January 2016, Vol. 54, No. 1

IEEE Communications Magazine

•Software Defined Wireless Networks

•Software Defined Radio – 20 Years Later

•Network and Service Management

•Ad Hoc and Sensor Networks

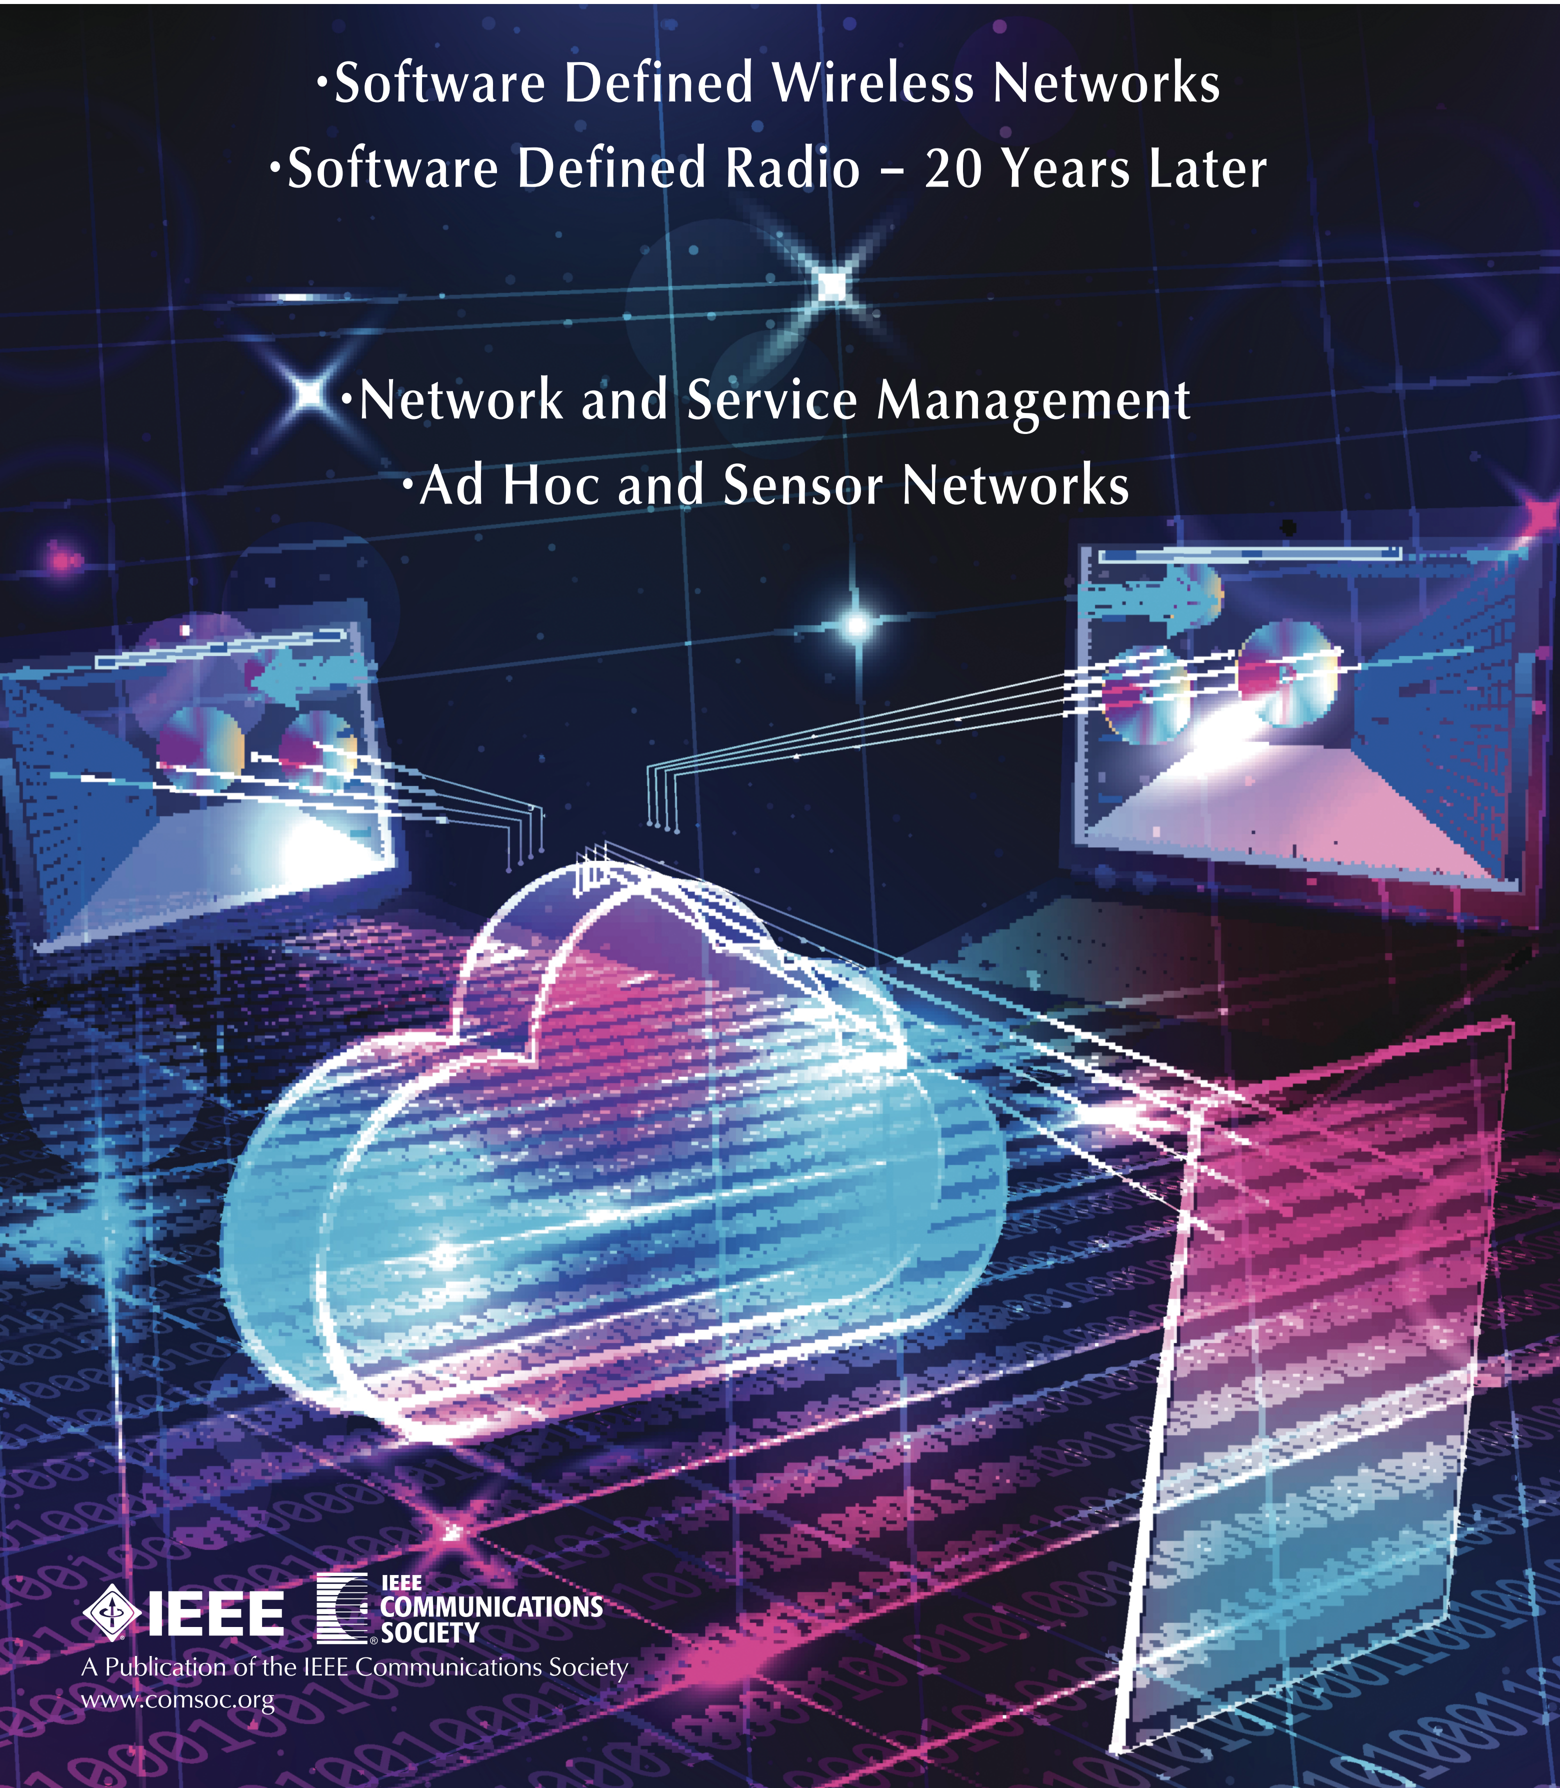

IEEE
IEEE Communications Society
A Publication of the IEEE Communications Society
www.comsoc.org

January 2016, Vol. 54, No. 1

IEEE Communications Magazine

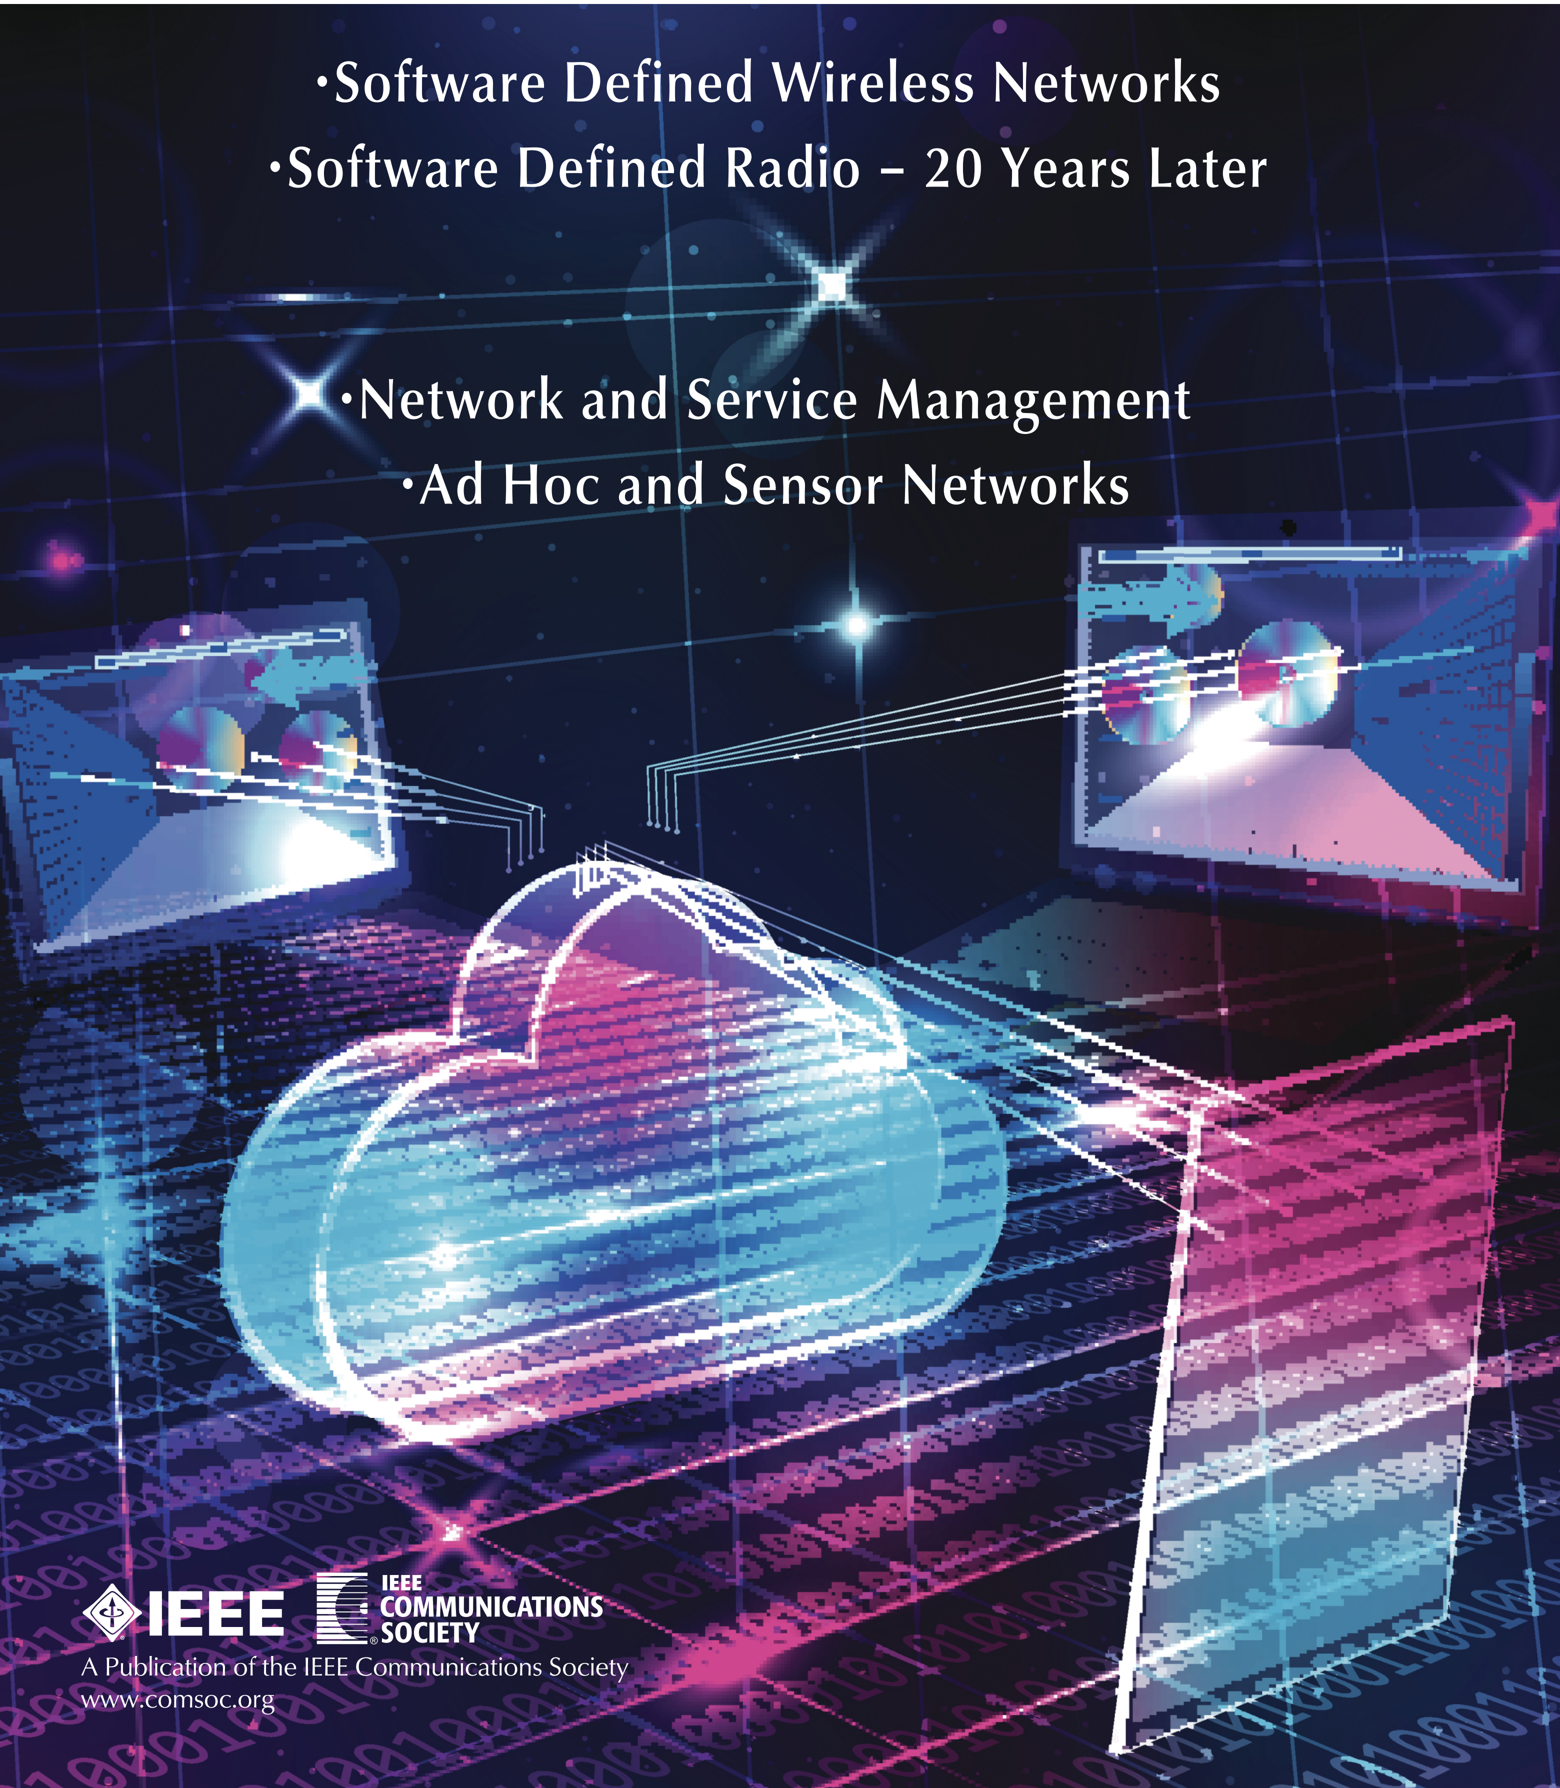

- Software Defined Wireless Networks
- Software Defined Radio – 20 Years Later
- Network and Service Management
- Ad Hoc and Sensor Networks

IEEE
IEEE COMMUNICATIONS SOCIETY
A Publication of the IEEE Communications Society
www.comsoc.org

There's a 5G moment of discovery out there.

We're here to help you find it.

In just a few years, the fifth generation of wireless communications will be a reality. It will allow data to be transmitted up to 100 times faster than today's 4G networks. But getting to that speed won't be easy. You'll need genuine insights to overcome enormous technical challenges. We can help. We have the industry's first and most comprehensive 5G software library. It can significantly streamline design feasibility because it incorporates an iterative design sequence with every piece of Keysight 5G test equipment.

HARDWARE + SOFTWARE + PEOPLE = 5G INSIGHTS

- Designed for testing 5G simulation to verification
- Software platforms and applications that work seamlessly across our 5G instruments
- Incorporate iterative design and rapidly move between stages of your 5G development flow
- Industry's first and largest 5G library

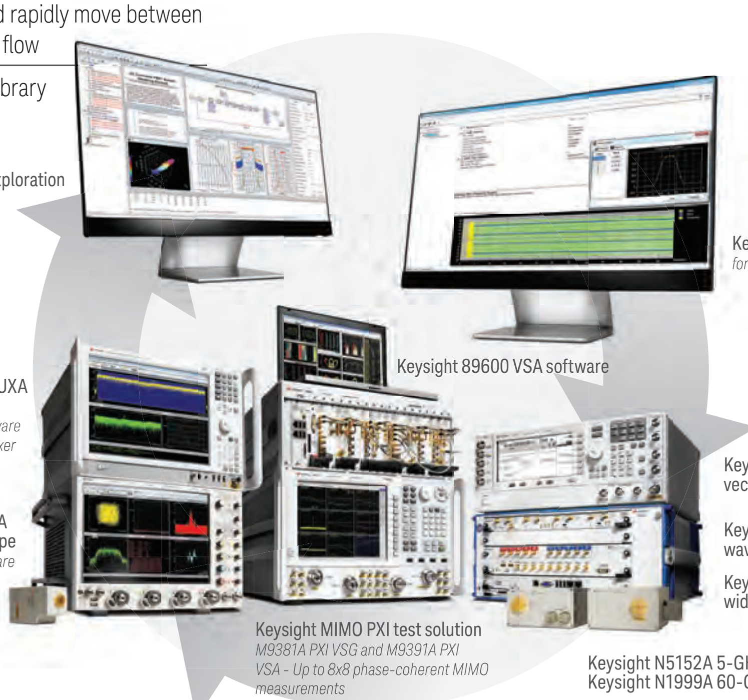

Download our white paper *Implementing a Flexible Testbed for 5G Waveform Generation and Analysis* at www.keysight.com/find/5G-Insight

USA: 800 829 4444 CAN: 877 894 4414

KEYSIGHT TECHNOLOGIES

Unlocking Measurement Insights

IEEE COMMUNICATIONS MAGAZINE

JANUARY 2016, vol. 54, no. 1
www.comsoc.org/commag

SOFTWARE DEFINED WIRELESS NETWORKS: PART 2

GUEST EDITORS: HONGLIN HU, HSIAO-HWA CHEN, PETER MUELLER, ROSE QINGYANG HU, AND YUN RUI

SOFTWARE DEFINED RADIO – 20 YEARS LATER: PART 2

GUEST EDITORS: JOSEPH MITOLA, PRESTON MARSHALL, KWANG-CHENG CHEN, MARKUS MUECK, AND ZORAN ZVONAR

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.

Network and Service Management

Series Editors: George Pavlou and Jürgen Schönwälder

Ad Hoc and Sensor Networks

Series Editors: Edoardo Biagioni and Silvia Giordano

Currently Scheduled Topics

Topic	Issue Date	Manuscript Due Date
Social and Mobile Solutions in Ad Hoc and Sensor Networking	July 2016	January 11, 2016
SDN Use Cases for Service Provider Networks	October 2016	January 31, 2016
5G Radio Access Network Architecture and Technologies	November 2016	February 1, 2016
Enabling Mobile and Wireless Technologies for Smart Cities	December 2016	February 29, 2016
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 February Issue

Underwater Wireless Communications

Optical Communications

Automotive Networking

Introducing industry-leading GaN RF technology for military applications—new from NXP, the global leader in RF power.

In the field, there's no time for second chances—or second best. As the largest, most deployed provider of RF power technology on Earth, NXP now brings its industry-leading technology, design and manufacturing prowess to the defense electronics industry. Our new 10 W GaN transistor provides breakthrough wideband and thermal performance, demonstrated in the available 200–2500 MHz reference design. Our RF power products are proven rugged and reliable in the harshest environments. US manufacturing and a 15-year product longevity warranty ensures long term security of supply. NXP's dedicated support team of military systems and applications experts will back you up every step of the way. Out here, who will you count on? Talk to us. Learn more **NXP.com/RFmilitary**

Moving the Society Forward

Harvey Freeman

After many years as a volunteer and almost as many as a Board of Governors member, I am honored to have been elected as the IEEE Communication's Society's President, and am privileged to serve for the next two years. Recessions, changes in technology, the increasing availability of technical information on the Internet, and the rise of social media, all contributed to the Society's decrease in membership from 60,000 in the year 2000 to 30,000 today. I, along with my new team, plan to reverse this trend and make the Society the go-to place for communications information, standards, technology, and community.

Byeong Gi Lee, a past ComSoc President and the Chair for the past two years of the Society's Strategic Planning Committee, has developed a set of four business plans to implement the ComSoc 2020 Report prepared in 2011. These plans will be a starting point for our work this year, with focus on the following four areas:

1. Education Programs, Content and Services.
2. Industry and Standards Strategies.
3. One-Stop Information and Communications Technology (ICT) Service System.
4. Member Activity and Service Supporting System.

In the education area, we will continue the good work to develop a world-class training and professional education program. We started this by developing a wireless training and certification program. Now we will expand this to other leading edge communications-related technologies such as 5G, Software Defined Networks/Network Function Virtualization, Cloud Computing, Big Data, Cybersecurity, and the Internet of Things. One effort will be aimed at setting up an education and training portal.

For industry, we intend to make ComSoc the preferred organization for communications and networking industry companies and professionals. As a first step toward this goal we have created two new roles: Vice President of Industry and Standards, and a Director of Industry Outreach. We have plans to add to this a new Director of Industry Communities. One important thrust will be to increase our membership from Internet companies like Facebook, Twitter, and LinkedIn by offering them products and services they will find useful.

We will position ComSoc as a focal point for the world's ICT community. One proposal is to implement this by simply taking open source wiki software to develop the service system. We will explore this avenue and others to offer this type of service.

The Member Activity and Service Supporting System should function as "the center" of all membership programs and services at a global level and serve as "the network" through which each individual member gets the service they want and connects with other members with similar interests. We may be able to use open-source CRM systems to implement this without developing something new ourselves.

In addition to bringing more industry professionals back to ComSoc, we will focus on providing additional benefits to ComSoc's younger members. We will work to increase Student Travel Grants, offer reduced membership dues during the first few years out of school, increase networking opportunities to assist them in their jobs and in finding new or more challenging positions, and establish a mentoring program.

These and other ideas that we have been discussing over the past year will take a lot of time and effort. One way that we can make it more manageable is to work together with some of the other IEEE Societies. We have started discussing this in our meetings with the four largest Societies (Computer, ComSoc, Signal Processing, and Power Engineering) held three times a year during the IEEE Board Series. Roger Fugii, the current Computer Society President, came to our Board meeting in December to see what we are currently involved in and to identify areas where we can work together.

Assisting me in these efforts are the five newly elected Vice-Presidents. Guoliang (Larry) Xue is the new Vice President of Conferences. Larry brings a wealth of conference experience to this position, having been involved with conference organizations and ComSoc services for many years. Luigi Fratta, the new VP of Technical Activities, has a diverse background and extensive experience in the communications field. Nelson Fonseca, our new VP of Publications, has been a ComSoc Editor-in-Chief as well as a previous Vice President. Stefano Bregni, Vice President of Member Relations, was reelected to a second term in this position. During the past two years he has led a number of successful initiatives, and will continue these activities as well as launch new ones. Finally, Rob Fish, a leading member of ComSoc's and IEEE's standards efforts, has also been reelected to a second term. For this term the Board of Governors added industry responsibilities to the position now call Vice President of Industry and Standards.

I am highly focused on bringing strong leadership, collaboration, and growth to the Society in the months to come. To do this I will need the help of the elected and appointed officers as well as that of the ComSoc membership. There are many areas where you can help. Please don't hesitate to contact me or any member of the ComSoc Board or Standing Committees with your ideas and offers of help.

Stefano Bregni

Robert S. Fish

Nelson Fonseca

Luigi Fratta

Guoliang (Larry) Xue

GCN
GLOBAL COMMUNICATIONS NEWSLETTER

January 2016
ISSN 2374-1082

CONFERENCE REPORT

Report of the Opto-Electronics and Communications Conference (OECC) 2015

By Xinwan Li, TPC Co-Chair of OECC2015, and Kan Wu, General Secretary of OECC2015, Shanghai Jiao Tong University, China

The OptoElectronics and Communications Conference (OECC) is one of the biggest worldwide conferences focused on the latest optical and photonic research areas. OECC 2015 was held in the Shanghai Everbright Convention Center, Shanghai, China, from June 28 to July 2, 2015. The conference provided a platform for researchers from around the world to share their latest research advances, promoting international collaboration.

The conference was sponsored by the IEICE Communications Society, the IEEE Photonics Society, the IEEE Photonics Society Shanghai Chapter, the IEEE Communications Society Shanghai Chapter, the China Institute of Communications, and the Optical Society of Korea.

There were 470 submissions in total, including 327 contributed submissions from 33 countries or regions around the world. A total of 205 papers were accepted, for an acceptance ratio of 62.7 percent. Each paper was reviewed by at least three peer reviewers. There were 152 oral presentations, 124 invited presentations, and 53 poster presentations, arranged into eight categories, two symposia, and two workshops, listed below.

Categories

C1 Optical Transmission Systems and Subsystems
C2 Optical Networking and Switching Technologies
C3 Optical Fibers, Cables, Devices, and Modules
C4 Optical Fiber Sensors and Microwave Photonics
C5 Laser Technologies and Applications
C6 Micro/Nano Photonic Devices and Integration
C7 Biomedical Optics
C8 Optoelectronic Materials for Communications, Display, and Energy

Symposia

S1 Optical Wireless Communications
S2 Frontier in Stimulated Brillouin Scattering
S3 Design and Fabrication Technology for Photonics Electronics Convergence
S4 Hybrid Nanophotonics

Workshops

W1 Optical Sampling and Photonic Analog-to-Digital Converters
W2 Photonics of Two-Dimensional Materials

Two workshops were organized in the afternoon of June 28. The opening ceremony was organized in the morning of June 29, followed by three plenary presentations. The first plenary presentation was given by Prof. B. Jalali from UCLA on the topic of 'Optical Information Capacity of Silicon'. The second plenary presentation

Plenary presentation.

Poster session.

was given by Dr. Atsushi Takahara from NTT Network Innovation Laboratories on the topic 'Next Challenges with Virtualization of Network Infrastructure'. The third plenary presentation was given by Prof. Chao-Yang Lu from the University of Science and Technology of China on the topic 'Recent Experimental Progress in Quantum Information Processing with Photons and Cold Atoms'.

In the technical sessions, there were 64 parallel oral sessions and one poster session during the conference period. Among 64 oral sessions, there were 10 symposia sessions covering the topics of optical wireless communications, stimulated Brillouin scattering, photonics electronics devices and nanophotonics.

The post deadline paper (PDP) presentations were organized on the evening on June 30. Four post deadline papers were accepted after the revision of technical committee members. They were "Record Field Demonstration of C-band Multi-Terabit 16QAM, 32QAM and 64QAM over 762km of SSMF" by Rahman, Coriant R&D GmbH, Germany, "What is the True Value of Dynamic Optical Path Switching?" by Kurosu, NIAIST, Japan, "First Demonstration of Holistically-Organized Metro-Embedded Cloud Platform with All-Optical Interconnections for Virtual" by Chen, BUPT, China, and "4–28 Gbaud PAM4 Integrated ROSA with High-Sensitivity APD Datacenter Provisioning" by Nakanishi, NTT Corporation, Japan.

The banquet took place on the evening of July 1. During the banquet, best student papers were awarded to three students. A review of the past 20 OECC conferences by sand painting was made successfully. Chinese traditional Sichuan face-changing was also performed.

(Continued on Newsletter page 4)

CONFERENCE REPORT

6th International FOKUS FUSECO Forum (FFF)

Digital Convergence and Seamless Connectivity for Everyone and Everything: Bringing 5G, SDN/NFV and M2M/IoT Together

By Thomas Magedanz, General Chair, Fraunhofer FOKUS/TU Berlin, Germany

FFF 2015 was the sixth event in the series and attracted approximately 200 telecommunication industry experts and academics from 32 countries. FFF 2015 provided a unique discussion platform for representatives from international network operators, standards development bodies, and leading academic organizations offering two full days of technical tutorials, interactive workshops, conferences, exhibition booths, and live demonstrations.

Tutorials and Interactive Workshops

Day one started with two parallel tutorials followed by three simultaneous workshop tracks. The tutorial 'Network Evolution Towards the 5G Environment' provided an overview of use cases and technology covering the main elements of future 5G technologies: radio, core, management, network function virtualization (NFV), and software defined networks (SDN). The second tutorial provided a detailed perspective on Internet of Things (IoT) and machine-to-machine (M2M), data processing, and cyber-physical systems (CPS) required for the Industrie 4.0, industrial Internet and smart cities.

The first workshop, 'Towards the 5G Environment', was divided into two sessions, one addressing the challenges and requirements for the radio and core network, and the other addressing the latest developments and benefits from network virtualization. The second workshop covered the challenges and requirements of information and communication technologies (ICT) for Industrie 4.0, industrial Internet, and the factories of the future, as well as the enablers for ICT within the industrial environment. The third workshop addressed the evolution of IoT and M2M platforms and services for different vertical domains. The benefits of both platform types were analyzed, and it was clear at the end of the workshop that they are key enablers for smart solutions and applications. Innovative M2M and IoT solutions and applications were also analyzed and discussed.

Conferences and Panel Discussions

On the second day the different tracks of the previous day were united under the main theme of this year's FFF by a full day conference entitled '5G Starts Now – Putting the Building Blocks for 5G Together', which consisted of four sessions. Each session was preceded by presentations from various international network operators, vendors, academics, and standardization bodies, concluding with a panel. The topics were: 'The Evolution Path to 5G, Including Radio, Core and Management Networks'; 'The Impacts of SDN and NFV on the Telecom Industry'; 'M2M, IoT and Data as Key Communication Drivers around the World'; and 'CPS: Driving Factories of the Future'.

Prof. Dr. Thomas Magedanz opening the IoT, M2M and Industrial Internet tutorial.

Live demonstration based on the Fraunhofer FOKUS Open5G-MTC Toolkit. Virtualization meets M2M/IoT.

Panel discussion on SDN and NFV.

With respect to 5G it was agreed that for its investment it is important to identify the best ways to monetize the technology and that investment in 5G should be targeted on enhancing the end user experience. This can only be achieved by the convergence of the different access methodologies. Openness will be one of the most important characteristics of the upcoming technology as long as it enables easy connectivity to APIs and innovative services.

Concerning SDN and NFV, standardization is one of the key challenges that needs to be addressed, especially with respect to interoperability and network management. According to the experts, the main purpose of standardization is to avoid proprietary solutions and enhance interoperability. The panel also indicated that important organizational changes and impacts need to be evaluated.

Regarding IoT, it was agreed that it enables the Internet of People and that entrepreneurs will benefit from it. IoT might be deployed without human presence, but it will require coverage. Therefore, coverage will become critical. Usability of the IoT technology for the end user, network simplification and life cycle management for IoT products will also be essential. Security, openness to service providers, and the resolution of data ownership also need to be added to the list.

With respect to best practices in smart cities and Industrie 4.0 communications, the participants agreed on the importance of analytics, connectivity, and low latency to enable digital manufacturing.

Live Demonstrations

Alongside the workshop seven live demonstrations were made. The following technologies and technical challenges were addressed: core networks, next generation M2M connectivity and mobile edge integration in a single solution; secure connectivity for the industrial environment; cost efficient edge-cloud solutions; hands-on examples of product prototyping using NFV and SDN technologies; interconnection of virtual and physical ecosystems in the industrial Internet; combination and integration of multi-da-

(Continued on Newsletter page 4)

CHAPTER REPORT

Activities of the Xi'an ComSoc Chapter

By Prof. Bin Song and Prof. Rong Sun, State Key Lab of Integrated Services, Xidian Univerisity, Xi'an, 710071, China

Prof. Nei Kato in lecture.

Distinguished Lecturer Prof. Nei Kato and the students of IEEE Communication Society, Xi'an Chapter.

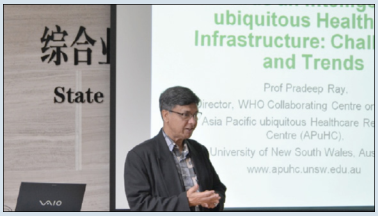

Prof. Pradeep Ray in lecture.

Distinguished Lecturer Prof. Pradeep Ray and the Members of IEEE Communication Society, Xi'an Chapter.

Prof. Zhu Han in lecture.

Distinguished Lecturer Prof. Zhu HAN and the members of IEEE Communication Society, Xi'an Chapter.

During the past half year, the Xi'an ComSoc Chapter arranged several seminars for students and young members. The discussion topics were mainly about wireless communications, electronic healthcare, communications security, and so on.

In addition to the seminars, we also invited three Distinguished Lecturers to conduct lectures for the students and young scholars.

On May 15, 2015, IEEE Distinguished Lecturer Prof. Nei Kato from the Tohuko University, Japan gave a lecture on 'Device-to-Device (D2D): Research Trends and Future Perspectives' to the students. More than 30 students and teachers attended the lecture.

On June 24, IEEE Distinguished Lecturer Prof. Pradeep Ray from the University of New South Wales, Sydney, Australia gave a lecture on 'Cooperative Service Management in Healthcare Sector: Emerging Trends and Future Challenges' to students and scholars from Xidian University and also to our chapter members.

Prof. Pradeep Ray also visited the study group of Prof. Yang Gang and had a discussion about the issues related to E-Health. The researchers hoped that they can have full collaboration with Prof. Pradeep Ray in the study of E-Health.

On July 6, 2015, IEEE Distinguished Lecturer Prof. Zhu Han from the University of Houston, USA gave a lecture on 'Cooperative Game Theory for Cognitive Radio'. Many students attended the lecture and had good discussions with Prof. Zhu Han.

The students developed many advanced study topics and study methods from those seminars and lectures.

SISTER SOCIETY NEWS

Geo2Tag: The Advanced Open Source Ecosystem for LBS Developers

By Mark Zaslavskiy, Kirill Krinkin and Sergey Balandin, Geo2Tag LBS Platform Development Team, FRUCT Ltd.

This review presents the new opportunity for research and developer teams working on LBS-enabled solutions and Internet of Things infrastructure. Geo2Tag is the world's most popular open source LBS platform[1]. Recently the IEEE IoT technical team listed Geo2Tag as a platform for making IoT solutions in city tagging scenarios[2].

The Geo2Tag platform enables simultaneous work on multiple different services and seamless creation of new services without server reloads. The platform provides a sophisticated spatiotemporal API for the services, which allows geo-fencing with altitude support, date/time filtering (including support of BC dates), data aggregation and visualization, and so on. At the same time service creation on Geo2Tag is an easy and cozy process, and developers have access to the large library of templates and functions implementing typical functionality. This, plus the use of an open source license, makes Geo2Tag the most cost-efficient and technically best available backend for commercial and non-commercial geo-apps.

Geo2Tag supports Platform as a Service Architecture (PaaS). It uses virtualization instruments such as Vagrant and Docker that enable easy integration with a majority of cloud services. The main services provided by the platform are:

•Storage, processing and visualization of spatiotemporal data.

•Simultaneous work on several location-based services inside one instance of the platform.

•Fast content-building startup by importing and pre-processing data from the open databases.

The Geo2Tag data model is as follows. Point is an atomic element of spatiotemporal data. Channel is formed as a set of points. The service is a set of channels. The user is an application or human actor that performs interaction with the platform.

The external interface for LBS developers is provided by REST requests over HTTP. API provides control over spatiotemporal data, geo-fencing functional and administrative actions for a service and the whole platform instance. The spatiotemporal data filtration has several criteria: date and altitude intervals, spatial region, and allowed channel set. All criteria can be applied independently. All REST requests for geo-data support the GeoJSON standard. API authorization is implemented by the use of the OAuth2 protocol.

Geo2Tag provides an API for high performance geo-data visualization and has a special web-service for displaying results of REST queries on a custom map based on a leaflet.js library with changeable map provider. With this solution geo-data elements (points) can be grouped into clusters, which makes data on the map easy to navigate and decreases the volume of used RAM.

Geo2Tag provides an open API for developing third-party plugins. A plug-in is an isolated extension for a platform REST API, which can perform background computations on services data and can be turned on and off without server restart. The plug-in

(Continued on Newsletter page 4)

GEO2TAG/*Continued from page 3*

system also contains common interfaces for developing open data plug-ins. These extensions allow the importing and processing of external open access web datasets.

Geo2Tag is used as a platform for various apps and use cases[3]. As an example, the Open Karelia museum system (www.openkarelia.org) networks a number of museums in Russian and Finnish and provides a common set of e-Tourism services. Other examples of Geo2Tag based services are: nearest doctor search, tracing public transport, car fleet services, and personal tracker.

Future development of Geo2Tag is focused on improving the usability of the REST API, getting more LBS developers on board. We continuously work to further improve mobile power consumption and performance and increase usability. More content is made available to developers via the API and plug-ins. We welcome you to start using Geo2Tag by visiting www.geo2tag.com, or contacting us via email (Mark.Zaslavskiy@fruct.org, Kirill.Krinkin@fruct.org, Sergey.Balandin@fruct.org).

[1]MetricsKey rating http://metricskey.com/open/open-source-location-based-services/ + is the 2nd in Google for "LBS platform" (after TomTom)

[2]The reference is at the bottom of the first page in PDF http://iot.ieee.org/iot-scenarios.html?prp=6.

[3]E. Balandina, S. Balandin, Y. Koucheryavy, and D. Mouromtsev, "IoT Use Cases in Healthcare and Tourism," *Proc. 17th IEEE Conf. Business Informatics (CBI 2015)*, Lisbon, Portugal, July 13–16, 2015. pp. 37–44.

OECC 2015/*Continued from page 1*

OECC 2015 was concluded with great success. A total of 470 submissions were obtained for the TPC to review. Among those submissions, the three largest submission totals by country/region were from mainland China, Japan, and the United States. Finally, 403 participants, including 206 non-Chinese nationality participants (~51%), participated in the conference this year. Three plenary talks, 124 invited presentations, 152 oral presentations, and 53 poster presentations were arranged. Four PDP presentations were selected; however, this year there were six no-show presentations, including two poster presentations.

OECC/PS 2016, to be held in Toki Messe, Japan, was promoted to all participants on July 1, 2015, a fantastic banquet night in Shanghai!

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 MEMBER RELATIONS
PEDRO AGUILERA, DIRECTOR OF LA REGION
MERRILY HARTMANN, DIRECTOR OF NA REGION
HANNA BOGUCKA, DIRECTOR OF EAME REGION
WANJIUN LIAO, DIRECTOR OF AP REGION
CURTIS SILLER, DIRECTOR OF SISTER AND RELATED SOCIETIES

REGIONAL CORRESPONDENTS WHO CONTRIBUTED TO THIS ISSUE

THOMAS MICHAEL BOHNERT (THOMAS.BOHNERT@ZHAW.CH)
EWELL TAN, SINGAPORE (EWELL.TAN@IEEE.ORG)

A publication of the IEEE Communications Society

www.comsoc.org/gcn
ISSN 2374-1082

FOKUS FUSECO FORUM/*Continued from page 2*

ta sources in an interoperable manner; and the combination of computer vision with cyber-physical systems for industrial purposes.

All demonstrations were based on the newest testbed toolkits from Fraunhofer FOKUS, namely Open5GCore (www.open5GCore.org), OpenSDNCore (www.openSDNcore.org), Open5GMTC (www.open5GMTC.org), OpenMTC (www.openMTC.org), 5G Playground (www.5G-playground.org), as well as the recently launched OpenBaton (www.openBaton.org), an open source ETSI NFV MANO Orchestrator.

CONCLUSION

Key takeaways of the event were:

•5G is not limited to the next generation of mobile communications, but is accompanied by a paradigm shift regarding softwarization and mobile-fixed convergence.

•Softwarization is the major innovation and the change is taking place today.

•SDN/NFV and network slicing are the enabler technologies for the transformation of the network infrastructure.

•Programmable platforms allow for the rapid introduction of new services and business agility.

•Outstanding business drivers will be the vertical markets, mainly the industrial Internet.

OUTLOOK: FOKUS FUSECO FORUM 2016

Based on the global relevance of the addressed topics and technologies, and following the continued success of the FOKUS FUSECO Forum series in the past, the 7th FUSECO Forum is planned for mid-November 2016 in Berlin, Germany. For more detailed information, please refer to: www.fuseco-forum.org

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

JANUARY

COMSNETS 2016 — 8th Int'l. Conference on Communication Systems & Networks, 5–9 Jan.
Bangalore, India
http://www.comsnets.org/index.html

IEEE CCNC 2016 — IEEE Consumer Communications and Networking Conference, 8–11 Jan.
Las Vegas, NV
http://ccnc2016.ieee-ccnc.org/

WONS 2016 — 12th Annual Conference on Wireless On-Demand Network Systems and Services, 20–22 Jan.
Cortina d'Ampezzo, Italy
http://2016.wons-conference.org/

ICACT 2016 — 18th Int'l. Conference on Advanced Communication Technology, 31 Jan.–2 Feb.
Phoenix Park, Pyeongchang, Korea
http://www.icact.org/

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 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/

IEEE ISPLC 2016 — 2016 IEEEInt'l.Symposium onPower Line Communications and Its Applications, 29 Mar.–1 Apr.
Bottrop, Germany.
http://www.ieee-isplc.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/

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

MAY

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

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

JULY

IEEE ICME 2016 — IEEE Int'l. Conference on Multimedia and Expo, 11–15 July
Seattle, WA
http://www.icme2016.org/

TEMU 2016 — Int'l. Conference on Telecommunications and Multimedia, 25–27 July
Heraklion, Greece
http://www.temu.gr/

AUGUST

EUSIPCO 2016, 29 Aug.–2 Sept.
Budapest, Hungary
http://www.eusipco2016.org/

SEPTEMBER

IEEE PIMRC 2016 — IEEE Int'l. Symposium on Personal, Mobile, and Indoor Radio Communications, 4–7 Sept.
Valencia, Spain
http://www.ieee-pimrc.org/

–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.

SOFTWARE DEFINED WIRELESS NETWORKS: PART 2

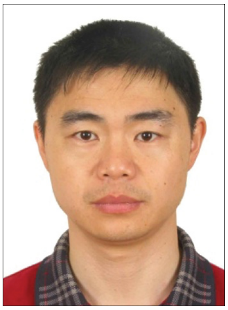
Honglin Hu


Hsiao-Hwa Chen


Peter Mueller


Rose Qingyang Hu


Yun Rui

The concept of software defined wireless networks (SDWN) has been rapidly evolving as the solution to meet the demand for dynamic wireless services in ubiquitous computing. In SDWN, the control plane and data plane are decoupled, network intelligence is logically centralized, and the underlying network infrastructure is abstracted from the applications. Thus it highly simplifies network administration and management, which facilitates the dynamic nature of future network functions and intelligent applications while lowering operating costs through simplified hardware, software, and management. This is the second part of the "Software Defined Wireless Networks (SDWN)" Feature Topic. In Part 1, which was published in November 2015, we provided an overview of the latest major developments and progress in SDWN architecture. In this Part II, we present articles that provide fruitful insights into selected topics such as Cloud-RAN management, flexible spectrum management, SDWN security, and SDN-based self organizing strategies. We believe that it complements Part I, and together they provide a holistic overview of features and trends in future SDWN development.

Future radio access networks (RANs) of cellular systems need to provide elastic service to dynamic traffic demands. The first article, "Software-Defined Hyper-Cellular Architecture for Green and Elastic Wireless Access" by Zhou *et al.*, reviews the state of the art which aims to renovate RANs from the perspective of control traffic, decoupled air interface, and software definability. The proposed and evaluated software-defined hyper-cellular architecture (SDHCA) is addressing enabling technologies such as the separation of the air interface, green base station operation, and base station function virtualization. Further, the authors summarize several future research directions on software-defined systems.

The second article, "Securing Software Defined Wireless Networks" by He *et al.*, delves into the important issue of security in SDWN. On one hand, SDWN enables new security mechanisms. On the other hand, new threats are introduced due to the separation of the control plane and data plane and its centralized logic. In lieu of these, the authors discuss threat vectors for SDWN, as well as design issues in making it secure. Security requirements of SDWN are analyzed, with security attacks and countermeasures summarized. Future research directions are often manifested.

The Cloud-RAN architecture enables dynamic configuration of base stations via virtualization technology. The third article, "Elastic Resource Utilization Framework for High Capacity and Energy Efficiency in Cloud-RAN" by Pompili *et al.*, proposes to utilize this introduced freedom to alleviate the issue of decreased energy efficiency of small cell networks by dynamically adapting to fluctuations in per-user capacity demand. The authors advocate for the need for co-location models for provisioning and allocation of virtual base stations (VBS), and propose different VBS architectures. The advantages of VBS clustering are demonstrated, which can enhance efficiency and capacity via collaborative communication techniques.

With the advent of SDWN technology, dynamic spectrum management becomes a feasible option. In the fourth article, "A Software-Defined Wireless Network Enabled Spectrum Management Architecture" by Wang *et al.*, spectrum management architecture design, which can reap the benefits of SDWN, is systematically investigated. Design principles and key challenges in realizing SDWN-enabled spectrum management architecture are discussed, and a general architecture with a new baseband virtualization design is developed. A prototype based on the IEEE 802.11 protocol is built and used to demonstrate the efficiency of the proposed architecture.

The fifth article, "Synergistic Spectrum Sharing in 5G HetNets: A Harmonized SDN-Enabled Approach" by Akhtar *et al.*, proposes a hierarchical architecture enabled by SDN that facilitates reliable and dynamic spectrum sharing in 5G cellular networks. The two key components of the proposed HSA framework are the macrocell BSs and the SDN controller. The task-sharing between the BSs and the controller harmonizes network operation and alleviates the SDN controller's scalability concerns. The article also presents an efficient resource management algorithm conceived for future 5G networks.

In the last article, "SDN Meets SDR in Self-Organizing Networks: Fitting the Pieces of Network Management" by Ramirez-Perez *et al.*, the authors combine the two popular paradigms of software defined networking (SDN) and software defined radio (SDR) under a unified management framework based on self organizing networks (SON). The proposed framework leverages programmable control planes and data planes, calls for the convergence of computing, communications, and networking research into one domain, and emphasizes the need for an open and extensible protocol interface that combines the main features of current protocols such as OpenFlow, MIH, and XMPP.

We are confident that these six articles will add value to current research activities and provide an overall direction for those researchers interested in this topic.

The Guest Editors would like to thank Sean Moore, the previous Editor-in-Chief, and Osman Gebizlioglu, the current Editor-in-Chief, for their guidance, feedback, and encouragement along the way. We also would like to thank the large number of people who significantly contributed to this Feature Topic, including the authors, reviewers, and *IEEE Communications Magazine* staff.

Biographies

Honglin Hu [SM] (honglin.hu@wico.sh) received his Ph.D. in 2004 from the University of Science and Technology of China (USTC), China. He then joined Siemens AG Communications in Munich. Now he is the vice director of the Shanghai Research Center for Wireless Communications (WiCO), and also is a professor at the Shanghai Advanced Research Institute (SARI), Chinese Academy of Sciences (CAS). He is a Finland Distinguished Professor (FiDiPro).

Hsiao-Hwa Chen [F] is currently a distinguished professor in the Department of Engineering Science, National Cheng Kung University, Taiwan. He obtained his B.Sc. and M.Sc. degrees from Zhejiang University, China, and a Ph.D. degree from the University of Oulu, Finland, in 1982, 1985, and 1991, respectively. He served as the Editor-in-Chief of *IEEE Wireless Communications* from 2012 to 2015. He is a Fellow of IET and an elected Member at Large of the IEEE Communications Society.

Peter Mueller joined IBM Research as a research staff member in 1988. His research expertise covers broad areas of human-machine interfaces, systems architecture, microwave technology, device physics, nano science, and modeling. His current field of research is in the areas of data center storage, security, and quantum technologies. As an active member of IEEE, he has authored and co-authored more than 100 papers and two books, has been granted 10 patents, and has served as a guest editor for many publications.

Rose Qingyang Hu is an associate professor in the Electrical and Computer Engineering Department at Utah State University. She has more than 10 years of R&D experience with Nortel, Blackberry, and Intel as a senior researcher. Her current research interests include next-generation wireless communications, wireless system design and optimization, multimedia QoS/QoE, communication and information security. She is an IEEE Communications Society Distinguished Lecturer (2015–2016), and she received best paper awards at IEEE Globecom 2012 and ICC 2015.

Yun Rui [SM] received the B.S. degree from Southeast University (SEU) in 2004 and the Ph.D. degree from the Chinese Academy of Sciences (CAS) in 2009, all in telecommunications engineering. From February 2011 to August 2011 he was a visiting fellow in the Department of Electronic Engineering, City University of Hong Kong. Since September 2011 he has been an associate professor at the Chinese Academy of Sciences.

Software-Defined Hyper-Cellular Architecture for Green and Elastic Wireless Access

Sheng Zhou, Tao Zhao, Zhisheng Niu, and Shidong Zhou

To meet the surging demand of increasing mobile Internet traffic from diverse applications while maintaining moderate energy cost, the radio access network of cellular systems needs to take a green path into the future, and the key lies in providing elastic service to dynamic traffic demands. To achieve this, it is time to rethink RAN architectures and expect breakthroughs.

This work is sponsored in part by the National Basic Research Program of China (973 Program: No. 2012CB316000), the National Science Foundation of China (NSFC) under grants No. 61461136004, No. 61201191, No. 61321061, and No. 61401250, and the Intel Collaborative Research Institute for Mobile Networking and Computing.

Abstract

To meet the surging demand of increasing mobile Internet traffic from diverse applications while maintaining moderate energy cost, the radio access network of cellular systems needs to take a green path into the future, and the key lies in providing elastic service to dynamic traffic demands. To achieve this, it is time to rethink RAN architectures and expect breakthroughs. In this article, we review the state-of-the-art literature, which aims to renovate RANs from the perspectives of control-traffic decoupled air interface, cloud-based RANs, and software-defined RANs. We then propose a software-defined hyper-cellular architecture (SDHCA) that identifies a feasible way to integrate the above three trends to enable green and elastic wireless access. We further present key enabling technologies to realize SDHCA, including separation of the air interface, green base station operations, and base station functions virtualization, followed by our hardware testbed for SDHCA. In addition, we summarize several future research issues worth investigating.

Introduction

Since their birth, cellular systems have evolved from the first generation analog systems with very low data rate to today's fourth generation (4G) systems with more than 100 Mb/s capacity to end users. However, the radio access network (RAN) architecture has not experienced many changes: base stations (BSs) are generally deployed and operated in a distributed fashion, and their hardware and software are tightly coupled. Facing the exponential growth of mobile Internet traffic on one hand, and the significant energy consumption of mobile networks on the other hand, breakthroughs are strongly expected in RAN architecture design and corresponding systematic control methods. In the conventional RAN architecture, even though there is little traffic requirement, the BSs cannot be switched off in order to maintain the basic coverage. This requires substantial static power to keep the BS active and to transmit the required signaling, thus causing energy waste. To meet the urgent need for green and elastic wireless access, it is envisioned that next-generation RANs should become increasingly software-defined, and the layout of their physical resources should break away from the fully distributed model.

Along with the above paradigm shift, emerging RAN architectures have been developed from three perspectives. The first is the new air interface architecture of cellular networks that features signaling and data separation, aimed at flexible and efficient control of small cells for throughput boosting and BS sleeping-based energy saving [1, 2]. To make cell coverage more adaptive to traffic dynamics, some control signaling functions should be decoupled from the data functions so that the data traffic service is provided on demand, while the control plane is always "on" to guarantee basic coverage. The second is renovating cellular networks into massive BSs with centralized baseband processing and remote radio heads (RRHs), which further evolves to have a cloud-based baseband processing pool [3] and BS functions virtualization. The third is inspired by software-defined networking (SDN) from wired networks, which separates the control and data planes to enable centralized optimization of data forwarding.

We believe that successfully delivering green and elastic mobile access relies on the deep convergence of the above three perspectives, of which the rationale is as follows. First, to realize the control-traffic decoupled air interface, flexible and efficient signal processing is required to reconstruct the frame components from the control and traffic layers of the air interface. The load of control signaling can also vary over time and space for current mobile networks, which requires the control coverage to be reconfigurable with adaptive power and spectrum resources in order to match the signaling load variations. To tackle these challenges, cloud-based RAN architectures can offer help by providing programmable BS functions and reconfigurability of radio elements. Furthermore, the fronthaul network in cloud-based RAN architectures can be aided by SDN to enable efficient data forwarding and flexible function splitting. Meanwhile, the control-data separation in SDN can be extended to wireless access by the control-traffic decoupled air interface. Naturally, the air interface separation should converge with cloud-based baseband processing under software-defined provisioning, exploiting its high flexibility and reconfigurability. In fact, recent studies have begun to investigate the integration of these perspectives [5, 10]. However, the problems of how to combine the

Sheng Zhou, Zhisheng Niu and Shidong Zhou are with Tsinghua University; Tao Zhao was with Tsinghua University and is now with Texas A&M University.

Trend	Literature	Features	Benefits
Decoupled air interface	• HCN [1] • Phantom Cell [2]	• Separation of control and data coverage • CBSs gather network control information	• Energy saving • Global optimization of network resources
Cloud-based RAN	• WNC [7] • C-RAN [3]	• BBU RRH separation • BBU consolidation • Virtual base stations	• Cost reduction • Improved flexibility
Software-defined RAN	• SoftRAN [8]	• Control data separation • Logically centralized controller • Control APIs	• Global utility optimization • Simplified network management
Integrated architectures	• OpenRAN [4] • CONCERT [5] • SDF [9] • Zaidi *et al.* [10]	• Integration of SDRAN and cloud computing • Integration of SDRAN and decoupled air interface	• Combined benefits • Realization acceleration

Table 1. Summary of new RAN architectures.

three perspectives into a converged architecture and how to realize the architecture in practice remain unclear.

To this end, in this article we propose a new software-defined hyper-cellular architecture (SDHCA) that realizes the separation of the air interface via a software-defined approach in a cloud-based infrastructure. We begin with reviewing the recent research on RAN architecture innovations from the aforementioned three perspectives. Then we present the overall design of SDHCA, which integrates air interface separation, cloud RAN, and SDN, emphasizing the major technical contributions of SDHCA that bring elastic and green mobile service. In the next section, our initial research efforts toward the key enabling technologies of SDHCA are presented, followed by the testbed implementation showing the feasibility of SDHCA. Finally, we outline future research directions that can ultimately facilitate SDHCA in practical RANs.

Recent RAN Architecture Developments

People have witnessed a rising interest in novel RAN architectures in the recent literature. We categorize them into three independent trends, and summarize their main features and benefits in Table 1. The first trend is signaling-data separation at the air interface. Among them the hyper-cellular architecture (HCA) [1] and the Phantom Cell concept [2] are typical examples. Under such architectures, the network coverage is divided into two layers: control coverage and traffic coverage. For instance, in HCA, BSs are classified into two types: control base stations (CBSs) and traffic base stations (TBSs). Specifically, CBSs take care of control coverage, which provides network access, system information broadcast, and so on. On the other hand, TBSs are meant for data traffic services to mobile users. With the decoupled air interface, TBSs can be switched on/off for significant energy savings without generating coverage holes. Besides, CBSs can gather network control information and globally optimize the on/off states of TBSs and the radio resource allocation. The idea of decoupled air interface has made its way into the cellular standard known as "dual connectivity" in Third Generation Partnership Project (3GPP) Long Term Evolution (LTE) Release 12 [6].

Industry and academia are also investigating integrating cloud computing technologies into RANs. Among the most representative are wireless network cloud [7] and Cloud RAN (C-RAN) [3]. They share the same idea of consolidating baseband units (BBUs) of BSs to a centralized computing cloud, while only leaving remote radio heads (RRHs) in the front-end. Cloud-based architecture can reduce energy consumption, as well as the RAN deployment and operational costs [3]. Besides, with virtualization, virtual base stations (VBSs) can be realized, opening up the RAN for flexible system configurations and operations.

Third, SDN has brought a rethinking of packet switching and routing in the Internet, and there have been emerging studies to bring SDN concepts such as control-data separation, centralized control, and software application programming interfaces (APIs) to RANs. SoftRAN [8] aims to enable software-defined RANs (SDRANs). It introduces the concept of the big BS, including a logically centralized controller and distributed radio elements. With the centralized controller, SDRAN optimizes the global utility over a local geographic area, and simplifies network management by software programming.

The recent literature has started to explore the integration of the above trends for future RANs. OpenRAN [5] is proposed to utilize a cloud computing resource pool and virtualization to implement the SDRAN architecture. CONCERT [5] builds a converged cloud and cellular system based on control-data decoupling. Arslan *et al.* [9] propose the concept of software-defined fronthaul (SDF) based on SDN and C-RAN. Zaidi *et al.* propose an integrated architecture that combines SDN concepts from SoftRAN and signaling-data separation [10]. These research works motivate us to explore the integration of the above three perspectives in order to take a further step in enabling green and elastic wireless access in future cellular systems.

Software-Defined Hyper-Cellular Architecture

As shown in Fig. 1, the SDHCA design is based on the deep integration of air interface separation, cloud RAN, and SDN. It exploits the cloud infrastructure, which can be divided into three subsystems: the RRH network, the front-haul network, and the virtual BS (VBS) cloud. The radio elements (RRHs) can merely deal with

The advantage of dual connectivity is that it requires minimum changes to the overall RAN architecture in LTE, and one can expect standard terminals to gradually support it. However, minimum changes also limit the degree of freedom in separation, which implies the separation scheme can be suboptimal.

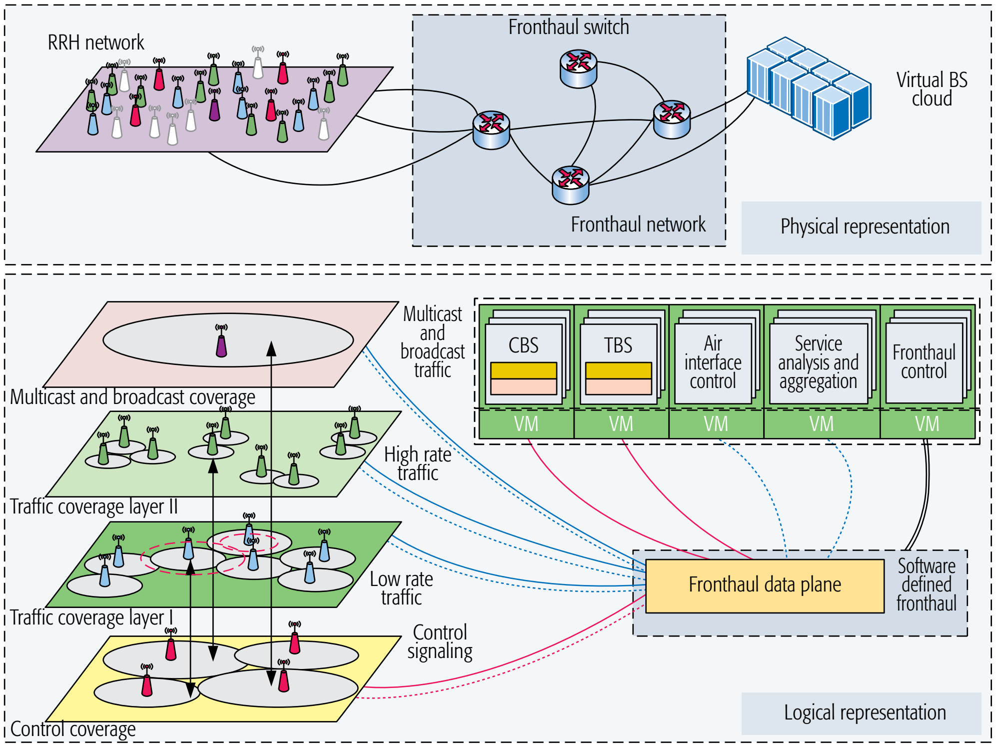

Figure 1. Software-defined hyper-cellular architecture.

RF transmission/reception, or have some baseband processing functions, and their roles can be dynamically configured as CBS or TBS, or put into sleep mode, according to the network status and hardware capabilities of the RRHs. The deployment of the RRHs can be done via conventional network planning mechanisms, satisfying the peak hour traffic of the network. From the perspective of logical functions, the proposed SDHCA provides one control coverage layer and multiple conceptual layers for different user traffic types. In this way, the cell coverage is "softer" and smarter to deliver greener wireless access. RRHs are connected to the VBS cloud via the fronthaul network, which is also software defined. In the cloud, the functions of CBS and TBS are realized as VBS applications in virtual machines (VMs).

The key features of SDHCA are summarized as follows.

Control Data Separation. The separation lies in three aspects. First, in the air interface, CBSs are in charge of control coverage, while TBSs are responsible for traffic coverage. Second, on the infrastructure level, the software in charge of the network functions is separated from the hardware that forwards or transmits the data. In particular, one RRH can be dynamically configured to act as a CBS or a TBS or even both (i.e., handling some traffic while acting as a CBS). Last but not least, the control and data planes of the SDF network are also decoupled.

CBS as the RAN Controller. CBSs take care of mobile users and TBSs underlying their coverage. Besides, CBSs also control the fronthaul network. In this way, the CBS has a global view of the RAN in a local geographic area, and optimizes the on-demand configuration and activation of TBSs so that the network resources, including spectrum resources and energy resources, can well match the dynamic traffic in an elastic way. When the traffic load changes, a CBS can also control the cell zooming behavior of active TBSs to balance the load, as shown in the red dashed circles on traffic layer 1 in Fig. 1.

Software-Defined Network Functions via Virtualization. The network functions, including air interface control, service analysis and aggregation, baseband sample generation, the fronthaul control plane, and so on, are realized by software applications running in VMs. The functions are thus easily programmed and updated, allowing for flexible and efficient network operations, potentially reducing the computing energy consumption.

Thanks to the above tightly integrated features, SDHCA provides flexible services to users exploiting spatial-temporal variations of the traffic demand, so the energy efficiency of the whole network can be greatly improved:

- The decoupled air interface produces flexible sleeping opportunities for TBSs. As shown in Fig. 1, the access control and other coverage-related signaling of different traffic layers are handled by CBSs, and the vertical arrows in Fig. 1 indicate that cells in the traffic layers are covered by a CBS in the control layer based on their positions. Therefore, the RRHs can be configured to TBSs on any traffic layer shown in Fig. 1, and they can be switched off without having coverage holes, leading to significant energy saving. Regarding active TBSs, the energy consumption related to control signaling can also be saved.
- The required compute resources in the BBU pool can adapt to the number of active RRHs,

and so do the virtualized fronthauling functions needed to support these active RRHs. For instance, the VMs that run the CBS and TBS realizations can be constructed or released on demand based on the active/sleeping status of the associated CBSs and TBSs, saving the energy of the computing cloud.

• Unlike conventional cost-inefficient solutions that rely on dark fibers to connect RRHs to the VBS cloud, the SDF enables flexible mapping between BBUs and RRHs [9], and efficient baseband function splitting [11] for equivalent baseband signal compression, so the high bandwidth requirement of the fronthaul is guaranteed at low cost.

In short, the proposed SDHCA is a viable solution to offer green and elastic mobile access.

Enabling Technologies

The enabling technologies of SDHCA are described in detail in this section. We first compare existing air interface separation solutions and discuss future separation-oriented air interface design. Then our initial research efforts to realize green BS operations with BS dispatching and BS sleeping are presented. Finally, possible solutions of BS functions virtualization in SDHCA are discussed.

Separation of the Air Interface

For possible air interface separation schemes, we observe that the naive "extreme separation" scheme is in fact unsuitable. Here, extreme separation means that only the transmission of user perceived data is handled by the TBS, while all other parts are processed at the CBS. The reason extreme separation fails is that modern cellular systems rely on pilot symbols to estimate the wireless channel to aid the decoding of user data bits. With pilots transmitted by the CBS and data bits transmitted from the TBS, this extreme separation approach will lead to inaccurate channel estimation and thus incorrect decoding of the data.

One possible method suggested in the literature is functionality separation [12], where functionality is defined as the essential sets of functions provided by the network to mobile users. It can be divided into five classes: synchronization, broadcast of system information, paging, multicast (for low-rate data transmissions, e.g., voice), and unicast (for high-rate data transmissions). Through functionality separation, the CBS is responsible for the former four classes, and the TBS for synchronization and high-rate data transmission. Careful analysis was conducted to ensure that the user equipment (UE) state transitions continue to work after the separation, although specific to the LTE standard. Based on functionality separation, an alternative separation scheme was proposed in HyCell [14], which is discussed in detail in the next section.

Besides, the separation of the air interface has also undergone discussion in 3GPP, called dual connectivity [6]. In dual connectivity, CBS and TBS are called PeNB and SeNB. respectively, and the separation scheme is specified in the aspects of the control plane (C-Plane) and user plane (U-Plane). In the C-Plane, only a CBS sends out radio resource control (RRC) messages to the user after coordination with the TBS. Two options are available for the U-Plane:

1. Independent Packet Data Convergence Protocol (PDCP) layers are used at the CBS and the TBS, and each node carries one bearer.
2. A slave bearer is split between the PDCP and radio link control (RLC) layers so that CBS processes the PDCP layer, and TBS processes the RLC and lower layers.

The advantage of dual connectivity is that it requires minimum changes to the overall RAN architecture in LTE, and one can expect standard terminals to gradually support it. However, minimum changes also limit the degree of freedom in separation, which implies the separation scheme can be suboptimal.

For future cellular systems, a new separation-oriented air interface design is preferable if one is allowed to break backward compatibility with earlier cellular standards, guided by the principle of separating control from traffic. It is expected to reduce redundancy and keep the protocol simple. It can also improve network efficiency through independent optimization of the air interface at the CBS and TBS according to their different characteristics. Besides, the air interface should be made easily programmable and upgradable via software, thus improving flexibility. The discussion of the aforementioned separation schemes is summarized in Table 2.

BS functions virtualization is part of network functions virtualization (NFV) in cellular systems. Proposed and standardized by ETSI, NFV aims to virtualize the network node functions and build virtual networks. With BS functions virtualization, BSs become VBSs, and their functions are software defined and can provide various APIs to network operators.

