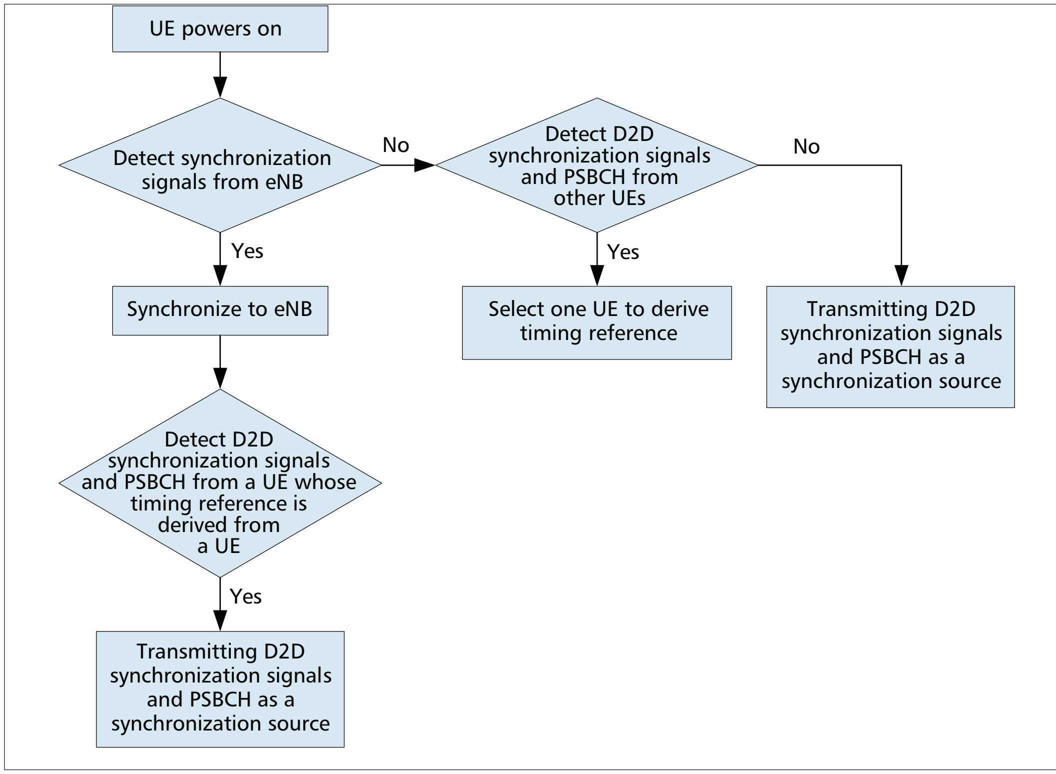

Figure 5. Principles of the D2D synchronization procedure.

tion signals and the PSBCH (where the synchronization source type is an eNB) and becomes a synchronization source. By providing the timing reference to out-of-coverage UEs, it is expected that the timing reference of this out-of-coverage UE aligns with the timing reference of the eNB.

- If no synchronization signal from an eNB can be detected, an out-of-coverage UE synchronizes to an in-coverage UE if the D2D synchronization signals and the PSBCH transmitted by an in-coverage UE can be detected. If multiple D2D synchronization signals and PSBCHs from different in-coverage UEs are detected, an out-of-coverage UE selects one appropriate in-coverage UE (and thus the D2D synchronization signals and the PSBCH) to derive the timing reference by taking the stratum level and signal strength of different PSBCHs into account.
- If no D2D synchronization signals and PSBCH from an in-coverage UE can be detected, an out-of-coverage UE synchronizes to another out-of-coverage UE if the D2D synchronization signals and the PSBCH transmitted by an out-of-coverage UE can be detected. In other words, the D2D synchronization signals and the PSBCH from an in-coverage UE have a higher priority than those from an out-of-coverage UE. If multiple D2D synchronization signals and PSBCHs from different out-of-coverage UEs are detected, the out-of-coverage UE selects one appropriate out-of-coverage UE (and thus the D2D synchronization signals and the PSBCH) to derive the timing reference by taking into account the stratum level and signal strength in different PSBCHs.
- If neither D2D synchronization signals and a PSBCH from an eNB nor D2D synchronization signals and a PSBCH from a UE (both in-coverage and out-of-coverage) can be detected, a UE transmits the D2D synchronization signals and the PSBCH, and becomes a synchronization source.

The performance of timing synchronization using the above D2D synchronization procedure is evaluated in Fig. 6. By randomly deploying 175 UEs (and among 175 UEs, there are 20 in-coverage UEs), we can observe from Fig. 6 that the timing references of 175 UEs can reach a consensus (i.e., differences between the timing references of 175 UEs and the eNB are less than 3 μs). This result confirms the effectiveness of the above D2D synchronization procedure.

Distributed Resource Access

As mentioned above, for Mode 1 D2D communication resource allocation and the Type 2 discovery procedure, the eNB schedules radio resources to in-coverage UEs. Thus, interference from D2D communications to cellular uplink transmissions and among different in-coverage D2D transmissions can be avoided. However, for Mode 2 D2D communication resource allocation and the Type 1 discovery procedure, radio resources are selected by UEs from the common resource pools. Although the resource pools for D2D communications can be pre-configured not to overlap the resources for cellular uplink transmissions, multiple UEs may select the same radio resources to result in interference. Unfortunately, in open-loop communications, a receiver UE is not able to inform the transmitter UE to change radio resources when interference occurs. Therefore, each transmitter UE is required to autonomously mitigate interference by avoiding occupying the same radio resources (i.e., resource collisions) with other UEs. For this purpose, each transmitter UE should "whiten interference."

To minimize the number of resource collisions, each out-of-coverage UE avoids selecting contiguous radio resources. Instead, each transmitter UE randomizes the radio resource occupation over the resource pools, which consequently randomizes/whitens interference. However, purely random radio resource selection largely increases the search space to detect needed data at the receiver side, which significantly increases receiver complexity. To tackle this engineering constraint, random radio resource selection is implemented by forming a number of time-frequency hopping patterns. When a time-frequency hopping pattern is selected by a transmitter UE, this pattern may be adopted for a period of time. Such a semi-persisted pattern considerably limits the search space to facilitate data detection at the receiver side.

Conclusion

In this article, a comprehensive overview of the state-of-the-art D2D communications in 3GPP Release 12 is presented, including reference architectures, discovery procedure, mobility/resource management, physical channels and signals, as well as physical procedures, to discuss practical issues and enable mechanisms for D2D communications over the B4G network. From this article, it is known that D2D communications in Release 12 adopt open-loop transmissions, which is very different from the conventional paradigms in UMTS and LTE/LTE-A adopting closed-loop communications. Nevertheless, Release 12 is just the beginning stage for the development of D2D communications, and a number of issues and performance optimizations still remain open for further investigations.

Acknowledgement

This study is conducted under the "III Innovative and Prospective Technologies Project" of the Institute for Information Industry, which is subsidized by the Ministry of Economy Affairs of the Republic of China, and is also supported by the Ministry of Science and Technology under the contract MOST 104-3115-E-009-004-.

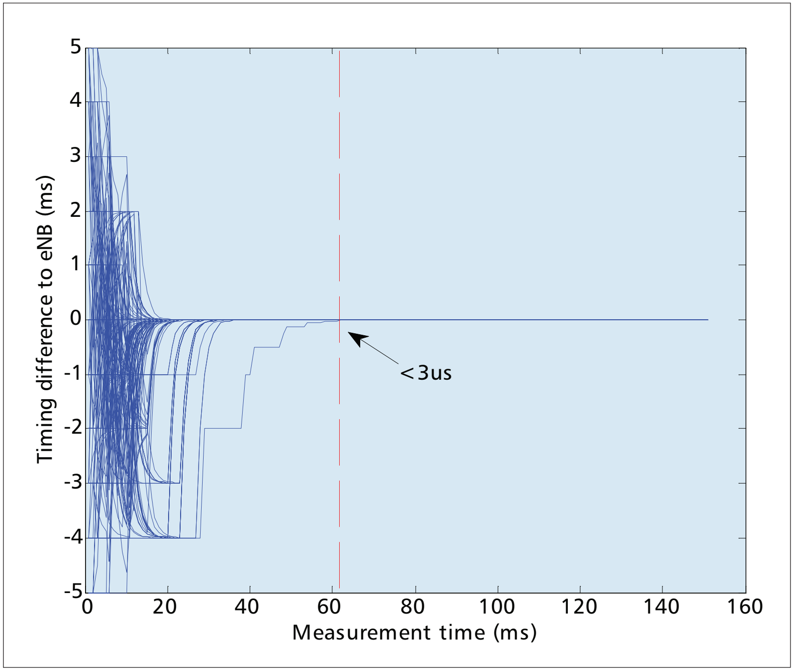

Figure 6. Timing differences between UEs and the eNB. In this simulation, initially, 175 UEs randomly distributed over a 1 km^2 area have a timing offset to the eNB. The timing offset is uniformly distributed over [–5, 5] ms. Transmission power of the D2D synchronization signals is 23 dBm. If the signal to interference and noise power of the received D2D synchronization signals exceeds 0 dB, the D2D synchronization signals can be successfully detected.

References

[1] K. Doppler *et al.*, "Device-to-Device Communication as an Underlay to LTE-Advanced Networks," *IEEE Commun. Mag.*, vol. 47, no. 12, Dec. 2009, pp. 42–49.

[2] D. Zhu *et al.*, "Downlink Resource Reuse for Device-to-Device Communications Underlaying Cellular Networks," *IEEE Signal Proc. Lett.*, vol. 21, no. 5, May 2014, pp. 531–34.

[3] P. Phunchongharn, E. Hossain, and D. I. Kim, "Resource Allocation for Device-to-Device Communications Underlaying LTE-Advanced Networks," *IEEE Wireless Commun.*, vol. 20, no. 4, Aug. 2013, pp. 91–100.

[4] M. Peng *et al.*, "Device-to-Device Underlaid Cellular Networks Under Rician Fading Channels," *IEEE Trans. Wireless Commun.*, vol. 13, no. 8, Aug. 2014, pp. 4247–59.

[5] D. Astely *et al.*, "LTE Release 12 and Beyond," *IEEE Commun. Mag.*, vol. 51, no. 7, pp. 154–160, July 2013.

[6] R. Ferrús *et al.*, "LTE: The Technology Driver for Future Public Safety Communications Release 12 and Beyond," *IEEE Commun. Mag.*, vol. 51, no. 10, Oct. 2013, pp. 154–61.

[7] 3GPP TR 23.703 V12.0.0, "Study on Architecture Enhancements to Support Proximity-based Services (ProSe)," Mar. 2014.

[8] 3GPP TR 36.843 V12.0.0, "Study on LTE Device to Device Proximity Services: Radio Aspects," Mar. 2014.

[9] 3GPP TR 36.837 V11.0.0, "Public Safety Broadband High Power User Equipment (UE) for Band 14," Jan. 2013.

[10] 3GPP TS 23.303 V12.0.0, "Proximity-based Services (ProSe)," Feb. 2014.

Biographies

Shao-Yu Lien (sylien@nfu.edu.tw) has been an assistant professor in the Department of Electronic Engineering, National Formosa University, Taiwan, since Feb. 2013. He has received a number of prestigious recognitions, including the IEEE Communications Society Asia-Pacific Outstanding Paper Award in 2014, the Scopus Young Researcher Award (issued by Elsevier) in 2014, the URSI AP-RASC 2013 Young Scientist Award, and the IEEE ICC 2010 Best Paper Award. His research interests include optimization techniques for networks and communication systems. Recently he has been focusing particularly on cyber-physical systems and 5G communication networks.

Chun-Che Chien (tez@iii.org.tw) received the Ph.D. degree from the Graduate Institute of Communication Engineering, National Taiwan University, Taipei, Taiwan in 2013. In 2004 he was an IC design engineer with the Multimedia Chip Design Team, Sunplus Technology Corp., Hsinchu, Taiwan. Since 2008 he has been a member of the Wireless Standardization Team at the Institute for Information Industry (III), Taipei, Taiwan. He currently takes part in standardization activities in 3GPP, focusing on next generation wireless communication technologies.

Min-Fan Tseng (farmingtseng@gmail.com) received B.S. and M.S degrees from the Department of Electrical Engineering and the Graduate Institute of Communication Engineering at National Taiwan University, Taipei, in 2008 and 2011, respectively. His research interests include network coding, cognitive radio networks, and information theory.

Tien-Chen Ho (tcho@ocu.edu.tw) received his B.S.E.E. degree from National Taiwan University in June 1973, the M.S. degree from the Computer Engineering Institute, National Chiao Tung University (NCTU) in June 1986, and the Ph.D. from the Department of Applied Mathematics, National Chung Hsing University, in June 1999. He was with the Central Taiwan Telecommunication Administration as an engineer, and a director of the Department of Engineering, South Taichung Business Unit, Chunghwa Telecom, in 2001. He was an assistant professor in the Department of Information Management, Overseas Chinese University, and is now an assistant professor in the Department of Information Technology, Overseas Chinese University.

LTE-Advanced in 3GPP Rel-13/14: An Evolution Toward 5G

While 5G will integrate the latest technology breakthroughs to achieve the best possible performance, it is expected that LTE-Advanced will continue to evolve, as a part of 5G technologies, in a backward compatible manner to maximize the benefit from the massive economies of scale established around the 3rd Generation Partnership Project LTE/LTE-Advanced ecosystem from Release 8 to Release 12.

Juho Lee, Younsun Kim, Yongjun Kwak, Jianzhong Zhang, Aris Papasakellariou, Thomas Novlan, Chengjun Sun, and Yingyang Li

The authors are with Samsung Electronics.

COMMUNICATIONS STANDARDS

Abstract

As the fourth generation (4G) LTE-Advanced network becomes a commercial success, technologies for beyond 4G and 5G are being actively investigated from the research perspective as well as from the standardization perspective. While 5G will integrate the latest technology breakthroughs to achieve the best possible performance, it is expected that LTE-Advanced will continue to evolve, as a part of 5G technologies, in a backward compatible manner to maximize the benefit from the massive economies of scale established around the 3rd Generation Partnership Project (3GPP) LTE/LTE-Advanced ecosystem from Release 8 to Release 12. In this article we introduce a set of key technologies expected for 3GPP Release 13 and 14 with a focus on air interface aspects, as part of the continued evolution of LTE-Advanced and as a bridge from 4G to 5G.

Introduction

The wireless cellular network has been one of the most successful communications technologies of the last three decades. The advent of smartphones and tablets over the past several years has resulted in an explosive growth of data traffic. With the proliferation of more smart terminals communicating with servers and each other via broadband wireless networks, numerous new applications have also emerged to take advantage of wireless connectivity.

After the introduction of the 4G LTE-Advanced [1, 2] standard in 3GPP Rel-10, LTE-Advanced has continued to evolve through several releases and has become a global commercial success. The research community is now increasingly looking beyond 4G and into future 5G technologies, both in standardization bodies such as 3GPP and in research projects such as the EU FP7 METIS. ITU-R has recently finalized work on the "Vision" for 5G systems, which includes support for explosive growth of data traffic, support for massive numbers of machine type communication (MTC) devices, and support for mission critical and ultra-reliable and low latency communications [3]. While today's commercial 4G LTE-Advanced networks are mostly deployed in legacy cellular bands from 600 MHz to 3.5 GHz, recent technology advancements will allow 5G to utilize spectrum opportunities below 100 GHz, including existing cellular bands, new bands below 6 GHz, and new bands above 6 GHz, including the so-called mmWave bands. There are coordinated efforts across the world to identify these new spectrum opportunities. There were decisions for new spectrums below 6 GHz at the World Radio-communication Conference (WRC)-2015, and further decisions for new spectrums above 6 GHz are expected at WRC-2019.

From the 5G technology roadmap perspective, we expect a dual-track approach to take place over the next few years in 3GPP. The first track is commonly known as the "evolution" track, where we expect the evolution of LTE-Advanced will continue in Rel-13/14 and beyond in a backward compatible manner, with the goal of improving system performance in the bands below 6 GHz. It is also our expectation that at least a part of 5G requirements can be met by the continued evolution of LTE-Advanced. For example, latency reduction with grant-less uplink access and shortened length of a transmission time interval (TTI) can make the over-the-air latency less than 1 ms. The second track is commonly known as the "new RAT" track in 3GPP, which is not limited by backward compatibility requirements and can integrate breakthrough technologies to achieve the best possible performance. The "new RAT" 5G system should meet all 5G requirements as it would eventually need to replace the previous generation systems in the future. The "new RAT" track is also expected to have a scalable design that can seamlessly support both above and below 6 GHz bands.

In this article, we focus on a set of important air interface features of LTE-Advanced in 3GPP Rel-13. We also discuss air interface features that are expected to be specified in 3GPP Rel-14. In the continued evolution of LTE-Advanced in Rel-13 and Rel-14, it is important to emphasize continuity and backward compatibility in order to leverage massive economies of scale associated with the current ecosystem developed around LTE/LTE-Advanced standards from Rel-8 to Rel-12.

Rel-13 includes three major technology categories. The first category is the enhancement of spectral efficiency, and its representative technology is full dimension MIMO (FD-MIMO) that aims to drastically increase spectral efficiency via the use of a large number of antennas at the base station. The second category is the utilization of additional frequency resources and includes licensed assisted access (LAA) for utilizing unlicensed spectrum while guaranteeing coexistence with existing devices and enhanced carrier aggregation (eCA) with up to 32 component carriers. In the third category are the technologies to support new services. A representative example is further cost reductions for MTC devices that can also support extended coverage. Other technologies in this category include enhancement of device-to-device (D2D) proximity services that was specified in Rel-12 for the support

 0163-6804/16/$25.00 © 2016 IEEE

of peer discovery as well as direct communication between proximity UEs, indoor positioning enhancements, and single-cell point-to-multipoint (SC-PTM) as a complementary tool for support of enhanced multimedia broadcast and multicast service (eMBMS).

Discussion on the evolution of LTE-Advanced in Rel-14 is already occurring. It is expected that there would be continued evolution of the features introduced in Rel-13, such as FD-MIMO and LAA. It is also expected that Rel-14 would introduce technologies for latency reduction, which is one of the most important aspects for improving the user experience but has not been improved much since the introduction of LTE. Technologies for vehicle-related services (V2X) such as vehicle-to-vehicle (V2V), vehicle-to-infra (V2I), and vehicle-to-pedestrian (V2P) have recently attracted significant attention from the cellular industry as another opportunity for LTE-Advanced technologies to be extended to support vertical industries, and are expected to be specified in Rel-14. As a technology for further improving spectral efficiency by allowing non-orthogonal downlink transmissions within a cell, the enhancement for downlink multiuser transmission using superposition coding was studied during Rel-13 for potential specification work in Rel-14.

The rest of this article is organized as follows. First, we describe FD-MIMO, LAA, and eCA with up to 32 component carriers, which would be the representative features in Rel-13 for improving spectral efficiency and utilizing additional frequency resources. Second, we introduce MTC as a representative example of Rel-13 technologies for support of new services. Third, we describe the features that are expected to be specified in Rel-14. The final section provides concluding remarks.

Full Dimension MIMO

FD-MIMO is one of the key candidate technologies considered for the evolution toward beyond 4th generation (B4G) and 5th generation (5G) cellular systems. The key idea behind FD-MIMO is to utilize a large number of antennas placed in a two-dimensional (2-D) antenna array panel to form narrow beams in both the horizontal and vertical directions. Such beamforming allows the enhanced NodeBs (eNB: 3GPP terminology for base station) to simultaneously transmit to multiple user equipments (UE: 3GPP terminology for mobile station) to realize high order multi-user spatial multiplexing.

Figure 1 depicts an eNB with FD-MIMO implemented using a 2-D antenna array panel, where every antenna is an active element allowing dynamic and adaptive precoding across all antennas. By utilizing such precoding, the eNB can simultaneously direct transmissions in the azimuth and elevation domains for multiple UEs. The key feature of FD-MIMO in improving the system performance is its ability to realize high order multi-user multiplexing.

3GPP has conducted several studies since December 2012 in an effort to provide specification support for FD-MIMO. The first step was a study item [4] for developing a new channel model for future evaluation of antenna technologies based on 2-D antenna array panels. The channel model provides the stochastic characteristics of a three-dimensional (3-D) wireless channel. Based on the new channel model, a follow-up study item [5] on FD-MIMO was initiated in September 2014 to evaluate the performance benefits of standard enhancements targeting the 2-D antenna array operation with up to 64 antenna ports over a standard-transparent approach such as vertical sectorization utilizing antenna elements in the vertical direction.

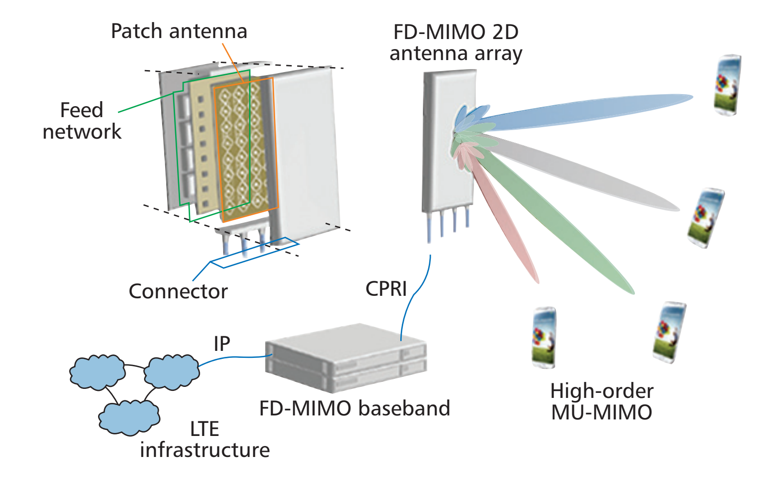

Figure 1. Conceptual diagram of a FD-MIMO system realizing high order MU-MIMO through a 2-D antenna array.

FD-MIMO has two important differentiating factors compared to MIMO technologies from previous LTE releases. First, the number of antennas can be increased beyond eight, e.g. to 64. As a result, FD-MIMO significantly improves beamforming and spatial user multiplexing capability. Second, specification support for FD-MIMO is targeted for antennas placed on a 2-D planar array. Using the 2D planar placement is also helpful in reducing the form factor of the antennas for practical applications.

Figure 2 summarizes the cell average throughput of FD-MIMO for various eNB antenna configurations (left) and various numbers of UEs per cell (right), where all UEs are distributed on the same horizontal plane, i.e. without UE distribution along the vertical direction. More details can be found in [6]. The antenna configuration is determined by the number of transmit antennas (N_T) and the 2-D antenna placement ($N_H \times N_V$) where N_H and N_V are the number of antennas on the horizontal and vertical directions, respectively.

The performance of MIMO utilizing up to eight transmit antennas was evaluated to be around 2.8 bps/Hz for cell average throughput and 0.1 bps/Hz for 5 percent-tile user throughput (cell edge performance) when there were 10 UEs per cell. Compared to these values, Fig. 2 (left) shows dramatic performance enhancements for FD-MIMO of up to 400 percent or more for both cell average throughput and 5 percent-tile user throughput. Additionally, Fig. 2 shows that although a one-dimensional horizontal antenna array provides the best performance, FD-MIMO with 2-D arrays also provides significant performance improvement while allowing for practical implementations.

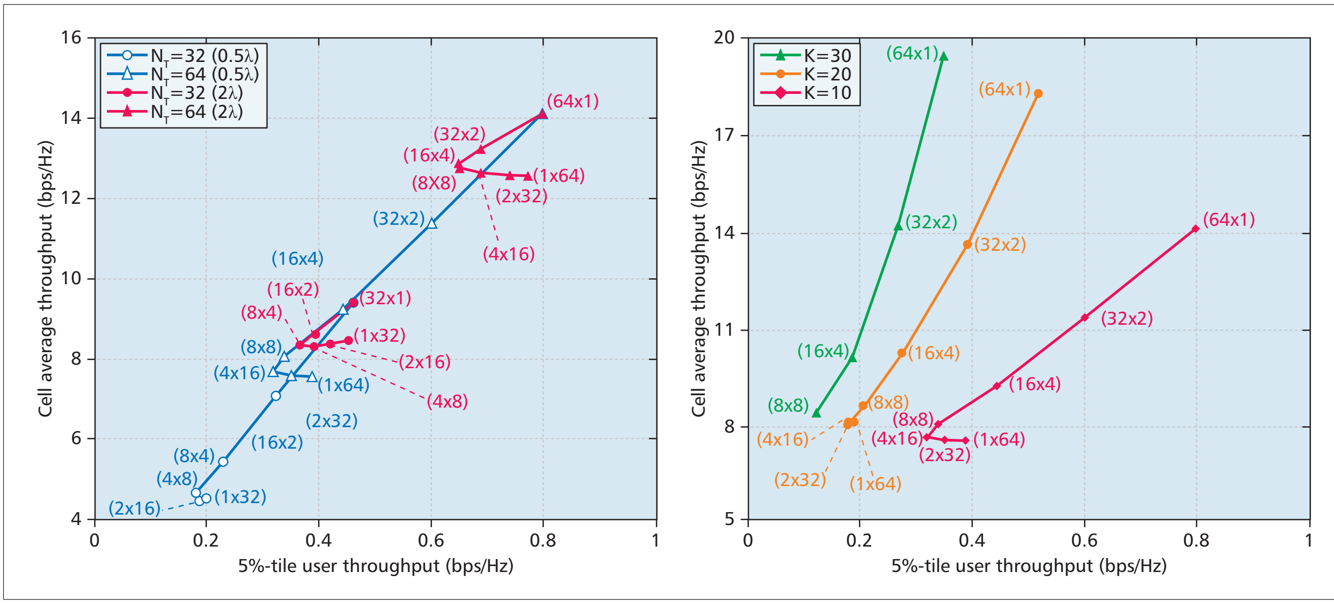

Figure 2. FD-MIMO cell average throughput for various antenna configurations (left) and various numbers of UEs in a cell (right).

Figure 2 (right) shows the cell average throughput for $N_T = 64$ antennas and for 10, 20, and 30 UEs per cell. It was assumed that the eNB transmits spatially multiplexed signals to all UEs in the cell. The results show that as the number of UEs per cell increases, the cell throughput of FD-MIMO also increases due to the associated increase in the order of multi-user spatial multiplexing.

Taking into account the significant performance benefit observed in the feasibility study and the available time for the specification work, LTE-Advanced in Rel-13 specified support for FD-MIMO with up to 16 transmit antenna ports at the eNB. The following enhancements were specified in Rel-13:

- Enhancement of downlink reference signals for measurements from a larger number of antennas on a 2-D array panel.
- Enhancement of channel state information (CSI) such as CSI reporting mechanism and codebook for spatial beamforming in both horizontal and vertical directions.
- Enhancement of demodulation reference signals to enable high order multi-user multiplexing.

FD-MIMO is expected to bring significant performance enhancements to future generations of cellular networks due to the wide range of deployment environments. FD-MIMO can be deployed not only for outdoor macro cells but also for smaller cells such as indoor, micro, and pico cells. Additionally, considering that the antenna spacing is inversely proportional to the carrier frequency, FD-MIMO systems can be deployed with smaller form factors for higher frequency bands.

Licensed-Assisted Access Using LTE

Due to the sharply increased demand for wireless broadband data, the use of unlicensed spectrum is now being considered as a potential complement to LTE systems operating in licensed spectrum. Although licensed spectrum affords operators exclusive control for providing guaranteed QoS and mobility, available bandwidths are typically limited and can be very costly to obtain. In particular, the 5 GHz unlicensed band has attracted considerable interest for LTE deployments due to the potentially large amount of globally available spectrum (>400 MHz). However, LTE operation in unlicensed spectrum needs to coexist with the operation of other radio access technologies (RATs), such as Wi-Fi, and this leads to unique design challenges compared to LTE operation in licensed bands. In June 2015, 3GPP completed the study on licensed-assisted access to investigate and identify possible designs to allow LTE to coexist with Wi-Fi in unlicensed bands [7], and started a work item for specification support.

In Rel-13, only downlink transmissions using the unlicensed band were specified, and the principle of uplink channel access and the necessary forward compatibility mechanism were developed so that the uplink transmission can be added in future releases without modifications to the downlink design. LAA operation in Rel-13 uses carrier aggregation to tightly integrate licensed spectrum and unlicensed spectrum. The primary cell (PCell) is maintained on licensed spectrum to provide control, system information, and continuity for high-QoS services, since unlicensed carriers are likely to be only intermittently available in order to satisfy the coexistence requirement of LAA with other RATs. The main deployment scenarios targeted by the LAA in Rel-13 include outdoor and indoor small cell deployments, with and without the presence of macro cell coverage, and with licensed and unlicensed carriers collocated in the same box.