Green Base Station Operations

As mentioned earlier, SDHCA can improve network energy efficiency by allowing flexible TBS sleeping and reducing the signaling related energy consumption on TBSs. We conduct a simulation study to evaluate the energy saving gain. We use a system-level simulator in accordance with 3GPP simulation requirements, and the BS power model is from the EARTH project [13]. The layout of 19 CBSs follows a standard hexagonal topology, while the distribution of TBSs follows a homogeneous Poisson point process with a total of 38 TBSs on average. Users, each with an FTP downloading source of the same volume, are randomly and uniformly dropped in the 19-hexagonal area, and we vary the number of users according to a daily traffic profile provided by EARTH [13], resulting in the average total number of users in each hour shown in Fig. 2a. In Fig. 2b, three schemes are compared. "HCA w/o sleep" corresponds to the case when TBSs are not allowed to sleep, while in "HCA w/ sleep" TBSs can dynamically go to sleep according to a simple load-threshold-based policy. "HetNet" corresponds to a conventional heterogeneous network, where we turn CBSs into macro BSs and TBSs into micro BSs from HCA w/o sleep. The time-frequency resource blocks (RBs) are fully used in all BSs in HetNet, while on CBSs, only those for control signaling are used, and the RBs for control signaling are muted on TBSs. The energy reduction from HetNet to HCA w/o sleep is due to the muted control-signaling-related RBs, while the energy saving from HCA w/o sleep to HCA w/ sleep is due to the dynamic TBS sleeping. Note that the separated air interface guarantees this flexible TBS sleeping. One of the

Scheme	Advantages	Drawbacks
Extreme separation: • TBS: user data • CBS: all others (including pilot)	–	Unsuitable
Functionality separation [12]: • CBS: synchronization, broadcast of system information, paging, multicast • TBS: synchronization, unicast	Analytic feasibility validation	LTE-specific
HyCell [14] • Joint functionality and logical channel separation • Synchronization only at CBS	• Generic to multiple standards • Testbed evaluation	Evaluation only on GSM/GPRS
Dual connectivity [6]: • C-Plane: CBS sends out RRC • U-Plane: –Option 1: independent PDCP –Option 2: slave bearer PDCP at CBS, RLC and lower at TBS	• Standard • Minimal change	Suboptimal
Separation-oriented air interface: • Separation in design phase • Independent optimization • Software defined	• Simple and efficient • Programmable	Backward compatibility breakage

Table 2. Comparison of air interface separation schemes.

fundamental issues in SDHCA or any air interface separation scheme is the channel condition acquisition, which is particularly important for BS dispatching and sleeping. When a new user arrives, the network should dispatch the best RRH to act as its serving TBS, which possibly requires waking up a sleeping RRH. However, getting to know the channel conditions of the sleeping RRH to the user is very challenging.

We thus propose a novel method based on machine learning to see the unobservable channel conditions of sleeping RRHs. The main idea is to build a mapping function from the user's observable channel conditions of active CBSs and TBSs to the unobservable channel conditions of sleeping RRHs. A neural network (NN)-based algorithm for RRH selection learning in SDHCA is designed, which combines the standard approaches in NN with crafted processing procedures including discrete Fourier transform (DFT), quantization by logarithmic treatment, and Lloyd's algorithm to form input features. We consider a single CBS with 80 antennas and 5 candidate single-antenna RRHs in sleep. The objective is to select a RRH from the five candidates with the best channel gain to a user, given that the user only has the instantaneous channel condition to the CBS, and historical channel conditions to the CBS and RRHs recorded by other users. The prediction accuracy is defined as the percentage of correct selection. Another performance metric is relative selection error between the predicted RRH's channel gain and the actual best channel gain. We compare our algorithm with other algorithms, and the results are shown in Table 3. The random selection (RS) method provides the baseline accuracy, which is 20 percent for all scatterer configurations. The simple K-nearest-neighbor (KNN) algorithm, which outputs the dominant choice among the K nearest neighbors in the channel space, increases the accuracy to about 55 percent. In comparison, the accuracy of the proposed NN-based channel learning algorithm with channel response as input (NN-CR) is around 70 percent. Moreover, there is an 8 percent gap between the accuracies of NN-CR and an NN-based algorithm using genuine user location as input (NN-LO). But note that in practical systems, location information sufficiently accurate for channel estimation is generally hard to obtain.

Base Station Functions Virtualization

BS functions virtualization is part of network functions virtualization (NFV) in cellular systems. Proposed and standardized by ETSI, NFV aims to virtualize the network node functions and build virtual networks. With BS functions virtualization, BSs become VBSs, and their functions are software defined and can provide various APIs to network operators.

Currently there are several proof-of-concept implementations of VBSs in cellular networks, typically implemented on hypervisor-based virtualization platforms, such as KVM and VMWare ESX, and one major challenge is providing real-time performance [3]. Modern cellular systems have stringent real-time requirements. For example, the LTE standard specifies that one subframe must be acknowledged after three subframes upon reception in frequency-division duplex (FDD) mode. It leaves a total of 3 ms budget for decoding and subframe generation. To fulfill the task, software optimization and hardware accelerators are employed in current implementations. In the software domain, real-time optimization of the whole software stack, including the host and guest operating system (OS, typically Linux) kernel, the hypervisor, and the guest applications, is necessary to power up modern cellular standards on general-purpose processor (GPP) platforms, which makes BS functions virtualization a daunting task.

Container virtualization has the potential to provide a lightweight but effective way to virtualize the BS functions and realize SDHCA. Compared to hypervisor-based virtualization, container virtualization eliminates the need for a guest OS. By reducing the intermediate virtualization layers, container virtualization provides better performance than hypervisor-based virtualization. Therefore, when using container virtualization to build VBSs, the need for real-time optimization can be relaxed.

Another issue of realizing BS functions on virtualization is inter-VM communication. Virtualization typically provides isolation across VMs for fault tolerance, but it also makes inter-application interactions difficult, because the network communication across VMs can degrade system performance dramatically. However, coordinated multipoint (CoMP) communication and CBS-TBS signaling rely on efficient inter-VBS communications. To tackle this issue, new mechanisms are needed to share resources such as disk, memory, or CPU cache to facilitate inter-VBS communications. Instruction set architecture enhancements can also be exploited as a complementary approach to improve the performance of VM communications.

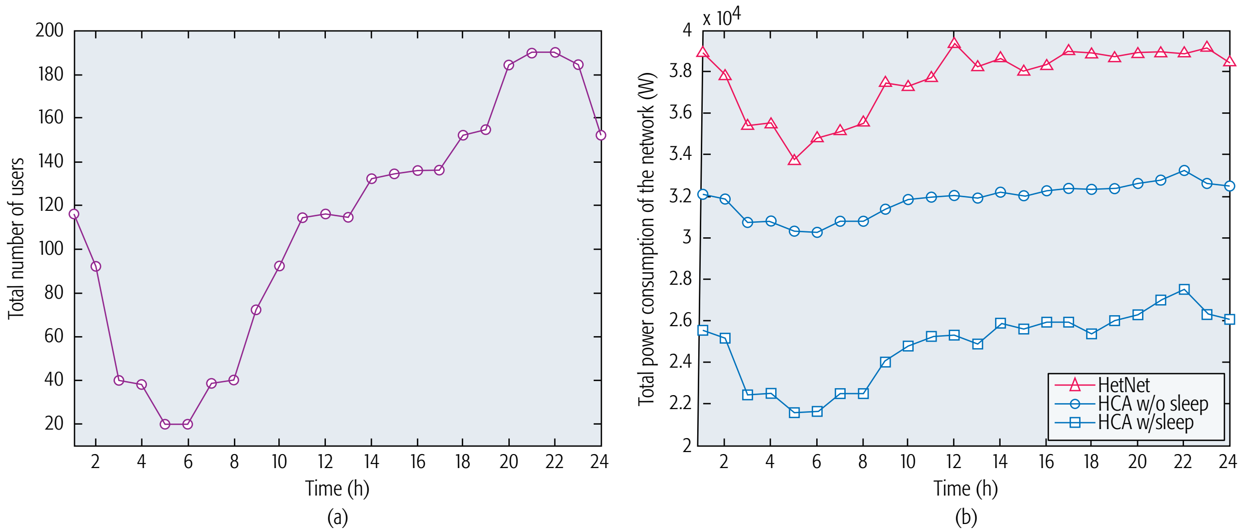

Figure 2. Simulation comparison of the network power consumption: a) average total number of users in the network for each hour; b) average total network power consumption for each hour.

No. of scatterers	10	15	20	25	30	35
RS	0.2/0.4062	0.2/0.3815	0.2/0.3665	0.2/0.3562	0.2/0.3488	0.2/0.3395
KNN ($K = 5$)	0.5601/0.1741	0.5532/0.1622	0.5420/0.1606	0.5306/0.1601	0.5188/0.1593	0.5029/0.1563
NN-CR	0.7483/0.0662	0.7249/0.0703	0.7140/0.0662	0.6871/0.0744	0.6700/0.0791	0.6528/0.0855
NN-LO	0.8268/0.0317	0.7972/0.0376	0.7845/0.0392	0.7788/0.0398	0.7584/0.0429	0.7443/0.0476

Table 3. Prediction accuracy/relative prediction error of different algorithms.

Hardware Testbed

We are prototyping a hardware testbed for systematic evaluation of SDHCA, which is realized via programmable software on GPP platforms. Figure 3 shows our current testbed called HyCell [14], which is based on OpenBTS,[1] an open source Global System for Mobile Communications (GSM)/general packet radio service (GPRS) BS application, and the USRP[2] hardware platform. The overall system structure including hardware interconnection and major software modules are shown in the left part of Fig. 3. In the testbed, the association between the BS servers and the USRP devices is static due to hardware limitations. The static assignment makes it relatively simple and easy to implement our testbed, while it cannot fully reveal SDHCA's potential adaptiveness to traffic dynamics.

Our testbed work presents an alternative separation scheme by jointly considering the functionality and logical channels of current standards. In this design, synchronization only resides at the CBS side, and the TBS is only in charge of high-rate data transmission. As a result, synchronization between the TBS and the UE is guaranteed by CBS-UE synchronization over the air. One major advantage of this synchronization scheme is that no user side modification is needed, so existing mobile terminals can seamlessly access the renovated cellular network. The scheme has been evaluated on our GSM/GPRS based testbed implementation, and is generic to multiple standards. Currently we are investigating its implementation over the LTE standard based on an open source platform OpenAirInterface.[3]

Besides, our testbed is able to demonstrate green BS operations including BS dispatching and BS sleeping. BS sleeping is utilized when the network load is light, and some TBSs can be switched off for energy saving, while load-balancing BS dispatching is beneficial for highly loaded networks or when some TBSs are in sleep mode and the remaining "active" network has high load. In the testbed, we also evaluate the delay overhead of the air interface separation and BS sleeping scheme, which are important factors to guarantee quality of service (QoS) while enjoying the energy saving benefits. In the performance measurements of Fig. 3, we have shown the delay overhead of signaling interaction between the CBS and the TBS when processing the channel request from a user terminal, of which the mean is about 0.36 ms and the standard variance is about 0.1 ms . The small values indicate the feasibility of the air interface separation in SDHCA. As for BS sleeping, we measure the close-down time of turning off a TBS, and its mean is about 52 ms. On the other hand, the setup time of waking up a TBS is several seconds; its optimization needs future study.

Regarding BS functions virtualization, we are prototyping a VBS pool on a commodity x86 server based on Docker[4] containers in

[1] http://openbts.org/

[2] http://www.ettus.com/

[3] http://www.openairinterface.org/

[4] https://www.docker.com/

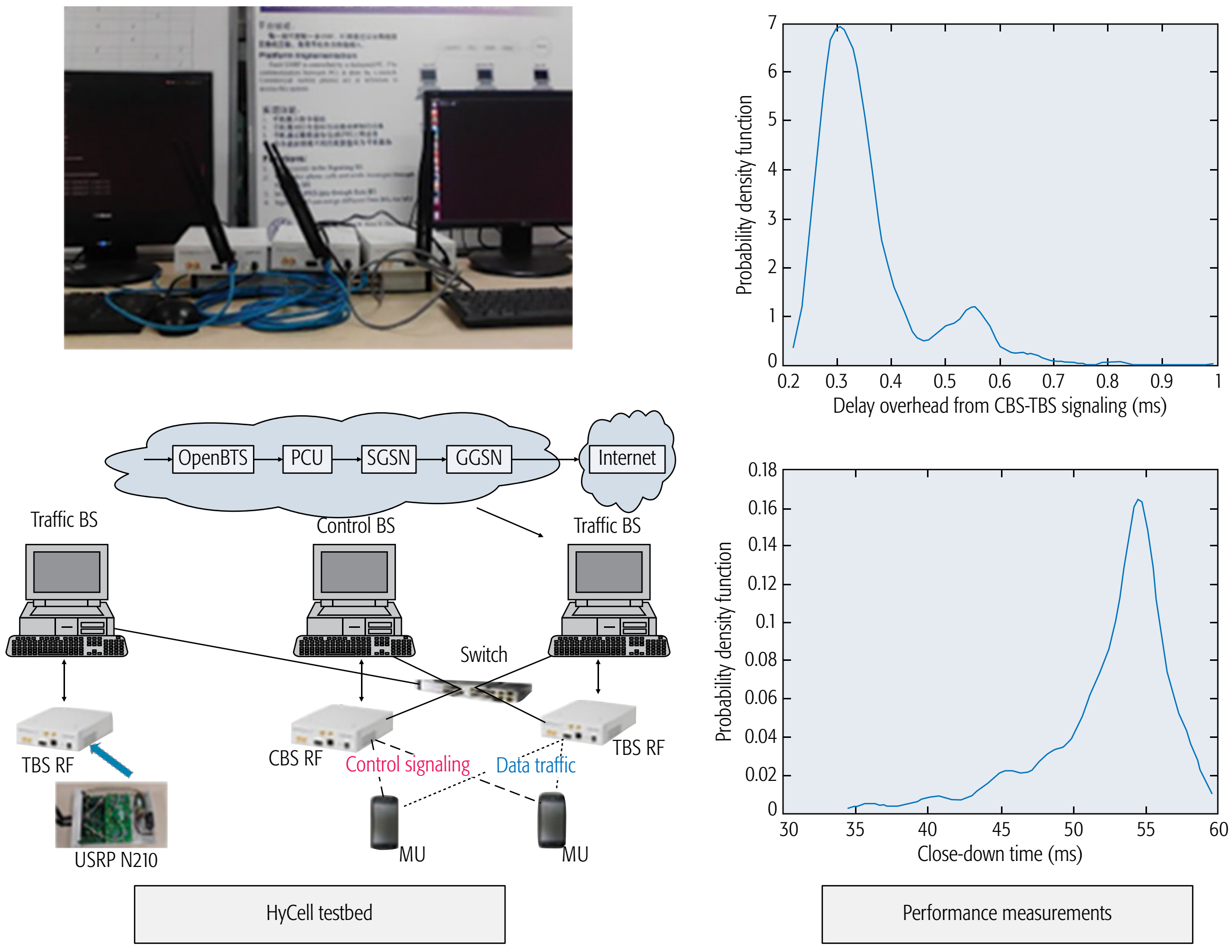

Figure 3. The HyCell testbed.

order to implement dynamic resource allocation algorithms for multiple VBSs and design the API for software-defined network control and management, including more sophisticated BS dispatching and sleeping algorithms. We envision a complete SDHCA testbed to demonstrate its concepts, compare different algorithms, and evaluate system performance as well as inspire new research directions.

The Way Forward

Synchronization is a big challenge in realizing SDHCA. Fine synchronization between different network elements is required to enable green BS operations. Current implementations typically make heavy use of GPS receivers. However, it incurs additional cost, and might not work well in indoor and underground environments. Alternative solutions such as fronthaul-based synchronization or radio-based synchronization should be investigated to tackle this challenge.

Inherited from the cloud-based RAN, SDHCA also faces the challenge of designing high-bandwidth fronthaul. Moreover, future fronthaul networks can have heterogeneous physical realizations, including wireless, fiber, high-speed Ethernet, and so on. As a promising solution for the next generation fronthaul, software-defined packet switching should be evaluated, and the question of how to realize it needs to be answered [15].

From the perspective of better satisfying mobile users' needs, SDHCA should be programmed to provide user-centric services by making the network aware of mobile users' states and allocating network resources to deliver their request contents. How to capture users' state and how to schedule users accordingly have yet to be addressed; for example, interference-aware RRH selection in SDHCA is a valuable but challenging issue.

By building virtual network functions on a software-defined platform, the same physical infrastructure can be shared by multiple network operators, including virtual operators. Dynamic RAN sharing relies on high-level abstraction of the RAN as a controllable entity. Technical problems lie ahead including RAN resource slicing and isolation, as well as RAN provisioning and orchestration. Besides, novel business models between multiple operators call for investigation, possibly via a game theoretic approach.

SDHCA enables green and elastic wireless access, and can further serve as a platform for RAN innovations. In this work we focuses on the RAN part of cellular systems. There are quite a few works trying to bring SDN and NFV concepts to the core network part. How to combine the RAN innovations with core network advances is an exciting research direction.

References

[1] Z. Niu *et al.*, "Energy Efficiency and Resource Optimized Hyper-Cellular Mobile Communication System Architecture and Its Technical Challenges," *SCIENTIA SINICA Informationis*, vol. 42, no. 10, 2012 (in Chinese), pp. 1191–203.

[2] H. Ishii, Y. Kishiyama, and H. Takahashi, "A Novel Architecture for LTE-B: C -plane/U-Plane Split and P hantom Cell Concept," *IEEE GLOBECOM Int'l. Wksp. Emerging Technologies for LTE-Advanced and Beyond-4G*, Dec. 2012, pp. 624–30.

[3] Ch.-L. I *et al.*, "Recent Progress on C-RAN Centralization and Cloudification," *IEEE Access*, vol. 2, 2014, pp. 1030–39.

[4] M. Yang *et al.*, "OpenRAN : A Software-Defined RAN Architecture via Virtualization," *Proc. ACM SIGCOMM 2013 Conf.*, Hong Kong, China, Aug. 2013, pp. 549–50.

[5] J. Liu *et al.*, "CONCERT : A Cloud-Based Architecture for Next-Generation Cellular Systems," *IEEE Wireless Commun.*, vol. 21, no. 6, Dec. 2014, pp. 14–22.

[6] 3GPP, "3GPP TR 36.842: Study on Small Cell Enhancements for E-UTRA and E-UTRAN–Higher Layer Aspects," Dec. 2013, v. 12.0.0, Rel. 12; http://www.3gpp.org/ftp/Specs/archive/36_series/36.842/36842-c00.zip

[7] Y. Lin *et al.*, "Wireless Network Cloud: Architecture and System Requirements," *IBM J. Research and Development*, vol. 54, no. 1, 2010, pp. 4:1–4:12.

[8] A. Gudipati *et al.*, "SoftRAN : Software Defined Radio Access Network," *Proc. 2nd ACM SIGCOMM Wksp. Hot Topics in Software Defined Networking*, Hong Kong, China, Aug. 2013, pp. 25–30.

[9] M. Y. Arslan, K. Sundaresan, and S. Rangarajan, "Software-Defined Networking in Cellular Radio Access Networks: Potential and Challenges," *IEEE Commun. Mag.*, vol. 53, no. 1, Jan. 2015, pp. 150–56.

[10] Z. Zaidi, V. Friderikos, and M. Imran, "Future RAN Architecture: SD-RAN through a General-Purpose Processing Platform," *IEEE Vehic. Tech. Mag.*, vol. 10, no. 1, Mar. 2015, pp. 52–60.

[11] J. Liu *et al.*, "Graph-based Framework for Flexible Baseband Function Splitting and Placement in C-RAN," *IEEE ICC 2015*, London, U.K., June 2015, pp. 958–1963.

[12] X. Xu *et al.*, "On Functionality Separation for Green Mobile Networks: Concept Study over LTE," *IEEE Commun. Mag.*, vol. 51, no. 5, May 2013, pp. 82–90.

[13] G. Auer *et al.*, "How Much Energy Is Needed to Run a Wireless Network?," *IEEE Wireless Commun.*, vol. 18, no. 5, pp. 40–49, Oct. 2013.

[14] T. Zhao *et al.*, "HyCell: Enabling GREEN Base Station Operations in Software-Defined Radio Access Networks," *IEEE ICC 2015 Wksp. Next Generation Green ICT*, London, U.K., June 2015, pp. 2868–73.

[15] J. Liu *et al.*, "Redesigning Fronthaul for Next-Generation Networks: Beyond Baseband Samples and Point-to-Point Links," *IEEE Wireless Commun.*, vol. 22, no. 5, Oct. 2015, pp. 90–97.

Biographies

Sheng Zhou [M] (sheng.zhou@tsinghua.edu.cn) received his B.S. and Ph.D. degrees in electronic engineering from Tsinghua University, China, in 2005 and 2011, respectively. He is currently an assistant professor in the Department of Electronic Engineering, Tsinghua University. From January to June 2010, he was a visiting student at Wireless System Lab, Electrical Engineering Department, Stanford University, California. His research interests include cross-layer design for multiple antenna systems, cooperative transmission in cellular systems, and green wireless communications.

Tao Zhao [S] (alick@tamu.edu) received his B.Eng. and M.Eng. degrees from Tsinghua University in 2012 and 2015, respectively. He is currently a Ph.D. student in the Department of Electrical & Computer Engineering, Texas A&M University, College Station. His research interests include software-defined wireless networks, cloud-based distributed systems, and cyber-physical systems.

Zhisheng Niu [F] (niuzhs@tsinghua.edu.cn) graduated from Beijing Jiaotong University, China, in 1985, and got his M.E. and D.E. degrees from Toyohashi University of Technology, Japan, in 1989 and 1992, respectively. During 1992–1994, he worked for Fujitsu Laboratories Ltd., Japan, and in 1994 joined Tsinghua University, where he is now a professor in the Department of Electronic Engineering and deputy dean of the School of Information Science and Technology. He is also a guest chair professor of Shandong University, China. His major research interests include queueing theory, traffic engineering, mobile Internet, radio resource management of wireless networks, and green communication and networks. He is now a Fellow of IEICE, a Distinguished Lecturer (2012–2015) and Chair of the Emerging Technology Committee (2014–2015) of the IEEE Communication Society, and a Distinguished Lecturer (2014–16) of the IEEE Vehicular Technologies Society.

Shidong Zhou [M] (zhousd@tsinghua.edu.cn) received the B.S. and M.S. degrees in wireless communications from Southeast University, Nanjing, China, in 1991 and 1994, respectively, and the Ph.D. degree in communication and information systems from Tsinghua University, Beijing, China, in 1998. From 1999 to 2001, he was in charge of several projects in the China 3G Mobile Communication R&D Project. He is currently a Professor at Tsinghua University. His research interests are in the area of wireless and mobile communications.

SDHCA enables green and elastic wireless access, and can further serve as a platform for RAN innovations. In this work we focuses on the RAN part of cellular systems. There are quite a few works trying to bring SDN and NFV concepts to the core network part. How to combine the RAN innovations with core network advances is an exciting research direction.

Securing Software Defined Wireless Networks

Daojing He, Sammy Chan, and Mohsen Guizani

Similar to its wired counterpart, SDWN is expected to introduce a wide range of benefits to the operation and management of wireless networks. Security is always important to any network. On one hand, SDWN enables new security mechanisms. On the other hand, some new threats are introduced due to the separation of the control and data planes and the consequent introduction of the logically centralized controller.

Abstract

Software defined wireless networking (SDWN) is a new paradigm of wireless networking, physically separating the data and control planes of various elements in the wireless infrastructure. Similar to its wired counterpart, SDWN is expected to introduce a wide range of benefits to the operation and management of wireless networks. Security is always important to any network. On one hand, SDWN enables new security mechanisms. On the other hand, some new threats are introduced due to the separation of the control and data planes and the introduction of the logically centralized controller. In this article, we discuss its security threat vectors as well as design issues in making it secure. Also, we analyze the security requirements of SDWN, and then summarize the security attacks and countermeasures in this area and suggest some future research directions.

Introduction

The protocol architectures of computer networks or telecommunications networks generally consist of a control plane and a data plane. The control plane manages the configuration of networking devices (i.e., switches or routers) and their forwarding functions. The data plane consists of protocols to execute the forwarding functions according to the rules configured by the control plane protocols. Traditionally, as shown in Fig. 1a, both control and data planes are implemented in each networking device. As a result, whenever device configurations or routing strategies need to be changed, the firmware of all involved networking devices have to be modified. This means high labor cost and long delay, which increase with the network size.

Software defined networking (SDN), as shown in Fig. 1b, is a new and promising networking paradigm in which the control and data planes are decoupled, network intelligence is logically centralized, and the underlying network infrastructure is abstracted from the applications [1]. It provides great advantages in simplifying network management such that network administrators have central programmable control of network traffic via controllers, and new functions can easily be supported without physical access to the networking devices. That is, SDN is a technology that enables efficient provisioning of future network services, and lowers the operating costs through simplified hardware, software, and management.

At the same time, mobile networks become more convergent as various wireless technologies, such as Long Term Evolution (LTE), WiMAX, and WiFi, are integrated into the network infrastructure. Such an infrastructure typically comprises networking devices from different vendors and involves multiple operators. Managing the interoperability of these devices with different configurations for different policies and security requirements imposes a challenge. Moreover, while mobile users roam between different networks managed by different operators, guaranteeing consistent security across multiple domains dynamically and efficiently adds complexity to network management. With its virtualized abstraction and programmability features, SDN can hide the complexity of wireless protocols and support granular policy control. Thus, it is natural to apply the SDN paradigm to wireless mobile networks, leading to software defined wireless networking (SDWN). It is expected that SDWN will also bring the benefits of cost-effective infrastructure upgrade, delivery of new service, and improvement of user experience to existing infrastructure. The conceptual architecture of SDWN is depicted in Fig. 2. Research work in SDWN is emerging, including SDN design for the cellular core infrastructure [2], supporting fine-grained policies in cellular networks through scalable architecture design [3], abstraction of multiple base stations into a single virtual big base station [4], and decoupling of protocol definition from the hardware and provision of the software abstraction layer to enable programmable MAC and physical layers [5].

From a security point of view, SDWN has both advantages and disadvantages. As for one advantage, it enhances network security with its capability of redirecting or filtering traffic flows based on packet contents or network states. Such functions normally require additional security modules in traditional networks (e.g., firewalls or intrusion detection systems). But they can be naturally supported in SDWN, just as in the case of SDN [6]. On the other hand, due to physical separation of the control and data planes, a disadvantage is that SDWN is vulnerable to more attack vectors than traditional network architectures. This means that the availability, authenticity, confidentiality, consistency, and integrity

Daojing He is with East China Normal University; Sammy Chan is with City University of Hong Kong; Mohsen Guizani is with the University of Idaho.

0163-6804/16/$25.00 © 2016 IEEE

of network and control traffic could be severely affected. Obviously, these security issues have to be addressed before SDWNs are adopted in production networks.

This article focuses on security issues of SDWNs. It is structured as follows. The next section describes the security threats to SDWN, which are inherited from SDN. We present some specific design issues of securing SDWN. We discuss the security requirements of SDWN. Then some possible attacks to SDWNs and their countermeasures are discussed. Subsequently, some future research directions are provided.

Security Threats and Challenges of SDWNs

In traditional networks, forwarding devices are distributed in different geographical locations. If an attack to multiple forwarding devices is made, it needs to be carried out in a cooperative manner; thus, launching such an attack is not straightforward. On the other hand, even though SDWNs bring the benefits of network programmability and logically centralized control, it is exactly these benefits that expose SDWNs to new threats or those treats that are harder to exploit in traditional networks. For example, just a single attack on the controller unit can compromise the entire network.

As shown in Fig. 3, a vector of threats to SDWNs have been identified in [1]. They are tabulated in Table 1 and described in this section.

1. Forged or faked traffic flows: Both forwarding devices and controllers are vulnerable to this attack. Either a non-malicious faulty device or an adversary could trigger this threat. An attacker can launch denial-of-service (DoS) attacks to exhaust the resources in forwarding devices and controllers. Certainly, this problem can be mitigated by an authentication mechanism. However, if the attacker has compromised an application server that holds the credentials of many users, it can easily inject forged flows, which are authorized, into the network.

2. Attacks on forwarding devices: Such attacks can easily devastate the network. One single forwarding device could be used to discard, slow down, or deviate network traffic. Even worse, forged requests could be injected to overload the controller.

3. Attacks on control plane communications: Such attacks can be used to generate DoS attacks or divert flows of network traffic for the purpose of data theft. Various weaknesses of the transport layer security/secure sockets layer (TLS/SSL) communications and the public key infrastructure have been reported [7]. As a result, the controller can be compromised. The security of those communications suffers from a single point of failure, which may be a self-signed certificate or a compromised certificate authority. For example, many implementations of SSL currently used in mission-critical systems suffer from man-in-the-middle attacks [8]. Moreover, the TLS/SSL model is not sufficient to establish trust between controllers and forwarding devices. Once an attacker has gained access to the control plane, it may be able to launch distributed DoS attacks by aggregating the resources of forwarding devices under its control.

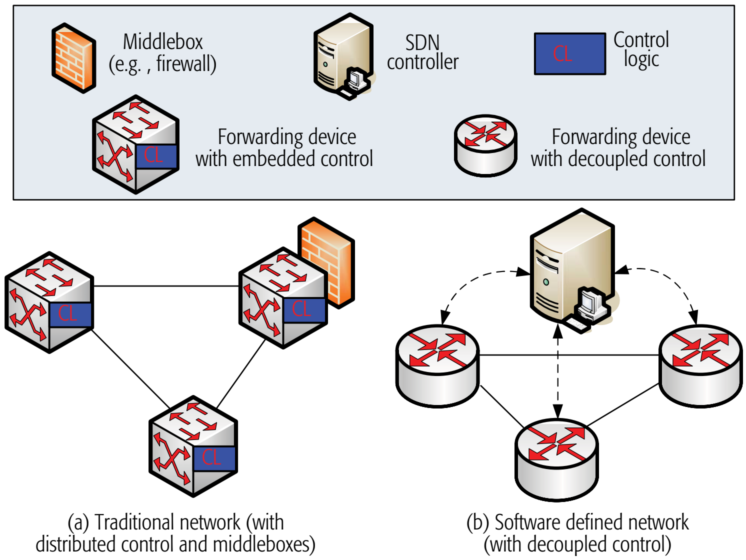

Figure 1. a) Distributed control plane in traditional networks; b) logically centralized control plane in SDN.

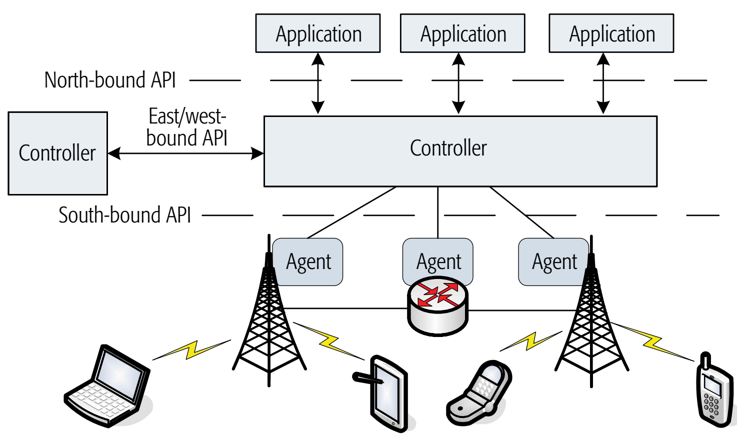

Figure 2. The conceptual architecture of SDWN.

4. Attacks on controllers: These could probably be the most severe threats to SDWNs. A malicious or faulty controller could compromise the entire network. Since it could be difficult to identify the exact combination of events that cause a particular malicious behavior, commonly used intrusion detection systems may not be applicable. Similarly, a malicious application can virtually do anything to the network, as the controller only provides abstractions that are used to issue configuration commands to the underlying infrastructure.

5. Lack of trust mechanisms between the controller and management applications: This is similar to threat 3 because trusted relationships cannot be established between applications and controllers. The major difference is how the certification is done since the techniques for certifying forwarding devices are different from those for applications.

6. Attacks on administrative stations: These machines are used in SDN to access the control-

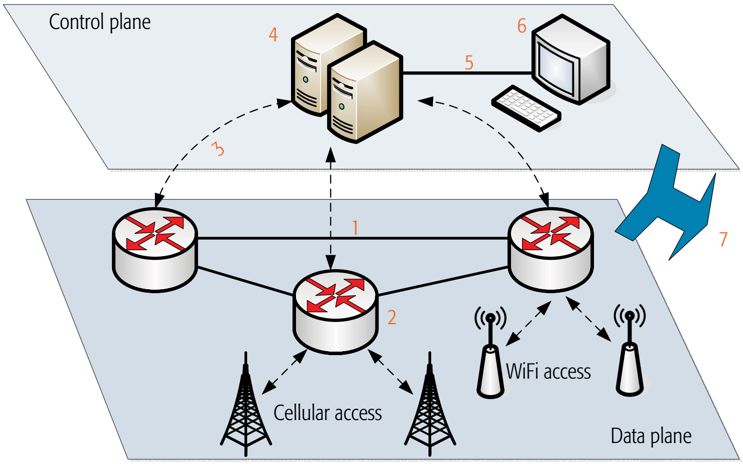

Figure 3. A vector of threats to SDWNs.

ler. In fact, they are already exploitable targets in traditional networks. For the case of SDN, the threat surface as seen from a compromised machine is even larger. For example, reprogramming the entire network from a single location becomes much easier.

7. Lack of trusted resources for forensics and remediation: Such resources help to understand the cause of a detected problem and carry out subsequent recovery to secure mode. Without reliable information from various components and domains of the network, it is difficult, if not impossible, to investigate and establish facts about incidents in question. Moreover, such information is useful only if it is trustworthy. Similarly, for remediation, reliable system snapshots are required for recovering network elements to a known working state quickly and correctly.

More Issues for Securing SDWN

Besides handling the above threats, the following issues also need to be taken into account when designing mechanisms to secure SDWN.

• **User Mobility**: Since users generally roam between networks using different access technologies, it becomes more difficult to detect anomalous activities and exchange security credentials.

• **Multiple Operators**: In a typical SDWN network, multiple operators are involved. This complicates the negotiation process between networks, raises privacy issuesm and possibly causes policy and quality of service (QoS) requirements conflicts. Also, interoperability needs to be supported.

• **Overhead**: Monitoring functions such as that provided by OpenFlow incur high overhead and have the weakness of incomplete sample information.

• **Compatibility**: Since different generations of mobile technologies are likely to be deployed at the same time by operators, any proposed security solutions should be backward compatible.

Security Requirements and Solutions

Although SDWN is a new networking paradigm, the standard network security requirements (i.e., confidentiality, integrity, availability, authenticity, authorization, non-repudiation, and consistency) are still applicable. At the same time, there are new requirements due to the specific characteristics of SDWN. Here, we discuss the security requirements of SDWN, which are also listed in Table 2 [10].

Authenticity: Authenticity refers to the property that entities are actually the ones they claim to be. This issue of forwarding devices in SDWN networks is similar to that in traditional networks. In both cases, the well established techniques for mutual authentication are applicable. Moreover, authentication and key establishment are closely related. Once two entities verify the authenticity of each other, they can establish some secret keys over the open wireless channel for subsequent secure information exchange. Traditional key generation and renewal algorithms should be adapted to take into account the characteristics of the wireless channel of SDWN. However, the authenticity of the controller and applications in SDWN is critical because the entire network can be compromised by a malicious controller or application. This issue is not relevant to traditional networks as they do not have the centralized controller and applications.

For SDWN, concerns have been raised about authentication and authorization mechanisms that simultaneously allow multiple entities to access network resources and provide appropriate protection of resources [11]. Generally, applications require different network privileges, thus a security model is needed to isolate applications and support network resources protection. Role-based authorization, such as FortNOX [12], is a potential solution. It resolves the issue of conflicting flow rules from two different applications. However, role-based authorization alone is not sufficient to deal with the complexity of SDWN to isolate applications or resources. The controllers are particularly vulnerable to attacks in the SDWN architecture open to unauthorized access and exploitation. Moreover, without a robust and secure controller platform, an attacker could possibly masquerade as a controller to carry out malicious activities. These threats can be mitigated by existing security technologies such as TLS to enable mutual authentication between the controller and forwarding devices. However, the use of TLS in the current specifications of OpenFlow [13] is only optional. A full specification for secure interface between controller and forwarding devices is needed.

Confidentiality: Confidentiality prevents information disclosure to unauthorized third parties. Its impact on SDWN networks and traditional networks is similar. Many techniques for the two common methods to ensure confidentiality, encryption, and access control have been developed for traditional networks and could be adapted to SDWN networks. Encryption of the communication channel between the controller and a forwarding device means that an attacker cannot recover the original plaintext even though it has access to the cipher-text. For example, TLS can be used to establish an encrypted channel. Access control ensures that only authorized entities have access to system data of the controller and forwarding devices. It may be enforced by the operating system.

Availability: Availability means that authorized users can access data, devices, and services whenever they have the need. In SDWN, if a forwarding device is not available due to either technical errors or DoS attacks, it can be mitigated by the controller to dynamically re-establish the network paths. Thus, similar to traditional networks, the issue of availability of forwarding devices is not as critical because network paths can be changed accordingly. However, non-availability of the controller in SDWN is critical. If the controller is unavailable due to a technical error, misconfiguration, or a DoS attack, the forwarding devices are only able to operate in a predefined way. For DoS attacks, possible solutions are rate-limiting mechanisms, discarding DoS-attack packets [14], or redundant controllers [15].

Integrity: It ensures that information has not been modified by any adversary. In SDWN, mainly the integrity of flow rules and protocol messages exchanged between the layers need to be ensured. The message authentication code is a commonly used approach to ensure integrity. The issue of flow rule integrity is critical in both SDWN networks and traditional networks since undesirable effects are caused by modified rules.

Consistency: This is about network traffic and control data. Generally, multiple applications could define flow rules and, as a result, could be inconsistent. One possible solution is to deploy a mediator between the controller and applications to resolve conflicting rules. FortNOX [12] is an example of such a mediator. This issue is critical in both traditional networks and SDWN networks because it can cause unpredictable behavior in both types of networks.

Fast Responsiveness: No matter whether security events are processed in reactive or proactive ways, it should be done in a timely fashion. This may involve efficient triggering and local optimization.

Adaptation: To take into account user mobility and dynamic network conditions, SDWN should be made adaptive by using mechanisms such as monitoring tools for network and user activities.

Num	Threats	Consequences in SDWN
1	Forged or faked traffic flows	Can be a door for DoS attacks.
2	Attacks on forwarding devices	The impact is potentially augmented.
3	Attacks on control plane communications	Communication with logically centralized controllers can be explored.
4	Attacks on controllers	Controlling the controller may compromise the entire network.
5	Lack of trust mechanisms between the controller and management applications	Malicious applications can now easily be developed and deployed on controllers.
6	Attacks on administrative stations	Now the impact is potentially augmented.
7	Lack of trusted resources for forensics and remediation	It is still critical to ensure fast recovery and diagnosis when faults happen.

Table 1. Threats inherited from SDN.

Property	Description
Confidentiality	To prevent information disclosure to unauthorized third parties.
Integrity	To ensure that information is not modified by any adversary.
Availability	To ensure that authorized users can access data, devices, and services whenever they have the need.
Authenticity	Entities are ensured to actually be the ones they claim to be.
Authorization	Only legitimate users can access resources.
Nonrepudiation	Users cannot deny any action that they have performed.
Consistency	To ensure that flow rules defined by different applications have no conflict.
Fast responsiveness	Security events should be processed in a timely fashion.
Adaptation	To take into account user mobility and dynamic network conditions.

Table 2. Security requirements.

Attacks and Countermeasures on SDWNs

SDWNs are subject to a variety of security attacks such as spoofing, tampering, repudiation, information disclosure, DoS, and elevation of privileges. Table 3 maps attacks to the properties that guard against them. In the following, we give a glimpse of recent developments of countermeasures to attacks by focusing on solutions to information disclosure and DoS attacks.

Information Disclosure: The objective of this attack is to exploit the use of flow aggregation to extract some network state information. An attacker can use such information to determine the nature and presence of services on a network, which may be useful in a later stage of an attack. Approaches to mitigate this attack should aim to prevent the internal system states from being disclosed in the observable system parameters. The following approaches are potential candidates.

1. Proactive Strategies: The establishment of proactive flow rules make the response time independent of the network states. Of course, this situation may be worsened by automatic flow aggregation techniques as an attacker might infer the presence of another connection, which is aggregated with its current one.

2. Randomization: The statistical uncertainty of an attacker can be increased while the strength of the attack can be reduced by increasing the variance of response times. For example, in OpenFlow, timeouts of the installed flow rules can be randomized to introduce unpredictable behavior. This prevents an attacker from having a coherent view of network states.

3. Attack Detection: Attacks based on timing analysis would exhibit distinctive and repetitive patterns. They could be exploited by controller applications to detect attacks and trigger countermeasures. Possible countermeasures include dropping suspicious traffic or adapting the forwarding strategies accordingly.

Denial of Service: DoS attacks can target forwarding devices and the controller, aiming to drain their resources so that the intended services are not available. One example is to send a large number of requests to the controller to install new flow rules in forwarding devices, leading to flow table overflow. The following are potential approaches to mitigate the attacks.

1. Packet Dropping and Timeout Adjustment: If attackers can be detected, flow rules can be

In SDWN, the controller sends policies to forwarding devices to instruct them how to deal with flows. Also, some policies may have to be sent to multiple forwarding devices. It is very important to keep these policies authentic and confidential because the wireless channel is insecure.

Attack	Security property
Spoofing	Authentication
Tampering	Integrity
Repudiation	Non-repudiation
Information disclosure	Confidentiality
Denial of service	Availability
Elevation of privilege	Authorization

Table 3. Attacks and security properties.

installed to identify malicious traffic and drop the misbehaving packets. On the other hand, if it is not possible to detect attackers, traffic prioritization and QoS mechanisms can be deployed to cope with the load. Moreover, flow timeouts can be adjusted to reduce the impact of DoS attacks.

2. Flow Aggregation: With this proactive strategy, multiple network flows are matched to a flow rule; the number of flow rules required to match traffic is thus reduced. It has the advantage that flow tables are less prone to overflows, and there is less load on the controller. Aggregated flow rules are suitable for networks such as backbone carriers, which deploy proactive strategies, but may not be applicable to enterprise networks, which deploy fine-grained control.

3. Access Control: This approach enforces access control lists stored as flow rules in forwarding devices. For example, traffic originating from a trusted domain is allowed to pass, while other incoming traffic is compared against a set of flow rules representing whitelists. This approach is particularly suitable for corporate networks, which have traffic coming from internal users or trusted external hosts.

Future Directions

SDWN influences defense against DoS attacks in negative ways. The middle-boxes or devices distributed within traditional networks are now located on top of the controller. Compared to packet processing based on hardware, it is much slower to process packets in software. The traffic overhead and network delay caused by the communications between the defense mechanisms and the forwarding devices could be the new attack surface. Thus, in the design of defense against DoS attacks, the computation and communication overhead must be taken into account in order to avoid introducing new security vulnerabilities.

The centralized and fine-grained control that comes with SDWN introduces a greater risk of outages due to errors made by network administrators. Misconfiguring of a controller by an administrator could significantly degrade network performance, even if the controller is functioning properly and the forwarding devices have no problematic flow rules installed. Detecting degraded network performance is challenging and needs further research. For example, even if Byzantine controller failures are assumed, it is difficult to determine what constitutes a fault for a properly functioning but misconfigured controller.

The threats to confidentiality, authenticity, integrity, availability, and consistency discussed in earlier sections are not totally new. Future work can leverage on established solutions to tackle these threats. Especially, there is increasing demand in the integration of public key infrastructures into SDWN in order to protect communications between different components of SDWN networks and to ensure authenticity of components.

Since a compromised controller can take control of the whole network, one simple attack is to suspend the functions of the controller. To deal such an attack, we can deploy multiple controllers in the network. Unfortunately, there are other possible attacks, such as issuing malicious commands from a compromised controller. Deploying multiple controllers would open more possible points of attacks. Therefore, more effective and secure mechanisms to protect controllers are needed.

In SDWN, the controller sends policies to forwarding devices to instruct them how to deal with flows. Also, some policies may have to be sent to multiple forwarding devices. It is very important to keep these policies authentic and confidential because the wireless channel is insecure. Work on designing policy distribution schemes that ensure the polices to be authentic and confidential simultaneously is needed.

Conclusion

SDWN, resulting from the extension of the SDN concept into wireless networks, will enjoy the benefits of cost-effective infrastructure upgrade, delivery of new services, and improvement of user experience to existing infrastructure. Similar to SDN, SDWN is vulnerable to new attacks due to physical separation of the control plane and data plane. Research work to address these issues have just commenced. We hope this article stimulates development of effective defense solutions to make SDWN attack-resilient.

Acknowledgment

This research is supported by a strategic research grant from City University of Hong Kong [Project No. 7004429], the Pearl River Nova Program of Guangzhou (No. 2014J2200051), the National Science Foundation of China (Grants: 51477056 and 61321064), the Shanghai Knowledge Service Platform for Trustworthy Internet of Things (No. ZF1213), the Shanghai Rising-Star Program (No. 15QA1401700), and the Specialized Research Fund for the Doctoral Program of Higher Education. Daojing He is the corresponding author of this article.

References

[1] D. Kreutz *et al.*, "Software-Defined Networking: A Comprehensive Survey," *Proc. IEEE*, vol. 103, no. 1, 2015, pp. 14–76.
[2] L. Li, Z. Mao and J. Rexford, "Toward Software-Defined Cellular Networks," *IEEE EWSDN*, Oct. 2012, pp. 7–12.
[3] X. Jin *et al.*, "SoftCell: Scalable and Flexible Cellular Core Network Architecture," *Proc. ACM CoNEXT*, 2013, pp. 163–74.
[4] A. Gudipati *et al.*, "SoftRAN: Software Defined Radio Access Network," *Proc. ACM HotSDN*, 2013, pp. 25–30.
[5] M. Bansal *et al.*, "OpenRadio: A Programmable Wireless Dataplane," *Proc. ACM HotSDN*, 2012, pp. 109–14.
[6] S. Ali *et al.*, "A Survey of securing Networks Using Software Defined Networking," *IEEE Trans. Reliability*, vol. 64, no. 3, Sept. 2015, pp. 1086–97.
[7] R. Holz *et al.*, "X.509 Forensics: Detecting and Localising the SSL/TLS Men-in-the-Middle," *Proc. ESORICS*, 2012, pp. 217–34.

[8] M. Georgiev *et al.*, "The Most Dangerous Code in the World: Validating SSL Certificates in Non-Browser Software," *Proc. ACM CCS*, 2012, pp. 38–49.
[9] S. Scott-Hayward, S. Natarajan, and S. Sezer, "A Survey of Security in Software Defined Networks," *IEEE Commun. Surveys & Tutorials*, DOI: 10.1109/COMST.2015.2453114, 2015.
[10] A. Y. Ding *et al.*, "Software Defined Networking for Security Enhancement in Wireless Mobile Networks," *Computer Networks*, vol. 66, June 2014, pp. 94–101.
[11] L. Schehlmann, S. Abt, and H. Baier, "Blessing or Curse? Revisiting Security Aspects of Software-Defined Networking," *Proc. CNSM*, 2014, pp. 382–87.
[12] P. Porras *et al.*, "A Security Enforcement Kernel for OpenFlow Networks," *Proc. ACM HotSDN*, Aug. 2012, pp. 121–26.
[13] ONF, "Software-Defined Networking: The New Norm for Networks," ONF White Paper, Apr. 2012.
[14] R. Kloti, V. Kotronis, and P. Smith, "OpenFlow: A Security Analysis," *Proc. IEEE ICNP*, Oct. 2013, pp. 1–6.
[15] B. Heller, R. Sherwood, and N. McKeown, "The Controller Placement Problem," *Proc. HotSDN*, Aug. 2012, pp. 7–12.

Biographies

DAOJING HE [S'07, M'13] (djhe@sei.ecnu.edu.cn) received his B.Eng.(2007) and M. Eng. (2009) degrees from Harbin Institute of Technology, China, and his Ph.D. degree (2012) from Zhejiang University, China, all in computer science. He is currently a professor in the School of Computer Science and Software Engineering, East China Normal University, P.R. China. His research interests include network and systems security. He is an Associate Editor or on the Editorial Board of some international journals such as *IEEE Communications Magazine* and *IEEE/KICS Journal of Communications and Networks*.

SAMMY CHAN [S'87, M'89] (eeschan@cityu.edu.hk) received his B.E. and M.Eng.Sc. degrees in electrical engineering from the University of Melbourne, Australia, in 1988 and 1990, respectively, and a Ph.D. degree in communication engineering from the Royal Melbourne Institute of Technology, Australia, in 1995. From 1989 to 1994 he was with Telecom Australia Research Laboratories, first as a research engineer, and between 1992 and 1994 as a senior research engineer and project leader. Since December 1994 he has been with the Department of Electronic Engineering, City University of Hong Kong, where he is currently an associate professor.

MOHSEN GUIZANI [S'85, M'89, SM'99, F'09] (mguizani@ieee.org) received his B.S. (with distinction) and M.S. degrees in electrical engineering, and M.S. and Ph.D. degrees in computer engineering from Syracuse University, New York, in 1984, 1986, 1987, and 1990, respectively. He is currently a professor and chair of the Electrical and Computer Engineering Department at the University of Idaho. He was previously a professor and associate vice president for graduate studies at Qatar University, Doha, Qatar. His research interests include computer networks, wireless communications and mobile computing, and optical networking. He currently serves on the Editorial Boards of six technical journals and is the Founder and EIC of the *Wireless Communications and Mobile Computing Journal* (Wiley; http://www.interscience.wiley.com/jpages/1530-8669/). He is a Senior Member of ACM.

Elastic Resource Utilization Framework for High Capacity and Energy Efficiency in Cloud RAN

Dario Pompili, Abolfazl Hajisami, and Tuyen X. Tran

Current radio access network architectures, characterized by a static configuration and deployment of base stations, have exposed their limitations in handling the temporal and geographical fluctuations of capacity demand.

This work was supported by the National Science Foundation (NSF) under Grant No. CNS-1319945.

Abstract

Current radio access network architectures, characterized by a static configuration and deployment of base stations, have exposed their limitations in handling the temporal and geographical fluctuations of capacity demand. Moreover, small cell networks have exacerbated the problem of electromagnetic interference and decreased the energy efficiency. Although there are some solutions to alleviate these problems, they still suffer from static provisioning of BSs and lack of inter-BS communication. Cloud RAN is a new centralized paradigm based on virtualization technology that has emerged as a promising architecture and efficiently addresses such problems. C-RAN provides high energy efficiency together with gigabit-per-second data rates across software defined wireless networks. In this article, novel reconfigurable solutions based on C-RAN are proposed in order to adapt dynamically and efficiently to the fluctuations in per-user capacity demand. Co-location models for provisioning and allocation of virtual base stations are introduced, and pros and cons of different VBS architectures are studied. Also, the potential advantages of VBS clustering and consolidation to support recently proposed cooperative techniques like cooperative multipoint processing are discussed.

Introduction

Over the last few years, the proliferation of personal mobile computing devices like tablets and smartphones, along with a plethora of data-intensive mobile applications, has resulted in a tremendous increase in demand for ubiquitous and high data rate wireless communications. The current practice to enhance data rate is to increase the number of base stations (BSs) and go for smaller cells to increase the band reuse factor. However, additional deployment and maintenance of a large number of BSs bring high inefficiencies due to excessive capital and operating expenditures. It has also been found that increasing the BS density or the number of transmit antennae will decrease energy efficiency (EE) due to the exacerbation of the electromagnetic interference problem and of the cooling requirements of cell site equipment [1].

On the other hand, the spatial distribution of users and the demand for capacity vary depending on the time of the day and week (the so-called *tidal effect*). In traditional cellular networks, each BS's spectral and processing resources are only used by the active users associated with that BS, causing idle BSs in some areas/times and oversubscribed BSs in others. The use of small cells is quite efficient in terms of power consumption as well as the utilization of spectral and processing resources when the capacity demand is high and evenly distributed in space. However, it becomes less so when the data traffic is low and/or uneven due to static resource provisioning and fixed power consumption. In this article, we discuss how the centralization of baseband units (BBUs), together with enabling virtualization of BSs while leveraging the paradigm of software defined wireless networking (SDWN), can be an effective way to address these challenges.

Cloud radio access network (C-RAN) is a new architecture for cellular networks where the BSs' computational resources are pooled in a central location; its main characteristics are:
- Centralized management of computing resources.
- Reconfigurability of spectrum resources.
- Collaborative communications.
- Real-time cloud computing on generic platforms.

C-RAN consists of three main parts:
1. Remote radio heads (RRHs) plus antennae, which are located at the remote site and are controlled by virtual BSs (VBSs) housed in centralized processing pools.
2. The BBU (VBS pool) composed of high-speed programmable processors and real-time virtualization technology to carry out the digital processing tasks.
3. Low-latency high-bandwidth optical fibers, which connect the RRHs to the VBS pool.

The communication functionalities of the VBSs are implemented (in software) on virtual machines (VMs) hosted over general-purpose computing servers that are housed in one or more racks of a small cloud data center. In a centralized VBS pool, as all the information from the BSs is resident in a common place, BSs can exchange control data at gigabit-per-second speed.

In this article, we propose a novel elastic resource utilization framework in which the

The authors are with the Department of Electrical and Computer Engineering, Rutgers University.

0163-6804/16/$25.00 © 2016 IEEE

VBS size, RRH density, and transmit power can be dynamically changed to meet fluctuations in per-user capacity demand. This *elasticity* brings significant improvement in user quality of service (QoS) as well as efficiency in energy and computing resource utilization within the C-RAN paradigm. Our solution includes a *proactive* and a *reactive* component: the former anticipates the fluctuation in per-user capacity demand and provisions the VBSs in advance for a certain (limited) horizon; the latter monitors the VM utilization and triggers over- or underprovisioning alerts when there is a mismatch between the expected resource utilization and the actual observation. We explore innovative models for VBSs that capture the effect of computing resource contention (CPU, memory, network interface) among co-located VBSs in racks or servers in the data center. We discuss pros and cons of different architectures ranging from the traditional all-in-one VBSs (like legacy BSs) to split PHY- and medium access control VBSs (MAC-VBSs), which is more suited to exploit specific hardware characteristics, minimize computing resource contention, and maximize resource utilization. We also present the novel idea of a *VBS-Cluster*, in which we merge VBSs serving a cluster into a unit VBS-Cluster while the RRHs' antennae in each cluster act as a single coherent antenna array distributed over a cluster region, and discuss its advantages.

The rest of this article is organized as follows. We present the state of the art; we describe the idea of elastic VBS and explain the proposed resource provisioning and allocation models; we introduce the VBS-Cluster idea, and explore some advantages that can be achieved through the cooperation of VBSs within a cluster; and finally, we draw our conclusions.

State of the Art

Centralized management of computing resources (i.e., BS pooling) renders information global, and hence enables cooperative communication techniques at the MAC and PHY layers that were previously not implementable due to the strict throughput/latency inter-BS coordination requirements. Examples of MAC- and PHY-layer enhancements include joint flow scheduling and load balancing, collaborative spatial multiplexing, interference alignment and cancellation, and advanced mobility management. Although work has been done on the aforementioned cooperative communication techniques that can benefit from the C-RAN characteristics, research on enabling technologies for C-RAN itself is at a nascent stage, so there are only a few works in this area.

In [2], a partitioning and scheduling framework is proposed that is able to reduce the compute resources by 19 percent. In [3], the authors present a solution for small cells that reconfigures the fronthaul based on network feedback to maximize the amount of traffic demand. The authors of [4] propose the concept of cell zooming, where the cell size is adaptively adjusted according to traffic load, user requirements, and channel conditions. The authors of [5] introduce a reconfigurable backhaul scheme to allow for a flexible mapping between the BBUs and radio access units (RAUs); through real-world experiments, they show that their proposed solution improves RAN performance and decreases energy consumption. In [6], the authors propose a cross-layer resource allocation model in which they optimize the set of selected RRHs and the beamforming strategies at the active RRHs in order to minimize the overall system power consumption. In [7], the authors explore the trade-off between full centralization and decentralization of BBUs, and provide an overview of the challenges for fifth generation (5G) networks and why cloud technology will be a key enabler for such networks. In [8], the authors propose low-complexity three-stage group-sparse beamforming algorithms to minimize the network power consumption in C-RAN. The authors of [9] consider the coordinated transmission problem to minimize the downlink power in C-RAN; in order to serve each user, they determine a set of RRHs and the precoding vectors for the RRHs to minimize the total transmission power subject to the fronthaul capacity constraint.

In contrast to prior works on C-RAN, we propose the idea of elastic VBSs and dynamic RRH density that adapt to the fluctuations in capacity demand on the fly through demand-aware dynamic VM provisioning and allocation. Moreover, we introduce the notion of VBS-Cluster, and propose innovative techniques where clustering, consolidating, and cooperation of VBSs improve the overall system performance.

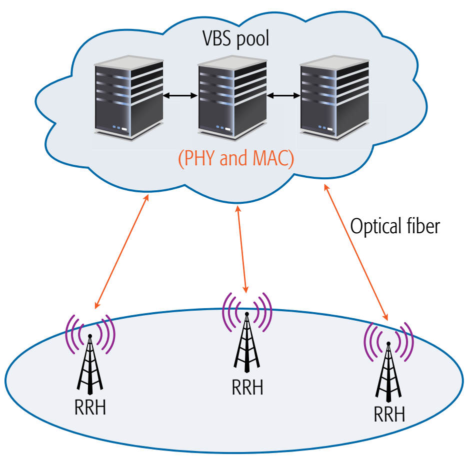

Figure 1. Cloud radio access network architecture, where the base stations are physically unbundled into virtual base stations and remote radio heads. Virtual base stations are housed in centralized processing pools and can communicate with each other at gigabit-per-second speeds.

In contrast to prior works on C-RAN, we propose the idea of elastic VBSs and dynamic RRH density that adapt to the fluctuations in capacity demand on the fly through demand-aware dynamic VM provisioning and allocation.

Demand-Aware Provisioning

The number of active users at different locations varies depending on the time of day and week. This movement of mobile network load is referred to as the *tidal effect*. Today, each BS's spectral and computing resources are only used by the active users in that BS's cell. Deploying small cells for *peak traffic* (i.e., for the worst case) leads to grossly underutilized BSs in some areas/at some times and is highly energy inefficient; conversely, deploying for the *average traffic* leads to

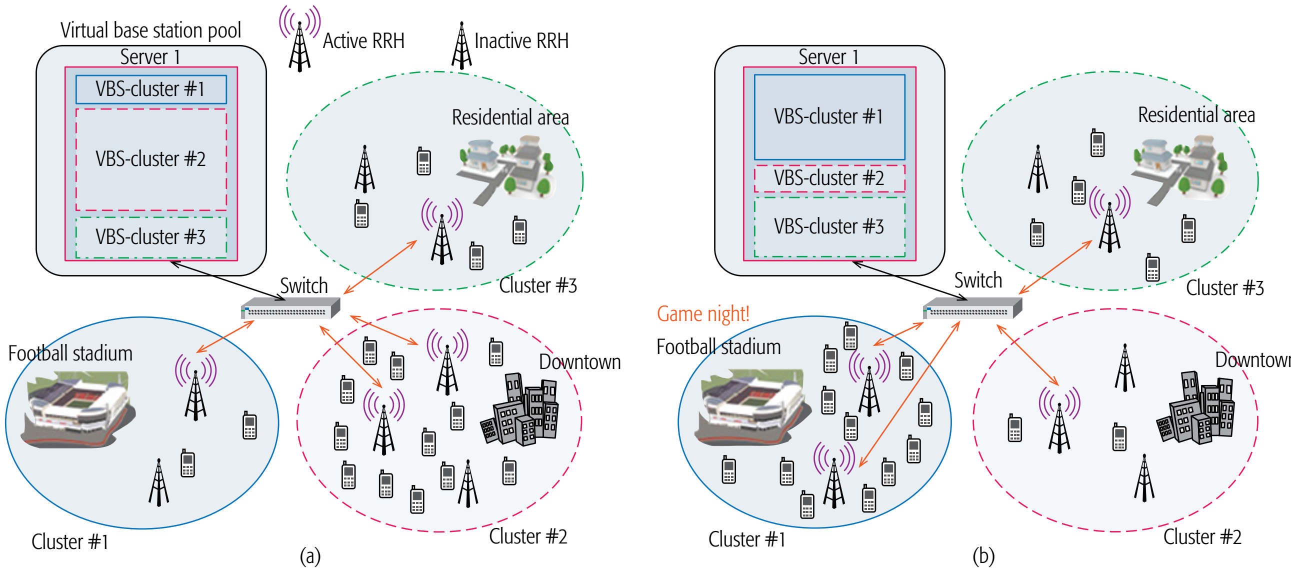

Figure 2. Virtualization in C-RAN allows for dynamic re-provisioning of spectral and computing resources (simplified here using different colored rectangles) to VBSs based on traffic demand fluctuation: a) and b) illustrate the movement of mobile network load from the downtown office area to the residential and recreational areas over the course of 24 hours, that is, during day and night, respectively; a) and b) also depict the corresponding changes in active RRH density and VBS size (note that active/inactive RRHs are identified by different icons, with or without wireless transmission).

oversubscribed BSs in some other areas/at other times. On the other hand, since traffic changes over time, there is no fixed cell size and transmission power that optimize the energy consumption; consequently, these system parameters can only be optimized for a fixed user density. Conversely, we propose a dynamic provisioning approach, at both the VBS and RRH sides, aimed at increasing the resource utilization and energy efficiency while providing a high level of QoS. As shown in Fig. 2, we cluster the neighboring RRHs and their corresponding VBSs, and change the density of active RRH and VBS size based on user density. Below we review the key features of our demand-aware provisioning approach.

Dynamic VBS resource provisioning: We advocate *demand-aware resource provisioning* in which VBSs are dynamically resized to meet the fluctuating traffic demand in the cellular network. As shown in Fig. 2a, during working hours, VBS-Cluster #2 will be provisioned with more computing resources compared to the ones serving a residential area (VBS-Cluster #3) or a stadium (VBS-Cluster #1). However, at night (Fig. 2b), the VBSs serving the stadium (e.g., on a game night) or a residential area will be provisioned with more resources than the ones downtown in order to meet the change in demand (VM resizing).

Size of a VM: All or some of the communication functionalities of a VBS (e.g., PHY, MAC, packet processing) are implemented in a VM. In order to achieve demand-aware dynamic VM provisioning, we introduce the notion of *size of a VM*, which is represented in terms of its processing power [CPU cycles per second], memory and storage capacity [bytes], and network interface speed [bits per second]. It primarily depends on the number of mobile users and type of data traffic (per-user capacity requirements) as well as on the computational complexity and memory footprint of the signal processing algorithms at the PHY layer, and the scheduling and frame processing algorithms at the MAC layer. In addition, the complexity of communication algorithms (e.g., the ones for inter-cell interference mitigation among densely deployed BSs in high-capacity RANs) may also affect the size of VMs. Therefore, to perform dynamic resource provisioning, a clear mapping from the number and type of mobile data users to the size of the VMs needs to be created. Hardware provisioning for VBSs must be such that the frame processing time be less than the frame deadline.

Known pattern vs. time-series prediction: Our solution for dynamic provisioning (or reprovisioning) of VBS resources to handle traffic fluctuations is composed of a *proactive* and a *reactive* component; in the former, the fluctuation in per-user capacity demand is predicted, and the computational resources are provisioned in advance for a limited time horizon. This anticipation is a result of *knowledge of known patterns* (e.g., day and night, weekdays and weekends, holidays, game schedules, etc.) or *predictions* based on advanced time-series analysis of historical traffic traces from the immediate as well as distant past. Once estimates of the number and combinations of different types of mobile data traffic are available, one just has to look up the closest profile and decide on the amount of resources to be provisioned for the VM.

Prediction uncertainties: Even though the proactive component allows for a smooth transition and greater optimization with respect to (w.r.t.) energy expenditure and resource utilization in the ensuing VM allocation procedure, it falls short in handling uncertainties. Some of the causes for uncertainties include unanticipated fluctuations in the number of users and per-user capacity demands in emergency scenarios arising out of natural (e.g., hurricanes, tsunamis) or man-made (e.g., industrial accidents, transportation system failures) disasters, unavailability of certain profiles, inaccuracies in the generated profiles, and mismatch between the generated profiles and reality due to hardware performance degradation. For these reasons, the reactive component monitors/profiles the CPU/memory/network utilization of the VMs and triggers over- or underprovisioning alerts when there is a "significant" mismatch between the expected resource utilization (based on the profile) and the actual observation.

A simple simulation scenario: To demonstrate the multiplexing processing gain (i.e., the increase in the region of feasibility) that can be achieved through dynamic resource provisioning, we simulated the following simple scenario: two BSs, one serving indoor users (or users in a downtown area with a large number of obstacles) and another serving outdoor users (say, a recreational area in a suburb). At each BS, we assume that each user's traffic belongs to one of the three following types with equal probability and in increasing order of priority:

- *Voice over IP* (very low and constant bit rate).
- *Light browsing* (bursty but low data rate).
- *Streaming/downloading* (high data rate).

We also assume that the cost of serving one user of each type at the downtown BS is higher than the corresponding cost at the suburb BS (this cost takes into account both the computational complexity as well as the memory footprint). We define *region of feasibility* as the total number of active users served by the BS pool with an acceptable blocking probability of 5 percent, which is a metric used in the context of voice calls. Here, for simplicity, we reuse the term to also convey an acceptable level of service degradation in data traffic. Figure 3 shows that dynamic resource provisioning (case 3) increases the region of feasibility (in terms of number of active users) by as much as 50 percent compared to the simplest static provisioning case (case 1). Note that knowledge of relative spatial distribution of users among BSs can help improve the feasibility region (case 2), but may also result in chronic over- and/or underprovisioning when the demand fluctuation is high. Greater benefits can be obtained when the distribution of users of different traffic types is unequal.

Dynamic RRH provisioning: Similar to what we mentioned above, deploying small cells (to provide enhanced spectrum resources for the peak traffic time) will make the network become energy inefficient due to the unavoidable energy costs when the capacity demand is low. For instance, circuitry, paging channel, cooling system, backhaul, and amplifiers all consume power so that even in a non-operational mode, BSs would consume a considerable amount of energy [10]. In traditional cellular networks, the cell planning and optimization, mobility handling, resource management, signal processing, and coverage are all done by each BS uniformly. In this case, even if the small cells have no traffic, they cannot be turned off [11]. Conversely, by decoupling BSs into VBS and RRH, the latter would only be responsible for providing spectral resources, and could be dynamically turned on and off as needed according to the traffic demand. Hence, in order to minimize energy consumption, we propose to optimize dynamically the number of active RRHs to adapt to the current traffic demand and user spatial distribution. For instance, as shown in Figs. 2a and 2b, due to the higher capacity demand during the day in cluster #2 (Fig. 2a), we provision it with more active RRHs than at nighttime (Fig. 2b) when we have lower capacity demand.

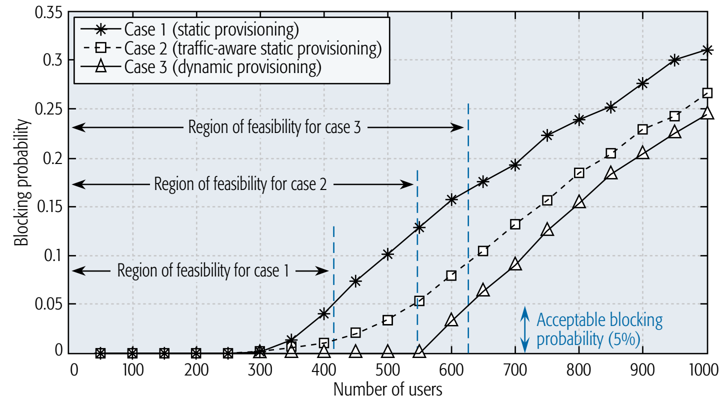

Figure 3. The benefit of dynamic computing-resource provisioning for VBSs at the remote data center: increase in the "region of feasibility" (w.r.t. active users).

Moreover, to minimize power consumption while ensuring a target data rate and user coverage, we need to adapt the transmission power of each cluster based on the density of its active RRHs. Both coverage and outage probability highly depend on the density and transmit power of the RRHs. This means that, given a fixed RRH density, we can minimize the transmit power of RRHs to target a certain coverage and outage probability. Since the RRH density of different clusters changes in time based on the capacity demand, we need to dynamically optimize the transmission power in each cluster. For instance, when the density of active RRHs becomes higher, each RRH has only a small coverage area, and users can receive acceptable signal-to-interference-plus-noise ratio (SINR) even when a lower output power is transmitted, which would save energy.

QoS-Aware VBS Allocation

Once the VMs holding the VBSs are provisioned, they have to be allocated to physical machines (PMs), that is, servers in the data center (called the centralized BS pool). The VM allocation has to be energy-, thermal-, and mobile-user-QoS-aware in order to fully realize the potential of C-RAN.

Thermal-aware VM consolidation: We advocate thermal-aware *VM consolidation* [12] for the VM-allocation problem. Thermal awareness, which is the knowledge of *heat generation* and *heat extraction* at different regions inside a data center, is essential to maximize energy and cooling efficiency as well as to minimize server system failure rates. Thermal-aware VM consolidation has the following three benefits:

1. The energy spent on computation can be saved by turning off the unused physical servers after VM consolidation.
2. The utilization of servers that are in the

In C-RAN, we are able to assign each cell to different clusters in order for them to cooperate with each other using different techniques. As associated VBSs of each cluster need a high-data-rate communication to perform cooperative techniques, they have to be allocated to the same server so to rely on high-speed inter-VBS connections.

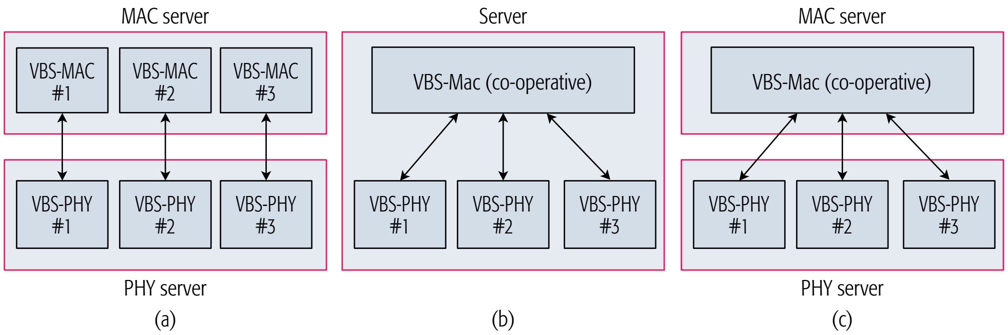

Figure 4. Split-VBS architectures: a) One-to-One mapping between PHY and MAC (different servers); b) Many-PHY-to-One-MAC (one server); c) Many-PHY-to-One-MAC (different servers). Note that the architectures in (a) and (c) (multi servers) can exploit the heterogeneity in datacenter server hardware.

"better cooled" areas of the data centers (i.e., with high heat extraction) can be maximized.

3. According to thermodynamics, heat can be extracted more efficiently (i.e., with a lower amount of work) by the cooling system from the consolidated server racks, which are hotter than non-consolidated racks.

In addition, consolidation on servers hosting VBSs allows efficient implementation of common functionalities such as signaling, channel state information (CSI) estimation for active users in a RAN, as well as for joint processing and scheduling techniques, such as ccordinated multipoint (CoMP) processing in 4G, for intercell interference mitigation.

Resource contention: Thermal and energy awareness alone, however, are insufficient for guaranteeing high VBS performance and maximizing energy and resource utilization efficiency. As multiple VMs share the same server resources (e.g., CPU, memory [RAM, cache], storage, and network interface), the performance of the corresponding VBSs in terms of per-user capacity and latency, and therefore the QoS of its mobile users, depend on the level of *contention for the computing resources* among co-located VMs. To factor in the effect of resource contention in VM allocation, we propose to classify the VBSs running a specific suite of algorithms for MAC- and PHY-layer functionalities as *CPU-*, *memory-*, and/or *network I/O-intensive*, and to develop co-location models that convey the degree of "compatibility" among co-located VMs. This way, we can incorporate the knowledge derived from co-location models into our VM allocation algorithm, thus making it QoS-aware.

Split-VBS architectures: To improve user QoS and resource utilization in C-RANs, we can deploy different architectures for VBSs. Figure 4 shows three possible/competing split-VBS architectures in addition to the traditional all-in-one VBSs in which the software modules for PHY and MAC are all implemented in one VM. The all-in-one architecture inherits characteristics from legacy BS designs, in which there is a one-to-one correspondence between MAC- and PHY-layer modules. *One of the primary motivations for this study is that the PHY and MAC layers are functionally quite different.*

One-to-one: PHY-layer processing requires vector execution techniques to accelerate signal processing, while MAC-layer processing requires multithread architecture and network accelerators for high-efficiency packet/protocol processing. In a data center with heterogeneous servers, exemplified as two separate servers in Fig. 4a (i.e., PHY server and MAC server), we can match the workload of BS-stack components with the capabilities of specific hardware.

Many-to-one: In general, communication between BSs can improve cellular system performance by exploiting the global and shared nature of information to make optimal decisions. For instance, in BS cooperation schemes, significant control information needs to be exchanged among neighboring BSs; however, cost, latency, and scarce interconnect capacity among BSs have been major impediments to the implementation of such schemes. We propose a split-VBS architecture, exemplified in Fig. 4b, in which the information of the MAC layers can be shared at gigabit-per-second speeds, making very-low-latency inter-BS communication possible. As a result, faster mobility management, more sophisticated interference suppression, and advanced cooperative multiple-input multiple-output (MIMO) techniques can be implemented to improve the user QoS. Finally, in order to take advantage of the heterogeneous processing pool as well as the high-speed inter-BS communication, we propose the architecture in Fig. 4c in which the VBS-MACs are merged together, and different physical servers are used for PHY- and MAC-layer processing.

Advantages of VBS Consolidation

In current distributed cellular systems, BSs can barely communicate with each other as the messages among BSs have to be exchanged through costly backhaul links. In C-RAN, as all the VBSs are located in a common rack of servers, they can exchange data with each other at gigabit-per-second speeds. Also, clustering the VBSs of the neighboring cells — together with enabling the coordination of the VBSs in the cluster — can greatly improve the system performance by exploiting the extra degrees of freedom, thus making optimal decisions.

We introduce the novel idea of a *VBS-Cluster*, according to which:

- All the VBSs associated with a certain cluster are merged together.

- The RRHs' antennae in each cluster act as a single coherent antenna array distributed over the cluster region.

Figure 5 shows two VBS-Clusters, #1 (on the left) and #2 (on the right), where the sizes of the clusters are 2 and 3, respectively. Since in C-RAN VBSs are implemented on VMs, the size of VBS-Clusters (in terms of number of VBSs) can also be changed based on the network requirements. In such a case, the serving VBS-Cluster sends a CLS-REQ message to the target cluster to check whether it is ready to change the cluster size. As a response, the target cluster sends back a CLS-RSP message to the serving cluster to report whether it approves or rejects the CLS-REQ. If the decision is to change the size, the serving VBS sends a VBS-REQ message to the candidate cluster as a request to join. At this point, the target cluster acknowledges the VBS-REQ by sending a VBS-ACK to the serving VBS, which is finally added to the cluster.

In C-RAN, we are also able to assign each cell to different clusters in order for them to cooperate with each other using different techniques. As associated VBSs of each cluster need high-data-rate communication to perform cooperative techniques, they have to be allocated to the same server to rely on high-speed inter-VBS connections. Moreover, as the number of active users in the cluster determines the size of the VBS-Cluster, resource allocation needs to be performed for each cluster. We present here a few scenarios where clustering the VBSs to enable cooperation improves system performance.

Mobility management: In 4G wireless networks, only hard handover (HHO) (in which the connection between the serving BS and user is terminated before the connection between the new BS and the user is started) is defined to support users' mobility. As studied in [13], the service disruption time caused by HHO can be 250 ms or longer, which is intolerable for real-time services like voice over IP (VoIP). Note that with small cells, users perform handover more frequently, leading to a decrease in the perceived QoS; such degradation of QoS is a consequence of the short interruption in communication during HHO caused by overhead generated for controlling and managing the handover procedure itself. On the other hand, soft handover (SHO), which is a code-division multiple access (CDMA)-based handover scheme, can avoid service disruption as a user is actively connected to multiple BSs *simultaneously*. This contrasts with non-CDMA systems, in which a user can *only* be connected to one BS at a time. In C-RAN architectures, as the VBSs are co-located in a common place and can communicate and exchange data as well as controlling signals with each other, we are able to connect a user to multiple VBSs regardless of the modulation/access scheme. This means that we are able to use SHO for *both* non-CDMA and CDMA systems. By clustering VBSs, a user is actively connected to the associated RRHs as long as it remains in a certain cluster; in this case, a handover is needed less frequently (i.e., *only* when the user wants/needs to change the VBS-Cluster). To support CDMA in the VBS-Cluster, additional network resources are used; also, the associated VBSs need to perform a time correlation operation to detect the signal. On the other hand, to support a non-CDMA system, VBSs need to know the CSI from all the users, and matrix multiplications need to be performed to detect the signal.

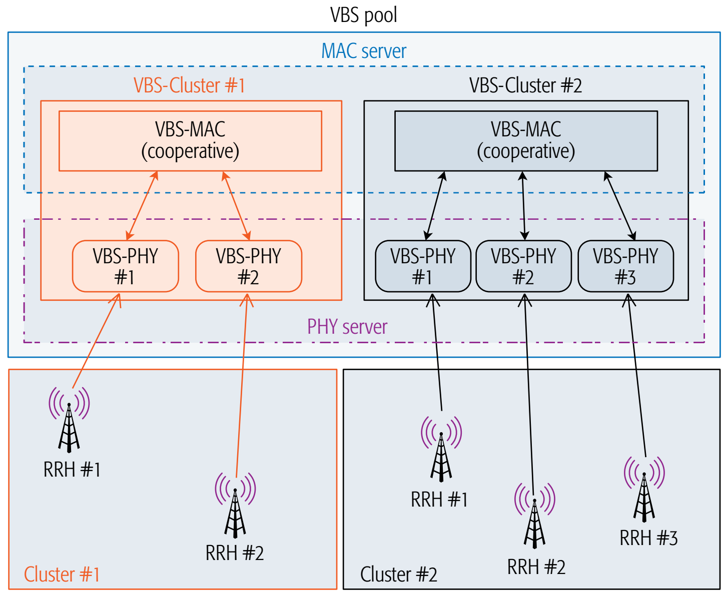

Figure 5. VBS-Cluster with a many-PHY-to-one-MAC architecture: VBSs associated with a cluster are merged together in the VBS-Cluster, and RRHs' antennae in each cluster act as a single coherent antenna array distributed over the cluster region.

Interference cancellation: In conventional cellular networks, each BS only serves users within its coverage area; thus, transmissions from neighboring cells interfere with each other, which decreases the SINR and spectral efficiency of cell edge users. A popular approach to address such interference is to employ CoMP, where neighboring cells are grouped together into clusters within which the BSs are connected to each other via the backhaul processing unit (BPU) [14]. In order to mitigate intra-cluster interference, the BSs in each cluster can perform coordinated beamforming and/or joint processing, which lead to improvements in spectral efficiency at the cost, however, of higher information exchange overhead among the BSs and more complex resource allocation. Although CoMP is able to reject the intra-cluster interference, it cannot mitigate the inter-cluster interference; consequently, cluster-edge users would still suffer from this type of interference. In addition, due to the distributed nature of the traditional cellular architecture, the latency and scarce interconnection capacity among the BSs have restrained the degree of cooperation among the BSs and the deployment of CoMP in practice.

These limitations can be overcome in C-RAN, where each cell can be associated with different clusters, and different clusters can communicate with each other at very high speeds. We envision a system employing CoMP over C-RAN to be highly capable of dynamically forming and reconfiguring user-centric clusters. In such a strategy, each scheduled user is in the center of its associated cluster, making it different from traditional static-clustering approaches where the cluster boundaries are fixed, and each cell belongs to one cluster only. This will eliminate cell-edge and

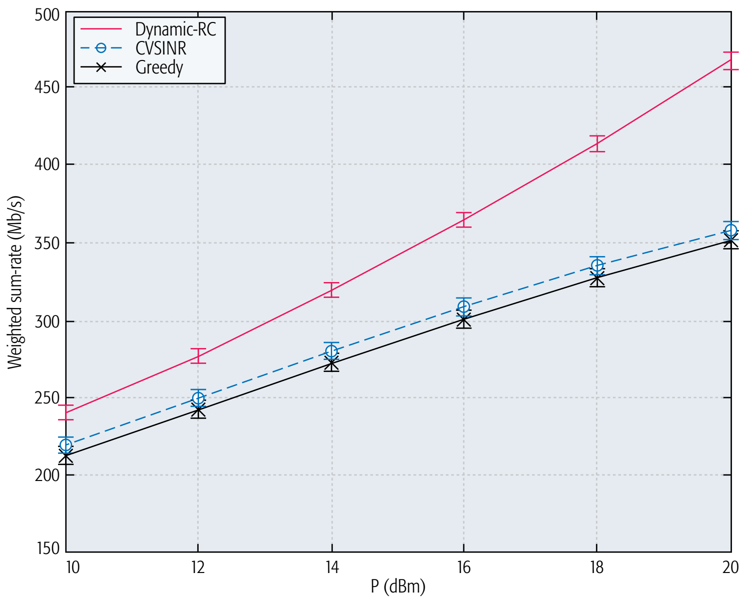

Figure 6. Improvement in downlink weighted sum rate (megabits per second) vs. RRH transmission power (dBm) of a C-RAN system employing dynamic clustering and cooperative beamforming. The competing strategies considered are dynamic-RC [15], a user-centric dynamic radio cooperation scheme; CVSINR, a user-centric heuristic clustering scheme with clustered virtual SINR beamforming; and greedy, a non-overlapping clustering scheme with zero-forcing beamforming.

cluster-edge users, mitigating both inter-cell and inter-cluster interference. The result of our work [15], shown in Fig. 6, demonstrates that a C-RAN system employing dynamic user-centric radio cooperation enables more effective beamforming techniques and outperforms traditional systems.

Technical challenges and open research issues: BSs have stringent real-time, low-latency, and high-performance requirements, to meet which the traditional virtualization technique is challenged. Specifically, in order to deploy a real-time VBS pool, the following requirements need to be met:

- Advanced real-time signal processing algorithms as well as high-performance low-power computing optimized for wireless signals.
- High-bandwidth, low-latency, low-cost BBU interconnection topology among physical processing resources in the baseband pool, which include the interconnection among the chips in a BBU, among the BBUs in a physical rack, and across multiple racks in the data center.
- Efficient and flexible real-time operating systems to achieve virtualization of hardware/resource management as well as dynamic allocation of physical processing resources to each VBS. This is needed to ensure latency and jitter control at the hardware level to support virtualization smoothly and efficiently.

Conclusion

We present novel reconfigurable solutions in the context of the cloud radio access network, a new centralized computing paradigm based on virtualization technology that has emerged as a promising architecture for broadband wireless cellular access. Such solutions adapt dynamically to fluctuations in per-user capacity demand, and offer higher energy efficiency and data rate (even in high-mobility scenarios). We advocate the need for co-location models for provisioning and allocation of VBSs, propose different VBS architectures, and discuss their pros and cons. Also, we present the advantages of VBS clustering, which can enhance energy efficiency and capacity in wireless cellular systems via advanced collaborative communication techniques.

References

[1] C. Li, J. Zhang, and K. Letaief, "Throughput and Energy Efficiency Analysis of Small Cell Networks with Multi-Antenna Base Stations," *IEEE Trans. Wireless Commun.*, vol. 13, no. 5, 2013, pp. 2505–17.
[2] S. Bhaumik *et al.*, "CloudIQ: A Framework for Processing Base Stations in a Data Center," *Proc. ACM MobiCom*, Aug. 2012, pp. 125–36.
[3] K. Sundaresan *et al.*, "FluidNet: A Flexible Cloud-Based Radio Access Network for Small Cells," *Proc. ACM MobiCom*, Sept. 2013, pp. 99–110.
[4] Z. Niu *et al.*, "Cell Zooming for Cost-Efficient Green Cellular Networks," *IEEE Commun. Mag.*, vol. 48, no. 11, Nov. 2010, pp. 74–79.
[5] C. Liu *et al.*, "The Case for Re-Configurable Backhaul in Cloud-RAN based Small Cell Networks," *Proc. IEEE INFOCOM*, Apr. 2013, pp. 1124–32.
[6] J. Tang, W. P. Tay, and T. Q. S. Quek, "Cross-Layer Resource Allocation with Elastic Service Scaling in Cloud Radio Access Network," *IEEE Trans. Wireless Commun.*, vol. 48, no. 9, 2015, pp. 5068–81.
[7] P. Rost *et al.*, "Cloud Technologies for Flexible 5G Radio Access Networks," *IEEE Commun. Mag.*, vol. 52, no. 5, May 2014, pp. 68–76.
[8] Y. Shi, J. Zhang, and K. Letaief, "Group Sparse Beamforming for Green Cloud-RAN," *IEEE Trans. Wireless Commun.*, vol. 13, no. 5, 2014, pp. 2809–23.
[9] V. N. Ha, L. B. Le, and N. Dao, "Energy-Efficient Coordinated Transmission for Cloud-Rans: Algorithm Design and Trade-Off," *Proc. IEEE CISS*, Mar. 2014, pp. 1–6.
[10] O. Arnold *et al.*, "Power Consumption Modeling of Different Base Station Types in Heterogeneous Cellular Networks," *Proc. FutureNetw*, June 2010, pp. 1–8.
[11] I. Chih-Lin *et al.*, "Toward Green and Soft: A 5G Perspective," *IEEE Commun. Mag.*, vol. 52, no. 2, Feb. 2014, pp. 66–73.
[12] E. Lee, H. Viswanathan, and D. Pompili, "VMAP: Proactive Thermal-Aware Virtual Machine Allocation in HPC Cloud Datacenters," *Proc. IEEE HiPC*, Dec. 2012.
[13] W. Jiao, P. Jiang, and Y. Ma, "Fast Handover Scheme for Real-Time Applications in Mobile WiMAX," *Proc. IEEE ICC*, June 2007, pp. 6038–42.
[14] D. Lee *et al.*, "Coordinated Multipoint Transmission and Reception in LTE-Advanced: Deployment Scenarios and Operational Challenges," *IEEE Commun. Mag.*, vol. 50, no. 2, Feb. 2012, pp. 148–55.
[15] T. X. Tran and D. Pompili, "Dynamic Radio Cooperation for Downlink Cloud-RANs with Computing Resource Sharing," *Proc. IEEE MASS*, Oct. 2015.

Biographies

Dario Pompili [SM] (pompili@cac.rutgers.edu) is an associate professor with the Department of Electrical and Computer Engineering at Rutgers University, where he directs the Cyber-Physical Systems Laboratory (CPS-Lab). He received his Ph.D. in ECE from the Georgia Institute of Technology in 2007. He previously received his Laurea (integrated B.S. and M.S.) and doctorate degrees in telecommunications and systems engineering from the University of Rome "La Sapienza," Italy, in 2001 and 2004, respectively. He is a recipient of the NSF CAREER '11, ONR Young Investigator Program '12, and DARPA Young Faculty '12 awards. He is a Senior Member of the ACM.

Abolfazl Hajisami (hajisamik@cac.rutgers.edu) started his Ph.D. program in electrical and computer engineering at Rutgers University in 2012. Currently, he is pursuing research in the fields of C-RAN, cellular networking, and mobility management under the guidance of Dr. Pompili. Previously, he received his M.S. and B.S. from Sharif University of Technology and Shahid Beheshti University, Tehran, Iran, in 2010 and 2008, respectively. His research interests are wireless communications, cloud radio access network, statistical signal processing, and image processing.

Tuyen X. Tran (tuyen.tran@cac.rutgers.edu) is working toward his Ph.D. degree in electrical and computer engineering at Rutgers University under the guidance of Dr. Pompili. He received his M.Sc. degree in electrical and computer engineering from the University of Akron, Ohio, in 2013, and his B.Eng. degree (Honors Program) in electronics and telecommunications from Hanoi University of Technology, Vietnam, in 2011. His research interests are in the application of optimization, statistics, and game theory to wireless communications and cloud computing.

A Software-Defined Wireless Networking Enabled Spectrum Management Architecture

Wei Wang, Yingjie Chen, Qian Zhang, and Tao Jiang

Abstract

Recent years have seen the proliferation in versatile mobile devices and application services that demand different data rates and latencies. Fixed channelization configuration in today's wireless devices fail to be efficient in the presence of such dynamic demands. In this regard, fine-grained spectrum management designs have been advocated by the research community to embrace the heterogeneity in devices and services. However, manufacturers hesitate to make hardware investments without comprehensive understanding of these designs. To break this stalemate, software-defined wireless networking (SDWN) has been pushed to market as a cost-effective paradigm. Motivated by recent innovations in SDWN, this article systematically investigates the spectrum management architecture design that reaps the benefits of SDWN while maintaining the features of fine-grained channelization. We shed light on design principles and key challenges in realizing the SDWN-enabled spectrum management architecture. With these principles and challenges in mind, we develop a general architecture with a new baseband virtualization design. We build a prototype that seamlessly integrates with the IEEE 802.11 protocol stack and commodity RF front-end. We demonstrate that the proposed architecture improves spectrum efficiency by emulating the upper layer behaviors using the traces captured in a campus WLAN.

Introduction

Recent years have witnessed a boom in versatile applications and heterogeneous wireless devices. A wireless device, such as a smartphone or laptop, may simultaneously run different types of applications, including video streaming services (Youtube, Netflix), cloud computing applications (Google photo auto backup, Dropbox, iCloud), and so on. These applications differ in required data rates and delay sensitivities. Moreover, the emerging of a new generation mobile internet devices, such as tablets, smartphones, and wearable devices, have augmented the heterogeneity in traffic demands.

The ever increasing heterogeneity in traffic demands has raised the stakes on developing new spectrum management architecture to utilize the limited spectrum resource in a cost-effective manner. The research community has realized that flexible channelization should be advocated to embrace the heterogeneous traffic demands. This vision is illustrated in Fig. 1, where mobile devices adopt different bandwidths according to their power constraints and service types. Wi-Fi access points (APs) communicate with smartphones on narrower channels to conserve power by using lower sampling rates, with tablets on medium-width channels to balance power and data rate requirements, and with laptops on wider channels to support traffic-intensive desktop services. Moreover, each mobile device runs versatile services with different latency and data rate requirements, which calls for finer-grained channelization.

To embrace the above vision in today's wireless LANs (WLANs), comprehensive understanding is required to properly manage the spectrum allocated to each AP [1]. Network operators need to configure each AP separately and even set service-specific preferences using vendor-specific commands. The configurations should be upgraded together with wireless protocols, which evolve continuously (once every few months) [2]. In addition to complex configurations, the operators should also have deep understanding of the impact of link dynamics and load changes. Ultimately, this situation has made the capital and operational expenses in spectrum management prohibitively high.

To fend off this ossification, both the telecommunication industry and the research community have placed considerable attention on a new paradigm, software-defined wireless networking (SDWN) [3], which creates a bundle of opportunities for managing current wireless networks in a cost-effective manner. Notably, SDWN simplifies network management by decoupling the control plane logic from the data forwarding plane. The control plane is logically centralized, while it can be implemented using a centralized controller or multiple controllers distributed in the network. As such, the logic of traditional networks is abstracted from within the hardware implementation and raised to a higher software level that can easily be manipulated by network operators.

Despite growing attempts and extensive

The authors systematically investigate the spectrum management architecture design that reaps the benefits of SDWN while maintaining the features of fine-grained channelization. We shed light on design principles and key challenges in realizing the SDWN-enabled spectrum management architecture.

Wei Wang is with Huazhong University of Science and Technology and Hong Kong University of Science and Technology; Yingjie Chen and Qian Zhang are with Hong Kong University of Science and Technology; Tao Jiang is with Huazhong University of Science and Technology.

0163-6804/16/$25.00 © 2016 IEEE

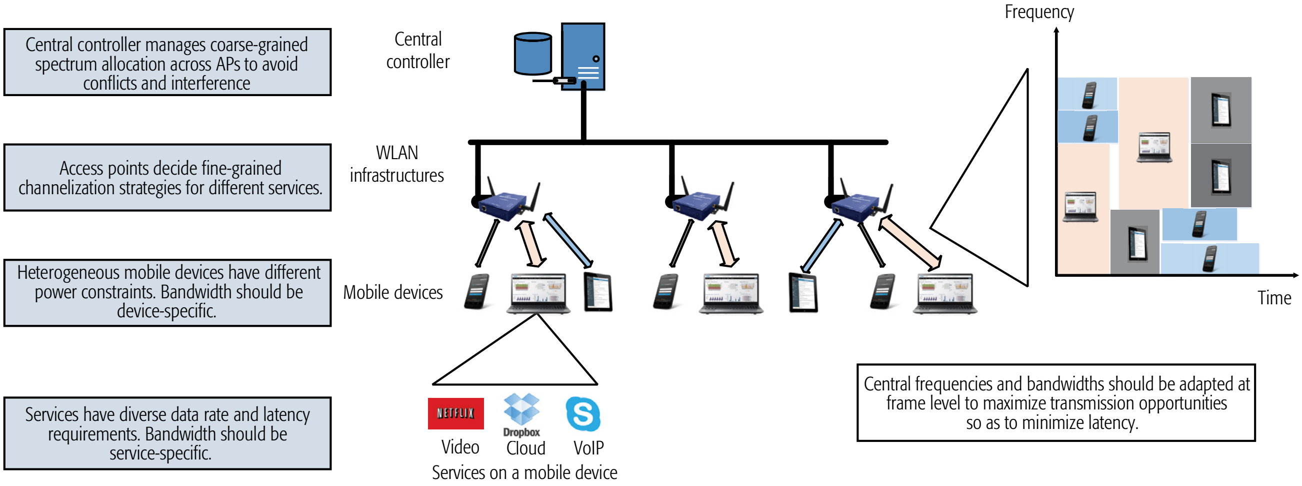

Figure 1. Illustration of fine-grained spectrum management.

efforts on SDWN-enabled management architectures [2, 4, 5], few have systematically investigated the spectrum management architecture to facilitate fine-grained channelization. The architecture should be carefully designed to ensure that the whole network stack operates reliably while performing fine-grained channelization at the lower layers, that is, it has to account for the mutual impact of the application requirements and the spectrum dynamics at the physical (PHY) and media access control (MAC) layers. Another challenge stems from the overhead incurred by spectrum adaptation, which mainly manifests itself in handling the spectrum agreement between transmission pairs and the coordination among multiple links.

In this article, we explore how the SDWN paradigm can be leveraged to efficiently overcome the limitations in current spectrum management architectures, and call attention to a clean-slate redesign of the fine-grained channelization mechanism. Specifically, we start at a deep dive into the features of the SDWN paradigm, and then review the challenges in realizing the SDWN-based architectures for spectrum management. We propose a spectrum management architecture that harvests the benefits of flexibility and programmability through the SDWN paradigm, while maintaining high efficiency in fine-grained channelization via a new baseband virtualization design. The benefits of our architecture are verified through experimental evaluation, and implications about future lines of research in architecture designs and applications are provided.

SDWN-Enabled Spectrum Management: Design Principles and Challenges

This section first presents an introduction to fine-grained spectrum management and the SDWN concepts. In particular, we highlight how the SDWN paradigm can be exploited to benefit spectrum management. We also investigate key SDWN-related technologies and review the state of the art. Then we discuss the challenges of realizing the SDWN-enabled spectrum management.

Fine-Grained Spectrum Management

This work considers a typical WLAN scenario as envisioned in Fig. 1, where a central controller interacts with APs in the WLAN through Ethernet backhaul. Each AP is associated with multiple clients, including different types of devices. Each client device may run a bundle of applications and services with different latency and data rate requirements.

In our vision, the bandwidth for each transmission is specified according to device and service types. On one hand, the channel bandwidth affects transmissions in many aspects. Wider bandwidth can provide higher throughput but require higher sampling rates for encoding and decoding, thereby consuming more power. Additionally, the transmission opportunities for links with wider bandwidth are lower as these links require larger amounts of vacant spectrum. On the other hand, today's wireless devices range from energy-constrained personal devices, such as smartphones and wearable devices, to powerful but data-hungry devices, such as laptops and personal computers. For each device, services have heterogeneous requirements on data rates and latency. Therefore, the transmission bandwidth of each AP should be dynamically adjusted to fit the requirements of devices and services.

Traditionally, such fine-grained spectrum management schemes are prohibitively complicated for network operators to realize on WLAN infrastructures, which may consist of devices from different manufacturers and are incrementally upgraded to support new protocols. To overcome this predicament, we borrow the SDWN architecture, the essential concepts of which are introduced in the following section.

SDWN Paradigm

SDWN is an emerging paradigm that has pioneered the introduction of programmability to network management. Architecturally, it contains three pillars, which are borrowed from the general SDN architecture [6]:

- **Decoupled control and forwarding planes**. Control logic is completely removed from network devices, which become simple data forwarding elements.

- **Network logic abstraction**. The logic of traditional networks is abstracted from within the hardware implementation into a higher level defined by software.
- **The presence of a programmable network controller entity**. The controller interacts with the underlying forwarding plane devices and coordinates their forwarding decisions.

By the clear separation of control and forwarding planes, SDWN exposes functions that have traditionally been deeply hidden in the network stack to higher levels [7]. The control plane interacts with higher layers via the *northbound* interfaces to understand operational tasks and network policies. The forwarding plane is controlled by the *southbound* interfaces, which refer to the interface and protocol between the controller and the SDWN-capable devices. These interfaces are generic, in that they are not tied to any particular system design or hardware platform architecture.

The control plane encodes the decision logic using a set of *rules* that compile higher-level policies and translate them into lower-level device configurations referred to as *actions* [6, 7]. Rules and actions are concrete representations of the separation of control and forwarding planes. In particular, rules determine the logic content, including scheduling and resource allocation, without dictating implementation details; actions, on the contrary, only specify functional behaviors such as signal processing operations, without knowing the logic content. Note that the SDWN paradigm follows on the heels of the SDN design principle to separate the control and forwarding planes, while SDWN differs from SDN in that wireless networks have distinct functions and lower-layer protocols, which should be carefully considered when implementing the SDWN architecture for distinct use cases in wireless networks.

Achieving SDWN-Enabled Spectrum Management: Designs and Challenges

The SDWN paradigm has promise to facilitate the spectrum management that can achieve the best of both worlds: we can maximize the spectrum efficiency and service quality through fine-grained channelization at the device side, meanwhile retaining simple management at the operator side. Architecturally, the SDWN-enabled spectrum management should provide the following capabilities:

- It is **self-configuring**, in that network operators do not need to understand the detailed signal processing procedures and when/how to apply lower-layer configurations.
- It **automatically translates** higher-level decisions into signal processing procedures and dynamically enforces the right configurations at MAC, PHY, and the RF front-end.
- It is **efficient**, in that it enables existing hardware to perform fine-grained channelization with merely lightweight overhead in both the time and frequency domains.
- It is **fully** compatible, in that it is integrated into existing infrastructures without modifying the existing protocol stack.

Realizing the above design goals requires systematic consideration from the high-level management architecture to the low-level baseband techniques. To make network management self-configuring and automatic, we need to properly design high-level management architecture. The SDWN paradigm has recently been applied to wireless network management in different aspects, from generic network architecture [3] and programmable forwarding plane [2] to specific use cases such as mobility management [4] and interference management [5]. These architectures follow the notion of separated data and control planes as envisioned by the SDWN framework, and introduce programmability and automation into wireless infrastructures to support a wide range of management techniques.

To make network management efficient and fully compatible, virtualization techniques at lower layers are the cornerstones to management architectures seamlessly integrating with the network protocol stack and devices based on software. To support spectrum adaptation for generic MAC/PHY protocols and RF front-ends, baseband virtualization techniques can be exploited to separate spectrum programmability from the general PHY modulation design [8]. Architecturally, spectrum adaptation functions are abstracted away from the PHY and RF front-end. As such, the protocol stack and the RF front-end are agnostic to the fine-grained spectrum dynamics as well as the underlying signal processing procedures, thereby adopting conventional configurations, such as modulation, pilot placements, and channel contention, without any modification [9].

Although the above innovations demonstrate significant benefits, in the case of fine-grained spectrum management we still face several design challenges that have not yet been fully explored:

Systematic integration. Both high-level management architecture and low-level baseband techniques should be seamlessly integrated with the conventional network protocol stack and hardware. It requires systematic investigation into how to precisely translate the upper-layer requirements into spectrum configurations at the PHY and RF front-end. Furthermore, spectrum management functions and interfaces to the network stack and RF front-end should be carefully defined.

Spectrum adaptation overhead. To achieve efficient channelization and spectrum adaptation, there are two practical hurdles:

- Out-of-band signal detection.
- Spectrum agreement.

As envisioned in Fig. 1, the spectrum band of one link is promptly adapted to maximize transmission opportunity. However, most existing approaches detect spectrum using spectrum virtualization techniques that are limited to signal detection within one channel (in-band signal detection). What prevents these techniques from out-of-band signal detection is that it results in frequency aliasing at the receiver. Additionally, before channel switching, senders and receivers need to agree on the transmission band, which is achieved by central coordination or separate control channels. Unfortunately, these approaches incur substantial overhead and are not prompt enough to respond to frame-level channel variance.

The protocol stack and the RF front-end are agnostic of the fine-grained spectrum dynamics as well as the underlying signal processing procedures, thereby adopting conventional configurations, such as modulation, pilot placements and channel contention, without any modification.

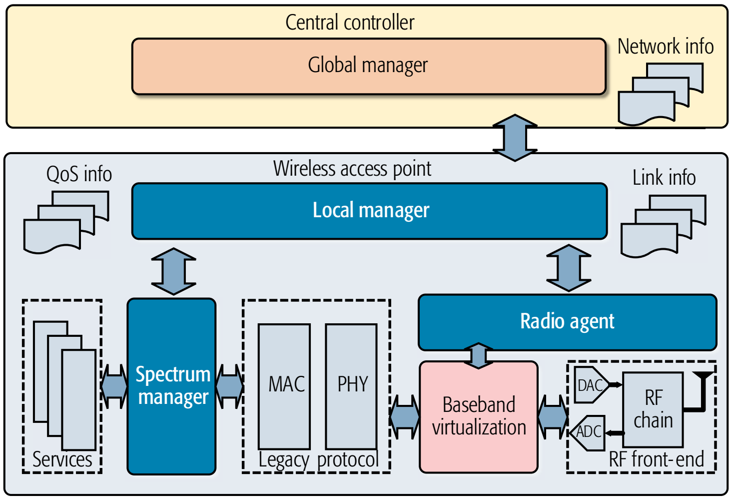

Figure 2. SDWN-enabled system architecture for spectrum management.

Spectrum Management Architecture

In this section, we develop an efficient spectrum management architecture that seamlessly integrates with the conventional network protocol stack and infrastructures.

Overview

Figure 2 outlines the architecture building blocks and their interfaces. Basically, the proposed architecture conforms to the SDWN framework in that it separates the control and forwarding planes. In particular, the key design components are described as follows.

- The control plane is realized using a two-tiered architecture: a top-level manager, referred to as the global manager, residing in the central controller, and a mid-level manager, referred to as the local manager, at or near each AP. Note that the local manager can be either a dedicated controller locating near the AP or a remotely programmable component within the AP. As such, the local manager can handle time-critical events with little latency and few load-intensive events, while the global manager can handle events that require global coordination [10]. The design rationale is that the central controller manages spectrum allocation across APs to avoid conflicts and interference, while delegating the traffic scheduling of applications and services to the local manager at each AP; that is, the local manager dynamically schedules which flow to transmit and its spectrum configurations.
- The control plane interacts with the network protocol stack and the RF front-end through a spectrum manager and a radio agent, which expose functions and information deeply hidden in the network protocol stack to the higher-level control plane.
- To support efficient and fine-grained spectrum adaptation, a decoupled baseband processing layer that employs baseband virtualization techniques is added between the legacy PHY/MAC layer and the RF front-end. By decoupling spectrum tuning and detection from packet decoding and scheduling, the legacy PHY/MAC protocols work independently, and the spectrum adaptation functionality can be integrated into existing devices without modifying the radio.

Our architecture focuses on the interaction and interfaces between the control and forwarding planes, while the management functions are simply implemented as part of the control plane. Note that we can also add a management plane to implement management functions and specify the management interface to configure the control plane.

Interface

At a high level, our architecture conforms to the SDN architecture [6] to define the southbound and northbound interfaces. In particular, the local and global managers interact with higher-layer applications through the northbound interface, and interact with lower layers and RF front-end through the southbound interface. We define interfaces based on their functionalities: we define the spectrum management interfaces to interact with the network protocol stack and the radio agent interfaces to interact with the RF front-end.

Spectrum manager interfaces. The spectrum manager acts as a central hub, coordinating the information flow between the control managers and the network protocol stack. On one hand, the spectrum manager extracts the quality of service (QoS) requirements of different types of service from the upper layers, while it retrieves the spectrum availability information from the MAC layer. All the collected information is forwarded to the local and global managers in order to assist them in traffic scheduling and spectrum adaptation. On the other hand, the control managers enforce the necessary configurations, including channel contention and traffic scheduling, into the network protocol stack through the interfaces of the spectrum manager.

Radio agent interfaces. As the spectrum adaptation and traffic scheduling decisions also rely on the link quality information, our architecture employs the radio agent to periodically pass the fine-grained link information such as channel state information (CSI) and clear channel assessment (CCA) to the control managers. Additionally, the radio agent controls the baseband and RF front-end to tune the bandwidth and central frequency of each transmission.

Global Manager

The global manager is the top-level manager residing in the central controller. Essentially, the global manager only handles global tasks that require coordination across APs, and offloads all AP-standalone functions to local managers. The global manager takes as input the network state from the spectrum manager interfaces and controls the channelization attributes of each AP through the local manager. The network state is described by an interference vector, which specifies the interference relations among APs, and a demand vector, which specifies the traffic load of each AP. The channelization attributes specify the bandwidth and central frequency of each AP.

Local Manager

The goal of the local manager is to dynamically schedule which flow to transmit and its spectrum configurations (i.e., the central frequency and bandwidth). The local manager takes the QoS

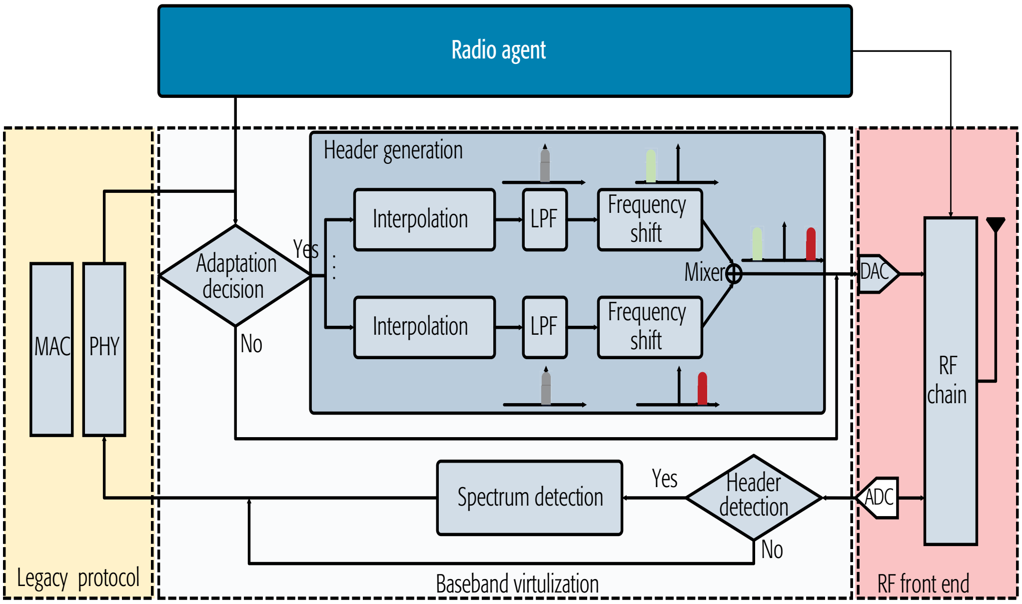

Figure 3. Architecture of baseband virtualization.

The goal of the local manager is to dynamically schedule which flow to transmit and its spectrum configurations, that is, the central frequency and bandwidth. The local manager takes the QoS requirements and the link information as input, and allocates the spectrum band assigned by the global manager to different flows.

requirements and link information as input, and allocates the spectrum band assigned by the global manager to different flows. Note that as the local manager is programmable, the QoS policy can be updated by the network administrator. In particular, the local manager can allocate the whole spectrum band to a single flow with high throughput requirement, or split the band into multiple orthogonal sub-bands, each of which acts as an independent channel to transmit a flow. Such fine-grained channelization can be realized via baseband virtualization, which is elaborated in the following section. To fulfill the throughput and latency requirements of each flow, the local manager prioritizes the traffic queue, and steers a flow's packets toward a particular sub-band for transmission. After arbitration, the sub-band allocation command is directed to the radio agent, while the channel contention and traffic scheduling commands are forwarded to the spectrum manager. Then the radio agent and spectrum manager compile these commands, and enforce corresponding configurations to the devices.

Basically, the local manager makes scheduling and allocation decisions based on service types. The service type is categorized based on required data rate and latency. The local manager tags each flow with a class identifier, with standardized values associated with corresponding characteristics, such as scheduling priority, packet delay budget, and packet error loss rate as defined in the Third Generation Partnership Project (3GPP) QoS Class Identifier (QCI) mechanism [11]. To enforce QoS requirements of different flows, the local manager controls the transmission properties, such as bandwidth, central frequency, data rate, and power, of a flow.

Baseband Virtualization

The global and local managers frequently change PHY configurations and need different basebands to support flows for multiple clients. To support such flexibility on commodity hardware, we need to build a software abstraction layer that decouples the tight connection between the PHY and RF front-end. This layer provides a virtual baseband that can be programmed by the control plane using the radio agent interfaces. The virtual baseband abstracts out the underly baseband dynamics and modifies the RF front-end for a given channelization configuration specified by the control plane. The baseband virtualization and its interactions with the radio agent are shown in Fig. 3. On one hand, the radio agent gathers link and channel statistics from lower layers and exports the statistics to the local and global managers. On the other hand, the radio agent exposes the functions of the virtual baseband to the local manager. The local manager controls the transmission attributes, including bandwidth, central frequency, data rate, and power, through the radio agent interfaces, and the virtual baseband takes the transmission attributes as input to reshape the baseband signals before feeding the signals to the RF front-end. The PHY/MAC exposes an interface to the virtual baseband to allow streams of complex digital baseband samples flowing between the layers. The virtual baseband receives spectrum adaptation decisions from the radio agent, and adds or removes extra preambles to those baseband samples for spectrum agreement and out-of-band detection. The prepended preamble exempts extra control frames for spectrum agreement and thus makes frame-level adaptation more prompt and efficient. As such, the functions of spectrum agreement and out-of-band detection are decoupled from packet encoding/decoding and channel contention, allowing the legacy PHY/MAC and RF-front-end to work independently without any modification. It is worth noting that although the management architecture is built on top of WLAN infrastructure and does not involve changes to the user terminal, the user terminal should be spectrum agile so that prompt spectrum adaptation can be performed.

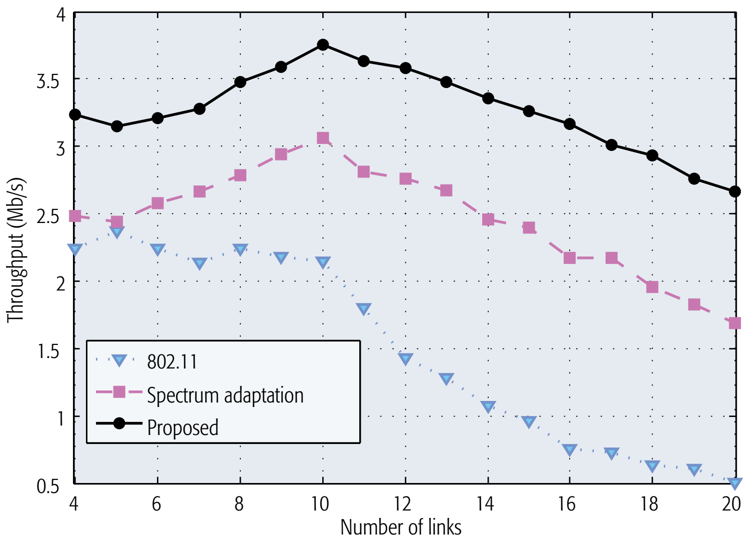

Figure 4. Throughput under various numbers of links.

Evaluation in a Campus WLAN Scenario

A primary motivation behind our spectrum management architecture is that the spectrum can be managed more efficiently when the fine-grained PHY functions are exposed to the upper-layer managers. The objective of this section is to demonstrate the efficacy of our architecture in promptly adapting spectrum according to link quality and upper-layer information.

Experimental setup. We realize the basic spectrum management blocks in existing OFDM PHY using commodity radios. We implement the entire baseband design directly in the Universal Software Radio Peripheral (USRP) hardware drive (UHD). Nodes in our experiments are USRP N210 devices equipped with RFX2450 daughterboards as the RF front-end, which operate in the 5.1–5.2 GHz range. Due to large processing delay of USRP hardware and limited power of general-purpose processors, the spectrum adaptation strategy cannot be performed in real time on USRP. To demonstrate the overall performance of the spectrum management architecture, we emulate the upper layer behaviors using the traces captured in a campus WLAN operating under IEEE 802.11n. In particular, we use Intel 5300 network interface cards (NICs) to send back-to-back frames and log the CSI and signal-to-noise ratio (SNR) traces at receivers. We vary the sender's and receiver's locations to measure 20 different links, with SNRs varying from –3 dB to 28 dB. Each link transmits 500 frames for every 20 MHz channel across an entire 80 MHz band.

To minimize unnecessary coordination overhead, the control plane triggers spectrum adaptation only when the channel availability or quality is unable to support reliable transmission. We adopt two metrics — *transmission opportunity* and SNR — to estimate the channel conditions. The transmission opportunity is defined to be the ratio of successful transmissions to the total number of transmission attempts. Only when the transmission opportunity or the SNR falls below the predefined threshold does the global manager reassign the central frequencies and bandwidths of APs. The central controller uses the following greedy spectrum adaptation strategy to assign channels. When an AP suffers from low transmission opportunity or low SNR, the local manager sends a spectrum adaptation request to the global manager via backhaul. The global manager goes through all possible channels, and selects the solution that maximizes the overall throughput in the WLAN. The global manager computes the overall throughput by measuring the SNR of each channel and mapping the SNR to corresponding data rate. Then it selects the spectrum adaptation solution that maximizes the overall throughput. The global manager calls off the adaptation if reassignment cannot improve the overall throughput. The global manager coordinates multiple APs and may simultaneously change the channel configurations of multiple APs. To be compatible with legacy Wi-Fi nodes, all nodes conform to the legacy data communication function (DCF) MAC (e.g., IEEE 802.11a/g/n/ac) to contend for channels. In particular, all nodes sense channel, backoff, and transmit as legacy nodes.

Baselines. In Fig. 4, we compare the proposed architecture (*proposed*) with two baseline approaches: 802.11 standard channelization (802.11) and the state-of-the-art spectrum adaptation approach [9] (*spectrum adaptation*). It is worthwhile noting that we take into account the spectrum adaptation overhead.

Figure 4 varies the number of links accessing the 80 MHz spectrum, which consists of four 20 MHz channels. The results show that on average, the proposed architecture outperforms 802.11 and spectrum adaptation approaches by 102 and 37 percent, respectively. When the number of links goes larger than 10, the throughput of all approaches decreases due to larger contention overhead. Figure 5 further compares the performance of all approaches when varying the total bandwidth. The number of links is set to eight. By leveraging the frequency diversity of multiple channels, the proposed architecture achieves higher throughput than the other two approaches when there is more than one channel.

The results demonstrate the merits of our architecture in making decisions by jointly considering the global interference relations as well as fine-grained PHY attributes. We show that it is feasible to harvest the benefits of flexibility and programmability using the SDWN paradigm, while still achieving high efficiency in fine-grained channelization. However, the preliminary implementation considers the bandwidth and central frequency of each flow; future spectrum management schemes can expose PHY functions, such as multiple-input multiple-output signal processing blocks, to high layers.

Concluding Remarks

This article has envisioned the crucial roles of the SDWN paradigm and baseband virtualization in achieving efficient fine-grained spectrum management in WLANs. Instead of modifying existing wireless devices or protocols, the SDWN-enabled architecture realizes the abstraction of fine-grained spectrum adaptation by employing decoupled components that seamlessly integrate with the network protocol stack or RF front-end. Through careful investigation of the pros and cons of existing approaches, we observe that the main hurdles in realizing the above vision lie in the systematic integration and spectrum adaptation overhead.

Under the design principles resulting from our investigation, we have presented a spectrum management architecture that reaps the merits of the SDWN paradigm and fine-grained spectrum adaptation. To efficiently support high-level SDWN-enabled architecture, we also devise a new virtual baseband that removes the need for extra spectrum coordination. The virtualized baseband is a clean-slate design that integrates with existing commodity radios without any hardware modification. We provide a case study on a campus scenario to demonstrate the benefits of the proposed architecture. We believe that the SDWN-enabled spectrum management architecture can contribute to future wireless networks by better supporting heterogeneous devices and versatile services.

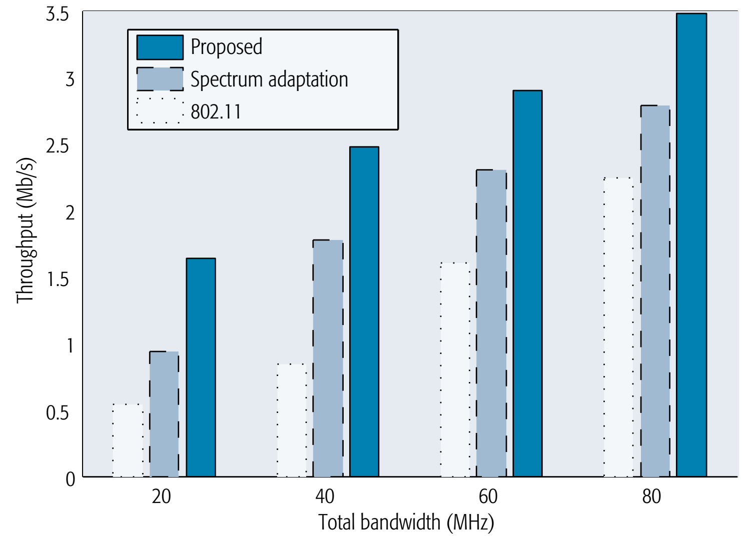

Figure 5. Throughput under various total bandwidths.

Acknowledgment

The research was supported in part by grants from 973 project 2013CB329006, the National Science Foundation of China with Grants 61502114, 61173156, 61428104, and 61401169, RGC under the contracts CERG 622613, 16212714, HKUST6/CRF/12R, and M-HKUST609/13, Joint Specialized Research Fund for the Doctoral Program of Higher Education (SRFDP) and Research Grants Council Earmarked Research Grants (RGC ERG) with Grant 20130142140002, as well as the grant from Huawei-HKUST joint lab.

References

[1] S. Rayanchu *et al.*, "FLUID: Improving Throughputs in Enterprise Wireless LANs through Flexible Channelization," *IEEE Trans. Mobile Comp.*, vol. 11, no. 9, 2012, pp. 1455–69.

[2] M. Bansal *et al.*, "OpenRadio: A Programmable Wireless Dataplane," *Proc. ACM HotSDN*, 2012, pp. 109–14.

[3] C. Bernardos *et al.*, "An Architecture for Software Defined Wireless Networking," *IEEE Wireless Commun.*, vol. 21, no. 3, 2014, pp. 52–61.

[4] L. Suresh *et al.*, "Towards Programmable Enterprise WLANs with Odin," *Proc. ACM HotSDN*, 2012, pp. 115–20.

[5] S. Kumar *et al.*, "Bringing Cross-Layer MIMO to Today's Wireless LANs," *ACM SIGCOMM Comp. Commun. Rev.*, vol. 43, no. 4, 2013, pp. 387–98.

[6] J. A. Wickboldt *et al.*, "Software-Defined Networking: Management Requirements and Challenges," *IEEE Commun. Mag.*, vol. 53, no. 1, 2015, pp. 278–85.

[7] H. Kim and N. Feamster, "Improving Network Management with Software Defined Networking," *IEEE Commun. Mag.*, vol. 51, no. 2, 2013, pp. 114–19.

[8] K. Tan *et al.*, "Enable Flexible Spectrum Access with Spectrum Virtualization," *Proc. IEEE DySPAN*, 2012, pp. 47–58.

[9] S. Yun, D. Kim, and L. Qiu, "Fine-Grained Spectrum Adaptation in WiFi Networks," *Proc. ACM MobiCom*, 2013, pp. 327–38.

[10] J. Schulz-Zander, N. Sarrar, and S. Schmid, "Aeroflux: A Near-Sighted Controller Architecture for Software-Defined Wireless Networks," *Proc. Open Networking Summit*, 2014.

[11] "3GPP TS 23.203 Policy and Charging Control Architecture," 2015.

Biographies

Wei Wang [S'10] (gswwang@cse.ust.hk) is currently a research assistant professor at Fok Ying Tung Research Institute, Hong Kong University of Science and Technology (HKUST). He received his Ph.D. degree from the Department of Computer Science and Engineering, HKUST, where his Ph.D. advisor was Prof. Qian Zhang. Before he joined HKUST, he received his Bachelor degree in electronics and information engineering from Huazhong University of Science and Technology, Hubei, China, in June 2010. His research interests include PHY/MAC in Wi-Fi networks, and privacy and fault management in wireless networks.

Yingjie Chen (ocgcyj@cse.ust.hk) received his M.Phil. degree from the Computer Science Department, HKUST, in 2012. He is currently a reseach assistant at HKUST. His research interests include PHY and MAC layer design in Wi-Fi networks, and mobile computing.

Qian Zhang [M'00, SM'04, F'12] (qianzh@cse.ust.hk) joined HKUST in September 2005, and is currently a full professor in the Department of Computer Science and Engineering. Before that, she was at Microsoft Research Asia, Beijing, from July 1999, where she was research manager of the Wireless and Networking Group. She has published about 300 refereed papers in international leading journals and key conferences in the areas of wireless/Internet multimedia networking, wireless communications and networking, wireless sensor networks, and overlay networking. She received the MIT TR100 (*MIT Technology Review*) World's Top Young Innovator award. She also received the Best Asia Pacific (AP) Young Researcher Award elected by IEEE Communication Society in 2004. Her current research is on cognitive and cooperative networks, dynamic spectrum access and management, as well as wireless sensor networks. She received her B.S., M.S., and Ph.D. degrees from Wuhan University, China, in 1994, 1996, and 1999, respectively, all in computer science.

Tao Jiang [M'06, SM'10] (taojiang@hust.edu.cn) is currently a Chair Professor in the School of Electronics Information and Communications, Huazhong University of Science and Technology. He received his B.S. and M.S. degrees in applied geophysics from China University of Geosciences, Wuhan, P. R. China, in 1997 and 2000, respectively, and his Ph.D. degree in information and communication engineering from Huazhong University of Science and Technology in April 2004. From August 2004 to December 2007, he worked at universities such as Brunel University and the University of Michigan-Dearborn. He has authored or co-authored over 200 technical papers in major journals and conferences and eight books/chapters in the areas of communications and networks. He has served or is serving as Symposium Technical Program Committee member for major IEEE conferences, including INFOCOM, GLOBECOM, and ICC. He was invited to serve as TPC Symposium Chair for IEEE GLOBECOM 2013, IEEEE WCNC 2013, and ICCC 2013. He has served or is serving as Associate Editor for some technical journals in communications, including *IEEE Transactions on Signal Processing*, *IEEE Communications Surveys & Tutorials*, *IEEE Transactions on Vehicular Technology*, and the *IEEE Internet of Things Journal*. He is a recipient of the NSFC for Distinguished Young Scholars Award in P. R. China.

Synergistic Spectrum Sharing in 5G HetNets: A Harmonized SDN-Enabled Approach

Auon Muhammad Akhtar, Xianbin Wang, and Lajos Hanzo

Given the anticipated 1000 times growth of tele-traffic, the conception of efficient spectrum management and sharing mechanisms becomes essential for 5G networks due to the scarcity of spectral resources. To meet this ambitious goal, we propose an SDN-based synergistic spectrum sharing technique, the harmonized SDN-enabled approach (HSA).

ABSTRACT

Given the anticipated 1000 times growth of tele-traffic, the conception of efficient spectrum management and sharing mechanisms becomes essential for 5G networks due to the scarcity of spectral resources. To meet this ambitious goal, we propose an SDN-based synergistic spectrum sharing technique, the harmonized SDN-enabled approach (HSA), which relies on distributed input reporting on spectrum availabilities in 5G heterogeneous networks. Conventional spectrum sharing mechanisms, particularly those using cognitive radio techniques, rely heavily on spectrum sensing and adaptive transmission-aided communications devices. However, this device-level spectrum sharing might lead to inaccurate sharing decisions due to the widely fluctuating signal quality observed by each device. Moreover, the network becomes vulnerable to malicious devices that report false sensing data to gain unfair access to the spectrum. These problems are further compounded by the fact that operational cellular networks have a somewhat rigid architecture relying on vendor-specific configuration interfaces, hence restricting the operators to have, at best, only indirect control of the network's operation and resources. In addressing these challenges, we adopt a synergistic approach that integrates the distributed sensing devices, cellular base stations, and SDN controller into an intrinsically amalgamated network. To alleviate any potential controller scalability and latency concerns, HSA harmonizes the network's operation by distributing the local decision making and network-wide policy making tasks among the base stations and the SDN controller, respectively. The article outlines the detailed network architecture and also presents an efficient resource management algorithm.

INTRODUCTION

With the advent of smart multimedia devices, mobile data traffic has seen an unprecedented growth over the last decade. It has been reported that the mobile data traffic in 2014 was 30 times higher than the entire global Internet traffic in 2000. This trend is predicted to continue as recent forecasts suggest that the global mobile data traffic between 2014 and 2019 will witness a further 10-fold increase [1]. The explosive growth of mobile traffic necessitates intensive global research efforts on 5th generation (5G) networks. However, the anticipated 1000-fold increase of network capacity to be supported by 5G will inevitably lead to significantly increased demand for spectral resources. Consequently, the conception of efficient spectrum management schemes becomes critical for 5G in achieving a dramatically increased network capacity. Additionally, 5G is also envisioned to have a heterogeneous network (HetNet) architecture, providing ubiquitous networking capabilities by using a diverse range of wireless networks and access technologies. The heterogeneous nature of 5G is mainly due to the inevitable network densification in the quest for high capacity and relying on a multi-tier architecture for providing ubiquitous coverage, bearing in mind the potential mixed use of legacy networks and the future 5G for reducing the deployment costs. As such, supporting the expected high network capacity in 5G in such diverse environments mandates the shared exploitation of the available spectral resources through dynamic coordination amongst devices and networks, which may be achieved by introducing software-defined networking (SDN) into 5G networks.

Spectrum sharing is capable of substantially improving the utilization efficiency of the available spectrum. Through spectrum sharing, mobile networks can exploit various unlicensed bands, including the industrial, scientific, and medical (ISM) band, visible light communication (VLC), [2] and the much-touted millimeter-wave (mm-Wave bands) [3]. It can be used to aggregate the channels of the licensed and unlicensed bands, thus supporting high-bandwidth applications to achieve data rates in the gigabit range in 5G. It also allows secondary (unlicensed) users (SUs) to access the licensed spectrum band, and this functionality can also be used for implementing device-to-device (D2D) communications and so forth. However, despite all the promise and interest generated by the research on spectrum sharing, widescale deployment of this technology remains yet to be seen. As detailed later, some of the reasons for the lack of practical implementa-

Auon Muhammad Akhtar and Xianbin Wang are with Western University; Lajos Hanzo is with the University of Southampton.

0163-6804/16/$25.00 © 2016 IEEE

tion include over-reliance on spectrum sensing, device-level decision making, and the lack of an efficient admission/eviction mechanism for SUs. Operators are hesitant to relinquish control of the spectrum by allowing SUs to make spectrum access decisions. Moreover, the current cellular networks have an inflexible architecture relying on vendor-specific equipment, which is remarkably complex and expensive, with an ever-increasing number of tunable parameters. These restrictions make it difficult for the operators to dynamically adapt their network parameters, which is another critical requirement of efficient and dynamic spectrum management.

In this article, we address the aforementioned challenges by designing a harmonized SDN-enabled approach (HSA) for 5G HetNets. To facilitate spectrum sharing, HSA adopts the principle of centralized management based on distributed inputs. The centralized management is facilitated by SDN and the cellular base stations (BSs), while the distributed inputs arrive from the SUs. Due to its open interface and ease of reconfiguration, SDN provides network operators with an ideal platform for complete global control over their entire network. However, it suffers from scalability and latency issues [4]. Especially in the context of wireless networks, latency may significantly impact the overall network performance. In order to address these issues, HSA harmonizes network operation by balancing the task distribution between the controller and the BSs. To this end, the controller is restricted to only managing the network rules and regulations by defining fine-grained operational policies, while the BSs are used for carrying out all the localized decision making based on the framework defined by the controller. It is only logical that the BSs at the edge of the access network carry out the local decision making, since they have the most up-to-date information about the rapidly varying wireless channel conditions. Instead of relying on spectrum sensing, we argue that the spectrum can be utilized more efficiently through database-assisted spectrum management. The article also presents a complete algorithm for optimal resource management.

The rest of the article is is organized as follows. We discuss related works and the challenges associated with current spectrum sharing techniques. Initially, we present a detailed architecture for software defined cellular networks. Following this, an optimal resource management algorithm is outlined and its performance is evaluated. Finally, we draw the conclusion to the article.

Related Works

Generally, spectrum sharing is a well researched area, especially in the context of cognitive radio (CR) networks [5]. However, despite all the research efforts spanning a period of more than a decade, CR-based spectrum sharing mechanisms have yet to be deployed on a large scale. One reason for this is the lack of a compelling business model that provides operators sufficient incentive to share their spectrum. Consequently, some recent research efforts have been focused on developing technologies that encourage network operators to share their spectrum. For example, in [6], the authors develop a scheme in which a mobile operator leases a channel in its licensed spectrum band to a privately owned WiFi access point (AP), and in return, the AP allows the mobile operator to offload its Internet traffic. Simulation results demonstrate significant throughput improvement for both the primary users (PUs) and SUs. The authors of [7] study resource allocation for indoor multihoming and multimode mobile terminals (MTs) in future HetNets. Specifically, the multihoming MTs aggregate resources from RF femtocells and VLC systems, while the multimode MTs only use one specific network at a time. It is shown that multihoming outperforms multimode under tight delay constraints. In [8], the authors propose both non-cooperative and cooperative architectures for cellular networks, and study the energy efficiency-spectral efficiency (EE-SE) trade-offs.

The SDN philosophy is a relatively new concept, and research in this area is still in its infancy. The idea of programmable networks has been around for many years. Nevertheless, the advent of the OpenFlow interface has given a new lease on life to SDN [4]. Consequently, some of the recent research efforts have shifted to introducing the SDN concept into cellular networks, with the aim of improving overall network performance and providing complete network control for the operator. In [9], the authors propose the concept of SoftCell, which aggregates traffic based on the service policy, BS locations, and user equipment (UE). By implementing packet classification through software-defined access switches and configuring optimal routing paths through specific middleboxes, SoftCell is shown to improve the flexibility of the cellular core networks of the future. In [10], Gudipati *et al.* introduce a software-defined centralized control plane of the radio access network, SoftRAN, which is based on a *virtual* big base station consisting of a centralized controller and the physical BS. It provides an architecture for coordinated radio resource management through the software defined centralized control plane of the RAN.

In the context of specific SDN applications in cellular networks, the authors of [11] study SDN-based mobile data offloading and load balancing. A partial offloading scheme is introduced in which the cellular network partially offloads the data traffic onto a WiFi network, under the constraint that the end-user quality of service (QoS) remains unaffected. Simulation results showed that SDN-based partial offloading and load balancing schemes significantly improve overall network performance. The authors of [12] introduce SDN-based efficient authentication handover and privacy protection for 5G networks. It is demonstrated that SDN-enabled security solutions are especially suitable for delay-constrained 5G networks.

In the context of specific SDN applications in cellular networks, the authors of [11] study SDN based mobile data offloading and load balancing. A partial offloading scheme is introduced, in which the cellular network partially offloads the data traffic onto a WiFi network, under the constraint that the end-user QoS remains unaffected.

Challenges in Spectrum Sharing

Current spectrum sharing approaches face several challenges. The most notable ones are outlined below.

Spectrum Sensing: Spectrum sensing is used to detect the presence of PUs and identify vacant spectral resources. The reliable detec-

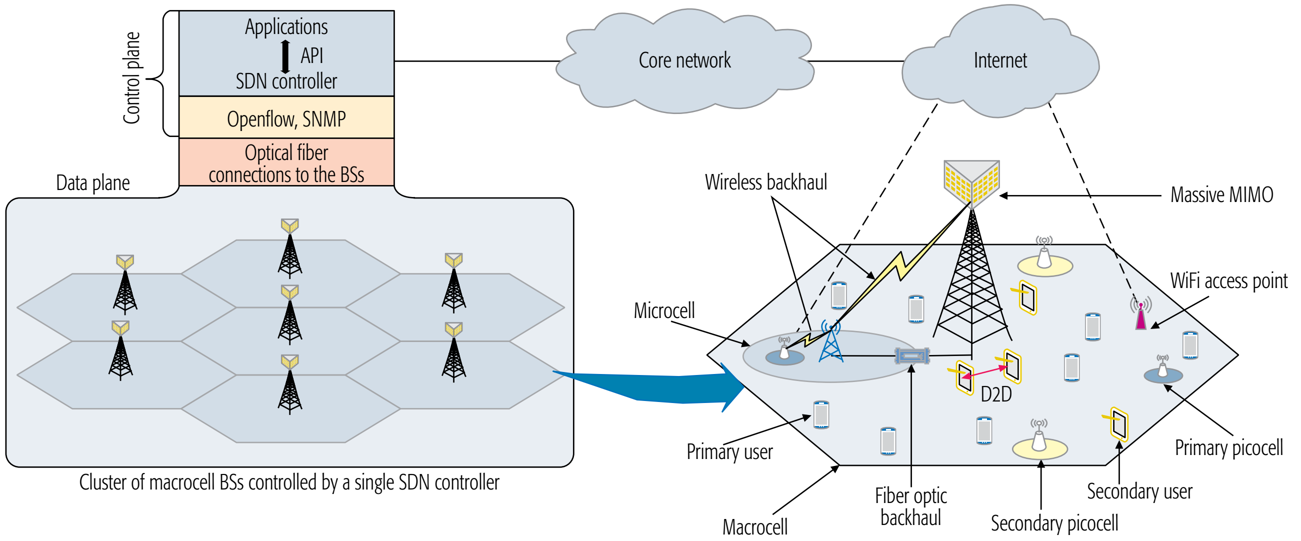

Figure 1. Software-defined 5G HetNet architecture that supports HSA. The macrocell BSs are directly connected to the controller, while the BSs of the smaller cells connect with the macrocells through a reliable backhaul link.

tion of a PU over wide bandwidth can be a challenge, not least because there is a wide range of transmission parameters employed by the PUs. The task becomes even more difficult due to the varying propagation losses and the interference generated by other SUs. Cooperative spectrum sensing has been shown to improve the performance of spectrum sensing techniques. However, as shown later, this performance improvement comes at the cost of substantial transmission overhead.

The two most important challenges associated with spectrum sensing are the problems of hidden terminals and missed opportunities [13]. The hidden terminal problem arises when the PU is passive, and the channel between the SU and the primary BS is in deep fade. Hence, the SU fails to detect the primary transmissions and mistakenly starts its transmissions. A suggested technique to tackle this problem is to allow the SUs roaming in the vicinity to opportunistically access the spectrum using the uplink frequencies in interference limited systems such as code-division multiple access (CDMA) [14]. However, in this scenario, if the SUs sense the spectrum, it will be deemed to be occupied by PUs (i.e., by mobile phones). Even a small number of PUs will render the spectrum unusable for the SUs, hence leading to missed spectrum access opportunities.

Decision Making: Most of the existing solutions conceived for opportunistic spectrum access rely on distributed decision making, allowing SUs to make decisions related to spectrum availability. The SUs invoke spectrum sensing algorithms to decide whether or not they should access the spectrum. As mentioned in the previous subsection, spectrum sensing has its associated drawbacks, and it is difficult to achieve high levels of accuracy in decision making. Moreover, such an approach also makes the network vulnerable to malicious devices that transmit false sensing data in order to selfishly access the spectrum. Lastly, allowing the SUs to make independent decisions is also problematic for operators, since they tend to improve the network's spectrum exploitation rate through complete control over their spectrum. Thus, to facilitate harmonious coexistence between the PUs and SUs, and to encourage the operators to be more open to the idea of spectrum sharing, all decision making related to spectrum access should be carried out in a centralized manner, for example, at the BS.

Admission Control: Admission control is another critical challenge for optimal spectrum sharing. Due to PU activities, the availability of frequency resources fluctuates quite erratically; hence, the quality of service (QoS) requirements of SUs can only be satisfied by relying on the efficient integration of admission and eviction control schemes into the medium access control (MAC) layer. This integration enables the MAC to adapt to the available communication resources. The design of an efficient admission/eviction control mechanism still remains an open research challenge.

Software-Defined 5G HetNets

In order to address the aforementioned challenges, we adopt HSA to design an SDN-based synergistic 5G HetNet relying on centralized management based on distributed inputs. The distributed inputs arrive from the SUs, which sense the spectrum and report their findings to their nearest BSs. The centralized management is carried out by the macrocell BSs and SDN controller. As mentioned previously, this task sharing between the BSs and the controller eases the burden imposed on the centralized SDN controller and harmonizes overall network operation. In order to minimize the dependence on spectrum sensing, HSA restricts the BSs to utilizing the SUs' inputs only as side information. Spectrum sharing decisions are based on the policies defined by a centralized SDN controller and the contents of a local database (LDB) maintained at each BS. The LDB is updated whenever a PU or an SU enters or leaves the network.