3GPP intends to define a single global solution framework ensuring that LAA can meet regulatory requirements for all different regions and can co-exist among operators or with other RATs operating on the same band. As a result, a major focus of the work in 3GPP was to perform extensive evaluations for the LAA design

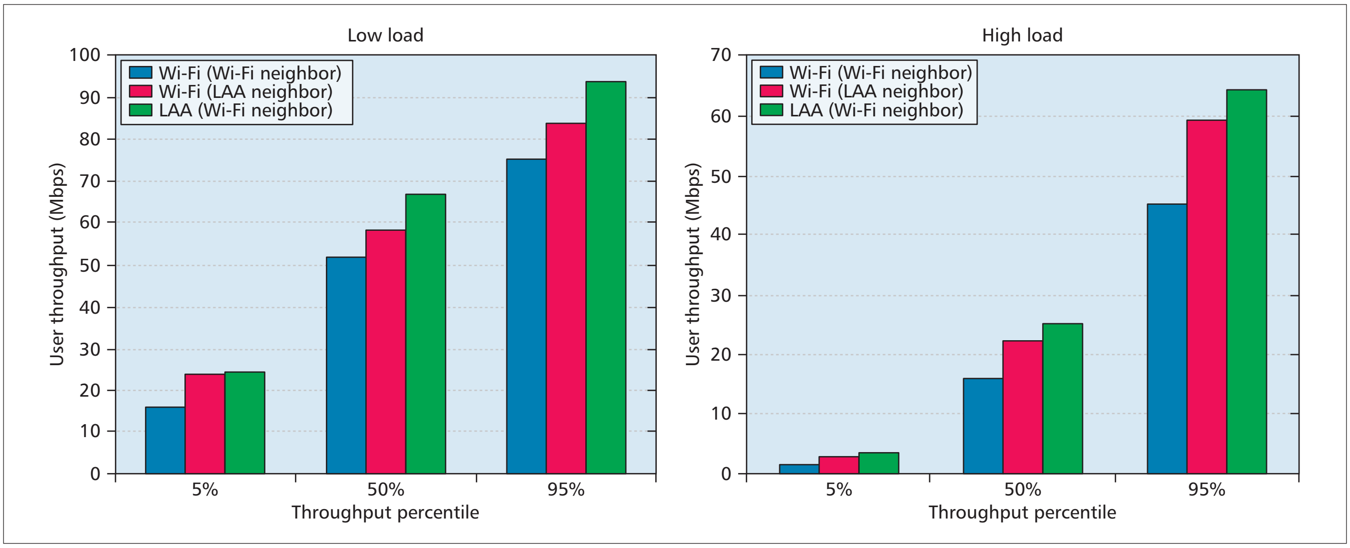

Figure 3. User throughput for coexisting Wi-Fi and LAA networks for low and high traffic loads (results from [6]), of which detailed simulation assumptions can be found in section A.1 of [6].

functionalities that ensure fair coexistence with Wi-Fi deployments. For example, one can consider a typical Wi-Fi network with two different Wi-Fi nodes. If one of them is replaced by a hypothetical LAA node, it should be ensured that the throughput and latency of the remaining Wi-Fi node should not be negatively impacted relative to the original Wi-Fi/Wi-Fi setup. Likewise, coexistence evaluations were also carried out for multi-operator LAA scenarios, since coordination and synchronization between nodes belonging to different operators cannot be universally assumed.

One key mechanism for enabling LAA/Wi-Fi and multi-operator LAA coexistence is listen-before-talk (LBT). LBT governs when a LAA cell may access the channel. For example, according to the European regulations in [8] for load-based equipment, clear channel assessment (CCA) must be performed prior to starting a new transmission. An extended CCA is performed if the medium is determined to be occupied during the CCA and transmission is postponed until the channel is considered clear. An LBT mechanism is expected to be a major part of the specification support for LAA. The intermittent nature of LAA transmissions also has significant implications on existing LTE functionalities such as radio resource management (RRM) measurements, automatic gain control (AGC) settings, coarse and fine time/frequency synchronization, and CSI measurements. Further, support for uplink operation on unlicensed spectrum requires careful study of the resource allocation and feedback mechanisms that may be impacted due to their current reliance on synchronized and fixed timelines.

Due to its capability to aggregate traffic across both licensed and unlicensed bands, and due to its advanced link management techniques, LAA is expected to improve the utilization efficiency of unlicensed spectrum. Consequently, LAA can also benefit other RATs by providing more opportunities for channel access. Based on the evaluation methodology defined in the 3GPP study item [7], Fig. 3 shows user throughput at low and high traffic loads when two neighboring Wi-Fi networks, each of which has four Wi-Fi nodes deployed on a single floor of 120 m x 50 m, coexist on a single unlicensed carrier, and when one operator's Wi-Fi nodes are replaced with LAA nodes.

Although the traffic and user densities of both networks are the same, the network utilizing LAA achieves significant gain in user-perceived throughput, due in part to fundamental LTE capabilities such as link adaptation based on explicit UE feedback and Hybrid ARQ (HARQ). It is additionally observed from Fig. 3 that LAA is a better neighbor to the remaining Wi-Fi network than the previous Wi-Fi network it replaced. This is because a more efficient packet transmission, combined with an adaptive channel utilization mechanism such as LBT, provides increased opportunities for networks coexisting with LAA to access the channel without contention. These benefits also extend to the case where LAA networks of different operators coexist as well. More evaluation results can be found in [7].

Carrier Aggregation Enhancements

As a natural approach for increasing the peak rate and improving the utilization efficiency of distributed frequency resources, carrier aggregation of up to five component carriers with common FDD or TDD duplexing was specified in Rel-10 to support a maximum combined bandwidth of 100 MHz. The combination of carrier aggregation and MIMO provides 3 Gb/s and 1.5 Gb/s peak rate on the downlink and the uplink, respectively. In Rel-12, carrier aggregation was extended to support aggregation of FDD carriers and TDD carriers, but the constraint of aggregating at most five carriers remained. This constraint limits commercial deployments, particularly considering the availability of the 5 GHz unlicensed band that can provide tens of 20 MHz carriers. Additionally, 3GPP is currently studying UE RF requirements for the introduction of CA with four downlink carriers, and it is expected that commercial needs would soon exceed the Rel-12 limitation of five downlink carriers.

Motivated by the above considerations, a new work item was approved for Rel-13 with the objective to specify CA operation for up to 32 carriers (or cells), which can support a peak rate of 25 Gb/s. The main specification impacts are on uplink control signaling and on the reduction of control channel decoding operations that a UE needs to perform. According to the Rel-12 design principle, the number of control channel decoding operations required for a UE increases almost linearly with the number of scheduling cells the UE can support. Rel-13 CA limits this increase by the eNB, essentially configuring the number of blind decoding operations a UE performs per carrier subject to a respective capability reported by the UE. The amount of uplink control information is increased to support HARQ-ACK information or channel state information for a large number of downlink carriers. Further, a UE can be configured to transmit the uplink control information on a secondary cell (SCell) in addition to the PCell to reduce the control signaling overhead of the PCell.

Machine-Type Communications

Support of machine-type communications through cellular networks is emerging as a significant opportunity for new applications in a networked world where devices, e.g. smart power meters, street lights, cars, home electronics such as refrigerators and TVs, and surveillance cameras, communicate with humans and with each other. LTE offers a proven technology with a large existing ecosystem for MTC UEs, but it is also associated with LTE-specific design challenges primarily due to the requirement on a network to simultaneously support UEs with significantly different capabilities and also support coverage enhancements.

Rel-12 specifications for MTC UEs achieved a cost reduction of approximately 50 percent relative to the lowest category LTE UEs (category 1 LTE UEs), and Rel-13 MTC UEs are expected to achieve an additional 50 percent cost reduction primarily through restrictions in transmission/reception within only six resource blocks (RBs) of a system bandwidth per TTI and a lower power amplifier gain, where the RB bandwidth is 180 kHz [9].

The absence of receiver antenna diversity and the possible reduction in a power amplifier gain can result in significant reductions in coverage even for Rel-13 MTC UEs that do not experience large path-loss. A key design target is to provide up to 15 dB coverage enhancement while minimizing the impact on network spectral efficiency and MTC UE power consumption. Coverage enhancement is mainly achieved by repetitions. In order to reduce the required number of repetitions, other physical layer techniques, such as the use of multiple contiguous TTIs to improve channel estimation accuracy and frequency error correction at the receiver, and frequency hopping to increase the frequency diversity gain, are also specified. Narrowband Internet of things (NB-IoT) is also being specified in Rel-13 as another approach for efficient support of low-throughput (~50 kbps) cellular IoT devices using a very narrow bandwidth of 180 kHz (one RB). NB-IoT can be deployed by refarming the 200 kHz GSM carriers, using a single RB in LTE systems, or using a part of the guard band in LTE systems [10].

Features Being Discussed for Inclusion in Rel-14

It is expected that LTE-Advanced in Rel-14 will continue enhancements on FD-MIMO and LAA. For FD-MIMO, the number of eNB transmit antenna ports may be increased to 32 to support larger arrays, and other enhancements may also be specified, including support for more robust FD-MIMO transmission via open loop operation with reduced feedback overhead. For LAA, the support of uplink transmission in the unlicensed band is expected. The rest of this section focuses on latency reduction, V2X, and downlink multiuser transmission using superposition coding, which are expected to be standardized in Rel-14 following the feasibility studies in Rel-13.

Latency Reduction

Latency is one of the most important performance metrics for evaluating wireless communication systems. LTE provides less than 10 ms user plane air latency, and it is now recognized to be a system that provides lower data latencies than previous generations of mobile radio technologies. However, taking into account various emerging applications, tighter latency requirements need to be met, e.g. 1 ms over-the-air latency is being considered as an important requirement of 5G communication systems. 3GPP has started a study [11] of possible technologies for latency reduction, and we expect that they will be specified in Rel-14.

Uplink data transmission involves a scheduling request (SR) by a UE, a resource grant by an eNB, and data packet transmission by the UE. The request-grant procedure occupies a large portion of the entire latency required for uplink data transmission, especially for the transmission of small size payloads such as TCP/IP ACK/NACK. Introducing a grant-less procedure, i.e. removing the request-grant procedure, would be helpful in achieving low latency for small packets. On the other hand, the request-grant procedure is still useful in serving large packets as it enables highly efficient utilization of the valuable spectrum resource.

Another approach gaining attention is shortening the TTI length. In the current LTE standard, the TTI length is 1 ms and is equal to the duration of a subframe, which consists of two slots and corresponds to 14 OFDM symbols. Reducing the TTI length to one slot, i.e. 0.5 ms, or to one OFDM symbol duration, i.e. 0.07 ms, can be good candidates because they allow easy multiplexing of signals with legacy and reduced TTI lengths. With the reduced TTI length, it would be natural to assume that the UE and eNB processing time can be proportionally reduced due to a smaller amount of data to process.

Figure 4 illustrates latency incurred in the existing procedures, and Table 1 shows the latency reduction gain achieved by a grant-less procedure and TTI shortening. The over-the-air latency is lower than 1 ms for a grant-less procedure with a TTI length of one OFDM symbol.

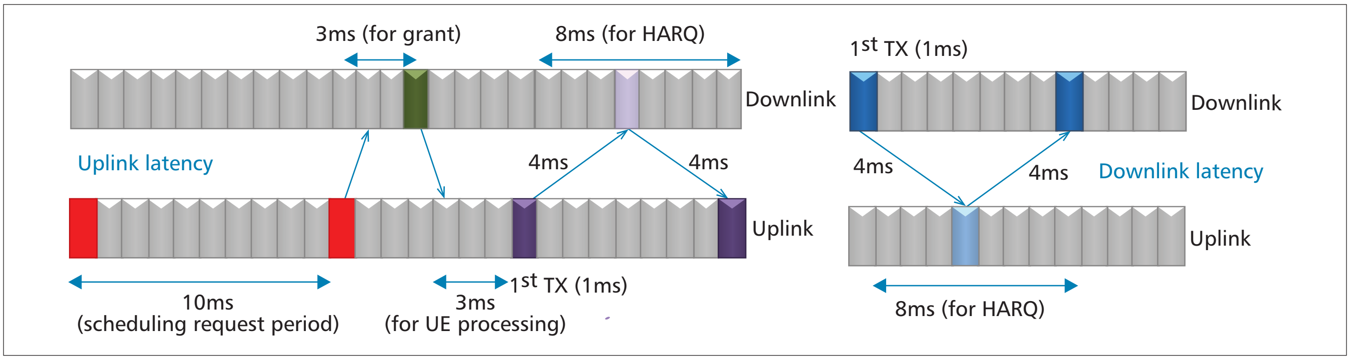

Figure 4. Uplink and downlink latency considering grant procedure and HARQ, where average HARQ delay is 0.8ms assuming a 10 percent BLER and a maximum of 1 retransmission.

TTI	SR	Grant	UE Proc	FA	TTI	Average HARQ delay	eNB Proc	Total uplink latency	
								Grant procedure	Grant-less procedure
1ms	6ms	3ms	3ms	0.5ms	1ms	0.8ms	1.5ms	15.8ms	4.8ms
0.5ms	3ms	1.5ms	1.5ms	0.25ms	0.5ms	0.4ms	0.75ms	7.9ms	2.4ms
0.07ms	0.43ms	0.21ms	0.21ms	0.04ms	0.07ms	0.06ms	0.11ms	1.13ms	0.35ms

Table 1. Latency analysis depending on TTI lengths, where UE processing time (UE Proc) is 1 TTI length for grant-less case and frame alignment (FA) time is half of TTI length.

V2X

A vehicle-centric communications network is one of the key enabling technologies for the emerging 'connected car' ecosystem, supporting a broad range of new services and applications, including automotive safety, autonomous vehicles, telematics, traffic control, and infotainment. 3GPP started studies of the use of LTE mobile networks to enable connectivity between vehicles (V2V), between vehicles and roadway infrastructure (V2I), and between vehicles and pedestrians (V2P) or other mobile users, jointly known as LTE V2X, as shown in Fig. 5.

As a wide array of different sensing and positioning technologies (e.g. mmWave radar, video, and high-precision GNSS) become standard automotive features, one key V2X service is the timely and reliable delivery of critical messages to improve safety and traffic congestion. However, the exchange of these messages can be challenging due to a large variation in message sizes, strict end-to-end latency requirements, a potential range of several hundred meters, and support of high Doppler spread when, for example, two vehicles are directly approaching each other while each one is traveling at a speed of 140km/h (equivalent Doppler speed of 280 km/h).

LTE V2X is 3GPP's response to increasing market interest as well as increasing expectations that regulatory bodies worldwide consider technology requirements and potential mandates in the next few years for vehicle communication networks. For example, Korea, Japan, the EU, and the US have allocated frequency spectrum in the 5.8-5.9 GHz range for dedicated short range communications (DSRC) to support intelligent transportation systems (ITS). In China, CCSA has also conducted studies of the feasibility of providing vehicle safety services over LTE, and it is expected that the National Regulatory Authority in China will allocate dedicated frequency spectrum for V2X. It is envisioned that LTE V2X could be deployed on licensed or shared spectrum, and it supports operation outside of the infrastructure network coverage with technologies such as enhanced D2D communications.

Alternative technologies have been developed for V2X applications and have been the subject of many academic and industry research projects and trials, with the most prominent being the IEEE WAVE and 802.11p standards. However, LTE V2X is a very attractive candidate due to its fundamental support for wide-area coverage, mobility, and spectrally efficient V2I broadcast services using eMBMS. Meanwhile, flexible and scalable resource allocation functionalities for V2V and V2P can be built upon the recently introduced proximity services (ProSe) functionalities (also known as D2D).

The feasibility study of V2X [12] is expected to evaluate the performance of different LTE-based solutions for providing V2V, V2I, and V2P services, with the goal of specifying identified enhancements starting in Rel-14. Support for V2V and V2P over D2D communication links between UEs is being studied with the highest priority, including potential resource allocation and channel estimation enhancements to support efficient and robust transmissions with low-latency. In addition, provisioning of V2X services over the link between the LTE network and the UE is also within the scope of the study, including the applicability of latency reduction and multi-cell multicast/broadcast enhancements to sufficiently meet industry and regulatory requirements for

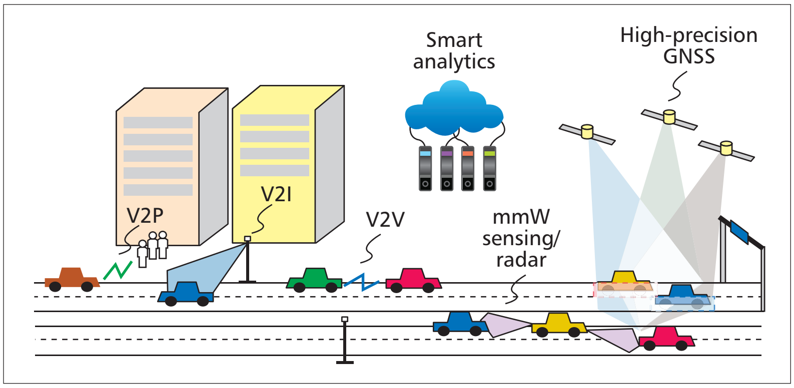

Figure 5. Overview of V2X scenarios and component technologies.

V2X. While the V2X study item is still continuing, taking into account the urgent request from industries, a work item to specify V2V support started in December 2015, with the targeted completion date of September 2016 [13].

Enhanced Downlink Multiuser Transmission Using Superposition Coding

Support of simultaneous non-orthogonal transmissions without spatial separation on the downlink has the potential to further improve system capacity. For example, downlink transmissions to a UE located at the cell boundary and to another UE located at the cell center can be scheduled using the same beam. The former UE would typically be allocated a large transmit power, and hence its interference to the latter UE can be cancelled before decoding the desired signal at the latter UE. Possible performance benefits and required specification support to assist intra-cell interference cancellation or suppression at the UE receiver were studied in Rel-13 [14] for preparation of potential specification work in Rel-14.

Conclusion

In this article we have introduced a set of features for the evolution of LTE-Advanced that are expected to be specified in 3GPP Rel-13/14, taking into account the latest status of discussions in 3GPP. These timely enhancements would allow the cellular industry to improve the efficiency of the network, and in the meantime continue to benefit from the massive economies of scale associated with the current ecosystem developed around LTE/LTE-Advanced standards from Rel-8 to Rel-12. We also expect that they would serve as the bridge from 4G to 5G, taking into account that 5G would consist of both the continued evolution of LTE-Advanced and the introduction of non-backward compatible breakthrough technologies.

References

[1] J.C. Zhang, S. Ariyavisitakul, and M. Tao, "LTE-Advanced and 4G Wireless Communications (Guest Editorial)," *IEEE Commun. Mag.*, vol. 50, no. 2, 2012, pp. 102–03.
[2] 3GPP TR 36.912, "Feasibility Study for Further Advancements for E-UTRA (LTE-Advanced)," 2010.
[3] ITU-R WP-5D, Document 5D/TEMP/625-E, "IMT Vision — Framework and Overall Objectives of the Future Development of IMT for 2020 and Beyond."
[4] 3GPP TR 36.873, "Study on 3D Channel Model for LTE," 2014
[5] 3GPP TR 36.897, "Study on Elevation Beamforming/Full-Dimension (FD) MIMO for LTE," 2015.
[6] Y. Kim *et al.*, "Full Dimension MIMO (FD-MIMO): The Next Evolution of MIMO in LTE Systems," *IEEE Wireless Commun.*, vol. 21, no. 2, 2014, pp. 26–33.
[7] 3GPP TR 36.889, "Study on Licensed-Assisted Access Using LTE," 2015.
[8] ETSI, EN 301.893, "Broadband Radio Access Networks (BRAN); 5 GHz High Performance RLAN; Harmonized EN Covering the Essential Requirements of Article 3.2 of the R&TTE Directive," 2015.
[9] 3GPP TR 36.888, "Study on Provision of Low-Cost Machine-Type Communications (MTC) User Equipments (UEs) Based on LTE," 2012.
[10] 3GPP RAN Tdoc RP-151621, "New WI Proposal: Narrow Band IOT," approved in RAN plenary #69 meeting, Sept. 2015.
[11] 3GPP RAN Tdoc RP-150465, "New SI Proposal: Study on Latency Reduction Techniques for LTE," approved in RAN plenary #67 meeting, Mar. 2015.
[12] 3GPP RAN Tdoc RP-151109, "New SI Proposal: Feasibility Study on LTE-based V2X Services," approved in RAN plenary #68 meeting, June 2015.
[13] 3GPP RAN Tdoc RP-152293, "New WI Proposal: Support for V2V Services Based on LTE Sidelink," approved in RAN plenary #70 meeting.
[14] 3GPP TR 36.859, "Study on Downlink Multiuser Superposition Transmission for LTE," 2015.

Biographies

Juho Lee (juho95.lee@samsung.com) is currently a Master with Samsung Electronics and is the lead executive in charge of research on standardization of wireless communications. He received his B.S., M.S., and Ph.D. degrees in electrical engineering from Korea Advanced Institute of Science and Technology (KAIST), Korea, in 1993, 1995, and 2000, respectively. He joined Samsung Electronics in 2000 and has been working on standardization of mobile communications for 3G and 4G such as WCDMA, HSDPA, HUSPA, LTE, and LTE-Advanced. He is also actively working on research and standardization for 5G. He was a vice chairman of 3GPP TSG RAN WG1 from February 2003 to August 2009, chaired LTE/LTE-Advanced MIMO sessions, and served as the rapporteur for the 3GPP LTE-Advanced Rel-11 CoMP work item.

Younsun Kim (younsun@samsung.com) received his B.S. and M.S. degrees in electronic engineering from Yonsei University, Korea, and his Ph.D. degree in electrical engineering from the University of Washington in 1996, 1999, and 2009, respectively. He joined Samsung Electronics in 1999, and has since been working on the standardization of wireless communication systems such as cdma2000, HRPD, and recently LTE/LTE-A. His research interests include multiple access schemes, network coordination schemes, advanced receiver algorithms, and advanced antenna techniques for next generation systems.

Yongjun Kwak (yongjun.kwak@samsung.com) received his B.S. degree in electronics and electrical engineering from POSTECH, Korea, his M.S. degree in electrical engineering from KAIST, Korea, and the Ph.D. degree in engineering sciences from Harvard University, USA in 1998, 2000, and 2012, respectively. He joined Samsung Electronics in 2000 and has been working on physical layer standardization of WCDMA, HSDPA, HUSPA, and LTE technologies. His research interests include multiple access, channel coding, system design for 5G wireless communication systems.

Jianzhong (Charlie) Zhang (Jianzhong.z@samsung.com) is currently a vice president and head of the Standards and Research Lab with Samsung Research America in Dallas, where he leads research and standards for 5G cellular systems as well as next generation multimedia networks. From August 2009 to August 2013 he served as the vice chairman of the 3GPP RAN1 working group and led development of LTE and LTE-Advanced technologies such as 3D channel modeling, UL-MIMO and CoMP, carrier aggregation for TD-LTE, etc. Before joining Samsung he was with Motorola from 2006 to 2007 working on 3GPP HSPA standards, and with Nokia Research Center from 2001 to 2006 working on the IEEE 802.16e (WiMAX) standard and EDGE/CDMA receiver algorithms. He received his Ph.D. degree from the University of Wisconsin, Madison.

Aris Papasakellariou (aris.papas@samsung.com) received his B.S. degree in physics from the University of Athens, Greece, and his M.S. and Ph.D. degrees from Rice University in 1991, 1993, and 1996, respectively. He joined Samsung Electronics in 2007 and has been working on 3GPP standardization since 2005.

Thomas Novlan (t.novlan@samsung.com) received his B.S. with high honors, M.S., and Ph.D., in electrical engineering from the University of Texas at Austin in 2007, 2009, and 2012, respectively. He joined Samsung Electronics in 2012 and is working on wireless communications research and LTE-A standardization for device-to-device and small cell-related technologies.

Chengjun Sun (chengjun.sun@samsung.com) received his B.S., M.S., and Ph.D. degrees in signal and information processing from Harbin Institute of Technology (HIT), China, in 1997, 1999 and 2002, respectively. He joined Samsung Electronics in 2003 and has been working on standardization of mobile communications such as TD-SCDMA, WCDMA, HSUPA, LTE, and LTE-Advanced. His research interests include link adaptation, duplexing, radio resource management, and carrier aggregation for next generation systems.

Yingyang Li (yingyang2.li@samsung.com) received his B.S. degree on communications from Nanjing University of Posts & Telecommunications (NUPT) in 2001, and his M.S. degree oi mobie communication from Beijing University of Posts & Telecommunications (BUPT) in 2004. He joined Samsung Electronics in 2004 and has been working on the standardization of mobile communications such as WCDMA, LTE, and LTE-Advanced. His research interests include multiple access schemes, duplexing, and carrier aggregation for next generation systems.

CALL FOR PAPERS
IEEE COMMUNICATIONS MAGAZINE
COMMUNICATIONS STANDARDS SUPPLEMENT

BACKGROUND

Communications standards enable the global marketplace to offer interoperable products and services at affordable cost. Standards development organizations (SDOs) bring together stakeholders to develop consensus standards for use by a global industry. The importance of standards to the work and careers of communications practitioners has motivated the creation of a new publication on standards that meets the needs of a broad range of individuals, including industrial researchers, industry practitioners, business entrepreneurs, marketing managers, compliance/interoperability specialists, social scientists, regulators, intellectual property managers, and end users. This new publication will be incubated as a Communications Standards Supplement in *IEEE Communications Magazine*, which, if successful, will transition into a full-fledged new magazine. It is a platform for presenting and discussing standards-related topics in the areas of communications, networking, and related disciplines. Contributions are also encouraged from relevant disciplines of computer science, information systems, management, business studies, social sciences, economics, engineering, political science, public policy, sociology, and human factors/usability.

SCOPE OF CONTRIBUTIONS

Submissions are solicited on topics related to the areas of communications and networking standards and standardization research, in at least the following topical areas:

Analysis of new topic areas for standardization, either enhancements to existing standards or of a new area. The standards activity may be just starting or nearing completion. For example, current topics of interest include:

- 5G radio access
- Wireless LAN
- SDN
- Ethernet
- Media codecs
- Cloud computing

Tutorials on, analysis of, and comparisons of IEEE and non-IEEE standards. For example, possible topics of interest include:

- Optical transport
- Radio access
- Power line carrier

The relationship between innovation and standardization, including, but not limited to:

- Patent policies, intellectual property rights, and antitrust law
- Examples and case studies of different kinds of innovation processes, analytical models of innovation, and new innovation methods

Technology governance aspects of standards focusing on both the socio-economic impact as well as the policies that guide it. These would include, but are not limited to:

- The national, regional, and global impacts of standards on industry, society, and economies
- The processes and organizations for creation and diffusion of standards, including the roles of organizations such as IEEE and IEEE-SA
- National and international policies and regulation for standards
- Standards and developing countries

The history of standardization, including, but not limited to:

- The cultures of different SDOs
- Standards education and its impact
- Corporate standards strategies
- The impact of open source on standards
- The impact of technology development and convergence on standards

Research-to-standards, including standards-oriented research, standards-related research, research on standards
Compatibility and interoperability, including testing methodologies and certification to standards
Tools and services related to any or all aspects of the standardization life cycle

Proposals are also solicited for Feature Topic issues of the Communications Standards Supplement.

Articles should be submitted to the IEEE *Communications Magazine* submissions site at

http://mc.manuscriptcentral.com/commag-ieee

Select "Standards Supplement" from the drop-down menu of submission options.

SEMANTICS FOR ANYTHING-AS-A-SERVICE

Nigel Davis

Kam Lam

Yuji Tochio

Scott Mansfield

A system is an assembly of interacting components where a component may itself be seen as a system. To interact, the components of a system must have a communication channel, a common protocol, some common semantics, and a shared representation of those semantics expressed at their respective ports. To achieve a shared representation requires some form of agreement and implementable standardization. To promote advancement, the standardization must embrace innovation as well as propagate domain insight via enabling constraints (like the shoulders of a road). Today there are many disconnected standardization and implementation activities. However, for successful broad advancement and improvement of solutions, it is vital that standardization is converged and the represented enabling constraints are applied to implementations.

Standardization of architectures and information models provide the necessary enabling constraints. To be successful, these information models and architectures need to be decoupled from the realization of the implementation. This allows for flexibility in deployment while accommodating consistency and coherence in the model. We all effectively do informal information modeling in our day to day work. Information modeling is essentially simply the process of identifying, defining, labeling, and recording, from some viewpoint, the key concepts (including things) in a domain of interest and the interrelationships between those key concepts. However, we rarely capture the model in any coherent fashion, and hence lose the opportunity to propagate key insights. Where aspects of the model are captured, the representation is often of a simple taxonomy or some other hierarchical form.

Even the most basic view of a domain of interest will yield a mesh of interrelated concepts where the mesh essentially embodies an intertwining of patterns and rule. Patterns such as hypergraph underpinning aspects of the ONF Common Information Model, component-system of the TM Forum FMO, control loop of the ONF architecture, spec/descriptor of the ETSI NFV, and encapsulation in many models, appear key to forming a robust canonical model of a domain. A canonical model, expressing the essence of the domain, in a language such as UML supported by tooling, is pruned and refactored to form views, again in UML, and interface representations in languages such as YANG or JSON schema. The canonical model is the axis for mapping between the views.

The canonical model will evolve with insights to become increasingly precise and stable. An efficient versatile implementation is achieved by taking advantage of the patterns in the domain expressed in the canonical model, using composition to apply specific properties (also in the information model) and specification/descriptors to constrain specific cases. The resulting balance of generalized recursive code and data-driving the recursions yields support for a continuum of subtly varied cases in the domain. One challenge with a canonical model is the opaqueness of the generalization. A domain specific graphical language that eases the representation of cases without losing the generalization will assist in representation of the concepts.