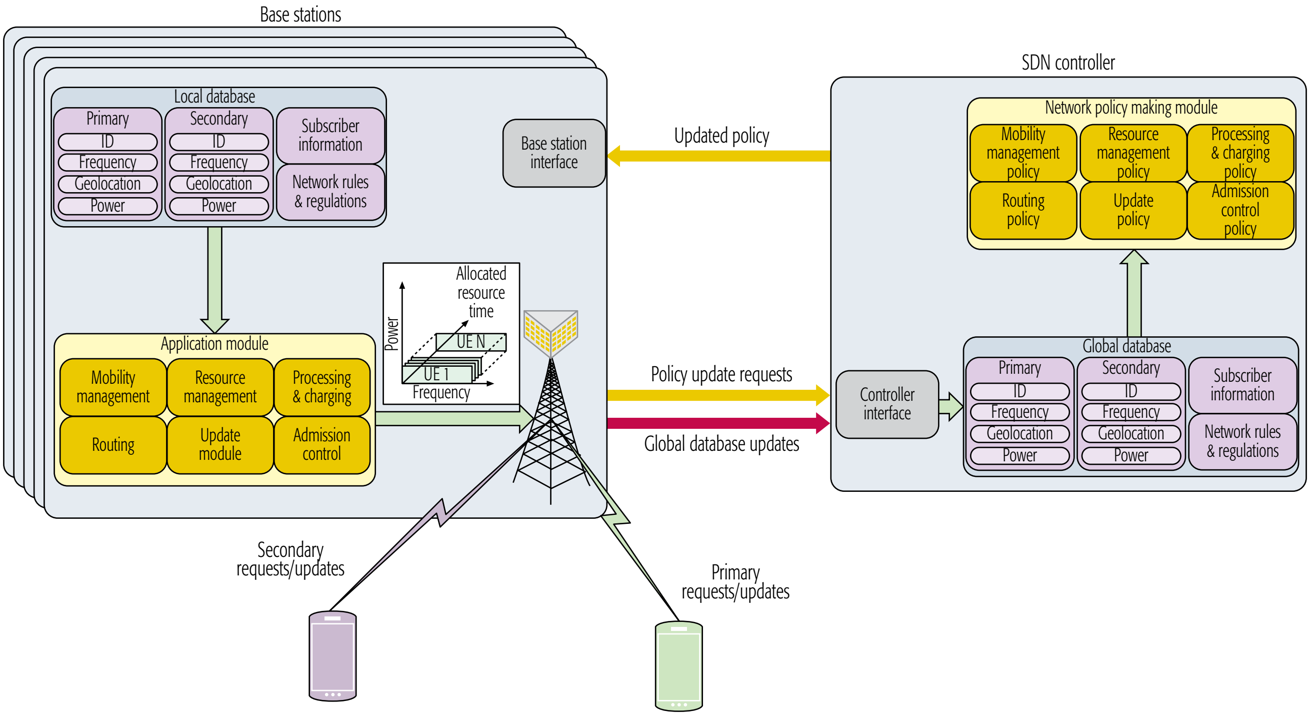

Figure 2. Operational architecture of the software-defined cellular network. The SDN controller decides the policies related to various network functionalities while the BSs implement the controller-defined policies.

The Network Architecture

Bearing in mind the 5G network vision, the proposed architecture, which supports HSA, consists of a HetNet environment, as shown in Fig. 1. It has a multi-layered architecture consisting of macrocells, microcells, and picocells, which communicate over the licensed spectrum band, and also incorporates the WiFi APs that communicate over the unlicensed band. The HetNet supports a multitude of features and technologies, such as traffic offloading, spectrum sharing, multihop, and D2D communications, and massive multiple-input multiple-output (MIMO). As shown in the figure, a cluster of macrocell BSs is controlled by a centralized SDN controller, which connects to the BSs through high-capacity fiber optic links and utilizes the OpenFlow protocol for controlling the data plane. Additionally, the Simple Network Management Protocol (SNMP) plugin shown in the figure enables the controller and the relevant applications to monitor and manage (management plane) the data plane devices, provided that those devices support an SNMP agent. In order to alleviate any potential scalability concerns, HSA restricts the controller to managing only a cluster of macrocell BSs. The BSs of the smaller cells in the HetNet regularly inform the macrocell about their load conditions through a wireless or fiber-optic-based backhaul link.

Figure 2 shows the operational architecture of the proposed framework, where multiple macrocell BSs are connected to the centralized SDN controller. Within the framework of HSA, the controller decides concerning the policies related to various access network functionalities, including resource allocation, handovers, transmit power allocation, and so on, and this feature forms the control plane of the radio access network. The base stations constitute the data plane of the network and implement the controller-defined policies. Below, we discuss these two components (i.e., the BSs and the SDN controller) in more detail. Unless stated otherwise, the term BS refers to a macrocell BS throughout this article.

The Macrocell Base Stations: In the proposed framework for 5G HetNets, each macrocell BS consists of an LDB and an application module. The LDB provides the BS with a view of cell load conditions and facilitates local decision making. It stores information about the users within the cell, and consists of a PU section and an SU section. Each of these sections is updated when the corresponding PU or SU accesses or leaves the spectrum slot. Additionally, the LDB also contains the subscriber information and network rules and regulations modules. The former module is related to subscriber attributes, such as the cellular network provider, device type, subscriber type, and recent usage. The network rules and regulations module contains the policies related to the different functionalities present in the application module, including packet processing and charging, resource management, mobility management, routing, update module, and admission control. The operator manages and enforces these policies through the centralized SDN controller. The BSs carry out all local decision making based on the regulations defined by the centralized controller.

SDN Controller: As shown in Fig. 2, the SDN controller incorporates a global database (GDB) and a network policy making module. The GDB contains information about all the users within the cluster and is updated regularly by the BSs. The controller utilizes the GDB to update the policies related to various functions of the network. The SDN controller grants greater control for network operators, since any modifications or updates to the network rules and regulations have to be made at a single control point. The controller updates the network policies either proactively (after a fixed amount of time) or reactively (a BS requests more resources due to cell overloading).

Resource Management Relying on HSA

The resource management process of the BSs in the proposed HSA-based 5G HetNets is shown with the aid of a flow chart in Fig. 3. The BSs reserve a channel for their communication with the SUs. Whenever the SUs have to access the spectrum, they rely on a low-complexity multiple access scheme, such as carrier sense multiple access (CSMA), in order to connect with the BS. On the other hand, the PUs, which have higher priority within the network, gain access by using the traditional multiple access schemes of cellular networks. As shown in Fig. 3, the BS initially waits for the users to request spectrum access. Upon receiving a spectrum access request, the BS first checks whether this request originated from a PU or an SU. Depending on the PU or SU status of the requesting user, a different course of action is followed.

Secondary User Access: In the proposed framework, each SU first senses the spectrum for the sake of detecting the presence of PUs. After completing the spectrum sensing process, the SUs report their findings to the BS, along with a request for spectrum access. The BS aggregates the inputs received from different SUs and utilizes this information to verify whether or not the contents of the LDB are up to date. If there is a mismatch, the BS checks by communicating with the corresponding spectrum users to validate the accuracy of the inputs gleaned from the SUs. The LDB is updated only if the SUs' inputs are proven to be accurate. Based on how often each SU's input is true/false, the BSs can also categorize the SUs as trusted/untrusted devices for security purposes. Upon successful completion of the aforementioned procedure, the BS estimates the location of each secondary source-destination pair and allocates an optimal spectrum band to each SU.

The optimal spectrum allocation is based on a number of factors, including the source-destination locations, the locations of the PUs, the specific frequencies already being used by the PUs and SUs, and so on. All of this information is present in the LDB. As an example, the decision could be based on the distance between the source and the destination. If the two are located close to each other, the BS could allocate, for example, mmWave frequencies. On the other hand, if the distance between the source and the destination is large, it would be better to allocate lower frequencies, since lower frequencies suffer from lower path loss and hence require less transmit power than mmWave frequencies. In addition to the distance between the source and the destination, the BS also takes into consideration the amount of interference imposed by the SUs on the primary network and vice versa.

Once the spectral resource has been allocated, the BS updates the LDB. If the cell becomes fully loaded after this resource allocation, the BS sends a GDB update along with a notification to the SDN controller requesting more spectral resources. The SDN controller exploits the information contained within the GDB to check if there are any lightly loaded cells available within the cluster. If such a cell exists, the controller reassigns some of its already allocated frequencies to the heavily loaded cells. On the other hand, if there are no such cells available, the controller checks to see if it is possible to offload some of the traffic to another network, such as WiFi or TV white space (TVWS). Finally, the SDN controller conveys its final decision and the corresponding instructions to the BS over the backhaul network. Depending on the specific type of granted resource (more frequencies allocated, offloading), the BS proceeds as shown in Fig. 3. If no new resources are granted, the BS frees up parts of the spectrum by removing the SUs that have been using the spectrum for the longest time.

Primary User Access: If a PU wants to use the spectrum, the BS first calculates the locations of the primary source and destination. Then it

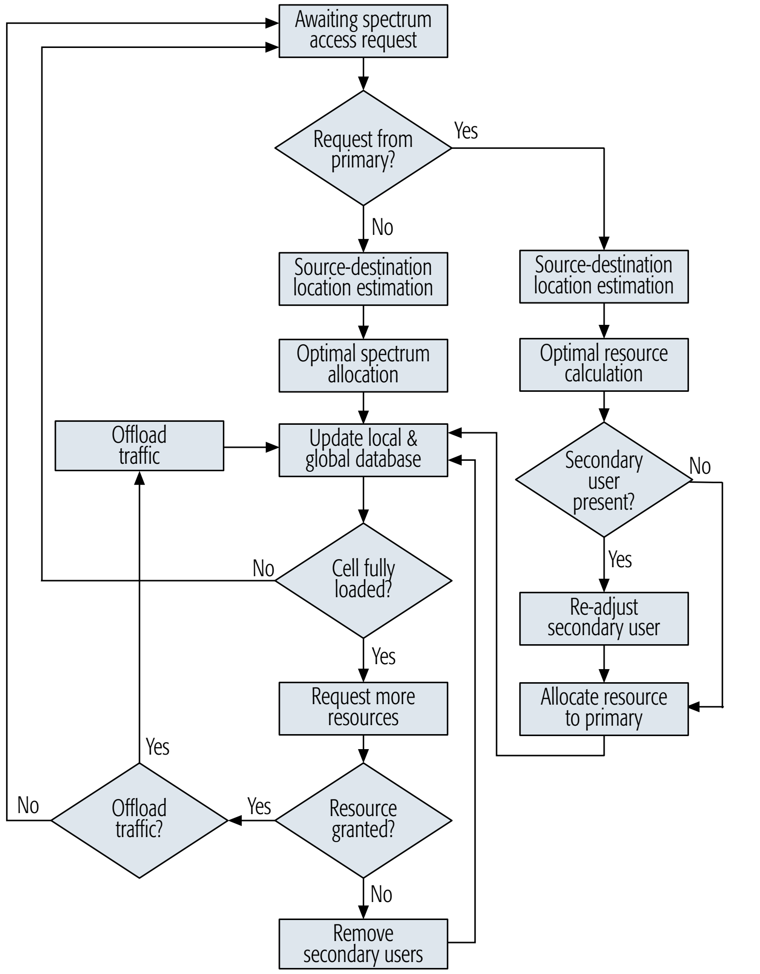

Figure 3. Resource management process at the BS. The proposed algorithm minimizes the dependence on spectrum sensing by relying on database assisted spectrum management.

calculates the optimal frequency for the associated primary transmissions. Since the PUs have the highest priority, the BS first calculates the optimal transmit frequency for the PU and then evaluates its impact on the SUs. If an SU is in the vicinity of the PU and uses the same frequency as the one calculated for the PU, the BS readjusts the transmission parameters for the SU (transmit power, frequency, etc.) and allocates the frequency deemed most appropriate to the PU. Following this, the LDB is updated, and the rest of the process follows a similar procedure as the one discussed in the previous subsection.

Before proceeding to our performance evaluation, it is worth mentioning that the transmission overhead involved in implementing the above-mentioned resource management algorithm remains either comparable to or lower than that involved in managing the spectrum through sensing. For example, as we show in the next section, with the aid of cooperative spectrum sensing, a large number of cooperating nodes has to achieve acceptable levels of successful primary detection. The high number of collaborating nodes significantly increases the overhead involved. By contrast, in the proposed design, a new SU entering the system only has to communicate with the BS to gain spectrum access. As far as updating the LDB is concerned, the BSs in current cellular systems already maintain this information about their users; hence, the cost of implementing and updating the LDB remains approximately the same. Additionally, the cost of updating the GDB can be minimized by updating only those values that have changed since the previous update.

Performance Evaluation

We commence our performance evaluation by highlighting the shortcomings of cooperative spectrum sensing algorithms. Figure 4 shows the probability of successful primary detection for various cooperative spectrum sensing schemes in a network consisting of a single PU and a varying number of SUs. We considered a scenario where the users were randomly distributed within a square shaped area of 1 km × 1 km, and the transmit power of the PU was fixed at 40 mW. Three commonly used spectrum sensing schemes, multi-peak cyclostationary detection, one-peak cyclostationary detection, and low-complexity energy detection, were compared [15]. It can be seen that the multi-peak technique outperforms the other two. However, even in the case of multi-peak detection, up to eight cooperating nodes are required to ensure a successful primary detection probability of 99 percent. This high number of cooperating nodes imposes both significant latency and increased transmission overhead since the nodes have to share their sensed data.

Figures 5 and 6 compare the performance of both sensing-based spectrum access and the proposed centralized spectrum access schemes within a single cell of radius 1 km. The performances of the two schemes are compared in terms of both the interference imposed on the PUs and the ratio of denied access requests (RADAR). Explicitly, RADAR is defined as the ratio of the denied secondary spectrum access requests to the total number of generated requests. For the communications, three primary users (PU1, PU2, and PU3) were assumed to be communicating at frequencies of 900 MHz, 1800 MHz, and 2.4 GHz. This setup may be considered synonymous to a HetNets scenario, where a macrocell, microcell, and/or picocell operate at different frequencies. With the aid of spectrum-sensing-based access, the SUs utilized the low-complexity energy detection technique for sensing the spectrum, where the spectrum was deemed to be idle if the received primary signal level was less than a predefined threshold. On the other hand, the proposed scheme utilized the geolocation information of the SUs and PUs for estimating the interference imposed on the PUs in different

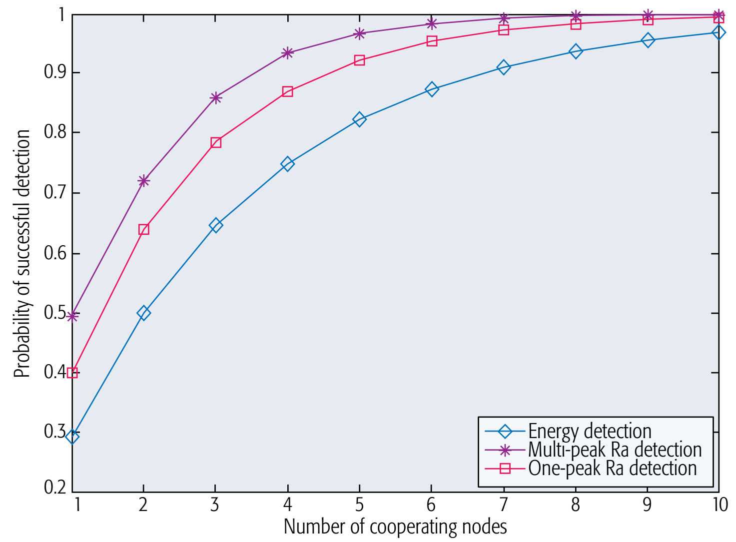

Figure 4. Probability of successful detection vs. the number of cooperating nodes. The figure highlights the fact that cooperative spectrum sensing requires a large number of cooperating nodes in order to achieve acceptable levels of performance.

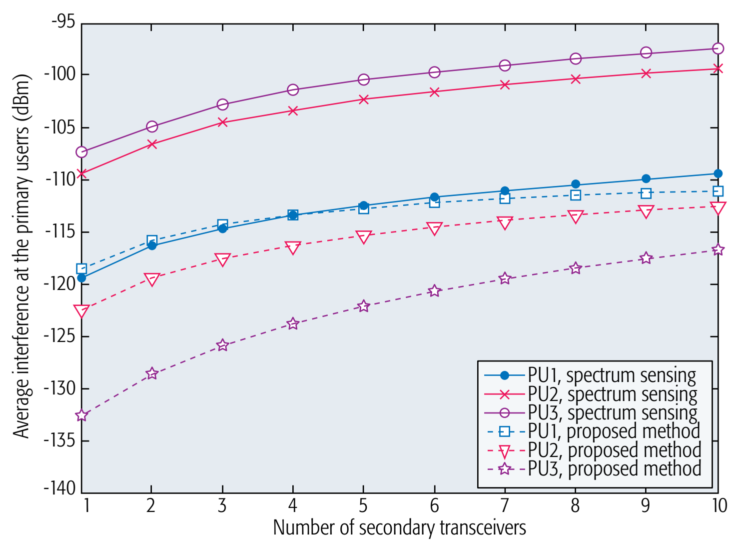

Figure 5. Average interference imposed on the PUs as a function of the number of secondary transceivers within the network. The figure shows that database-based spectrum management significantly outperforms spectrum-sensing-based resource management algorithms.

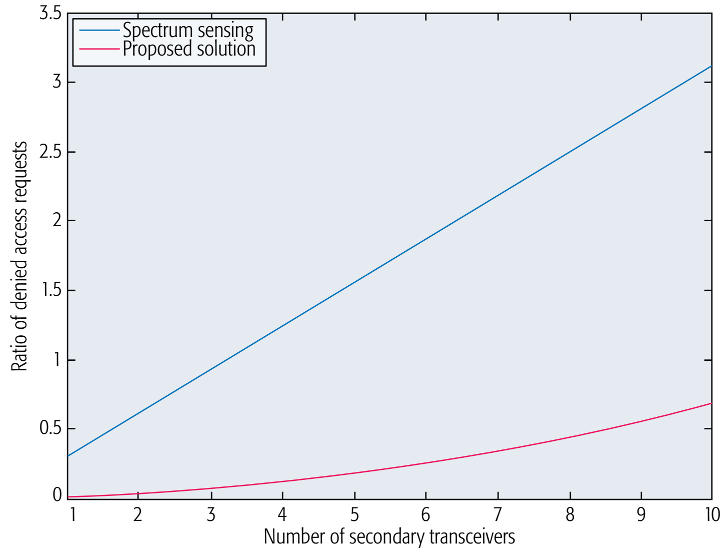

Figure 6. Average number of denied spectrum access requests as a function of the number of secondary transceivers within the network. The figure emphasizes the fact that the proposed resource management algorithm improves spectrum utilization by accommodating a larger number of secondary users.

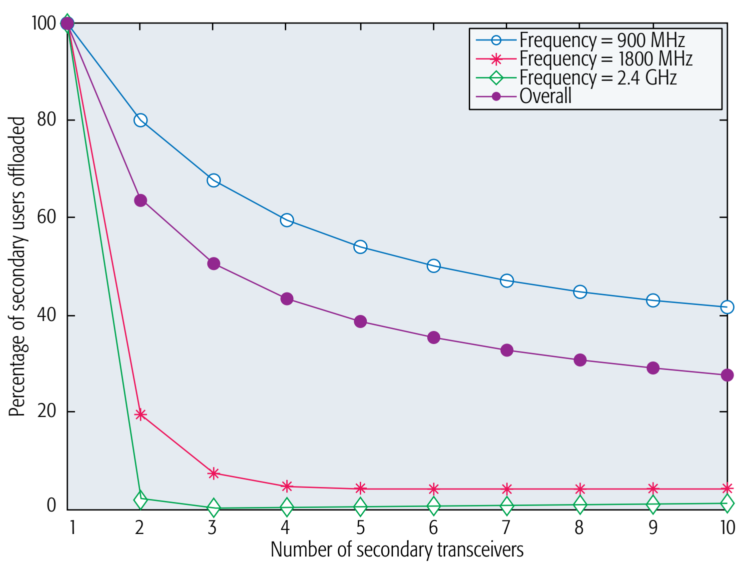

Figure 7. Percentage of secondary users offloaded from the primary network as a function of the number of secondary transceivers within the network.

frequency bands. The requested spectrum was allocated to an SU if the total sum of the interference at the PU imposed by the SU transmissions remained below a predefined threshold.

Observe from Fig. 5 that the proposed method significantly outperforms the baseline method in terms of the interference imposed on the PUs. While the average interference inflicted on PU1 remains roughly the same for both methods, it can be seen that the interference imposed on PU2 and PU3 is significantly reduced. The reason for this trend is that with the aid of the proposed method, the BS first attempts to assign the new SUs in the lower frequency bands so that they transmit at lower power. For this reason, PU1, communicating at the lowest frequency of 900 MHz, experiences the highest interference. If it is not possible to accommodate the user in the lowest frequency band, the BS searches for access opportunities at 1800 MHz and then at 2.4 GHz. In addition to minimizing interference, the proposed method also reduces the relative frequency of SUs being denied spectrum access. This can be seen from Fig. 6, which shows that the SUs are granted more frequent access to the spectrum with the aid of the proposed scheme.

Finally, Fig. 7 demonstrates the benefits of inter-network cooperation through offloading. The figure plots the percentage of SUs that are offloaded from each cellular frequency band onto a WiFi network. To obtain this result, we extended the scenario considered in Figs. 5 and 6 by allowing the BS to offload SUs onto a WiFi network operating at 5 GHz. The offloading scheme operates as follows: For each frequency band, the BS selects the SUs that cause the highest interference for the PUs. Then the selected SUs are offloaded onto the WiFi network only if their transmissions cause a tolerable amount of interference for the SUs that have already been offloaded. The process continues until the BS has completed checks of all the SUs within the cellular network. As shown in Fig. 7, the highest percentage of SUs is offloaded from the most crowded frequency band, that is, 900 MHz (Fig. 5). Overall, it can be seen that for supporting dense deployments (e.g., 10 secondary transceivers), up to 43 percent of the SUs are offloaded.

Conclusions

Current spectrum sharing solutions are mostly based on cognitive spectrum sensing, which is prone to inaccurate decisions due to the uncertainties imposed by the wireless medium. This drawback of current spectrum sharing mechanisms, coupled with their heavy reliance on device-level spectrum management, has discouraged network operators from embracing the idea of spectrum sharing. In this article, we envision a hierarchical architecture enabled by SDN that facilitates reliable and dynamic spectrum sharing in 5G cellular networks. The two key components of the proposed HSA framework are the macrocell BSs and the SDN controller. The BSs form the front-end of the network and interact directly with the user devices as well as make decisions related to spectrum access. The SDN controller forms the back-end of the network and enables the network operator to define fine-grained policies for network management. This task sharing between the BSs and the controller harmonizes network operation and alleviates the SDN controller's scalability concerns. The article also presents an efficient resource management algorithm conceived for future 5G networks. Finally, our simulation results quantify the performance gains achieved by the proposed framework.

References

[1] "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2014–2019 White Paper," http://tinyurl.com/mokcut3, 2015 (accessed 17 Aug.,-2015).
[2] X. Li, R. Zhang, and L. Hanzo, "Cooperative Load Balancing in Hybrid Visible Light Communications and WiFi," *IEEE Trans. Commun.*, vol. 63, no. 4, Apr. 2015, pp. 1319–29.

[3] T. S. Rappaport *et al.*, *Millimeter Wave Wireless Communications*, Prentice Hall, 2014.
[4] B. A. A. Nunes *et al.*, "A Survey of Software-Defined Networking: Past, Present, and Future of Programmable Networks," *IEEE Commun. Surveys & Tutorials.*, vol. 16, no. 3, 3rd qtr. 2014, pp. 1617–34.
[5] J. Wang, M. Ghosh, and K. Challapali, "Emerging Cognitive Radio Applications: A Survey," *IEEE Commun. Mag.*, vol. 49, no. 3, Mar. 2011, pp. 74–81.
[6] M. B. Pandian, M. L. Sichitiu, and H. Dai, "Optimal Resource Allocation in Random Access Cooperative Cognitive Radio Networks," *IEEE Trans. Mobile Comp.*, vol. 14, no. 6, June 2015, pp. 1245–58.
[7] F. Jin, R. Zhang, and L. Hanzo, "Resource Allocation under Delay-Guarantee Constraints for Heterogeneous Visible-Light and RF Femtocell," *IEEE Trans. Wireless Commun.*, vol. 14, no. 2, Feb. 2015, pp. 1020–34.
[8] X. Hong *et al.*, "Cognitive Radio in 5G: A Perspective on Energy-Spectral Efficiency Trade-off," *IEEE Commun. Mag.*, vol. 52, no. 7, July 2014, pp. 46–53.
[9] X. Jin *et al.*, "Softcell: Scalable and Flexible Cellular Core Network Architecture," *Proc. ACM CoNEXT '13*, New York, NY, 2013, pp. 163–174.
[10] A. Gudipati *et al.*, "Softran: Software Defined Radio Access Network," *Proc. 2nd ACM SIGCOMM Wkshp. Hot Topics in Software-Defined Networks*, New York, NY, 2013, pp. 25–30.
[11] X. Duan, A. M. Akhtar, and X. Wang, "Software-Defined Networking Based Resource Management: Data Offloading with Load Balancing in 5G Hetnet," *EURASIP J. Wireless Commun. Neworks.*, vol. 2015:181, no. 1, June 2015, pp. 1–13.
[12] X. Duan and X. Wang, "Authentication Handover and Privacy Protection in 5G Hetnets Using Software-Defined Networking," *IEEE Commun. Mag.*, vol. 53, no. 4, Apr. 2015, pp. 28–35.
[13] K. Pelechrinis *et al.*, "Cognitive Radio Networks: Realistic or Not," *ACM SIGCOMM Comp. Commun. Rev.*, vol. 43, no. 2, Apr. 2013, pp. 44–51.
[14] R. Saruthirathanaworakun and J. M. Peha, "Dynamic Primary-Secondary Spectrum Sharing with Cellular Systems," *Proc. 5th Int'l. Conf. Cognitive Radio Oriented Wireless Networks & Commun.*, Cannes, France, June 2010, pp. 1–6.
[15] J. Wang, A. Pervez, and M. Nekovee, "A Review and Verification of Detection Algorithms for DVB-T Signals," *Proc. 7th Intl. Commun. Sys. Networks and Digital Signal Processing Symp.*, Newcastle upon Tyne, U.K., July 2010, pp. 469–73.

Biographies

Auon Muhammad Akhtar (aakhtar8@uwo.ca) received his B.E. degree in electrical engineering, majoring in electronics, from Air University Islamabad, Pakistan, in 2007. He received his M.Sc. degree in mobile and personal communications from Kings College London in 2009. Later that year, he joined the Centre for Telecommunications Research, Kings College London, where he completed his Ph.D. in 2013. Currently, he is working as a postdoctoral fellow at Western University, London, Ontario, Canada. His research interests are in the areas of software-defined networking, resource allocation in 5G networks, cross-layered optimization, and green and cognitive cellular and ad hoc networks.

Xianbin Wang [S'98, M'99, SM'06] (xianbin.wang@uwo.ca) is a professor and Canada Research Chair at Western University, Canada. He received his Ph.D. degree in electrical and computer engineering from National University of Singapore in 2001. Prior to joining Western, he was with Communications Research Centre Canada (CRC) as a research scientist/senior research scientist between July 2002 and December 2007. From January 2001 to July 2002, he was a system designer at STMicroelectronics, where he was responsible for system design for DSL and Gigabit Ethernet chipsets. His current research interests include 5G networks, adaptive wireless systems, communications security, and locationing technologies. He has over 200 peer-reviewed journal and conference papers on various communication system design issues, in addition to 24 granted and pending patents and several standard contributions. He is an IEEE Distinguished Lecturer. He has received many awards and recognitions, including Canada Research Chair, CRC Presidents Excellence Award, Canadian Federal Government Public Service Award, Ontario Early Researcher Award, and three IEEE Best Paper Awards. He currently serves as an Editor/Associate Editor for *IEEE Wireless Communications Letters*, *IEEE Transactions on Vehicular Technology*, and *IEEE Transactions on Broadcasting*. He was also an Editor for *IEEE Transactions on Wireless Communications* between 2007 and 2011. He has been involved in a number of IEEE conferences, including GLOBECOM, ICC, WCNC, VTC, ICME, and CWIT, in different roles such as Symposium Chair, Tutorial Instructor, Track Chair, Session Chair, and TPC chair.

Lajos Hanzo [F] (lh@ecs.soton.ac.uk) FREng, FIET, Fellow of EURASIP, D.Sc. received his degree in electronics in 1976 and his doctorate in 1983. In 2009 he was awarded an honorary doctorate by the Technical University of Budapest, and in 2015 by the University of Edinburgh. During his 38-year career in telecommunications he has held various research and academic posts in Hungary, Germany, and the United Kingdom. Since 1986 he has been with the School of Electronics and Computer Science, University of Southampton, United Kingdom, where he holds the Chair in telecommunications. He has successfully supervised about 100 Ph.D. students, co-authored 20 Wiley/IEEE Press books on mobile radio communications totalling in excess of 10,000 pages, published 1500+ research entries at IEEE Xplore, acted both as TPC and General Chair of IEEE conferences, and presented keynote lectures, and has been awarded a number of distinctions. Currently he is directing a 60-strong academic research team working on a range of research projects in the field of wireless multimedia communications sponsored by industry, the Engineering and Physical Sciences Research Council (EPSRC) UK, the European Research Council's Advanced Fellow Grant, and the Royal Society's Wolfson Research Merit Award. He is an enthusiastic supporter of industrial and academic liaison, and he offers a range of industrial courses. He is also a Governor of the IEEE VTS. During 2008–2012 he was Editor-in-Chief of IEEE Press and also a Chaired Professor at Tsinghua University, Beijing. His research is funded by the European Research Council's Senior Research Fellow Grant. For further information on research in progress and associated publications please refer to http://www-mobile.ecs.soton.ac.uk. He has 22,000+ citations.

SDN Meets SDR in Self-Organizing Networks: Fitting the Pieces of Network Management

Carlos Ramirez-Perez and Victor Ramos

The authors propose an SON-based management framework that leverages on programmable control and data planes, and emphasizes the need for an open and extensible protocol interface that combines the main features of current protocols like OpenFlow, MIH, and XMPP.

Abstract

LTE self-organizing networks (SONs) automate several management mechanisms related to network planning, configuration, and optimization, which require network programmability at both the control plane and data plane. This requirement is a design target shared by software defined networking (SDN) and software defined radio (SDR) paradigms, which are intended to achieve open programmable and reconfigurable networks. Hence, the convergence of the programmable-control-plane nature of SDN with the programmable-data-plane nature of SDR into an SON-based architecture will complement and enhance their respective scopes. We propose an SON-based management framework that leverages programmable control and data planes, and emphasizes the need for an open and extensible protocol interface that combines the main features of current protocols such as OpenFlow, MIH, and XMPP.

Introduction

The proper management and orchestration of evolving networking technologies, rather than technology evolution by itself, is perhaps the key to achieving sophisticated network environments. The general goal of network management is to provide configuration capabilities flexible enough to allow rapid creation, deployment, and operation of control services, which in turn manipulate the behavior of the data plane as needed. Network management has been explored for at least 20 years. In 1995, Prof. M. Schwartz examined several management issues which he referred to as "*very interesting, complex research questions that will have to be addressed*" [1]. Most of the questions raised therein are now pragmatically covered by radio access networks such as 3GPP LTE and IEEE 802.11 WLANs. Nowadays, the LTE architecture provides the necessary stratum to deploy a sophisticated management framework, which leverages paradigms and tools pertaining to SDN, network function virtualization (NFV), and cloud computing. Additional network flexibility and adaptability may be gained by incorporating programmable network elements based on SDR technology. Recently, the work in [2] proposes an architectural design aimed at increasing the flexibility of network service deployment and facilitating operation and maintenance procedures. It allows SDR to provide device and function support for NFV, which in turn provides infrastructure support to SDN mechanisms.

An advanced network management framework should harness network programmability in order to maximize cohesion between the control plane and data plane, while minimizing coupling. The 1990s witnessed the emergence of architecture models that laid out design issues that were considered revolutionary due to the networking know-how of that time. SDN may be considered the latest version and innovation in the realm of programmable networks [3]. In fact, the SDN approach leads to a real and pragmatic change in network management. A sign of this is the proliferation of controllers and organizational initiatives such as the Open Networking Foundation (ONF) and the OpenDaylight platform.

The term SDN was coined in an *MIT Technology Review* article about the OpenFlow project at Stanford. However, the earliest use of the term *software-defined* was in the context of programmable radios or SDR. Intended to enhance service quality, SDR is about going from static and fixed radio hardware devices toward the total programmability of layer-1 (L1) functions. Bringing L1 processing into the software domain enables a high flexibility to customize the data plane behavior, and exposes new capabilities for management and automation that may be exploited at higher levels of abstraction. A remaining challenge for SDN is how to go from handling the data plane behavior by programming the control plane, toward directly programming the data plane. Here is where a promising application field for SDR technology arises. However, neither SDN nor SDR specify any compelling application. In this work, we describe the potential benefits of merging the programmable-data-plane nature of SDR with the programmable-control-plane nature of SDN in the context of LTE self-organizing networks. SON-based management frameworks represent a prevalent perspective for LTE-based systems [4], but deploying SON mechanisms requires nimble network architectures with a high degree of programmability. By highlighting some advanced features of LTE-based mobile networks, we review the architecture, use cases, and functionalities of SONs. Then we propose an SON-based multi-layered management framework composed of key technologies and protocols.

The authors are with Universidad Autónoma Metropolitana.

0163-6804/16/$25.00 © 2016 IEEE

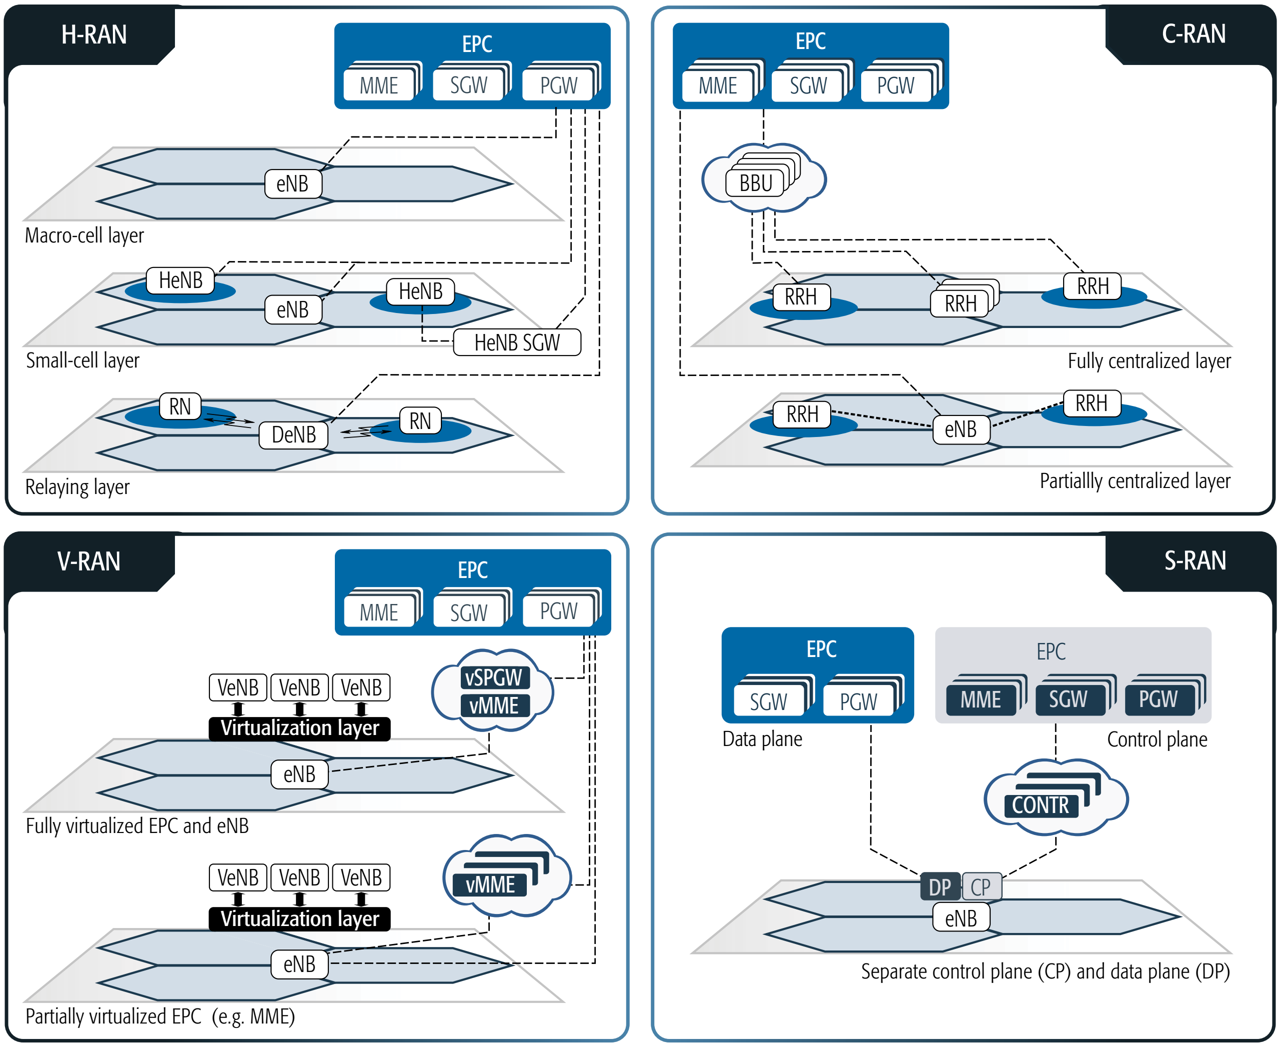

Figure 1. Mobile network architectures enabled by LTE. Heterogeneous radio access networks (H-RAN); centralized radio access networks. (C-RAN); virtual radio access networks (V-RAN); and software-defined radio access networks (S-RAN).

LTE-Based Mobile Network Architectures

Given the widespread acceptance of LTE, it will have to address innovative system-level design goals related to the adoption of novel business models, content delivery services, energy efficiency, and sustainability. In addressing these issues, future mobile network architectures (MNAs) will extend the LTE basic infrastructure by incorporating technologies such as HetNets, SDN, NFV, and cloud computing. Following the C-RAN notation, which refers to the cloud-based network architecture, we use the terms H-RAN, S-RAN, and V-RAN to refer to architectures arising from merging LTE with HetNets, SDN, and NFV, respectively. Figure 1 depicts the main structural features of the MNAs described below.

Heterogeneous Radio Access Networks (H-RAN)

The key LTE components are the evolved UMTS terrestrial radio access network (E-UTRAN) and the evolved packet core (EPC). In its most elemental form, the E-UTRAN consists of base stations called evolved NodeBs (eNBs). A special type of E-UTRAN base station known as a home eNodeB (HeNB) is a low-power node supporting the same functionalities as an eNB. HeNBs provide additional coverage and service capacity still in the licensed spectrum, but within residential or enterprise environments. The E-UTRAN may deploy a home eNB gateway (HeNB GW), which serves as a concentrator for the control plane to support a large number of HeNBs in a scalable manner. The E-UTRAN also may include relay nodes (RNs), another type of low-power base stations. An alternative method to improve throughput at hot spots is by physically deploying remote radio heads (RRHs) separated from the cell site cabinets and connected to base band units (BBUs). A small-cell is an umbrella term for low-powered radio access nodes, including femtocells, picocells, and microcells. The service capacity of macro-cell networks may be extended by adding small-cells, deploying HeNBs, RNs, or RRHs. Such node aggregation results in heterogeneous networks or HetNets, which involve the coordinated usage of one complex and dynamic layer of low-power network nodes, under the coverage of a structured layer of macro-cells. The work in [5] describes issues that must be considered during the design, implementation, and management of Hetnets. Also, it provides solid reason why the extreme heterogeneity of future mobile networks requires long standing models as well as rethinking the conventional wisdom.

Software-Defined Radio Access Networks (S-RAN)

The architecture in [6] applies the SDN principles in LTE by decoupling network control from S-GW and P-GW, which are mostly data-plane entities within the EPC. On one hand, with this modification significant control functionalities can be implmented in a logically centralized manner. On the other hand, it enables the deployment of S-GW and P-GW as virtual forwarding engines over commercial off-the-shelf (COTS), such as the Advanced Telecommunications Computing Architecture (ATCA), which is the leading carrier-grade COTS platform, according to recent marketing trend analysis. Hence, the deployment of such an SDN-based architecture is not closely tied to the use of data centers. From a functional perspective, the mobility management entity (MME) is a carrier-grade reliable multiprocessing serving platform, which makes fast decisions at the control plane. Based on the observation that the EPC is a fit-for-purpose closed system, the S-RAN approach is intended to provide maximum flexibility, openness, innovation, and programmability to future mobile carriers.

Centralized Radio Access Networks (C-RAN)

The modular design of the eNB integrates the BBU and the RRHs. The BBU takes care of baseband processing functions and control, while the RRH is responsible for radio signal transmission and reception, to and from UEs. In order to optimize BBU utilization, the strategy of the C-RAN architecture is to centralize different BBU resources into a virtual cluster commonly referred to as a BBU pool. Basically, the C-RAN architecture comprises the BBU pool, RRH mesh networks, and high speed transport networks connecting them. With the centralized BBU pool, resources can be managed and allocated on demand. Notice that these ideas are quite similar to the principles of cloud computing. Nonetheless, a generic public-cloud is not suitable to meet the specific performance requirements, QoS, and security in wireless systems. Thus, the concept of carrier-cloud has been defined, where ATCA-based technologies and platforms play a crucial role. In [7], the authors comprehensively review the fundamentals of the C-RAN architecture, and they point out that in order to support different wireless standards, hardware solutions based on multi-mode base stations may be deployed using SDR-based reconfigurable processing boards.

Virtual Radio Access Networks (V-RAN)

In LTE networks, the usual targets for virtualization are the eNBs. However, control entities such as the MME may also be implemented in virtual environments. Virtualization allows multiple virtual eNBs to share the common hardware resources, but also provides a better support for multi-tenancy and multi-provider coexistence over the same physical substrate. Virtualization of network components encompasses several levels of programmability and control. For example, the work in [8] explains that RAN virtualization may be performed at either the hardware level or the flow level, with the latter being better suited to support spectrum-sharing-based virtualization models, improving resource multiplexing. The authors also describe a base station virtualization scheme based on the network virtualization substrate (NVS). Aimed at OFDMA-based systems such as LTE and WiMAX, NVS allows multiple entities to share the same spectrum. Some authors in [8], as part of the European FP7 FLAVIA project, were designing programmable link processors, exposing interfaces for service customization and performance optimization.

LTE Self-Organizing Networks

Each MNA described above may be particularly suitable to achieve some network-wide goals. Therefore, future network architectures will have to establish several trade-offs to efficiently compose hybrid MNAs. This requires that difficulties should be overcome to administer network operations while ensuring network performance. The principles for LTE operations, administration, and maintenance (OAM) were created considering the ITU-T Recommendation M.3010 approved in 2000, and are described in the 3GPP Technical Specification 32-Series. The M.3010 recommendation defines fundamental concepts of telecommunications management network (TMN) architectures, and describes a reference model for partitioning the overall management functionality, structuring functions into logical layers.

Considering that network management should coordinate a wide span of organization and optimization mechanisms at different network levels and time scales, the automation of management processes is widely accepted as a proper approach to deal with such a complex task [4]. The Recommendation on Self-Organizing Network (SON), published by the NGMN Alliance in 2008, is the base for LTE automation of management processes. Basically, the idea behind SON is to achieve automation by adding certain capabilities or intelligence to network elements, which enable them to automatically execute management procedures to maintain network performance in an operational state, ideally the optimal. SON-based architectures may be centralized, distributed, or hybrid, depending on where the SON functionalities are implemented. In the centralized approach, the OAM system communicates with other entities such as eNBs to collect measurements and provide parameter settings. In a decentralized architecture, the SON functionality is entirely distributed on the network entities. The hybrid architecture is a compromise between the centralized and decentralized architectures.

SON use cases are categorized into functional areas along the key OAM areas of planning, deployment, optimization, and maintenance, which are commonly referred to as self-*configuration*, self-*optimization*, and self-*healing*. The use cases for self-configuration specified by 3GPP are dynamic configuration of the S1 and X2 interfaces, automatic neighbor relation (ANR) func-

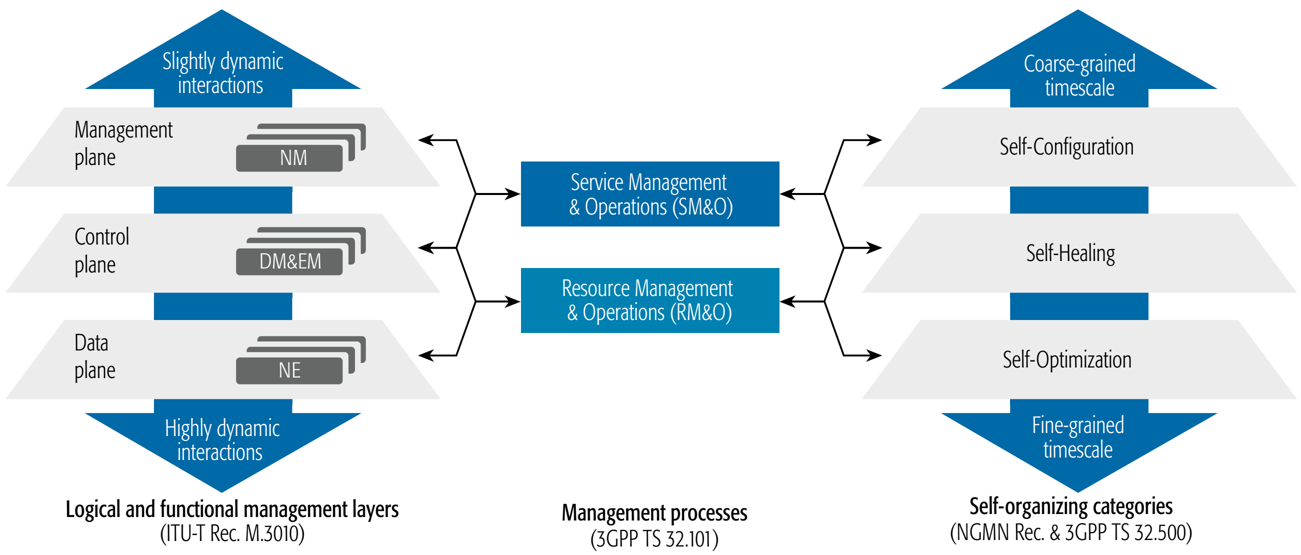

Figure 2. Functional layered networks.

tion, physical layer cell ID (PCI) selection, and transport network layer (TNL) address discovery. While self-configuration applies to the pre-operational state of network elements, self-optimization is necessary during the operational state of the network. The self-optimization use cases considered in the specifications are ANR function, mobility load balancing (MLB), and mobility robustness optimization (MRO). When the ANR function is employed as a self-optimization function, cells can be dynamically added/removed from a neighbor list. Cell degradation detection, diagnosis, and prediction are functionalities related to the self-healing category.

Figure 2 shows a layered structure attempting to capture the hierarchical composition of logical and functional management entities, and their relationship with SON categories. The main management entities include the network managers (NMs), domain managers (DMs), element managers (EMs), and network elements (NEs). We distribute each management entity into the management, control, and data planes to make explicit the level of relevance that these entities have on each plane. We also depict in Fig. 2 the service management and operations (SM&O) processes and the resource management and operations (RM&O) processes, which represent horizontal functional processes.

SDR Meets SDN on Self-Organizing Networks

Going from human-driven management to network self-management involves rethinking current networking know-how. Besides, the deployment and orchestration of SON functionalities require a comprehensive systematic approach. Here, we describe the integration of SDN and SDR for self-management under the LTE scope. SDR and SDN paradigms clearly formulate a common design objective, which is going from closed-nature hardware-based to open programmable and reconfigurable systems. Thus, their convergence into an LTE SON-based architecture will expose interfaces to retrieve, exchange, and update system configurations, as well as to supply and query context information, all of this in an automated way. Figure 3 depicts the multi-layered management framework we propose, whose main components are described below.

Software Appliance Layer (SAL)

This layer is a software appliance factory, in charge of creating, deploying, and orchestrating management functions, aimed at achieving system-level objectives related to high-level networking paradigms, such as user-centric networking (UCN) or information-centric networking (ICN). The SAL is equipped with tools for abstract composition and chaining of self-configuration, self-healing, and self-optimization services. Conceptual abstractions related to each SON category may be defined using a high-level declarative specification language, which must be later translated into imperative programming language, enabling low-level control of individual network elements. Nonetheless, network programming languages are needed to simplify this task. In [9], the *Frenetic* project is presented, which is designed to ease the specification, composition, and updating of network management policies in a consistent way.

The SAL must be aware of what, when, and where to place SON software appliances. Distributing appliances to manipulate traffic as needed, alludes to the use of middleboxes that, according to RFC 3234, are any intermediary device, in the data path between end-to-end hosts, performing functions apart from those standard at routers. Thus, SON appliances created at the SAL may be regarded as software middleboxes. By dynamically placing these software middleboxes, the forwarding path of specific traffic flows can be established and updated. SDN promises valuable opportunities for middlebox policy enforcement using logically centralized management. The system design challenges addressed in [10] not only demonstrate the feasibility of using SDN

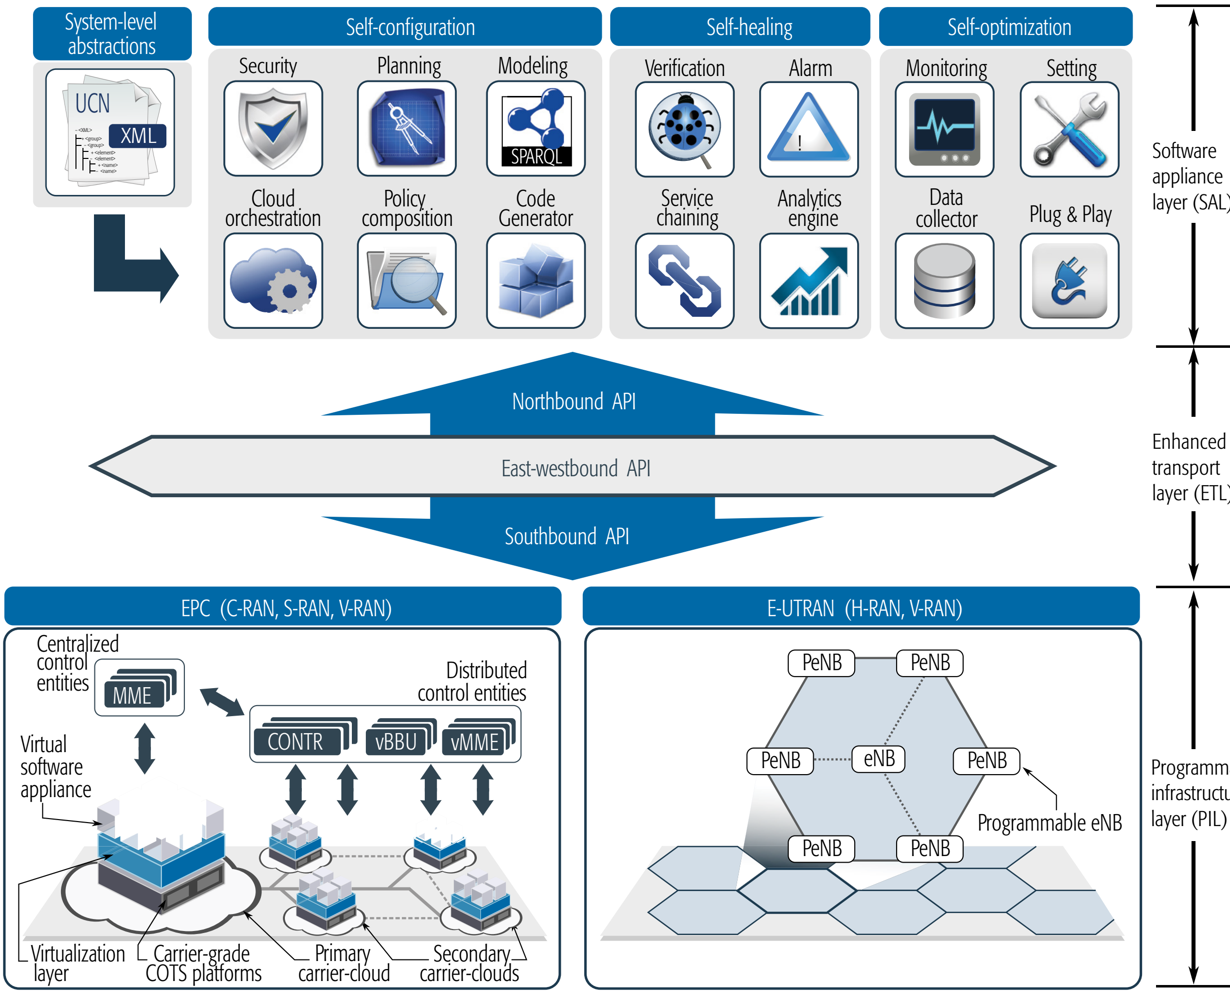

Figure 3. SON-oriented architecture.

to simplify middlebox traffic steering, but also represent a significant step toward broadening the scope of SDN, beyond conventional L3 tasks.

Enhanced Transport Layer (ETL)

As mentioned in [2], to achieve the goals envisaged for 5G systems, the core network will have to reach unprecedented levels of flexibility and intelligence; eventually, this requirement will expand toward the radio access networks. Thus, it is necessary to account with a transport service robust enough to allow quick communication between intelligent entities located at the core and at the radio network. Next we describe the main building blocks of the transport layer in our proposal.

Northbound Interface: According to the 3GPP management reference model, the interface-northbound (Itf-N) is a management interface between network managers (NMs) and element managers (EMs). We depict the network manager, element manager, and network element entities in Fig. 2. The Itf-N acts as the front-end hiding the mapping of generic information models used by NMs and vendor-specific parameter representation used by NEs. In SDN, there is an analogous concept for the northbound interface [11]. SDN northbound applications manage and control the network by programming SDN controllers and requesting services from them. Since there may be several applications featuring different purposes, ranging from virtual network provisioning to granular traffic steering mechanisms, the development of information models to composed orchestration systems is needed. To date, SDN northbound APIs remain completely closed or vendor-specific, because standardization initiatives remain stagnant. Notice that SDN-based and SON-based architectures are trying to solve a similar set of problems looking for solutions to orchestrate network management systems. Combining the structured self-organizing functions described earlier with the flexibility provided by SDN developments will greatly benefit the automation of management tasks.

Eastbound-Westbound Interface: An additional interface among network managers and domain managers is called Itf-P2P, which is also considered in the 3GPP management reference model. While Itf-N may be considered as an interface for management procedures taking a vertical (north-to-south) path through the network, it could be considered that the Itf-P2P horizontally exchanges management information. Thus, Itf-P2P may be regarded as an eastbound-westbound interface. The IEEE 802.21 standard, commonly referred to as media independent handover (MIH), may be particularly suitable to be employed as an eastbound-westbound interface. The primary role of MIH is to facilitate handover procedures by making available information for network discovery and selection. The MIH function (MIHF) is a logical entity that communicates with other layers using service access points (SAPs). The media independent services (MIS) provide event (MIES), command (MICS), and information services (MIIS). Prior to providing the MIS from one MIHF entity to another, the entities need to be properly configured using the MIH capability discovery, registration, and event subscription functions. The discovery procedure is used by an MIH user to find local or remote MIS, then the registration procedure makes it possible to request access to specific services. The event subscription mechanism allows the MIH users to subscribe to events originating from a local or remote MIHF. Access networks may host one or more points of service (PoS), which support remote MIES and MICS. PoS may be paired with points of attachment (PoAs) or placed deeper beyond the access network. To provide information services, MIIS servers may be located as independent entities on the network side. SDN and MIH frameworks may be considered in an integral manner, since the former focuses on data path control and resource management, and the latter focuses on mobility and radio resource management. In fact, as described in the ONF Wireless & Mobile Working Group documentation, OpenFlow may leverage on MICS to effectively perform flow reassignments. Also, recent research activities focus on analyzing how to enhance SDN controllers with MIH services [12].

Southbound Interface: Naturally, a complementary interface for the northbound and westbound-eastbound interfaces is a southbound interface. The 3GPP management reference model identifies this interface from EMs to NEs, and only indicates the possible use of protocols such as CORBA IIOP, NETCONF, SNMP, and SOAP. In the context of SDN, it is well known that the OpenFlow protocol is used as the main southbound protocol [13]. With OpenFlow, L3 forwarding rules can be dynamically controlled at the data plane, and therefore it does not provide management functions. In this regard, the OpenFlow Management and Configuration Protocol (OF-CONFIG) provides standard methods that cen be used to set parameters for communication between OpenFlow controllers and switches. The OF-CONFIG protocol enables the remote configuration of datapaths using as transport the NETCONF protocol. Due to the extensibility of NETCONF, which uses XML-based data encoding, it could be possible to develop an integral southbound interface.

Implementation and Deployment: The three interfaces described above are either being defined or maturing, thus the implementation of an integral solution is still an open issue. In an SON-based management system, it would be natural to use real-time messaging to coordinate management tasks between distributed network elements. Hence, the Extensible Messaging and Presence Protocol (XMPP) may be used as the basis for the development of an enhanced and extensible transport interface for management. According to RFC6120, XMPP is an application profile of the Extensible Markup Language (XML) enabling near-real-time exchange of messages, called stanzas, between network entities. There are three types of stanzas: message, presence, and info/query (IQ). The extensibility and scalability of XMPP are perhaps the main features making it suitable enough to develop a robust and integral management platform. On one hand, the base protocol is extended by means of an open and collaborative process that, if approved, may be considered as an XMPP Extension Protocol (XEP). On the other hand, the push-based communication model used in XMPP solves scalability problems associated with traditional pull-based approaches. Since the transport interface should be properly secured, the XMPP usage may be complemented with an authentication and authorization protocol like *Diameter*. According to RFC 6733, the Diameter base protocol is intended to provide an authentication, authorization, and accounting (AAA) framework for applications such as network access or IP mobility. In LTE, the Diameter protocol is employed for subscription and authentication data transfers between MMEs and home subscriber service (HSS) entities. With Diameter, specific commands can be defined with their own attribute value pairs (AVPs), which are the basis for the extensibility of the Diameter protocol. Together, XMPP and Diameter may compose an event-driven management system, featuring high levels of scalability, flexibility, and extensibility. Table 1 summarizes some features of MIH, OpenFLow, DIAMETER, and XMPP.

To shed some light on the deployment of the ETL, we use the scenario shown in Fig. 4, wherein several programmable eNBs (PeNB) are systematically deployed within the coverage area of an eNB to increase its service capacity. Consider the PeNBs as distributed management agents implementing the protocol stack shown on the left side of Fig. 4. Within the protocol suite, upper protocols pertaining to SON, ANDS, and SDN management applications are enhanced with the XMPP and DIAMETER frameworks, and also might be supplemented by MIH services. Both XMPP and Diameter require reliable transport protocols (L4), such as TCP or SCTP. At the middle section of the stack, the MIH reference model is coupled between management applications and the LTE radio resource control (RRC) layer, which is the interface for the MAC/PHY protocols. Due to the programmability of the PeNBs, they support the software implementation of both the 802.11-based MAC/PHY and the LTE-based protocols, which might be managed by RRC layer. Later we provide further details about the technological design of PeNBs.

To exemplify a use case for such protocol composition, assume a hypothetical scenario where the PeNB labeled as *penb04* infers an imminent congestion condition based on its own traffic measurements. The goal is to notify this situation to the eNB to somehow obtain assistance. Besides, notice that the nearest neighbors, namely *penb03* and *penb05*, are in sleep mode due to energy policies. Since all the network ele-

Criteria	Protocol			
	MIH	OpenFlow	Diameter	XMPP
Main goal	Provides a framework for network information management	Remote configuration of L3 forwarding related tasks	Secure authentication, authorization, and accounting	Near-real-time exchange of XML-based structured data
Standardization (entity/year)	IEEE/2008	ONF/2008	IETF RFC 6733/2012	XSF/2004 (RFC 6120)
Architecture	Peer-to-peer	Client-server	Peer-to-peer	Decentralized client-server P2P
Network entities	MIH users MIH function (MIHF)	OF controllers OF-capable switches	Clients Servers Agents	Clients Servers Bots Gateways
Neighboring list	Information element container list of networks	–	Peer and peer-routing tables	XEP-0083: nested roster groups XEP-0144: roster item exchange
Neighbor discovery	DNS-based (RFC-5679)	–	Diameter peer discovery SRVLOC and DNS	XEP-0030: service discovery XEP-0115: capabilities advertisement
Fault-tolerance	L2+ media-dependent MIH protocol ACK service	OpenFlow channel: main connection and multiple auxiliary connections	Hop-to-hop and end-to-end failover and failback procedures via alternate paths	XEP-0198: stream management
Element manager (EM)	MIH point of service (PoS)	OF configuration points (CPs)	Diameter network access server (RFC 7155)	–
Data modeling	Type-length-value (TLV) or RDF	–	Attribute-value-pairs (AVPs)	–
Encoding language	RDF/XML (SPARQL supported)	–	XML-based structured syntax	XML-based structured syntax
Subscription support	Not defined, but required.	–	–	XEP-0163: personal eventing protocol (PEP) XEP-0060: publish-subscribe
Extensibility support	–	–	Extensible Authentication Protocol encapsulation	XEP-0001: XMPP extension protocols

Table 1. Comparison of several protocols considered for management of SON.

ments in our example are XMPP entities, they need a jabber identifier (JID), which consists of `localpart@domainpart/resourcepart`. The domain part, which typically maps to a fully qualified domain name (FQDN), is the only required component. In Fig. 4 we show the FQDN for the eNB labeled as *enb0A*, running an XMPP server. The LTE identifiers, along with several domain name system (DNS) procedures, are described in the 3GPP TS 29.303 document. The eNB FQDN identifier represents the domain portion of the JIDs for the PeNBs, while the local part corresponds with their respective labels. Upon detection of the congestion condition, a publishing event is triggered, in which the XMPP client hosted in *penb04* automatically constructs an IQ stanza such as the one shown in the right part of Fig. 4. Notice that the `MIH_Link_Going_Down` indication is embedded in the IQ stanza. This example illustrates how MIH primitives might be encapsulated into XMPP IQ stanzas. Further details of XMPP stanzas and MIH primitives are beyond the scope of this article. However, it is worth remarking that the implementation issues of the MIH mechanisms are out of the scope of the IEEE 802.21 standard. Therefore, we consider that XMPP and several of its extensions, such as the XMPP PUB-SUB (XEP-0060), the Personal Eventing Protocol (XEP-0163), and the Presence Information and Service Discovery (XEP-0030), may jointly provide standardized implementation solutions.

In our example, *penb03*, *penb05*, and *enb0A* are subscribed, and therefore authorized to learn about the information published by *penb04*. Thus, when *penb04* publishes its congestion event, the publish-subscribe service pushes an event notification to all these entities. After being notified, *penb03* and *penb05* leave sleep mode to enter into an expectant state looking for instructions coming from the eNB. Another possibility is when *enb04* publishes its congestion event and it also announces information about MAC/PHY transmission parameters, which may be used by *penb03* and *penb05* to speed up traffic steering mechanisms, like handover. Meanwhile, regarding the eNB, it may take different actions on the notification. For example, further information queries may reveal additional context of the operational state of *penb04,* making it possible to identify which optimization, healing, or configuration procedures are required, as well as whether these procedures need to be deployed. Of course, this may involve the analysis of considerable amounts of data and making decisions based on such analysis. In the example above, we can observe that management services need to be notified about relevant changes in the network state in a timely manner. It is at this point where big data analytics and cognitive networks will be paramount.

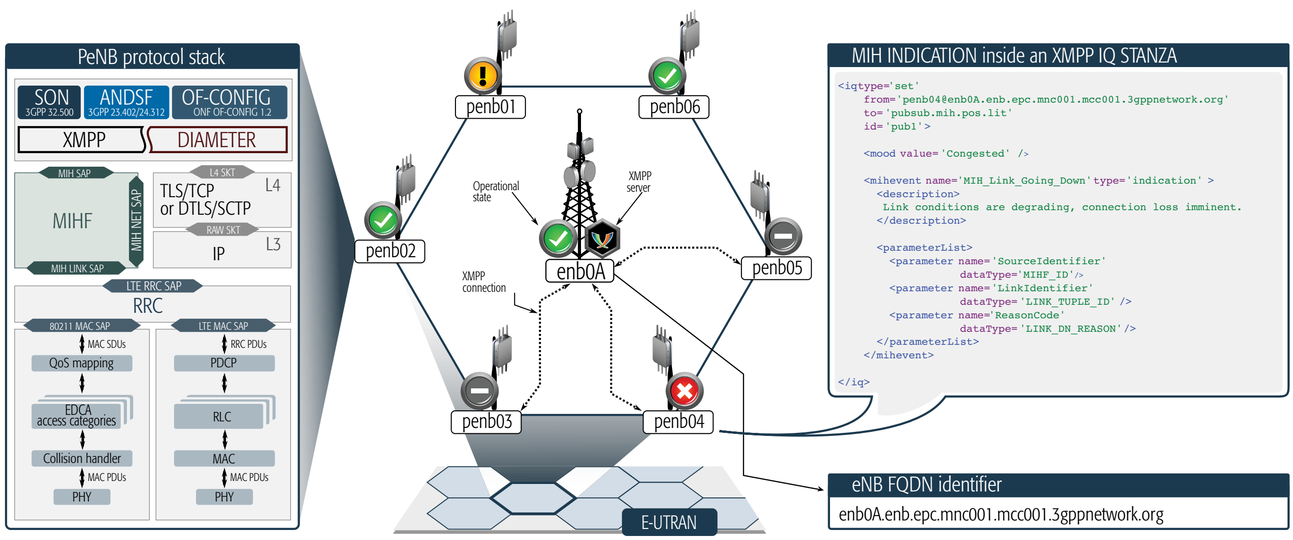

Figure 4. An instance of PeNBs deployed within the coverage area of an eNB.

Programmable Infrastructure Layer (PIL)

As mentioned, the SAL is a fabric of SON-oriented software appliances that are dynamically instantiated and deployed. Also, the ETL provides the means for reporting network events, possibly at different time scales, and exchanging context information. Now we discuss the programmable infrastructure layer (PIL), which extends the LTE basic infrastructure. The PIL is the last piece to enhance the responsiveness of the management framework with network adaptiveness. The bottom-left side of Fig. 3 depicts the core network merging the C-RAN, S-RAN, and V-RAN architectures. The EPC is represented as a hierarchical network of primary and secondary carrier-clouds. While the former are centralized data centers featuring highly-available storage and communication systems, the latter allow the migration of virtual entities over distributed high performance processing units. Secondary carrier-clouds provide support for dense Cloud-RAN with a large number of RRHs, but they also make it possible to establish and update forwarding paths applying SDN-based control functionalities. Introducing NFV and SDN into the core networks provides the necessary elasticity to dynamically place SON-based software appliances where needed, making it possible to improve load balancing and latency issues. Additional network flexibility and adaptability may be gained by strategically deploying programmable network elements built on the principles underlying the H-RAN, S-RAN, and V-RAN architectures. The bottom-right side of Fig. 3 shows the E-UTRAN elements, which include conventional eNBs and programmable eNBs (PeNBs).

Motivation and Deployment: Before delving into the design details of PeNBs, we describe the potential benefits of their deployment by considering the interaction of heterogeneous radio technologies such as WiFi, small-cells, and relay nodes (RNs), which are key building blocks in the architectural design of 5G networks. First, several research activities such as the licensed-assisted access (LAA), LTE in unlicensed spectrum (U-LTE), and LTE Wi-Fi Link Aggregation (LWA), are intended to improve the LTE and Wi-Fi interworking, guaranteeing their fair coexistence. However, with these technologies sharing the spectrum in extremely dense and heterogeneous network deployments, their coupled interaction may not be enough. Second, as explained in [14], RNs can effectively improve service coverage and system throughput, especially when multiple RNs are deployed. However, the performance of relay transmissions is greatly affected by the collaborative strategy employed to select pairing and relaying schemes. This selection process is complex, since it involves the decision of when, how, and with whom to collaborate, taking into account the variable channel conditions. Notice that both cases are addressed from the perspective of interaction. Nonetheless, we analyze them from another perspective, which goes from interaction toward integration. As mentioned in [2], heterogeneous networks take full advantage of the complementary characteristics of different network tiers of access technology, but various service requirements lead to a series of problems. The idea behind the PeNBs is to have network elements capable of changing their behavior by toggling between several radio technologies. This will help overcome many challenges posed by HetNets to NFV and SDN.