Naudts *et al.* describe the current standardization ecosystem and the challenges faced while trying to develop solutions in an agile, software-defined, and virtualized fashion. The article provides suggestions on how the various actors in the ecosystem can leverage each other's strengths to collaboratively, rapidly, and iteratively develop standards and solutions. Building on the ecosystem described above, Wac *et al.* provide an end-to-end service management framework that provides the ability to virtualize the proprietary interfaces of various network components and populate a comprehensive network information model. This framework provides the opportunity to coordinate previously unconnected silos and provide an environment that supports a holistic end-to-end view of any service across an operator's network. Manzalini and Crespi focus on the network edge, describing how virtualization of the devices at the edge will provide the flexibility needed to support innovative services. Garay *et al.* give an overview of the advancements being made in network function virtualization and service delivery. The article delves into the details of important standards in the service description space, and provides a discussion on ways to address the challenges in assembling a solution. Moberg and Vallin provide a very practical example of how to utilize a YANG-based service model along with a YANG-based device model to provide the encapsulation necessary to isolate the end-to-end service information from the device decomposition. This separation allows end-to-end services to be shielded from device differences and vise-versa.

This Feature Topic provides a view into how the standards world is using information and data modeling to describe the abstract representations and the detailed structured data needed by the emerging end-to-end service-based ecosystem. Driving the innovation to provide this new ecosystem is the creation of common semantics, consistent models, and interoperable APIs.

Biographies

NIGEL DAVIS is involved in standards/specification activities related to management-control, focused on architecture and information modeling, and has leadership, editor, and key contributor roles primarily in the TM Forum and ONF. He is a Ciena Fellow and a Distinguished Fellow of the TM Forum. He is also the co-chair of the Information Model Project (IMP) of the ONF. He holds several patents, and is recognized for his track record of innovation and influence dealing with network resource management-control.

KAM LAM is a CMTS at Alcatel-Lucent. He is active in many standards development organizations. He is the editor or co-editor of many standards specifications in the areas of network management architecture, FCAPS functions, and information models. He serves as the rapporteur of ITU-T Study Group 15 Question 14, which is responsible for transport resource management and control standardization. He is also the co-chair of the Information Model Project (IMP) of the ONF.

YUJI TOCHIO is a research engineer at Fujitsu Laboratories Limited focusing on transport network standardization. He is a document editor on a number of recommendations related to OAM, network equipment, and network management for packet transport networks, such as Ethernet and MPLS-TP, in ITU-T Study Group 15. He is also contributing to the standardization work for IETF and ONF related to information and data modeling.

SCOTT MANSFIELD (scott.mansfield@ericsson.com) is a principal engineer at Ericsson focusing on wireline network standardization. He serves as associate rapporteur of ITU-T Study Group 15 Question 14, concentrating on equipment management and information modeling. He is also a document editor in the ITU-T on a number of recommendations related to modeling (including generic, Ethernet, MPLS-TP, and OTN equipment recommendations).

Deploying SDN and NFV at the Speed of Innovation: Toward a New Bond Between Standards Development Organizations, Industry Fora, and Open-Source Software Projects

The authors examine the evolving roles of open-source software communities, industry fora, and standards development organizations, and places them in an NFV/SDN context. They sketch the differences between these roles and provide guidelines on how the interaction between them can turn into a mutually beneficial relationship that balances the conflicting goals of timely development on the one hand and technical excellence, openness, and fairness on the other, to reach their common goal of creating flexible and efficient telecommunications networks.

Bram Naudts, Wouter Tavernier, Sofie Verbrugge, Didier Colle, and Mario Pickavet

Abstract

Standards development organizations (SDOs) exist to assure the development of consensus-based, quality standards. These formal standards are needed in the telecommunications market to achieve functional interoperability. The standardization process takes years, and then a vendor still needs to implement the resulting standard in a product. This prevents service providers (SPs) who are willing to venture into new domains from doing so at a fast pace. With the development of software-defined networking (SDN) and network function virtualization (NFV), open-source technology is emerging as a new option in the telecommunications market. In contrast to SDOs, open-source software (OSS) communities create a product that may implicitly define a de-facto standard based on market consensus. Therefore, SPs are drawn to OSS, but they face technical, procedural, legal, and cultural challenges due to their lack of experience with open software development. The question therefore arises, how the interaction between OSS communities, SDOs, and industry fora (IF) can be organized to tackle these challenges.

This article examines the evolving roles of OSS communities, IF, and SDOs, and places them in an NFV/SDN context. It sketches the differences between these roles and provides guidelines on how the interaction between them can turn into a mutually beneficial relationship that balances the conflicting goals of timely development on the one hand and technical excellence, openness, and fairness on the other, to reach their common goal of creating flexible and efficient telecommunications networks.

The authors are with Ghent University – iMinds.

Introduction to an Ever-Evolving Telecommunications Market

Based on the number of subscribers and the multibillion dollar industry that surrounds it, we can resolutely state that fixed and mobile network architectures are very successful. These architectures are fit-for-purpose closed systems based on standardized interfaces. Every component performs specific functions, and each of the dozens of interfaces has a unique definition that has been standardized via an often long, formal, and consensus-based procedure. However, as customer demand evolves and new technologies emerge, the complex nature of these architectures starts to become a hindrance to sustainable growth. First, SPs will have to deal with higher capital expenditures and operational expenditures at a time when average revenue per user is decreasing [1]. As a result, some SPs will delay or refrain from investing further while those who do invest in new services or features face long time-to-market periods as they push an entire industry to standardize the newly developed features and then wait for vendors to actually implement them [1]. Furthermore, even when these new features are standardized and implemented, it may not be possible to realize them with existing equipment, as even though these can be controlled through standardized interfaces, there is little possibility to extend them through the use of open interfaces such as extensible application programming interfaces (APIs).

Therefore, SPs are looking for alternatives that can reduce the time-to-market and cost of new products and services. Three complementary, self-reinforcing drivers can bring them closer to that goal. First, the shift toward SDN offers the opportunity to learn from the experience of previous and ongoing management domain endeavors so as to be able to move to the next level of insight in realizing truly open and extensible interfaces. Additionally, there is an opportunity to migrate from multiple operations systems silos and many specialized operations functions in SP networks toward operations support systems that provide an overall solution architecture for operating services delivered across current and new technologies. Second, NFV can decrease the dependence on expensive network equipment vendor solutions, by replacing network functions with software implementations running on low-cost multi-purpose hardware. The advantages of NFV are most relevant for location independent network functions as better service scalability can be realized through sharing of resources. Third, by investing in OSS, a de-facto market-based standard can be created while the software is developed, and the time-to-market can be reduced by providing a workflow that allows for rapid deployment of software updates to very flexible hardware platforms. However, OSS development also faces challenges such as poor interoperability and high integration costs.

These de-facto market-based standards compete with the telecommunications market's long and often successful tradition of consensus-based standards that are developed within SDOs and

 0163-6804/16/$25.00

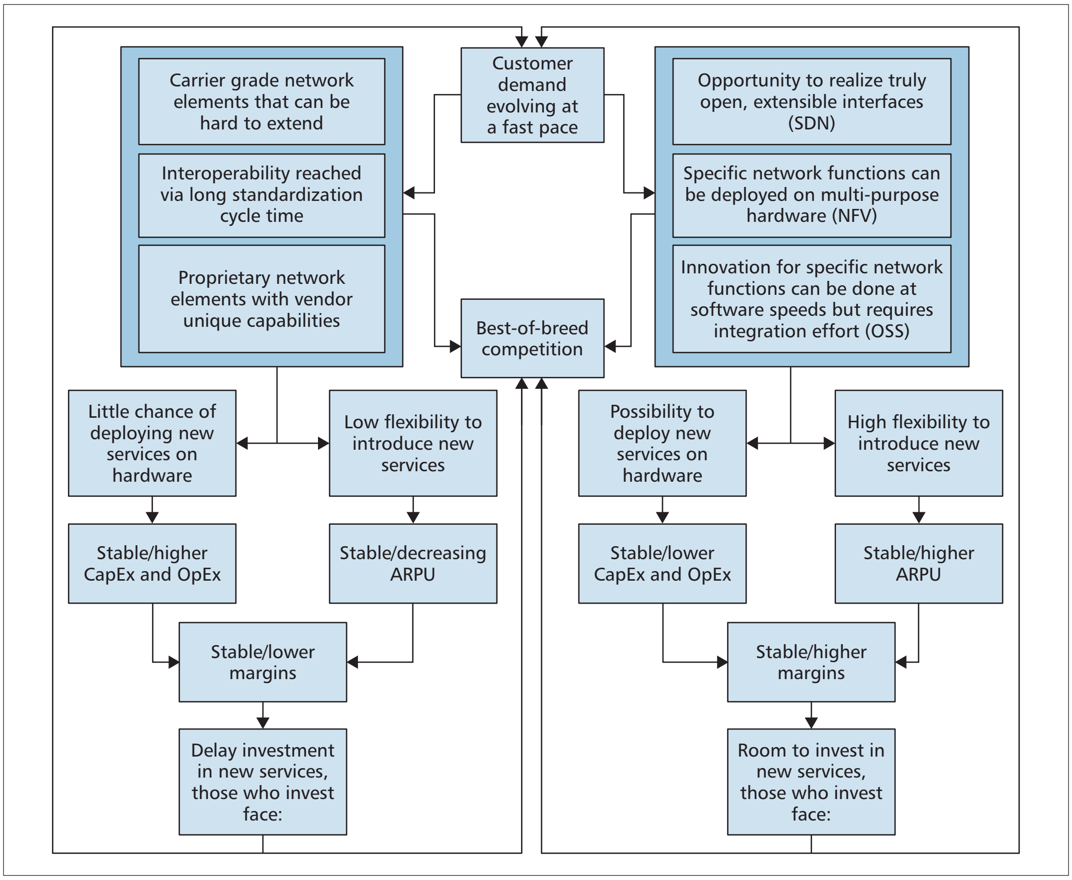

Figure 1. The operator's perspective: benefits and drawbacks of continuing with conventional methods versus the benefits and drawbacks of migrating to SDN/NFV and OSS.

The general trend toward open source, particularly open APIs, and the interest of SPs in these can be seen as a reaction to the lack of attention to the operational reality that SPs face day in and day out and the domination of vendors and academics in the decision-making processes within the SDOs.

IF. The general trend toward open source (OS), particularly open APIs, and the interest of SPs in these can be seen as a reaction to the lack of attention to the operational reality that SPs face day in and day out and the domination of vendors and academics in the decision-making processes within the SDOs [2]. Even though the strength of the carrier voice varies across SDOs/IF, some SDOs recognize this challenge. The Internet Engineering Task Force (IETF), for example, has a network working group that addresses the perceived gap between operators and the IETF whose objective is to help ensure that operational realities inform the development of key standards [3]. According to a survey conducted by that working group among network operators, the culture within the SDO was given as one of the four major obstacles to participation (time, money, and awareness are the other three) [3]. While the IETF is open to participation by anyone, almost half of the respondents avoid that organization because they do not feel their operator input is welcomed [3]. By not engaging, network operators write themselves out of the process, leading to the disparity that operators are expected to deploy technologies of which they do not even know that the standards are being developed. A recent counter example to the lack of involvement of SPs is the standardization process of NFV at the European Telecommunications Standards Institute (ETSI) which was initiated by Internet Service Providers (ISPs).

Without a doubt, SDOs are needed to produce high quality, relevant technical and engineering documents that create flexible and efficient telecommunications networks. However, standards become less relevant if they trail behind the pace of technology evolution. As such, if the trend toward OSS projects continues, the question arises how SDOs/IF can remain relevant in their role of enabling innovation. The goal of this article is to describe how the interaction between OSS communities, SDOs, and IF can be improved. The remainder of this article is structured as follows. After introducing an overarching SDN/NFV architecture and describing the most relevant roles in the ecosystem, we discuss the differences between the market-based standards formed in OSS communities and the consensus-based standards developed by SDOs/IF. We then formulate guidelines on how these can work together to reach a mutually beneficial relationship.

SDN/NFV Architecture Overview and Main Ecosystem Roles

This section sketches the main functional components and layers in the control architecture of a modern telecom network supporting NFV and links them to the main ecosystem roles, in order to provide the necessary context for the discussion on the interaction between OSS communities, SDOs, and IF. The International Telecommunications Union Telecommunication Standardization Sector (ITU-T) describes

Modern network architectures are structured into multiple functional layers of smaller components. This modular approach reduces complexity, enhances component reusability, and enables multiple migration paths toward future architectures.

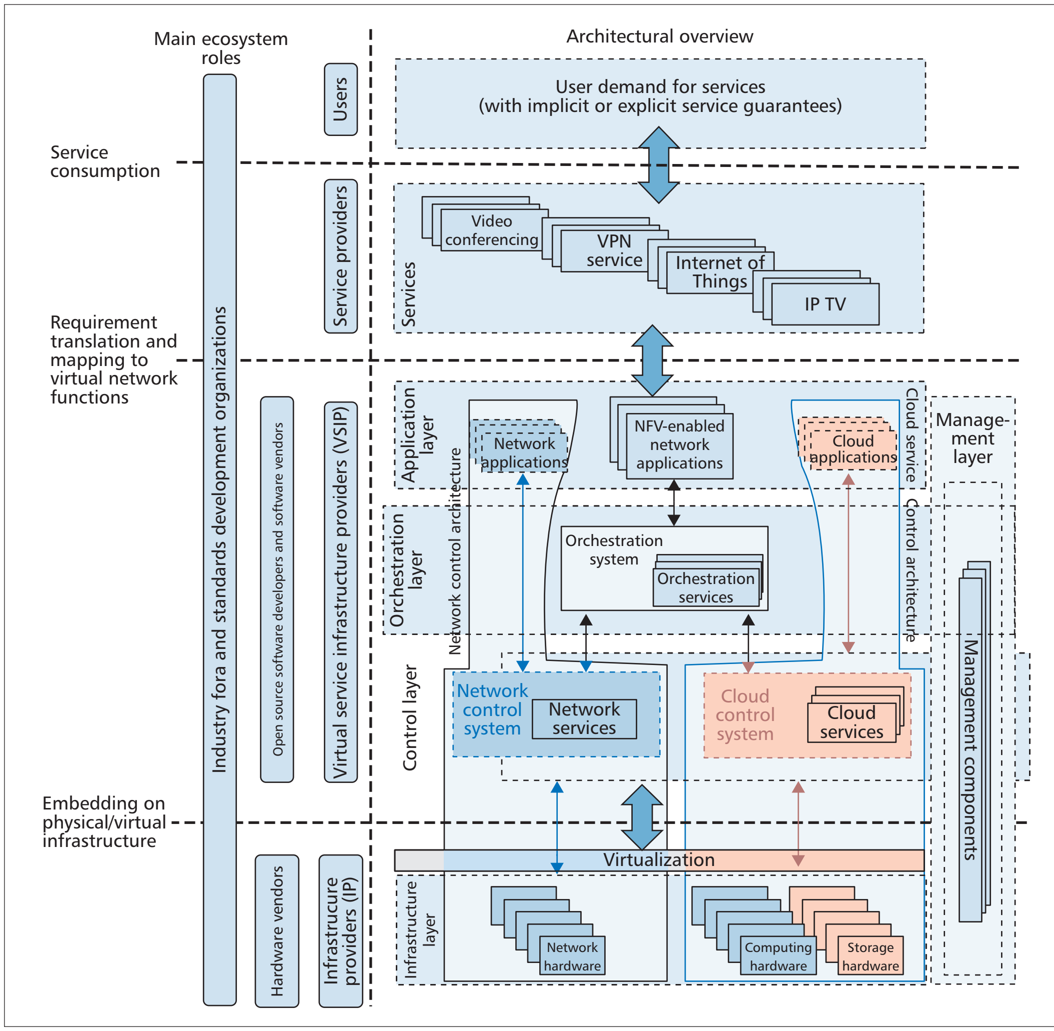

Figure 2. Architectural overview of network- and cloud control platforms.

the requirements to reach carrier grade service for an independent, scalable control plane in future, packet-based networks [4]. The requirements include reachability, scalability, flexibility, reliability, manageability, service, security, interworking, routing, and forwarding.

Modern network architectures are structured into multiple functional layers of smaller components. This modular approach reduces complexity, enhances component reusability, and enables multiple migration paths toward future architectures. Recent softwarization and virtualization tendencies have only further accumulated the decomposition of functional components and layers within architectures. By decoupling the forwarding from control functionality, SDN transforms previously monolithic switches/routers into multiple independent components. Server and network virtualization mechanisms in turn introduce additional functional splits that isolate the data plane functionality of its underlying hardware platform (interested readers should consider [5] and references therein). When network functions (NFs) such as firewalls (FW) or deep packet inspectors (DPI) are decoupled from their underlying hardware platform, and are realized in software that might be executed by commercial off-the-shelf (COTS) hardware, we are speaking about NFV.

SDN and NFV are fully complementary paradigms [6]. SDN is centered on the software-based control of network resources to provide services, while NFV focuses on the creation and life cycle support of some classes of service resources, i.e. virtualized NFs. Indeed, a software-based control architecture might be used to provide network services that consist of either traditional network hardware, virtualized network resources, or combinations of both. In fact, such a combination might be conceived by considering two existing control areas:

- The (software-driven[1]) control of communication networks.
- The control of cloud (service) platforms.

Both control architectures are depicted in the architectural overview of Fig. 2, which is based on [7].

The first (in blue, left) is in charge of controlling the network of switching and routing equipment; the second (in orange, right) is in charge of creating and exposing cloud networks, i.e. a network of reusable computing and storage servers for the purpose of building web services, for example. The control architecture of both domains follows a roughly similar three-layered approach, as depicted in Fig. 2. At the lowest layer, infrastructure resources form the physical foundation on top of which services are provided.

[1] In the context of this article we focus on SDN-controlled networks, although traditional distributed routing protocols could also be considered as the control layer of communication networks.

Communication networks rely on network hardware such as switches and routers; cloud infrastructures rely on (interconnected) computing and storage hardware (servers). A second layer, the control layer, interconnects the components of the infrastructure layer via their north-bound interface (e.g. OpenFlow for network control) in order to provide control-level services such as topology management or datastore services.

The virtualization layer enables a decoupling of functionality from its underlying hardware. At the computing device level, virtualization enables one device to be segmented in multiple logical devices. At the network level network virtualization enables isolation of network resources across different network hardware devices into virtual networks or slices.

At the highest layer, components of the application layer build further on control layer services to program client applications. A traffic engineering application might be defined on top of the SDN-control layer, while a Hadoop cluster might be an application on top of the cloud platform. The orchestration system has a complete view on available networking as well as on computing and storage resources, and is used for services that require a combination of these resources. The orchestration components are able to make an informed decision on which infrastructure should be used. The provisioning process itself can then be further delegated to the already existing network and cloud control system. Orthogonal to the horizontal layers, management functionality might be required to configure any of the components at the infrastructure, control, or application layer, for example to ensure policies or security-related options.

A number of stakeholders are involved in the realization of this SDN/NFV-driven architecture. We discuss stakeholder responsibilities and interactions in the remainder of this section. On the left side of Fig. 2, the most relevant ecosystem roles are represented. These roles are accomplished by the actors that actively participate in the exchange of value. Most actors will perform more than one role at the same time. For example, traditional ISPs fulfill the role of infrastructure provider, virtual service infrastructure provider, and service provider.

Users: Users, i.e. end/enterprise users, retail, or over-the-top providers, request, and consume a diverse range of services. In general, users have no strong opinion about how the service is delivered as long as their quality of experience expectations are satisfied.

Service Providers (SPs): SPs accommodate the service demand from users by offering one or multiple services, including over-the-top service and X-play services (e.g. triple play). The service provider realizes the offered services on a (virtualized) infrastructure via the deployment of virtualized network functions (VNFs).

Virtual Service Infrastructure Providers (VSIPs): VSIPs [8] deliver virtual service infrastructure to SPs, meeting particular service level requirements by combining physical network and cloud resources into service infrastructure meeting particular SLA requirements implemented through NFV-enabled network applications. These network applications might involve resources (or network functions) that are either implemented in traditional network hardware, or as virtualized NFs. These are the result of an orchestration system that interacts with the network control system as well as the cloud control system.

Infrastructure Providers (InPs): InPs own and maintain the physical infrastructure and run the virtualization environments. By virtualizing the infrastructure, they open up their resources to remote parties for deploying VNFs. The reusable physical resources comprise all possible resource options (computing, storage, and networking), and they span the entire service delivery chain from the end-user gateway and set-top-box over the access, aggregation, and core network up to the cloud.

Hardware Vendors: Hardware vendors provide the physical devices that are deployed by the infrastructure providers. The shift away from specialized equipment toward reusable, industry-standard high-volume servers, switches, and storage devices can reduce the total costs of infrastructure providers as they cost less than manufacturer-designed hardware and increase flexibility. The hardware must provide an interface toward the controller systems.

Software Vendors: Software vendors, including OSS developers, deliver the implementation of the logic that is used to optimally deploy the services on the physical infrastructure. Today a patchwork of specialized software products exists to realize that functionality. The most relevant software for the SDN/NFV architecture are those that focus on the following:

- The acceleration of packet processing on commodity hardware.
- Virtual machine technologies and software container-based technologies.
- Network virtualization software for virtualizing SDNs.
- SDN and cloud control software.
- Software for the orchestration of VNFs.
- Software implementations of VNFs.
- Software for monitoring, management, automated roll-out, configuration, and specification of VNFs.

For each of these, OSS communities have developed or are developing viable alternatives to proprietary software. We do not list all of these OSS projects due to space constraints (interested readers should consider [10] and the references therein).

Standards Development Organizations and Industry Fora: The networking industry today is very much standards-driven to make a product or service safe (safety standards) and interoperable (interface standards), while making the industry as a whole more efficient. The purpose of SDOs/IF such as ITU-T, ETSI, the Open Networking Foundation (ONF), IETF, the TM Forum, and the Metro Ethernet Forum (MEF) is to standardize the concepts that emerge in the ecosystem via coordination of the different actors in the development of new technical standards, as well as the revision and amending of existing standards when needed. Participants from across the ecosystem contribute to the development of these standards.

Next, we look into the details of the roles of

Software vendors, including OSS developers, deliver the implementation of the logic that is used to optimally deploy the services on the physical infrastructure. Today a patchwork of specialized software products exists to realize that functionality.

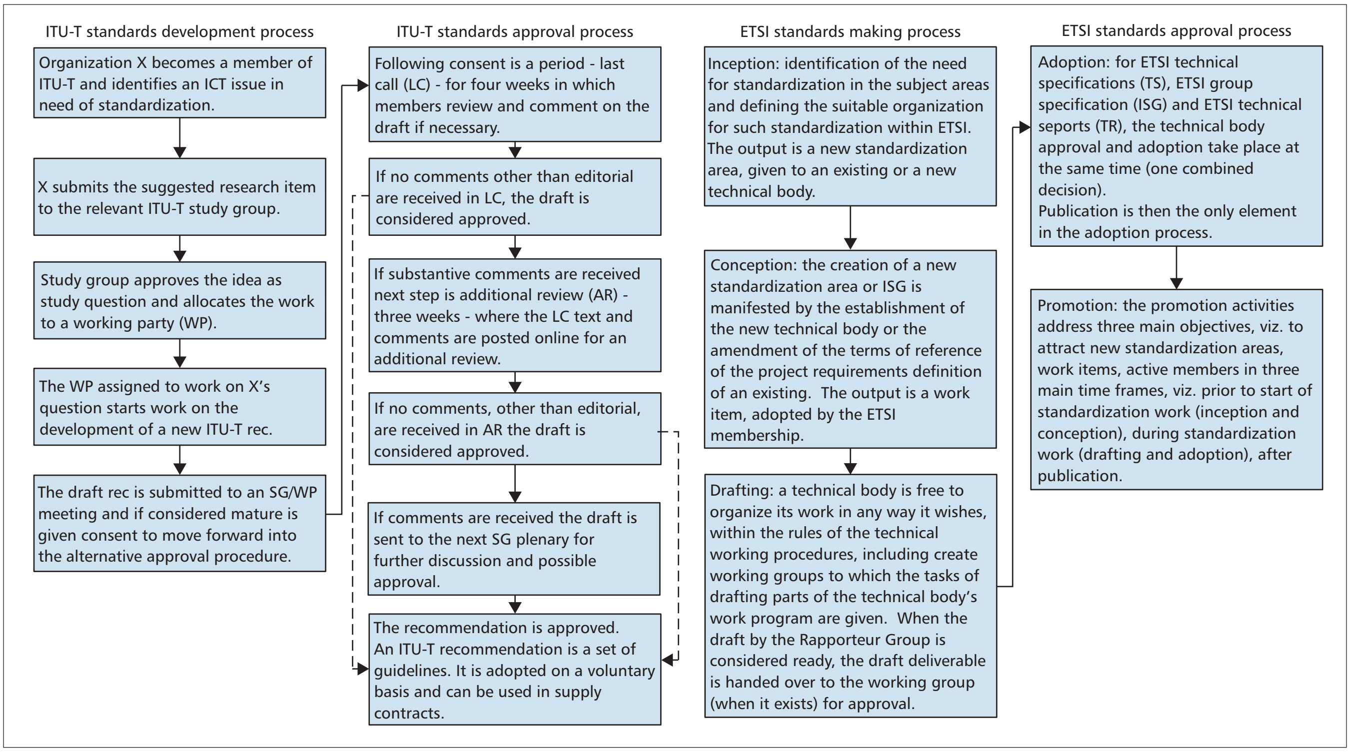

Figure 3. Overview of the standards development process and approval process of ITU-T and ETSI.

OSS communities on the one hand, and SDOs and industry fora on the other, in the development of standards.

Standards Developed by SDOs vs De-Facto Standards as a Result of the Work Done in OSS Communities

ETSI defines a standard as a document, established by consensus and approved by a recognized body, that provides, for common and repeated use, rules, guidelines, or characteristics for activities or their results, aimed at achievement of the optimum degree of order in a given context [9]. In general, five steps can be recognized in the standards process:
- Identification of the need.
- Assignment to the relevant body/group.
- Drafting and submission of the standard.
- Approval.
- Adoption and distribution.

The specific implementation differs between SDOs/IF, as illustrated in Fig. 3 for ITU-T and ETSI.

In practice, this requires significant time and effort due to:
- The difficulty of creating specifications of high technical quality.
- .The need to consider the interests of all of the affected parties.
- .The importance of establishing widespread community consensus.
- .The difficulty of evaluating the utility of a particular specification for the Internet community [11].

This is in sharp contrast with today's rapid development of networking technology, which demands the timely development of standards.

An OSS project, on the other hand, must deliver a working product. During the development, a de-facto market-based standard is created (development and standardization are executed as parallel processes). The agile development model, which is tied closely together with OSS projects, results in smaller incremental releases with each release, building on previous functionality. This approach takes into account that user demand is dynamic and that plans are short-lived. The OSS community decides on a way to implement a feature and, once it is included in the OSS project, it can be deployed at once. As a result, the opportunity exists to reduce the time-to-market. Similarly, SDOs/IF could apply an agile development approach in specification development to reduce their cycle time. For example, the authors in [12] state that the cycle time of a paper standard compared to an OSS project can be shortened by at least a factor of two.

SDOs focus on the design of norms or requirements of technical systems to achieve a technical goal that can only be met when multiple partners agree, and preferably subsequently adopt the proposed norm. Most SDOs follow a rigid specification mechanism, which once published, can only be corrected, changed, or extended in rather discrete steps following a rigorous process of validation and agreement. This makes SDO-based standards slow to adapt to a changing environment or problem statement. On the contrary, OSS projects are able to almost continuously adapt and integrate new code contributions driven by contributors in order to solve important current issues. While OSS communities can contribute to the goals of operators to reduce the costs of services and time-to-market, it should also be clear that the number of failed or dormant OSS projects is also notable [13]. Operators that want to contribute to OSS com-

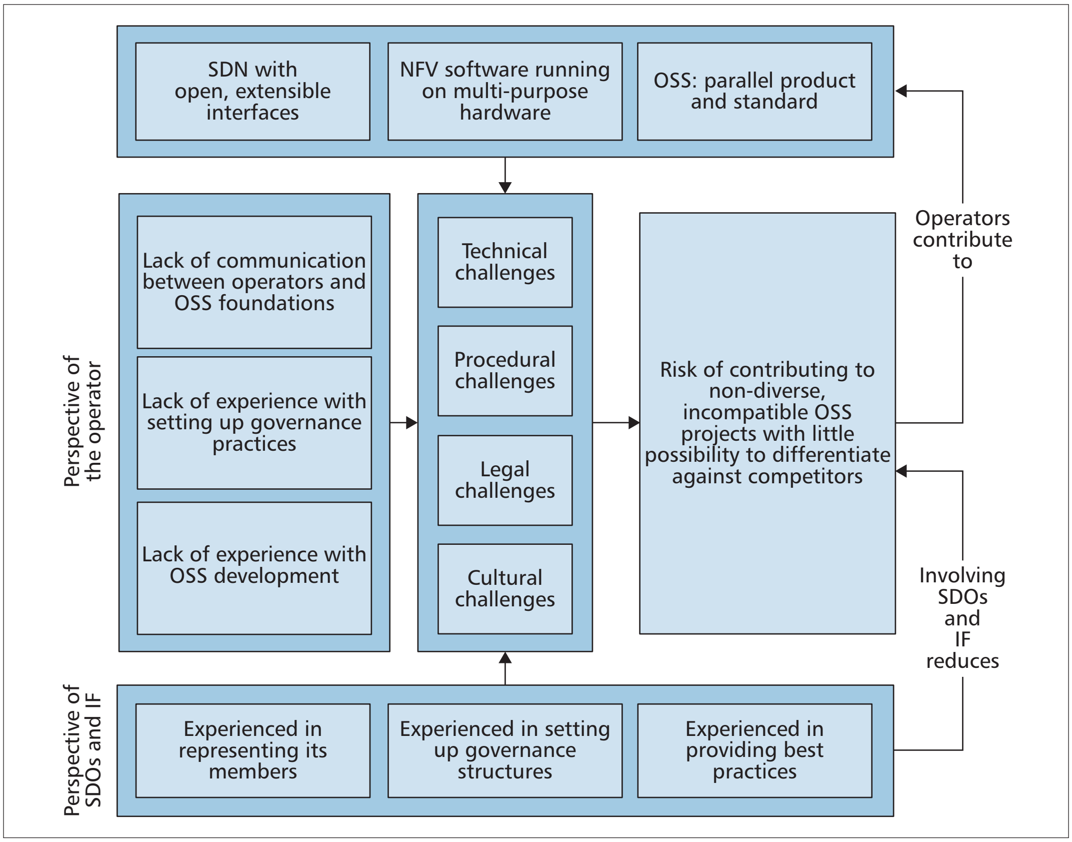

Figure 4. Interaction between operators, SDOs, IF and OSS foundations.

> Operators typically work with product managers, while OSS communities focus on use cases and feature sets. Changing a company's culture is not a simple challenge, as internal resistance from people who fear losing their job can be severe when not properly managed.

munities must therefore also overcome a variety of challenges [14]: technical, procedural, legal, and cultural.

Technical: OSS development can be disorganized as developers work on the parts that interest them most. Less tempting, but necessary, parts such as writing code documentation, automated tests, and manuals may as such receive less attention Also, to overcome fragmentation, OSS projects need to be able to interconnect and fit into a larger architecture. These technical challenges, while being pertinent to reach success, may receive less attention due to community diversity.

Procedural: From a procedural perspective, OSS cannot prevent companies from dominating a project and pushing though their own approach. This is a result of the lack of governance structure that ensures quality in the development and integration as well as the procedures for its assessment.

Legal: The choice of license may affect interoperability and the possibility for SPs and vendors to differentiate themselves. Permissive licenses, such as the Apache License Version 2, do not impose special conditions on the second redistribution, while strong licenses impose conditions in the event of wanting to redistribute the software. These conditions are intended to ensuring compliance with the license's conditions following the first distribution. Under the General Public License (GPL) of the GNU project, for example, it is only possible to redistribute code licensed under a compatible license, while under the Apache License Version 2, a project may be forced to develop proprietary extensions based on the material.

Cultural: As the center of value shifts from hardware toward software, the operator's culture and skillset must evolve as well (interested readers should consider [15] and the references therein). Operators typically work with product managers, while OSS communities focus on use cases and feature sets. Changing a company's culture is not a simple challenge, as internal resistance from people who fear losing their job can be severe when not properly managed.

To summarize this section, we wish to point at the conflicting goals of timely development of products and services on the one hand, and technical excellence, openness, and fairness on the other. Moving in one direction often leads to compromising in the other. Therefore, the next section focuses on how SDOs, IF, and OSS communities can work together to balance these conflicting goals and reach the common goal of creating flexible and efficient telecommunications networks.

Guidelines for Improving Interaction Between OSS Communities, SDOs, and IF

Both SDOs and IF should engage with the OSS community to tackle the technical, procedural, legal, and cultural challenges that operators face in contributing to OSS. The causes of these challenges can be backtracked to a lack of communication, governance practices, and inexperience with OSS development.

The fundamental reason behind the existence of SDOs/IF is to avoid miscommunication and to establish impartial third-party governance practices. The competencies that SDOs/IF have

In parallel with the standards development process, code should be developed to support extensibility and modularity, and allow agile workflows for each of the modules independently. The NFV Proof of Concept-zone is an example of how function demonstration can be encouraged.

developed by performing these functions can provide an answer to the challenges that operators face when contributing to OSS. However, without change, the relevance of the interaction between SDOs, IF, and OSS will remain negligible. Attempts to bridge the gap via an alternative SDO model are therefore emerging. The ONF is an early example, which is dedicated to the promotion and adoption of SDN through open standards development. Initially established to promote the OpenFlow protocol via market development, the ONF now covers a broad range of specifications activities that encompass SDN architecture, the open common information model of network resources, a data model, and API development (including NETCONF, YANG, etc.). For example, to enable SDN control and network programmability, and allow SDN to be applied to a wide range of network resources, the ONF has a major effort to establish a consistent description of network resource functionality, capabilities, and flexibility. This resource description is provided by an information model that is independent of implementation details (including the protocol), providing the foundation with the derivation of a coherent suite of interface protocol-specific data models. Promoting a common industry-wide open model has been an informal collaboration among the ONF, ITU-T SG15, and the TM Forum. Between this, data model/API development, associated OS projects, and usage of OS tooling, ONF links these areas together in creating a bridge between "paper specifications" and "software development." Open Source SDN (OSSDN) is one example of how the ONF supports and sponsors OSS development by supplying people, monetary support for the maintenance and development of the community, and the hiring of a community manager. The Atrium project, which integrates OSS components and tries to make it easier for network operators to deploy SDN, is a direct outcome of that support. Another example is the Open Platform for NFV (OPNFV), a project operating under the Linux Foundation in close collaboration with ETSI's NFV ISG (among others), which has as its purpose the establishment of an integrated, open-source reference platform that uses the open-source NFV building blocks that already exist. A final example is the ETSI NFV Proof of Concept-zone, which promotes multi-vendor open ecosystems integrating components from different players.

To return to the goal of this article, we conclude the article by formulating a set of guidelines, based on lessons learned from alternative SDO models, which provide an outline toward what SDOs/IF can do to tackle the previously described challenges.

•SDOs/IF and OSS communities should establish open communication to reach more engagement in compatible projects. As an example, OpenMANO is an open source project (initiated by Telefónica) that provides a practical implementation of the reference architecture for management & orchestration under standardization at ETSI's NFV ISG (NFV MANO).

•SDOs/IF should emphasize software development and function demonstration more in its culture and structure by aligning their processes with the OSS development practices. In parallel with the standards development process, code should be developed to support extensibility and modularity, and allow agile workflows (e.g. hackathons) for each of the modules independently. The NFV Proof of Concept-zone is an example of how function demonstration can be encouraged.

•SDOs/IF should help OSS communities with the development of governance structures to guarantee technical excellence, openness, and fairness among the contributors to OSS projects. First, SDOs/IF should provide internal project governance in terms of developing the practices as well as the procedures that guarantee an effective development, integration, release, maintenance, and update process. and help in setting up the essential legal, business, management, and strategic processes. Second, SDOs/IF should offer cross-project governance to avoid:

- Unintentional competition between OSS projects that aim for the same goal (assuring project diversity).
- OSS projects that each deliver part of an overall solution, and which cannot be used together (assuring interoperability).

This is particularly challenging as these governance structures and processes differ among SDOs. In fact, it would also require an SDO/IF requirement upon an overall (modular) management/control architecture for software development in the domain of interest, with supporting guidelines, processes, and common open source tooling. This would assure consistency when diverse teams work independently on a part of the solution (e.g., technology-/application-/etc.-specification modules).

•SDOs/IF should guide operators, which are typically not so familiar with the world of OSS, among the plethora of OSS projects, and help them find the projects that best fit their needs and are worth contributing to. Examples are the Atrium and OPNFV projects, which integrate several OSS projects to speed up adoption.

•SDOs/IF should gather end users together, facilitate their discussions, and help operators with the definition of use cases and feature sets in a way that is implementable by an OSS project. As an example, OPNFV helps operators understand how to articulate their use cases as functional gaps in OSS projects.

•SDOs/IF should provide best practices in OSS development via training and learning materials, for example, by providing advice on best practices with regard to OSS licenses. SDOs/IF can help to make OSS credible for both operators and vendors (by preserving their ability to differentiate). For instance, OPNFV is licensed under an Apache 2.0 license, which explicitly grants patent rights where necessary to operate, modify, and distribute the software.

•SDOs/IF should overlook the integration of OSS projects and point toward development gaps while establishing and maintaining communication with other SDOs and IF. An example is the TM forum Catalyst proof of concepts, which bring together service providers and suppliers to work collaboratively. Another initiative, started by MEF and the TM Forum, is the UNITE program, to ensure a more open and rapid alignment of SDO work.

Summary

In this aticle we argued that margin pressure and the lack of possibilities for SPs to introduce new services has spurred their interest in:

- Emerging technologies such as SDN and NFV that provide an opportunity to reduce cost and increase flexibility.
- Other collaboration models such as OSS projects that can reduce the time-to-market.

By linking the most relevant ecosystem roles on the proposed overarching SDN/NFV architecture, we illustrated the general trend toward OS, particularly extensible APIs, in the SDN/NFV network space. Next, we focused on how these evolutions are changing the role of SDOs and IF, and how the OSS development methods affect how new standards are proposed, developed, and implemented. On one side of the spectrum, consensus-based standards developed by traditional SDOs tend to have a longer cycle time than the pace at which technology evolves. On the other side, OSS projects lead to a de-facto market-based consensus in a shorter cycle time. As such, SDOs may gradually lose their relevance in enabling innovation, and operators might turn to OSS communities to realize innovation. However, SPs that wish to contribute to OSS communities face technical, procedural, legal, and cultural challenges. We argue that the fundamental reason behind the existence of SDOs/IF is to resolve these challenges. Based on lessons learned from the interaction that is starting to happen between SDOs, IF, and OSS communities, we formulated a list of guidelines to improve interaction between both worlds and improve the relevance of SDOs/IF in innovation and increase the technical excellence, openness, and fairness of OSS projects.

Acknowledgements

We are grateful to the anonymous reviewers for their constructive comments and suggestions, which helped us to hone the core message of this article. This work has been partially funded by the European Commission under the 7th Framework research program project UNIFY.

References

[1] K. Pentikousis *et al.*, "Mobileflow: Toward Software-Defined Mobile Networks," *IEEE Commun. Mag.*, vol. 51, no. 7, July 2013, pp. 44–53.

[2] R. Bush, "Into the Future with the Internet Vendor Task Force: A Very Curmudgeonly View of Testing Spaghetti: A Wall's Point Of View," *ACM SIGCOMM Computer Commun. Rev.*, vol. 35, no. 5, Oct. 2015, pp. 67–68.

[3] C. Grundemann and J. Zorz, "Operators and the IETF," IETF Secretariat, Apr. 2015, available at https://datatracker.ietf.org/doc/draft-opsawg-operators-ietf/.

[4] ITU-T Recommendation Y2.621, "Requirements for an Independent, Scalable Control Plan in Future, Packet-based Networks," Aug. 2011, available at https://www.itu.int/rec/T-REC-Y.2621-201108-I/en.

[5] S. Sezer *et al.*, "Are We Ready for SDN? Implementation Challenges for Software-Defined Networks," *IEEE Commun. Mag.*, vol. 51, no. 7, July 2013, pp. 36–43.

[6] ONF, "TR-518 Relationship of SDN and NFV," Oct. 2015, available at: https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/onf2015.310_Architectural_comparison.08-2.pdf.

[7] R. Szabo, "D2.2 Final Architecture," Nov. 2014, available at: https://www.fp7-unify.eu/files/fp7-unify-eu-docs/Results/Deliverables/UNIFY%20Deliverable%202.2%20Final%20Architecture.pdf.

[8] S. Latré *et al.*, "The Fluid Internet: Service Centric Management of Virtualized Future Internet," *IEEE Commun. Mag.*, vol. 52, no. 1, Jan. 2014, pp. 140–48.

[9] E. Soederstroem, "Standardising the Business Vocabulary of Standards" *Proc. 2002 ACM Symp. Applied Computing*, Mar. 2002, pp. 1048–52.

[10] W. Tavernier *et al.*, "Can Open-Source Projects (Re)-Shape the SDN/NFV-Driven Telecommunications Market?," *it-Information Technology*, vol. 57, no.5, Oct. 2015, pp. 267-276

[11] S. Bradner, "The Internet Standards Process – Revision 3" RFC 2026, IETF, Oct. 1996

[12] D. Ward, "Open Standards, Open Source, Open Loop," *IETF Journal*, Mar. 2015.

[13] S. Lee, "Measuring Open Source Software Success," *Omega*, vol. 37, no.2, June 2009, pp. 426–38.

[14] T. Kredil *et al.*, "Open Source: How to Speed NFV and SDN Deployment," *TM Forum Quick Insigths*, July 2015.

[15] C. Meirosu *et al.*, "DevOps for Software-Defined Telecom Infrastructures, draft-unify-nfvrg-devops-03.txt, IETF Secretariat, Work in Progress, Oct. 2015, available at: https://datatracker.ietf.org/doc/draft-unify-nfvrg-devops/.

Biographies

Bram Naudts (bram.naudts@intec.ugent.be) received his M.S. degree in applied economics in 2011 from Ghent University (Belgium), where he is currently pursuing a Ph.D. in techno-economics in the Department of Information Technology (INTEC). His research interests include techno-economic evaluation of communication network architectures and services, with a focus on software-defined networking and network function virtualization in research projects such as SPARC and UNIFY.

Wouter Tavernier (wouter.tavernier@intec.ugent.be) received a M.S. in computer science in 2002, and a Ph.D. degree in computer science engineering in 2012, both from Ghent University. He joined the IBCN group of Ghent University in 2006 to research future Internet topics. His research focus is on software-defined networking, network function virtualization, and service orchestration in the context of European research projects such as TIGER, ECODE, EULER, UNIFY, and SONATA.

Sofie Verbrugge (sofie.verbrugge@intec.ugent.be) received M.S. and Ph.D. degrees from Ghent University, Belgium. Since 2008 she has been working as a researcher affiliated with iMinds, a research institute to stimulate ICT innovation in Flanders, where she coordinates the techno-economic research within the IBCN group. In October 2014 she was appointed an associate professor at Ghent University. Her main research interests include cost modeling and telecom service and network deployment planning.

Didier Colle (didier.colle@intec.ugent.be) is a full professor at Ghent University. He received a Ph.D. degree in 2002 and a M.S. degree in electrotechnical engineering in 1997 from the same university. He is group leader in the iMinds Internet Technologies Department. He is co-responsible for the research cluster on network modelling, design and evaluation (NetMoDeL). This research cluster deals with fixed Internet architectures and optical networks, green-ICT, design of network algorithms, and techno-economic studies.

Mario Pickavet (mario.pickavet@intec.ugent.be) is a full professor at Ghent University, where he is teaching courses on discrete mathematics and network modeling. He is co-leading the research cluster on network modeling, design and evaluation (NetMoDeL). His main research interests are Internet architectures and optical networks, green ICT, and design of network algorithms. He is involved in several European and national research projects. He has published approximately 300 international publications in journals and in conference proceedings.

SDOs/IF should gather end users together, facilitate their discussions, and help operators with the definition of use cases and feature sets in a way that is implementable by an OSS project. As an example, OPNFV helps operators understand how to articulate their use cases as functional gaps in OSS projects.

e2eUberIM: End-to-end Service Management Framework for Anything-as-a-Service

To leverage existing efforts to provide new, high-quality services, the authors propose an innovative end-to-end service management framework called e2eUberIM. This framework contains a scalable, flexible and volatile information model (UberIM) and a real-time, "organic" communication and storage processes (e2eUber). It leverages the proprietary interfaces of different vendors, enabling the resulting enhanced operations system environment to automatically and in real-time optimize the operational overheads while maximizing the user experience, and quickly field new innovative, composed "any-services."

Katarzyna Wac, Mark Cummings, and Jayanta Dey

Katarzyna Wac is with the University of Copenhagen.

Mark Cummings is with Orchestral Networks .

Jayanta Dey is with Wipro's Media and Telecom & Product Engineering Services.

Some of the early material discussed in this paper was developed in the context of the Next Generation Converged Operations Requirements (NGCOR) Project in Next Generation Mobile Networks (NGMN) led by Mr. Klaus Martiny of Deutsche Telekom. Some of this material was developed in discussions with Caroline Chappell, Principal Analyst, Cloud and NFV, Heavy Reading.

COMMUNICATIONS STANDARDS

Abstract

Telecommunication network operators (TNOs) are seeing an exponential growth in mobile and fixed traffic, and on a growing scale need to manage (in some cases even manually) operational data silos. These silos necessitate a high degree of expensive operation that limits the ability to rapidly compose new innovative services, leads to high and exponentially growing integration expenses, and impairs the user experience for deployed services. All this is reducing long-term TNO profits. To meet this challenge, solutions employing network function virtualization (NFV) and software defined networking (SDN) have been proposed. However, different vendors of these solutions continue to deploy proprietary interfaces that are not interoperable with the new technology systems nor with legacy systems, thus further increasing the number of operational silos. To leverage existing efforts to provide new, high-quality services, we propose an innovative end-to-end service management framework called e2eUberIM. This framework contains a scalable, flexible, and volatile information model (UberIM) and a real-time, "organic" communication and storage process (e2eUber). It leverages the proprietary interfaces of different vendors, enabling the resulting enhanced operations system environment to automatically and in real-time optimize operational overhead while maximizing the user experience, and quickly field new innovative, composed "any-services" (XaaS). We document the e2eUberIM requirements and its design choices, as well as provide a set of implications for its successful implementation in operational, live TNO networks.

Introduction

Network function virtualization coupled with automated management is rapidly evolving into a game changer for next generation communication service providers due to the changing technology environment; growing market demand flexible, agile, and cost efficient operation of existing services; and rapid development and deployment of new innovative services. However, management information silos, hampering flexibility and resulting in operational inefficiency, are eroding the possible gains. It is estimated that even small telecommunication network operators (TNOs), providing wireline and cellular Internet services to specific regions, could be paying upward of $400M of what we refer to as *operations and "integration tax"* every year [1]. For larger, e.g. nationwide TNOs with over 200 different inventory management systems, the costs could be three times that.

TNOs, by their very nature, have long life system elements spread over large geographies with fragmented business support systems (BSS), network elements (NE, e.g. base stations, switches, routers, network core components), element management systems (EMS, controlling the NEs from a specific vendor), and information systems [2] interfacing with operations support systems (OSS) (Fig. 1). Data flows at the user information level; however, at the operational level, there is a profusion of non-compatible interfaces stemming from organizational divisions, histories of mergers, but most importantly, product strategies of vendors. Although the TNOs contribute to and advocate the development of standards, the vendors seek to have at least partially proprietary interfaces enabling them to continually differentiate their products, innovate, and maintain margins. As a result, the number of non-compatible systems is increasing [2, 3] and is being maintained by large numbers of specialized engineers conducting highly manual operational tasks.

As networks are growing in complexity and volatility, the "integration tax" is growing exponentially (Fig. 2). This, combined with subscriber saturation, results in falling TNO profits.

The network suffers problems in quality of service (QoS), which in turn results in a degraded quality of experience (QoE) for the end-users. The QoS for a TNO's service is *"a collective effect of service performances, which determines the (objective) degree of satisfaction of a user"* [4], and it embraces service dependability, its speed and accuracy.

QoE is *"the overall acceptability of service, as perceived subjectively (i.e. qualitatively) by the user"* [4], for the delivered service. There is neither an exact, conclusive definition of QoE nor factors influencing it, especially for mobile services [5]. For example, for VoIP service, a response-time of at most 150 ms and data loss of 3 percent satisfies the QoS requirements of its user [4], and they may be happy or not with the service, depending on how it was experienced in a given context and service needs (e.g. urgent call vs. just a chat).

To attempt to overcome these QoS/QoE challenges, resources are over-provisioned in the network infrastructure, incurring increasing costs. Current solutions limit the speed of development and deployment of new innovative services, leading to diversion of revenue to over the top (OTT) providers [6].

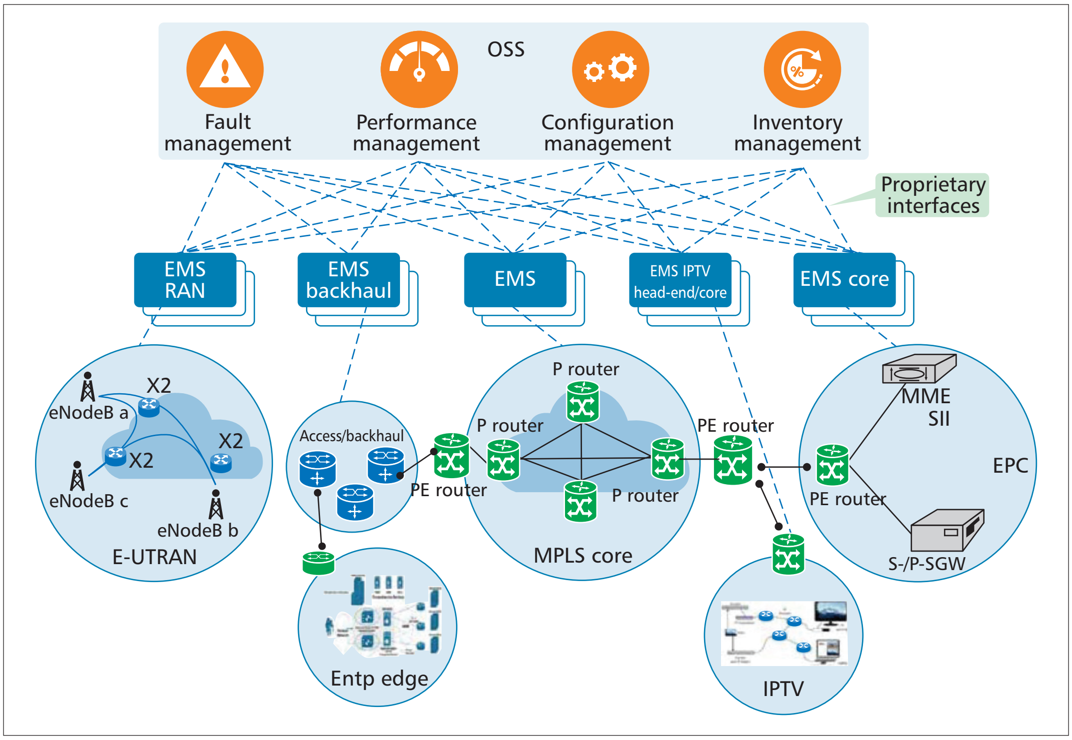

Figure 1. TNO's current co-existing management and information systems.

To manage the existing EMS and OSS complexity, auxiliary solutions have been proposed, e.g. based on network function virtualization (NFV) or software defined networking (SDN) [7]. NFV virtualizes network functions previously carried out by dedicated hardware. The current most active industry-based initiative, led by the European Telecommunications Standards Institute (ETSI NVF ISG), focuses on carrier grade cloud-based solutions for TNO data centers. SDN focuses on moving the network control plane functionality from special-purpose hardware to software. With service chaining, a chain of network (virtual) functions can be composed, for the purpose of collectively ensuring that a provided service includes all the features and quality assurances defined in an associated service level agreement (SLA).

These NFV and SDN solutions promise cost savings and increases in flexibility; however, given the long life of installed legacy systems, the transition may take decades. In this transition period, NFV and SDN will produce more incompatible operations data silos, increasing the integration tax.

To leverage the existing efforts to provide new, high-quality services, we propose an innovative end-to-end service management framework called *e2eUberIM* (Fig. 3). The framework contains a scalable, flexible, and dynamic information model (*UberIM*) and real-time, "organic" communication and storage processes (*e2eUber*). It leverages the proprietary interfaces of different vendors, enabling the resulting enhanced operations system environment to automatically and in real-time optimize operational overhead, while maximizing the user QoS/QoE, and quickly field new innovative, composed "any-services" (XaaS).

Related Work

Description languages in use today include Extensible Markup Language (XML), JavaScript Object Notation (JSON), Simple Network Management Protocol (SNMP), Management Information Base (MIB) and Object Identifiers (OID), Common Object Request Broker Architecture (CORBA)/Interface Definition Language (IDL), or Web Service Definition Language (WSDL). Some examples of protocols include Simple Network Management Protocol (SNMP), REpresentational State Transfer (REST), Simple Object Access Protocol (SOAP), and proprietary protocols. Recently, the Internet Engineering Task Force (IETF) proposed Yet Another Next Generation (YANG) and Network Configuration Protocol (NETCONF), where all the network components have their data model defined in YANG and use a NETCONF interface [8]. However, such a change is not easy to achieve due to the incompatibility of many legacy and proprietary systems. Overall, there is a growing industry interest in modeling enabled by YANG/NetConf while they continue to mature. However, the silo problem lies in the vendor implementations that are likely to follow the historical pattern of being 30-40 percent standardized and 60-70 percent proprietary. With respect to information models, ETSI NFV ISG has launched a Multi Standard Development Organizations (SDO) project seeking to develop a comprehensive information model involving 3GPP Service & Systems Aspects group (SA5), TeleManagement Forum (TMF), and others.

Finally, TMF has a project underway, "Future Mode of Operation" (FMO) [9], seeking to provide a foundation for composable services based on underlying virtualization.

Overall, there is a growing industry interest in modeling enabled by YANG/NetConf while they continue to mature. However, the silo problem lies in the vendor implementations that are likely to follow the historical pattern of being 30-40 percent standardized and 60-70 percent proprietary.

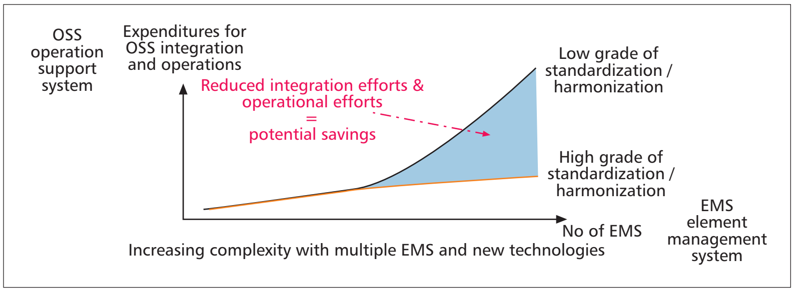

Figure 2. TNO's increasing operations and "integration tax."

From the research perspective, there exists a proposal for an NFV-based solution by NTT-Japan [10], i.e. vConductor, which also considers an information model for a TNO. Similarly, Kt-Infra-Lab focuses on an SDN-only management framework [11], without defining the information model. On the other hand, the *IEEE Communications Magazine* Guest Editorial on 5G [12] focuses on SDN-enabled and NFV-enabled 5G solutions, focusing on network handovers, users' mobility, and terminal management functions. Overall, scientists contributing to research on the end-to-end management frameworks for TNOs do not research holistic solutions, assuming interoperability and focusing on specific network components.

To summarize, by themselves XML, neither YANG/NETCONF, nor platforms emerging from an extensive research project provide end-to-end solutions to the operational silo problem. Additionally, solutions must link to legacy systems, and enable new and currently unforeseeable technologies. Toward this end, we propose e2eUberIM to address these shortcomings. In doing so, we fulfill the TMF's FMO vision at different M&O levels.

The e2eUberIM Requirements

This section presents the requirements for the e2eUberIM as derived from the literature, regulatory and standardization efforts, and the authors' experience with TNO M&O.

Requirements Overview

User Experience: e2eUberIM must maintain the user's QoS/QoE by leveraging any underlying resource suitable for the service, especially while the underlying infrastructure under-delivers.

Composable Services: e2eUberIM must enable the rapid development and deployment of innovative, composed services, i.e. anything-as-a-service (XaaS) as combinations of existing atomic service components, potentially with new elements and service functions.

Completeness: UberIM must combine all the data models of all the current and any new network operations subsystems so as to encompass the whole network by translating from all the other structures into a unified one. It must have a "picture" of the whole network at any point of time.

Embracing Heterogeneity: e2eUberIM must embrace the high heterogeneity and volatility of the infrastructure, and encompass all the proprietary extensions that vendors add and change with time, embracing the legacy interfaces and standards (e.g. EMS/OSS and BSS). It must be able to communicate with any sub-system in a network, across any interface, with any message format based on any data model, described in any language. The change of communication schemas must be done in real-time.

Localized Operations & Management (O&M): e2eUberIM must embrace the high geographical distribution and volatility of the infrastructure, and have localized and global O&M functions, all harmonized along the holistic "picture" of the whole network at any point in time.

Expressiveness: UberIM must describe the full range of potential variability in the momentary, i.e. just-in-time of service delivery requirements. In so doing, it must overburden neither communications facilities nor local processing facilities by delivering not needed information.