Design and Implementation: Now we provide some insights into the internal design of PeNBs. As Fig. 5 illustrates, each PeNB is provided with an SDR-based multiprocessor platform, which contains specialized processing units that execute L1, L2, and L3 processes. Then a virtualization layer abstracts the hardware resources, allowing for multiple virtual eNBs (veNBs), which can be provisioned, controlled, and inspected as needed. Any veNB must support the same set of functions implemented by conventional eNBs, which moves forwarding and packet processing toward the BBU, and radio signal processing to the RRH. Notice that if eNBs become virtual entities, L1, L2, and L3 become threads executed by their corresponding processing units, presumably in an optimized way. Currently, WiFi APs

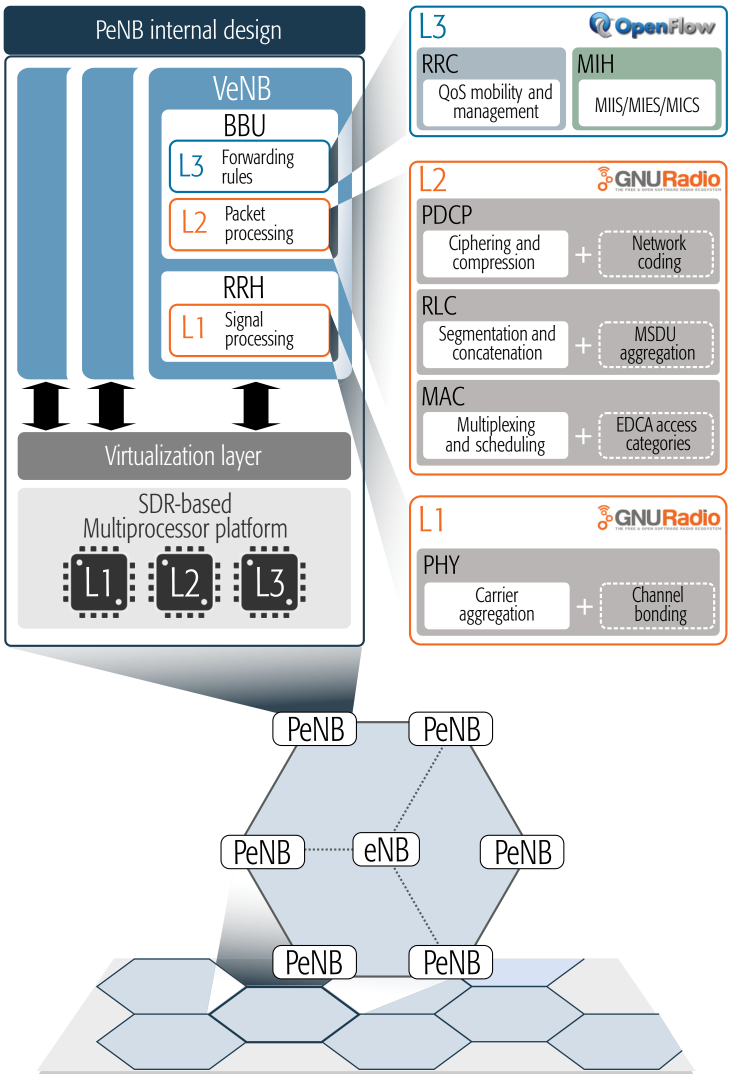

Figure 5. Internal design of PeNBs.

and LTE eNBs implement L2/L1 stacks as firmware in closed-nature specialized hardware platforms. Hence, they might be regarded as control entities featuring constrained capabilities to set up operational parameters on the fly. However, since L1 and L2 procedures are clearly delineated and intrinsically related, it is feasible to implement various wireless L2/L1 protocol stacks over general-purpose operating systems, enabling higher levels of reconfigurability and adaptability [15]. In Fig. 5, we indicate some analogous L2/L1 functions performed by WiFi APs and LTE eNBs. For example, while the LTE Packet Data Convergence Protocol (PDCP) sublayer supports ciphering and header compression, network coding mechanisms have been widely studied in WiFi-based mesh networks to improve throughput. The LTE RLC protocol performs segmentation and concatenation of PDCP packets. Similarly, WiFi supports mechanisms such as the MAC service data unit aggregation (A-MSDU) to send multiple data frames into a single larger one. Besides, QoS control at the LTE MAC layer is supported by scheduling and prioritizing mechanisms, which dynamically allocate radio resources to service data flows associated with a QoS class identifier (QCI). Analogous QoS features are supported in the 802.11e-based MAC layer, where the enhanced distributed channel access (EDCA), a contention-based channel access mechanism, delivers prioritized traffic flows using four different access categories (ACs). Finally, in addressing the challenges to realize the programmable L2/L1 stacks, current software development platforms such as GNU Radio may be helpful. Even if such a platform mainly provides signal processing blocks, it may be adapted to support software implementations of MAC-layer processing using packet-oriented semantics. Equally important to software development platforms are the programmable hardware platforms such as USRPs, developed by Ettus Research (www.ettus.com), which combine the hardware programmability of FPGAs with the software programmability of processors.

To complement the L2/L1 with upper-level functionalities, at the top of the protocol stack depicted in Fig. 5, we consider the integration of the MIH services with the LTE radio resource control (RRC) layer, which is the main air interface protocol for the control plane signaling messages. Also, at this level of abstraction the OpenFlow protocol, along with its management protocol OF-CONFIG, may compose a robust L3 interface for enhanced mobility support. At the core of the OpenFlow processing pipeline are the match/action procedures. Under the scope of OpenFlow, these operations are constrained to a particular paradigm. However, dynamically performing a wider range of traffic manipulation functions, such as traffic monitoring, marking, and shaping, requires higher levels of customization, potentially in software. Hence, programmable data planes provide the potential to expand the current capabilities of OpenFlow switches, making it possible to program custom match/action primitives.

Concluding Remarks

The ever increasing demand for mobile and wireless applications and services calls for a robust, flexible, and dynamic management platform. This requires control and data planes featuring high levels of programmability. Even if the mass market of current mobile communications relies on closed-nature *hardware-defined* systems, the adoption of NFV, SDN, SDR, and HetNet technologies will narrow the bridge for future systems to be self-organized. One of the main goals of this work is to promote the symbiosis between SDR and SDN within the context of LTE self-organized networks. This requires the convergence of computing, communications, and networking research into one domain, which hopefully will result in innovative architectural models that will leverage the features of protocols such as OpenFlow, MIH, XMPP, and Diameter. Considering the relevance of management architectures, there is a compelling need to establish evaluation frameworks or benchmarks, intended to measure the effectiveness of emerging architectures. This would make it possible to expose the effective success or failure of the methodologies employed for the definition, instantiation, deployment, and verification of specific SON functionalities or general system-level policies.

References

[1] M. Schwartz, "Network Management and Control Issues in Multimedia Wireless Networks," *IEEE Pers. Commun.*, vol. 2, no. 3, June 1995, pp. 8–16.

[2] S. Sun *et al.*, "Integrating Network Function Virtualization with SDR and SDN for 4G/5G Networks," *IEEE Network*, vol. 29, no. 3, May/June 2015, pp. 54–59.

[3] N. Feamster, J. Rexford, and E. Zegura, "The Road to SDN: An Intellectual History of Programmable Networks," *SIGCOMM Computer Communications Review*, vol. 44, no. 2, Apr. 2014, pp. 87–98.

[4] H. Hu *et al*, "Self-Configuration and Self-Optimization for LTE Networks," *IEEE Commun. Mag.*, vol. 48, no. 2, Feb. 2010, pp. 94–100.

[5] J. Andrews, "Seven Ways that HetNets are a Cellular Paradigm Shift," *IEEE Commun. Mag.*, vol. 51, no. 3, Mar. 2013, pp. 136–44.

[6] K. Pentikousis, Y. Wang, and W. Hu, "Mobileflow: Toward Software-Defined Mobile Networks," *IEEE Commun. Mag.*, vol. 51, no. 7, July 2013, pp. 44–53.

[7] A. Checko *et al.*, "Cloud RAN for Mobile Networks — A Technology Overview," *IEEE Commun. Surveys Tutorials*, vol. 17, no. 1, Sept. 2014, pp. 405–26.

[8] X. Costa-Perez *et al.*, "Radio Access Network Virtualization for Future Mobile Carrier Networks," *IEEE Commun. Mag.*, vol. 51, no. 7, July 2013, pp. 27–35.

[9] N. Foster *et al.*, "Languages for Software-Defined Networks," *IEEE Commun. Mag.*, vol. 51, no. 2, Feb. 2013, pp. 128–34.

[10] Z. A. Qazi *et al.*, "SIMPLE-fying Middlebox Policy Enforcement Using SDN," *SIGCOMM Comp. Commun. Rev.*, vol. 43, no. 4, Aug. 2013, pp. 27–38.

[11] H. Kim and N. Feamster, "Improving Network Management with Software Defined Networking," *IEEE Commun. Mag.*, vol. 51, no. 2, Feb. 2013, pp. 114–19.

[12] C. Guimaraes *et al.*, "Empowering Software Defined Wireless Networks through Media Independent Handover Management," *IEEE Global Commun. Conf. (GLOBECOM)*, Dec. 2013, pp. 2204–09.

[13] N. McKeown *et al.*, "OpenFlow: Enabling Innovation in Campus Networks," *SIGCOMM Computer Commun. Rev.*, vol. 38, no. 2, Mar. 2008, pp. 69–74.

[14] Y. Yang *et al.*, "Relay technologies for WiMax and LTE-Advanced Mobile Systems," *IEEE Commun. Mag.*, vol. 47, no. 10, Oct. 2009, pp. 100–05.

[15] K. Tan *et al.*, "Sora: High-Performance Software Radio Using General-purpose Multi-Core Processors," *Commun. ACM*, vol. 54, no. 1, Jan. 2011, pp. 99–107.

Biographies

Carlos Ramirez-Perez (carap@xanum.uam.mx) received the B.Eng. degree in electronics and the M.Sc. degree in information technologies, both from the Universidad Autónoma Metropolitana (UAM), Mexico City, in 2009 and 2011, respectively. He is currently working toward his Ph.D. degree in sciences and information technologies at UAM. His research interests are big data analytics for cognitive autonomic networks and software-defined networking.

Victor Ramos (vicman@xanum.uam.mx) is a researcher with the Department of Electrical Engineering at the Universidad Autónoma Metropolitana (UAM) Iztapalapa at Mexico City. In 2004 he received the Ph.D. degree in computer science, and in 2000 the DEA degree in networking and distributed systems, both from the Université de Nice-Sophia Antipolis, France. In 1995 he received the B.Eng. degree in electronics from UAM, Mexico City. From 2007 to 2011 he was the Head of the Department of Electrical Engineering of UAM, and from 2006 to 2007 the Head of the Networking and Telecommunications Research Team. His research interests are in performance evaluation of communications protocols, vehicular ad hoc networks, RFID tag identification, software-defined networking, and network coding for wireless mesh networks.

Software Defined Radio – 20 Years Later: Part 2

Joseph Mitola

Preston Marshall

Kwang-Cheng Chen

Markus Mueck

Zoran Zvonar

Following the Part 1 that appeared in the September 2015 issue of *IEEE Communications Magazine*, this Part 2 of the Radio Communications Series features several papers illustrating key trends in Software Defined Radio.

The last two decades have brought multiple advances in enabling technologies for SDR including radio frequency, analog to digital conversion, and digital signal processing. The effort has notably split along two distinctive paths. The first path has focused on state-of-the-art commercial software-defined platforms whose software and hardware capabilities classify them as platforms to enable rapid prototyping and advance experimental research in wireless networking. The second path was focused on developing volume oriented products, primarily in consumer spaces such as cellular systems.

The first article in this issue, "Addressing Next-Generation Wireless Challenges with Commercial Software-Defined Radio Platforms", provides a comprehensive overview of commercially available platforms and provides key comparisons that can help researchers identify the right platform for the testbed of interest. Continuing in the same direction, the second article, "Software Defined Radio: Revolutionizing Communication System Design and Prototyping" by Wyglinski *et al.*, further revisits concepts behind the most popular platforms such as Universal Software Radio Peripheral (USRP), the GNU Radio project, and selective technical computing software.

The next group of papers in this issue offers forward looking concepts for SDR. The third article in the issue by the University of Arizona team expands on the concept of metacognition and the opportunities to cognitive radio. Following on the tutorial exposition of metacognition, the authors explore performance improvements using a meta cognitive engine.

Sagar *et al.* also look forward to the next 20 years of SDR with Software-Defined Access (SDA) for heterogeneous wireless networks. SDN has been illusive commercially for many reasons, including the lack of a control architecture that is independent of the specific radio access technology (RAT) such as GSM, WiFi, or LTE. Particularly in the early stages of 5G, service providers confronted with a mix of deployed hardware will benefit from the architecture thinking, framework, and demonstration of this final article in the series.

We would like to thank all of our colleagues who significantly contributed to this two-part series, including the authors, reviewers, and *IEEE Communications Magazine* editorial team. It was a privilege bringing this special issue to *IEEE Communications Magazine*, and we look forward to the next productive 20 years of technology development.

Biographies

Joseph Mitola III (joseph.mitola@federatedwireless.com) is President and CEO of Hackproof Technologies Inc. Previously he held positions of technology leadership with Stevens Institute of Technology, U.S. DoD, The MITRE Corporation, ITT, Harris, E-Systems, and the Executive Office of the President of the United States. He received a B.S. E.E. from Northeastern University; an M.S.E. from The Johns Hopkins University; and a Licentiate and Doctorate in teleinformatics from KTH, The Royal Institute of Technology, Stockholm, Sweden.

Preston Marshall (pres@google.com) has more than 30 years of experience in communications systems. Currently he is with Google Inc. Previously he was deputy director at the USC Information Science Institute. Before that he was with the Defense Advanced Research Projects Agency (DARPA) Strategic Technology Office, and served as a program manager for the DARPA Wireless and Cognitive Radio and Network programs. He holds a B.S.E.E and M.S. in information science from Lehigh University, and a Ph.D. from Trinity College, Dublin.

Kwang-Cheng Chen (ckc@ntu.edu.tw) is the distinguished professor at National Taiwan University, and an IEEE Fellow. He has contributed essential technology to IEEE 802.11, Bluetooth (adaptive frequency hopping as the first widely used cognitive radio technology), and LTE and LTE-A wireless standards. He received the 2011 IEEE COMSOC WTC Recognition Award, the 2014 IEEE Jack Neubauer Memorial Award, and the 2014 IEEE COMSOC AP Outstanding Paper Award.

Markus Mueck (markus.dominik.mueck@intel.com) received the Dipl.-Ing. and ing. dipl. degrees from the University of Stuttgart, Germany, and the Ecole Nationale Supérieure des Télécommunications (ENST), Paris, France, respectively, in 1999. In 2006 he received the doctorate degree of ENST in communications. He is with Intel Deutschland GmbH, Munich, Germany, and he acts as an ETSI board member supported by INTEL, as general chairman of the ETSI RRS Technical Body (Software Radio and Cognitive Radio Standardization).

Zoran Zvonar (zoran.zvonar@mediatek.com) is the senior director and MediaTek Fellow, leading advanced technology development. Previously he was with Analog Devices (ADI) and was the recipient of the company's highest technical honor of ADI Fellow. He was named an IEEE Fellow in 2010. He received the Dipl. Ing. and the M.S. degree from the School of Electrical Engineering, University of Belgrade, and the Ph.D. degree in electrical engineering from Northeastern University, Boston.

Addressing Next-Generation Wireless Challenges with Commercial Software-Defined Radio Platforms

George Sklivanitis, Adam Gannon, Stella N. Batalama, and Dimitris A. Pados

Abstract

We review commercially available software-defined radio platforms and classify them with respect to their ability to enable rapid prototyping of next-generation wireless systems. In particular, we first discuss the research challenges imposed by the latest software-defined radio enabling technologies including both analog and digital processing hardware. Then we present the state-of-the-art commercial software-defined radio platforms, describe their software and hardware capabilities, and classify them based on their ability to enable rapid prototyping and advance experimental research in wireless networking. Finally, we present three experimental testbed scenarios (wireless terrestrial, aerial, and underwater) and argue that the development of a system design abstraction could significantly improve the efficiency of the prototyping and testbed implementation process.

Introduction

Since the early 1990s, the software-defined radio (SDR) or "software-radio" architecture, conceived by the seminal work in [1], changed the landscape of radio engineering by leveraging the flexibility provided by programmable software-reconfigurable hardware. Defining and programming radio communication functionalities in software that control heterogeneous hardware platforms such as general-purpose processors (GPPs), digital signal processors (DSPs), and field programmable gate arrays (FPGAs) was envisioned as a compelling solution for the development of flexible and reconfigurable wireless networks.

As of today, commercially available SDR platforms are capable of tuning in software critical physical layer communication parameters such as carrier frequency, bandwidth, modulation, and data rate. However, the ever growing demand for communication bandwidth, end-to-end reliability, and self-awareness under dynamically changing channel conditions requires SDR systems that exhibit self-organization and reconfiguration capabilities across all the layers of the network protocol stack [2]. Arguably, such systems would significantly benefit testbed developments in emerging research areas such as cognitive radio, multiple-input multiple-output (MIMO) communications (i.e., systems with multiple transmit/receive antennas), full-duplex multi-user MIMO, and massive MIMO (also known as large-scale antenna systems).

Existing hardware technologies such as FPGAs, DSPs, and GPPs enable individually, modular digital signal processing. However, rapid simulation, a necessary feature for modern SDR systems [3], rapid prototyping, and real-time experimental testing of next-generation/reconfigurable cross-layer optimized wireless network protocols require the development of comprehensive software environments that are able to i) adopt a holistic hardware-software approach to wireless system design; ii) provide heterogeneous multiprocessing hardware capabilities to modularize algorithmic deployment and match computational needs of signal processing software modules to computational capabilities of the available hardware platforms; iii) provide multiple layers of abstraction to tame hardware complexity and allow rapid iteration between simulation and prototyping.

In an attempt to address latency and throughput needs in software-defined wireless networking, recent software developments in both GPP-centric and FPGA-centric SDR architectures follow parallel paths with little or no cross-fertilization between the two software communities. Consequently, there are a plethora of software tools and pertinent libraries that are tightly integrated with specific hardware platforms, without offering consistent abstractions toward a unified SDR system design and simulation philosophy. Software tools are distinguished into open and closed source (proprietary) programming environments, each providing different levels of flexibility and integration with hardware. Furthermore, distributing processing across heterogeneous hardware platforms toward, for instance, optimizing the implementation of resource-demanding cross-layer networking protocols requires developers to master multiple software tools. At the same time, hardware configurations vary across different SDR vendors, while software portability is not guaranteed across different SDRs. Hence, facilitating the transition from concept to simulation and then to prototype can become quite a tedious

The authors review commercially available software-defined radio platforms and classify them with respect to their ability to enable rapid prototyping of next-generation wireless systems. They discuss research challenges imposed by the latest software-defined radio enabling technologies including both analog and digital processing hardware, and they present the state-of-the-art commercial software-defined radio platforms.

The authors are with the State University of New York at Buffalo.

0163-6804/16/$25.00 © 2016 IEEE

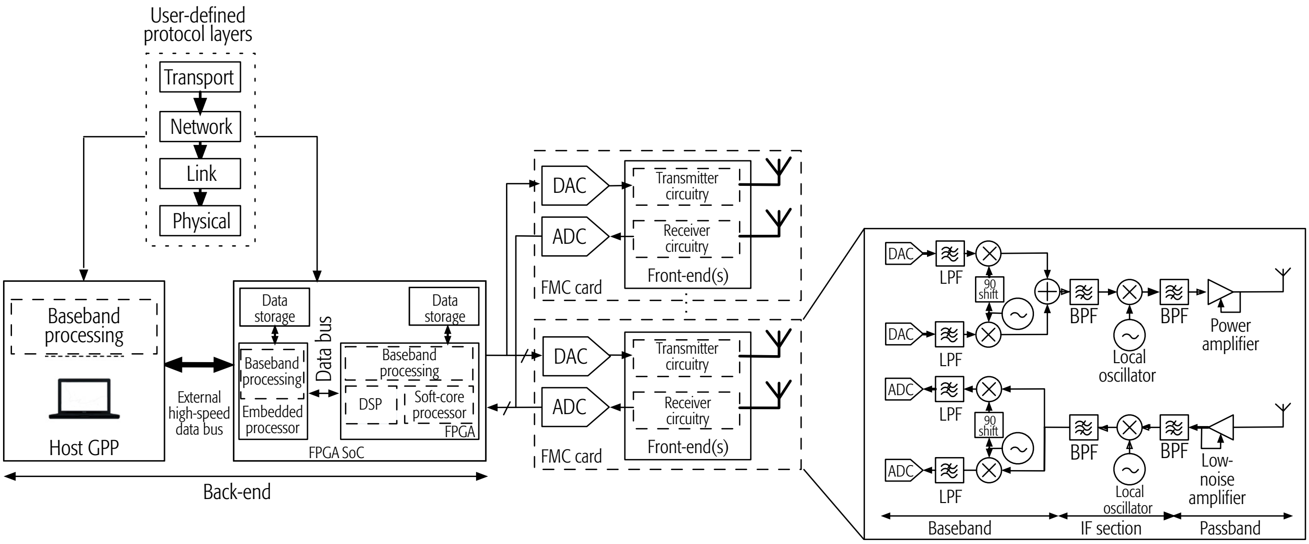

Figure 1. Generic software-defined radio architecture. Static analog transmit/receive circuitry (front-end) is interfaced with programmable hardware processing platforms (back-end) through ADCs/DACs. Analog-to-digital converters (ADCs)/digital-to-analog converters (DACs) and analog front-end radio circuitry may be featured under the same compact card, referred to as an FPGA mezzanine card (FMC).

and daunting process for developers who must trade off development time for system optimization and vice versa, until they finely tune to the wireless system's desired performance for different communication standards.

In light of the above, in this article we present current SDR challenges with respect to rapid prototyping and testing of reconfigurable software-defined wireless networks, review the existing state-of-the-art SDR platforms, and discuss their limitations with respect to projected challenges of the next-generation programmable wireless networking experimental testbeds and applications.

The rest of the article is organized as follows. We first provide an overview of a generic SDR architecture and discuss in detail the emerging challenges toward rapid experimental assessment and testing of novel wireless networking protocol proposals. Then we list and classify state-of-the-art commercial SDRs according to the capabilities and limitations of the underlying hardware and software platforms. Finally, we highlight key future challenges and directions dictated by the real-world experimental requirements of three diverse (i.e., terrestrial, aerial, underwater) wireless networking testbeds.

Software-Defined Radio Architecture

Software-defined radio proposes a paradigm shift from inherently inflexible dedicated-functionality hardware radio platforms by combining analog static or parameterizable (front-end) circuits and software reprogrammable digital hardware (back-end) components that are easily reconfigurable via software updates. Software-defined architectures are therefore ideal for rapid prototyping, and testing of new military applications and commercial standards. In the context of this work we use the term front-end to describe the analog signal processing stages between the antenna and analog-to-digital converters (ADCs) or digital-to-analog converters (DACs), and the term back-end to refer to software reprogrammable digital processing platforms such as GPPs, DSPs, and FPGAs.

Figure 1 illustrates a generic SDR architecture adopted by the majority of commercially available SDRs, consisting of an analog front-end interfaced with ADC/DAC converters, an FPGA, and a GPP. The SDR front-end consists of analog circuitry responsible for up/down-conversion of analog information signals directly to either the passband or baseband, respectively (homodyne, zero-intermediate frequency [IF] architecture) or an IF ([super-] heterodyne architecture). Bandpass and lowpass filters, and amplifiers in the front-end are used for signal conditioning. The analog front-end is interfaced with high sample rate/resolution ADCs that sample the baseband signal, and DACs that convert digital samples to analog waveforms for transmission. As a result, analog passband signals that arrive at the receiver antenna(s) are first bandpass filtered, then amplified via a low-noise-amplifier (LNA), down-converted directly to baseband or optionally to an IF, lowpass filtered, and finally amplitude leveled/power normalized by an automatic gain controller (AGC) before being sampled at the Nyquist rate by the ADC. The reverse process is followed at the transmit chain of the front-end, where incoming complex baseband signals from the DAC are filtered, up-converted, and amplified for passband transmission.

Baseband signal processing on the ADC output and DAC input digital samples is handled by the SDR back-end. User-defined protocol functionalities (Fig. 1) at different layers of the network protocol stack exhibit variable latency and memory requirements for multiple communication standards, so the wireless system designer may decide to split the execution of certain functionalities to heterogeneous hardware platforms (e.g., GPP, FPGA). The high parallelism offered by the FPGA is usually leveraged to accelerate the implementation of computationally demanding signal processing operations (e.g., filters) on incoming/outgoing data from/to the ADC/DAC at the expense of increased implementation complexity. Digital data can be transferred to and stored long-term in an external storage medium such as a secure digital (SD) card or an onboard synchronous dynamic random access memory (SDRAM) for faster access, while upper layer (e.g., link and network layer) functionalities may be handled by either a software co-processor implemented in the FPGA, an embedded DSP, or a GPP (as depicted in Fig. 1). GPPs are well suited for the implementation of highly branching programs and offer short development times by exploiting high-level software programming languages. However, real-time operating systems at GPPs offer low resolution in strict real-time data flow constraints. Typical SDR designs implement physical layer functionalities and handle the data at the packet level at a GPP. Digital samples are then transferred to the FPGA through an external high-speed data bus connection (e.g., Gigabit Ethernet). GPPs are either external host PCs or embedded system-on-chip (SoC) processors, sometimes even incorporated in the same integrated circuit (IC) package with the FPGA.

Software-Defined Radio Challenges in Next-Generation Wireless Systems Prototyping

Next-generation wireless networking protocols and sophisticated network topologies such as multi-user MIMO are difficult to model and test in a software simulation environment. For example, software simulation mostly relies on simplified channel fading models that do not incorporate critical real-world networking conditions or hardware impairments. As a result, researchers can either loosely emulate in software the performance of new signal processing and wireless networking algorithms, or experimentally validate them in heterogeneous multiprocessing hardware platforms consisting of application-specific ICs (ASICs), FPGAs, GPPs, and DSPs. Figure 2 illustrates the trade-off between reconfigurability and development time for FPGA, ASIC, DSP, GPP, and hybrid GPP/FPGA-centric SDR architectures.

ASIC implementations have a static application-specific architecture and provide tailored processing units to optimize computational efficiency and power consumption for dedicated functionalities. On the other hand, FPGAs are field-programmable for different applications and provide rapid reconfiguration between different signal processing designs, thus enabling the development of software programmable SDR platforms at the expense of increased power consumption and circuit area. Reconfigurable SDR platforms aim to minimize the utilization of ASIC modules such as analog filters, amplifiers, digital down and up converters (DDC and DUC) due to their limited flexibility, or allow their parameterization and runtime reconfiguration through primitive functions that are activated by a processor such as a DSP or a GPP. DSPs offer the best trade-off between processing power and power consumption by

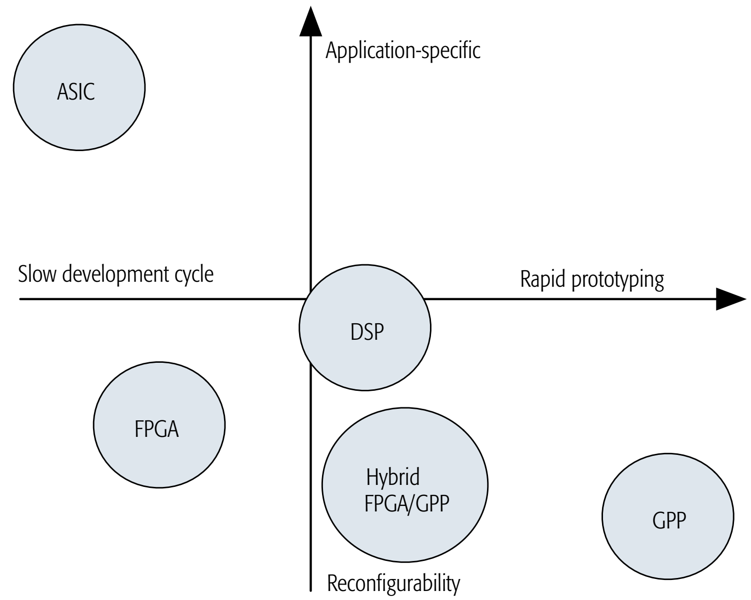

Figure 2. Trade-off between reconfigurability and development time for FPGA, ASIC, DSP, GPP, and hybrid GPP/FPGA-centric SDR architectures.

providing optimized features especially targeted for digital signal processing operations (e.g., combined multiply-accumulate operations). They are usually embedded in FPGA hardware platforms to efficiently address signal processing tasks that can be pipelined (i.e., sequenced and repeated for each sample in a buffer). In addition, GPPs are real-time reprogrammable processing alternatives that can handle processing of a wide variety of applications at low implementation complexity. Furthermore, multi-core architectures of DSPs and GPPs can enhance processing performance by executing multiple operations in parallel.

The rest of this section reviews and discusses the major emerging hardware and software challenges in existing commercially available SDRs toward accelerating experimental assessment and testing of novel wireless networking protocols. In this context, we study and classify SDR systems according to the following criteria:

- Level of flexibility and efficient interaction between analog front-end and digital back-end hardware technologies
- Multiprocessing capabilities in heterogeneous hardware platforms
- Level of abstraction between software environments and hardware platforms

Analog/Digital Hardware: The analog nature of the transmission medium (air, water, soil, etc.) as well as the requirements for multi-band support to accommodate multiple standards challenge the design of reconfigurable SDR architectures with respect to both the design of the analog front-end circuitry and their respective physical interface with back-end digital processing platforms. As an example, the (super)-heterodyne transceiver architecture depicted in the circuitry design of Fig. 1 fails to reconfigure its fixed narrowband components, such as channel selective filters, to meet the broadband requirements of multiple standards. On the other hand, the homodyne (zero-IF) transceiver architecture

The majority of existing SDR software tools are tightly integrated with specific hardware platforms with few or no abstractions available to the wireless system designer. As a result, optimization of the system design flow at a high level depends on software-hardware co-design for objects executed in processors and FPGAs, respectively.

minimizes the number of analog (fixed) components and is a flexible design alternative to the (super)-heterodyne architecture at the expense of increased baseband interference and phase noise (DC offset due to local oscillator leakage and IQ imbalance, respectively), which can be partially compensated for in the digital domain. Flexible receiver design implementation with minimal analog components requires ADCs with high dynamic range to be able to detect weak desired signals in the presence of possibly strong undesired in-band signal interference. Furthermore, the effective (useable) bandwidth of the SDR system of Fig. 1 is defined as the minimum of:

- The radio's analog bandwidth, which should not exceed the ADC/DAC sample rate
- The FPGA processing bandwidth, which depends on the FPGA clock rate
- The host-GPP bandwidth

Typical commercial SDR platforms either include a high-speed 1/10 Gigabit Ethernet (GigE) connection to address the host GPP-FPGA communication or implement FPGA-based soft-core processors or select wireless-focused embedded hard processors (e.g., ARM). In this way, the FPGA is able to handle GPP-like branching logic. Therefore, performance efficiency for multiple standards entails effective interaction and data transfer optimization between different hardware platforms at the SDR back-end.

Heterogeneous Multiprocessing: Multiprocessing is an SDR intrinsic requirement that aims to enhance SDR computational efficiency by off-loading complex signal processing operations such as finite impulse response (FIR) filtering or fast Fourier transform (FFT) to heterogeneous hardware platforms. However, software development of such architectures suffers from increased implementation complexity and lack of a standard methodology for partitioning the implementation of signal processing functionalities to heterogeneous hardware platforms [4]. As an example, GPPs and DSPs have memory architectures well suited for streaming data and rely on sequential execution of a single instruction stream. They are single-instruction multiple-data (SIMD) units where each processor is constrained to execute the same instruction at a single program but operate over multiple data streams in parallel. Therefore, GPPs and DSPs are suitable for simple computational tasks such as modulation/demodulation and encoding/decoding, higher-level logic applications, and network and medium access layer functions. On the other hand, FPGAs are well suited for implementation of logical functions that can be separated and run simultaneously or for execution of tasks that use non-standard data-type representations (i.e., 12- or 14-bit ADC/DAC); however, they exhibit increased development time and complexity. Data-type consistency, sampling rate adaptation, and resource mapping and scheduling of software applications for heterogeneous hardware are additional architectural considerations in the development of heterogeneous multiprocessing platforms. Exchanging a common data stream between multiple hardware architectures that use different data-type representations, say between a GPP that utilizes floating-point SIMD instruction sets and a DSP or FPGA that uses fixed-point representations, requires data-type conversions. Data-type conversions (e.g., conversion of a 12/14-bit ADC output to 16-bit) short at GPP input lead to increased processing load. Furthermore, sampling rate adaptation is necessary for interfacing multiple subsystems with different clock rates as well as for fine tuning to the sampling rate of multiple communication standards. Implementation of reconfigurable sampling rate adaptors suffers from the design of highly dynamically adaptive and computationally efficient FIR filters. Thus, existing SDR architectures assign the implementation of sampling rate adaptors to the FPGA. Finally, automatic instead of ad hoc resource mapping for heterogeneous multiprocessing may enable adaptation of protocol execution speeds that can be traded off during runtime to satisfy pertinent latency requirements of the specific communication standard.

Hardware Abstraction: The majority of existing SDR software tools are tightly integrated with specific hardware platforms with few or no abstractions available to the wireless system designer. As a result, optimization of the system design flow at a high level depends on software-hardware co-design for objects executed in processors and FPGAs, respectively. As an example, a researcher willing to develop and experimentally evaluate a complex modulation waveform or channel/source coding scheme is required to develop and optimize signal processing components, as well as decide on the optimal execution architecture according to the available hardware platforms and the respective software tools. Nevertheless, the implementation of abstract hardware application programming interfaces (APIs) that offer granular control and separate the execution architecture from signal processing may allow trading off prototyping time for optimization of platform resources and vice versa. In addition, the availability of dedicated software tools for particular hardware platforms in combination with hardware abstractions may give the wireless system designer the opportunity to integrate optimized intellectual property (IP-protected) processing components developed by domain specialists or third party developers, and rely only on abstract characteristics such as memory use and execution speed. Multiple levels of granularity in hardware control may also accelerate the transition from simulation to prototyping by providing bit level visibility into the design and the capability to quickly test its functional behavior. FPGA software development will definitely benefit from such abstractions as currently, long compile times, synthesize times, and place-and-route times are prohibitive for trial-and-error performance assessment of new wireless designs.

Commercial Software-Defined Radio Platforms

For small-scale laboratory testbed setups, commercially available SDR platforms offer low-cost hardware and software solutions for rapid experimental assessment of programmable wireless

	Ettus/NI USRP		Nutaq		Mango WARP		Per Vices		Epiq	
	GPP	FPGA	GPP	FPGA	GPP	FPGA	GPP	FPGA	GPP	FPGA
GNU Radio	•	N/A	•	N/A	○	N/A	•	N/A	•	N/A
Mathworks	•	○	•	•	•	•	○	○	○	○
NI LabVIEW	•	•	○	○	○	○	○	○	○	○
• Compatible option; ○: Incompatible option; N/A: option is not available.										

Table 1. Software/hardware compatibility for state-of-the-art commercially available SDR platforms.

GNU Radio is an open source software framework that follows a component-based design, where signal processing chains are broken into primitive components/blocks, therefore enabling code reusability and rapid block reconfiguration. Each block is assigned to a dedicated processor thread, while data exchange between blocks is achieved through shared memory buffers.

networks. In the following section, we discuss the strengths and limitations of both software frameworks and hardware architectures with respect to rapid prototyping and testing of next-generation wireless systems. We also provide a compatibility overview between the available software tools and heterogeneous SDR platforms in Table 1.

Software Frameworks

GNU Radio: GNU Radio is an open source software framework that follows a component-based design, where signal processing chains are broken into primitive components/blocks, therefore enabling code reusability and rapid block reconfiguration. Each block is assigned to a dedicated processor thread, while data exchange between blocks is achieved through shared memory buffers. GNU Radio applications are only supported by GPPs (Table 1) as well as embedded processors that support floating point SIMD instruction sets. New GNU Radio applications, called flow graphs, are programmed in Python and C++. The API of the GNU Radio framework enables integration of optimized signal processing blocks (GNU Radio IP) such as modulators and demodulators. The framework allows simulation of the performance of new physical layer communication designs through an XML-based graphical user interface (GUI), called GNU Radio Companion (GRC). With respect to performance acceleration of pure GPP-centric designs, GNU Radio provides a programming tool, named VOLK (vector-optimized library of kernels) [5], which enables vectorized mathematical operations and is independent of the processor's architecture. VOLK offers an abstraction layer to hardware-specific SIMD implementations, which vary across different processor families and vendors. Packet-based processing in upper layers is enabled by stream tagging and asynchronous message passing software APIs. More specifically, stream tags enable sample streams to carry metadata information (e.g., information on packet boundaries), while message passing enables asynchronous parameter reconfiguration in flow graph blocks regardless of their location in the data flow (upstream or downstream blocks). As a result, GNU Radio provides rapid software simulation and testing of new GPP-centric wireless system designs.

Mathworks/LabVIEW: Mathworks provides rapid radio prototyping solutions [6] by building a bridge between simulation at the Simulink environment and execution on heterogeneous hardware platforms, such as an FPGA interfaced with a DSP or GPP. More specifically, Simulink follows a graphical high-level modeling design approach that enables fast building and simulation of new designs based on Mathworks libraries. Simulation models are then linked to C-based DSPs or GPPs through the Mathworks RealTime Workshop tool, which enables automatic translation of Simulink simulation models into C-code. Simulink has the ability to interface with System Generator, a Xilinx DSP design tool, thus reducing the FPGA software development time. System Generator contains platform-specific sets of Xilinx FPGA IP blocks such as FFT, filters, and memory blocks that are guaranteed to exhibit equivalent cycle accuracy to the IP blocks available at Simulink. Mathworks therefore provides the required hardware abstractions to the SDR system designer to effectively integrate heterogeneous processing elements in the Simulink model-based environment and thus accelerate the transition from simulation to real-world testing. However, the available hardware APIs are constrained by the FPGA provider and technology. Another Mathworks tool that requires no prior experience with low-level FPGA or register-level (RTL) programming is the hardware-description-language (HDL) coder. The HDL coder enables automatic conversion of floating-point MATLAB or Simulink simulation models into fixed-point FPGA designs that are ready to be synthesized. Similar programming approaches are followed by the LabVIEW software tools, which target rapid development of heterogeneous multiprocessing system designs, that is, hybrid GPP-FPGA-centric architectures restricted to Ettus/National Instruments SDR platforms (Table 1).

Hardware Architectures

Ettus/NI: Ettus/National Instruments (NI) Universal Software Radio Peripherals (USRPs) are a family of heterogeneous hardware SDR platforms. They are classified into the Networked (N)-series, the Bus (B)-series, the Embedded (E)-series, and the X-series. USRP N- and X-series consist of an FPGA-based motherboard that is interfaced with a single-input single-output (SISO) daughterboard and multiple daughterboard(s) (MIMO capable), respectively, through high sample rate ADCs and DACs. Table 2 lists the frequency range (DC-6.0 GHz) and bandwidth capabilities (10–160 MHz) of Ettus daughterboards, which allow for flexible, interchangeable analog front-end circuitry for a variety of applications. FPGA-based signal processing designs are either controlled by host GPPs via external high-speed data bus connec-

Daughterboard	Function	Low frequency	High frequency	Max RF Bandwidth per channel
UBX[1]	Transceiver	10 MHz	6 GHz	40 MHz 160 MHz
WBX[1]	Transceiver	50 MHz	2.2 GHz	40 MHz 120 MHz
CBX[1]	Transceiver	1.2 GHz	6.0 GHz	40 MHz 120 MHz
SBX[1]	Transceiver	400 MHz	4.4 GHz	40 MHz 120 MHz
XCVR	Transceiver[2]	2.4 GHz 4.9 GHz	2.5 GHz 5.9 GHz	36 MHz (RX) 48 MHz (TX)
Basic RX	Receiver	1 MHz	250 MHz	100 MHz
Basic TX	Transmitter	1 MHz	250 MHz	100 MHz
LFRX	Receiver	DC	30 MHz	30 MHz
LFTX	Transmitter	DC	30 MHz	30 MHz
DBSRX2	Receiver	800 MHz	2.3 GHz	60 MHz
TVRX2	Receiver	50 MHz	860 MHz	10 MHz

[1] The N-series family of USRPs is compatible only with 40 MHz bandwidth daughterboards, while the X-series can support both 40 MHz and 120/160 MHz daughterboards.
[2] TVRX2 supports half-duplex operation only.

Table 2. Ettus USRP daughterboards.

tions (Table 3, USRP X, N, and B-210), by soft-core processor cores implemented at the FPGA (e.g., 32-bit AeMB, MicroBlaze soft-processor cores), or by wireless-focused embedded (SoC) processors such as ARM Cortex-A9, that enable small-form-factor, portable SDR solutions (Table 4, USRP E310). Transceiver architectures such as homodyne or (super)-heterodyne vary across different daughterboards, while ADC, DAC sample rates and bit resolutions vary across different models of the USRP-series. The USRP family is compatible with the majority of software frameworks, as seen in Table 1, through the USRP-Hardware-Driver (UHD) software API that acts as a host communication driver for controlling Ettus SDRs. Furthermore, Ettus introduces a novel Network-on-Chip mechanism (RFNoC) [7], which enables integration of heterogeneous processing components into a GNU Radio flow graph by providing a standard method of consistently routing data and distributing processing throughout complex heterogeneous hardware platforms (i.e., FPGA and GPP). As a result, a researcher is able to minimize development time by integrating modular IP blocks that are executed in heterogeneous hardware. At the same time, she is able to maintain system design flexibility during runtime for performance demanding protocols by moving timing critical and computationally complex functionalities (e.g., lower medium access control and physical layer) at the FPGA, and implementing high-level control (e.g., upper medium access control) functionalities at the GPP. However, testbed requirements pertinent to specific applications entail proper selection and combination of daughterboards, USRP hardware back-end, and software.

Nutaq: Table 3 depicts two heterogeneous hardware SDR platforms provided by Nutaq, called PicoSDR (FPGA interfaced with an embedded PC or host PC) and ZeptoSDR (SoC-FPGA platform). Instead of using interchangeable analog front-ends or integrated radio chipsets, Nutaq SDRs use an FPGA mezzanine card (FMC) that features ADCs/DACs on the same compact card with analog front-end circuitry. The FMC card's transceiver follows a homodyne architecture with software selectable bandpass and baseband filters that enable easy runtime adaptation of the analog front-end to multiple communication standards. Furthermore, the separation of SDR back-end from the ADC/DAC and front-end circuitry enables hardware interoperability to deal with challenges imposed by different channel environments, on the condition that software APIs are backward compatible with different analog components. Nutaq provides an FPGA framework for embedded applications development (board-software-development kit [BSDK]) that includes custom-built hardware IPs and software APIs to enable efficient interaction with both Mathworks and GNU Radio software tools. In addition, Nutaq provides a set of custom blocks to control and handle real-time data exchange between the host or embedded processor and FPGA, thereby enabling easy synthesis and testing of heterogeneous multiprocessing designs through Simulink's model-based design approach.

Others: Tables 3 and 4 present additional tabletop and small-form-factor SDR platforms. In particular, the third revision of the wireless open access research platform (WARP) provides an FPGA-based SDR architecture with two integrated ADC/DAC and homodyne radio transceiver chipsets. The WARP platform allows expansion to four radio transceivers through FMC expansion ports, while rapid software development of new wireless designs is enabled through the WARPLab framework. WARPLab provides online and offline processing capabilities by allowing rapid physical layer prototyping with MATLAB at the host PC. Real-time online processing requires user modifications to existing WARPLab reference designs, which are implemented in MATLAB and System Generator. Epiq provides low-power portable SISO and MIMO-capable SDR platforms with proprietary IPs to reduce development time. More specifically, Maveriq's package features an FPGA interfaced with two homodyne wideband radio transceivers, an embedded Intel processor and an internal hard drive for recording and playback operations. Sidekiq and Matchstiq are small-form-factor SISO-capable alternatives that provide host and embedded multi-core processing capabilities, respectively. Per Vices offers both a SoC-based wideband SDR platform (Table 3, Crimson) equipped with four integrated homodyne radio transceivers, as well as a low-cost peripheral-component-interconnect-express (PCIe)-based alternative (Table 3, Noctar), which can be controlled and reprogrammed by custom web or Python-based GUIs. Both Epiq and Per Vices SDR platforms can be interfaced with GNU Radio; however, they currently lack the software abstractions to accelerate complex

SDR platform	Hardware	Number of TX/RX antennas	Low frequency	High frequency	Max. RF BW per IQ channel	ADC speed (MS/s, bits)	DAC speed (MS/s, bits)	Ext. data bus host connections	Maximum host throughput
Ettus USRP B-210	Xilinx Spartan 6 FPGA with integrated radio chipset	2/2	70 MHz	6 GHz	56 MHz	61.44,12	61.44,12	USB 3.0	1.96 Gb/s
Ettus USRP N-series[1]	Xilinx Spartan 3A-DSP FPGA	1/1	DC	6 GHz	40 MHz	100,14	400,16	GigE	0.8 Gb/s
Ettus USRP X-series[1]	Xilinx Kintex 7 FPGA	2/2	DC	6 GHz	160 MHz	200,14	800,16	Dual 1/10 GigE PCIex4	6.4 Gb/s
Nutaq PicoSDR[2]	Xilinx Virtex 6 FPGA	2/2	300 MHz	3.8 GHz	28 MHz	80,12	80,12	Dual GigE PCIex4	6.4 Gb/s
Nutaq ZeptoSDR	Xilinx Zynq 7020 SoC (Dual-core ARM Cortex-A9)	1/1	300 MHz	3.8 GHz	28 MHz	80,12	80,12	GigE	0.8 Gb/s
Mango WARP v3	Xilinx Virtex 6 FPGA	2/2	2.4 GHz 4.9 GHz	2.5 GHz 5.8 GHz	40 MHz	100,12	170,12	Dual GigE	0.8 Gb/s
Epiq Maveriq	Dual-core Intel Atom processor Xilinx Spartan 6 FPGA	2/2	70 MHz	6 GHz	50 MHz	50,12	50,12	GigE Dual USB	0.8 Gb/s
Epiq Sidekiq	Xilinx Spartan 6 FPGA	1/2	70 MHz	6 GHz	50 MHz	61.44,12	61.44,12	PCIex1 USB 2.0	1.6 Gb/s
Per Vices Noctar	Altera Cyclone IV FPGA	1/1	100 kHz	4.4 GHz	200 MHz	125,12	250,16	PCIex4	6.4 Gb/s
Per Vices Crimson	Altera Arria V ST SoC (Dual-core ARM Cortex-A9 MP)	4/4	100 kHz	6 GHz	322 MHz	370,16	2500,16	Dual 1/10 GigE USB	6.4 Gb/s

[1] Larger-scale antenna system setups with USRP X or N-series require additional hardware such as an OctoClock(-G), which is a clock distribution system for coherent operation of multiple SDRs under external clock reference.
[2] PicoSDR is also available in a 4 × 4 MIMO antenna configuration and a 2 × 2 embedded-PC (Quad-core i7) configuration.

Table 3. State-of-the-art tabletop SDR platforms.

algorithmic developments on either the FPGA or GPP.

Future Research Challenges and Next-Generation SDR Applications

A platform-independent system design flow can expedite rapid testing and fast standardization of novel signal processing algorithms and networking protocols. Such a high-level system design flow should also be capable of supporting the throughput and latency requirements of next-generation mobile communication protocols such as Long Term Evolution-Advanced (LTE-A) or large-scale MIMO systems, advances in the 802.11 family of networking standards as well as state-of-the-art programmable wireless networks (e.g., cognitive radio networks with dynamic spectrum access capabilities) that require fast and intelligent adaptation at all layers of the protocol stack. Existing software-hardware SDR architectures offer considerable flexibility and high performance for experimentation with new concepts at the physical layer, but lack adequate and coherently designed abstractions to define either networking protocols with cross-layer interactions across multiple layers of the protocol stack or decision making mechanisms to control such interactions [8]. Thus, state-of-the-art medium access protocols that need to comply with the IEEE 802.11 distributed coordination function (DCF) and enhanced distributed channel access (EDCA) timings, as well as routing protocols are implemented "from scratch" for different SDR testbed configurations. In this context, we describe the research challenges arising from three wireless networking testbeds (wireless terrestrial, aerial, and underwater) and argue for the potential future benefits of wireless network-specific APIs and abstractions to easily control network reconfiguration.

All-Spectrum Cross-Layer Optimized Cognitive Networks

In the context of multihop cognitive underlay networks, where primary spectrum licensees coexist with unlicensed secondary users, state-of-the-art proposals require reprogramming and optimization of the entire wireless protocol stack for every specific SDR platform. For example, the work in [9–11] implements a distributed algorithm that maximizes secondary network throughput, while at the same time avoiding interference to primary users. Implementation and deployment of the proposed cognitive algorithm at each SDR node requires complex optimization and low-latency interaction across the network and physical layers of the protocol stack to enable real-time adaptive decisions on the optimal channel waveform and routing path. Thus, implementation and deployment efforts would significantly bene-

SDR platform	Hardware	Number of TX/ RX antennas	Low freq.	High freq.	RF BW per IQ channel	ADC speed (MS/s, bits)	DAC speed (MS/s, bits)	Form factor
Ettus USRP E310	Xilinx Zynq 7020 SoC (Dual-core ARM Cortex-A9) with integrated radio chipset	2/2	70 MHz	6 GHz	56 MHz	61.44, 12	61.44, 12	133 × 68 × 26.4 mm
Epiq Matchstiq S10	Xilinx Spartan 6 FPGA Quad-core ARM Cortex-A9	1/1	70 MHz	6 GHz	50 MHz	61.44, 12	61.44, 12	114.3 × 40. × 27.9

Table 4. Portable small-form-factor SDR platforms.

fit from future SDR architectures that effectively separate decision and control from data processing so that the optimization of the protocol stack becomes SDR-platform-independent.

Software-Defined Autonomous Airborne Networks

The design and evaluation of airborne networks [12] in the context of network-centric warfare suffer from high network dynamics, such as bandwidth efficiency, link reliability, and security, across wireless nodes that are either geographically or hierarchically dispersed. Existing approaches rely on either unrealistic channel/network simulations or small-scale laboratory SDR setups with static wireless protocol designs. Thus, an SDR architecture with self-reconfigurable functionalities that can easily be controlled through a modular and declarative programming interface may enable rapid real-world deployment of novel complex communication techniques [13]. Such a radio architecture may also decouple network implementation from the underlying hardware through available wireless network APIs, and thus allow easy heterogeneous network integration and autonomous reconfiguration for next-generation airborne networks.

Real-Time Reconfigurable Underwater Acoustic Networks

Spectral efficiency of underwater acoustic networks is currently limited by the spatially and temporally variable characteristics of the underwater acoustic channel. As a result, research efforts focus on tailored protocol designs for different layers of the wireless protocol stack that are well suited to the deployment environment (e.g., ocean, lake). At the same time, existing commercially available underwater acoustic modems rely on fixed hardware designs that prevent real-time reconfiguration at any layer of the protocol stack or experimental testing of novel algorithmic developments [14, 15]. As a result, the definition of abstractions in a software defined radio architecture to handle cross-layer interactions will allow next-generation software-defined acoustic modems to decide intelligently and adapt their communication parameters to maximize spectral efficiency.

Conclusions

We have presented a comprehensive overview of the SDR design challenges in the context of next-generation programmable wireless networks, reviewed commercially available hardware and software platforms, and classified them according to their capability to promote rapid prototyping, testing, and evaluation. Finally, we have referred to three experimental testbed environments (wireless terrestrial, aerial, and underwater) and argued that although there is a wide variety of hardware and software frameworks for flexible SDR prototyping, the selection of the appropriate platform is still connected to the application and communication standard requirements. Thus, SDR-platform-independent system design flow is necessary to allow rapid prototyping, testing, evaluation, and standardization of novel wireless networking proposals.

Acknowledgments

This work was supported in part by the National Science Foundation under grants CNS-1422874, CNS-1126357, CNS-1117121, and CNS-1055945.

References

[1] J. Mitola, "Software Radio Architecture," *IEEE Commun. Mag.*, vol. 33, no. 5, May 1995, pp. 26–38.
[2] M. Mueck *et al.*, "ETSI Reconfigurable Radio Systems: Status and Future Directions on Software Defined Radio and Cognitive Radio Standards," *IEEE Commun. Mag.*, vol. 48, no. 9, Sept. 2010, pp. 78–86.
[3] T. Ulversoy, "Software Defined Radio: Challenges and Opportunities," *IEEE Commun. Surveys Tutorials*, vol. 12, no. 4, 4th qtr. 2010, pp. 531–50.
[4] D. Efstathiou, J. Fridman, and Z. Zvonar, "Recent Developments in Enabling Technologies for Software Defined Radio," *IEEE Commun. Mag.*, vol. 37, no. 8, Aug. 1999, pp. 112–17.
[5] T. Rondeau, N. McCarthy, and T. O'Shea, "SIMD Programming in GNU Radio, Maintainable und User-Friendly Algorithm Optimization with volk," *SDR, Wireless Innovation Forum Europe*, June 2012.
[6] C. Moy and R. M., "High-Level Design Methodology for Ultra-Fast Software Defined Radio Prototyping on Heterogeneous Platforms," *Advances in Electronics and Telecommun.*, vol. 1, no. 1, Apr. 2010, pp. 67–85.
[7] J. Malsbury and M. Ettus, "Simplifying FPGA Design with a Novel Network-on-chip Architecture," *Proc. 2nd Wksp. Software Radio Implementation Forum*, ser. SRIF '13, 2013, pp. 45–52.
[8] E. Demirors *et al.*, "RcUBe: Real-Time Reconfigurable Radio Framework with Self-Optimization Capabilities," *Proc. 12th IEEE Int'l. Conf. Sensing, Commun., and Networking*, June 2015.
[9] K. Gao *et al.*, "Cognitive Code-Division Channelization," *IEEE Trans. Wireless Commun.*, vol. 10, no. 4, Apr. 2011, pp. 1090–97.
[10] L. Ding *et al.*, "All-Spectrum Cognitive Networking through Joint Distributed Channelization and Routing," *IEEE Trans. Wireless Commun.*, vol. 12, no. 11, Nov. 2013, pp. 5394–5405.
[11] G. Sklivanitis *et al.*, "All-Spectrum Cognitive Channelization Around Narrowband and Wideband Primary Stations," *Proc. IEEE GLOBECOM*, Dec. 2015.
[12] K. Kwak *et al.*, "Airborne Network Evaluation: Challenges and High Fidelity Emulation Solution," *IEEE Commun. Mag.*, vol. 52, no. 10, Oct. 2014, pp. 30–36.
[13] W. Su *et al.*, "On the Capacity of Airborne MIMO Communications," *IEEE GLOBECOM*, Dec. 2012, pp. 4629–34.
[14] G. Sklivanitis *et al.*, "Receiver Configuration and Testbed Development for Underwater Cognitive Channelization," *Proc. 48th Asilomar Conf. Signals, Systems and Computers*, Nov. 2014.
[15] E. Demirors *et al.*, "Design of A Software-Defined Underwater Acoustic Modem with Real-Time Physical Layer Adaptation Capabilities," *Proc. ACM Int'l. Conf. Underwater Networks & Systems*, Nov. 2014.

Biographies

George Sklivanitis [S'11] (gsklivan@buffalo.edu) received his Diploma degree in electronic and computer engineering from the Technical University of Crete, Greece, in 2010. He is currently working toward his Ph.D. degree in electrical engineering at the State University of New York at Buffalo, and his research interests span the areas of signal processing, software-defined wireless communications and networking, cognitive radio, and underwater acoustic communications. In 2014 he was the winner of the Nutaq Software-Defined Radio Academic US National Contest and in 2015 he received the Graduate Student Excellence in Teaching award from SUNY Buffalo.

ADAM GANNON [S'14] (adamgann@buffalo.edu) received his Bachelor of Science degree in electrical engineering in 2013 from the State University of New York at Buffalo, where he is currently pursuing his Ph.D. degree in electrical engineering. His research interests include signal processing, software-defined and cognitive radio, and wireless communication in airborne and space environments.

STELLA N. BATALAMA [S'91, M'94] (batalama@buffalo.edu) received her Diploma degree in computer engineering and science from the University of Patras, Greece, in 1989 and her Ph.D. degree in electrical engineering from the University of Virginia, Charlottesville, in 1994. In 1995 she joined the Department of Electrical Engineering, State University of New York at Buffalo, where she is presently a professor. From 2009 to 2011, she served as the Associate Dean for Research of the School of Engineering and Applied Sciences, and since 2010, she has served as chair of the Electrical Engineering Department. During the summers of 1997–2002 she was visiting faculty at the U.S. Air Force Research Laboratory (AFRL), Rome, New York. From August 2003 to July 2004 she served as acting director of the AFRL Center for Integrated Transmission and Exploitation (CITE), Rome, New York. Her research interests include small-sample-support adaptive filtering and receiver design, cooperative communications, cognitive networks, underwater communications, covert communications, steganography, compressive sampling, adaptive multiuser detection, robust spread-spectrum communications, and supervised and unsupervised optimization. She was an Associate Editor for *IEEE Communications Letters* (2000–2005) and *IEEE Transactions on Communications* (2002–2008).

DIMITRIS A. PADOS [M'95, SM'15] (pados@buffalo.edu) received his Diploma degree in computer science and engineering from the University of Patras and his Ph.D. degree in electrical engineering from the University of Virginia, Charlottesville. Since August 1997, he has been with the Department of Electrical Engineering, State University of New York at Buffalo, where he currently holds the title of Clifford C. Furnas Professor of Electrical Engineering. His research interests are in the general areas of communication theory and systems and adaptive signal processing with applications to interference channels and signal waveform design, secure wireless communications, and cognitive acoustic-to-RF modems and networks.

Revolutionizing Software Defined Radio: Case Studies in Hardware, Software, and Education

Alexander M. Wyglinski, Don P. Orofino, Matthew N. Ettus, and Thomas W. Rondeau

SDR has increasingly become an invaluable research, development, and educational tool within the telecommunications sector with respect to rapidly prototyping new algorithms and paradigms in actual radio hardware and evaluating them in real-world over-the-air conditions. Due to advances in microprocessor technology, radio frequency hardware, and software, SDR has matured into a reliable tool that is now part of almost every communication engineer's toolbox.

Abstract

SDR has increasingly become an invaluable research, development, and educational tool within the telecommunications sector with respect to rapidly prototyping new algorithms and paradigms in actual radio hardware and evaluating them in real-world over-the-air conditions. Due to advances in microprocessor technology, radio frequency hardware, and software, SDR has matured into a reliable tool that is now part of almost every communication engineer's toolbox, and it has changed the way the telecommunication sector produces innovative solutions to technical challenges. In this article, we explore four case studies that highlight SDR as a reliable tool in industry, academia, and government. Specifically, we study four examples that illustrate: advances in low-cost, reliable, and versatile SDR platforms, open source and universal SDR software development environments, powerful technical computing environments employing SDR hardware for real-world experimentation, and educational paradigms for synthesizing digital communications and digital processing concepts using SDR technology. By understanding the impact of these case studies, we intend to provide some insight on how the SDR revolution has changed the way the world designs and implements telecommunication systems.

20 Years of SDR: A Revolution in the Making

Software defined radio, or SDR, has increasingly captured the attention of the telecommunications sector over the past several decades with its promise of rapid design cycles, flexible real-time operations, reusable hardware for different transceiver implementations, ease of manufacturing, and upgrading, and accessibility to many communication system engineers, technologists, and researchers. This SDR vision, which has revolutionized the telecommunications sector over the past 20 years, has been fueled by significant advances in digital processing technologies, analog-to-digital and digital-to-analog converters, software tools, and radio hardware. Consequently, SDR technology is finally beginning to fulfill its promise, and become a mainstream, powerful, and accessible communication system and network prototyping tool.

In fact, at the time of the writing of this article, many companies, research laboratories, and universities are using SDR to support a wide range of communications-related activities.

Although SDR technology possesses significant potential to make communication system prototyping more efficient and accessible to the telecommunications community, there were several challenges that needed to be resolved in order to realize this potential. In order to enable the widespread use of SDR technology for prototyping communication systems and networks, the following conditions need to be achieved:

- *Affordable SDR hardware platforms* that possess sufficient computational horsepower, operate across a wide range of carrier frequencies and bandwidths, and are portable
- *Availability of SDR software development environments* that provide the communication technologist with a high level of control of the SDR platform, a rich set of modules, algorithms, and features, and a substantial level of accessibility to the software (i.e., shallow learning curve)
- *Support between SDR hardware and powerful technical computing software* that enables communication technologists to use reliable software models and tools in experiments using real-world SDR hardware in real time
- *Established SDR-based engineering undergraduate curricula* that introduce hands-on SDR design, prototyping, and experimentation to the next generation of communication technologists

Fortunately, the latest advances in SDR technology have recently achieved these conditions, making SDR more accessible to the telecommunications sector for use in prototyping communication systems and networks.

In this article, we present four case studies that highlight the SDR revolution by focusing on how it was fueled by advances in SDR hardware platforms, SDR software development environments, technical computing software solutions for SDR, and undergraduate educational pedagogy using SDR systems. It is expected that this article will provide the reader with a better understanding of current SDR technology, the many different layers that constitute these complex systems, and their capabilities with respect to prototyping communication systems.

The rest of this article is organized as fol-

Alexander M. Wyglinski is with Worcester Polytechnic Institute; Don P. Orofino is with MathWorks Inc.; Matthew N. Ettus is with Ettus Research; Thomas W. Rondeau is with Rondeau Research LLC

	Ettus USRP N200/N210	ZedBoard w/ Xilinx Zynq-7000 FPGA & AD-FMCOMMS5-EBZ	NooElec NESDR Mini SDR USB Stick	Ettus USRP E300
Interface to host Computer	Gigabit Ethernet	Dual FMC Connectors	USB 2.0	AXI4-MM interface to an embedded dual-core ARM Cortex-A9 processor
RF front-end Instantaneous bandwidth RF frequency coverage	USRP daughterboards 25–50 MHz DC to 6 GHz (determined by daughterboard)	Integrated RFIC 56 MHz 70 MHz to 6 GHz	Integrated RFIC 3.2 MHz 24 to 1766 MHz	Integrated RFIC 56 MHz 70 MHz to 6 GHz
MIMO	1 × 1 per unit, up to 8 × 8 using multiple units	4 × 4	N/A	2 × 2
Full duplex	Yes	Yes	Rx only	Yes
ADC	Dual 14-bit 100 MS/s	Dual/quad 12-bit 61.44 MS/s	8-bit 3.2 MS/s	Dual/quad 12-bit 61.44 MS/s
DAC	Dual 16-bit 400 MS/s	Dual/quad 12-bit 61.44 MS/s	None	Dual/quad 12-bit 61.44 MS/s
FPGA RFNoC-compatible Cost	Xilinx Spartan 3A DSP No $$	Xilinx Zynq-7000 No $$$	None No $	Xilinx Zynq-7000 Yes $$$$

Table 1. Family of available SDR platforms.

lows. We provide an overview of commonly used SDR platforms, including the popular Universal Software Radio Peripheral (USRP) by Ettus Research LLC. Software development tools used to prototype SDR platforms are an essential element of the SDR revolution with respect to implementing digital transceiver designs and algorithms. Consequently, we examine several such tools, including the widely used GNU Radio project. We provide an overview of technical computing software used in enabling over-the-air experimentation via SDR technology, with a focus on the MATLAB technical computing software environment. Given all of these advances, the transformation of engineering undergraduate pedagogy with respect to digital communications by using SDR platforms for prototyping designs and algorithms is covered. Finally, several concluding remarks are given.

Building an SDR Platform for the Masses

The revolution in SDR hardware has always been tightly coupled to advances in computing technology, analog-to-digital converters (ADCs), and digital-to-analog converters (DACs). Specifically, the ability of SDR hardware to satisfactorily serve as a reliable interface between the digital world of bits, packets, and user applications and the surrounding analog RF environment of electromagnetic spectrum significantly depended on how quickly sample and continuous signal information can be converted between the two domains, as well as how fast the sample signal information can be processed. Although the technology for enabling SDR hardware has been around for decades, the cost of this technology only recently became affordable to most of the wireless community. Platforms such as the WARP Radio [1], various NUTAQ systems (formerly Lyrtech), and the Universal Software Radio Peripheral (USRP) by Ettus Research (now part of National Instruments) [2] have enabled real-time over-the-air experimentation in areas such as:

- Wireless networking
- Spectrum monitoring
- Dynamic spectrum access
- Global System for Mobile Communications (GSM), wideband code-division multiple access (WCDMA), and Long Term Evolution (LTE) mobile telephony base stations
- RADAR
- Radio astronomy and RADAR astronomy
- Wildlife tracking
- RF test equipment
- Magnetic resonance imaging (MRI)
- Motion tracking
- Radio navigation and global positioning
- Satellite communications
- RF identification (RFID)
- Wireless security research

To obtain some insight into the capabilities of present-day SDR platforms, let us study a range of systems and their technical specifications as shown in Table 1. The Ettus USRP N210 (Fig. 1a) is a widely used modular SDR platform consisting of an RF front-end (RFFE), a field programmable gate array (FPGA), and a general-purpose processor (GPP). This SDR platform is based on the initial framework devised by Matt Ettus in 2003 when he started work on the USRP in order to help lower the barrier of entry to SDR. The RFFE is tasked with bridging the digital world of sequences of timed samples and the antenna via a direct conversion architecture. A combination of low noise amplifiers (LNAs), switches, variable attenuators for gain control, local oscillators (LOs), and lowpass filters, coupled with ADCs and DACs, enabled the interfacing between the analog RF domain and the baseband digital domain. The FPGA performs all of the high-speed baseband digital signal processing (DSP) operation on the samples coming in and out, as well as all precision timing and synchronization functions, which allow

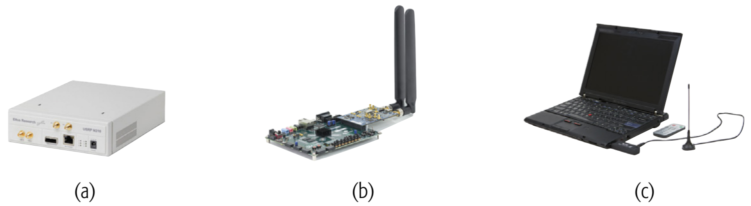

Figure 1. Three examples of widely-used SDR platforms: a) Ettus Research USRP N210 Platform; b) ZedBoard with Xilinx Zynq-7000 FPGA with analog devices FMC-based RF I/O; c) NooElec NESDR Mini SDR USB stick.

for such capabilities as time-division duplexing (TDD), time-division multiple access (TDMA), and multiple-input multiple-output (MIMO) operations across multiple devices. Finally, the FPGA contains the logic to interface with the GPP. To abstract away the interfacing and control of the USRP devices and instead present the user with a small set of primitives that can be used to build real-time communication systems, the USRP Hardware Driver (UHD) was developed in order to provide a single application program interface (API).

The combined ZedBoard-based Xilinx Zynq-7000 FPGA with AD-FMCOMMS5-EBZ SDR platform (Fig. 1b) is a joint venture between Analog Devices and Xilinx to provide a high-performance SDR evaluation system for the wireless community. The system-on-chip (SoC) solution available on the Xilinx Zynq-7000 FPGA enables powerful, versatile computational performance derived from both the onboard ARM processor and FPGA fabric. Thus, the Zynq-7000 can support a wide range of digital functions on this single platform. The AD-FMCOMMS5-EBZ from Analog Devices is the fifth generation of a family of high-speed analog RFFEs designed to showcase the latest generation of high-speed data converters, especially compute-intensive FPGA-based radio applications. In particular, the AD-FMCOMMS5-EBZ is designed around the AD9361 2×2 RF Agile Transceiver, which is capable of supporting instantaneous bandwidths of up to 56 MHz across 70 MHz to 6 GHz. As opposed to the USRP N210 SDR platform, the Xilinx Zynq-7000 FPGA/AD-FMCOMMS5-EBZ SDR platform is a standalone solution that does not require a GPP-based host (a Linux-based operating system can be supported on the Zynq-7000). To support standalone applications, Ettus Research released the USRP E300, which uses much of the same components as the Xilinx Zynq-7000 FPGA/AD-FMCOMMS5-EBZ SDR platform and possesses similar specifications and performance characteristics.

From a performance perspective, the USRP N210, Xilinx Zynq-7000 FPGA/AD-FMCOMMS5-EBZ, and USRP E300 SDR platforms are all very capable systems that can implement a wide range of solutions for different applications. On the other hand, these solutions range in cost from hundreds to thousands of dollars, which might be prohibitively expensive for relatively simple applications, such as satellite communication signal reception, wireless spectrum sensing, or applications requiring numerous SDR platforms. Consequently, the SDR market has also witnessed the advent of numerous low-cost low-complexity SDR platforms such as the NooElec NESDR Mini SDR USB Stick (Fig. 1c). On the order of tens of dollars, these simple SDR receivers plug into the USB port of a laptop computer and perform a wide range of operations based on the available software packages installed on the host computing platform.

As described previously, the USRP E300 SDR platform is a highly capable standalone wireless system that can support sophisticated digital communications and DSP functionality. The RFFE based on the AD9361 enables the USRP E300 to support wireless communications across a large part of the frequency spectrum, while both the ADC and DAC accurately interface the analog world with the digital baseband domain of the FPGA and the ARM processors on the Zynq-7000. Until recently, efficient and effective utilization of SDR computing hardware resources by the larger wireless community has always been a key technical challenge that prevented the widescale use of SDR within the wireless sector. Nevertheless, a new programming system called the RF Network-On-Chip (RFNoC) hopes to remedy this situation.

RF Network-On-Chip

RFNoC is a new programming system for FPGAs developed at Ettus Research with the goal of easing large SDR designs in FPGAs. This architecture allows users to easily integrate custom modules, such as modulators, demodulators, processors, and protocol stacks, without having to become experts on FPGA design. The basic concept behind RFNoC is that rather than treating the entire FPGA as a single monolithic sea of gates, users instead operate with a network of functional units called computation engines (CEs). This network dramatically reduces the complexity of large designs and allows for the dynamic runtime flexibility that many applications, especially cognitive radios, require.

Each computation engine has the exact same interface to the network, so they are easily interchangeable. This network of computation elements can scale across multiple FPGAs, including those of different types, and this allows for portability of CEs across all of the third generation USRP devices. It also makes dynamic reconfiguration of the FPGA a much easier task and makes it much easier to meet timing requirements in the FPGA.

The network fabric (a crossbar switch) connects the computation engines, radios, and external network interfaces. Figure 2 shows the FPGA internals of an X300 device, with external 10G Ethernet and PCIe interfaces. An E300 device would look the same, but instead of Ethernet and PCIe, it has interfaces to its on-chip ARM CPU.

The network exists inside the FPGA, but it also transparently routes across multiple FPGA devices that can be connected directly or through Ethernet switches and the like. This allows the user to easily create large systems consisting of many FPGAs and many MIMO-synchronized radios easily.

Open Source Solutions for SDR Prototyping

As discussed in the previous section, suitable digital/computing hardware combined with versatile RFFEs are some of the necessary ingredients for enabling the widescale adoption of SDR technology by the wireless community. Nevertheless, the implementation of a communication system also requires sufficient software support and tools that are accessible to the user in order to enable the rapid development and evolution of ideas and designs. A variety of open source SDR software projects exist in order to meet the needs of the community, including OSSIE [3], CubicSDR, ALOE [4], and the widely used GNU Radio project [5]. In order to obtain better understanding of these open source software frameworks for building and studying communications systems, we explore the revolution of open source software for SDR by examining the history and capabilities of GNU Radio.

GNU Radio

In 2001, Eric Blossom started the GNU Radio project with the goal of providing a framework for building SDR applications with free software. It has attracted a large community of users and developers from around the world, and has become the design environment of choice for much research in the area of SDR and cognitive radio. GNU Radio works by using a pluggable architecture, where blocks of signal processing algorithms are placed together into a graph such that samples flow through the graph, with each block operating independently on these samples to produce the radio application. With this concept of "drag-and-drop" signal processing, designs can easily be modified. Furthermore, GNU Radio comes with a variety of graphical plotting tools and simulation tools such as various channel models in order to provide straighforward ways to simulate and observe the behavior of a new design. In order to understand how software tools interface with the SDR hardware, as well as how information is passed between the hardware and software domains, let us study how GNU Radio operates in handling the design and implementation of various communication systems.

Development Models: When it comes to building communication and other signal processing systems, GNU Radio has a few modes that help enable the movement of data through the system. An intuitive way of handling the signal processing is to move the data as a stream of samples. GNU Radio has supported the streaming model since its inception, and this solves a number of communications issues. However, as interest in packet-based communications has been steadily increasing, handling packets as an infinite stream of samples has become problematic, with decreased efficiency experienced when trying to handle packet boundaries within the stream of samples. Consequently, GNU Radio also implements a message passing system, where messages of protocol data units (PDUs), which may represent a packet, frame, fragment, or similar, are transferred as a single unit. Boundary conditions

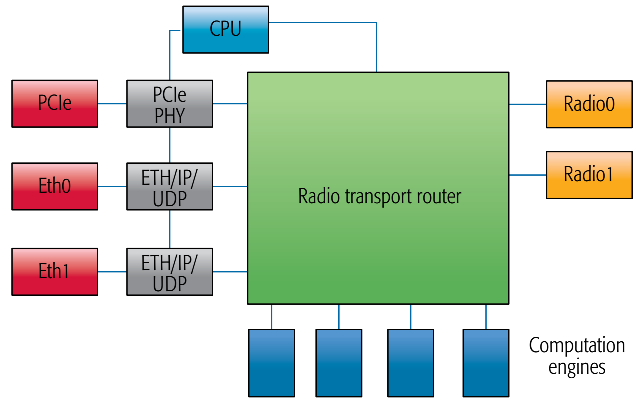

Figure 2. Schematic of the Ettus USRP E300 architecture.

are much more easily handled in this way and thus enable the focus to shift to efficient processing of the data within the PDU. The message passing system is also a useful way to signal and pass meta-data or control data between blocks.

GNU Radio has a third model for moving data around, called the tagged stream system, which is specifically designed for passing meta-data. A block can tag an item in the data stream with information about the sample, such as the time the sample was created, the signal-to-noise ratio measured at that sample, or even information about the state of the system such as the frequency and gain settings of the RFFE. The tags move synchronously with the data and are handled appropriately through sample rate changes. Each of the three models of moving data around have their uses and application spaces, and most of the real-world modems built in GNU Radio use each concept to some extent.

The radio communications research community has a number of problems that it is simultaneously addressing, such as examining existing standards to explore other use cases, improvements and efficiency issues, and security risks. Furthermore, the research community is looking at new models of communications that may or may not be tied to existing systems or methods. Finally, research is ever progressing to address the next level of challenges for wireless data communications, such as the rapidly developing push for 5G standards. GNU Radio itself as a project is not directly interested in specific standards. Instead, GNU Radio develops the architectural framework and API that enables the development and study of these issues through the out-of-tree (OOT) module project concept. On the other hand, there are several software efforts that focus on specific standards, implementations, and general technical computing research efforts, which are covered in the next section.

When SDR Meets Technical Computing Software

Once the hardware device and software framework for an SDR platform has been developed, wireless prototyping of various communication systems by the user can commence. Wireless prototyping is a workflow for the design, verification, and prototyping of radio systems. Traditionally, Barriers to the wireless prototyping workflow are vanishing. Furthermore, the wireless community is at an inflection point in the pursuit of wireless prototyping given the availability of a wide range of affordable SDR platform options.

wireless prototyping required mastery of many distinct tools, languages, and interfaces, with very little tool integration and workflow usability. This posed a high cost of entry and limited wireless prototyping to highly motivated and well funded organizations. Technical computing software solutions, such as MATLAB [6] and SystemVUE [7], enable this workflow; often, model-based design is the preferred approach for design iteration. Technical computing software has made significant progress in recent years, with comprehensive support for cellular, wireless local area networks, and other wireless waveforms, enabling engineers to prototype commercially relevant systems in hours instead of months. Barriers to the wireless prototyping workflow are vanishing. Furthermore, the wireless community is at an inflection point in the pursuit of wireless prototyping given the availability of a wide range of affordable SDR platform options (refer to Table 1).

Wireless Prototyping Workflow

A compelling vision for the wireless prototyping workflow embodies four steps and places significant demands on hardware integration with technical computing software.

System Simulation: The initial step of the workflow executes all algorithms on the desktop in a convenient and interactive fashion with synthesized data. Desktop execution sidesteps the constraints of the target embedded system, making it easier to explore algorithm alternatives, identify execution errors, and tune parameters as a simulation is underway.

RF Integration: The second step configures the RF I/O so that the desktop simulation receives and transmits data using the target hardware. This enables the simulation to include sensor noise, quantization, fading, and power levels that typically influence design decisions. RF signals can be recorded for use during repeated testing and verification using real-world data, or streamed to the desktop in real time. This enables desktop testing of wireless systems and typically implicates multicore, GPU, and other acceleration techniques that increase simulation throughput.

Incremental Deployment: The third step generates code for elements of the design, replacing desktop simulation with streaming execution on target hardware. The highest-rate elements in the front-end of the radio are typically moved to target hardware first, while the rest of the design remains on the desktop. For an SDR, these elements are often destined for execution on an FPGA. Additional elements are transitioned to and validated on the target hardware iteratively, using bit error rate (BER) or other quality metrics. This step places high importance on automatic code generation and data transfer between desktop and target hardware.

System Validation: The final step executes the design on the target hardware and validates it for correctness relative to simulation results obtained in the first step. Synchronous (gated execution) and asynchronous (full-speed streaming) execution of hardware can be employed, using techniques such as FPGA-in-the-loop, to gain confidence in final system operation.

This workflow enables practitioners to minimize "time to next insight" throughout system design. The productivity afforded by each step has proven significant on its own, based on feedback from academic, research, and commercial wireless prototyping communities for the MathWorks technical computing environment. Gaining productivity from a technical computing environment should not require adopting the entire workflow.

An unexpected benefit of pursuing SDR with a technical computing platform is the use of RF modeling tools integrated into the platform. Predicting the imperfections of RF hardware is difficult, due in part to the complexity of nonlinear RF transceivers and the impact of antenna arrays; for example; the design of compensation and equalization algorithms requires simulation. RF modeling tools enable higher fidelity simulation of the complete wireless system.

An Example of Wireless Prototyping: 4G LTE Cellular Communications

As an illustration of how straightforward it is to perform the wireless prototyping workflow, let us consider the example of migrating a fourth generation (4G) LTE implementation from the MATLAB technical computing software environment to an actual SDR hardware platform. Referring to Fig. 3, the LTE test waveform is synthesized using an LTE resource grid that is initially designed and then analyzed in MATLAB. This figure also summarizes many details within this test signal. The LTE waveform is then transmitted using an SDR platform consisting of a Xilinx Zynq-7000 FPGA with an Analog Devices FMC-based RFFE, received by a second SDR platform, and processed using MATLAB in order to analyze the equalization of the physical downlink control channel (PDCCH) symbols and channel magnitude frequency response.

Note that throughout this experience, the communication systems engineer did not need to leave the MATLAB technical computing software environment. Instead, once an implementation for the 4G LTE system has been verified via computer simulation, it can readily be applied to actual SDR hardware for experimentation. Thus, this example highlights the capability of SDR technology with technical computing software. As shown in the following section, this functionality can also be leveraged as a powerful educational tool for teaching wireless communications.

Project-Based Learning with SDR

SDR technology has become a viable instructional resource for the teaching of undergraduate courses in digital communications. The relatively low cost of the SDR hardware can fit within the equipment budget of an academic department. Moreover, the increasing amount of software support for SDR systems, the reliability of SDR/software integration, and the growing familiarity of undergraduate students with the array of available technical/scientific software tools have all contributed to a decrease in the learning curve associated with the prototyping of communication systems using SDR. Finally, the availability of SDR-based instructional resources for educators in the form of ready-to-use models

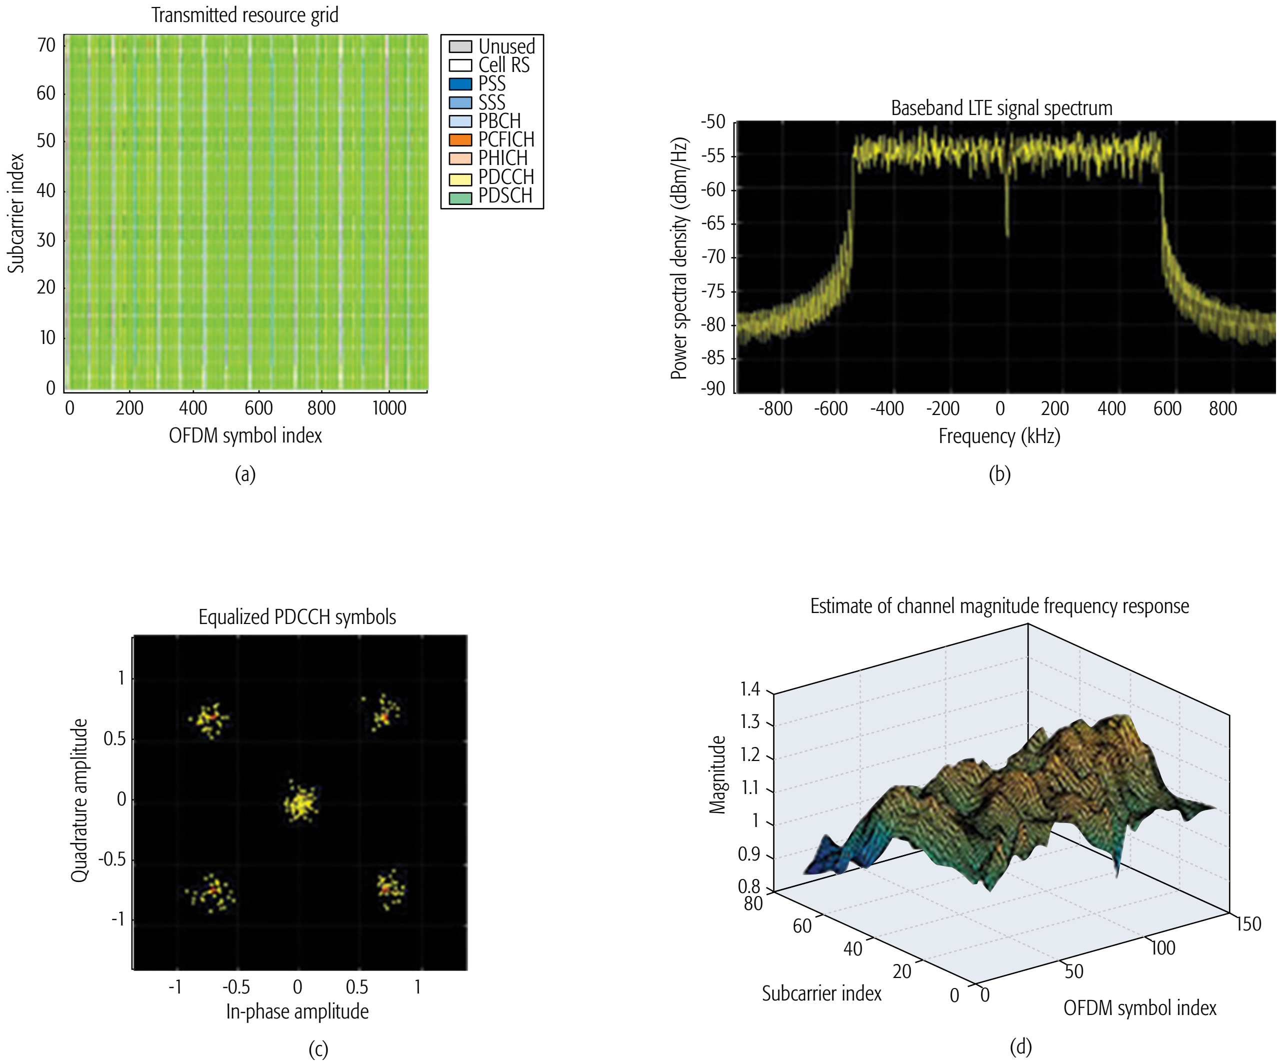

Figure 3. Results obtained using MATLAB, LTE system toolbox and Zynq SDR support from the communications system toolbox.

and demostrations, laboratory guides, and textbooks has enabled the introduction of SDR into the classroom [8].

Although the technology and resources for SDR-based education are finally becoming widely available to both instructors and students, it is very important to understand the context in which SDR can be deployed within a classroom environment in order to maximize the educational experience. Excellent coverage of lessons learned via the pedagogical usage of SDR technology at Penn State, Worcester Polytechnic Institute (WPI), the United States Naval Academy, Indiana University/Purdue University Fort Worth, University of Utah, and Virginia Tech was presented in a 2014 Feature Topic on Education of *IEEE Communications Magazine* [9]. Based on these different university experiences, the use of SDR in the classroom as a tool for teaching, reinforcing, and synthesizing concepts in digital communications can be decomposed into three steps (Fig. 4):

Step 1. Sample-based perspective of digital communications

Step 2. Insights on the fundamental building blocks of these systems

Step 3. Open-ended communication system design experiences

Understanding the Analog/Digital Divide

Most approaches for teaching undergraduate digital communications focus on the study of the transceiver from a binary perspective that gradually works toward the RFFE. Conversely, when it comes to working with SDR hardware, the key challenge in getting a digital communication system working is understanding how the ADC and DAC operate, since without the correct samples nothing else will function properly [10]. Thus, it has been observed that a digital communications course using SDR yields the best outcomes when students start the course from a sample-based perspective. Thus, students need to understand the functions of both the ADC and DAC, decimation and interpolation, and other discrete time signals and systems fundamentals.

Several key concepts with respect to sampling covered in SDR-based undergraduate courses include the following.

Signal Bandwidth: Given that the ADC and DAC on the SDR platform often operate at a fixed sampling rate, students need to understand

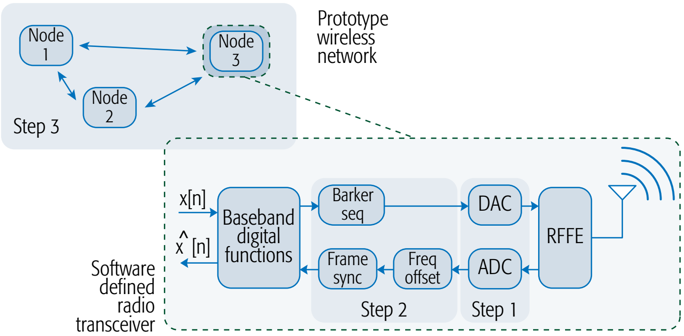

Figure 4. Illustration of the educational paradigm employed in teaching digital communications using SDR hardware.

how the digital data can be interpolated to or decimated from a target sample rate.

Appropriate Sampling Instants: Students must understand the physical issues associated with the ADC in terms of sampling at the correct instants, as well as performing decimation that results in the target samples being discarded.

Sampling Underflow and Overflow: There are several data bottlenecks in the overall setup of the SDR system, such as Gigabit Ethernet and available processing power. Thus, students are exposed to scenarios when too much data is received by an SDR system such that some of the information is lost (i.e., overflow) as well as when not enough data is provided to the USRP N210 such that gaps appear in the transmission (i.e., undeflow).

Open-Ended Design Experience

Once the sampling concept and fundamental modules needed for the construction of a functional digital communication transceiver have been covered by an SDR-based course, this provides several educational opportunities for presenting concepts in medium access control and wireless networking. One approach that could be employed is the *open-ended course design project*, where students work in teams of two or three on a task with specific objectives but loosely defined constraints on how to achieve them. The purpose of these projects is to synthesize the concepts already taught in class, combined with the aforementioned modules, in order to pursue a project-based experience that also promotes teamwork and hands-on learning in order to yield a real-world solution to a problem.

Step 3 of Fig. 4 shows how communication nodes can serve as the building blocks for relatively complex network architectures with sophisticated medium access protocols, such as:

- Wireless ad hoc networks capable of bootstrapping from scratch
- Jamming-resistant multihop wireless networks
- Scaled-down cellular networks

These kind of projects expose students to the challenges of coordinating multiple wireless transceivers in order to perform a variety of different operations. Some of these challenges include bidirectional communications, establishing contact between two wireless nodes, medium access control, and timing of network operations. Consequently, students obtain substantial insight into a wireless network consisting of digital communication transceivers from the bottom up. In particular, these projects using actual SDR hardware to prototype communication systems and wireless networks have the ability to help synthesize and reinforce digital communication concepts while exposing students to real-world issues encountered during transceiver and network prototyping.

Conclusion

Advances in SDR technology have revolutionized the way the telecommunications sector conducts research, development, and educational activities. Low-cost, accessible, and reliable SDR hardware coupled with open source SDR development environments and powerful technical computer software capable of interfacing with SDR platforms have significantly transformed the way we all think of prototyping new communication systems and networks. With hands-on SDR-based communications and networking pedagogy being introduced in engineering undergraduate curricula, the skill set needed to wield these SDR tools is becoming more widely available among the next generation of telecommunication technologists. Twenty years ago, many of the advantages and capabilities of SDR technology that we take for granted today were unrealizable. Given the rate at which advances are being made in this sector, it is expected that this revolution in SDR technology will continue for another 20 years.

References

[1] P. Murphy, A. Sabharwal, and B. Aazhang, "Design of WARP: A Wireless Open-Access Research Platform," *Proc. 2006 14th Euro. Signal Processing Conf.*, Sept. 2006, pp. 1–5.
[2] R. Dhar *et al.*, "Supporting Integrated MAC and PHY Software Development for the USRP SDR," *Proc. 1st IEEE Wksp. Networking Technologies for Software Defined Radio Networks*, Sept. 2006, pp. 68–77.
[3] M. Robert *et al.*, "OSSIE: Open source SCA for Researchers," *Proc. SDR Forum Tech. Conf.*, vol. 47, 2004.
[4] I. Gomez, V. Marojevic, and A. Gelonch, "Aloe: An Open-Source SDR Execution Environment with Cognitive Computing Resource Management Capabilities," *IEEE Commun. Mag.*, vol. 49, no. 9, Sept. 2011, pp. 76–83.
[5] E. Blossom, "GNU Radio: Tools for Exploring the Radio Frequency Spectrum," *Linux J.*, vol. 2004, no. 122, p. 4, 2004.
[6] J. Proakis, M. Salehi, and G. Bauch, *Contemporary Communication Systems Using MATLAB*, Cengage Learning, 2012.
[7] D. Silage, *Digital Communication System Using System VUE*, Firewall Media, 2006.
[8] D. Pu and A. M. Wyglinski, *Digital Communication Systems Engineering with Software-Defined Radio*, Artech House: Norwood, MA, USA, 2013.
[9] S. G. Bilen *et al.*, "Software-Defined Radio: A New Paradigm for integrated Curriculum Delivery," *IEEE Commun. Mag.*, vol. 52, no. 5, 2014, pp. 184–93.
[10] R. G. Machado and A. M. Wyglinski, "Software-Defined Radio: Bridging the Analog-Digital Divide," *Proc. IEEE*, vol. 103, no. 3, 2015, pp. 409–23.

Biographies

Alexander M. Wyglinski [S'99, M'05, SM'11] (alexw@wpi.edu) received his B.Eng. and Ph.D. degrees from McGill University, Montreal, Quebec, Canada, in 1999 and 2005, respectively, and his M.Sc.(Eng.) degree from Queens University, Kingston, Ontario, Canada, in 2000, all in electrical engineering. He is an associate professor of electrical and computer engineering with Worcester Polytechnic Institute, Massachusetts, and the director of the Wireless Innovation Laboratory (WI Lab). 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, the Naval Research Laboratory, the Office of Naval Research, the Air Force Research Laboratory Space Vehicles Directorate, 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.

DON P. OROFINO is director of engineering for signal processing at MathWorks, leading product development for the analysis, design, and implementation of signal processing, communications, and computer vision systems. He holds 18 patents related to engineering system simulation and automatic code generation. Prior to joining MathWorks in 1995, he worked at Hewlett-Packard Imaging Systems (Andover, Massachusetts). He earned a Ph.D. in Electrical and Communications Engineering at Worcester Polytechnic Institute.

MATTHEW N. ETTUS is a core contributor to the GNU Radio project, a free framework for software radio, and is the creator of the Universal Software Radio Peripheral (USRP). In 2004, he founded Ettus Research to develop, support, and commercialize the USRP family of products. Ettus Research was acquired by National Instruments in 2010, and he continues as its president. USRPs are in use in over 100 countries for everything from cellular and satellite communications to radio astronomy, medical imaging, and wildlife tracking. In 2010, the USRP family won the Technology of the Year award from the Wireless Innovation Forum. In the past he has designed Bluetooth chips, GPS systems, and high-performance microprocessors. Before that, he received B.S.E.E. and B.S.C.S. degrees from Washington University, and an M.S.E.C.E. degree from Carnegie-Mellon University. In 2011, he was named an eminent member of Eta Kappa Nu. He is based in Mountain View, California.

THOMAS W. RONDEAU holds a Ph.D. from Virginia Tech in eectrical engineering, graduating in 2007. He is the current maintainer and lead developer of GNU Radio and is a visiting scholar with the University of Pennsylvania. He also works as a consultant on GNU Radio and wireless technology through his firm Rondeau Research, LLC. His Ph.D. dissertation on aritifical intelligence applied to wireless communications received the Council of Graduate Schools' Distinguished Dissertation for Math, Science, and Engineering. His research interests span areas of communications theory, signal processing, and software design, which are all part of his larger interests in software and cognitive radios.

Metacognition and the Next Generation of Cognitive Radio Engines

Hamed Asadi, Haris Volos, Michael M. Marefat, and Tamal Bose

Much of the previous research on cognitive radio has focused on developing algorithms based on artificial neural networks, the genetic algorithm, and reinforcement learning, each with its pros and cons. In this research, we present a new approach based on metacognition. We believe that the metacognitive framework can be the foundation for the next generation of CRs and further the performance improvements in CR.

ABSTRACT

Much of the previous research on cognitive radio has focused on developing algorithms based on artificial neural networks, the genetic algorithm, and reinforcement learning, each with its pros and cons. In this research, we present a new approach based on metacognition. We believe that the metacognitive framework can be the foundation for the next generation of CRs and further the performance improvements in CR. In this work, we present the elements involved in metacognitive radio, discuss the challenges in their development, present solutions to the challenges along with a possible meta-CR architecture, and show results from our implementation. Each cognitive engine (CE) algorithm has strengths and limitations that make it more suitable for certain operating scenarios (channel conditions, operating objective, available hardware, etc.) than other algorithms. A meta-CE can adapt faster and improve performance by exploiting the characteristics and expected performance of the individual CE algorithms. It understands the operational scenarios and utilizes the most appropriate algorithm for the current operational scenario by switching between the algorithms or adjusting them as necessary.

INTRODUCTION

Conventional radios typically use a fixed set of settings (modulation, channel coding, etc.) selected by their operator. A cognitive radio (CR) automatically adapts its configuration to meet its goal under the operating conditions. CR was first described by Mitola; his ideal CR can optimize its own capabilities and self-determine the operation goal by observing its operator and the environment [1].

To make CR possible, we developed an intelligent agent (IA) called the cognitive engine (CE) that enables the radio to have the desired learning and adaptation capabilities [2]. An IA is a system that senses its environment (operating channel scenario), acts by using a communication method based on its experience, and monitors its own performance to learn its capabilities.

COGNITIVE ENGINE CHALLENGES

Since Mitola's seminal work [1], CEs have been researched extensively, and several algorithms and approaches were developed. That said, there are remaining challenges that our work seeks to address. First, CEs are typically designed around one or two primary algorithms or frameworks. Examples include the genetic algorithm (GA), the case-based reasoning (CBR) framework, and the epsilon-greedy (ε-greedy) exploration technique. The chosen algorithms and their parameters make them more suitable for certain operating scenarios than others. An operating scenario is a combination of the channel conditions and the operating objective (e.g., maximum throughput, maximum energy efficiency). For different operating scenarios, either the existing algorithms' parameters need to be adjusted or new algorithms need to be used. For example, in a low signal-to-noise ratio (SNR) environment a conservative algorithm may be preferred that focuses on energy-efficient methods with low error rates vs. in a high SNR environment where an aggressive algorithm may be preferred that prioritizes high spectrum-efficiency methods. Consider a complex system, such as massive multiple-input multiple-output (MIMO), with numerous modes; the problem of identifying the best[1] adaptation technique for the channel conditions becomes more challenging. The first contribution of the article is the introduction of a framework that facilitates the selection of the right learning algorithms and their parameters based on the operating conditions.

Moreover, the "predictability" of a CE's performance is of paramount importance to the adoption of a CE radio. A radio that provides estimates of its expected performance is easier to certify than a radio that cannot provide this information. "Predictability" means that given the channel environment, the CE can predict its performance range based on its experience. If the CE has significant experience in the channel environment, the predicted performance range will be narrow; otherwise (limited experience), the predicted performance range will be wide. Performance predictability and the suitability of each CE algorithm given the operating scenario motivates the development of automated methods for evaluating a CE's adaptation capabilities and expected performance, which is the second contribution of this article. A CE that can estimate its own performance is preferred because other parts or operators of the system will know what to expect from the CE and plan accordingly.

[1] Best in terms of the time and total data sent until the optimal or a near-optimal communication method is found.

The authors are with the University of Arizona.

The Metacognitive Solution

We use a metacognitive engine (meta-CE) framework that inherently addresses the mentioned challenges. From the outside, a meta-CE appears to be a regular CE, that is, a CE that provides expected performance estimates (through confidence intervals). On the inside, a meta-CE is made of one or more CEs. A meta-CE has the following features:

- Knows at which operating conditions each CE is best
- Knows the expected performance of each CE and in turn knows its own performance when it chooses a certain CE
- Automates the process of evaluating and selecting the appropriate CE for the operating conditions

Considering all possible channel conditions, the meta-CE performs better than only using one individual CE. The meta-CE, as discussed in the sequel, can also provide knowledge indicators (KIs), which measure how much a CE has learned and how close it is to reaching the state at which there is nothing more to be learned (for that specific operating scenario).

This article presents the metacognitive concept in detail with initial innovations and discussion of future directions for maturing the meta-CE framework. First, we present a brief background on CEs and the origin of the metacognition concept. We then present the elements of metacognition along with key results. Finally, we provide a discussion on future work and concluding remarks.

Background

Brief Background on Cognitive Engines

For nearly two decades, CE designers have been continuously working to understand and develop better learning techniques for CR. They have typically borrowed ideas from machine learning and artificial intelligence [3] to design their CEs. Notable examples include artificial neural networks (ANNs), GA, and CBR [3]. In addition, they used reinforcement learning techniques such as the ε-greedy, softmax, and Gittins index [4, 5]. Furthermore, other techniques such as particle swarm [6] and ant colony optimization are also used to create CEs.

Different types of CEs have their own advantages and disadvantages. Some perform really well in high SNR conditions; others are more effective in low SNR conditions. In addition, providing predictable and more confident performance levels is the most important aspect of designing various CE algorithms. Therefore, the meta-CE is being proposed to provide the mentioned level of prediction for various types of CEs. The first effort was Gadhiok *et al.* [7], who proposed a primitive architecture of metacognition. Moreover, we proposed a more advanced and general meta-CE in [8, 9], which can classify various CE algorithms based on the operating conditions (objective, channel condition, radio capabilities, etc.). The proposed meta-CE employs a general performance characterization method to evaluate the performance of individual CE algorithms. Also, the meta-CE can identify distinct operating scenarios based on the performance level of the individual CEs.

More Details on the Metacognition Concept

The motivation for incorporating metacocognition into a CE comes from psychology. Metacognition is defined as thinking about thinking [10]. There are three interrelated components to metacognition: metacognitive knowledge, monitoring, and control [11]. These components are integrated with the primary cognition, which refers to object-level thoughts or an individual CE algorithm's process. *Metacognitive knowledge* derives from the beliefs individual CEs generate about their decisions (e.g., "the CE does not consider all the communication methods when it is facing low power"). *Metacognitive monitoring* is the process by which an agent evaluates its own thoughts for comparison (e.g., "failing to consider all the available communication methods can lead to choosing an inappropriate one for the channel conditions"). *Metacognitive control* refers to the regulation of the agent's thinking (e.g., "if the CE does not have enough power or time to consider all the communication methods, it needs to stay with the most robust communication method it currently knows"). Achtziger and colleagues [12] note that metacognitive monitoring and control processes connect secondary thinking (meta thoughts) with primary thinking about an object; monitoring processes represent information flowing from the object level to the meta level; and control processes represent information flowing from the meta level to the object level. Metacognitive processes serve a self-regulatory function when monitoring permits secondary thinking to inform the state of primary thinking about an object, and the control processing permits primary thinking to be informed by secondary thinking. Research shows [12, 13] that secondary thoughts can play an independent role in judgment and behavior when they modulate (e.g., increase, decrease, or reverse) the impact of primary cognition. An important outcome of metacognition is the confidence that people perceive in their thinking.

The scope of our work is to implement the needed components of metacognition as a meta-CE. Figure 1 shows the connection between the three metacognition components and the operation of the meta-CE: the meta knowledge translates to calculating the KIs that indicate the status of the learning progress, and learning curves and performance characterization that analyze the performance of the individual CEs. Meta monitoring translates to channel characterization, comparing the performance results of the individual CEs, and monitoring the meta-CE's real-time performance. Finally, metacognitive control translates to selecting the most appropriate CE and adjusting its parameters for the current operating scenario. The following sections elaborate on how each concept affects the meta-CE's operation and our current implementation.

The motivation for incorporating metacocognition into a CE comes from psychology. Metacognition is defined as thinking about thinking. There are three interrelated components to metacognition: metacognitive knowledge, monitoring, and control. These components are integrated with the primary cognition, which refers to object-level thoughts or an individual CE algorithm's process.

Elements of Metacognition

Metacognitive Engine Operations

Figure 2 shows that the operation of the meta-CE can be summarized in six steps. First, any available prior information such as results from previous experience and learning outcomes is used to provide initial estimation of the expect-

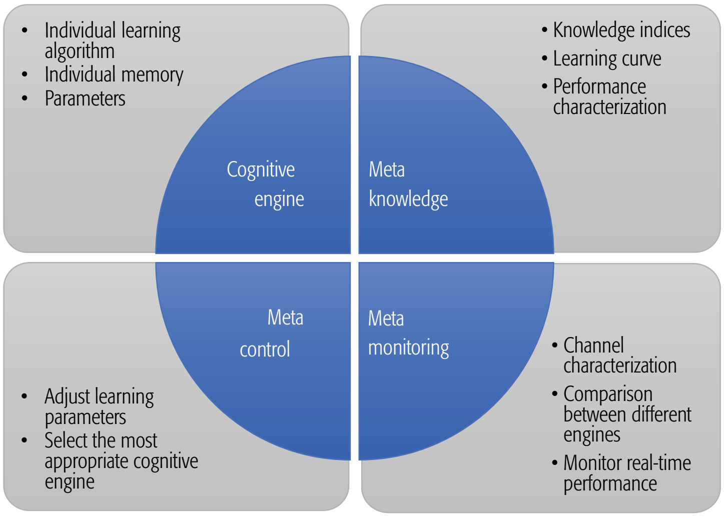

Figure 1. The psychological metacognitive components and their application to a metacognitive radio engine's functions.

ed performance. Second, the performance characterization module uses the prior information and real-time data from the currently selected CE to update the CE's performance and knowledge metrics. Third, the Best Cognitive Algorithm Classification (BCAC) compares the available CEs and identifies the most appropriate CE for each operating region. Fourth, the current operating scenario is updated with the current channel conditions and the radio's operating objective. Fifth, the meta-CE selects the most appropriate CE algorithm for the current operating scenario based on the metacognitive knowledge gathered by the previous steps. Sixth, the selected CE algorithm is used to control the communications link while the operating conditions remain unchanged. Finally, once the CE is done with link adaptation, it feeds its performance data (monitoring) to the CE performance characterization block, and the process repeats.

Following this overview of the meta-CE operations, we now provide more detail on how to realize a meta-CE.

Metacogntive Engine Development

An ideal meta-CE needs to develop the three aforementioned interrelated components that have distinct functionalities (Fig. 1). These components are:

- Meta knowledge
- Meta monitoring
- Meta control

A prerequisite for developing these components is a systematic method for evaluation of performance of CEs [14]. The challenges in this systematic evaluation are due to the following. First, the performance is affected by a multitude of parameters, such as channel parameters, operating objectives, learning techniques, and radio capabilities. Second, the evaluation method must be able to compare performance in different scenarios (channel conditions, objectives, etc.) against each other meaningfully. We have addressed these challenges through utilizing the concept of a "learning curve." A learning curve is a graphic representation of a CE's learning rate.

Third, it is necessary to be able to identify the operating environments (channel) and match them to the performance of the CEs. In theory, there are infinitely many operating conditions (e.g., there are infinite SNR values); however, classification into only a finite number of regions matter (e.g., low SNR, medium SNR, and high SNR). Once the meta-CE can characterize and identify these regions, it can match each region to the CE that performs best in a particular region. We treat this "region to CE matching" process as a classic pattern classification problem, which is BCAC.

A different challenge in the development of meta-CE relates to the metacognitive control component. The role of this component is to make a decision based on the information provided by the other components. An ideal meta-CE should be able to adapt the learning parameters or switch among possible algorithms in order to provide the most effective and efficient way of learning.

Metacognitive Knowledge

Metacognitive knowledge is related to knowledge or belief about an individual CE's decisions and utility. Metacognitive knowledge allows a meta-CE to make sense of performance and understand the state of individual CEs. In order to generate/develop metacognitive knowledge, we need to define metrics and methods that serve this purpose. One of the most utilized indicators is the learning curve. The meta-CE will create each CE's learning curve(s) by observing its operation after numerous operating sessions. The learning curve is characterized by using statistical inference based on the achieved performance by a CE's decisions. The performance is derived from multiple repetitions of the CE algorithm in the same or similar channel conditions.

Figure 3 presents a plot of the learning curves of three CE algorithms. The plot is generated by observing the mean and variance of the achieved rewards of each CE. The reward quantifies how well the operating objective of the CE is met; in this case, the reward is the achieved throughput. Figure 3 shows examples of different CE learning behaviors. However, before explaining these three CEs' algorithms, let us see what kind of information this plot provides for us. For example, CE Algorithm 2 has zero average performance between time steps 0 and 50. Moreover, after time step 180, it quickly reaches 38 Mb/s, which is the maximum possible in this example. On the other hand, CE Algorithms 1 and 3 provide a slower but steadier performance increase. This is because CE Algorithm 2 is more aggressive with risk taking in applying the communication method options. The standard deviation shown by the error bars indicate that CE Algorithm 2 not only has a zero average performance for the first 40 time steps, but also is consistently zero since its standard deviation is also very low. Between time steps 50 and 200 (approximately), its standard deviation increases and peaks, which means that there is a great variability in performance between those steps. This is an intuitive result which captures the fact that the perfor-

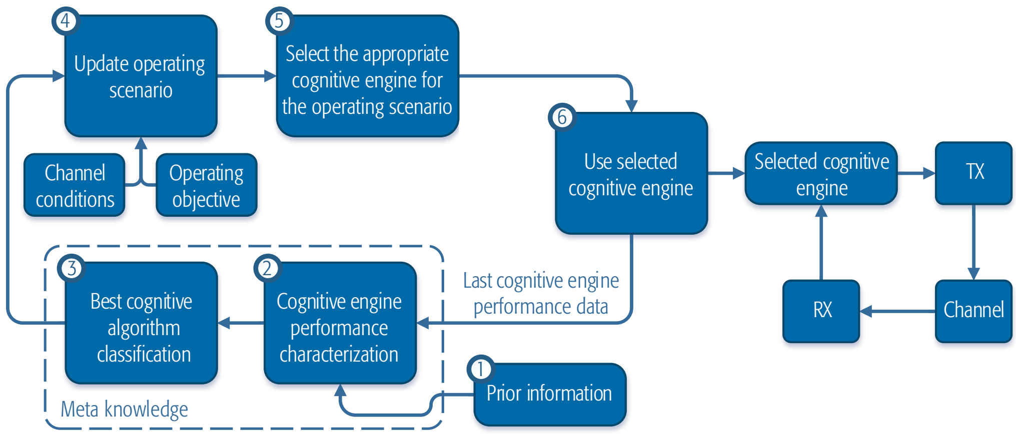

Figure 2. The metacognitive engine operation graph.

In order to generate/develop metacognitive knowledge, we need to define metrics and methods that serve this purpose. One of the most utilized indicators is the learning curve. The meta-CE will create each CE's learning curve(s) by observing its operation after numerous operating sessions.

mance becomes non-zero and reaches near maximum somewhere between the 50th and 200th time step.

Two of the three CE algorithms shown in Fig. 3 are based on the ε-greedy strategy, and the third one is based on the Gittins index. The ε-greedy strategy randomly explores the different communications methods with probability ε and uses the communication method with the highest average throughput with probability 1-ε. The ε parameter controls how aggressively exploration is performed. A higher ε will cause preference for exploration over using existing data. It may result in arriving at an optimal or near-optimal option faster. However, the high exploration rate may also cause reduced overall returns because of the higher exploration cost.

CE Algorithm 1 uses an ε = 0.05 and does not utilize a priori estimates of the maximum possible throughput. CE Algorithm 2 is based on the Gittins index strategy [15] with a normal reward process and a discount factor equal to 0.9. The strategy is simply to use the method with the highest Gittins index, which is based on the reward statistics of each method and must be estimated only when those statistics change (i.e., only when each method is used). Gittins proved that exploration vs. exploitation can be optimally balanced using a dynamic allocation index-based strategy. This strategy maximizes the total sum of rewards collected over a long-term horizon. CE Algorithm 3 uses a more aggressive ε = 0.5; however, ε decreases by 0.001 at each time step. This allows CE Algorithm 3 to decrease the exploration rate as more information is collected about each communication method.

Although observing outcomes of a CE algorithm provides much information about how it learns, the actual amount of knowledge that is obtained by a CE algorithm can vary. For instance, if we achieve a wireless link of 10 Mb/s, we may assume that we have an excellent outcome; however, if we know that the actual potential of this link is more than 100 Mb/s, we know that there is much work to do; we will have to either achieve 100 Mb/s or determine that it is not practically possible. Hence, we need some metrics to show us the true amount of knowledge for each CE algorithm.

In a previous publication [9], we proposed several KIs to estimate the amount of knowledge obtained by different CE algorithms. Different types of KIs focus on various aspects of the learning process. Some KIs reflect the amount of trials attempted by the algorithm, and some KIs depict the level of achieved performance in respect to the maximum achievable performance. To be able to have a general estimation of the amount of obtained knowledge, we utilize the concept of information entropy. This new KI behaves similarly to the CCI indicator presented in [9]. The meta-CE observes the information entropy of the expected performance for each communication method. The variation of the entropy represents the changes in the learned information by a CE. The lower the entropy, the more knowledge is already obtained by a CE algorithm.

Figure 4 illustrates the amount of knowledge obtained by our CE algorithms. It is shown that the level of knowledge CE Algorithm 2 obtained until time step 50 (approximately) matches CE Algorithm 3; however, they have drastically different performance (as shown in Fig. 3). Although CE Algorithm 3 appears to have a performance edge, the obtained knowledge suggests that both CE algorithms significantly improved. This was evident at the 150th time step where CE Algorithm 2 outperformed CE Algorithm 3.

Metacognitive Monitoring

Metamonitoring is the ability to distinguish among different communication scenarios, and use the metacognitive knowledge gained to dynamically determine the best CE algorithm and parameters for communicating in the current scenario. A meta-CE needs to distinguish between the operating scenarios in order to match the obtained metacognitive knowledge of the CE algorithms to the operating scenarios. The goal is to find the best cognitive algorithm for each operating scenario. This is achieved by BCAC in our implementation. Our approach is to facilitate BCAC by using two different classification techniques: offline classification using a support vector machine (SVM) and online classification using *k*-means clustering. Both algorithms rely on a distance function between the feature points. For classification, we use selected features from the channel scenarios.

To classify operating scenarios based on the performance of CE algorithms, we need the ability to distinguish and identify various operating scenarios. In order to achieve this, the meta-CE should extract distinctive features of operating scenarios. Generally, an operating scenario is characterized by the channel conditions, operation objective, and the number of possible communication methods the radio's hardware supports.

In machine learning and statistics, classification is used to group data into sub-populations according to their similarities. This is usually done using a training data set that contains examples belonging to known categories. In our case, a training set is a set of data used in various areas of channel scenarios to discover potentially predictive relationships.

The offline and online methods have their own advantages and disadvantages; their main difference is that the offline method waits until all training data are collected in order to process them, and the online method processes the data as they arrive. Furthermore, by preprocessing the offline training data we are able to create examples of channel conditions and the appropriate CE to be used in those channel conditions. This step makes the offline method a supervised method since it is trained using examples of the desired relationships. On the other hand, the online method is unsupervised since it is not possible to identify the desired relationships until after the data are processed. The offline method is more accurate if the training data cover all the likely scenarios. However, if new scenarios arise, the offline method has to wait until a new batch of data is collected and the classification results are updated. The online method starts adapting as new data become available and gets more accurate over time. Finally, the offline method shifts the computation cost to times when the meta-CE may not be needed and perhaps when the radio is connected to a power source or can offload that processing to the network or cloud. On the other hand, the online method consumes computing and power resources on an ongoing basis.

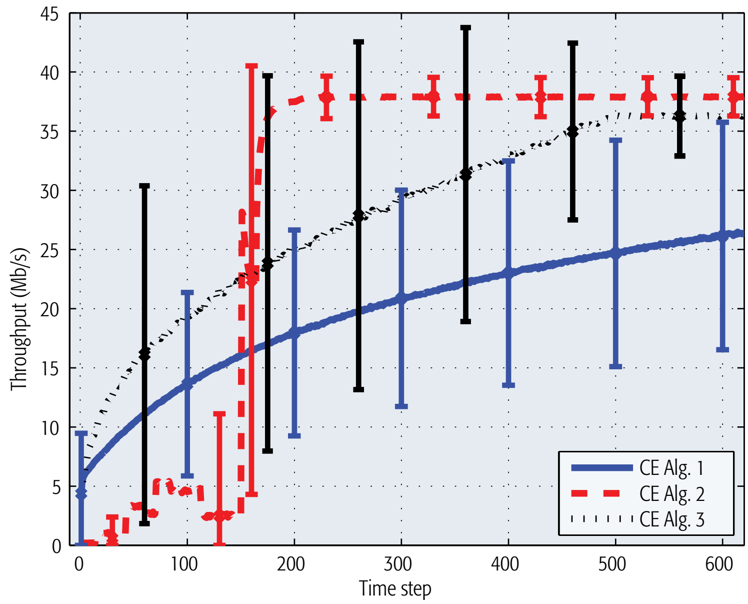

Figure 3. Learning curves based on mean and standard deviation.

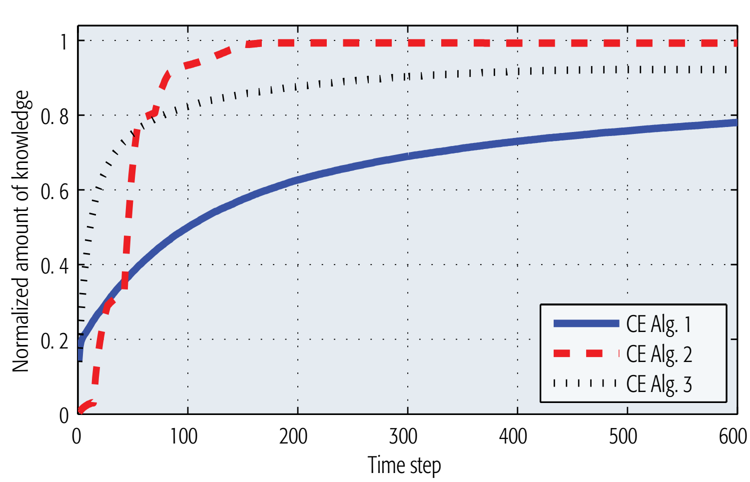

Figure 4. Learning curve based on the estimated amount of knowledge.

The following provides more detail on our BCAC implementation. For the offline supervised learning algorithm, the training process works by first mapping data to a feature space so that these data points can be categorized. Then the algorithm finds a function that maps the input data points to distinct classes. In our simulations, we use a predefined training dataset of 200 different operating scenarios. Then the meta-CE uses the classifier's mapping and the features of the ongoing channel conditions to assign the most appropriate CE algorithm for operation. For online classification, generally there is no data available at the start unless prior information and results are available that can be used to initialize the classifier. When no prior information or results are available, at the beginning the meta-CE chooses a CE algorithm randomly (uniformly distributed) among the available algorithms. As the CEs start to operate, the meta-CE collects information about their behavior. The meta-CE compares each operating scenario with similar previously experienced scenarios. The scenarios' similarity is determined by the mahalanobis distance of the channel scenario feature vectors.

The metacognitive monitoring component operates by regularly monitoring the real-time performance of the selected CE algorithm. Therefore, the classification algorithms used to categorize the operating scenarios are continuously updated. The metacognitive control component uses the current monitoring information to make its critical decision for switching among the available CE algorithms or adjusting the learning parameters of a particular CE algorithm.

Metacognitive Control

The metacognitive control component refers to the regulation of the learning process. This regulation can be done by changing the exploration parameters of the learning algorithms, switching between the different algorithms, or by combining two or more CE algorithms to make a new one. In our implementation, the metacognitive control switches among the available CE algorithms by using the provided information from the other components (metacognitive monitoring and metacognitive knowledge), which is based on current operating scenarios.

For instance, at the very first starting point, the metacognitive control chooses among the CE algorithms randomly since it does not have any information about their performance. However, by collecting more data through the monitoring and meta knowledge components, the metacognitive control identifies the CE algorithm that can provide the best outcome (reward) for the ongoing operating scenario.

Implementation Results

To demonstrate the benefits of metacognition in CR, we present a simple example that assumes a 4 × 4 MIMO system with the following modulation options: quadrature phase shift keying (QPSK); 8PSK; 16, 32, 64, 128, and 256-quadrature amplitude modulation (QAM); with error correction rates 1, 7/8, 3/4, 2/3, 1/2, 1/4, 1/6, and 1/8; and multi-antenna techniques: VBLAST, space-time block code (STBC), and maximum ratio combining (MRC). Furthermore, we consider an SNR range of 0–50 dB and the $\log_{10}$ of the eigen spread (κ) of the channel matrix in the range of 0–12. The CR also has 12 channels available with varying SNR and bandwidth (either 1.25 or 2.5 MHz). At each time step, the CR can transmit over only one of the available channels.

In Fig. 5, we show the results from each CE algorithm and the meta-CE algorithm using offline classification for BCAC. The meta-CE algorithm simply selects the CE algorithm that was found to have the best adaptation performance for the current operating scenario. Best performance is defined as the total throughput achieved during the adaptation session. We used a support vector machine (SVM) as an offline classifier to be trained by 200 operating channel scenarios. We assume that the channel conditions remain static for 100 time steps; therefore, each adaptation session's duration is also 100 time steps. As a result, at time step 0, 100, …, 500 the algorithms have to readjust to the new channel conditions. It was found, that the meta-CE selects with a probability of 92 percent the CE algorithm that is better suited to the given channel conditions. Figure 5 depicts the total data transferred for all 500 time steps (we assume that each time step takes 0.1 ms), which clearly shows the benefit achieved by the meta-CE. The meta-CE transfers a total of 900 kbits vs. 700 kbits transferred by the best individual CE algorithm.

In Fig. 6, we compare the performance of meta-CE, when the online BCAC method is used, with the performance of the individual CEs. The operating objective of the CEs in this example is maximizing throughput. We use the regret concept to compare the CEs. Regret is defined as the difference between the expected reward sum (throughput) using optimal decisions and the actual reward sum that each CE algorithm obtains through its decisions.

Here we compare the results from the offline and online BCAC methods (Figs. 5 and 6). With the offline method (Fig. 5), the meta-CE selects the best CE algorithm 92 percent of the time. When the meta-CE uses the online method (Fig. 6), since it begins with no information, at the start its accuracy is 33 percent with a normalized regret equal to 0.39. This means that the performance of the system with the meta-CE at time step 0 is 61 percent (1 – *regret*) of the ideal throughput. As the online meta-CE learns at each time step, its decision accuracy improves and its regret decreases. Therefore, the meta-CE reaches 85 and 95 percent of the ideal performance after 50 and 120 trials, respectively.

Our results demonstrate that selecting a CE algorithm based on the operating conditions using our meta-CE framework has significant performance advantages. This results in more efficient channel utilization, allowing more users to share the same bandwidth.

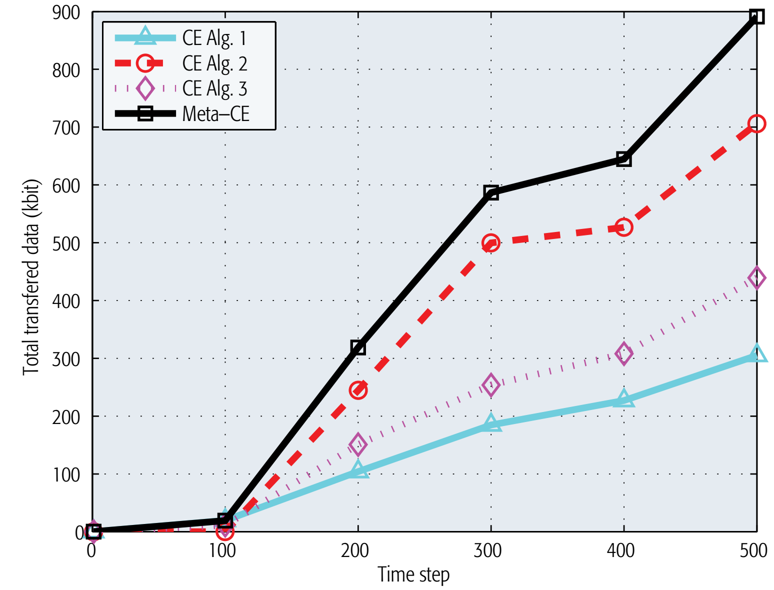

Figure 5. Average metacognitive engine performance (offline classification).

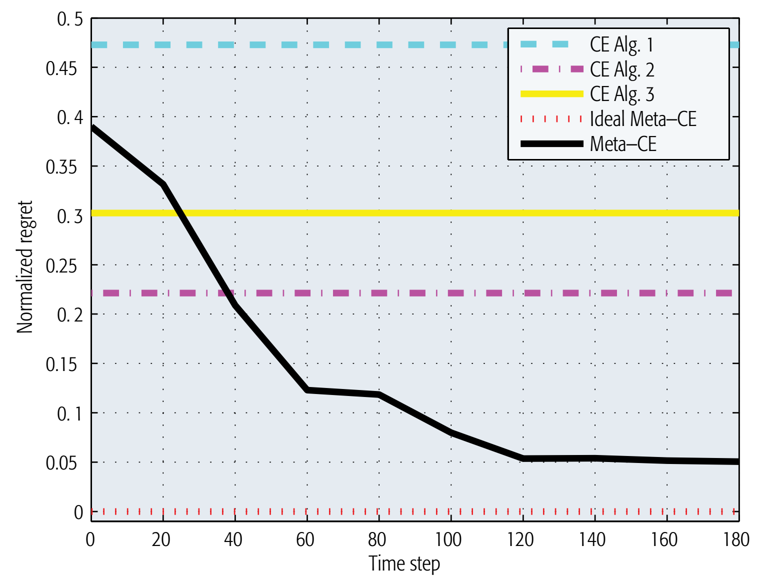

Figure 6. Average metacognitive engine regret (online classification).

A meta-CE inherently solves the problem of characterizing and evaluating the performance of a CE since this process is built-in to the meta-CE. Using this information, a meta-CE can provide estimates of its expected performance. In this work, this information is derived from the performance of individual CEs.

Future Work

Our future plans include the following: first, develop more efficient and mature methods for operating scenario identification and CE selection. The current algorithm classification methods are resource intensive, and their accuracy can be improved. We plan to investigate more appropriate channel metrics to be used for channel scenario identification. Second, assessing the knowledge acquired by a CE and the rate at which that knowledge is acquired is also important. We will therefore work on quantifying the experience level of a CE. Third, we plan to develop a meta-CE that can construct its own CE algorithm by using a set of primitive CE operational elements.

Conclusions

The work presented here shows the evolution from cognitive engines to metacognitive engines for cognitive radio. This evolution can pave the way for another generation of CEs with better performance. A meta-CE inherently solves the problem of characterizing and evaluating the performance of a CE since this process is built into the meta-CE. Using this information, a meta-CE can provide estimates of its expected performance. In this work, this information is derived from the performance of individual CEs. The methods presented in this article are the building blocks for a meta-CE framework, and only the beginning of the development of more advanced and sophisticated techniques and processes for metacognition in radios.

Acknowledgments

This project was partially supported by the Broadband Wireless Access and Applications Center (BWAC); NSF Award No. 1265960.

References

[1] J. Mitola, III, *Cognitive Radio: An Integrated Agent Architecture for Software Defined Radio*, Ph.D. dissertation, KTH, Stockholm, Sweden, May 2000.

[2] S. Haykin, "Cognitive Radio: Brain-Empowered Wireless Communications," *IEEE JSAC*, vol. 23, no. 2, 2005, pp. 201–20.

[3] A. He *et al.*, "A Survey of Artificial Intelligence for Cognitive Radios," *IEEE Trans. Vehic. Tech.*, vol. 59, no. 4, May 2010, pp. 1578–92.

[4] H. I. Volos and R. M. Buehrer, "Cognitive Engine Design for Link Adaptation: An Application to Multi-Antenna Systems," *IEEE Trans. Wireless Commun.*, vol. 9, no. 9, Sept. 2010, pp. 2902–13.

[5] H. I. Volos and R. M. Buehrer, "Cognitive Radio Engine Training," *IEEE Trans. Wireless Commun.*, vol. 11, no. 11, 2012, pp. 3878–89.

[6] Z. Zhao *et al.*, "Cognitive Radio Adaptation Using Particle Swarm Optimization," *Wireless Commun. and Mobile Computing*, vol. 9, no. 7, 2009, pp. 875–81.

[7] M. Gadhiok *et al.*, "Metacognition: Enhancing the Performance of a Cognitive Radio," *2011 IEEE Int'l. Multi-Disciplinary Conf. Cognitive Methods in Situation Awareness and Decision Support*, Feb. 2011, pp. 198–203.

[8] H. Asadi *et al.*, "Metacognitive Radio Engine Design and Standardization," *IEEE JSAC*, vol. 33, no. 4, Apr. 2015, pp. 711–24.

[9] H. Asadi *et al.*, "On Quantifying the Experience Level of a Cognitive Engine," *Proc. SDRWInnComm 2015 Wireless Innovation Conf. Wireless Commun. Technologies and Software Defined Radio*, Mar. 2015, pp. 9–15.

[10] J. H. Flavell, "Metacognition and Cognitive Monitoring: A New Area of Cognitive Developmental Inquiry," *American Psychologist*, vol. 34, no. 10, 1979, pp. 906–11.

[11] J. Dunlosky and J. Metcalfe, *Metacognition*, 1st ed., SAGE Publications, Sept. 2008.

[12] A. Achtziger *et al.*, "Metacognitive Processes in the Self-Regulation of Goal Pursuit," P. Brinol & K. G. DeMarree, Eds., *Social Metacognition*, 2012, pp. 121–39,.

[13] R. E. Petty and P. Briol, "Persuasion: From Single to Multiple to Metacognitive Processes," *Perspectives on Psychological Science*, vol. 3, no. 2, 2008, pp. 137–47.

[14] Y. Zhao *et al.*, "Performance Evaluation of Cognitive Radios: Metrics, Utility Functions, and Methodology," *Proc. IEEE*, vol. 97, no. 4, Apr. 2009, pp. 642–59.

[15] J. C. Gittins, *Multi-Armed Bandit Allocation Indices*, Wiley, 1989.

Biographies

Hamed Asadi (hasadi@arizona.edu) received his B.S. degree in computer engineering from Isfahan University of Technology, Iran, in 2006. He is currently a Ph.D. candidate at the Department of Electrical and Computer Engineering, University of Arizona. His research interests include wireless communication, cognitive-radio engine design, computational intelligence, machine learning, intelligent systems, and dynamic radio frequency mapping. He has also previously done research in image processing and computer vision. He also regularly serves as a reviewer for wireless communication conferences and journals. He received a Best Paper Award from the Wireless Innovation Forum in 2015.

Haris Volos [S'02, M'10, SM'15] (volos@ieee.org) is currently a research assistant professor with the Department of Electrical and Computer Engineering at the University of Arizona. He is an expert in cognitive radio engines with several publications on the topic. His other research interests include channel modeling, software-defined radios, and dynamic radio frequency mapping. He has received several awards such as the Fred W. Ellersick MILCOM Award for the Best Paper in the Unclassified Technical Program in 2010, the Paul E. Torgersen Research Excellence Award in 2011, and a Best Paper Award from the Wireless Innovation Forum in 2015.

Michael M. Marefat [SM] (marefat@arizona.edu) is an associate professor at the Department of Electrical and Computer Engineering, University of Arizona. He is a member of ACM and AAAI. He received his Ph.D. and MS degrees in electrical engineering from Purdue University in1991 and 1988, respectively, and his B.S. in electrical and computer engineering and B.A. in mathematical sciences from Rice University. He is the author/co-author of about 100 refereed research articles. He has directed about 20 M.S. theses and Ph.D. dissertations.

Tamal Bose (tbose@arizona.edu) is the department head and a professor of the Department of Electrical and Computer Engineering at the University of Arizona. He is also the national director of the NSF sponsored multi-university multi-industry wireless research center BWAC. He is the author of *Digital Signal and Image Processing* (Wiley, 2004,) co-author of *Basic Simulation Models of Phase Tracking Devices Using MATLAB* (Morgan & Claypool, 2010), and co-author of *Partial Update Least-Square Adaptive Filtering* (Morgan & Claypool, 2014). His research interests lie in the domain of signal processing for wireless communications.

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.

Software Defined Access for HetNets

Vidya Sagar, R. Chandramouli, and K. P. Subbalakshmi

The authors present software defined access (SDA), an architecture for HetNets that incorporates various elements of SDN and software defined wireless networking (SDWN). SDA introduces a novel logical control path across radio interfaces and up to mobile devices. Unlike SDN and SDWN, SDA can be deployed without changing network elements of radio access technologies.

Abstract

We present software defined access (SDA), an architecture for HetNets that incorporates various elements of SDN and software defined wireless networking (SDWN). SDA introduces a novel logical control path across radio interfaces and up to mobile devices. This control path allows SDA to regulate and configure data path via different RATs. Unlike SDN and SDWN, SDA can be deployed without changing network elements of RATs. The userspace implementation of client side utility is also presented to empower mobile devices with SDN-like flow management. This client with cloud hosted SDA servers presents an alternative for operator-independent and RAT-agnostic standalone SDA deployment. To demonstrate some of SDA's capabilities, we also present various testbed experimental measurement results.

Introduction

The proliferation of smart devices and mobile applications has brought web-based services closer to the masses. It has changed the way we order pizza, use taxis, and bank. Most importantly, it has transformed our means of personal and professional communication. This has given birth to ever increasing demands of mobile data traffic. Further enhancements in device, communication, and web technologies are going to fuel this even more. This growing demand has already put network infrastructure under a lot of stress. Due to fierce competition, average revenue per user (ARPU) is lagging behind this burgeoning demand. Increasing demand and decreasing revenue present an imminent threat to networks and operators.

To handle this situation, multiple solutions are being adopted. These solutions aim to improve network utilization, boost spectral efficiency, and enhance interworking among different radio access technologies (RATs), such as Long Term Evolution (LTE) and WiFi. We enumerate some of the widely discussed solutions below:

Traffic Engineering: Traditionally, networks offer minimal intelligence at the core with simple functionalities (e.g., routing). Traffic is usually routed through a fixed or the shortest path. This inefficient scheme causes congestion on some network paths, whereas others remain underutilized. Routing with awareness of load on different paths can improve overall system utilization. Moreover, network flows carrying traffic usually have diverse service requirements. For example, some applications and their flows are delay-tolerant; therefore, routing them through a non-shortest path while meeting their quality of service (QoS) requirements can reduce load on the shortest path. Such intelligent traffic engineering with awareness of load and flow requirements can improve network utilization and guarantee flow-specific QoS requirements.

Network Virtualization: To reduce operational cost, operators have increasingly embraced network virtualization. It facilitates shared network resources among multiple different operators with better service guarantees. Software defined networking (SDN) is an approach that virtualizes network elements (e.g., switches and routers) and shares them among different flows, services, and operators. SDN decouples the control plane from the data path, and uses a centralized controller to configure the data path through the control plane [1]. Network functions virtualization (NFV) is another approach, which virtualizes different network servers (name server, etc.) with the help of IT virtualization techniques. Although independent of each other, SDN with NFV can provide efficient traffic engineering and better service orchestration.

Heterogeneous Deployment: Increasing mobile traffic demands spectral efficiency. Current RATs already operate near theoretical limits of spectral efficiency at the physical layer. Further increments in spectral efficiency are expected from diversity gains (e.g., transmission or reception diversity). Heterogeneous networking (HetNet), with low-power intermediate nodes (e.g., femtocells) and small cells, can leverage these diversities and boost spectral efficiency even more. Additionally, HetNets are useful in reducing coverage holes and extending cell coverage.

MultiRAT Architecture: Multiple different RATs in a HetNet deployment operating at different frequencies is considered part of evolving fifth generation (5G) networks [2]. Although HetNet deployment can provide better connectivity to the core network through individual RATs, spectrum scarcity and growing demand advocate enhanced interworking among these multi-RAT HetNets. Such convergence requires an architecture to incorporate multiple RATs and facilitate interworking. Software defined wireless networking (SDWN) is expected to address this architectural need. SDWN interworks with multiple RATs and extends the control path to the infrastructure plane consisting of base stations, access points, and so on. Using an

This research was supported by the National Science Foundation under award no.1353340

The authors are with Stevens Institute of Technology.

0163-6804/16/$25.00 © 2016 IEEE

SDN-like interface with applications, SDWN can also configure RATs to support application-specific demands. Different technologies mentioned above suggest that a futuristic network will have converged multi-RAT HetNet deployment and an SDWN/SDN-like architecture with enhanced interworking and intelligent traffic engineering [3]. In this work, we address such a converged multi-RAT HetNet architecture.

SDMN advocates changes to various RAT elements. Adopting these changes across different RATs will be an expensive and slow process. Therefore, we propose a readily deployable alternative called software defined access (SDA).

This work is organized as follows. We present an overview of SDN and SDWN architectures. Thereafter, we review a few relevant MultiRAT interworking mechanisms. Then we introduce the proposed SDA architecture and discuss different SDA procedures enabling MultiRAT interworking. The prototype implementation is explained, and experimental results are presented to demonstrate operator-independent QoS-aware offloading. Finally, we conclude with pointers to future work.

Overview: SDN, SDWN, and SDA

Figure 1 shows SDA along with SDN and SDWN. For simplicity, we have presented SDN, SDWN, and SDA independent of each other. However, SDN is generally considered to be part of SDWN, and we foresee SDWN adopting some of SDA's functionalities too. To understand the differences between these architectures, we first identify the following five planes.

Application Plane: The application plane consists of different application servers. NFV techniques can be employed to virtualize these servers for better service orchestration.

Control Plane: SDN has introduced this new plane consisting of an SDN controller. We place the SDWN controller and SDA controller in this plane. These controllers are responsible for intelligent traffic policing and employ the control path to enforce these policies. The SDN control path is limited to the core network, whereas SDWN extends the SDN control path to the RAT infrastructure (e.g., base station, access point), and the proposed SDA introduces a novel control path across the air interface up to a mobile device. All these programmable controllers can also interact with the application plane and adapt policies to accommodate different application-specific requirements. OpenFlow SDN architecture has standardized northbound application programming interfaces (APIs), which govern the interface between the control and application planes [1].

Data Plane: This plane consists of a data path with switches and routers constituting the legacy network core. The SDN control path terminates in this plane. OpenFlow SDN architecture suggests southbound APIs to administer traffic policies between the control and data planes [1].

Infrastructure Plane: This plane consists of different RAT elements (base station, access points, etc.). SDWN attempts to extend the SDN control path up to this plane. Although the SDWN architecture is still evolving, there are a few relevant architectures for SDWN. OpenRoads extends OpenFlow by using Simple Network Management Protocol (SNMP) [4]. On the other hand, RAT-specific southbound APIs are employed for HetNet SDN [5]. OpenRadio outlines refactoring radios to devise a programmable data plane for SDWN [6].

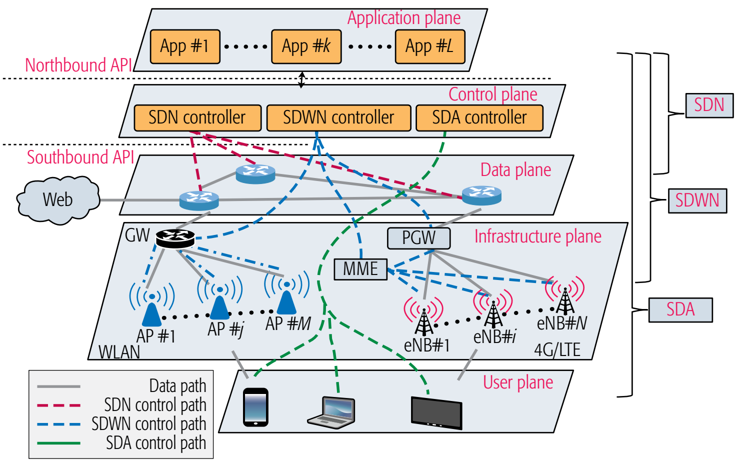

Figure 1. SDN, SDWN, and SDA.

User Plane: The user plane is the last plane consisting of end-user equipments (UE) (e.g., tablets and smartphones). Due to the prevalence of multiple radios in these devices, they can connect to the network core via multiple RATs. By selecting different radios for different applications, these devices can act as wireless switches for traffic originating or terminating at them. Moreover, a data path over a given RAT can be exploited to configure a data path over other RATs. In multi-RAT deployments, a temporal logically separate control path can also be devised from multiple data paths. The proposed SDA architecture employs such a logical control path between UE and SDA controller.

Existing Interworking Methods

Interworking among different RATs will be a key component in emerging networks. We now present an overview of existing multi-RAT interworking mechanisms and their limitations in order to motivate the need for an innovative solution.