Distributed & Organic Operation: e2eUberIM must operate efficiently and effectively over the distributed and centralized data stores (storing the data models) and must have an organically flexible distributed coordination process. Organically flexible implies:

- Local system components may discover their neighbors, and negotiate with them interoperable interface details (OACS: objectives, algorithms, constraints and structures).
- Global control (with established components' OACS) within a given domain, and enable establishing a hierarchy of global control elements where a group of such elements is controlled by an element one step higher in the hierarchy.

The data stores must also be organic, i.e. able to change their information/data models in real-time without interrupting system operation.

Volume and Velocity of Data: The e2eUberIM data store has to be able to store large amounts of data and handle high transaction volumes originating from single EMSs, OSSs, and other network components within its span of control.

Hybrid Input: e2eUberIM must operate based on the hybrid input, i.e. machine-based and human, manual-based input. It must employ algorithms to autonomously, and, whenever needed, in conjunction with a manual input, orchestrate (control, manage, optimize) the network in real-time.

Abstraction: e2eUberIM must enable fast, cost-effective deployment of NFV, SDN, and other solutions on top of the existing EMS/OSS infrastructures.

Openness: e2eUberIM must enable easy and effective deployment of open API's to allow operators and third parties to deploy their algorithms. New components and services must be fast to design, test, and deploy.

Management: e2eUberIM must be effective and efficient to be managed itself.

Performance: e2eUberIM must be flexible across different system layers, highly scalable, and get just the right information (no more, no less), to the right point (system component), at the right time.

Speed: e2eUberIM must operate in real-time, having a view of the network in close to real-time.

Security: e2eUberIM must have designed-in adequate security measures, not added as an afterthought.

Cost: e2eUberIM must enable a reduction of costs of operation of TNO infrastructures.

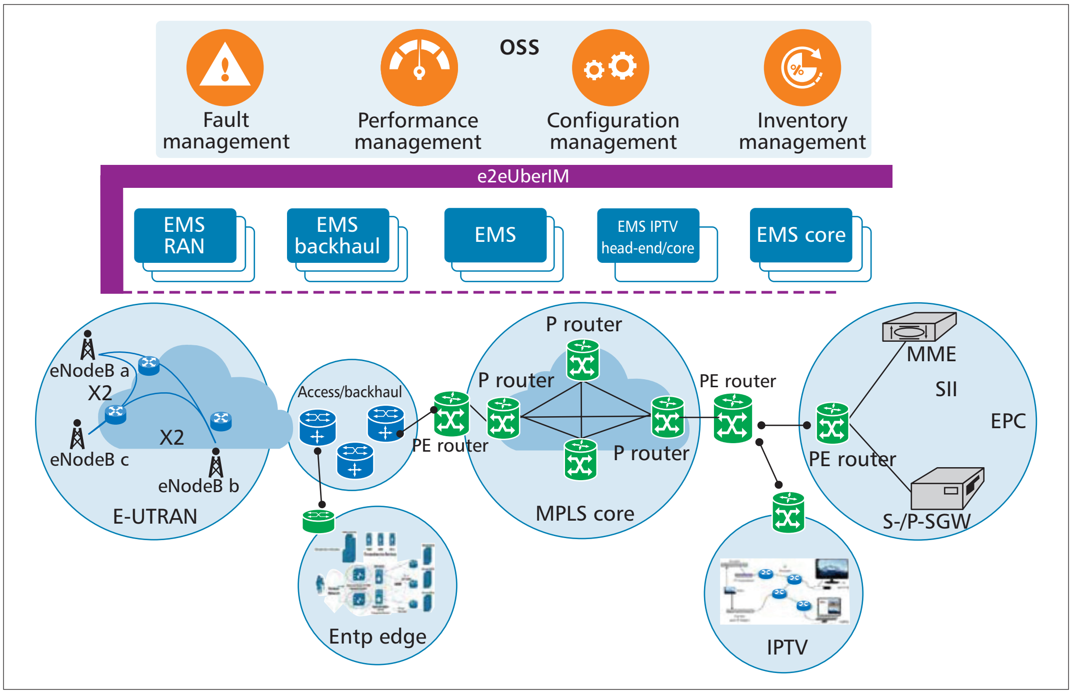

Figure 3. The e2eUberIM placement in the current infrastructure.

e2eUberIM must combine in a latency-aware manner the information from some otherwise incompatible systems, and based on that information, make timely decisions about which changes are needed for meeting the O&M objectives.

The baseline for the comparison is current OSS integration/operation costs and their projected growth if e2eUberIM is not implemented.

"Speed," "Localized O&M", and "Openness" Requirements

The latency of any operation is a function of the number of instructions executed, hardware latency, data storage accesses, and network accesses and transfers. The latency challenge stems from the fact that some EMSs can take multiple minutes to respond to a transaction, and when combining multiple EMSs, OSSs, and BSSs, these delays cascade.

We further differentiate speed requirements for O&M aspects:

- *Control* requires extremely low latency, milliseconds to minutes.
- *Management* requires higher latencies, minutes to hours.
- *Orchestration* is a combination of the high latency portion of control and the low latency of management.
- *Planning* implies a high latency, days to months.

e2eUberIM must combine in a latency-aware manner the information from some otherwise incompatible systems, and based on that information, make timely decisions about which changes are needed for meeting the O&M objectives.

Additionally, the challenge is to manage the localized O&M and applications and to deliver the needed (and only the needed) information to the O&M functions and applications at the required location and at the required time. This challenge is twofold:

- If the data is delivered to a single network location, the large volume of data may require significant time for the point to process it; the data is 'stalled' and cannot be acted upon.
- There are many O&M functions and applications that need access to only specific parts of the data, and if they cannot access it, the actions will be based on old information.

The latency requirement, propelled by NEs/O&M functions/applications' distribution and heterogeneity, is one of the most significant determinants of the e2eUberIM architecture. They imply the following:

- *Precision information:* the information model must clearly indicate exactly which information is needed, where, and for whom.
- *Precision communication:* the communication process must deliver that exact information exactly where and when it is needed.

The above requirements are further enclosed in the "openness" requirement, which focuses on a standardized way of information computing, storage, and communications. It must enable a high level of granularity for information storage and communication.

The e2eUberIM Design Choices

This section presents the UberIM and e2eUber design choices' motivations.

UberIM and e2eUber Designs

The UberIM is derived from the Umbrella Information Model (UIM) [13] developed by 3GPP SA5 and TMF, bringing together wireline and cellular standards communities. The current UIM is XML-based; however, it covers only the standardized components, i.e. 30–40 percent of the interfaces from the cellular radio access network (RAN) and evolved packet core (eCore) EMSs and OSSs. This ratio is similar for fixed network components. The proposed UberIM encompasses the proprietary extensions and implementations that vendors deploy and add in time. UberIM is defined at a conceptual level, independent of implementations or

Interoperability level		UberIM	e2eUber	TNO component/process involved
6	Conceptual	Fine-grained, adaptable UberIM	Service factory, low latency communications process	Composed services; XaaS: anything-as-a-service
5b	Dynamic	Adaptable UberIM	Low latency communication process	Management of network and environmental change
5a		UberIM	Low latency communication process	Management of environmental change
4	Pragmatic	UberIM	High latency communication process	Process: planning
3	Semantic	UberIM	Message format	EMSes
2	Syntactic	Data model	Message format	NEs
1	Technical	—	Protocol	—
0	None	—	—	—

Table 1. Conceptual interoperability model mapping on e2eUberIM and the TNO components.

data transport protocols (conforming RFC 4333 [14]). The data model (DM) is at a lower level of abstraction and it includes implementation details.

The e2eUber encompasses the information communication (and storage) processes of TNO [2] that further rely on UberIM for information exchange. We explain the e2eUberIM design based on the Conceptual Interoperability Model (LCIM) [15] having layers each corresponding to a different level of interoperability for a given system. The authors have broken one of the layers down into two sub-layers, as motivated further. The generalized model has been adapted to TNOs' infrastructures as follows. Level 0 represents non-interoperability, while level 7 represents the conceptual, implementation-independent level of interoperability allowing for composable services (Table 1).

Level 0 represents the stand-alone systems that have **no interoperability**, i.e. the information and management silos at their worst. Current RAN and wireline EMSs have level 0, 1, 2, and portions of 3. Level 0 does not map onto the e2eUberIM.

Level 1 represents technical interoperability, where **communication protocols** exist for data exchange between elements, e.g. the current 3GPP SA5 northbound interface from the EMS to the OSS (ItfN), SNMP (although having rigid MIB format, not expressive enough), or NETCONF (not covering the legacy interfaces, proprietary extensions). e2eUberIM operates on this level via the e2eUber protocol for a communication process.

Level 2 represents syntactic interoperability focusing on a common structure to exchange information, i.e. **a common data format**, e.g. for NEs. For example, ItfN specifies 30–40 percent of the data model and corresponding message formats in XML, while YANG has limitations in expressiveness and does not cover the legacy interfaces or proprietary extensions, such as JSON, SNMP MIB, CORBA/IDL, or WSDL. e2eUberIM operates on this level via the UberIM data model and e2eUber message format for services' integration.

Level 3 represents a definition and deployment of the common information exchange reference model, where the **meaning of the data** is shared between components/services, e.g. EMSs and OSSs. For example, Unified Modeling Language (UML) is used in 3GPP to define information models (partially proprietary and not fully interoperable). e2eUberIM operates on this level via having processes to define and deploy a common information model: UberIM spanning across multiple data models and e2eUber message format for services' interoperability.

Level 4 represents pragmatic interoperability, where the interoperating systems are aware of the **methods and procedures** that each system is employing, i.e. context of the systems' use of data. For example, the current 3GPP SA5 northbound interface partially fulfills this goal via its communication protocols. The e2eUberIM operates on this level via UberIM and e2eUber high-latency communication process for global harmonization of the M&O activities.

Level 5 represents dynamic interoperability, where the operational system is "aware" and acts upon the environmental and other **assumptions and constraints** that affect its data interchange. For example, a TNO network with 100 million subscribers with 1 million base stations, at any given moment has between 10,000 and 100,000 base stations with impaired backhaul, and a centralized solution to manage that would fail because of the combination of scale and volatility. Only a harmonious combination of local and global M&O activities communicating in an effective way can deal with the volatility and scale of today's TNOs.

e2eUberIM operates on this level twofold. **Level 5a** enables interoperability with the environments via UberIM as an information model and a low latency communication process, e2eUber, focusing mainly on the local M&O activities defined via a data tuple (OACS). **Level 5b** adds the adaptation capability to the UberIM model (e.g. changes in NE resulting in changing local data models), while a low latency communication process, e2eUber, enables (locally) to manage the environmental changes. e2eUberIM operates in real-time, without the disruption to the live TNO network and services.

Level 6 represents the ultimate conceptual model enabling service compose-ability (XaaS), where the **assumptions and constraint**s of the meaningful implementation-independent abstractions of reality are aligned. This mode of operation is part of the TMF's intent in its "Future Mode of Operation" vision [10].

e2eUberIM operates on this level via a fine-grained, adaptable UberIM information model and e2eUber fine-grained, low latency communication processes, together forming a **service factory**. At this level, the e2eUberIM can compose *anything-as-a*-service (XaaS, predefined via a data tuple OACS) from the pre-existing logical components and then, enable fine-grained management by converting these compositions back into non-compatible data models of the actual NE's in the network. A composition function of e2eUberIM enables an assembly of fast service prototypes, testing these against snapshots of operational data, conversion of the prototypes into the deployable product, deployment of this product, its tracking and evaluation in the environment where it operates. A composition function of e2eUberIM has an open API and fine-grained, low latency services, and moreover, it enables equally to plan in months, prototype and test in weeks, and deploy the services in days to run for years, as well as plan in hours, prototype and test in minutes, deploy in seconds, and run for minutes.

With the Layer 6 solution, composed services can be deployed despite the underlying silos. Thus, e2eUberIM enables TNOs to handle the scale, complexity, and volatility of the EMSs.

Fulfilling the "Speed," "Localized O&M", and "Openness" Requirements

The proposed e2eUberIM fulfills the above requirements as follows.

Local Communications: A network element (NE) interacts approximately 95 percent of the time with only a very small subset of all the NEs, i.e. its physical and virtual neighbors. If the NEs or a subsystem associated with that NE holds the detailed data associated with it, the task of continuing all the current M&O for that NE is manageable with acceptable latency. The challenge is to communicate some of this data to the neighbors that require it when they require it. The local UberIM, i.e. the local information model, supports the local e2eUberIM, i.e. the local communications model. In this way the local UberIMs are local "virtualizers"/"orchestrators", one for each EMS or OSS providing local virtualization, adaptation, and coordination. They observe their neighbors, and based on policies they respond to specific events (e.g. alert a human operator, take an automated action) and provide adaptation and virtualization of the interface they are attached to, and translate the information contained in the device they are associated with into UberIM, a consistent over-arching end-to-end information model.

Global Communications: For the small number of interactions that are global in nature (approximately 5 percent), the e2eUberIM communication process provides the data to the global applications (e.g. orchestration, planning) that require this data at their location(s), whenever they require it.

The global UberIM information model supports the global e2eUberIM communications model. One global "virtualizer"/'orchestrator" for the whole system communicates with the local "virtualizers"/"orchestrators", and based on policies, responds to specific events. It is responsible for global adaptation and coordination.

Openness: e2eUberIM provides an open API based on 3GPP's and the TM Forum's Umbrella Information Model (UIM), enabling third parties to easily connect. It allows TNOs to cost-effectively take advantage of emerging specialized systems such as third party analytics and inventory systems and eases development activities. The TNO can use the in-house developed services itself and/or sell to other TNOs (generating revenue).

Conclusions and Future Work

The current TNO infrastructure relies on ever increasing numbers of non-interoperable silos demanding significant effort for O&M while their ever increasing number of subscribers demands the best possible user experience for their services at decreasing prices. This is not a sustainable situation, and TNOs are recognizing it. Overall, they strive for a holistic view of the O&M data, controlled operational costs, rapid deployment of 'any-service' (XaaS), and improved user QoE. The e2eUberIM solution proposed in this paper addresses these needs and meets the TNOs' needs for end-to-end network coordination and optimization. It also enables the vendors to continue to innovate and differentiate their products and value price, as well as assures that the networks can deal effectively with rising scale, complexity, and volatility.

The novelty of the e2eUberIM framework relates to virtualizing the proprietary interfaces of various network components, collecting the operations data from them, and converting it into a single comprehensive view of the network (UberIM), communicated effectively among the network components (via e2eUber). As a result, the existing OSS/BSSs are enhanced, shortcomings that exist in automated network functioning due to lack of end-to-end operational data are filled in, and an open API allows operators and third party providers to bring their special expertise to bear on problems of coordination, optimization, and end-to-end orchestration. In turn, previously impossible functions or functions requiring much manual coordination can be automated. In doing so, the e2eUberIM enables operational cost reduction. For example, the operations staff currently engaged in manually linking/coordinating the over 200 inventory management systems (mentioned in the Introduction) can now be freed to work on automation applications previously impossible that make the network operate more efficiently. That implies moving them from an expense function to a revenue function. At the same time, the efficiency improvements coming from end-to-end orchestration via e2eUberIM can be translated into reduced provisioning of the infrastructure, resulting in reduced CAPEX. As a corollary, the efficiency improvements may also result in reduced power consumption.

The e2eUberIM has already been successfully demonstrated in three high profile public demonstrations under the auspices of TeleManagement Conferences (Nice, France, June 2013, and San Jose, CA, USA, October 2013). Our current and

The current TNO infrastructure relies on ever increasing numbers of non-interoperable silos demanding significant effort for O&M while their ever increasing number of subscribers demands the best possible user experience for their services at decreasing prices. This is not a sustainable situation, and TNOs are recognizing it.

The e2eUberIM has already been successfully demonstrated in three high profile public demonstrations. Our current and future work includes 'in-situ' TNO-based deployment of e2eUberIM and the collection of data providing evidence for its efficiency and effectiveness.

future work includes 'in-situ' TNO-based deployment of e2eUberIM and the collection of data providing evidence for its efficiency and effectiveness.

References

[1] GSMA Report, Comparison of Fixed and Mobile Cost Structures, available from http://www.gsma.com, 2012.

[2] E. Bertin, N. Crespi, and T. Magedanz, *Evolution of Telecommunication Services: The Convergence of Telecom and Internet: Technologies and Ecosystems*, v.776. Springer, 2013.

[3] V. Gonçalves and P. Ballon, "Adding Value to the Network: Mobile Operators' Experiments with Software-as-a-Service and Platform-as-a-Service Models," *Telemat. Informatics*, vol. 28, no. 1, 2011, pp. 12–21.

[4] ITU-T, Vocabulary for Performance and QoS, ITU-T. P.10/G.100, 2011.

[5] S. Ickin *et al.*, "Factors Influencing Quality of Experience of Commonly-Used Mobile Applications," *IEEE Commun. Mag.*, vol. 50, no. 4, 2012, pp. 48–56.

[6] A. Ghezzi, M. N. Cortimiglia, and A. G. Frank, "Strategy and Business Model Design in Dynamic Telecommunications Industries," *Technol. Forecast. Soc. Change*, vol. 90, no. 1, 2015, pp. 346–54.

[7] T. Wood *et al.*, "Toward a Software-Based Network: Integrating Software Defined Networking and Network Function Virtualization," *IEEE Network*, vol. 29, no. 3, May/June 2015, pp. 36–41.

[8] J. Schonwalder, M. Bjorklund, and P. Shafer, "Network Configuration Management Using NETCONF and YANG," *IEEE Commun. Mag.*, vol. 48, no. 9, 2010, pp. 166–73.

[9] TMF Forum, IG1118 OSS/BSS Futures: Preparing the Future Mode of Operation R14.5.1, technical report, 2014.

[10] W. Shen *et al.*, "vConductor: An NFV Management Solution for Realizing End-to-End Virtual Network Services," *AP-NOMS*, 2014, pp. 1–6.

[11] S. Baik, C. Hwang, and Y. Lee, "SDN-based Architecture for End-to-End Path Provisioning in the Mixed Circuit and Packet Network Environment," *AP-NOMS*, 2014, pp. 1–4.

[12] D. Soldani *et al.*, "5G Networks: End-to-End Architecture and Infrastructure," Guest Editorial, *IEEE Commun. Mag.*, vol. 52, no. 11, Nov. 2014, pp. 62–64.

[13] 3GPP TS 28.620, Telecommunication Management; Fixed Mobile Convergence (FMC) Federated Network Information Model (FNIM) Umbrella Information Model (UIM), 2014.

[14] A. Pras, "On the Difference between Information Models and Data Models," IETF RFC 3444, 2003.

[15] W. Wang, A. Tolk, and W. Wang, "The Levels of Conceptual Interoperability Model: Applying Systems Engineering Principles To M&S," *Spring Simulation Multiconference*, article 168, 2009.

Biographies

Katarzyna Wac (wac@stanford.edu), Ph.D. is an associate professor of computer science at the University of Copenhagen, Denmark, and the University of Geneva, Switzerland, affiliated with Stanford University since 2013. Her research interests include quality of experience-aware mobile services with special emphasis on support for adaptive multimedia protocols and quality of service mechanisms, especially in the mobile healthcare domain.

Mark Cummings (mark@orchestralnetworks.com), Ph.D., is currently the CTO of Orchestral Networks. He has actively participated in 3GPP SA5, NGMN, TMF, and IEEE 802. Previously he was the principal inventor of the first SDR patent granted, creator of the first successful SDR semiconductor start-up (Morphics, now owned by Intel) and multimode/multiband antenna company (SkyCross). He has more than 200 publications and six patents, and he holds an MBA in communications management (Wharton) and a Ph.D. in information systems (Tohoku Imperial University, Japan).

Jayanta Dey (jayanta.dey@wipro.com) is the VP and CTO of Wipro's Media and Telecom & Product Engineering Services. He holds a B.E. (Hons) in electronics engineering (BITS Pilani) and an MBA (NorthEastern University). He specializes in fault tolerant operating systems, distributed systems, network and service management, SDN/NFV, and wireless technologies. He has been responsible for building strategic partnerships with telecom clients globally, as well as representing Wipro in various industry forums, including OSS/J, LTE Small Cell, and TM Forum.

CALL FOR PAPERS
IEEE COMMUNICATIONS MAGAZINE
NEW WAVEFORMS AND MULTIPLE ACCESS METHODS FOR 5G NETWORKS

BACKGROUND

With expected 2020 initial commercialization, 5G mobile communications is gathering increased interest and momentum around the world. Following discussions on the 5G vision and key requirements (such as high data-rate, low latency, and massive connectivity), various candidate technologies have been proposed and investigated. The candidate enablers for 5G mobile communications include massive antenna technologies (from legacy cellular frequency bands up to high frequencies) to provide beamforming gain and support increased capacity, new waveform (or a new radio access technology (RAT)) to flexibly accommodate various services/applications with different requirements, new multiple access schemes to support massive connections, and so on.

The International Telecommunication Union (ITU) has categorized the usage scenarios for International Mobile Telecommunications (IMT) for 2020 and beyond into 3 main groups: Enhanced Mobile Broadband, Massive Machine Type Communications, and Ultra-reliable and Low Latency communications. In addition, they have specified target requirements such as peak data rates of 20 Gb/s, user experienced data rates of 100 Mb/s, a spectrum efficiency improvement of 3X, support for up to 500 km/h mobility, 1-ms latency, a connection density of 106 devices/km^2, a network energy efficiency improvement of 100X and an area traffic capacity of 10 Mb/s/m2. While all the requirements need not be met simultaneously, the design of 5G networks should provide flexibility to support various applications meeting part of the above requirements on a use case basis.

There is also increased interest in the use of spectrum above 6 GHz for 5G mobile communications. Several researchers in academia and industry have explored the feasibility of using mmWave frequencies for 5G mobile communications, considering frequencies up to 100 GHz. This has also been supported by regulatory bodies with ITU-R investigating the spectrum between 6 – 100 GHz for possible global harmonization and usage by 2020 and regulatory bodies such as FCC in the US and OFCOM in UK, starting a notice-of-inquiry (NOI) for using mmWave spectrum for mobile communications.

There are several considerations in the design of new waveforms and access methods for 5G such as:

- Support for various applications and usage scenarios outlined above
- Support for large bandwidths (100 MHz) and mmWave frequencies
- Improved spectral efficiency over state-of-the-art waveforms such as OFDM
- Support for MIMO and other techniques
- Low power and low complexity design

The goal of this FT (Feature Topic) is to give a comprehensive view and comparison of different waveform, modulation and access techniques being proposed for 5G communication systems. This will enable researchers to provide their viewpoints and share the latest research results, which will feed into the 5G ecosystem development as standardization and regulatory bodies explore spectrum and access techniques for 5G.

Original contributions are invited on the latest advancements on 5G waveform and multiple access design methods. The topics of interest within the scope of this issue include (but are not limited to) the following:

- Non-OFDM waveforms for 5G such as FBMC
- OFDM variants for 5G with filtering and windowing
- Waveforms for efficient large bandwidth signal processing
- Waveform features to support mmWave spectrum (such as reduced PAPR)
- Low power and energy efficient waveforms
- Non-orthogonal Multiple Access schemes
- Flexible numerology to support diverse applications for 5G

SUBMISSIONS

Articles should be tutorial in nature and written in a style comprehensible and accessible to readers outside the specialty of the article. Complete guidelines for prospective authors can be found at http://www.comsoc.org/commag/paper-submission-guidelines. The guest editors reserve the right to reject papers they unanimously deem to be either out of scope of this Feature Topic or otherwise extremely unlikely to be accepted after a peer review process.

It is important to note that *IEEE Communications Magazine* strongly limits mathematical content, and the number of figures and tables. Mathematical equations should not be used (in justified cases up to three simple equations are allowed). Article length (introduction through conclusions, excluding figures, tables and their captions) should not exceed 4,500 words. Figures and tables should be limited to a combined total of 6. The number of archivable references is limited to 15. Non-archivable references (website URLs, web-posted papers and reports, unpublished/to be published papers) should not be used. All articles must be submitted through the IEEE Manuscript Central site (http://mc.manuscriptcentral.com/commag-ieee) to the "November 2016 / 5G Waveform" category by the submission deadline according to the following schedule:

IMPORTANT DATES

- Manuscript Submission Deadline: April 2, 2016
- Decision Notification: July 1, 2016
- Final Manuscript Due Date: August 15, 2016
- FT Publication Date: November 2016

GUEST EDITORS

Charlie Jianzhong Zhang
Samsung Research America, USA
jianzhong.z@samsung.com

Jianglei Ma
Huawei, Canada
jianglei.ma@huawei.com

Geoffrey Ye Li
Georgia Institute of Technology, USA
liye@ece.gatech.edu

Wei Yu
University of Toronto, Canada
weiyu@ece.utoronto.ca

Nihar Jindal
Google, USA
niharjindal@google.com

Yoshihisa Kishiyama
NTT DoCoMo, Japan
kishiyama@nttdocomo.com

Stefan Parkvall
Ericsson, Sweden
stefan.parkvall@ericsson.com

An Edge Operating System Enabling Anything-as-a-Service

The authors argue that SDN and NFV, together with cloud and edge-fog computing, can be seen as different facets of a systemic transformation of telecommunications and ICT, called softwarization. The first impact will be at the edge of current telecommunications infrastructures, which are becoming powerful network and service platforms. The edge operating system (EOS) software architecture is proposed as the means to get there.

Antonio Manzalini and Noel Crespi

Abstract

This article argues that SDN and NFV, together with cloud and edge-fog computing, can be seen as different facets of a systemic transformation of telecommunications and ICT, called softwarization. The first impact will be at the edge of current telecommunications infrastructures, which are becoming powerful network and service platforms. The edge operating system (EOS) software architecture is proposed as the means to get there. In fact, the main feature of EOS is to bring several service domains, such as cloud robotics, Internet of Things, and Tactile Internet, into convergence at the edge. The development of EOS leverages available open source software. A use case is described to validate the EOS with a proof-of-concept.

Communications Standards

Context and Drivers

We are witnessing a period of rapidly growing interest on the part of industry and academia in software-defined networks (SDN) [1] and network function virtualization (NFV) [2]. The growing interest in these paradigms (re-proposing principles have been well known) is most probably motivated by the novelty of the overall context, specifically their techno-economic sustainability and high-level performance. These advances are mainly due to the technological milestones achieved in the last two decades: the impressive diffusion of fixed and mobile ultra-broadband, the increasing performance of chipsets and hardware architectures, the ever-growing availability of open source software, and the cost reductions (determined also by a shift in how IT services are provided).

This article argues that, thanks to these techno-economic trends, SDN and NFV principles will soon impact not only current telecommunications networks, but also service and application platforms. In fact, SDN and NFV, together with cloud, edge and fog computing, can be seen as facets of a broad innovation wave (called softwarization) that is accelerating the ongoing migration of "intelligence" toward the users.

In view of that, it is argued that the first impact of softwarization will be at the edge, which is defined as the peripheral part of current infrastructures, ranging from the distribution and access segments up to the direct proximity to users (e.g., home, office).

While cloud computing is a well known paradigm, already exploited from an industrial point of view, the concepts of edge and fog computing require, at least for this article, a short definition. Concerning the former, we basically refer to ETSI [3] which defines mobile edge computing as the method of providing IT and cloud-computing capabilities within the radio access network (RAN) in close proximity to mobile subscribers. Fog computing pushes the edge computing paradigm up to the end users terminals (e.g., smart phones) and other devices, which will be able to store pieces of data and to execute service components locally.

Softwarization will be a radical change of paradigm. Current telecommunications infrastructures have been exploited with purpose-built equipment designed for specific functions. In the future, network functions and services will be virtualized software processes executed on distributed horizontal platforms mainly made of standard hardware resources.

Standard hardware and open source software will play a strategic role in this profound transformation, by fuelling open innovation while reducing the investments required to deploy said infrastructures. For example, OpenStack [4] is an open source platform designed to provide cloud services. Several pre-standardization bodies and fora regard OpenStack as an ideal candidate for developing orchestration features in NFV infrastructures. Other examples of open source software are the SDN controllers that have been released to date, and ONOS [5].

SDN and NFV open source software adoptions will fuel innovation and reduce software costs. This does not not necessarily mean CAPEX and OPEX reductions for operators and service providers. In fact, most open source software products may eventually require contracted third party support to become exploitable in production environments for commercial, industrial, financial, and public service applications.

On the other hand, the entire value chain will change radically: the "thresholds" for new players to enter the telecommunications and ICT (Information and Communication Technology) markets [6] will be lowered.

Cost savings alone will not be enough to assure the future sustainability of the telecommunications industry: it is key also to enable innovative service paradigms. Two often-mentioned examples are "immersive communications" and "anything as a service," service paradigms that are posing challenging requirements for future telecommunications infrastructures.

"Immersive communications" looks beyond the "commoditization" of current communication paradigms (e.g., voice, messaging, etc.) by addressing new advanced forms of social communications and networking (e.g., artificially intelligent avatars, cognitive robot-human interactions, etc.).

"Anything as a service" is about providing (anytime and anywhere) wider and wider sets of ICT services by means of new terminals, even

Antonio Manzalini is with Telecom Italia, Strategy and Innovation — Future Centre.

Noel Crespi is with Institut Mines-Telecom.

going far beyond our imagination (e.g., intelligent machines, robots, drones, and smart things) [7]. Imagine, for example, services for improving industrial and agricultural efficiency, for enabling decentralized micro-manufacturing, for improving efficiency in private-public processes, and for creating and maintaining smart environments.

In summary, softwarization is a systemic transformation. It is not just about the introduction of another technology or network layer in current infrastructures. Rather, it goes beyond the networks to also impact the service platforms and the future role of terminals. In this respect, beyond the technological aspects, softwarization implies business sustainability and strategic regulatory issues.

The outline of this article is as follows. We outline the main enabling technologies, and we describe the software architecture of the edge operating system (EOS). We then describe a use-case and elaborate on the design and development of the EOS, leveraging open source software. Closing remarks are provided in the last section.

Enabling Technologies of Telco Softwarization

SDN and NFV are two of the most-discussed technologies capable of enabling the softwarization of telecommunications.

SDN relies on the separation of control and data-forwarding functions. In principle, this is applicable to any node of a telecommunication network (e.g., a switch, a router, or other transmission equipment). Another key characteristic of SDN is the possibility of executing the above-mentioned (control) software outside of the equipment boundaries, for example on dedicated IT servers or even in a data centre (e.g., cloud computing). Control programmability (via APIs) is a third relevant aspect of SDN.

NFV is about the virtualization of network functions and their dynamic allocation and execution on (almost) general purpose processors (e.g., x86), shared over multiple customers, datastreams, and applications. SDN and NFV are not directly dependent, but they are mutually beneficial. In fact, when coupled, they amplify their potential innovation impact on telecommunications infrastructures.

If software-hardware decoupling and the virtualization of functions and services can be seen as the "common denominator" of softwarization, the potential differentiation and evolution of cloud toward edge and fog computing represents other interesting and synergistic expressions of the same overall transformation.

TPC is a serious performance bottleneck for video and other large files (as it requires receiver acknowledgement) and throughput is inversely related to round trip time (RTT) or latency. It is impossible to provide HD-quality streams if the servers are not relatively close to the users. At the same time, with just best effort traffic it will not be possible to achieve the low latency requirements posed by services such as caching or interactive applications.

In fact, the "last mile" connection between a user and the ISP is a significant bottleneck. According to the FCC's Measuring Broadband America report [8], during peak hours "Fiber-to-the-home services provided 17 milliseconds (ms) round-trip latency on average, while DSL-based services averaged 44 ms."

It should be mentioned that PON and DSL delays are intrinsic in the access protocols. Achieving lower delays means either changing said protocols or locating all of the necessary data at the subscriber, including content caches and databases.

In the latter direction, fog computing pushes the edge computing paradigm even further, up to the end users' terminals and devices, which are storing data and locally executing pieces of service logic. This will further amplify the diffusion of applications and the migration of "intelligence" toward the users.

In summary, it is very likely that techno-economic drivers and emerging technologies will create the conditions for exploiting very powerful network and service platforms at the edges of current infrastructures. Such platforms will be able to carry out a substantial amount of storage and real time computation, thereby supporting a wide range of innovative communications and ICT services.

> It is very likely that the techno-economic drivers and the emerging technologies will create the conditions for exploiting very powerful network and service platforms at the edges of current infrastructures. Such platforms will be able to carry out a substantial amount of storage and real time computation.

The Edge Operating System

The edge of current telecommunications infrastructures (i.e., the access areas up to the direct proximity to users) will become powerful network and service platforms. The EOS software architecture proposed by this article is the means to get there. EOS will provide the typical services of an operating system, e.g., abstractions, low-level element control, commonly-used functionalities, message-passing between processes, management of packets of processes, etc.

For the basic design of the EOS, we took our inspiration from the architecture of the robot operating system (ROS) [9], an open source, widely adopted meta-operating system for robotic systems. Among the merits of ROS that have been adopted by EOS is the variety of processes (called nodes), executed on a number of different hosts, connected at runtime with logical topologies.

Moreover, another main reason for that design choice is the observation that a robot, generally speaking, can be considered a dynamic aggregation of resources such as sensors, actuators, and processing-storage capabilities, implementing a cognitive loop. These are the same categories of resources that will populate the edges of current infrastructures, named infrastructure elements (IE).

Obviously the domain contexts ROS and EOS applications are quite different; in fact, the design of the EOS software architecture has been extended to meet the edge requirements. In particular, a physical IE (Fig. 1) has been defined to include any dynamic combination of sensors, actuators, processing-storage resources, and data forwarding capabilities. Sensors, actuators, robots, drones, routers, and terminals can all be seen as particular physical IEs. This generalization will help in structuring the functional model of the EOS.

From a functional perspective, IE will provide a set of services, leveraging the concept of the self-managed cell reported in [10]. For example, the set of services may include: discovery services

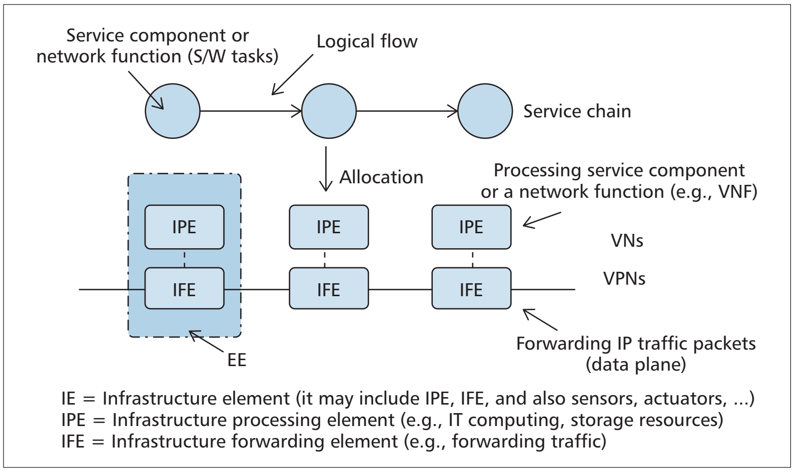

Figure 1. Edge elements.

to discover local resources and components as part of the physical IE; policy services to manage the policies specifying IE behavior; or even cognitive services to implement a certain level of cognition (Fig. 2), even more when coupled with sensing and actuating capabilities.

Generalizing, we can say that cognition (implemented through artificial intelligence methods, deep learning techniques, heuristics, algorithms, etc.) will allow IEs to learn and reason about how to behave in response to goals in a complex context, or at least be able to optimally execute their service and network functions.

Key Characteristics of the Robot Operating System

The ROS is a widely adopted meta-operating system for robots. The full source code of ROS is publicly available and currently runs only on Unix-based platforms. A robotic system built using the ROS consists of a number of processes (ROS nodes), potentially on a number of different hosts, connected at runtime in a peer-to-peer topology. ROS has a lookup mechanism (ROS master) to allow processes to find each other at runtime.

ROS master acts as a name-service, storing topics and service registration information for ROS nodes. Nodes communicate with the master to report their registration information. As these nodes communicate with the master, they can receive information about other registered nodes and make connections as appropriate (bypassing messages and structuring data).

Nodes connect to other nodes directly; the master only provides lookup information, much like a DNS server. Nodes that subscribe to a topic will request connections from nodes that publish that topic, and will establish that connection over an agreed upon connection protocol (e.g., standard TPC/IP sockets). This is represented in Fig. 3.

An ROS node sends a message by publishing it to a given topic, which is simply a string such as "map." A node that is interested in a certain kind of data will subscribe to the appropriate topic. There may be multiple concurrent publishers and subscribers for a single topic, and a single node may publish and/or subscribe to multiple topics.

Although the topic-based publish-subscribe model is a flexible communications paradigm, it is not appropriate for synchronous transactions, which can simplify the design of some nodes. Therefore, ROS developers have introduced the concept of services, defined by a string name and a pair of strictly typed messages, one for the request and one for the response. This is analogous to web services, which are defined by URIs and have request and response documents of well defined types.

The special characteristics of the ROS architecture allow for decoupled operation, wherein names are the primary means by which larger and more complex systems can be built. This decoupling is one of the main reasons why we have taken most of our inspiration from ROS when designing the EOS. One of the EOS requirements, in fact, is to allow the flexible and scalable operations of complex and dynamic systems, built by aggregations of IEs.

Assumptions of the Edge Operating System

EOS leverages the concept of services as representing a sort of "unifying" abstraction across physical edge resources, across multiple infrastructure domains, and across different service levels. A service provides a function (e.g., from ISO-OSI L2 to L7, so it could also be a network function, or a middle-box), it exports an API, it is available anywhere and anytime (location-time independent), is scalable, elastic, and resilient, and it can be composed with other existing S/W components (e.g., to create a service chain). Services are executed into one or more infrastructure virtual slices, which are made of a set of logical resources (e.g., virtual machines, containers) connected through virtual networks.

The allocation and orchestration of logical resources, in charge of executing a service chain, requires solving constraint-based double optimization problems. Not only do VMs have to be properly allocated (to avoid hot-spots), but also the traffic crossing the VMs has to be properly routed (to avoid congestion).

The term orchestration has long been used in the IT domain to refer to the automated tasks involved with arranging, managing, and coordinating higher-level services provisioned across different applications and enterprises [11]. In the SDN-NFV, orchestration is concerned with lower-level (i.e., network) services, with a comprehensive management of both IT and network logical resources.

EOS software adopts a publish-subscribe model (Fig. 3) [13] as a basic way to distribute software task execution requests. Each software task execution request is coded as a tuple and written on the tuple space, named blackboard, while a take operation is used by IEs to offer their process capability.

Functional Architecture of the Edge Operating System

This sub-section describes the main characteristics of the EOS, whose high-level architecture is reported in Fig. 4.

The main elements of the EOS are listed below.

- An EOS node is an S/W module that can be executed on top of any operating system (e.g., Linux-based OS, Android, Robot Operating System, etc.) of an IE. Similar to the ROS, any EOS node communicates with an EOS master to whom it registers (e.g., services that it can provide) and updates the status (e.g., resource utilization) of its associated IEs. This data is stored in the EOS master data base (EOS MDB). IE nodes are interconnected on the data plane via fixed and virtual radio links (these links could be either local, in a single edge domain, or across a WAN).

- An EOS master is dynamically allocated to a specified edge domain. It is responsible for the local creation (and deletion) of the slice(s) where service chains are executed in order to provide the requested services. It has to interact with a higher-level EOS orchestrator and with other EOS masters, in case the service chain has to be allocated across multiple edge domains. The EOS MDB stores the data related to the IEs belonging (assigned) to the specified domain.
- The master blackboard is a sort of virtual repository shared among the EOS master and the EOS nodes. The EOS master publishes (using the pub primitive) the task/component of the service chain that has to be allocated. In turn, EOS nodes subscribe (using the sub primitive) to the S/W task/component if the associated IE can provide the logical resources to execute it (i.e., can serve the specific task of the chain). Multiple subscriptions are possible, so in a next stage the EOS master will make an optimized selection of the IEs to whom the S/W task/component will eventually be allocated.
- The collector abstraction [12] has been introduced to make master blackboards recursive, thus overcoming scalability limitations. In this sense, a collector can be seen as an agent acting as an ensemble of IEs together with their shared blackboard. A collector thus can act toward other collectors as a super-IE, as it can take tasks on its blackboard from other overloaded collectors.
- The EOS master includes a capability called the selection method, which makes it possible to select the proper IEs to assign the execution of the service chain tasks. Selection is done according to specific criteria, for example through the minimization of specified KPI, as with end-to-end application latency. Interacting with the

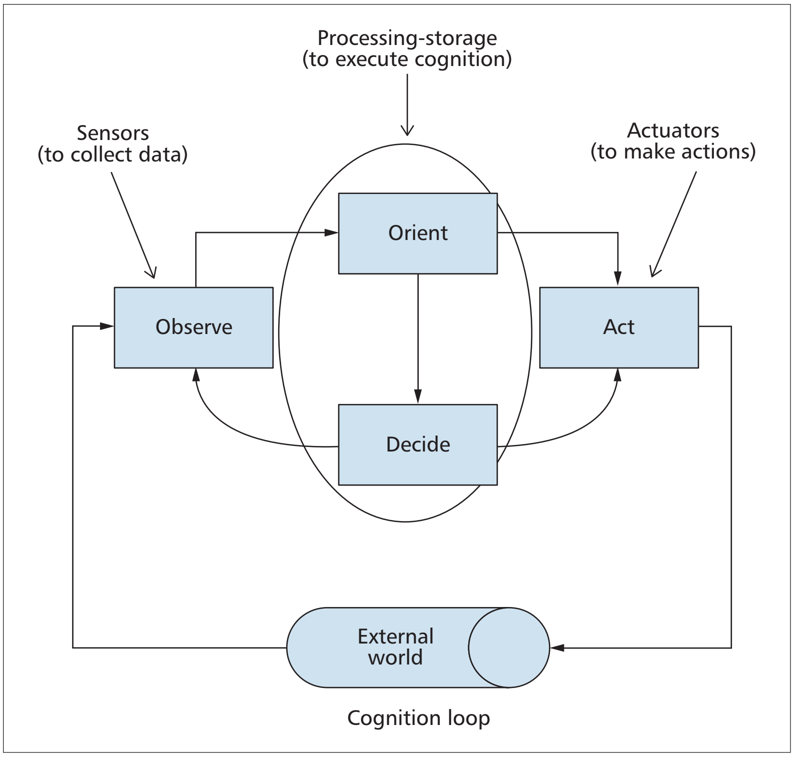

Figure 2. Cognitive loop implementable in an IE.

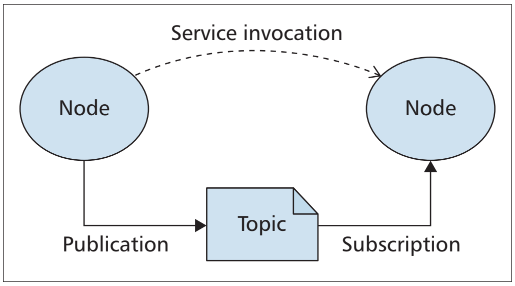

Figure 3. ROS: pub-sub model.

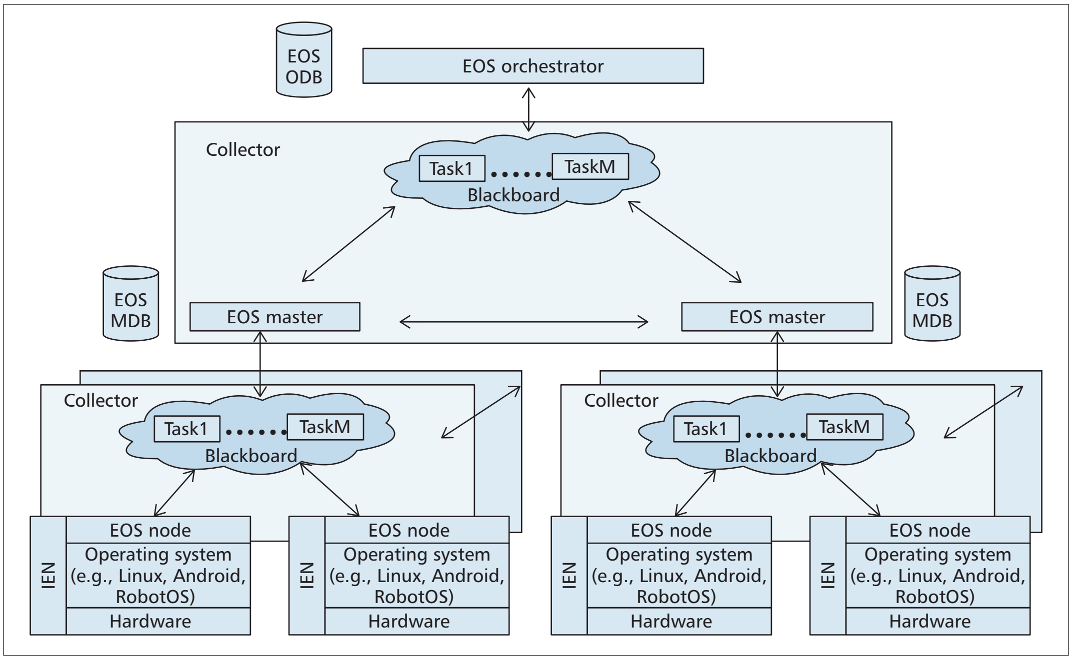

Figure 4. High-level functional architecture of the EOS.

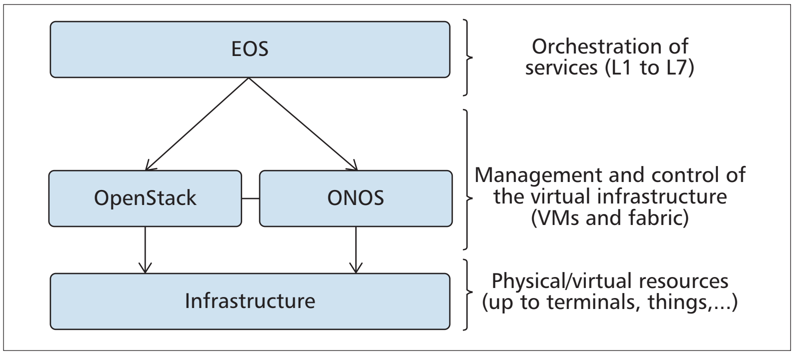

Figure 5. EOS prototype.

EOS node, the EOS master can configure the logical resources hosted by the IEs at run time.

- The EOS orchestrator is a higher-level S/W module that is set up to receive the service request that must span multiple edge domains. It decomposes the request in a service chain, selecting and interacting with the appropriate EOS masters for the end-to-end allocation of IEs across multiple domains. The EOS master can communicate with the EOS orchestrator to whom it registers and updates its status. This data is stored in the EOS orchestrator data base (EOS O-DB).
- The orchestrator blackboard is a higher-level virtual repository shared between the EOS orchestrator and the EOS masters. As above, for the master blackboard, the collector concept can be applied at this level also for the orchestrator blackboard.

The EOS software architecture can be seen as an expression of the integration of the SDN functional architecture defined in ONF (e.g, EOS nodes are like controllers) and the NFV reference architecture framework by ETSI NFV (e.g., an EOS master has VIM capabilities).

Next we briefly describe an example that shows the functioning of EOS. At the startup of an IE, the EOS node sends the EOS master a description of the associated IE services and the status (including the configuration) of the resources. The description can adopt a variety of formats, e.g., YANG modeling [13].

The users' service requests are sent to the EOS master through their terminals (which run EOS nodes). If a service request can be executed just locally, within the EOS master domain, then it is simply decomposed into a sequence of service components and required network functions (i.e., a service chain).

The EOS master then publishes (pub primitive) the software tasks of the service chain. EOS nodes subscribe (sub primitive) to said tasks if the related IEs can execute them. At the end, the EOS master must make an optimized selection of the IEs (selection method). On the other hand, if the EOS master realizes that the service request cannot be executed locally, it forwards it to the EOS orchestrator. In turn, the EOS orchestrator decomposes the service request and publishes it on its blackboard. The flow of actions then proceeds as above within each edge domain.

The EOS is a distributed software architecture where the states of the resources are store in distributed DBs. The well-known CAP theorem [14] will dictate some limitations. In fact, it states that any networked shared-data system can have at most two of the following three properties: 1) consistency (C) equivalent to having a single up-to-date copy of the data; 2) high availability (A) of that data (for updates); and 3) tolerance to network partitions (P).

The general idea is that two of the three properties have to be privileged (CP favors consistency, AP favors availability, and with CA there are no partitions), a trade-off that will be needed then for storing/configuring the states of the infrastructure while achieving specific performance levels. End-to-end latency (or delay) and partitioning are deeply related, and such relations become more important in the case of a widely distributed infrastructure. This situation contributes even more to the requirement of minimizing the application end-to-end latency. These areas require further investigation.

Use Case: Mobile Cognitive Machines

The use case described in this section aims at both the definition of main challenges and requirements for EOS and the feasibility demonstration of a prototype. The main concept of the use case is about the pervasive adoption of mobile cognitive machines (Fig. 2) provisioning any sort of ICT services.

Already today we are witnessing growing interest in using drones, robots, and autonomous machines in agriculture, industry, security, and several other domains. For example, the advent of robots remotely controlled via 5G connections would create a tremendous impact on Industry 4.0. Also, the contexts of the Tactile Internet and cyber physical systems envision several applications for cognitive machines.

In all these contexts, among the major requirements for the EOS there will be, for example, the optimal allocation of logical resources while minimizing end-to-end network and application latencies. Let's see how this requirement has been taken into account in the EOS prototype design and development.

The EOS prototype leverages available open source software complemented with the development of other required software modules. In particular, the two main pieces of open source software are OpenStack and ONOS. The former will be used to manage the virtual machines executing the network and service functions of the virtual infrastructure; the latter will be in charge of managing the fabric of connections, while executing the control applications. EOS can be seen as an overarching operating system that runs on top of both OpenStack and ONOS.

The software architecture of the EOS prototype is shown in Fig. 5, where the circle represents code additions to OpenStack. These code additions mainly address the capability of OpenStack to handle chains of VMs (i.e., service chains) and the Nova-scheduler of OpenStack, which currently uses algorithms (i.e., Filter&Weight) for scheduling VMs in isolation, without considering the status of the underlying network links.

Looking at Fig. 5, from a purely architectural viewpoint, EOS looks similar to XOS [15]. On the other hand, there are major differences that should be highlighted.

XOS is a service orchestration layer that manages scalable services running in a central office re-architected as a datacenter (CORD). On the other hand, EOS is a highly pervasive software architecture, i.e., it is extended up to the terminals (current and future ones), smart things, and elements of capillary networks (e.g., aggregations of sensors, actuators, etc.).

XOS unifies management of a collection of services that are traditionally characterized as being NFV, SDN, or cloud specific. EOS also addresses services that can also be executed at the edge (also leveraging edge and fog computing). This difference dictates profound implications, i.e., EOS is implemented with a scalable software architecture leveraging a trade-off between top-down and bottom-up intelligence.

In the specific use case shown in Fig. 6, a mobile cognitive machine is seen as an IE, with its own operative system. The control system of a mobile cognitive machine usually comprises many ROS nodes. For example, one node controls a laser range-finder, one node controls the wheel motors, one node performs localization, one node performs path planning, one node provides a graphical view of the system, one node for the cognitive service logics, and so on. Other remote ROS nodes may be required to provide other services. In this perspective, it can be argued that ROS nodes can be seen as service components of a service chain executed over the EOS.

Let's focus on a service request to allow a mobile cognitive machine to perform some articulated task (e.g., at the scene of a disaster) with ultra-low reacting times (e.g., a mobile robot being controlled remotely to act in an environment that is changing dynamically).

These requirements are dictating the need to execute said cognitive service by using a proper balance of local, edge, and centralized processing-storage resources. In fact, the machines' reaction times very much depend on IT response time and network latency, and even small changes in the area's layout, or delays in the actuation commands, can lead to catastrophic failures.

EOS, with the software architecture described previously, will be able to exploit this intelligence. For example, the selection method makes it possible to select the proper IEs to assign the execution of the service chain tasks, minimizing end-to-end network and application latencies.

Concluding Remarks

Broadband diffusion and ICT performance acceleration, coupled with cost reductions, are boosting innovation in several industrial and society sectors, thus creating the conditions for a socio-economic transformation, called softwarization. In particular, softwarization of telecommunications will make possible virtualizing network and service functions and executing them in software platforms fully decoupled from the physical infrastructure.

This article has focused attention on the edge of telecommunications infrastructures, arguing that softwarization will transform it in very powerful software platforms, enabling anything-as-a-service. EOS software architecture is proposed to achieve this, even in the short term. In fact, the development of EOS leverages available open source software. A use case was described to validate an EOS prototype with a proof-of-concept.

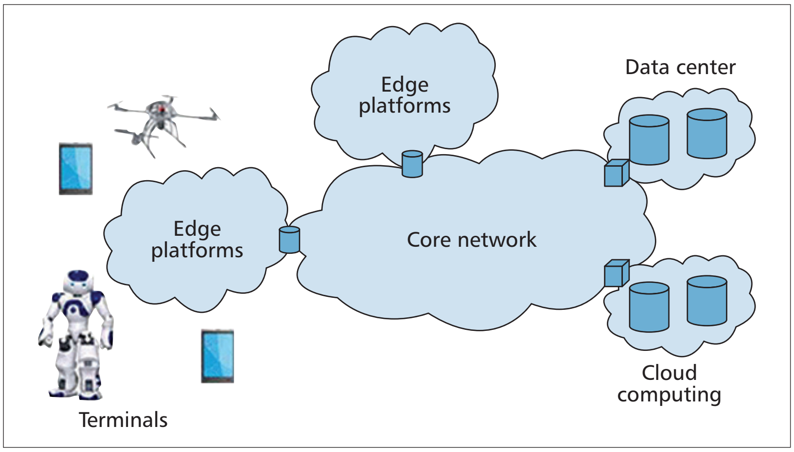

Figure 6. Use-case.

References

[1] "Software-Defined Networking: The New Norm for Networks," *Open Netw. Found.*, Palo Alto, CA, USA, white paper, available: https://www.opennetworking.org/.

[2] "Network Functions Virtualisation," *Eur. Telecommun. Std. Inst. (ETSI)*, Sophia-Antipolis Cedex, France, white paper, available: http://portal.etsi.org/NFV/NFV_White_Paper.pdf.

[3] "Mobile-Edge Computing," *Eur. Telecommun. Std. Inst. (ETSI)*, Sophia-Antipolis Cedex, France, white paper, available: https://portal.etsi.org/Portals/0/TBpages/MEC/Docs/Mobile-edge_Computing_-_Introductory_Technical_White_Paper_V1%2018-09-14.pdf.

[4] "OpenStack," available: http://www.openstack.org/.

[5] "ONOS," available: http://onosproject.org/.

[6] A. Manzalini *et al.*, "Software-Defined Networks for Future Networks and Services," presented at the *IEEE Workshop SDN4FNS*, Trento, Italy, Nov. 2013, white paper, available: http://sites.ieee.org/sdn4fns/whitepaper/.

[7] A. Manzalini and A. Stavdas, "The Network is the Robot," *Communications & Strategies, ERN Economics of Networks eJournal*, Dossier. No. 96 — 4th Quarter 2014.

[8] "Measuring Broadband America," FCC report, available: https://www.fcc.gov/measuring-broadband-america.

[9] "Robot Operating System," available: http://wiki.ros.org/ROS/Introduction.

[10] M. Sloman and E. Lupu, "Engineering Policy-Based Ubiquitous Systems," published by Oxford University Press on behalf of The British Computer Society, 2005; doi:10.1093/comjnl/bxh000.

[11] B. Martini, M. Gharbaoui, and P. Castoldi, "Cross-Functional Resource Orchestration in Optical Telco Clouds," *ICTON 2015*.

[12] R. Alfano and G. Di Caprio, "Turbo: An Autonomous Execution Environment with Scalability and Load Balancing Features," *IEEE Wksp. Distributed Intelligent Systems: Collective Intelligence and Its Applications*, 2006.

[13] "A YANG Data Model for System Management," Internet Engineering Task Force, Aug. 2014, available: https://tools.ietf.org/html/rfc7317.

[14] E. Brewer, "Towards Robust Distributed Systems," *Proc. 19th Ann. ACM Symp.Principles of Distributed Computing (PODC 00)*, ACM, 2000.

[15] "Central Office Re-architected as a Datacenter," available: https://wiki.onosproject.org/pages/viewpage.action?pageId=3441030.

Biographies

Antonio Manzalini (antonio.manzalini@telecomitalia.it) received the M. Sc. Degree in electronic engineering from the Politecnico of Turin. In 1990 he joined CSELT, which then became part of Telecom Italia. He is the author of a book on network synchronization (for SDH), and his RT&D achievements have been published in more than 100 papers. He was active in ITU as Chair of ITU-T Questions, leading standards developments on transport networks. Since 2000 he has had leading roles in several EURESCOM and European Projects. In 2008 he was awarded with the International Certification of Project Manager by PMI. In 2013 and 2015 he led two projects on software defined networks (SDN) funded by EIT – Digital. He is currently a senior manager in the Strategy & Innovation Dept. (Future Centre) of Telecom Italia. His main interests are innovative architectures and services enabled by SDN and NFV, primarily for 5G. He is General Chair of the IEEE initiative on SDN.

Noel Crespi holds masters degrees from the Universities of Orsay (Paris 11) and Kent (UK), a diplome d'ingénieur from Telecom ParisTech, and a Ph.D. and habilitation from Paris VI University (Paris-Sorbonne). In 1993 he joined CLIP, Bouygues Telecom and then went to Orange Labs in 1995. He took leading roles in the creation of new services with the successful conception and launch of Orange's prepaid service, and in standardization. In 1999 he joined Nortel Networks as a telephony program manager, architecting core network products for EMEA region. He joined Institut Mines-Telecom in 2002, and is currently a professor and program director, leading the Service Architecture Lab. He is also an adjunct professor at KAIST, an affiliate professor at Concordia University, and is the scientific director of the French-Korean laboratory ILLUMINE.

SERVICE DESCRIPTION IN THE NFV REVOLUTION: TRENDS, CHALLENGES AND A WAY FORWARD

The authors review the proposals of service description in the main initiatives related to the NFV arena. They elaborate on key novel challenges and evaluate how the different proposals solve them. They propose a straw man model of service and resource description addressing these challenges and defining the features that would serve as design directions for future initiatives and updates in this topic.