The Third Generation Partnership Project (3GPP) has adopted several measures to facilitate LTE-WiFi interworking for LTE-Advanced (LTE-A) systems. MAPCON allows UE to maintain simultaneous multiple data paths over different RATs [7]. The access network discovery and selection function (ANDSF) and IFOM mechanisms are proposed to enable RAT selection and flow mobility across different RATs [8]. These technologies rely on DualStack Mobile IPv6 (DSMIPv6), but it requires kernel-level modification to UEs. Proxy Mobile IPv6 (PMIPv6)-based flow mobility requires changes in RAT elements [9]. Apart from UE or RAT modifications, these technologies also mandate coordination among operators. Cellular and non-cellular services are usually marketed separately, and very often users subscribe to these services separately too. These 3GPP mechanisms may not be helpful for users with subscriptions to non-coordinating operators.

We attempt to address these shortcomings by introducing SDA. We propose SDA as a readily deployable architecture to facilitate operator-in-

We propose SDA as a readily deployable architecture to facilitate operator-independent multi-RAT interworking with minimal changes in existing RAT elements. We believe that the operator-independent architecture will also help operators by reducing RAT complexity and promoting better utilization of current RAT infrastructure.

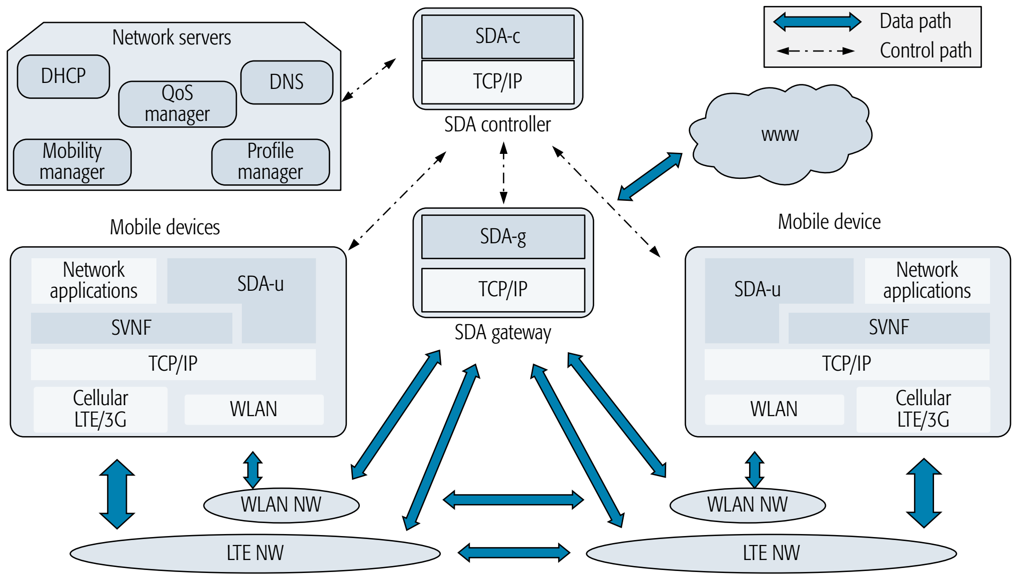

Figure 2. Proposed SDA architecture.

dependent multi-RAT interworking with minimal changes to existing RAT elements. We believe that the operator-independent architecture will also help operators by reducing RAT complexity and promoting better utilization of current RAT infrastructure.

SDA ARCHITECTURE

Figure 2 presents the proposed SDA architecture. It consists of two new network entities, SDA-Gateway (SDA-g) and SDA-Controller (SDA-c). On UE, it makes use of SDA userspace client (SDA-u) and SDA virtual network function (SVNF). SDA-g and SDA-c are assumed to be accessible via all the different RATs available to UEs. Details of different elements are presented below.

SDA Virtual Network Function: The SVNF is a system utility that provides a virtual network interface. It acts as a bidirectional conduit between applications and SDA-u. It maintains an IP address assigned by SDA-c. This RAT-independent IP address allows applications to maintain their sockets and flows during handover and coverage loss.

SDA-u: SDA-u is a client side utility running as a service on UEs. It is responsible for authenticating mobile devices with SDA-c. It also updates SDA-c with user preferences. It complies with traffic policies provided by SDA-c. SDA-u is responsible for configuration of the data path between SVNF and SDA-g. SDA-u splits various outgoing flows into subflows specific to different RATs. It also retrieves incoming flows from RAT-specific subflows. Other responsibilities of SDA-u include loss recovery, delay adjustment, and encryption of the data path.

SDA-g: SDA-g works as a proxy server and a gateway for the SDA network. On the data path, it performs functions similar to the SDA-u. It also reports different subflow characteristics (e.g., delay, loss, and throughput) to SDA-c.

SDA-c: It is the centralized SDA controller. SDA-c authenticates UEs and authorizes them in accordance with their profile and preferences. It establishes different data subflows between SDA-u and SDA-g. SDA-c enforces a multitude of traffic policies (e.g., offloading, aggregation) over these subflows. Additionally, it monitors the data path using measurement reports from SDA-u and SDA-g. It can also detect several impairments (interference, congestion, etc.) through these measurements and adapt traffic policies to ensure better QoS.

DATA AND CONTROL PATH

For the data path, SDA relies on multiple subflows. These subflows are network flows bound to different RATs. SDA-u and SDA-g disseminate outgoing user data over these subflows. They also retrieve incoming flows from these subflows. These dissemination policies of user flows over subflows are governed by SDA-c. SDA-c classifies traffic on a per flow basis and enforces different policies for different flows. These policies are essentially routing policies with weighted matrices, and provide mapping between user flows and RAT-specific subflows. SDA realizes flow offloading from a RAT by removing mapping of flows from that RAT or associated subflow. Similarly, aggregation can be achieved by selecting multiple RATs or subflows with appropriate weights.

For the logical control path, SDA depends on multiple UDP flows. It uses three-way handshake and *UDP hole-punching* to traverse through Network Address Translation (NAT). Communication over this control path dynamically selects the appropriate RAT to achieve better reliability and have the least impact on data flows. This logical flow is initiated by SDA-c during the registration process after successful authentication. Although this work does not consider symmetric and restricted cone NATs, we assume three-way handshake can be replaced with an appropriate procedure to enable UDP NAT traversal [10]. The SDA three-way handshake is as follows. SDA-c assigns a local port number and a connection identifier to SDA-u. SDA-c sends this information over an initial TCP connection used for authentication. Thereafter, SDA-u opens a UDP socket bound to *localhost* and communicates to SDA-c on the assigned UDP port with the connection identifier. SDA-c learns the public IP information (public IP address and port number) of SDA-u and sends an acknowledgment to SDA-u over a new control socket. It concludes three-way handshake. Whenever SDA-u registers to a new RAT, it makes use of the previously established control path to initiate this three-way handshake over the new RAT. It helps SDA-c to maintain a list of public IP information linked to the same SDA-u over different RATs. Finally, dynamic RAT selection in this logical control path is achieved by selecting appropriate destination IP information in downlink and source IP information in uplink communication.

Interworking with SDN/SDWN

SDA introduces one new node in the control plane (i.e., SDA-c) and one in the data plane (i.e., SDA-g). SDA can achieve operator independence by hosting these servers outside operators' networks. Additionally, SDA-c can interact with the SDN/SDWN controller of the underlying RATs for enhanced interworking and better resource management. We foresee SDWN/SDN to accommodate SDA controllers into their controllers. However, SDA can also be deployed independently without SDN/SDWN as a cloud service (i.e., *SDA-as-a-service*). We illustrate our prototype implementation of SDA-as-a-service later in this article.

Call Flow in SDA Systems

Figure 3 illustrates a sample call flow with registration, measurement, and flow offload procedures in SDA. We consider UE with LTE and WiFi radios. We assume LTE as the primary RAT and WiFi as the preferred RAT for flow offloading. Therefore, SDA-u will register and start communication over LTE; then the availability of WiFi will cause flow offload from LTE to WiFi. As shown in Fig. 3, SDA-u sends a *RegistrationRequest* to SDA-c using a TCP connection over LTE. This request consists of identity and authentication data as well as user preferences. SDA-c makes use of this information to authenticate and identify an appropriate QoS level for SDA-u. To configure the data path, SDA-c sends an *AddSession* message to SDA-g. This message includes *ConfigureSession*, which provides configuration for data and control path as well as periodic measurement reports. *ConfigureSession* also embeds a *PolicyUpdate* message to configure the transmission policy. The transmission policy enables SDA-g/SDA-u to schedule packets among different RATs. As discussed earlier, SDA transmission policies are essentially routing policies. In a standalone implementation of SDA, these policies do not interact with RAT-specific radio resource scheduling. Furthermore, SDA-g acknowledges successful configuration to SDA-c with an *AddSessionACK* message. Thereafter SDA-c creates a UDP socket for the control path and sends a *RegistrationACK* message to SDA-u indicating successful registration. This *RegistrationACK* also includes *ConfigureSession* consisting of configurations for control, data path, and measurement sessions. It also conveys an embedded *PolicyUpdate*

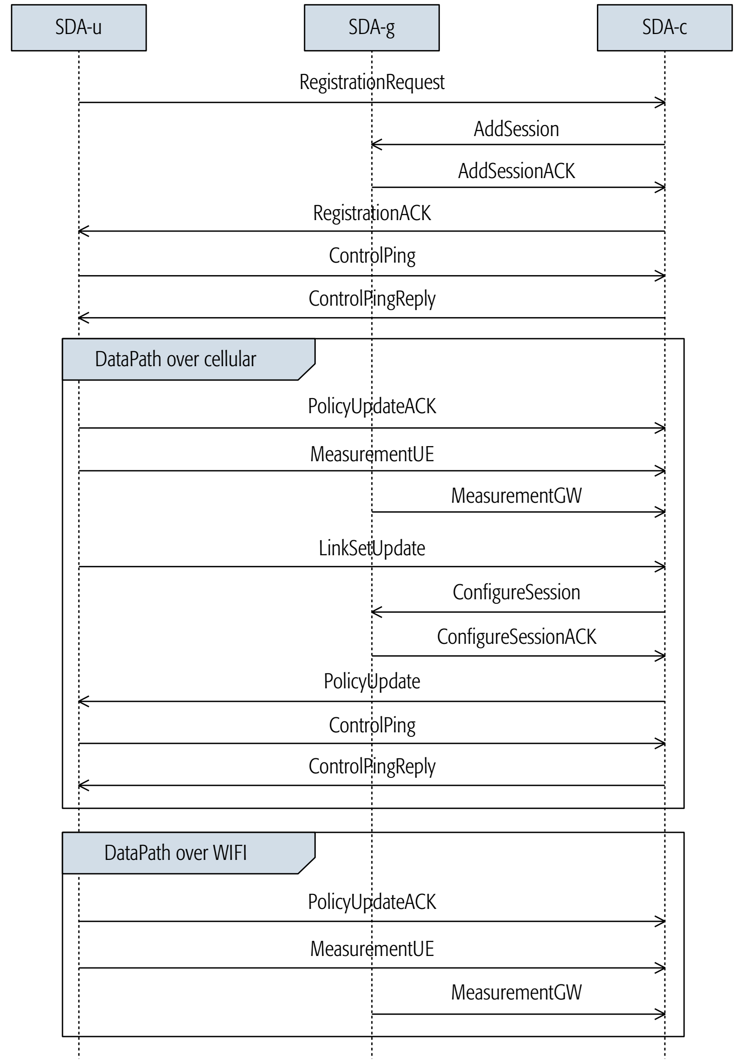

Figure 3. Sequence diagram for offloading.

message along with scheduling policy and initiates a three-way handshake for initial control path configuration. After successful registration, SDA-u sends *ControlPing* over UDP using the default RAT to SDA-c. SDA-c learns the public IP address (and port number) of SDA-u and sends an acknowledgment back by sending *ControlPingReply* over the control path. Finally, SDA-u sends *PolicyUpdateACK* to SDA-c and terminates the TCP session with SDA-c. Thereafter, user data is carried over LTE, periodic measurements are reported by SDA-u and SDA-g.

When WiFi becomes available, SDA-u sends *LinkSetUpdate* message to SDA-c with WiFi network information over the control path. Since WiFi is the preferred network, SDA-c triggers offload of data flows. It updates SDA-g about the new data path by sending *ConfigureSession*. SDA-g replies back with *ConfigureSessionAck* to SDA-c, indicating the completion of configuration. SDA-c sends a *PolicyUpdate* message to SDA-u with the new scheduling policy. This is followed by a three-way handshake to update the logical control path over WiFi. Once the control path over WiFi is configured, SDA-u sends a *PolicyUpdateACK* over WiFi, and all data flows are offloaded to WiFi.

Also, SDA-u and SDA-g send periodic measurement reports assisting SDA-c in monitoring.

Apart from priority-based offloading and registration, the above example also illustrates measurement procedures in SDA. SDA-c can configure these measurement sessions to identify a variety of conditions (e.g., coverage loss, congestion). This measurement information can be used to facilitate QoS aware traffic engineering.

Prototype Implementation

In this section, we discuss the standalone deployment of SDA-as-a-service. We make use of SDA-c and SDA-g servers hosted on a cloud infrastructure. SDA-u registers to SDA-c with its credentials and preferences. SDA-c facilitates UDP-based control and data paths to SDA-u. SDA-c also assigns an IP address that does not change with changes in the IP addresses of the underlying RATs. SDA-c makes use of different measurement reports from SDA-u and SDA-g to identify the optimal scheduling policy of user data over RAT-specific subflows.

Cloud hosted SDA servers provide the necessary framework for operator-independent SDA deployment with no changes in underlying RATs. To make SDA-as-a-service a readily deployable solution, we present the design of SDA-u implemented as a userspace application.

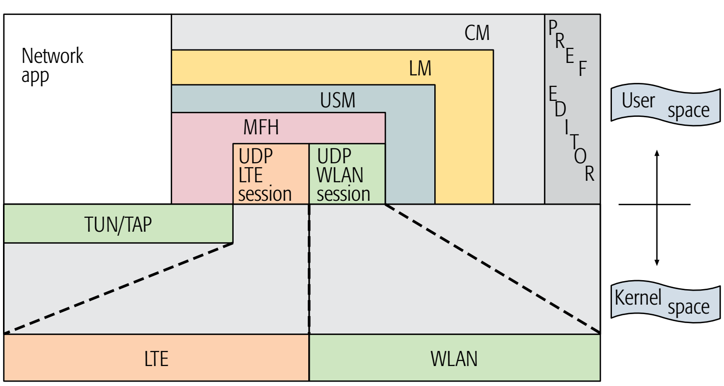

Figure 4. SDA-u design.

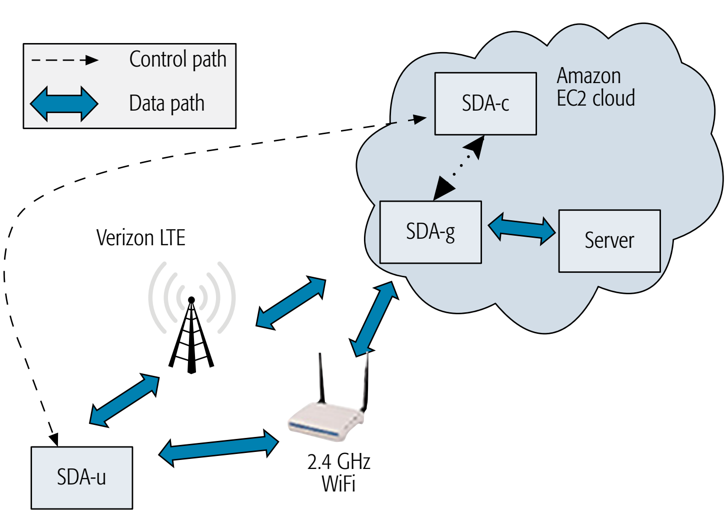

Figure 5. SDA-as-a-service experimental setup.

SDA-u Design

Our SDA-u design depends on basic functionalities (e.g., socket) provided by many operating systems (e.g., Unix/Linux, Windows). Figure 4 illustrates the design of SDA-u. Each component of SDA-u is explained below.

Tun/Tap: Tun/Tap is used as an SVNF here. It carries IP configuration provided by SDA-c. The Tun/Tap interface forwards all the outgoing packets sent by other network applications to SDA-u. It also delivers all the packets sent by SDA-u to the respective applications.

Multiple Flow Handler: The MFH is a core module for the data path. The MFH is common between SDA-g and SDA-u. It reads outgoing flows from Tun/Tap and employs different functions (e.g., encapsulation, encryption) to process them. Processed flows are forwarded to different UDP sessions/subflows bound to different RATs. It communicates with SDA-g and ensures multilink reliable UDP connectivity.

UDP Session Manager: The USM maintains and manages different UDP sessions with SDA-g. The USM reestablishes UDP sessions upon link failure or handover. It also takes care of the maintenance of the control path.

Link Manager (LM): The LM maintains IP connectivity with various RATs. It triggers RAT-specific procedures for registration and association. It requests the USM to reestablish subflows upon changes in IP configuration. It also identifies link losses/failures and requests the USM to close associated subflows.

Client Manager: The CM is the main controller of the SDA-u software system. It takes care of authentication and association with SDA-c and also reports measurement information requested by SDA-c. It enforces different policy updates from SDA-c.

Preference Editor: PrefEdit is a user interface. PrefEdit allows a user to modify preferences, RAT priorities, and RAT-specific procedures.

This design does not require kernel updates, and can be implemented and installed as an application. User space implementation of SDA-u along with cloud hosted SDA-c and SDA-g servers presents a readily deployable solution to facilitate operator-independent multi-RAT interworking without altering the underlying RATs.

Experimental Results

Figure 5 presents our SDA-as-a-service experimental setup. We implemented SDA-u on a mini-ITX unit with a 1.86 GHz Intel Atom processor, 4 GB RAM, and an Intel Advanced-N 6235 wireless card. SDA-c and SDA-g servers are deployed on Amazon ec2 cloud. All SDA entities were implemented on Ubuntu 14.04. A Netgear WNDR3400v2 access point with a university LAN is used to provide the WiFi data path. The LTE data path is established using a Pantech UML 290 card over Verizon's LTE network. The IPERF utility is used to evaluate throughput over this setup. To demonstrate the capability of SDA-as-a-service, we emulated congestion in the LTE data path. SDA-c identifies congestion and offloads flow from LTE to avoid throughput degradation. We call it QoS-aware offloading.

Figure 6 presents QoS-aware offloading under SDA-as-a-service. This figure presents a timeline plot of TCP throughput performance over the SDA data path. In this experiment, we assume LTE as the preferred network. For the first 20 s, the SDA data path is catered for by the LTE data path. During this period the WiFi data path is not used, although SDA-c monitors it through measurement reports. Meanwhile, the average throughput performance observed is 7.7 Mb/s. After 20 s, NETEM and IPTABLE utilities are used to reduce the LTE throughput to 500 kb/s. The SDA-c detects it as core network congestion and offloads traffic to the WiFi network. Up to the 40th second this simulated congestion is maintained, and traffic is served through the WiFi data path with 3.6 Mb/s of average throughput. When throttling of LTE is removed, SDAc detects this event. Since the preferred RAT (i.e., LTE) can now serve a larger throughput, SDA-c offloads flow to LTE. In this last stage, SDA data path is routed over LTE with an average throughput of 7.3 Mb/s. In this experiment, offloading delay is measured as 0.8 s for LTE to WiFi offloading and 1.3 s for WiFi to LTE offloading. LTE and WiFi networks exhibit different round-trip delays, 180 ms and 35 ms, respectively. When flow is offloaded from WiFi to LTE, due to increased delays, TCP throughput degrades initially, and thereafter recovers and saturates the LTE data path. This is why throughput drops after 40 s before ramping up to the appropriate LTE throughput.

These experimental results demonstrate the capability of standalone implementation of SDA (i.e., SDA-as-a-service) to offload flow to and from different RATs as well as the capability to detect congestion in the data path and adapt accordingly.

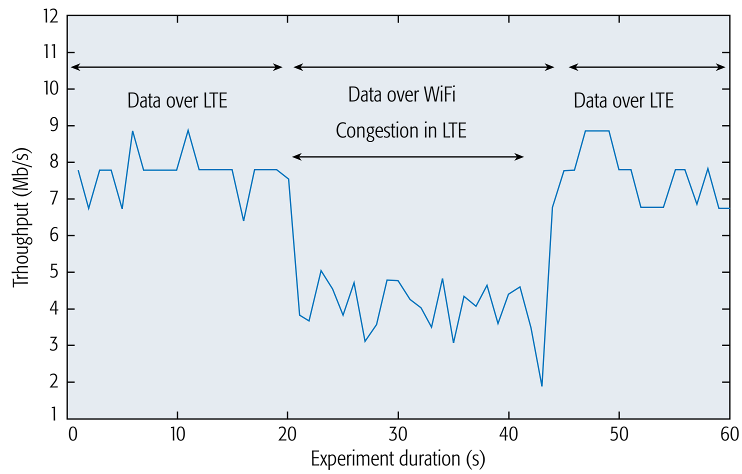

Figure 6. SDA data offloading performance.

Advantages of SDA

As discussed above, *SDA-as-a-service* implementation makes SDA agnostic to network operators and the underlying RATs. We foresee SDA as a complementary to futuristic SDWN architectures. Some of the advantages of SDA are the following.

Logical Control Path: Extension of the logical control path up to mobile devices will provide better means of control and reconfigurability to SDWNs.

User-Centric Architecture: SDA introduces UEs into an SDWN-like architecture. This can help SDWNs to converge into a user-centric architecture. Such an architecture can bring fine-grained control over various network elements with an awareness of user preference and can provide superior QoS.

QoS-Aware Flow Mobility: SDWN-based flow offloading often relies on fixed preferences for RATs (e.g., offload to WiFi). Such offloading can result in poor user experience if offloaded flow encounters poor network conditions (e.g., interference). By adopting SDA like operator-independent offloading, SDWN can cater for seamless and QoS-aware offloading.

Application/Service Awareness: SDA provides integration of user preferences and application service requirements in SDWN/SDN architecture. For example, it can enable opportunistic multi-RAT aggregation for data-hungry applications.

Conclusion

We have presented SDA, a novel architecture to facilitate interworking in multi-RAT HetNets. SDA also provides better reconfigurability and control to SDWNs without changing any network elements. A standalone deployment of SDA, SDAas-a-service, is also introduced. Experimental results show the efficacy of the prototype deployment.

In the future, we will implement SDA-as-a-service on smart devices and evaluate the impact of these mechanisms on battery life.

References

[1] N. McKeown *et al.*, "Openflow: Enabling Innovation in Campus Networks," *SIGCOMM Comp. Commun. Rev.*, vol. 38, no. 2, Mar. 2008, pp. 69–74.

[2] B. Bangerter *et al.*, "Networks and Devices for the 5G Era," *IEEE Commun. Mag.*, vol. 52, no. 2, Feb. 2014, pp. 90–96.

[3] R. Wang, H. Hu, and X. Yang, "Potentials and Challenges of C-RAN Supporting Multi-Rats Toward 5G Mobile Networks," *IEEE Access*, vol. 2, 2014, pp. 1187–95.

[4] K.-K. Yap *et al.*, "Openroads: Empowering Research in Mobile Networks," *SIGCOMM Comp. Commun. Rev.*, vol. 40, no. 1, Jan. 2010, pp. 125–26.

[5] S. Tomovic *et al.*, "SDN-Based Concept of QoS Aware Heterogeneous Wireless Network Operation," *2014 22nd Telecommun. Forum Telfor*, Nov. 2014, pp. 27–30.

[6] M. Bansal *et al.*, "Openradio: A Programmable Wireless Dataplane," *Proc. 1st ACM Wksp. Hot Topics in Software Defined Networks*, ser. HotSDN '12. New York, NY, 2012, pp. 109–14.

[7] 3GPP, "Multi Access PDN Connectivity and IP Flow Mobility," TS 23.861, Feb. 2010.

[8] 3GPP, "IP Flow Mobility and Seamless Wireless Local Area Network (WLAN) Offload; Stage 2," TS 23.261, Mar. 2012.

[9] A. de la Oliva *et al.*, "IP Flow Mobility: Smart Traffic Offload for Future Wireless Networks," *IEEE Commun. Mag.*, vol. 49, no. 10, Oct. 2011, pp. 124–32.

[10] R. Mahy, P. Matthews, and J. Rosenberg, "Traversal Using Relays Around NAT (TURN): Relay Extensions to Session Traversal Utilities for NAT (STUN)," IETF RFC 5766, Apr. 2010.

Biographies

Vidya Sagar (vsagar@stevens.edu) is a Ph.D candidate in the Department of Electrical and Computer Engineering, Stevens Institute of Technology. His research interests include cognitive, cloud and social networks.

R. Chandramouli (mouli@stevens.edu) is the Thomas Hattrick Chair Professor of Information Systems in the Department of Electrical and Computer Engineering, Stevens Institute of Technology, and co-founder of Dynamic Spectrum, LLC and jaasuz.com. His research spans the areas of cognitive wireless networking, text mining, social media security, and analytics. His projects have been supported by the U.S. National Science Foundation, National Institute of Justice, Department of Defense (DoD) agencies, and industry

K. P. Subbalakshmi (ksubbala@stevens.edu) is a professor in the Department of Electrical and Computer Engineering, Stevens Institute of Technology. Her research interests include cognitive radio networks, network security, media forensics, and social networks. She is the Founding Chair of the Special Interest Group on Security of the IEEE TCCN. Her work has been supported by the National Science Foundation, National Institute of Justice, DoD, and industry. She is also co-founder of Dynamic Spectrum LLC and Jaasuz.com.

Network and Service Management

George Pavlou

Jürgen Schönwälder

This is the 20th issue of the series on Network and Service Management, which is typically published twice a year, in January and July. The series features articles on the latest developments in this well established discipline, highlighting recent research achievements and providing insight into both theoretical and practical issues related to the evolution of the discipline from different perspectives. The series provides a forum for the publication of both academic and industrial research, addressing the state of the art, theory, and practice in network and service management.

A key annual event of the network and service management community was the 11th International Conference on Network and Service Management (CNSM 2015) (http://www.cnsm-conf.org/2015/), which was co-sponsored by IFIP and IEEE and organized in cooperation with ACM. CNSM 2015 took place November 9–12 in Barcelona, Spain, and it hosted among other things the technical program committee meeting of the IEEE/IFIP Network Operations and Management Symposium (NOMS 2016), which will take place April 25–29 in Istanbul, Turkey. Other noteworthy events in the first half of 2016 are the 2nd IEEE Conference on Network Softwarization (NetSoft 2016) taking place June 6-10 in Seoul, Korea, and the 10th Conference on Autonomous Infrastructure, Management and Security (AIMS 2016) taking place June 20–24 in Munich, Germany.

During the CNSM 2015 conference, a number of discussions took place focusing on how to improve the network and service management taxonomy, which has been used by several conferences in this discipline. The discussions were fueled by a questionnaire produced by the European FLAMINGO research project in order to collect input about the relevance of topics. A total of 154 people provided input, 49 from industry and 105 from academia. While there is much commonality in how these groups judge the importance of topics, there are also notable differences. One topic area where judgments differ is the management of the Internet of Things and its services. Participants from academia consider topics related to the Internet of Things much more important than participants from industry.

We again experienced excellent interest for the 20th issue, having received 28 submissions in total. We received at least three independent reviews for all the articles, finally selecting four articles, resulting in an acceptance rate of 14.3 percent. It should be mentioned that the acceptance rate for all the previous issues has ranged between 14 percent and 25 percent, making this series a highly competitive venue to publish. We intend to maintain our rigorous review process in future issues, thus maintaining the high quality of the published articles.

The first article, "Cloudifying the 3GPP IP Multimedia Subsystem for 4G and Beyond: A Survey" by Abu-Lebdeh, Sahoo, Glitho, and Chouati, explains that the 3GPP IMS does not meet requirements for 4G and beyond, while cloudifying it will allow its use as a service delivery platform. In this context it provides a critical overview of the relevant architectures proposed so far.

The second article, "Management and Orchestration Challenges in Network Function Virtualization" by Mijumbi, Serrat, Gorricho, Latre, Charalambides, and Lopez, first presents the ETSI NFV management and orchestration framework, then presents representative projects and vendor products, and finally identifies open challenges and opportunities for future research.

The third article, "Transparent Reallocation of Control Functions in IMS Deployments" by Garcia-Reinoso, Vidal, Bellavista, Soto, and Aranda-Gutierrez, presents a solution to transfer users between IMS elements so that an operator can perform redistribution of load and be able to accommodate instantaneous load generated by users.

Finally, the fourth article, "Self Healing in Mobile Networks with Big Data" by Khatib, Barco, Munoz, de la Bandera, and Serrano, considers the amount of control and measurement data used by self-healing systems in mobile networks which can be characterized as big data, and subsequently surveys some relevant use cases along with their big data solutions.

We hope that readers of this issue again find the articles informative, and we will endeavor to continue with similar issues in the future. We would finally like to thank all the authors who submitted articles to this series and the reviewers for their valuable feedback and comments on the articles.

Biographies

George Pavlou (g.pavlou@ucl.ac.uk) is a professor of communication networks in the Department of Electronic and Electrical Engineering, University College London, UK, where he coordinates research activities in networking and network management. He received a Diploma in Engineering from the National Technical University of Athens, Greece, and M.Sc. and Ph.D. degrees in computer science from University College London, UK. His research interests focus on networking and network management, including areas such as traffic engineering, quality of service management, policy-based systems, autonomic networking, information-centric networking, and software-defined networks. He has been instrumental in a number of European and UK research projects that produced significant results with real-world uptake, and has contributed to standardization activities in ISO, ITU-T, and IETF. He has been the technical program chair of several conferences, and in 2011 he received the Daniel Stokesbury award for "distinguished technical contribution to the growth of the network management field."

Jürgen Schönwälder (j.schoenwaelder@jacobs-universiy.de) is a professor of computer science at Jacobs University Bremen, Germany, where he is leading the Computer Networks and Distributed Systems research group. He received a doctoral degree from the Technische Universität Braunschweig, Germany. His research interests include network management, distributed systems, network measurements, embedded networked systems, and network security. He is an active member of the Internet Engineering Task Force (IETF), where he has edited more than 30 network management related specifications and standards. He has been co-chairing the ISMS working group of the IETF, and he currently serves as co-chair of the NETMOD working group. Previously, he chaired the Network Management Research Group of the Internet Research Task Force (IRTF). He has been involved in several European research projects, and he has served as in various roles for IEEE and IFIP sponsored conferences. He currently serves on the editorial boards of the *Springer Journal of Network and Systems Management*, and the *Wiley International Journal of Network Management*. Previously, he served on the editorial board of *IEEE Transactions on Network and Service Management*.

Cloudifying the 3GPP IP Multimedia Subsystem for 4G and Beyond: A Survey

Mohammad Abu-Lebdeh, Jagruti Sahoo, Roch Glitho, and Constant Wette Tchouati

Abstract

4G systems have been continuously evolving to cope with the emerging challenges of human-centric and M2M applications. Research has also now started on 5G systems. Scenarios have been proposed and initial requirements derived. 4G and beyond systems are expected to easily deliver a wide range of human-centric and M2M applications and services in a scalable, elastic, and cost-efficient manner. The 3GPP IMS was standardized as the service delivery platform for 3G networks. Unfortunately, it does not meet several requirements for provisioning applications and services in 4G and beyond systems. However, cloudifying it will certainly pave the way for its use as a service delivery platform for 4G and beyond. This article presents a critical overview of the architectures proposed so far for cloudifying the IMS. There are two classes of approaches; the first focuses on the whole IMS system, and the second deals with specific IMS entities. Research directions are also discussed. IMS granularity and a PaaS for the development and management of IMS functional entities are the two key directions we currently foresee.

Introduction

Mobile systems have been undergoing a rather fast evolution recently. Fourth generation (4G) systems have provided increasingly higher bandwidth, lower latency, and more features to meet the more stringent requirements of human-centric and machine-to-machine (M2M) communications since their inception during the second half of the last decade. This constant innovation has paved the way for the growth of future human-centric and M2M applications, and is now leading us to the 5G era.

METIS is a European project that aims to lay the foundation of the 5G concept to fulfill the requirements of the beyond-2020 connected information society and support new usage scenarios. It identifies five service and application scenarios that 5G will have to support: *amazingly fast great service in a crowd, best experience follows you, super real-time reliable communications,* and *ubiquitous things communicating* [1]. Several requirements are derived from these scenarios, such as much higher bandwidth, much lower latency, and much more stringent reliability and scalability than offered today by the evolved 4G systems. For instance, 5G systems are expected to attain 10 to 100 times higher user data rate and 5 times lower end-to-end latency [1]. Another example is the requirement of cost efficiency, which was not a primary concern in 4G. This is certainly due to the recent emergence of new technologies such as cloud computing that can easily enable cost efficiency.

The Third Generation Partnership Project (3GPP) IP multimedia subsystem (IMS) [2] is a strong candidate for application and service provisioning in 4G and beyond because it will enable a smooth migration. It was specified as the application and service delivery platform for 3G networks and was then used at the inception of 4G as the de facto service platform. However, it does not meet all of the requirements of 4G and beyond.

Cloud computing has emerged as a paradigm for delivering computing resources (e.g., servers and storage) as a utility. It promises many benefits including elasticity, efficiency in resource usage, easy application and service provisioning, and cost reduction. It has established the foundations for the emergence of network functions virtualization (NFV), which aims to transform network architectures through the implementation of network functions (e.g., IMS) in software that can run on industry standard hardware. Cloud and NFV technologies can certainly aid in tackling the IMS shortcomings when it comes to the requirements of 4G and beyond mobile and wireless systems.

There are several approaches for integrating IMS and cloud technologies. Gouveia *et al.* [3] illustrate these approaches by presenting scenarios in a 4G network setting. In the first group of scenarios, IMS is re-engineered using cloud technologies. In the second group, IMS is used to access applications and services implemented in clouds. In this article, "cloudifying IMS" means re-engineering IMS using cloud technologies. This corresponds to the first group of scenarios. Readers interested in the use of IMS to access applications and services implemented in the cloud can consult [4].

This article is a survey on IMS cloudification for 4G and beyond. It provides a critical review of the architectures for cloudyfying IMS that have been proposed in the literature and further

The 3GPP IMS was standardized as the service delivery platform for 3G networks. Unfortunately, it does not meet several requirements for provisioning applications and services in 4G and beyond systems. However, cloudifying it will certainly pave the way for its use as a service delivery platform for 4G and beyond.

This article is an extended version of a paper presented at NTMS 2014 under the title "Cloudifying the 3GPP IP Multimedia Subsystem: Why and How?"

Mohammad Abu-Lebdeh, Jagruti Sahoo, and Roch Glitho are with Concordia University; Constant Wette Tchouati is with Ericsson.

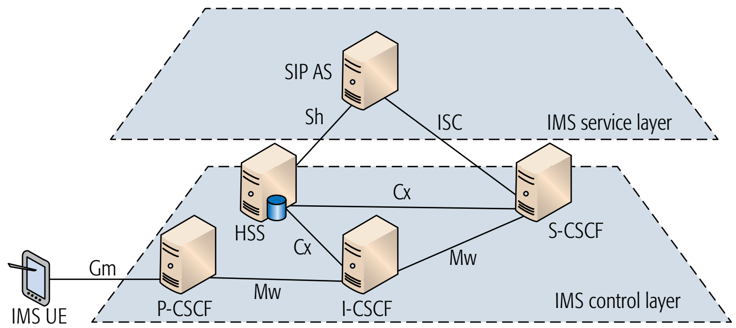

Figure 1. Simplified IMS architecture.

discusses research directions. NFV-based architectures are included in our review. The architectures proposed for IMS cloudification thus far focus on either the entire IMS system or on specific entities. We start by introducing IMS, cloud computing, and NFV, and also outlining the requirements of IMS cloudification for 4G and beyond. The third section reviews the architectures that focus on the entire IMS system. In the fourth section, we discuss the architecture that focuses on specific IMS entities. The fifth section focuses on research directions, and we conclude in the final section.

Background Information on IMS, Cloud Computing/NFV, and Requirements for IMS Cloudification

IMS

IMS is an overlay control layer on top of an IP transport layer required for seamless and robust provisioning of IP multimedia services to end users. It is made up of a service layer and a control layer. The service layer includes application servers, such as a presence server. The key functional entity of the control layer is the call state control function (CSCF). It uses the Session Initiation Protocol (SIP) to control multimedia functions.

Figure 1 depicts a simplified architecture for an IMS network. There are three types of CSCF: proxy-CSCF (P-CSCF), interrogating-CSCF (I-CSCF), and serving-CSCF (S-CSCF). The P-CSCF is the first point of contact for the IMS user equipment (UE) within an IMS network. It acts as a stateful SIP proxy when routing SIP signaling messages going to and from an IMS UE. It is allocated to the IMS UE and does not change for the duration of the registration. I-CSCF is the first contact point for external IMS networks. It is a stateless SIP proxy that selects an S-CSCF for IMS UE and routes incoming SIP signaling messages to the selected S-CSCF. The S-CSCF is the central node of the signaling plane of an IMS network. It acts as a stateful SIP registrar and proxy in an IMS network. As a SIP registrar, it registers IMS users and maintains the binding between the public user identity and the user profile. It also interacts with the home subscriber server (HSS) via the Cx reference point to obtain users' profiles. As a SIP proxy, the S-CSCF forwards specific types of SIP messages to the appropriate application server.

The HSS is another key component of the architecture. It is the central database of the mobile network that contains user-related information, such as subscription, location, and identification information. It supports the network entities' functions (e.g., mobility) and service provisioning. Several IMS functional entities at both the IMS service and control layers interact with it using the diameter protocol.

The SIP application server (AS) is a SIP-based server that implements the logic of IMS services. The SIP AS interacts with HSS to obtain users' profiles via the Sh reference point. An example of an IMS service is the presence service, which accepts, stores, and distributes presence information via SIP messages.

The 3GPP IMS specification provides scalability through the distribution of components such as the CSCF and HSS. However, despite this provision, scalability remains a key issue in IMS, as articulated in [5]. This is due to the fact that SIP is a text-based protocol. Signaling delay may not be sustainable when several CSCFs and application servers are deployed. In addition to the scalability issue, there is actually no provision in IMS for meeting the cost efficiency requirement of 4G and beyond mobile and wireless communications.

Cloud Computing and NFV

Cloud computing has emerged as a viable delivery model for IT resources. It leverages visualization technology to enable on-demand network access to a shared pool of configurable resources (e.g., networks, servers, storage, applications, and services) with self-service provisioning and administration. It has three main service models: infrastructure as a service (IaaS), platform as a service (PaaS), and software as a service (SaaS).

IaaS offers end users computing resources such as processing, storage, and network as a service over a network. End users can dynamically provision and de-provision resources according to their need. Service providers use PaaS to provision applications and services that are offered as SaaS on a pay-per-use basis to end users or other applications. PaaS eases the provisioning process by adding levels of abstraction to the infrastructure offered as IaaS. PaaS solutions vary widely in the capabilities they offer. However, they all have the basic capability to deploy applications on IaaS.

The NFV technology offers a new way to design, deploy, and manage network services. It decouples network functions that are implemented in software from the underlying proprietary hardware and runs the software as applications (i.e., virtual network functions [VNFs]) on commercial off-the-shelf (COTS) hardware [6]. The shift toward software-based network functions leads to flexibility as the VNFs can easily be deployed in various locations, updated, and scaled without the need to change the hardware.

NFV was developed to benefit the networks from virtualization technology to consolidate and run VNFs on COTS hardware such as servers and switches. It promises many benefits to the telco industry such as flexibility, openness, network services agility, and reduced capital expenditures (CAPEX) and operational expenditures (OPEX) [6].

Although related, cloud computing and NFV are different concepts. Cloud computing refers to the concept of delivering the computing resource as a service, whereas NFV focuses on migrating the network functions to run on COTS hardware. However, by leveraging cloud computing, NFV can take advantage of the benefits of cloud computing and bring it to the telco industry. The benefits include elasticity, resource efficiency, and even more reduced CAPEX and OPEX than NFV on its own.

The NFV architectural framework [7], as it is being standardized by the European Telecommunications Standards Institute (ETSI), is depicted in Fig. 2. It comprises NFV infrastructure (NFVI), VNFs, and NFV management and orchestration layers. NFVI provides the environment in which VNFs can execute. It provides the compute capabilities comparable to an IaaS, although usually with much more stringent performance requirements. It also supports the dynamic network connectivity between VNFs, which can be achieved by leveraging emerging technologies such as software-defined networking (SDN). The virtualized infrastructure manager performs resource management and allocation. The VNF manager handles VNF life cycle management (e.g., instantiation, scaling, and termination). The VNF orchestrator is mainly responsible for the life cycle management of the network services, which usually includes several VNF instances.

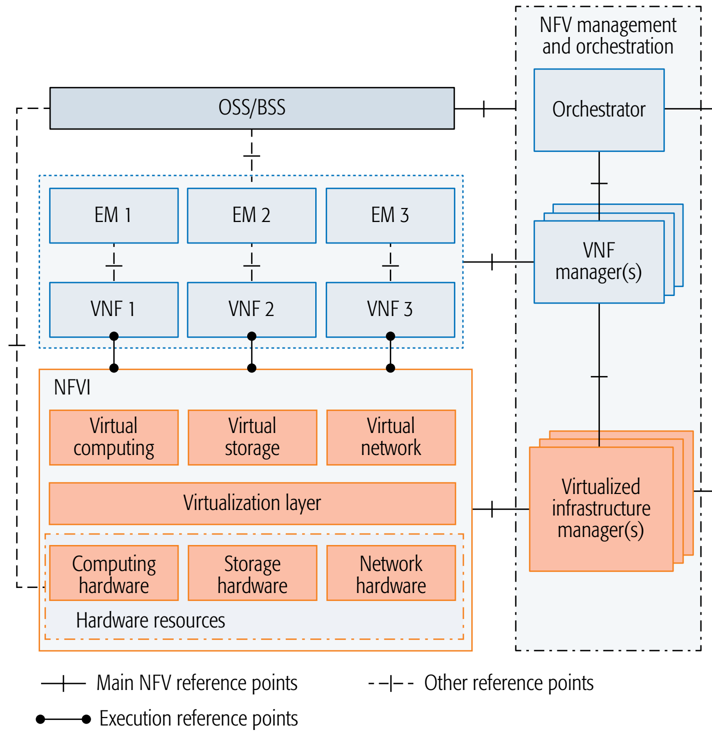

Figure 2. NFV architectural framework [8].

Requirements

The IMS was designed for 3G with human-centric applications in mind; however, 4G and beyond aim to cater for both human-centric applications and M2M applications (e.g., smart grid). This calls for a redesign of the IMS, and cloud computing is the ideal basis since it enables features such as scalability and efficiency in resource usage. We consider the following requirements to be the most pertinent for cloudifying the IMS for 4G and beyond.

Elastic Scalability: IMS today relies solely on using pre-allocated and overprovisioned functional entities to meet the expected demand peak. New capacity requires significant efforts to manually add new equipment to the system. On the other hand, a cloudified IMS should take advantage of the elasticity of the cloud to adapt dynamically to the growing or shrinking of the load requirements by adjusting the allocated resources in a fine-grained manner. Additionally, it should be able to smoothly handle a massive number of IMS UEs. Indeed, 10–100 times more devices are expected to be connected to 5G compared to today.

Latency: 4G and beyond will support a wide variety of human-centric and M2M applications that will tolerate different values of latency. Some of these applications can tolerate latencies on the order of a few seconds, while others have stricter latency requirements than what exist today. For instance, teleprotection is a mission-critical application for power utilities. It includes real-time monitoring and alerting functionalities that require transferring the messages with about 8 ms delay on the application layer [1]. The cloudified IMS should be able to support applications that require different levels of latency. This includes applications that have much stricter latency requirements

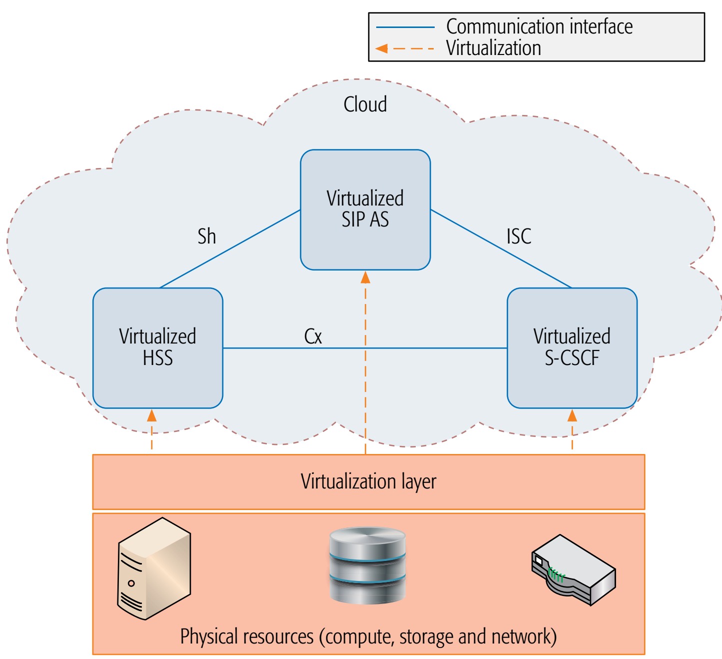

Figure 3. Simplified virtualized IMS.

compared to today. It should also be able to maintain the required latency under a high load.

Resource Efficiency: Today, IMS is installed with overprovisioning of resources to accommodate the peak demand. However, the shift toward on-demand capacity makes resource efficiency more critical, since inefficiency would be translated directly into higher running cost (i.e., OPEX) with the pay-per-use pricing model.

Architectures		Requirements			
		Elastic scalability	Laten-cy	Resource efficiency	Follow-me
Virtualized IMS	[9]	Partly	NO	Partly	NO
IMS as a Service (Virtualized-IMS)	[10]	Partly	NO	NO	NO
IMS as a Service (Split-IMS)	[10]	Partly	NO	NO	NO
IMS as a Service (Merge-IMS)	[10]	Partly	NO	NO	NO
HSS	[11]	NO	NO	NO	NO
HSS	[12]	NO	NO	NO	NO
Presence	[13]	Partly	NO	NO	NO
Presence	[14]	NO	NO	NO	NO

Table 1. Summary of the evaluated approaches vs. the identified requirements for cloudifying the IMS for 4G and beyond.

Follow-Me: The basic idea behind the "follow-me" concept is that cloud services follow end users as they move [8]. Mobile operators will use multiple IaaSs that are geographically distributed and interconnected [8]. IMS and IMS services could be deployed in different locations to offer better user experience. Therefore, as soon as the end user moves, the optimal application server for providing the IMS service may change. In the future, the service should follow the end user, and should always be accessed from the application server and through the IMS functional entities that ensure the best user experience. Nowadays, P-CSCF and S-CSCF entities are allocated to the IMS UE and do not change for the duration of the registration. Through this period, end users access their IMS services through these assigned entities. Therefore, to have service mobility in this model, the IMS UE should deregister from the assigned IMS entities and then register again, which causes service interruption.

Requirements are unfortunately often in conflict, and our proposed requirements are no exception to that tendency. Appropriate tradeoffs will need to be made when new architectures are designed. Let us illustrate this by demonstrating the conflicts between elastic scalability, latency, and resource efficiency. It is clear that today's granularity level (i.e., 3GPP functional entities) is an impediment to elastic scalability. However, refining that level of granularity through the splitting of the functional entities will usually lead to additional cost (e.g. management complexity, inter sub-functional entities communications). These costs may (or may not) offset the gains expected from the refinement. In addition, the splitting may prevent latency requirements from being met. Therefore, optimal splitting becomes the key. We further elaborate on this in the research directions section.

Approaches that Deal with the Entire IMS

This section reviews the approaches that focus on the whole IMS system in the light of the requirements set forth above. In these approaches, a common pool of resources is dynamically allocated to IMS functional entities. Figure 3 provides an illustration. The physical computation, storage, and networking resources are virtualized. This allows for an IMS with a set of interacting virtual functional entities (i.e., functional entities that rely on virtualized resources). Table 1 summarizes the review findings.

Virtualized IMS

In [9], Lu *et al.* propose a cloud platform for the IMS core network that runs IMS entities on cloud-based virtual machines (VMs). The proposed platform supports dynamic resource allocation and disaster protection. The proposed resource allocation algorithm can dynamically allocate and de-allocate virtual central processing unit (vCPU) and memory resources to VMs according to the current workload. The algorithm aims to allow the platform to satisfy the carrier-grade response time requirement, achieve high resource utilization, and reduce cost.

Additionally, the algorithm assumes that each VM boots with an initially allocated vCPU and memory. Each VM also has a predefined maximum amount of vCPU and memory that can be allocated. When the resource utilization exceeds a predefined threshold, the system adds one vCPU or more memory if the VM has not reached the maximum allowed resources. If the physical machine (PM) that hosts the VM does not have enough resources to scale the resources of the VM, the algorithm performs live migration of the VM to another PM with enough resources. The algorithm can also elastically scale the number of active PMs in the cloud infrastructure automatically according to the workload. It aims to achieve high resource utilization and reduce power consumption costs.

The proposed resource allocation algorithm can elastically scale IMS vertically to adapt to the workload, whereas horizontal scalability is not tackled. However, the stateful architecture for many of the IMS functional entities (e.g., S-CSCF) hinders the implementation of horizontal scalability. For instance, it would be difficult to terminate an S-CSCF instance when it handles an ongoing call because this would require transferring the stored state to another S-CSCF instance. The authors propose a resource allocation algorithm to achieve high resource utilization. However, resource efficiency may not be maximal since the optimal splitting is not considered and the default splitting (i.e., IMS functional entities as defined today) is used. The authors also do not evaluate the latency achieved by their architecture. Furthermore, they do not tackle follow-me requirement. However, it remains an issue in the proposed design due to the static assignment of IMS functional entities for a specific IMS UE in the registration process.

IMS as a Service

Carella *et al.* [10] propose three architectures for cloud-based virtualized IMS using NFV: Virtualized-IMS, Split-IMS, and Merge-IMS. In the Virtualized-IMS architecture, each IMS functional entity is implemented as software that runs on a single VM. The interfaces with external components are not changed. The Split-IMS moves the state of the subscribers, which is maintained in many IMS entities (e.g., P-CSCF and S-CSCF), to an external functional entity called Shared-Memory. This makes the IMS entities stateless. A load balancer is positioned as an entry point for the new stateless entities to distribute the load.

Additionally, the Merge-IMS architecture groups the four main entities of IMS (i.e., P-CSCF, S- CSCF, I-CSCF, and HSS) and deploys them into one VM called IMS-VM. It introduces the IMS-Locator entity, which assigns the subscribers to a particular IMS-VM instance during the registration process. All HSS entities in IMS-VM instances share the same database to store subscriber information.

The Virtualized-IMS architecture can scale using the procedures already standardized by 3GPP to some extent. However, the scalability is limited due to the stateful architecture. The Split-IMS architecture separates functional entities' logic and state, so the logic can scale easily by instantiating new stateless entities and adding them to the load balancer. However, the scalability is not fine-grained. For instance, to scale HSS, a full-fledged HSS (e.g., storage and all reference points) should be instantiated. It is also important to verify the optimality of the proposed splitting, and how it affects the performance and resource efficiency. Moreover, the Merge-IMS architecture scales by creating a new IMS-VM that has all items (i.e., full-fledged IMS). Thus, elastic scalability is limited since it is difficult to scale in due to the granularity level (i.e., IMS-VM) and stateful architecture.

The authors do not provide performance metrics to evaluate the architectures' latency. They also do not tackle the optimal splitting of the IMS functional entities. Therefore, it is hard to assess whether resource efficiency could be met. Although the follow-me requirement is not tackled, none of the proposed architectures can satisfy it without re-architecting the IMS.

Approaches that Deal with Specific IMS Entities

This section reviews the approaches that focus on specific IMS entities. The main IMS entities that have attracted the attention of researchers are the HSS and the presence service. This is probably due to the fact that they are much less complex than other nodes such as the CSCF. Table 1 summarizes the review findings.

HSS

Few works propose virtualized and cloud-based HSS architectures. Yang *et al.* [11] propose the distribution of the HSS into a resource layer and a management layer. The resource layer is implemented in the cloud, and simulations are performed to demonstrate performance gains. Although their proposed solution enables independent scaling of resource and management layers, elastic scalability is not tackled. It is also not possible to evaluate resource efficiency as the optimal splitting is not tackled. The performed experiment shows that the latency is high. Furthermore, follow-me is not tackled in their work.

In [12], Paivarinta *et al.* use a home location register (HLR) to evaluate whether cloud technologies can meet the carrier-grade requirements. The HLR was the primary subscriber database for mobile networks up to the 3GPP Release 4 standards, and today is considered a subset of the HSS. The proposed architecture uses the HBase NOSQL database as HLR storage and deploys it on Amazon EC2 IaaS. It utilizes the telecommunication application transaction processing (TATP) benchmarking tool to measure the performance of the HBase database under load, which is typical in telecommunications. Unfortunately, the authors do not tackle elastic scalability. They also only discuss the storage of HLR and do not discuss the HLR application logic. Thus, it is not possible to evaluate the resource efficiency. In addition, the performed experiment shows that the latency increases proportionally to the throughput. The follow-me requirement has not been tackled.

Presence Service

In [13], Belqasmi *et al.* propose an early architecture for a virtualized presence service for the future Internet. Although scalability is not tackled, it is ensured through the use of presence service substrates. However, the authors do not tackle the level of granularity of the substrates. It is therefore rather difficult to assess whether the architecture could scale in a fine-grained manner and whether resource efficiency could be ensured. The latency and follow-me requirements are not tackled.

Quan *et al.* [14] also focus on presence service. They propose a cloud-based implementation of presence service. The Eucalyptus cloud open source software is used, and the whole presence server is deployed on a VM. The authors do not tackle elasticity scalability, optimal splitting, and follow-me. Moreover, the evaluation shows that the architecture's latency is high and increases proportionally to the throughput.

Research Directions

Research on IMS cloudification has started. This section provides insightful directions for future studies. In this article, we focus on two research issues as illustrations. In the first section, we discuss challenges related to the IMS granularity level. This discussion includes both architectural and algorithmic issues. The second section focuses on PaaS for IMS.

Reconsidering the Granularity Level of IMS

Each IMS network functional entity as defined by 3GPP contains a set of functions as one deployable and scalable unit. These entities are often stateful, which hinders elastic scalability and resiliency in the cloud. We believe it is worthwhile, in the cloud environment, to investigate the possibility of having finer granularity for IMS network functional entities to achieve finer control, elastic scalability, and better resiliency.

A good starting point may be to separate the functional entities' logic and data (or state). It should be noted that 3GPP has also stipulated this separation [15] primarily for data consistency purposes, but has also mentioned better scalability as a potential advantage. This brings about the challenge of leveraging cloud technologies (e.g.,

Each IMS network functional entity as defined by 3GPP contains a set of functions as one deployable and scalable unit. These entities are often stateful, which hinders elastic scalability and resiliency in the cloud. We believe it is worthwhile, in the cloud environment, to investigate the possibility of having finer granularity for IMS network functional entities to achieve finer control, elastic scalability, and better resiliency.

A key open issue in PaaS is the aspiration for a standard language to describe the SaaS services. This language should be able to describe the structure of these services and their management aspects. It should support the deployment and management across multiple IaaSs so that functional entities could be deployed at different locations.

distributed cache) to ensure equivalent performance characteristics.

A next step will be to consider decomposing the IMS network functional entities' logic into smaller sub-functional entities, leading to finer control over the distinct functions. However, this decomposition may not come cost-free. Indeed, it increases the management complexity and may have a negative impact on latency. The cloud can help to alleviate the management complexity by leveraging PaaS to automate IMS's life cycle. However, many challenges have to be addressed at the PaaS level to make it a reality. We elaborate more on this in the next section. As for latency, placement algorithms are a potential avenue to effectively minimize the latency and cross-network traffic. For instance, the algorithms may place the related functions on VMs hosted on the same physical server, so they communicate through a virtual switch, which leads to lower latency compared to the communication over the network.

The decomposition also gives rise to architectural and algorithmic challenges. At the architecture level, if the new sub-functional entities interact with each other, there is a need to design new interfaces. The interfaces should be very lightweight to minimize the extra cost induced by the communication. On the other hand, they also need to be reliable and scalable.

At the algorithmic level, there is a need to identify the optimal granularity for the sub-functional entities that can achieve the intended benefits (if possible). A key challenge is to determine the fine-grained atomic operations of each coarse-grained IMS network functional entity (e.g., HSS and presence server), and the degree of relationship between these operations and the associated cost (e.g., memory and processing). Resource inefficiency could be translated into the cost of unused resources for given operations. This could, for instance, be translated into a graph theory problem with a weighted undirected graph formed by representing atomic operations as vertices. In this model, two vertices would be joined by an edge if they are related and need to communicate. It could then be solved by formulating it as an optimal clustering problem where each cluster is represented as a set of vertices. The objective would be to maximize the sum of intra-cluster communication cost, minimize the sum of resource costs of all clusters, and minimize the sum of the inter-cluster communication cost. Of course, there would be constraints such as latency. The optimal clustering problem can be solved for each coarse-grained functional entity independently. It can be shown that the optimal clustering problem is non-deterministic polynomial-time hard (NP-hard) when the number of atomic operations is large and hence requires efficient heuristics to solve it. The design of these heuristics is an important research direction. More importantly, clustering algorithms such as hierarchical clustering and K-means clustering can be modified to solve the optimal granularity problem.

Toward a PaaS for IMS

The telco industry could leverage PaaS to deliver IMS network functional entities (e.g., CSCFs, HSS, presence) or a subset (e.g., only HSS) as SaaS services with pay-per-use pricing to end users (i.e., IMS UE) or even to other SaaS services. The PaaS would automate the life cycle of the functional entities from deployment to management (e.g., monitoring, auto-scaling, and auto-healing) and orchestration. For telco, the PaaS would need to run on multiple geographically distributed IaaSs that are interconnected by a wide area network (WAN). This would help ensure the service continuity and reduce latency by deploying closer to end users.

A key open issue in PaaS is the aspiration for a standard language to describe the SaaS services. This language should be able to describe the structure of these services (i.e., functional entities, relationships, requirements, etc.), and their management aspects (e.g., deployment, monitoring, scaling). It should support the deployment and management across multiple IaaSs so that functional entities could be deployed at different locations. PaaS could use the services' description to automate their life cycle. The topology and orchestration specification for cloud applications (TOSCA) [16] may be a good starting point. It is standard to describe cloud applications by means of topology templates and management plans. However, the current TOSCA version (1.0) does not support all management aspects needed in telco, such as monitoring.

Another research challenge is the elastic scaling of the SaaS services offered by the PaaS. These services often consist of multiple interconnected functional entities that could be distributed across multiple IaaSs. The traditional scaling approaches in PaaS usually scale the overloaded entity itself without considering the impact on other entities in the service. In telco, these approaches would not be sufficient and efficient since there could be a need in many cases to scale and optimize other entities in the service. In fact, there is a need for new smart scaling approaches that consider the end-to-end service (i.e., all entities in the service) and are aware of service requirements (e.g., latency and resiliency) and the surrounding environment (e.g., resource availability and network traffic status). These approaches should evaluate the impact of scaling and then decide accordingly what to scale, where to scale (the same IaaS or across multiple IaaSs), and what to optimize, aiming to meet the service requirements.

The PaaS includes management and orchestration functions (e.g., monitoring, fault management, and scaling). These functions are responsible for automating the life cycle of the functional entities. However, there will be many challenges in designing them in a distributed environment. One challenge is related to the architecture and whether it is centralized or distributed. The centralized architecture is simpler, but will suffer from nontrivial latency in detecting problems and making decisions. On the other hand, the distributed architecture has lower latency. However, it is more complex and gives rise to the challenge of maintaining end-to-end service visibility. Another challenge is related to the capacity of these functions. The number of functional entities that need to be managed changes over time as the services scale elastically. Thus, the capacity of these functions should elastically scale to adapt to system workload. This requires clear definitions of the key performance indicators needed to manage the capacity.

Another open issue is network orchestration. To the best of our knowledge, current IT PaaS solutions use networks with best effort delivery. On the other hand, quality of service (QoS) is a requirement in Telco to guarantee the performance (e.g., latency) required by the applications (e.g., multimedia applications). Indeed, telco PaaS should leverage the IaaS network and WAN capabilities to interconnect the deployed functional entities (perhaps across multiple IaaSs) using a transport network that meets specific requirements (e.g., latency and bandwidth). This requires that both IaaS and WAN support advanced networking capabilities (e.g., QoS) and exposes them via northbound interfaces.

Conclusions

In this article, we identify the most pertinent requirements for cloudifying the IMS for 4G and beyond. We also review the architectures proposed thus far for the cloudification of IMS using the identified requirements. These architectures are classified into two categories: the first focuses on the whole IMS system, and the second deals with specific IMS functional entities. Our evaluation shows that the existing literature does not meet the requirements of cloudifying the IMS for 4G and beyond. Subsequently, we outline some interesting research issues that still need to be resolved. We discuss the possibility of decomposing IMS functional entities to achieve elastic scalability and better resilience in cloud settings. We also discuss the main challenges resulting from this decomposition, such as the need for new communication interfaces and optimal granularity. Furthermore, we identify many challenges at the PaaS level. One challenge is the lack of a standard language that can describe the IMS structure and management aspects. Another challenge is the design of elastic management and orchestration functions in a distributed environment.

Acknowledgment

This work is supported in part by Ericsson and the National Science and Engineering Research Council (NSERC) of Canada.

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] 3GPP TS 23.228, "IP Multimedia Subsystem (IMS); Stage 2." v. 12.8.0, 2015.

[3] F. Gouveia et al., "Cloud Computing and EPC/IMS Integration: New Value-Added Services on Demand," *Proc. 5th int'l. ICST Mobile Multimedia Commun. Conf.*, 2009, p. 51.

[4] W. Zhang *et al*., "Architecture and Key Issues of IMS-Based Cloud Computing," *Proc. 2013 IEEE Sixth Int'l. Conf. Cloud Computing*, 2013, pp. 629–35.

[5] P. Agrawal *et al.*, "IP Multimedia Subsystems in 3GPP and 3GPP2: Overview and Scalability Issues," *IEEE Commun. Mag*., vol. 46, no. 1, Jan. 2008, pp. 138–45.

[6] B. Han *et al*., "Network Function Virtualization: Challenges and Opportunities for Innovations," *IEEE Commun. Mag*., vol. 53, no. 2, Feb. 2015, pp. 90–97.

[7] ETSI GS NFV, "Network Functions Virtualization (NFV); Architectural Framework." v. 1.2.1, 2014.

[8] T. Taleb and A. Ksentini, "Follow Me Cloud: Interworking Federated Clouds and Distributed Mobile Networks," *IEEE Network*, vol. 27, no. 5, Sept. 2013, pp. 12–19.

[9] F. Lu *et al.*, "A Virtualization-Based Cloud Infrastructure for IMS Core Network," *Proc. 2013 IEEE 5th Int'l. Conf. Cloud Computing Technology and Science*, 2013, vol. 1, pp. 25–32.

[10] G. Carella *et al*., "Cloudified IP Multimedia Subsystem (IMS) for Network Function Virtualization (NFV)-Based Architectures," *Proc. 2014 IEEE Symp. Computers and Communication*, 2014, vol. Workshops, pp. 16.

[11] T. Yang *et al.*, "A New Architecture of HSS Based on Cloud Computing," *Proc. 2011 IEEE 13th Int'l. Conf. Commun. Technology*, 2011, pp. 526–30.

[12] R. Paivarinta and Y. Raivio, "Applicability of NoSQL Databases to Mobile Networks: Case Home Location Register," *Cloud Computing and Services Science*, I. Ivanov, M. van Sinderen, and B. Shishkov, Eds., Springer, 2012, pp. 225–42.

[13] F. Belqasmi, N. Kara, and R. Glitho, "A Novel Virtualized Presence Service for Future Internet," *Proc. 2011 IEEE ICC Workshops*, 2011, pp. 1–6.

[14] W. Quan et al., "Research of Presence Service Testbed on Cloud-Computing Environment," *Proc. 2010 3rd IEEE Int'l. Conf. Broadband Network and Multimedia Technology*, 2010, pp. 865–69.

[15] 3GPP TS 23.335, "LTE; User Data Convergence (UDC); Technical Realization and Information Flows; Stage 2." v. 11.0.0 Release 11, 2012.

[16] OASIS Committee Specification, "Topology and Orchestration Specification for Cloud Applications," v. 1.0, Nov. 2013.

Biographies

Mohammad Abu-Lebdeh received his B.Sc. degree in computer engineering from An-Najah National University, Palestine, and his M.Sc. degree in electrical and computer engineering from Concordia University, Canada. He is currently pursuing his Ph.D. degree in information and systems engineering at Concordia University. In the past, he worked for several years as a software engineer. His current research interests include cloud computing, service engineering, and next generation networks.

Jagruti Sahoo holds a Ph.D. degree in computer science and information engineering from National Central University, Taiwan. She is currently a postdoctoral fellow at Concordia University. In the past, she worked as postdoctoral fellow in the Department of Electrical and Computer Engineering, University of Sherbrooke, Canada. Her research interests include cloud computing, network functions virtualization, content delivery network,s and wireless sensor networks.

Roch Glitho holds a Ph.D. (Tekn. Dr.) in tele-informatics (Royal Institute of Technology, Stockholm, Sweden), and M.Sc. degrees in business economics (University of Grenoble, France), pure mathematics (University of Geneva, Switzerland), and computer science (University of Geneva). He is an associate professor and Canada Research Chair at Concordia University. He is also an adjunct professor at several other universities including Telecom Sud Paris, France, and the University of Western Cape, South Africa. In the past, he has worked in industry and has held several senior technical positions (e.g., senior specialist, principal engineer, expert) at Ericsson in Sweden and Canada. His industrial experience includes research, international standards setting, product management, project management, systems engineering, and software/firmware design. He has also served as an IEEE Distinguished Lecturer, Editor-In-Chief of *IEEE Communications Magazine*, and Editor-In-Chief of *IEEE Communications Surveys & Tutorials Journal*.

Constant Wette Tchouati is a senior system developer in the Cloud & IP Business Unit, Ericsson Canada. He holds an M.Sc. in computer engineering from Ecole Polytechnique de Montreal, Canada, an M.Eng. In electronics and communication from Ecole Polytechnique de Yaounde, Cameroon, and an M.B.A. in management from HEC Montreal. He joined Ericsson in 2000, and has been involved in research, product development, innovation projects, and new business development in different areas including cloud architectures, data management and analytics, IoT/M2M, IMS, and social media.

We identify man challenges at the PaaS level. One challenge is the lack of a standard language that can describe the IMS structure and management aspects. Another challenge is the design of elastic management and orchestration functions in a distributed environment.

Management and Orchestration Challenges in Network Functions Virtualization

Rashid Mijumbi, Joan Serrat, Juan-Luis Gorricho, Steven Latré, Marinos Charalambides, and Diego Lopez

NFV continues to draw immense attention from researchers in both industry and academia. By decoupling NFs from the physical equipment on which they run, NFV promises to reduce CAPEX and OPEX, make networks more scalable and flexible, and lead to increased service agility. However, despite the unprecedented interest it has gained, there are still obstacles that must be overcome before NFV can advance to reality in industrial deployments, let alone delivering on the anticipated gains.

Abstract

NFV continues to draw immense attention from researchers in both industry and academia. By decoupling NFs from the physical equipment on which they run, NFV promises to reduce CAPEX and OPEX, make networks more scalable and flexible, and lead to increased service agility. However, despite the unprecedented interest it has gained, there are still obstacles that must be overcome before NFV can advance to reality in industrial deployments, let alone delivering on the anticipated gains. While doing so, important challenges associated with network and function MANO need to be addressed. In this article, we introduce NFV and give an overview of the MANO framework that has been proposed by ETSI. We then present representative projects and vendor products that focus on MANO, and discuss their features and relationship with the framework. Finally, we identify open MANO challenges as well as opportunities for future research.

Introduction

In recent years, telecommunication service providers (TSPs) have experienced constant dwindling in revenue. This has been attributed, in part, to two main factors. On one hand, the seemingly insatiable traffic demands of subscribers require physical network expansions, which are achieved at increased capital expenditures (CAPEX) and operating expenditures (OPEX) [1]. On the other hand, competition both among themselves and from over-the-top service providers means that TSPs cannot respond to the increased costs with increased subscriber fees.

Network functions virtualization (NFV) [2] has been identified as a potential solution to these problems. The main concept of NFV is the decoupling of network functions (NFs) from capacity (the physical infrastructure on which they run). Breaking the bond between NFs and hardware promises several advantages. First, there is a potential for significant reductions in OPEX through more efficient operations, since most maintenance and updates to NFs can be performed remotely and at scale. In addition, the increased flexibility can lead to more efficient utilization of resources and hence reductions in CAPEX, since TSPs could use the existing network capacity for more user traffic. Finally, NFV may lead to better service agility by allowing TSPs to deploy and/or support new network services faster and less expensively.

These expectations have made NFV a burgeoning research field. The most notable NFV activities are being led by the European Telecommunications Standards Institute (ETSI). The main objective of the ETSI NFV group[1] is to develop standards for NFV as well as share experiences of its development and early implementation. To this end, they have defined the NFV problem, some use cases, a reference architecture, and a management and orchestration (MANO) framework, among other items [3].