Jokin Garay, Jon Matias, Juanjo Unzilla, and Eduardo Jacob

COMMUNICATIONS STANDARDS

ABSTRACT

The telecommunications landscape has been undergoing a major shift in recent years. Initially Software Defined Networking (SDN), and then Network Function Virtualization (NFV) have opened up new ways of looking at the increasingly demanding service provider scenario. The description of the service to be provided will be a key point determining the success in the integration and interoperability of the different proposals. However, the refined understanding of the future scenario and its requirements has recently introduced unique challenges in the path to fully achieve the benefits of the NFV vision. In this paper we review the proposals of service description in the main initiatives related to the NFV arena. Then we elaborate on key novel challenges and evaluate how the different proposals solve them. Finally, we propose a straw man model of service and resource description addressing these challenges and defining the features that would serve as design directions for future initiatives and updates in this topic.

INTRODUCTION

Network Function Virtualization (NFV) [1] offers the promise of flexible and efficient service delivery to network operators [2], leveraging the benefits of virtualization technologies to break the strong coupling in current networks between the services offered and the resources supporting them. From the traditional approach, mainly built around hardware appliances, NFV proposes an evolution to consolidate the required functions into industry standard high volume servers, switches, and storage. In the NFV vision, virtualized network functions (VNFs) are dynamically deployed over the infrastructure to create and manage network services.

There are currently different trends pushing for this vision, either specifically launched in the wake of NFV or converging from related areas. One of the foundations underlying the different proposals is the description of the service to be provided, and its evolution will be a key point determining success in the integration and interoperability of the different proposals.

In this paper we review the proposals of service description in the main initiatives related to the NFV arena. Then we elaborate on novel key challenges, appearing as understanding of the future scenario and its requirements refines and evolves, and evaluate how the different proposals solve these challenges. Finally, we propose a straw man model of service and resource description, addressing these challenges and providing design directions for future initiatives and updates in this topic.

The rest of the paper is organized as follows. We cover the approaches for service description. We detail the challenges and how the initiatives meet them. Finally, we describe the proposed model and end the paper with our conclusions

RELATED WORK

In order to provide a short yet comprehensive view of the alternatives in service description, we have focused on the following:

- The definition from the European Telecommunications Standards Institute (ETSI) as initial proponents of NFV.
- The work carried out in the Internet Engineering Task Force (IETF) for the impact of the RFCs from this organization.
- The approach in OpenStack as the de facto open source standard in cloud computing.
- The standards from the Organization for the Advancement of Structured Information Standards (OASIS) for its orientation to interoperability.

Each of these alternatives targets somewhat different problem spaces and defines its own set of requirements, yet the service description is a key part of all of them.

ETSI NETWORK FUNCTIONS VIRTUALIZATION

The ETSI NFV Industry Specification Group (ISG) has been leading the way since the publication of the seminal white paper[1] that launched the NFV idea and called for action. Since then the ISG has published the architecture [3] defining the main components and the Management and Orchestration (MANO) framework [4], among many other documentation.

According to ETSI, network service (NS) is the "composition of network functions and defined by its functional and behavioral specification." Following this approach, the NS can be defined as a set of VNFs and/or physical network functions (PNFs), with virtual links (VLs) interconnecting them and one or more virtualized network function forwarding graphs (VNFFGs) describing the topology of the NS. The VNFFG in turn contains network forwarding paths (NFPs) that describe a traffic flow in the NS based on policy decisions. Figure 1(a) shows the elements included in a NS; Fig. 1(b) represents a single NS with multiple VNFFGs and NFPs defined.

The processes of service deployment and overall lifecycle management rely on the information elements describing the NS and its components, both as templates in a service catalog and as records of running instances. Both the

The authors are with the University of the Basque Country UPV/EHU.

[1]Network Functions Virtualisation: Introductory White Paper (2012): http://portal.etsi.org/nfv/nfv_white_paper.pdf

VNFs and VLs contain the resource requirements that will be used in the orchestration of the NS, together with the VNF-FGs to have the complete connectivity information. In contrast, the PNFs contain the requirements for the attached VLs, as the PNF, by definition, covers its own resource requirements and cannot be deployed in locations other than its own. The VNF also describes its operational behavior requirements for life cycle management.

IETF Service Function Chaining

The IETF Service Function Chaining (SFC) Work Group (WG) focuses on the definition of a new approach to service delivery and operation, built around the idea of an abstract view of the required service functions and the order in which they are to be applied. Currently they have published the problem statement, and the architecture as RFCs [5, 6].

The architecture is defined around three main components that are deployed in an SFC domain and which compose the service, as depicted in Fig. 2 and detailed below. It relies on the SFC encapsulation, which includes metadata to be carried between service functions and the identification of the path to be followed for the service function forwarder.

- Service functions (SFs): Functions responsible for specific treatment of received packets.
- Service function forwarders (SFFs): Responsible for forwarding traffic to/from one or more connected SFs, as well as to other SFFs.
- Service classification functions (SCFs): Used to select which traffic enters an SFC domain. The initial classification determines the SFC that must process the traffic, and subsequent classification can be used to alter the sequence of SFs applied.

As per the WG charter, the focus of SFC is oriented to the operation and composition of the service itself, aiming more for interoperability of the SFs from different vendors than for defining a detailed model of the service components or the processes to manage the service deployment and life cycle.

Inside the parallel organization Internet Research Task Force (IRTF), there is also the Network Function Virtualization Research Group (NFVRG). Currently it is focused on near-term work items that do not have in their scope the definition of a service description.

OpenStack

OpenStack is a major player in the cloud computing technology field. The project aims for simple implementation, massive scalability, and a rich set of features. Initially oriented toward the Infrastructure-as-a-Service model, it was a natural alternative to the infrastructure layer in NFV. In recent years, OpenStack has been extending its features to address several challenges, core to NFV as well as other cloud computing use cases, such as orchestration and advanced networking capabilities.

Heat is the OpenStack component for orchestration and defines the Heat orchestration template (HOT)[2] to describe the infrastructure for a service, called a cloud application in OpenStack. A HOT has three main components.

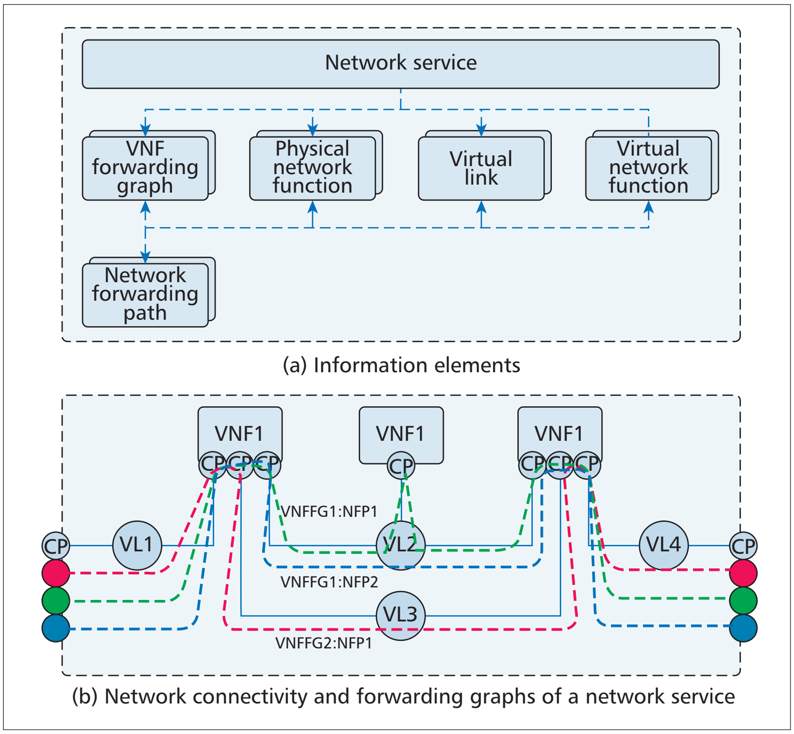

Figure 1. ETSI NFV network service.

- Input parameters to specify the information that has to be provided when the template is instantiated, thus allowing the customization of the instances to be deployed.
- Resources to define the actual resources (in OpenStack compute instances, networks or storage volumes) that will have to be instantiated, allowing also the definition of dependencies between them so the deployment sequence can be controlled.
- Output parameters to define which values from the instantiation process will be fed back to the requester of the deployment.

The application life cycle is also managed from Heat, which will keep track of the resources assigned to the deployed template, although this information is not explicitly included in the information model.

OASIS TOSCA

Topology and Orchestration Specification for Cloud Applications (TOSCA)[3] is an standard from OASIS that targets interoperable deployment and life cycle management of cloud services when the applications are ported over alternative cloud environments.

The core TOSCA specification defines a language and metamodel to describe services, its components, relationships and management procedures. The major elements defining a service are depicted in Fig. 3 and detailed as follows:

- A topology template defines the structure of a service as a set of node templates and relationship templates that together define the topology model as a (not necessarily connected) directed graph.
- Node and relationship templates specify the properties and the operations (via interfaces) available to manipulate the component. Relationships link different nodes and can

[2] Heat orchestration template (HOT) specification: http://docs.open-stack.org/developer/heat/template_guide/hot spec.html

[3] Topology and Orchestration Specification for Cloud Applications (2013): http://docs.oasis-open.org/tosca/TOSCA/v1:0/os/

have diverse meanings (e.g. a relationship between a "process engine" node and an "application server" node could mean "hosted by").
- Plans define the process models that are used to create and terminate a service as well as to manage a service during its entire lifetime.

The TOSCA NFV profile[4] defines an NFV data model using TOSCA language aligned with the ETSI definition. ETSI network services are represented as TOSCA service templates, whereas node templates are used for the rest of the information elements (including virtual links) and relationships to describe which elements are connected through each link.

[4] TOSCA Simple Profile for Network Functions Virtualization (2015): http://docs.oasis-open.org/tosca/tosca-nfv/v1.0/csd01/tosca-nfv-v1.0-csd01.pdf

Challenges

One of the challenges emerging in the NFV scenario is the consideration of hierarchical orchestration, as there are different aspects leading to this approach. First, we must consider the scalability of the orchestration process. The services will most likely be deployed over infrastructures covering significant geographical scopes, mixing resources deployed over access, aggregation, and core networks, and even reaching resources managed by third parties. Relying on a single orchestrator to handle such a wide variance of resources would significantly hinder the scalability of the process. Also, the different capabilities of the involved infrastructure domains (e.g. operator points of presence, transit networks, data centers, etc.) would be more efficiently used by specialized orchestration processes rather than a common, global orchestration.

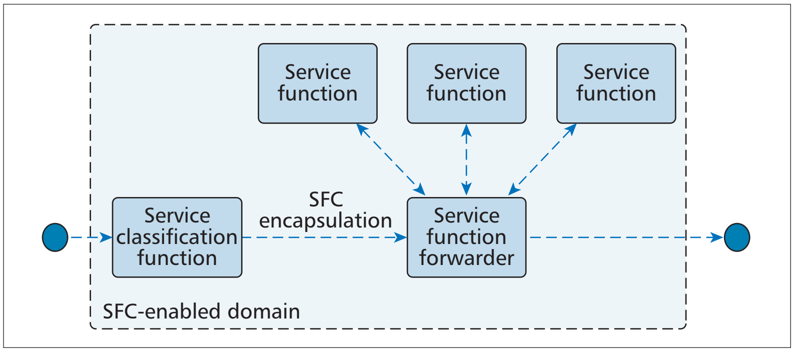

Figure 2. Service function chain architecture.

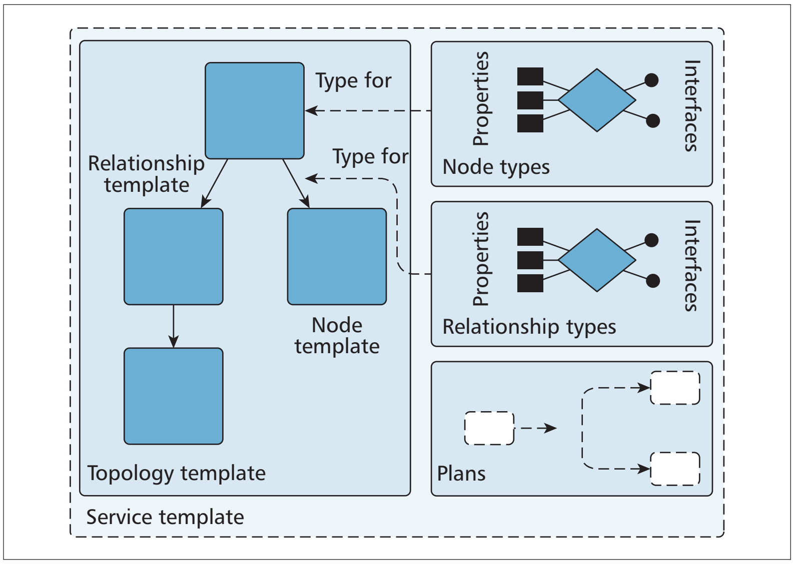

Figure 3. TOSCA service template elements and relations.

A related challenge is the intrinsic multidomain nature of service deployment. Taking the entire, end to end service, few use cases will be confined within the boundaries of a single domain. All the services available for end users outside the domain must consider how the users will access the service. As for the hierarchical orchestration scenarios discussed before, they need to take care of the interconnection of the segments of the service deployed across the different domains. Finally, any infrastructure involving distinct business or administrative domains will also face a similar situation.

Another challenge to be considered is the expected variability of the service across its life cycle, and in the deployment process itself. Network services will be built as a mix and match of the different available network functions, but this process will probably be carried out in multiple steps by different actors. To the original definition from the user, the service provider could transparently add accounting or security related functions, which would ultimately be expanded by functions supporting resiliency or scalability in the infrastructure. Multidomain scenarios, as presented before, could also require interdomain adaptation functions. Finally, internal re-optimization processes, triggered for example by policies, service modification requests, or external changes (e.g. infrastructure updates, auto-scaling) would also require updating the service.

Finally, there is the importance of the associated resource model. The embedding process [7] is one of the hot topics in NFV, and relies on the alignment of the resource requirement information from the service with the resource description of the infrastructure. This premise must be kept even in the face of the aforementioned challenges. Hierarchical orchestration would also imply the necessity of carrying over certain resource assignment information if embedding decisions are taken at the different levels.

ETSI NFV MANO details resource requirements in the NS, but does not extend to the resource model. It is more oriented to a single-layer approach with resource requirements described at the lowest level of detail (CPU, PCIe parameters, etc.). The defined approach for PNFs slightly disrupts the homogeneity of this information, as it also contains resource requirements for the links. The interdomain scenario is not specifically addressed in the NS, which references endpoints that have no information element defined. The different elements detailing connectivity information (VLs, VNF-FGs, and NFPs) add complexity to the additional operations needed when considering hierarchical scenarios or service variability. Finally, resource assignment information is not fully covered. The ETSI model includes resource reservation for the overall NS and the reference of the virtual infrastructure managers that will manage each VL.

IETF SFC focuses mainly on the operation of the service, not how it is described, and treats the SFs as black boxes, considering the chain-

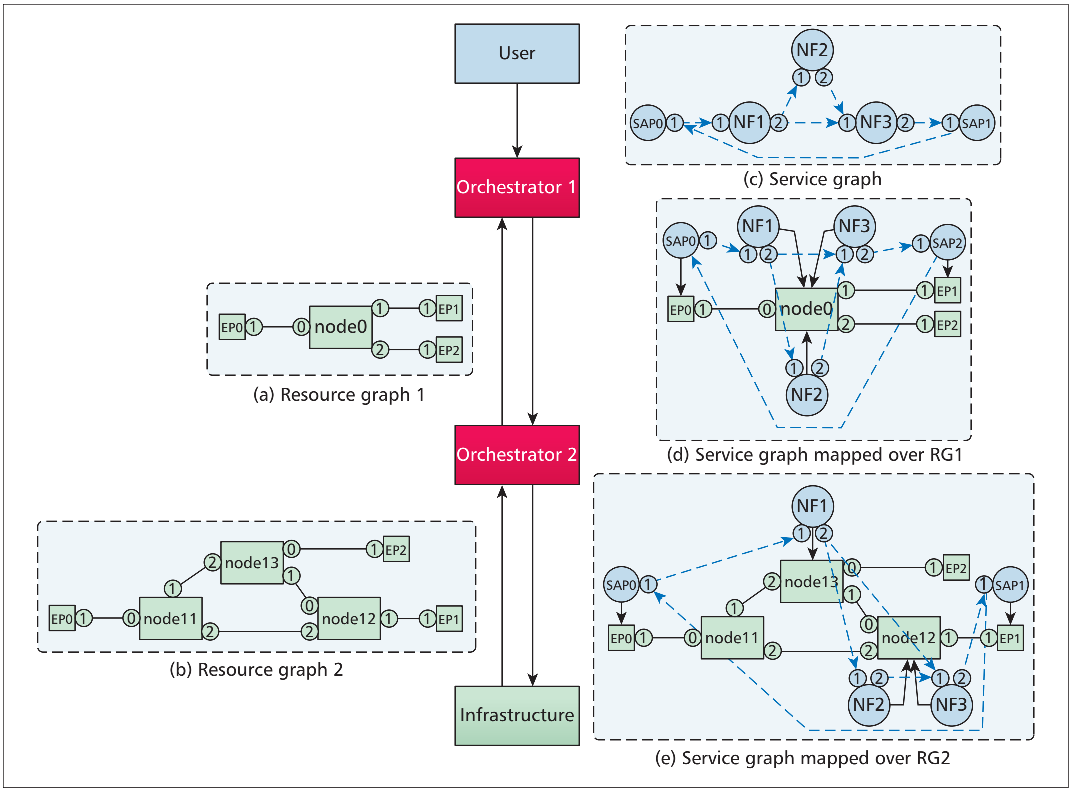

Figure 4. Examples of service graph, resource graph and mapped service graph.

IETF SFC focuses mainly on the operation of the service, not how it is described, and treats the SFs as black boxes, considering the chaining of the SFs and the criteria to invoke them as specific to each domain. One of the targets is for it to be topologically independent from the underlying forwarding topology, thus it does not target resource description.

ing of the SFs and the criteria to invoke them as specific to each domain. One of the targets is for it to be topologically independent from the underlying forwarding topology, thus it does not target resource description. The role of the SFC addresses the requirements of interdomain connections, as it rightly defines the first step of the service is to decide what, in fact, must be processed by the service.

The first consideration about OpenStack stems from its trajectory. Originated in the cloud scenario, it was at its inception more compute-oriented, so its networking capabilities are still not as advanced and lack the required detail (e.g. the connectivity description is more oriented toward networks rather than links). Also, the focus has been more on pure deployment and operation rather than orchestration as defined in the NFV architecture, so its capabilities in this area are still evolving.

Finally, TOSCA, on the one hand, shares with OpenStack its cloud oriented origins, as well as its focus on deployment and operation. On the other hand, the core specification is a metamodel, whereas the detailed specification maps the information elements of ETSI over the TOSCA model, so the same considerations as for ETSI NFV would apply.

Straw-Man Service and Resource Graph Models

One conclusion of the described challenges is that the description of the service and the resources are strongly coupled and will be processed multiple times, so their structure should be closely aligned. Also, this would simplify an explicit mapping of services to resource elements as a result of the embedding. Moreover, the necessity of considering jointly the network function placement and the path calculation in the embedding process [8], as well as the deployment of services including both compute and networking elements [9], points towards a common representation of both. Multiple and variate processing also calls for a careful placement of the information according to its uses, concentrated in the affected elements.

Following these ideas, we propose a service graph (SG) and corresponding resource graph (RG) model described in Fig. 4 and detailed next. It has been designed to provide support for functionalities such as resource orchestration or service deaggregation, on the one hand, and features such as interdomain, scalability, and dynamicity, on the other.

Service Graph

The SG is a directed graph representing the service. It is depicted in Figs. 4(c), 4(d), and 4(e), with round blue nodes joined by dashed arrows, and they contain the following elements:

- Network functions (NFs) are nodes in the SG representing the functions composing the service, including any specification and deployment information, as well as support for life cycle management operations. NFs contain ports for detailing connectivity descriptions (i.e. links connecting the node can be related to a port).
- Service access points (SAPs) are nodes in the SG representing the attachment of the SG to other elements outside the domain. Examples could be "Company branch A, office 1," "All users with Gold service," "Internet," "Service 147685 from Domain x." Optionally,

Information about connectivity or traffic handling policies is associated with SAPs and SLs. Resource requirements are associated with each of the elements, as well as possible placement constraints. Once the embedding process has been carried out, each element will also include the resources assigned from the RG.

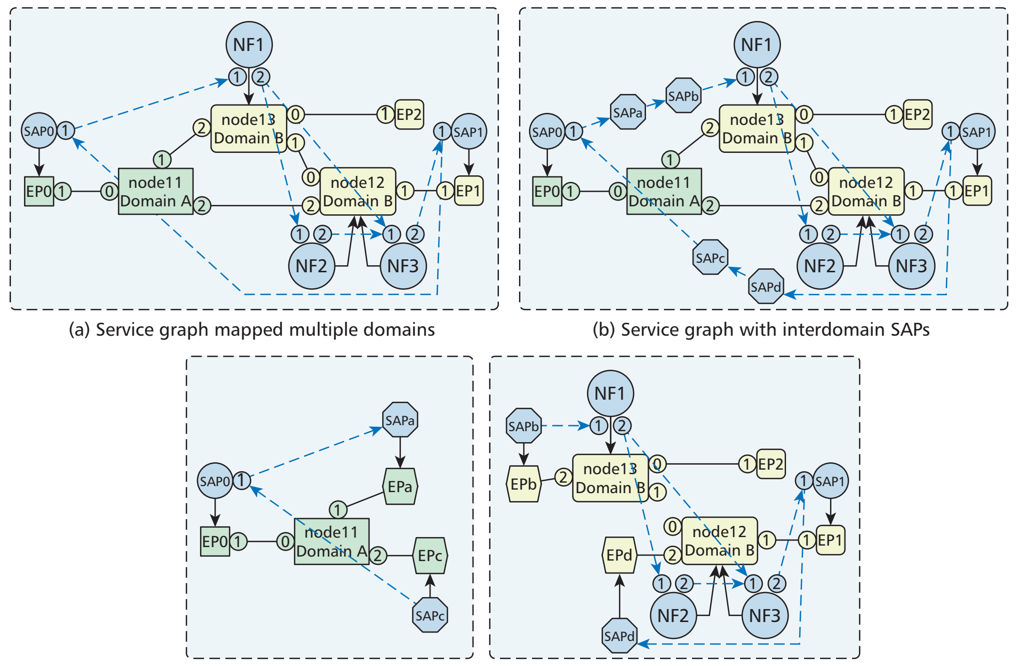

(a) Service graph mapped multiple domains
(b) Service graph with interdomain SAPs
(c) Service graph for Domain A
(d) Service graph for Domain B

Figure 5. Examples of SG deployed over multiple domains.

SAPs can perform classification functions to select the traffic entering the service.

- Service links (SLs) are edges in the SG representing logical connectivity. Optionally, SLs can perform classification functions to select the traffic entering the link.

Information is attached to the SG or its components based on its scope (i.e. applicable to the whole service or specific elements). Examples of information pertaining the SG would be placement constraints, e.g. privacy related (so resources are not shared among SGs or isolation is guaranteed), geographical, etc., and service level agreements (SLAs) applying to the service, such as service availability, end to end latency, etc. This allows an efficient access to the information, either for retrieving it during the embedding process, updating it during the service lifecycle or splitting it when the service is deployed across several domains.

Resource requirements are associated with each of the elements, as well as possible placement constraints (e.g. geographic or legislative restrictions). Scalability and resiliency requirements are also attached to the corresponding elements (or the graph as a whole), whereas information about connectivity or traffic handling policies is associated only to SAPs and SLs. Once the embedding process has been carried out, each element will also include the resources assigned from the RG, as represented in the mapped SGs in Figs. 4(d) and 4(e) for orchestrators 1 and 2 respectively. This placement of the information follows the approach outlined in the previous paragraph, so any process that must handle information from the service at any point of the lifecycle, can efficiently access it based on the affected elements. Progressive refinement of the service is carried out adding or deaggregating NFs (i.e. substituting single NFs for SGs) and reconfiguring the VLs connecting them.

In hierarchical or multidomain scenarios, the SG is split into subgraphs according to the orchestrator or domain responsible for the assigned resources, as represented in Fig. 5: NFs and SAPs are assigned to the corresponding subgraph and VLs connecting elements in different subgraphs are split in two, terminating now in newly introduced SAPs representing the interdomain connectivity.

These SAPs contain all the required information to configure the interdomain connection in each of the corresponding domain endpoints. As each element contains all the related information, the only processing required in the graph splitting is for the information related to the whole graph, on the one hand, and for the SLs traversing interdomain links, on the other hand.

Resource Graph

The RG describes the resources that will be used to deploy the requested services (potentially a directed graph, but could be simplified as undirected in infrastructures with symmetrical links) It provides a homogeneous representation of the infrastructure, in terms of both capacities and capabilities, and is composed of the following elements, represented in Fig. 4 (isolated in (a) and (b), with a mapped SG in (d) and (e)) with green square nodes joined by lines:

- Infrastructure nodes (INs): Nodes in the RG that, depending on their capabilities, can have NFs deployed on them, provide network connectivity or general traffic processing capabilities.
- End Points (EPs): Nodes in the RG that represent a reference point that defines the attachment of the RG to other elements outside the domain in the context of the infrastructure.

- Infrastructure links (ILs): Edges in the RG that represent the connectivity available between the INs.

The resource model included in the RG mirrors that of the SG and is built around three main abstractions, compute, networking, and storage, described in terms of capacities (which are finite and consumed by the requests) and capabilities (which further characterize the resources and are not consumed by the requests). Examples of capacities are the number of vCPUs an IN can handle or the bandwidth of an IL. Examples of capabilities are redundancy for a link, the delay matrix for a node abstracting a network, or the presence of a hardware accelerator for SSL. Support for PNFs, for example, is included in this model based on a capability describing the type of NFs the IN can handle. The elements in the RG are assigned to domains, which group all the elements managed by the same orchestrator entity and determine the splitting to be performed.

Mirroring the top-down process of the SG, which can be extended and split into several subgraphs, we envision the RG being subject to several aggregations and abstractions in a bottom-up process.

In a hierarchical orchestration scenario, each orchestrator would construct a composed RG based on the input of the different infrastructure domains as input for its embedding process. This same aggregation could, in turn, be exposed to the next level, or be aggregated and presented to the orchestrator above, hiding the details of the inter-domain connections, as shown between orchestrators 1 and 2 in Figs. 4(a) and 4(b). In the top-down process, the details would be added back to support splitting the graph between the different domains, as exemplified in Figs. 5(b), 5(c), and 5(d).

Moreover, in these scenarios we envision the resource description to be different in each level, increasing the abstraction in the higher layers, aligned with the service description. In this vision, the two extremes would be a descriptive/quantitative view where the resource description would be oriented to what the resources are (more meaningful in the lower layers, e.g. one IN can offer x vCPUs) and a functional/qualitative view where the resource description would be oriented to what the resources can do (more meaningful in the upper layers, e.g. one IN can support service X for up to 1,000 users).

In this way, in domains with hierarchical orchestration processes, the RG in the higher-level orchestrators would have a wider scope and abstract away the finer grain details of the underlying resources, whereas the RG in the lower-level orchestrators would have a narrower scope and fine grain detail of the resources, thus fostering scalability. Current models target resource description at the lowest level, but in upper orchestrator levels the RG should follow the level of abstraction in the corresponding SG embedding process.

Deployment Process

The mapping between SG and RG elements is stored in the SG elements (both nodes and edges) with the following considerations:

- NFs are mapped to INs. This assignment is maintained during deaggregation of NFs. For shared NFs, the mapping could be extended to specify an instance of NF running in that IN.
- SAPs are mapped to one or multiple EPs.
- SLs are mapped to an IN internal connection (if both ends of the SL are mapped to the same IN), a single IL, or a sequence of them forming a path. A single SL can be mapped into several ILs (or paths) simultaneously to provide multipath links and offer more possibilities for embedding. In such a case, the embedding process must also define the flow space corresponding to each of the available paths that will be active at the same time (as opposed to resiliency, where only a single link/path would be active).

The necessity of SAPs and EPs to also be part of the embedding process can be seen as excessive. The mapping could be more straightforward and require a lighter embedding process than for NFs or SLs, just picking one element from a list of available EPs or even a direct assignment, such as selecting the EP corresponding to the requesting user. Nevertheless, the replication of the SAP and EP elements allows for a coherent and complete description of the SG and RG and prevents changes in the infrastructure (e.g. adding one more endpoint) impacting the service definition. If an SAP is defined so it is mapped to all endpoints of a certain type (for example, all Wi-Fi hotspots in one area) that would just be a change in the infrastructure that would be handled by the orchestration process (not having to modify the service to follow infrastructure changes and vice versa). Also, depending on the scenario, the SAP to EP mapping could be impacted by the embedding process. For example, if the infrastructure on which the service will be deployed has several possible EPs offering connection to the Internet (as an SAP example), the one selected must be reachable (optimally) from the INs in which the NFs are deployed.

In all the aggregation scenarios, the explicit declaration of the mapping between SG and RG allows for clearly identifying the scope for the re-orchestration based on the mapping already done by the layer above, and the relation between the RG exposed to the layer above and the RG received from the layer below. In the example in Fig. 4, if node12 and node13 represent an aggregation of resources (same as node0 represents the aggregation of node11, node12 and node13), the responsible orchestrator would need to refine the placement, choosing a IN among those aggregated in node13 for NF1, and among those aggregated in node12 for NF2 and NF3 (not necessarily the same IN for both but from the same set). Placement of NF1 in a IN aggregated in node12 would not be allowed as it contravenes the placement done by the upper level orchestrator. Following the proposed model, the boundaries within which each orchestrator can function independently are clearly set, thus reducing the need for interaction between the orchestrators in the management of the service lifecycle.

Finally, the SG and RG model allows for two different ways of supporting resiliency: in the SG

In all the aggregation scenarios, the explicit declaration of the mapping between SG and RG allows clearly identifying the scope for re-orchestration based on the mapping already done by the layer above, and the relation between the RG exposed to the layer above and the RG received from the layer below.

NFV is increasingly being recognized as the future direction in service provider scenarios. Multiple efforts are working to bring all the necessary components to the required level of maturity, so the expected benefits can begin to be reaped from actual deployments.

for resiliency managed by the orchestrator or the RG for resiliency managed by the infrastructure. In the first case, the output of the embedding process determines primary and secondary resources for deployment of the SG, and all of them will be assigned in the SG. In the second case, the elements included in the RG offer resiliency and the specific details are hidden from the orchestrator (e.g. a single link in the RG represent multiple different links in the infrastructure, so if any one of them fails, the infrastructure switches to a backup link without this change being propagated to the orchestrator). The orchestrator would select those resources offering the capability and signal the layer below that such capacity must be used and to what extent. Service elasticity, in a similar way, could be achieved in two different ways: either requesting updates to the SG or embedding the elasticity requirements in the SG. In the former, the SG would be updated to modify the resource requirements of NFs or SLs (for scaling up and down) or to add/remove additional NFs and SLs (for scaling out and in). In the latter, the elements of the SG would contain the triggers for the scalability (e.g. in the form of SLAs or specific metrics) and the sequence of actions to perform, so the orchestrator could handle the elasticity. As the proposed model attaches the information to the corresponding element, in both cases the orchestrator would find all the relevant information grouped and would not need to parse any other element but those added/removed in the first case, and those for which a trigger has been met in the second.

Conclusion

NFV is increasingly being recognized as the future direction in service provider scenarios. Multiple efforts are working to bring all the necessary components to the required level of maturity, so the expected benefits can begin to be reaped from actual deployments. A key piece of the puzzle is the service description, required to allow for the different components to interoperate and address the upcoming challenges, and it is still an open topic. This paper presented some novel challenges to be addressed and contributed with a strawman model addressing them, thus fostering the refinement of the service description models toward their successful completion.

Acknowledgment

This research was partly funded by the Spanish Ministry of Economy and Competitiveness under the "Secure Deployment of Services over SDN and NFV Based Networks" project S&NSEC TEC2013-47960-C4-3-P, and by the European Commission under the FP7 UNIFY (Unifying Cloud and Carrier Networks) project CNECT-ICT-619609. This work has been produced within the Training and Research Unit UFI11/16 supported by the UPV/EHU.

References

[1] B. Han *et al.*, "Network Function Virtualization: Challenges and Opportunities for Innovations," *IEEE Commun. Mag.*, vol. 53, no. 2, Feb 2015, pp. 90–97.

[2] E. Hernandez-Valencia, S. Izzo, and B. Polonsky, "How will NFV/SDN Transform Service Provider OpEx?," *IEEE Network*, vol. 29, no. 3, May/June 2015, pp. 60–67.

[3] ETSI, "Network Functions Virtualisation (NFV); Architectural Framework," ETSI, Tech. Rep. GS NFV 002 v1.2.1, Dec. 2014, accessed: 2015/09/01, available: http://www.etsi.org/deliver/etsi_gs/NFV/001_099/002/01.02.01_60/gs_NFV002v010201p.pdf.

[4] ——, "Network Functions Virtualisation (NFV); Management and Orchestration," ETSI, Tech. Rep. GS NFV-MAN 001 V1.1.1, Dec. 2014, accessed: 2015/09/01, available: http://www.etsi.org/deliver/etsi_gs/NFV-MAN/001_099/001/01.01.01_60/gs_NFV-MAN001v010101p.pdf.

[5] P. Quinn and T. Nadeau, "Problem Statement for Service Function Chaining," Internet Requests for Comments, RFC Editor, RFC 7498, Apr. 2015, available: http://www.rfc-editor.org/rfc/rfc7498.txt.

[6] J. Halpern and C. Pignataro, "Service Function Chaining (SFC) Architecture," Internet Requests for Comments, RFC Editor, RFC 7665, Oct. 2015, available: http://www.rfc-editor.org/rfc/rfc7665.txt.

[7] A. Fischer *et al.*, "Virtual Network Embedding: A Survey," *IEEE Commun. Surveys & Tutorials,* vol. 15, no. 4, Feb. 2013, pp. 1888–906.

[8] R. Cohen *et al.*, "Near Optimal Placement of Virtual Network Functions," *2015 IEEE Conf. Comp. Commun. (INFOCOM)*, Apr. 2015, pp. 1346–54.

[9] J. Matias *et al.*, "Toward an SDN-Enabled NFV architecture," *IEEE Commun. Mag.*, vol. 53, no. 4, Apr. 2015, pp. 187–93.

Biographies

Jokin Garay (jokin.garay@ehu.eus) received his B.S. and M.S. degrees in telecommunication engineering in 2003 from UPV/EHU. After a period in the private sector, he returned to the university to work as a researcher and pursue a Ph.D. His research interests include software defined networking, network functions virtualisation, and cloud computing.

Jon Matias (jon.matias@ehu.eus) received his B.S. and M.S. degrees in telecommunication engineering from the University of the Basque Country (UPV/EHU) in 2003, and his Ph.D. from the same university in 2016. He currently works as a researcher in the Communications Engineering Department of the same university. His research interests include software defined networking, network functions virtualization, broadband access networks, and security.

Juanjo Unzilla (juanjo.unzilla@ehu.eus) holds B.S. and M.S. degrees in electrical engineering (1990), and Ph.D. in communications engineering (1999), and is a professor in the Communications Engineering Department at UPV/EHU, where he teaches subjects related to telecommunications networks and services. He is a member of the I2T Research Group, where he participates in several national and European R&D projects. His research interests include SDN and NFV, network security, and techno-economic models for access networks. Among his other interests are models and metrics for knowledge transfer from universities to enterprises.

Eduardo Jacob [SM] (eduardo.jacob@ehu.eus), after spending a few years in the private sector, first as a network manager and later as an R&D project leader, returned to UPV/EHU, where is a professor and leads a research group at the university that is participating in several national and European R&D projects. Among his other interests are industrial applications of SDN and NFV for resiliency, experimental network infrastructures, and cyberphysical systems.

CALL FOR PAPERS
IEEE COMMUNICATIONS MAGAZINE
INTEGRATED COMMUNICATIONS, CONTROL, AND COMPUTING TECHNOLOGIES FOR ENABLING AUTONOMOUS SMART GRID

BACKGROUND

The electricity grid is facing four major challenges — increasing electricity demand, ageing grid infrastructure, ever-increasing penetration of renewables, and significant uptake of electric vehicles. To address these challenges, it is of vital importance to integrate modern communications, control, and computing technologies into one of the most complicated systems on earth, the electricity grid, for building an autonomous smart grid.

The realization of the autonomous smart grid will require collaborative and sustained efforts from the Societies of Power Electronics, Power & Energy, Control, Communications, and Computing over the years to come. This FT aims to facilitate this sustained effort by disseminating general ideas extracted from cutting-edge research results spanning multiple disciplines. Potential authors will be able to share various viewpoints and the newest findings from research and ongoing projects relevant to autonomous smart grid. Topics of interest include, but are not limited to:

- Machine-to-machine control and communications in smart grid
- Vehicle-to-grid control and communications
- Electric vehicle management in smart grid
- Distributed and autonomous control in smart grid
- Smart metering and autonomous demand response
- Internet of Things and its applications in smart grid
- Mobile computing and communications for smart energy management system

SUBMISSIONS

Authors should refer to the submission requirements specified in the "Information for Authors" section of the *IEEE Communications Magazine* guidelines (http://www.comsoc.org/commag/author-guidelines) to prepare their manuscripts. The Guest Editors reserve the right to reject papers they unanimously deem to be either out of scope of this Feature Topic or otherwise extremely unlikely to be accepted after a peer review process. Papers should be submitted through the IEEE Manuscript Central site (at http://mc.manuscriptcentral.com/commag-ieee) to the "December 2016 / Autonomous Smart Grid" " category by the submission deadline according to the schedule shown below.

IMPORTANT DATES

- Submission Deadline: April 1, 2016
- Notification Due Date: August 1, 2016
- Final Version Due Date: September 15, 2016
- Feature Topic Publication Date: December, 2016

GUEST EDITORS

Wei-Yu Chiu
Yuan Ze University, Taiwan
wychiu@saturn.yzu.edu.tw

Hongjian Sun
Durham University, UK
hongjian.sun@durham.ac.uk

John Thompson
Edinburgh University, UK
john.thompson@ed.ac.uk

Kiyoshi Nakayama
NEC Labs America, USA
knakayama@nec-labs.com

Shunqing Zhang
Intel Corporation, China
shunqing.zhang@intel.com

COMMUNICATIONS STANDARDS

A Two-Layered Data Model Approach for Network Services

Connectivity services are ubiquitous to enterprises, and many enterprises are looking to outsource basic networking services traditionally implemented using on-premise network equipment. The rising expectations on service providers to rapidly change the definition of services and the ability to introduce new types of network elements is leading to exploding complexity in the orchestration layer. The severity of this problem is such that the ability to introduce new services and new device vendors in the network is reduced due to the time and cost associated with such changes.

Carl Moberg and Stefan Vallin

Abstract

Connectivity services are ubiquitous to enterprises, and many enterprises are looking to outsource basic networking services traditionally implemented using on-premise network equipment. The rising expectations on service providers to rapidly change the definition of services and the ability to introduce new types of network elements is leading to exploding complexity in the orchestration layer. The severity of this problem is such that the ability to introduce new services and new device vendors in the network is reduced due to the time and cost associated with such changes.

We illustrate that a two-layered data model approach using the YANG language can help overcome these challenges. We use an example describing the process of implementing an IP MPLS VPN service in a network comprising sets of provider edge (PE) routers and customer edge (CE) routers. The example includes a service model and decomposition logic using data transformation, and we show the resulting configuration.

Introduction

Connectivity services are becoming ubiquitous to enterprises, and can be said to be one of the basic requirements for organizations to perform well in a connected world. This has created enormous opportunities for service providers that are able to scale their networks to meet the rising demands for cheap communication services across many geographical markets. As a result, the size of networks in terms of endpoints, as well as the amount of traffic they are required to carry, are expected to rise sharply over time.

At the same time, many enterprises are looking to outsource the networking services that they have traditionally implemented using on-premise network equipment. Examples of such services are end-point security using packet filtering, application acceleration, and collaboration tools. This shift in the location of applications inevitably leads to additional complexity for the service provider offering to host and manage such services on behalf of customers.

There are two factors contributing to the above challenges [1]:

- The wide variety of services and the fact that they change at an increasingly rapid pace as service providers compete in the market.
- The challenge of keeping device-related configuration consistent and aligned with the intent of the deployed services along their lifecycle.

Current orchestration and activation solutions lack formal mapping between the service and device layer configuration models. These systems have historically not been built on data modeling technologies, but have been purpose-built for specific services or types of services. This hard-coding of service definitions and device types assume that the structure and syntax of the input data for service lifecycle operations does not change significantly over time, but that the definition of what constitutes a service and how that is provisioned and activated is relatively static. The rising expectations on service agility (the ability to rapidly change the definition of what services are) and device agility (the ability to introduce new types of network elements from new vendors) leads to exploding complexity in the integration code and dependency on manual steps. The severity of this problem is such that the ability to express new services and introduce new vendors in the network is reduced due to the time and cost associated with such changes.

For example, experience from a large North American service provider shows that the delivery time for point-to-point VPN services for enterprise customers is counted in weeks, even when disregarding time to deliver the physical infrastructure. This delivery time is attributed to numerous manual steps where, for example, work orders are written in natural language in desktop document formats and where network engineers subsequently interpret the orders and translate them into hand-written command line interface (CLI) commands based on which vendors was deployed for a certain geographic area.

We will illustrate a two-layered data model approach using the YANG [2] language that can help overcome these underlying challenges. Splitting the management solution into a service layer and device layer is not new, but is a common pattern in network management software. However, in most deployments these layers are characterized by the following:

- The *service layer* is most often only described in an information-model, not a data model.
- The *device layer* has various application program interfaces (APIs) and command line interfaces (CLIs). APIs are most often RPC-oriented.

These interfaces create a gap between the service instances and corresponding device-configurations, a mapping from informal and often incomplete models to sets of CLI commands or operations to an RPC-based APIs.

Our suggested solution is characterized by the following:

- *Unified YANG data modeling* for both services and devices.

Carl Moberg and Stefan Vallin are with Cisco Systems–Tail-f Engineering, Stockholm, Sweden

- *Transformation* from service models to device models.

The principle has been illustrated in previous work by Vallin *et al.* [3]. It provides the following benefits:
- Using a data modeling language to express service-models introduces strict formalism.
- A single data modeling language for device configuration removes hurdles around the variety of APIs and CLIs.
- Applying the same data modeling language in both layers enables formal transformation and validation techniques.

This leads us to suggest a way forward by decomposing the problem and applying known solutions to the resulting sub-problems.

Applying Data Models

We are suggesting the introduction of two distinct layers of formal data models:
- A *service layer* hosting the entities that represent the order-able services and their components.
- A *device layer* representing the configuration of the participating network elements.

We suggest using the same data modeling language in both layers, allowing for easy application of data transformation. By opening up for the use of transformation, we can leverage well known technologies, allowing for conversion of a set of data values from the data format of a source data system into the data format of a destination data system. Our system uses the YANG data modeling language for these purposes.

The application of data modeling languages provides a number of additional features that we can leverage in our system, including:

Correctness: The specification (data model) is the implementation. With traditional systems, there is always a significant risk that software developers misinterpret informal service specifications and implement them incorrectly.

Completeness: Model mapping specifies only how a service is created; the other operations are automatically generated by the system. With other systems, only a subset of the create, read, update, and delete (CRUD) operations are typically implemented completely. This leads to a need for manual configuration to complement the tasks supported by the system.

Service Agility: The northbound and southbound interfaces, as well as database schemas of the orchestration system, are automatically derived from the data models. With other approaches, these are manual tasks that become increasingly complex as the number of services and device types grows.

The process of defining a new service type starts with asking these fundamental questions:
1. What input parameters do we require when we create the service? That is, what are the parameters sent from a higher-level system such as an order management system or self-service portal?
2. What is the resulting set of device configurations to implement the service instance in the network?
3. What does the mapping between the service parameters and the device configuration look like?

```
module: l3vpn
   +--rw vpn
      +--rw l3vpn* [name]
         +--rw name        string
         +--rw endpoint* [id]
            +--rw id            string
            +--rw as-number     uint32
            +--rw ce
            |  +--rw device    -> /ncs:devices/
                                  device/name
            |  +--rw local
            |  |  +--rw interface-name?     string
            |  |  +--rw interface-number?   string
            |  |  +--rw ip-address?         inet:
                    ipv4-address
            |  +--rw link
            |     +--rw interface-name?     string
            |     +--rw interface-number?   string
            |     +--rw ip-address?         inet:
                    ipv4-address
            +--rw pe
               +--rw device    -> /ncs:devices/
                                  device/name
               +--rw link
                  +--rw interface-name?     string
                  +--rw interface-number?   string
                  +--rw ip-address?         inet:
                    ipv4-address
```

Listing 1. Service Model, Tree Representation.

Informal and Formal Service Models: The answer to the first question above is captured in the service model. The first version of this model is normally an informal description of the data as an unordered list on paper. We use an IP VPN using MPLS transport as an example:
- VPN name.
- BGP AS number.
- A list of end-points defined as:
  - CE device, interface identifier, and IP address.
  - PE device, interface identifier, and IP address.

In this first version we list the PE devices as input. However, if we make sure that our system knows about CE-to-PE links, then the PE devices do not have to be identified in the input, but can be inferred from the topology. This serves to illustrate the power of model-driven systems: our system can handle both scenarios. The decision boils down to the question of how system borders are defined. In this example we use a model where explicit CEs and PEs will be used.

The next step is to transform the informal service model into a formal model expressed in the standardized YANG data modeling language. This activity should be done with a team of stakeholders including the product definition owner, networking subject matter experts, and YANG experts. The steps involved include working through the informal service definition as exemplified above, expressing them using the YANG language and associated best practices. The exact details vary based on the requirements from the product definition owner, the nature of the systems that will operate on the data models, and any expectations on modularity and extensibility of the models over time.

Device Models and Configurations: To answer the second question above: What are the resulting device configurations in the network corresponding to a service instance? This is defined by subject matter experts (in this case network engineers) who

```
module l3vpn {

  namespace "http://example.com/l3vpn";
  prefix l3vpn;

  import ietf-inet-types {
    prefix inet;
  }

  import tailf-ncs {
    prefix ncs;
  }

grouping endpoint-grouping {
    leaf interface-name {
      type string;
    }

    leaf interface-number {
      type string;
    }
    leaf ip-address {
      type inet:ipv4-address;
    }
  }

  container vpn {

    list l3vpn {
      description "Layer3 VPN";

      key name;
      leaf name {
        type string;
      }

      list endpoint {
        key "id";
        leaf id{
          type string;
        }

        leaf as-number {
          description "AS used within all VRF of the VPN";
          mandatory true;
          type uint32;
        }
        container ce {
          leaf device {
            mandatory true;
            type leafref {
              path "/ncs:devices/ncs:device/ncs:name";
            }
          }
          container local {
            uses endpoint-grouping;
          }
          container link {
            uses endpoint-grouping;
          }
        }

        container pe {
          leaf device {
            mandatory true;
            type leafref {
              path "/ncs:devices/ncs:device/ncs:name";
            }
          }
          container link {
            uses endpoint-grouping;
          }
        }
      }
    }
  }
}
```

Listing 2. YANG Service Model.

have experience with manually configuring devices to implement a VPN. Their contribution is the target configuration that should be the result of a service instance. This then is the input to the mapping process described below.

This step can be done either after or before informally defining the service model and its attributes. It might seem that device configurations is too low-level of a concern during service implementation; it is not. Rather, it is the most important step, since it is required to make the service operational in the network. Customer expectations are not satisfied just because the service exists in the supporting systems, and bills are sent. The intent of the service must be configured and activated in the network.

Data Model Mapping: The third step in the process of defining a new service type is to associate each attribute in the service model with the corresponding location in the device models as the basis for our data model transformation. We refer to this as mapping the service model into device models. There are two ways to specify this mapping in our system:

- **Declarative templates:** In many cases a straightforward template that maps service attributes one-to-one to device configuration models is sufficient.
- **Programmatic data model mapping:** In some cases algorithmic expressions are needed to map service attributes to device configuration models. In these cases a programmatic mapping approach can be used.

There is a fundamental difference between data model mapping and the traditional approaches using explicit procedures, workflows, device configuration templates, and other imperative techniques. Data model mapping provides a single mapping describing how service configuration data sets shall be transformed to the device configuration data sets. At run-time this can cover any operation, including instance modifications and deletion. Traditional techniques using workflows or CLI script templates, for example, have a fundamental limitation in that every possible operation on a data set needs an explicit definition. This is not true for data model mapping approaches.

An important question is whether to start with an informal service model and work our way down through a formal service model toward the device configuration (top-down), or start with the device configurations and work our way up to a service model (bottom-up). The preferred direction should be guided by the context at hand and take into account aspects such as, for example, requirements from systems consuming the service models, and if it is at all possible to change the resulting configuration.

An Examples of YANG as the Data Modeling Language for IP MPLS VPN

This example describes the process of implementing a VPN connectivity service in a network comprising sets of PE routers running Cisco IOS-XR and CE routers running Cisco IOS. The process described typically requires a couple of working days for a network engineer. The service model consists of 90 lines of YANG and the service mapping template is 180 lines of XML. The result is a fully functional VPN provisioning system with service-aware northbound interfaces including REST, CLI, Web UI, and a database schema for internal use. It also includes southbound interface drivers toward the network elements. The system has automatically derived CRUD operations including adding, updating, and removing end-points, changing BGP AS numbers while the service is running, and decommissioning of service

```
interface GigabitEthernet0/2
 description Link to PE
 ip address 10.1.1.1 255.255.255.252
exit
interface GigabitEthernet0/9
 description Local network
 ip address 192.168.0.1 255.255.255.0
exit
router bgp 65001
 neighbor 10.1.1.2 remote-as 100
 neighbor 10.1.1.2 activate
 redistribute connected
```

Listing 3. CE configuration.

```
vrf volvo
 address-family ipv4 unicast
  import route-target
   65001:1
  exit
  export route-target
   65001:1
  exit
 exit
exit
interface GigabitEthernet 0/0/0/1
 description link to CE
 ipv4 address 10.1.1.2 255.255.255.252
 vrf volvo
exit
router bgp 100
 vrf volvo
  rd 65001:1
  address-family ipv4 unicast
  exit
  neighbor 10.1.1.1
   remote-as 65001
   address-family ipv4 unicast
    as-override
   exit
  exit
 exit
exit
```

Listing 4. PE configuration.

instances with automatic clean-up of associated device configurations.

In this example we will work top-down, starting with the service model for the VPN and mapping the service parameters to a known set of configuration parameters in the devices. This is then the input data for our transform.

Service Model: The VPN service attributes according to our informal description is structured as a list of the entries, each with the following content:
- VPN name (the list key)
- A list of VPN end-points each with the following structure:
  - –Endpoint identifier (list key)
  - –BGP AS number
  - –Parameters related to a CE device:
    - * CE device identifier
    - * Customer interface identifier
    - * Customer IP address
    - * Uplink interface identifier
    - * Uplink IP address
  - –Parameters related to a PE device:
    - * PE device identifier
    - * Downlink interface identifier
    - * Downlink IP address

```
<!-- CE template for Cisco IOS routers -->
<name>{/endpoint/ce/device}</name>
<config tags="merge">
  <interface xmlns="urn:ios">
    <GigabitEthernet tags="nocreate">
      <name>{link/interface-number}</name>
      <description tags="merge">Link to PE</description>
      <ip tags="merge">
        <address>
          <primary>
            <address>{ip-address}</address>
            <mask>255.255.255.252</mask>
          </primary>
        </address>
      </ip>
    </GigabitEthernet>
    <GigabitEthernet tags="nocreate">
      <name>{local/interface-number}</name>
      <description tags="merge">Local network</description>
      <ip tags="merge">
        <address>
          <primary>
            <address>{ip-address}</address>
            <mask>255.255.255.0</mask>
          </primary>
        </address>
      </ip>
    </GigabitEthernet>
  </interface>
  <router xmlns="urn:ios">
    <bgp>
      <as-no>{../as-number}</as-no>
      <neighbor>
        <id>{pe/link/ip-address}</id>
        <remote-as>100</remote-as>
        <activate/>
      </neighbor>
      <redistribute>
        <connected>
        </connected>
      </redistribute>
    </bgp>
  </router>
</config>
</device>
```

Listing 5. CE XML mapping template.

Listing 1 shows the formal structure of the YANG service data model based on the tree-output from the pyang compiler; Listing 2 shows the full YANG module.

Device Configuration: We start by looking at the configuration in the CE router related to the IP VPN service instance in a network. Listing 3 is the subset of the CE-configuration; Listing 4 is the subset of the PE-configuration directly related to the IP VPN. The highlighted strings are all data that will be referenced from the service instances. This gives us all the output data required in the next step, where we map the service structure into applicable locations in the resulting device configuration according to the above.

Service Mapping: The next step is to define the mapping of service attributes to device configuration parameters. This example will use the declarative template model for the mapping. Our system uses an XML-based templating language to represent the output structure. Listing 5 and Listing 6 are the example XML

```
<!-- PE template for Cisco IOS-XR routers -->
  <vrf xmlns="http://tail-f.com/ned/cisco-ios-xr">
    <vrf-list>
      <name>{string(/name)}</name>
      <address-family>
        <ipv4>
          <unicast>
            <import>
              <route-target>
                <address-list>
                  <name>{../as-number}:1</name>
                </address-list>
              </route-target>
            </import>
            <export>
              <route-target>
                <address-list>
                  <name>{../as-number}:1</name>
                </address-list>
              </route-target>
            </export>
          </unicast>
        </ipv4>
      </address-family>
    </vrf-list>
  </vrf>
  <interface xmlns="http://tail-f.com/ned/cisco-ios-
      xr" tags="nocreate">
    <GigabitEthernet>
      <id>{link/interface-number}</id>
      <description tags="merge">link to CE</description>
      <ipv4 tags="merge">
        <address>
          <ip>{ip-address</ip>
          <mask>255.255.255.252</mask>
        </address>
      </ipv4>
      <vrf tags="merge">{string(/name)}</vrf>
    </GigabitEthernet>
  </interface>
  <router xmlns="http://tail-f.com/ned/cisco-ios-xr"
       tags="merge">
    <bgp>
      <bgp-no-instance>
        <id>100</id>
        <vrf tags="merge">
          <name>{string(/name)}</name>
          <rd>{../as-number}:1</rd>
          <address-family>
            <ipv4>
              <unicast>
              </unicast>
            </ipv4>
          </address-family>
          <neighbor>
            <id>{../ce/link/ip-address}</id>
            <remote-as>{../../as-number}</remote-as>
            <address-family>
              <ipv4>
                <unicast>
                  <as-override>
                  </as-override>
                </unicast>
              </ipv4>
            </address-family>
          </neighbor>
        </vrf>
      </bgp-no-instance>
    </bgp>
  </router>
</config>
```

Listing 6. PE XML mapping template.

templates for the CE-devices and PE-devices, respectively, and the highlighted syntax are references into the service model using XPath [4] expressions.

The XPath expressions in the highlighted strings use relative and absolute paths to reference parts of the service model. When an instance of the service is configured, our system transforms the service instance data by resolving the paths in the device templates in the context of the service instance data. This process produces a set of output data in a device-structure format. This can then be transformed into ordered sequences of commands or operations, depending on the nature of the protocol used toward the devices.

The transformation context also provides some additional data, including for example the device identifier captured in the *device* element, and the per-device namespace that is used to match on which subset of the template to apply. This allows us to have a single XML template defining the output data structures for many vendors and OS versions at the same time.

Conclusion

We have illustrated that a two-layered approach using the YANG data modeling language in both the service layer and device layer can help overcome current challenges to deploying service configurations in networks. The use of data modeling languages allows for application of declarative transformation technologies to the decomposition step between the service layer and device layer. This significantly reduces the complexity in the integration layer between the two, allowing for rapid development of services over time.

We have shown this using an example consisting of a simple IP MPLS VPN service data model and how it is applied to CE and PE routers running Cisco IOS and IOS-XR. The declarative decomposition is expressed using an XML-based device template language using XPath to reference values in service instances.

Acknowledgment

Our big thanks go to the Tail-f engineering team.

References

[1] W. Enck *et al.*, "Configuration Management at Massive Scale: System Design and Experience," *Proc. 2007 USENIX: 21st Large Installation System Administration Conf. (LISA '07)*, 2007.
[2] M. Bjorklund, "YANG — A Data Modeling Language for the Network Configuration Protocol (NETCONF)," RFC 6020, Oct. 2010.
[3] S. Wallin and C. Wikström, "Automating Network and Service Configuration Using NETCONF and YANG," *Proc. 25th Int'l. Conf. Large Installation System Administration*, 2011.
[4] W3C "XML Path Language (XPath) 2.0 (2nd Edition)," 2010.

Biographies

Carl Moberg (camoberg@cisco.com) is a technology director with Tail-f, a Cisco company specializing in OSS and NMS systems, and network management standards. He joined Cisco from the Tail-f Systems acquisition. Carl spends his time across standards work, large scale implementations, and product-oriented technology strategy. He has presented extensively on the subject of model driven network management, with a focus on NETCONF and YANG and is a member of the IETF YANG Model Coordination Group.

Stefan Vallin (svallin@cisco.com) contributes more than 20 years of network and service management experience to Tail-f, a Cisco company specializing in OSS and NMS systems, and network management standards. He has presented at a wide variety of international network engineering and management conferences. Before joining Tail-f, he worked as a OSS solutions architect at Data Ductus. Dr. Vallin also served as a network management specialist at Ericsson and a research engineer at the University of Linköping. He holds a M.Sc. degree in computer science from Linköping University, and received his doctorate in network management from Luleå University of Technology.