However, while a lot of progress has been made, there are still many technical challenges, which must be overcome before the gains anticipated from NFV can come to fruition. Among them, MANO challenges have drawn special attention. The reason behind this interest is that MANO is a critical aspect in ensuring the correct operation of the NFV infrastructure (NFVI) as well as virtual network functions (VNFs). MANO provides the functionality required for the provisioning of VNFs, and related operations such as the configuration of VNFs and the infrastructure on which these functions run. It includes the orchestration and life cycle management of physical and/or virtual resources that support the VNFs [4]. Just like the decoupled NFs, NFV demands a shift from management models that are device-driven to those that are aware of the orchestration needs of NFs running in a virtualized environment. We believe that for NFV to succeed, the main MANO challenges should be addressed at the current specification phase, rather than later when real large-scale deployments commence.

In this article, we survey current efforts that address NFV MANO. In particular, in the next section, we begin by summarizing the MANO framework that has been proposed by ETSI. We then overview representative projects and vendor products that focus on NFV MANO. We classify these projects and products in two ways. First, we map their functionality to the functional blocks of the ETSI MANO framework, and then we study their features based on four criteria:

- Management approach (centralized, distributed, policy-based, self-managed)

[1] http://www.etsi.org/technologies-clusters/technologies/nfv

Rashid Mijumbi is with Waterford Institute of Technology; Joan Serrat and Juan-Luis Gorricho are with Universitat Politécnica de Catalunya; Steven Latré is with the University of Antwerp and iMinds; Marinos Charalambides is with University College London; Diego Lopez is with Telefonica I+D. Part of this work was done while R. Mijumbi was still with Universitat Politècnica de Catalunya.

0163-6804/16/$25.00 © 2016 IEEE

- Management functions supported (fault, configuration, accounting, performance, and security: FCAPS)
- Scope (functions, services, network)
- The integration of management with that of software defined networking (SDN) and cloud computing, both of which are complimentary and/or enablers of NFV

Finally, we identify open challenges and discuss opportunities for future research, before concluding the article.

ETSI MANO Framework

The ETSI MANO framework [4] is shown in Fig. 1. The functional blocks in the framework can be grouped into three main entities:

- NFV architectural layers
- NFV management and orchestration
- Network management systems

These entities, as well their constituent functional blocks, are connected together using a set of defined reference points.[2] The NFV architectural layers include the NFVI and VNFs. NFVI is the combination of both hardware and software resources that makes up the environment in which VNFs are deployed, while VNFs are implementations of NFs that are deployed on those virtual resources.

NFV Management and Orchestration

The NFV MANO consists of three functional blocks: the virtual infractructure manager (VIM), NFV orchestrator (NFVO), and VNF manager (VNFM); and four data repositories: NS catalog, VNF catalog, NFV instances, and NFVI resources.

VIM: A VIM manages and controls NFVI physical and virtual resources in a single domain. This implies that an NFV architecture may contain more than one VIM, with each of them managing or controlling NFVI resources from a given infrastructure provider. In principle, a VIM may be specialized in handling a certain type of NFVI resource (e.g., compute-only or storage-only), or could manage multiple types of NFVI resources (e.g., nodes in the NFVI).

VNFM: Each VNF instance is assumed to have an associated VNFM. The VNFM is responsible for the management of the life cycle of VNFs. A VNFM may be assigned the management of a single or multiple VNF instances of the same or different types, including the possibility of a single VNFM for all active VNF instances for a certain domain.

NFVO: The NFVO aims to combine more than one function to create end-to-end services. To this end, the NFVO functionality can be divided into two broad categories: resource orchestration and service orchestration. The first is used to provide services that support accessing NFVI resources in an abstracted manner independent of any VIMs, as well as governance of VNF instances sharing resources of the NFVI. Service orchestration deals with the creation of end-to-end services by composing different VNFs, and the topology management of the network service instances.

Data Repositories: Data repositories are databases that keep different types of information in the NFV MANO. Four types of repositories can be considered:

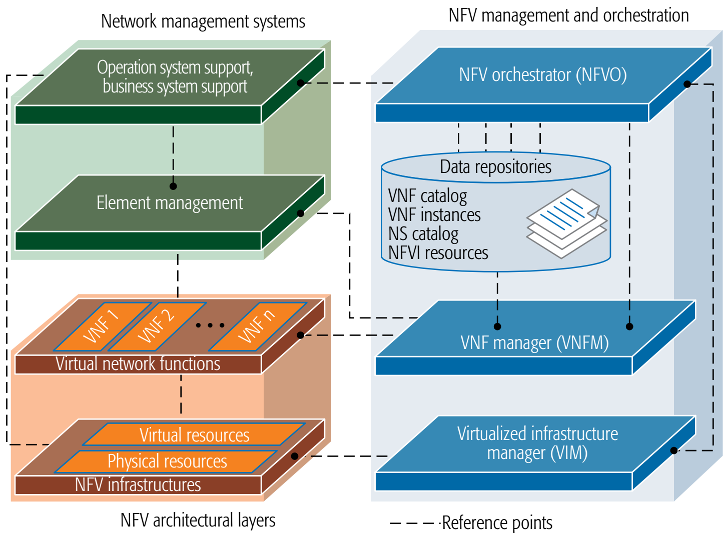

Figure 1. ETSI NFV MANO Framework.

- The NS catalog is a set of predefined templates, which define how services may be created and deployed, as well as the functions needed for the service and their connectivity.
- The VNF catalog is a set of templates that describe the deployment and operational characteristics of available VNFs.
- The NFVI resources repository holds information about available/allocated NFVI resources.
- The NFV instances repository holds information about all function and service instances throughout their lifetimes.

Network Management Systems

NFV is not intended to require a drastic change in the current mechanisms of network service provisioning, and is aimed at a gradual transition from network infrastructures based on physical nodes, easing the integration in a heterogeneous environment. Therefore, network management systems will continue to have a key role in NFV, coordinated with the MANO entities by means of a clear separation of roles. MANO entities will deal with those aspects related to the virtualization mechanisms, while network management functions are expected to take care of the features associated with the semantics of the specific network services being provided by the composition of VNFs and, potentially, physical nodes. These network management systems include element management (EM), operation system support (OSS), and business system support (BSS).

Projects Related to NFV MANO

CloudNFV

CloudNFV[3] is an open platform for implementing NFV based on cloud computing and SDN. The CloudNFV architecture, illustrated in Fig. 2, is made up of three main elements: active virtualization, NFVO, and NFVM. Active virtualization is a data model (based on TM Forum's SID [5]), which represents all aspects of services, functions, and resources. It is made up of

[2] A reference point is a conceptual point at the conjunction of two communicating functional entities. A detailed description of all the reference points in the ETSI NFV MANO framework can be found in [4].

[3] www.cloudnfv.com/WhitePaper.pdf

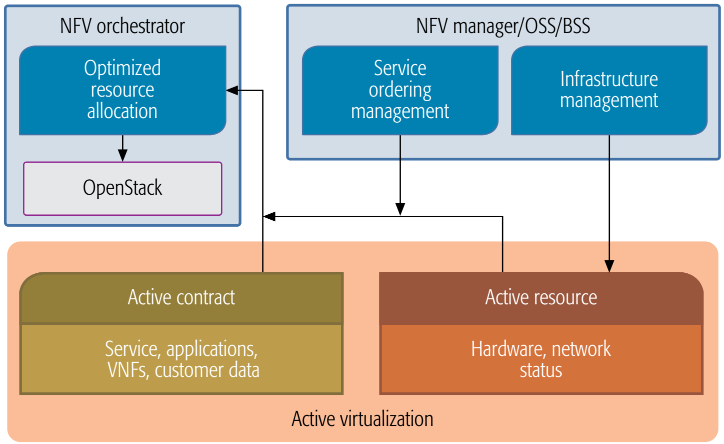

Figure 2. CloudNFV Architecture.

an active contract and active resource. Active resource describes the status of all resources in the infrastructure, while active contract includes all service templates that define the characteristics of all the available NFs. The orchestrator has policy rules, which, combined with service orders and the status of available resources, determine the location of the functions that make up the service as well as connections between them.

After service deployment all resources report their status and traffic to the active resource. The management processes running against active resources allow reflection of this status using management information bases (MIBs). The main difference between the ETSI NFV MANO and CloudNFV is that unlike the former, the latter considers both management and orchestration as applications that can run off a unified data model.

ExperiaSphere

ExperiaSphere[4] is a MANO model for NFV, which is based on a combination of open source tools. ExperiaSphere is founded on the concept of service models. These define how resources expose service features and how the functions decompose to resources. The service models are then used by a broker who selects one or more required service models to create a service instance. Once the service instance is created, its status is tracked throughout its life cycle. To achieve these objectives, ExperiaSphere is based on two principles: structured intelligence and derived operations.

Structured intelligence uses an integration of Universal Service Definition Language (USDL) and topology and orchestration specification for cloud applications (TOSCA) [6] to define the relationship between service elements, service goals, and the infrastructure. Derived operations allow virtualized services and resources to be managed as if they were physical. The management functions operate on virtual elements of the service, using variables defined per element but derived from the state of real resources.

OpenMANO

OpenMANO [7] is an open source project led by Telefonica, which is aimed at implementing the ETSI NFV MANO framework, and addressing the aspects related to performance and portability by applying Enhanced Platform Awareness (EPA)[5] principles. As shown in Fig. 3, the OpenMANO architecture consists of three main components: openmano, openvim, and a graphical user interface (GUI). In addition, there are two command line interfaces (CLIs) used to interact with openmano and openvim.

openvim is a lightweight, NFV-specific VIM implementation directly interfacing with the compute and storage nodes in the NFVI, and with an openflow controller in order to create the infrastructural network topology, and enforce the EPA principles mentioned above. It offers a REST-based northbound interface (openvim application programming interface: API) to openmano, where enhanced cloud services are offered including the life cycle management of images, flavors, instances, and networks. The openvim API extends the OpenStack API to accommodate EPA. OpenMANO has a northbound interface (openmano API) based on REST, where MANO services are offered including the creation and deletion of VNF templates, VNF instances, network service templates, and network service instances.

OPNFV

OPNFV[6] is an open source project founded and hosted by the Linux Foundation, and composed of TSPs and vendors. The objective is to establish a carrier-grade integrated open source reference platform that may be used to validate multi-vendor interoperable NFV solutions. OPNFV plans to validate existing standard specifications, contribute improvements to relevant upstream open source projects, and develop necessary new functionality within both OPNFV and upstream projects. In particular, it is focused on implementing the NFV requirements provided by ETSI. To this end, the first outcome of the project, referred to as OPNFV Arno, was released in June 2015. Arno is an initial build of the NFVI and VIM components of the ETSI architecture.

ZOOM

ZOOM[7] is a TM Forum project aimed at enabling the deployment of services by automating the provisioning process through improved OSS/BSS models. To achieve this, the project regularly conducts a range of hands-on technology demos each of which is developed from what they call a *catalyst project*. Each catalyst project is sponsored by one or more network operators and equipment and software vendors in a real-world demo. The project currently runs about nine catalysts with a focus on NFV aspects including end-to-end automated management of hybrid networks, and demonstrating the impact and value of dynamic security orchestration in an NFV environment.

Pre-Standardization NFV MANO Products

CloudBand

Alcatel-Lucent's CloudBand[8] is an NFV platform comprising software and hardware stacks with two elements: a node and a management system. The CloudBand node provides the com-

[4] http://www.experiasphere.com/

[5] https://01.org/sites/default/files/page/openstack-epawpfin.pdf

[6] https://www.opnfv.org/

[7] https://www.tmforum.org/collaboration/catalyst-program/current-catalysts/

puting, storage, and networking hardware to host cloud services, while the management system is the MANO element of CloudBand. The management system aggregates distributed cloud resources — the nodes — to provide a view of the entire NFVI as a single pool. It orchestrates, automates, and optimizes VNFs across the service provider's network and data centers.

Ensemble Service Orchestrator

Overture's Ensemble Service Orchestrator (ESO)[9] is an NFV service and VNF life cycle management and orchestration system. The system coordinates and connects virtual resources to physical network elements to create virtualized services across multiple networking layers. ESO supports the placement of VNFs in centralized as well as distributed data centers. It uses the OpenStack cloud controller to manage the virtual computing environment, including virtual machines, virtual switches, and data center switches. ESO is the key component of Overture's Ensemble Open Service Architecture (OSA), which is a framework for service MANO.

OpenNFV

HP's OpenNFV[10] is an NFV platform that leverages open source technology to provide an open end-to-end NFV and SDN infrastructure. OpenNFV is aligned toward providing solutions to each of the functional blocks defined in the ETSI NFV reference architecture. With regard to MANO, OpenNFV includes three solutions; NFV director, NFV manager, and Helion OpenStack.

The NFV director is an NFVO that can be used to automate the deployment and monitoring of a VNF ecosystem. Its aim is to ensure that each VNF can efficiently run on heterogeneous hardware platforms and virtualization environments. VNF managers are responsible for the VNFs life cycle actions (e.g., by deciding to scale up or down). The Helion OpenStack provides an open source cloud platform for running VNFs.

Open Network Strategy

Cisco's Open Network Strategy (OPN)[11] includes a services orchestrator, a VNFM, and an SDN controller, all of which are aimed at providing implementations for some of the functional blocks of the ETSI MANO framework. The services orchestrator is responsible for providing the overall life cycle management at the network service level. It uses model-based workflows that enable the design of services based on predefined service elements. The VNFM provides VNF life cycle management, including the creation, provisioning, and monitoring of both Cisco and third-party VNFs. Finally, the SDN controller is responsible for connecting the virtualized services to the service provider VPNs, the Internet, or both.

Planet Orchestrate

Planet Orchestrate is part of Cyan's Blue Planet Suite,[12] which is an SDN and NFV platform aimed at service orchestration, automation, SDN control, and multi-vendor management capabilities. Its functionality is based on the requirements of ETSI's NFV MANO framework. It uses TOSCA templates and information models to define service components and their relationships.

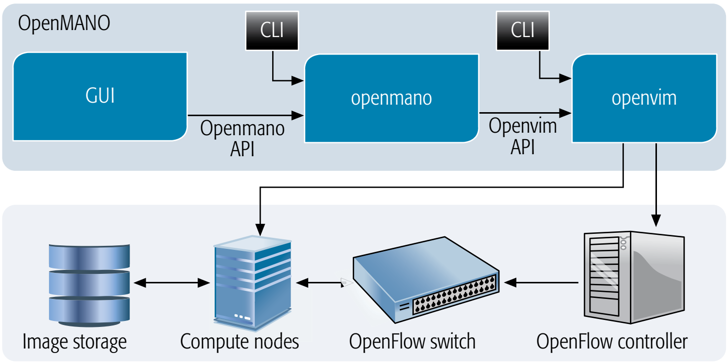

Figure 3. Telefonica's OpenMANO Architecture.

Planet Orchestrate can perform VNF management and orchestration functionality. The NFV orchestration engine performs placement of VNFs and supports distributed NFVI to optimize use of NFV resources. On the management side, it supports the performance, availability, and security demands of service provider applications. Performance monitoring and alarm/event reporting is provided for the NFVI and virtual functions. Intrinsic knowledge of the topology and the mapping between application and virtual resources enables fault isolation and recovery as well as high availability and resiliency.

Summary: In Tables 1 and 2, we summarize current activities toward NFV MANO. In Table 1, we map the functionalities of each project or product to the functional blocks of the ETSI NFV reference architecture. We can observe that most projects or products choose to rely on existing infrastructures and cloud systems such as OpenStack for achieving the NFVI and some form of data modeling and storage to model and store VNFs. On the MANO side, it can be noted that almost all propose a solution for each of the three functional blocks, VIM, VNFM, and NFVO. The difference, however, is in the functionality. This can be observed in Table 2, which summarizes the management and orchestration functionality of each project/product based on their description. For this purpose, we define four functionality categories. The management approach classifies them based on whether they are centralized, distributed, policy-based, or automated (self-managed). The management function classifies them according to their support for five of the basic management functions — FCAPS. We define the management scope to include functions, services, and networks. An approach is classified as a network management one if it proposes functionality to manage network nodes and links, while service management applies to an approach that manages both functions and their connectivity or chaining to form a service. Finally, since SDN and cloud computing are very important technologies with regard to NFV, we also categorize a project/product based on its ability to manage the interactions between SDN and/or cloud and NFV.

[8] http://resources.alcatel-lucent.com/asset/180265

[9] http://www.overturenetworks.com/products/network-virtualization/

[10] http://www8.hp.com/us/en/cloud/nfv-architecture.html

[11] http://www.cisco.com/c/en/us/solutions/collateral/service-provider/network-functions-virtualization-nfv/white-paper-c11-732123.html

[12] http://www.cyaninc.com/products/blue-planet-sdn-platform

	NFV architectural layers		NFV MANO framework				Traditional management systems	
	NFVI	VNFs	VIM	VNFM	NFVO	Data repositories	OSS/BSS	EM
CloudBand	Nuage, RedHat, CloudBand	VNF Modelling (TOSCA)	CloudBand node	CloudBand Management System	CloudBand Management System	✓		
CloudNFV	Active eesource	Active Contract	Infrastructure manager	OSS/BSS	✓	Active Contract	✓	OSS/BSS
ESO		✓	Ensemble network controller (ENC)	ESO	ESO	Database		
Experia-Sphere	Resource somain	TOSCA, USDL	Infrastructure manager	State-action service life cycle management	State-action service life cycle management	Derived operations	State-action service life cycle management	Derived operations
OpenMANO	✓	✓	Openvim		OpenMANO			
OPN			SDN Overly Controller	✓	Services orchestrator			
OpenNFV	✓		HP Helion Open-Stack Carrier Grade	✓	HP NFV director	HP NFV director	✓	
OPNFV	✓		✓					
Planet Orchestrate				✓	✓			
ZOOM			✓	✓	✓	Shared catalog	Order, SLA, and billing management systems	

Table 1. Mapping of state-of-the-art projects and products to the ETSI MANO.

Challenges and Research Opportunities

Resource Management

The servers used to host VNFs have a finite amount of memory, compute, and storage capacity. And since in practice these servers may be distributed across multiple domains, inter-domain link capacity will also be finite. Therefore, to achieve the economies of scale expected from NFV, physical resources should be efficiently managed. Dynamism, scalability, and automation are important features of such resource management. In this context, we identify three main challenges as listed below.

NFV PoP Locations: The first item is determining the locations of NFV points of presence (PoPs) [3]. In cases where the VNFs will be hosted in operator network nodes, it is necessary to decide which subset of the nodes can be used as NFV PoPs. As this does not have to be done often, it could be formulated as an optimization problem with the objective of considering the (latency to the) location of subscribers, and the setup and maintenance costs of both servers as well as fronthaul links in the case of virtualized radio access networks.

Function Placement: In order to compose a service, its constituent functions must be deployed. Decisions must be made on where functions should be placed among the available PoPs. This problem is related to the virtual network embedding (VNE) [8], and similar approaches may be applied. To this end, it may be formulated as a mathematical problem with such objectives as load balancing and energy conservation. However, any such formulation should be able to take the function chaining and/or precedence requirements into consideration to avoid network congestion.

In addition, while most current NFV proofs of concept have been based on each VNF being hosted by a dedicated VM, such an approach would not scale, especially for light functions such as those that are part of customer premises equipment. In this case, it would be more efficient to host multiple functions in a single VM by use of docker containers. In this case, there would be a need for scheduling approaches (e.g., [9]) for allocating the VM resources.

Dynamic Resource Management: One of the selling points of NFV is the ability to scale resources dynamically. There must be capabilities to increase or reduce the amount of resources allocated to specific functions or VMs. While current virtualization or cloud platforms allow for this, many of them require a manual trigger by the user or resource owner. Therefore, automation and self-allocation mechanisms that allow the network to dynamically manage resources are critical to the success of NFV.

Distributed Management

Current MANO approaches mainly focus on centralized solutions, which pose scalability limitations, especially in scenarios where services span multiple administrative domains. This is mainly attributed to the communication overhead and the delay incurred by the collection and analysis

		Cloud-Band	Cloud-NFV	ESO	Experia-Sphere	OpenMA-NO	OPN	Open-NFV	OPNFV	Planet Orchestrate	ZOOM
Management approach	Centralized	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Distributed										
	Policy-based	✓	✓	✓	✓	✓	✓	✓		✓	✓
	Self-managed	✓		✓	✓		✓	✓		✓	✓
Management function (FCAPS)	Fault			✓			✓		✓	✓	✓
	Accounting					✓			✓		
	Performance	✓	✓	✓		✓	✓	✓		✓	✓
	Security	✓								✓	✓
Management scope	Functions			✓		✓	✓		✓		✓
	Services	✓	✓	✓	✓	✓	✓			✓	✓
	Network			✓						✓	✓
Managing related areas	SDN	✓				✓	✓		✓		
	Cloud	✓	✓		✓			✓	✓	✓	✓

Table 2. Characteristics of state-of-the-art projects and products.

of data associated with a large number of heterogeneous sources, which prevents these processes from being executed frequently. As a result, the lag in learning the state of services and resources does not allow for online reconfiguration operations. To better react to demand dynamics but also to changing service requirements, efficient monitoring mechanisms need to be implemented, which feed distributed management entities the necessary information to perform dynamic configuration changes. A communication protocol to support lightweight coordination among distributed decision makers, aiming to optimize the usage of resources and the performance of services, is another key research issue.

Management of SDN

While NFV and SDN are not dependent on each other, they are closely related and complementary. Individually, NFV and SDN introduce high levels of dynamism and variability, which curtails the visibility and control of human operators. Therefore, traditional management approaches must be improved to accommodate each of them. While some claim to be based on SDN, all the surveyed projects and solutions focus on managing virtualized compute infrastructure/resources and functions. In the same way, from the SDN perspective, the focus is on managing networks in a programmatic way. We are not aware of management solutions that combine both, which is a key research area. In addition, the management of SDN itself still has open questions [10] such as the number of controllers, their location, and avoiding conflicts in cases where more than one controller manages a given forwarding element. In this case, ideas from policy-based management that have been proposed in most of the surveyed MANO solutions may be extended to SDN.

Security in the Cloud

As SDN and NFV focus on the remote programmability of network resources and their functions, it opens up an important set of new potential threats and attacks that, if successful, could have a far greater impact than in a non-NFV environment. The ETSI NFV security group recently drafted a document that defines the possible security threats that NFV brings to the table.[13] The document is a statement regarding possible security problems but does not yet provide a recommendation to tackle them. Partly because of this, security is currently underdeveloped. Of the surveyed projects and solutions, only CloudBand proposes a comprehensive security solution including anomaly prediction, detection, and isolation, as well as providing security as a service. In the case of ZOOM and Planet Orchestrate, security support is claimed based on a set of best practices and integration with another product, respectively.

Therefore, real security support is lacking in all NFV products, despite the relevant new threats. Important security challenges in this area are detecting and blocking possible intrusion. Specifically, in multi-vendor environments, there are new security concerns of one TSP competitor having access to another TSP's data/configuration. In such a case, isolation between them is important.

Management across the Board

Most solutions provide a way to perform configuration. A limited number add performance management as well, and, as discussed above, only a few provide — albeit limited — security management too. In most solutions, accounting management is completely overlooked. As such, there are currently no ways to track network utilization to ensure that individual parties can be appro-

[13] NFV Security; Problem Statement. Bob Briscoe (Rapporteur). Draft Group Specification published, Oct 2014.

While remarkable progress has been made by ETSI in defining the NFV MANO framework and the constituent (intra-operator) interfaces, there is still much to be done in terms of defining interfaces aimed at supporting interoperability between different vendors with different functions.

priately billed for their use. This is very contradictory given the openness that NFV promises, especially in introducing a more multi-vendor world where different parties can coexist on the same device through virtualization.

More generally, the management of the entire service life cycle is still missing. One of the unique selling points of NFV is that it promises to automate the entire process of setting up and removing a service (chain), including configuration, performance optimization, response to faults, and billing. With support for accounting missing in all products, this promise has not yet been delivered. One of the reasons is that accounting is often heavily intertwined with legacy solutions. Providing support for all FCAPS functionality is therefore highly challenging, but at the same time very important for NFV to really bring a change to the telco world.

Programmability and Intelligence

Given that NFV envisions the deployment and maintenance of complex services across heterogeneous physical resources, a rich set of programmable interfaces should be developed, which will extend the current SDN functionality beyond the scope of controlling simple connectivity resources. Based on the abstraction of the flow, SDN solutions control the distribution of traffic in the network according to forwarding rules. Additional abstractions that apply to computation and storage resources are required so that network functions can be instantiated across multiple vendor technologies, but also interfaces that will allow dynamic (re-)programming of the configuration of those functions and control (e.g., their placement).

A related research challenge concerns the level of intelligence that can be achieved. This is defined by the way in which services are programmed (e.g., declaratively) and the degree to which parameters can be configured. A MANO system should be intelligent enough so that (re-) configuration operations can be automated to a large extent, especially those that react to runtime events. In this respect, intelligent mechanisms that automatically transform high-level policy to operational parameters and perform configuration integrity checks are of paramount importance.

Interfacing and Interoperability

One of the main goals of NFV is to break the bond between equipment vendors and TSPs, and the services they provide. One key requirement for this is support for interoperability. While remarkable progress has been made by ETSI in defining the NFV MANO framework and the constituent (intra-operator) interfaces, there is still much to be done in terms of defining interfaces aimed at supporting interoperability between different vendors with different functions. Interoperability problems can already be observed in all the surveyed projects and solutions. For example, while they all are "based on the ETSI MANO framework," each of the projects and solutions surveyed uses a custom model and/or representation for functions and services. This would mean that unless clear interfaces are defined, it is impossible to chain functions from different operators into a single service. This is because while ETSI proposes VNF and network service descriptors as templates for definition of functions and services, it does not define a data model to realize descriptors.

In this direction, the Alliance for Telecommunications Industry Solutions (ATIS)[14] has been focused on inter-carrier interoperability, new service descriptions, and automated processes. While ATIS has recently proposed seven inter-provider use cases aimed at the same, these are generally generic descriptions. In particular, they do not define any technical requirements or solutions that could be used to enable the use cases.

Conclusion

In this article, we present an overview of the NFV MANO framework recently proposed by ETSI as well as representative projects attempting to realize the framework. Other research projects as well as industry products focused on NFV MANO have been surveyed, identifying their functionality as well as their mapping to the ETSI NFV MANO. Based on these, we identify and discuss open challenges as well as opportunities for research with regard to MANO in NFV.

We have observed that while ETSI completed the first phase of work, the proposed MANO framework still lacks details and standards on the implementation for both the managers as well as the interfaces. As a result, most of the pre-standardization solutions are in fact customized and based on proprietary solutions, which will likely lead to interoperability issues. In addition, while some of the identified challenges such as security are being considered in some of the projects and/or industrial products, others such as resource management — in particular automated resource management — have not received significant attention yet. It is our opinion that the success of NFV will depend, in part, on the availability of mechanisms that are able to autonomously manage network and function resources.

Acknowledgment

The authors are indebted to the Editor-in-Chief and the Series Editors for coordinating the review process, and to the anonymous reviewers for their insightful comments and suggestions. This work has been supported in part by FLAMINGO, a Network of Excellence project (318488) supported by the European Commission under its Seventh Framework Programme, the Science Foundation Ireland Research Centre CONNECT (13/RC/2077), and project TEC2012-38574-C02-02 from Ministerio de Economia y Competitividad.

References

[1] A. Checko *et al.*, "Cloud Ran for Mobile Networks – A Technology Overview," *IEEE Commun. Surveys & Tutorials*, vol. 17, no. 1, 1st qtr. 2015, pp. 405–26.

[2] R. Mijumbi *et al.*, "Network Function Virtualization: State-of -the-Art and Research Challenges," *IEEE Commun. Surveys & Tutorials*, vol. PP, no. 99, 2015, pp. 1–1.

[3] ETSI NFV ISG, "ETSI Network Functions Virtualisation (NFV) Industry Standards (ISG) Group Draft Specifications," http://docbox.etsi.org/ ISG/NFV/ Open, Dec. 2014, accessed May 26, 2015.

[4] ETSI GS NFV-MAN 001, "Network Functions Virtualisation (NFV); Management and Orchestration," http://www.etsi.org/, Dec. 2014, ETSI ISG.

[14] https://www.atis.org/NFV/index.asp

[5] K. Ogaki *et al*., "Integrating Heterogeneous It/Network Management Models Using Linked Data," *Proc. 2013 IFIP/IEEE Int'l. Symp. Integrated Network Management*, May 2013, pp. 768–71.
[6] J. Cardoso *et al*., "Cloud Computing Automation: Integrating USDL and TOSCA," *Proc. Advanced Info. Sys. Eng.*, ser. Lecture Notes in Computer Science, C. Salinesi, M. Norrie, and s. Pastor, Eds.. Springer, 2013, vol. 7908, pp. 1–16.
[7] D. R. Lopez, "OpenMANO: The Dataplane Ready Open Source NFV MANO Stack," *Proc. IETF 92 Meeting Proc.*, Dallas, TX, Mar. 2015.
[8] N. Chowdhury *et al*., "Virtual Network Embedding with Coordinated Node and Link Mapping," *Proc. IEEE INFOCOM 2009*, April 2009, pp. 783–91.
[9] R. Mijumbi *et al*., "Design and Evaluation of Algorithms for Mapping and Scheduling of Virtual Network Functions," *Proc. IEEE Conf. Network Softwarization*, Univ. College London, April 2015.
[10] S. Kuklinski and P. Chemouil, "Network Management Challenges in Software-Defined Networks," *IEICE Trans. Commun.*, Special Section on Management for Flexible ICT Systems and Services, vol. E97-B, no. 99, Jan. 2014, pp. 2–9.

Biographies

Rashid Mijumbi obtained a degree in electrical engineering from Makerere University, Uganda, in 2009, and a Ph.D. in telecommunications engineering from the Universitat Politécnica de Catalunya (UPC), Spain, in 2014. He is currently a postdoctoral researcher in the Telecommunications Software and Systems Group (TSSG) at Waterford Institute of Technology, Ireland. His research interests are in management of networks and services. His current focus is on management of resources for virtualized networks and functions, cloud computing, and software defined networks.

Joan Serrat is a full professor at UPC. He received a degree in telecommunication engineering in 1977 and a Ph.D. in the same field in 1983, both from UPC. He has been involved in several collaborative projects with different European research groups, through both bilateral agreements and participation in European funded projects. His topics of interest are in the field of autonomic networking and service and network management.

Juan-Luis Gorricho received a telecommunication engineering degree in 1993 and a Ph.D. degree in 1998, both from UPC. He is currently an associate professor at UPC. His recent research interests are in applying artificial intelligence to ubiquitous computing and network management, with special interest in using smartphones to achieve recognition of user activities and locations; and applying linear programming and reinforcement learning to resource management in virtualized networks and functions.

Steven Latré received an MSc. degree and a Ph.D., both in computer science, from Ghent University, Belgium. He is currently an assistant professor at the University of Antwerp, Belgium. His research activity focuses on autonomous management and control of both networking and computing applications. He has been involved in several national and European research projects. His recent work has focused on quality of experience optimization and management, distributed control, and network virtualization.

Marinos Charalambides is a senior researcher at University College London. He received a B.Eng. in electronic and electrical engineering, an M.Sc. in communications networks and software, and a Ph.D. in policy-based management, all from the University of Surrey, United Kingdom, in 2001, 2002, and 2009, respectively. He has been working in a number of European and U.K. national projects since 2005, and his research interests include software-defined networking, policy-based management, in-network caching, and online traffic engineering.

Diego Lopez received an M.S. from the University of Granada in 1985 and a Ph.D. degree from the University of Seville in 2001. He joined Telefonica I+D in 2011 as a senior technology expert and is currently in charge of the Technology Exploration activities within the GCTO Unit. He is focused on network virtualization, infrastructural services, network management, new network architectures, and network security. He actively participates in the ETSI ISG on NFV (chairing its Technical Steering Committee), the ONF, and the IETF WGs connected to these activities.

While ETSI completed the first phase of work, the proposed MANO framework still lacks details and standards on the implementation for both the managers as well as the interfaces. As a result, most of the pre-standardization solutions are in fact customized and based on proprietary solutions.

Transparent Reallocation of Control Functions in IMS Deployments

Jaime Garcia-Reinoso, Ivan Vidal, Paolo Bellavista, Ignacio Soto, and Pedro A. Aranda Gutirrez

The authors present a novel solution to transfer users between IMS network elements. This solution enables a home operator to perform an appropriate redistribution of load among the call session control functions of the IMS, which can be deployed as virtualized network functions or over dedicated machines.

Abstract

In this article we present a novel solution to transfer users between IMS network elements. This solution enables a home operator to perform an appropriate redistribution of load among the call session control functions of the IMS, which can be deployed as virtualized network functions or over dedicated machines. This way, the operator is enabled to adequately accommodate the instantaneous load generated by users to the available control resources, which may dynamically be activated or deactivated as necessary, as well as to enhance the resilient operation of the IMS deployment. Additionally, our solution does not require any change to the IMS specifications, and at the same time, the procedures are transparent to the end-user applications running at the IMS terminals. Finally, we describe some results obtained with a proof-of-concept implementation of the procedures presented in the article in order to show the viability and correctness of our proposal.

Introduction

Nowadays, telco operators own a huge and complex infrastructure to provide their services to their end users. To maximize the usage of such infrastructure, and to accommodate the instantaneous requirements of their customers, these network operators usually deploy load balancers to select the proper network device to serve incoming requests. On the other hand, the traffic load generated by end users is variable during the day, where the high differences between the peak and valley zones impose different network requirements. In such scenarios, a telco operator would like to add or remove active machines in order to reduce costs. This reduction of costs is even clearer when operators use virtual network functions of third-party providers. By using virtualization, telco operators could dynamically request the instantiation of virtual machines in order to adapt their network resources to the traffic load. However, even using load balancers to allocate the incoming end users' traffic to (possibly virtual) network element functions, the change in number of available active elements generates different loads on such devices. In that case, it would be desirable to transfer load between network elements to maximize their utilization while minimizing the number of active network devices.

The IP multimedia subsystem (IMS) framework defined by the Third Generation Partnership Project (3GPP) is a next generation network architecture to provide multimedia services over IP networks. This architecture is designed to be scalable throughout the redundant instantiation of these entities and the usage of load balancing mechanisms, such as those supported with the DNS. In this respect, the considerations described before about the transfer of the load between functional elements would be beneficial for the IMS core elements too. Additionally, with such mechanisms, it would be straightforward to add resilience to both virtual and physical IMS deployments, as they would enable transfer of the load of failing core elements to other existing elements without disrupting the service provided to the end users. IMS is playing a relevant role in 3G/4G telco support infrastructures, even more with the introduction of Voice over Long Term Evolution (VoLTE), with many general concepts that are going to influence and persist in future 5G networking. In particular, the elastic provisioning of virtualized IMS-based functions is still considered a challenging issue.

In this article, we propose a new architecture to transparently transfer users among the functional elements that implement the IMS call session control functions. With our approach, a network operator can better adapt its active control resources to the instantaneous load generated by users, also considering roaming users from other operators, with evident advantages in terms of efficiency and economic costs. It would also be possible to use this solution to add resilience to IMS after failures, transferring all the state of the failed control element to a new machine.

IMS and Related Work on Its Efficient Deployment

IMS is a standardization effort supported by the majority of relevant telco players' consortia and initiatives, such as 3GPP, 3GPP2, Internet Engineering Task Force (IETF), and Open Mobile Alliance (OMA), in order to create a common signaling framework for the provisioning of value-added services in operator networks. The IMS architecture incorporates a set of key functions

Jaime Garcia-Reinoso, Ivan Vidal, and Ignacio Soto are with Universidad Carlos III de Madrid; Paolo Bellavista is with the University of Bologna; Pedro A. Aranda Gutiérrez is with Telefónica I+D.

 0163-6804/16/$25.00 © 2016 IEEE

related to quality of service (QoS) provisioning, charging, integration, security, and roaming. The definition of the IMS network entities and interfaces can be found in [1]. A simplified overview of the IMS architecture is depicted in Fig. 1, the main functional entities of which are the proxy-/interrogating-/serving-call session control functions (P-/I-/S-CSCF), the application servers (ASs), the home subscriber server (HSS), and the IMS terminal or user equipment (UE).

The UE controls session setup and media transport via a specific Session Initiation Protocol (SIP) profile defined in IETF and 3GPP IMS-related standards [2]. Each UE is associated with an IMS user identified by a globally unique identifier in the form of a SIP URI [3]. CSCFs are the core entities for providing session control in the IMS. In particular, the P-CSCF is the entry point to the IMS infrastructure for the SIP messages of a UE. Each UE maintains a security association with its P-CSCF to exchange SIP messages with the IMS core in a secure way. If the UE is in a visited network (i.e. the UE is roaming), it can use a P-CSCF placed in its operator network (i.e. the home network) or in the visited network. The I-CSCF is the entry point in a network for incoming SIP session setup messages and for UE registration messages coming through a P-CSCF. It interacts with the HSS to locate the particular S-CSCF assigned to a UE (allocation is based on user profiles) and routes incoming messages to it by acting as a stateless SIP proxy. The S-CSCF performs functionalities related to user registration and session control. In particular, it registers users, interacting with the HSS to obtain authentication and user profile information, and routes SIP messages to different ASs depending on message type and filters/triggers specified in user profiles; these are IMS initial filter criteria. ASs implement the logic to provide services to end users, simplifying the introduction of new IMS-based services. Finally, IMS uses the standard procedures defined in [1] (e.g., based on Domain Name Service [DNS] or Dynamic Host Configuration Protocol [DHCP]) to obtain the IP addresses of SIP servers, such as CSCFs and ASs.

In the literature we can find some proposals to improve the efficiency of IMS deployments. In [4] the authors propose a mechanism by which new registration requests are distributed among available S-CSCF nodes considering their load. At the cost of more complexity, the proposal in [5] provides better load distribution because, in addition to new registration requests, registered users without open sessions can be moved to a different S-CSCF according to load balancing requirements. In [6, 7] the authors propose solutions to distribute the load among ASs, but in both cases this is only done for new sessions. In [8, 9] the authors propose solutions that consider load balancing and also node failure recovery for P-CSCF and I-CSCF nodes in IMS, allowing reallocation of control nodes for users with open sessions. Nevertheless, both proposals require modifications to the standardized behavior and interfaces of IMS nodes (i.e., a change in IMS specifications), which are not needed in our proposal. A paper focused on scalability problems of the IMS core is [10], which describes interesting solutions for effective load balancing, adaptive replication, and elastic quality provisioning via dynamically replicated IMS core components, but considering only new sessions and not the reallocation of active sessions our proposal enables.

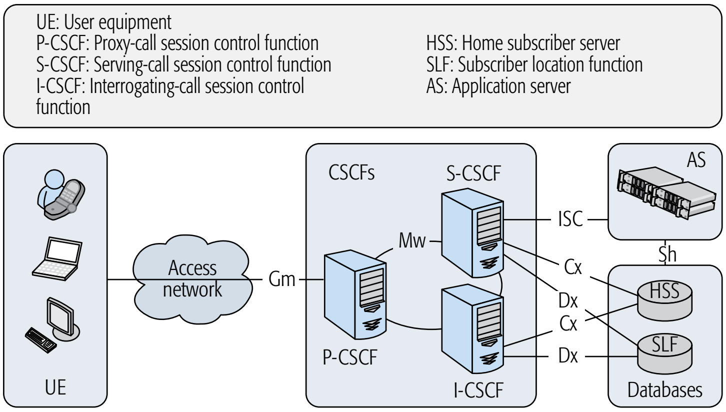

Figure 1. Simplified overview of the IMS architecture.

Finally, another interesting research direction (e.g., [11]) is the virtualization of IMS components to create flexible environments in which to add and remove those components to adapt to the load. The components can be deployed on public/private/hybrid cloud infrastructures. A requirement to maximize the benefit from this environment is effective management of the dynamic transfer of users, even with active sessions, among IMS functional elements, which is the key contribution of our article.

Transparent Reallocation of IMS Control Functions

This section describes the functional entities that are proposed in this article to enable an appropriate allocation of users to the available CSCFs in IMS deployments. These entities, interoperating with the IMS, allow a home operator to dynamically change the allocation of the P-CSCF or S-CSCF of any of its users, transparent to the end-user applications running at the UE. This way, our architecture enables the home operator to maintain an appropriate distribution of the load among the existing CSCFs. Figure 2 shows the new functional elements of our solution (highlighted in gray), as well as their relationship with the elements of the IMS architecture. It is important to remark that our solution does not add new interfaces or functionality to standard IMS elements, so it does not require modifications to the IMS specifications.

The *control function discovery* (CFD) is an application server acting as a SIP user agent. The SIP URI of this AS is configured in the user profile, in initial filter criteria that will be matched during the registration. This way, after successful registration or re-registration of the user in the IMS, and following the regular IMS procedures, the CFD AS will receive the information corresponding to this particular registration, which will include, among other things, the value of the registration expiration interval and the addresses

It is important to remark that our solution does not add new interfaces or functionality to standard IMS elements, so it does not require modifications to the IMS specifications.

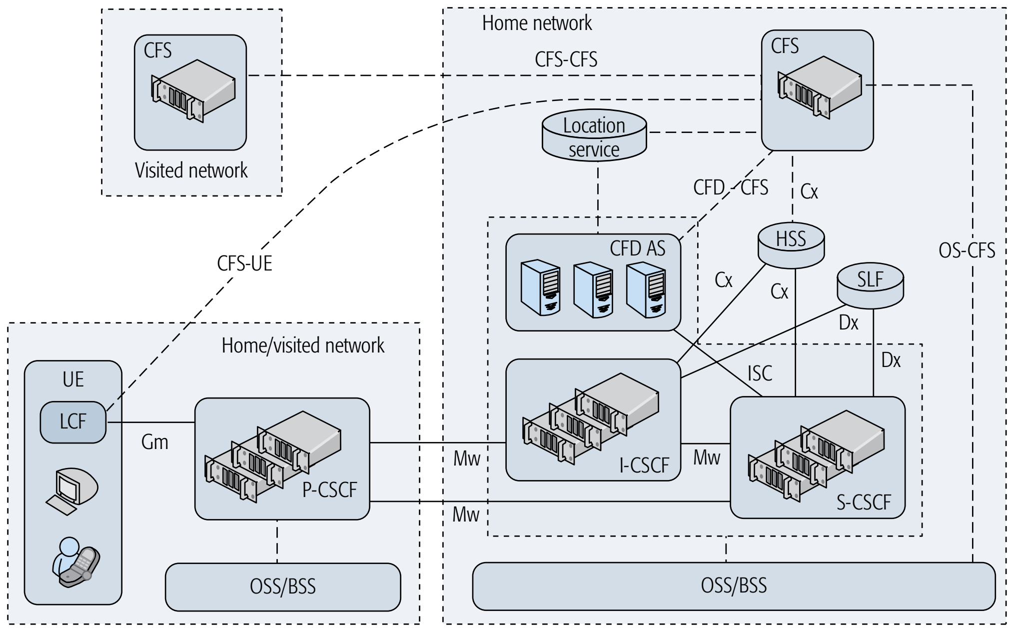

Figure 2. Overview of the proposed architecture.

of the P-CSCF and S-CSCF allocated to the user. The CFD AS stores this information in a *location service*, which is a new database in the home network domain that maintains all the information related to the IMS control functions serving the users of the home operator. For the sake of scalability, there may be several CFD AS entities in the home network to appropriately process the load of notifications about the registration status of users. However, we want to highlight that our solution only requires a reduced set of these entities with respect to the number of CSCFs, as CFD ASs are only contacted after successful registration and re-registration of a user, and are not involved in any other IMS signaling procedures (e.g., session setup, event subscription, and notification.).

The *control function selection* (CFS) is the key component of our solution that enables the change of the P-CSCFs and S-CSCFs allocated to the users of the home operator. This functional entity is the contact point of our solution with the operations support systems/business support systems (OSS/BSS) of the home network domain. The OSS/BSS can activate or deactivate control functions (P-CSCF, S-CSCF and I-CSCF entities) as needed to scale the IMS control resources to the instantaneous load generated by the users, and triggers the CFS when a redistribution of the load among the existing CSCFs is required. This may happen, for instance, to enforce a transfer of the load of an underutilized CSCF to other existing CSCFs prior to its deactivation, or to achieve the resilient operation of an IMS deployment under a failing or overloaded control node. To support the appropriate operation, the CFS implements a Cx interface, as defined for the IMS in [12], which enables the communication with the HSS. This interface will be used in S-CSCF reallocation procedures, as explained later.

The CFS has access to the information stored by the CFD AS in the location service of the home network domain. Thus, upon receiving an indication from the OSS/BSS, via the reference point os-cfs, to perform a load transfer from a P-CSCF (or S-CSCF) to other existing CSCF entities, it may retrieve the information about the users served by the former CSCF from this location service. With this information, the CFS can initiate a reallocation procedure to satisfy the request received from the OSS/BSS, transferring a subset of the users from the initial CSCF to the target CSCFs.

Our architecture includes a local control function (LCF) in the UE. This is integrated as an extension to the IMS stack of the terminal,[1] and maintains status information about the SIP dialogs established by the user. The LCF supports a new reference point cfs-ue, which can be used by the CFS to trigger the transfer of the user to a new CSCF. The LCF is in charge of executing the signaling procedures that are necessary to perform this transfer. Moreover, it carries out these procedures transparently to any end-user application running at the UE, which are kept unaware of the CSCF change.

In the following, we illustrate the procedures used in our solution to change the allocation of the P-CSCF or the S-CSCF for a given user of the home operator, and also to enable appropriate load balancing among I-CSCF and CFD AS entities.

Change of the P-CSCF Allocation

This section describes the procedures that are performed in our solution to change the P-CSCF allocated to a set of users. Prior to the execution of these procedures, we assume that the users have already registered to the IMS, and may have established a set of multimedia sessions by means of SIP dialogs. The procedures are outlined in Fig. 3.

In the scenario shown in this figure, the OSS/

[1] The required functionality can be implemented as software in the UE, and it can be delivered integrated with the required software to access the IMS services.

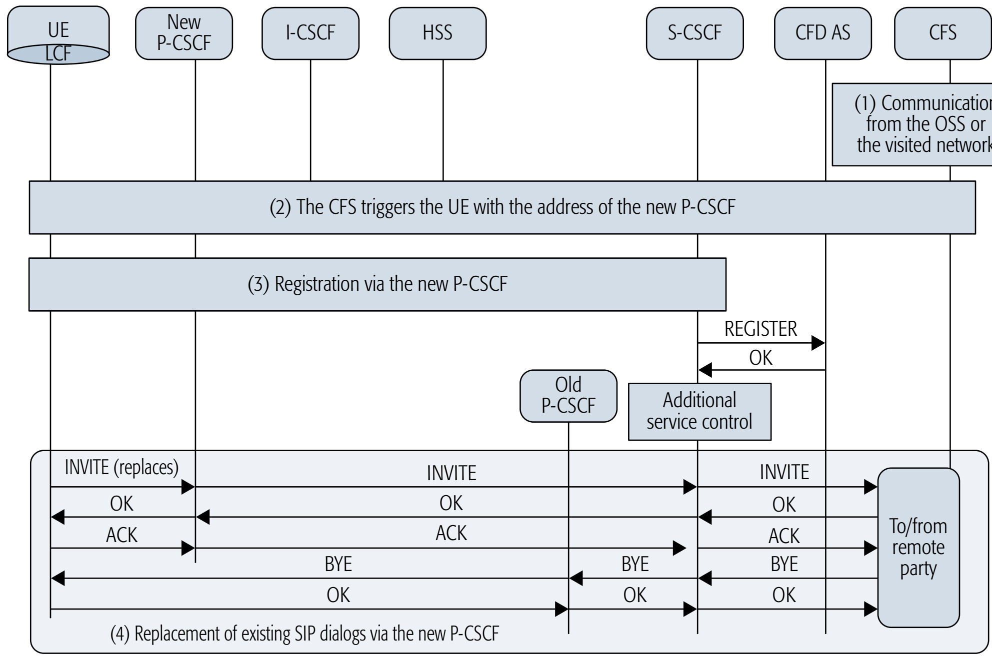

Figure 3. IMS procedures to change a P-CSCF allocation.

The operator may instruct the LCF to delay the replacement of each SIP dialog established by the terminal to a given deadline. This would allow distributing over time the signaling load corresponding to dialog replacement procedures, and even reducing this load by allowing the dialog to be terminated by the end user, via the old P-CSCF, before the deadline.

BSS in the home network contacts the CFS to perform a load transfer from a P-CSCF to other existing P-CSCFs (step 1). Next, the CFS retrieves from the location service the information about the users served by the affected P-CSCF, and determines the set of candidate users that will be transferred to each target P-CSCF. After defining the set of candidate users for each target P-CSCF, the CFS starts a reallocation procedure to transfer each identified user to its corresponding target P-CSCF. For this purpose, it communicates with the UE of each user to be transferred (step 2), using the reference point *cfs-ue*.[2] As a result of this communication, the LCF at each UE can independently initiate the IMS signaling procedures that are necessary to enforce the allocation of the new P-CSCF to the user. In particular, the LCF executes a new IMS-level registration using the address of the new P-CSCF (step 3). After successful registration of the user, and following the regular IMS procedures, the S-CSCF allocated to the user performs service control functionalities by sending a SIP REGISTER request to any AS specified in the user profile for the registration event. This way, a REGISTER request is received by a CFD AS, which updates the information contained in the location service to reflect the address of the new P-CSCF allocated to the user. Finally, to complete the transfer to the new P-CSCF, it is necessary to replace all the SIP dialogs established by the user via the old P-CSCF with new SIP dialogs through the new P-CSCF. This process is done by the LCF, according to the procedures defined in [13], using the status information corresponding to the SIP dialogs of the user. To clarify this concept, step 3 in Fig. 3 shows an example, where the LCF replaces a single SIP dialog established by the user.

Finally, our solution does not impose or recommend any specific mechanism to select which specific users will be reallocated to each target P-CSCF. Instead, the final decision is uniquely subject to the policy rules established by the home operator. Moreover, the proposed solution does not define the precise instant of time when each candidate user is to be contacted to carry out the reallocation. On the contrary, the time schedule of the reallocation procedures is dictated by the operator, who might contact all the candidate users simultaneously or spread the requests sent to UEs over time. Additionally, the operator might provide the LCF (through the *cfs-ue* interface) with diverse policies to govern user transfer. As an example, the operator may instruct the LCF to delay the replacement of each SIP dialog established by the terminal to a given deadline. This would allow distributing over time the signaling load corresponding to dialog replacement procedures, and even reducing this load by allowing the dialog to be terminated by the end user via the old P-CSCF before the deadline.

Change of the S-CSCF Allocation

The procedures to change the S-CSCFs allocated to users are similar to those described in the previous section, and are illustrated in Fig. 4. After receiving a trigger from the OSS/BSS to perform a redistribution of the load of a given S-CSCF to other existing S-CSCFs (step 1), the CFS retrieves the set of users served by the specified S-CSCF. Following the policy rules defined by the home operator, the CFS determines the set of candidate users that will be transferred to each target S-CSCF. Then, for each candidate user, the CFS contacts the HSS through the reference point Cx, and changes the S-CSCF assigned to the user (step 2). This can be done using a Multimedia-Authentication-Request (MAR) Diameter command, according to [12].

Next, the CFS communicates with the UE of

[2] The communication between the CFS and the LCF should be done out of the SIP signaling path, as this communication can be the result of a failed or overloaded P-CSCF, and could be implemented, for example, using HTTP.

Analogous to the P-CSCF reallocation procedure, note that the actual decisions on which candidate users are to be transferred from a given S-CSCF to another, and the precise time schedule of the reallocation process, are controlled by the policy rules established by the operator.

each of the users, using the reference point cfs-ue, instructing the LCF to start the IMS procedures needed to enforce the allocation of the new S-CSCF to the user (step 3). Analogous to the previous case, the LCF first initiates an IMS-level registration using the new S-CSCF. Step 4 in Fig. 4 illustrates an example of the registration process, where the S-CSCF performs user authentication procedures. In the example, the registration proceeds in two phases. In the first phase, the UE generates a SIP REGISTER request, which is routed via the P-CSCF allocated to the user and an I-CSCF belonging to the home network. Then the I-CSCF contacts the HSS to discover if there is an S-CSCF allocated to the user. As the user has already been allocated the new S-CSCF in step 2, the HSS returns the address of the new S-CSCF to the I-CSCF. Consequently, the new S-CSCF receives the REGISTER request and contacts the HSS to download the data that is necessary to authenticate the user. Then the S-CSCF answers back the REGISTER request with a SIP Unauthorized response, containing a challenge to be satisfied by the UE. The response to this challenge requires a second REGISTER transaction between the UE and the S-CSCF, as shown in Fig. 4.

After successful registration of the user, the new S-CSCF performs the regular IMS service control functionalities, which result in the update

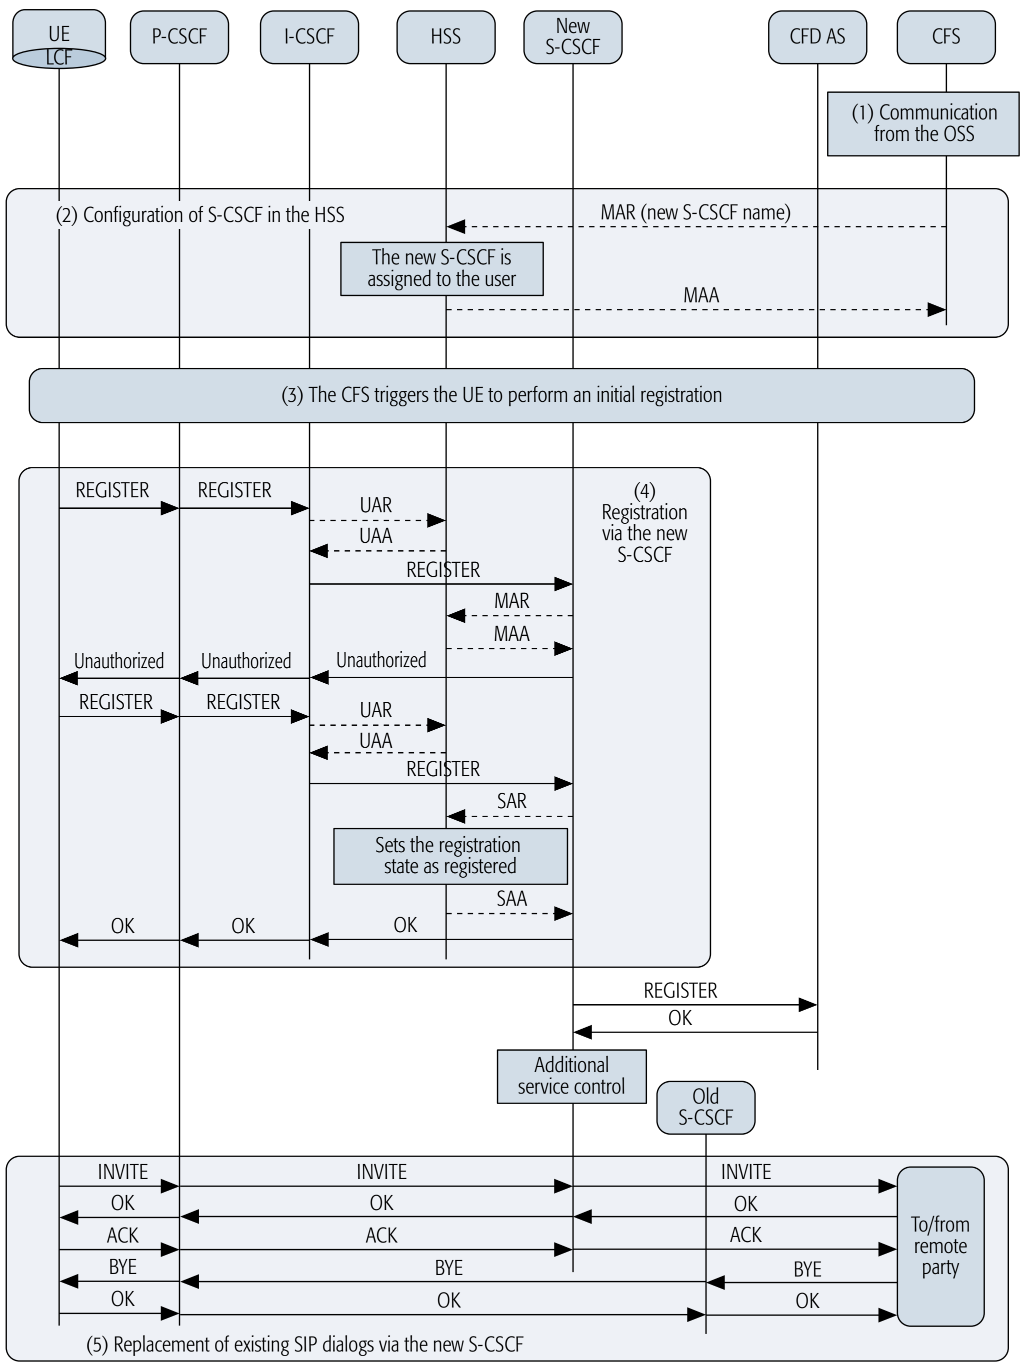

Figure 4. IMS procedures to change an S-CSCF allocation.

of the information stored in the location service by a CFD application server, to reflect the address of the new S-CSCF assigned to the user. Finally, the LCF initiates the signaling procedures to replace the SIP dialogs established by the user to use the new S-CSCF (step 5).

Analogous to the P-CSCF reallocation procedure, note that the actual decisions on which candidate users are to be transferred from a given S-CSCF to another, and the precise time schedule of the reallocation process, are controlled by the policy rules established by the operator.

Management of CFD AS and I-CSCF Entities

The application servers of our solution, that is, the CFD AS entities, can also be activated or deactivated as necessary by the OSS/BSS for resilience and scalability purposes. In our proposal, the CFD service is identified by a SIP URI, containing a domain name that, according to the procedures specified in [14], can be resolved by the S-CSCF in the DNS to the different addresses of the application servers executing the CFD service. The OSS/BSS will trigger the CFS when a CFD AS is activated or deactivated, so that it can change the configuration of the DNS in the home network accordingly, this way enabling an appropriate load balancing among the existing CFD AS entities.

On the other hand, besides the IMS control functions that are allocated to the end users (i.e., P-CSCFs and S-CSCFs), an IMS deployment also includes a number of I-CSCFs for scalability reasons. The selection of an I-CSCF during the IMS registration and session setup follows the SIP procedures specified in [14], and is based on DNS usage. In a dynamic environment, the OSS/BSS could also activate or deactivate an I-CSCF entity as necessary. In this case, the OSS/BSS would trigger the CFS, which in turn would change the configuration of the DNS in the home network to reflect the addition or deletion of the I-CSCF address. This will enable an appropriate (for low-medium frequency updates such as the ones of interest for many application domains [10]) DNS load balancing among the available I-CSCFs.

Considerations about Roaming Users

According to the IMS specifications [1], there is the possibility that the P-CSCF allocated to a roaming user is located in the visited network. In that case, there may be users served by a P-CSCF that are not registered in the network of the operator owning the P-CSCF. Our solution covers this specific use case, by allowing communications between CFS entities in the home and visited networks through the inter-domain reference point *cfs-cfs*.

After the successful registration of a user who is roaming in a visited network, a CFD AS in the home network of the user stores the address of the visited P-CSCF in its location service. Additionally, when this CFD AS detects a new address of a visited P-CSCF, it contacts the CFS in the home network (hereafter referred to as home CFS) via the reference point *cfd-cfs*. If the visited network implements the solution presented in this article, the home CFS communicates with the CFS in the visited network (hereafter referred to as visited CFS), through the reference point *cfs-cfs*, and registers its desire to receive notifications about the state of the P-CSCF allocated to the user.

Now, if a P-CSCF is overloaded, fails, or needs to be removed, the operator owning the P-CSCF may contact any operator that has expressed its interest in receiving notifications on the P-CSCF status. To this end, the visited CFS (i.e., the CFS in the domain of the affected P-CSCF) communicates with the home CFS via the *cfs-cfs* interface. The visited CFS may explicitly indicate the reason for the communication, and request of the home CFS information about the visiting users that are served by the P-CSCF (e.g., the number of these users). Taking into account the information from the different CFS entities that have been contacted, and according to the policy rules established by the visited operator, the visited CFS can then request any home CFS entity to perform a reallocation process of its visiting users. The trigger from the visited CFS may include a set of alternative P-CSCFs and the assignment of load that should be transferred to each of these P-CSCFs. With this information, and attending to the policy rules defined by the home operator, the home CFS then executes the reallocation process following the procedures previously described.

Finally, we want to note that there is always the possibility of having roaming users assigned to a P-CSCF belonging to a network domain that does not implement the solution described in this article. In this case, the visited CFS would not have access to the information about the visiting users from their home network domain, and would not be able to request the execution of a reallocation process for these users. However, this issue can easily be prevented if the service level agreement established with other operators allows use of the set of P-CSCFs of the operator only by those operators that implement the control function discovery and selection procedures described in this article (users can still be provided service with a P-CSCF in their home networks).

Experimental Results

To assess our proposed solution, we have deployed a testbed using the FOKUS OpenIMS core[3] and the SIPp open software.[4] The former includes the call session control functions of the IMS (i.e., P-CSCF, S-CSCF, and I-CSCF) and an HSS, while the latter is used to emulate an IMS UE. The OpenIMS core is deployed over three virtual machines: the first, including the HSS together with three CSCFs (P-CSCF1, S-CSCF1, and I-CSCF1); the second, executing P-CSCF2 and S-CSCF2; and the third, including P-CSCF3 and S-CSCF3. Two additional virtual machines are used to run the SIPp scripts corresponding to the UEs.

To illustrate the benefits that can be achieved reallocating users to CSCF entities, we designed a first experiment with 20 users. All these users are served by P/S-CSCF1 and are configured to generate a certain aggregate session setup rate. In the experiment, we measured the average session setup delay for different session setup rates. The results are shown in Table 1. As can

According to the IMS specifications [1], there is the possibility that the P-CSCF allocated to a roaming user is located in the visited network. In that case, there may be users served by a P-CSCF that are not registered in the network of the operator owning the P-CSCF. Our solution covers this specific use case.

[3] http://www.fokus.fraunhofer.de/en/fokus_testbeds/open_ims_playground/components/osims/index.html

[4] http://sipp.sourceforge.net

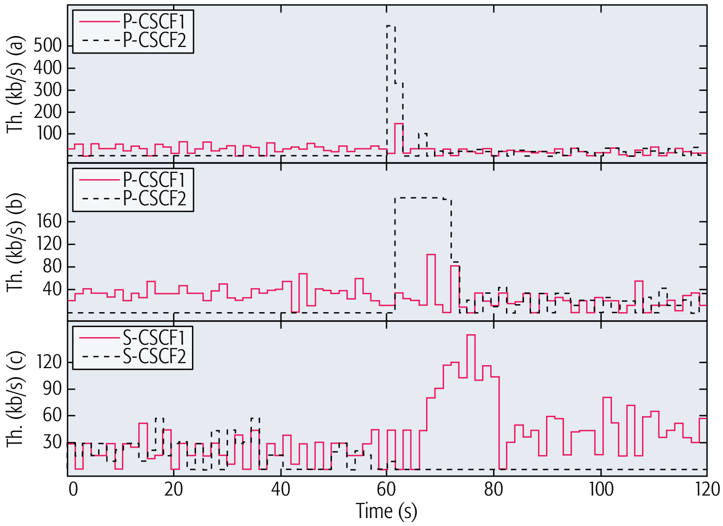

Figure 5. Results of the experiments.

Aggregate session setup rate (sessions/s)	Average delay of the session setup (s)
0.5	0.068
1	0.073
2.5	0.7
5	3.1
10	10.97

Table 1. Average delay of the session setup for different values of aggregate session setup rates.

be observed, in our particular testbed, when the session initiation rate is over 5 requests/s, the session setup delay is above the grade-of-service parameter for call setup delays recommended by the International Telecommunication Union — Telecommunication Standardization Sector (ITU-T) of 3 s [15]. This shows the benefits of distributing the rate of sessions established by users among CSCFs, as supported by our reallocation procedures.

In the second experiment, 20 registered users allocated to P-CSCF1 establish IMS sessions with 2 registered users using P-CSCF3. The aggregate session setup rate is 1 session/s. Using the SIPp scripts, we execute the procedures to change the allocation of 10 users from P-CSCF1 to P-CSCF2, as illustrated in Fig. 3. This change is triggered manually after approximately 60 s. In this particular setup, all the users are transferred in parallel, which imposes a high transitory load for both P-CSCFs. Figure 5a presents the result of this experiment in terms of total throughput at the two P-CSCFs, where we can observe that once the users are transferred to the new P-CSCF, the load is distributed between them. Note that the throughput shown corresponds to SIP signaling and is also a measurement of the load at the CSCFs to process those SIP messages. Although the specific policy to reallocate users between functional elements is under the control of the home operator, for the sake of completeness we also consider in this validation a simple algorithm to schedule the transfer of users between P-CSCFs. In such an algorithm, the users are contacted in sequence with a configurable transfer rate. In the third experiment, where we configure a transfer rate of 2 users/s, we are able to significantly reduce the peak throughput during the transfer stage, as shown in Fig. 5b.

In the fourth experiment, we show a scenario where the load of two S-CSCFs are below the minimum configured threshold, so all the users assigned to one of them can be transferred to the other. In this case, there are 10 users registered in S-CSCF1, while another 10 are registered in S-CSCF2. The 20 users generate an aggregate rate of 1 session/s, as in the previous experiment. After approximately 60 s, the users registered in S-CSCF2 are transferred to S-CSCF1 at a transfer rate of 2 users/s, using the mechanisms illustrated in Fig. 4. As can be seen in Fig. 5c, the users and their active sessions are transferred to the new S-CSCF between 60 s and 80 s, causing a higher load in S-CSCF1 during that period. After 80 s, the users have been successfully transferred to S-CSCF1, which receives a rate of 1 session setup per second.

Conclusions

In this article we have provided a solution to transparently transfer users between IMS control elements, which can be used to appropriately adapt the utilization of these elements to the instantaneous load generated by users as required by the operator, and add resilience to the IMS infrastructure at the same time. The results achieved with our proof-of-concept implementation show the viability and correctness of the proposed solution. The solution guidelines described in our article also have general validity and applicability to next generation telco support infrastructures, which, for example, will extensively exploit the opportunities of virtualization and dynamic re-allocation of their control functions. We plan to extend our proposal by including the possibility to manage other IMS functional elements like ASs.

Acknowledgments

The work in this article has been partially supported by the European FP7 Trilogy 2 project (grant agreement CNECT-ICT-317756).

References

[1] 3GPP, "IP Multimedia Subsystem (IMS); Stage 2," TS 23.228 v. 12.6.0 Release 12, Sept. 2014.

[2] 3GPP, "IP Multimedia Call Control Protocol Based on Session Initiation Protocol (SIP) and Session Description Protocol (SDP); Stage 3," TS 24.229 v. 12.6.0 Release 12, Oct. 2014.

[3] J. Rosenberg *et al.*, SIP: Session Initiation Protocol, RFC 3261, July 2002.

[4] I. Abdalla and S. Venkatesan, "Notification Based S-CSCF Load Balancing in IMS Networks," *Proc. Wireless Telecommun. Symp.*, 2011, pp. 1–5.

[5] L. Xu *et al.*, "De-Registration Based S-CSCF Load Balancing in IMS Core Network," *IEEE ICC '09*, 2009, pp. 1–5.

[6] Y-J. Yoo *et al.*, "Service Triggering Algorithm for Optimal Load Balancing In Heterogeneous Multiple AS (Application Server) System," *Proc. Int'l. Conf. Info. and Commun. Tech.. Convergence*, 2010, pp. 10–15.

[7] J. Zha *et al.*, "Research on Load Balance of Service Capability Interaction Management, *Proc. 3rd IEEE Int'l. Conf. Broadband Network and Multimedia Tech.*, 2010, pp. 212–17.

[8] C. Makaya *et al.*, "Service Continuity Support in Self-Organizing IMS Networks," *Proc. 2nd Int'l. Conf. Wireless Commun., Vehic. Tech., Info. Theory and Aerospace & Electronic Sys. Tech.*, 2011, pp. 1–5.

[9] A. Dutta *et al*., "Self Organizing IP Multimedia Subsystem," *IEEE Int'l. Conf. Internet Multimedia Services Architecture and Applications*, 2009, pp. 1–6.
[10] A. Corradi, P. Bellavista, and L. Foschini, "Enhancing Intra-Domain Scalability of IMS-Based Services," *IEEE Trans. Parallel Distrib. Sys*., vol. 24, no. 12, Dec. 2013, pp. 2386–95.
[11] P. Bellavista *et al*., "QoS-Aware Elastic Cloud Brokering for IMS Infrastructures," *Int'l. Symp. Computers and Commun.)*, 2012, pp. 157–62.
[12] 3GPP, "IP Multimedia (IM) Subsystem Cx and Dx Interfaces; Signalling Flows and Message Contents," TS 29.228 v. 12.3.0 Release 12. Oct. 2014.
[13] R. Mahy, B. Biggs, and R. Dean, "The Session Initiation Protocol (SIP) 'Replaces' Header," IETF RFC 3891, Sept. 2004.
[14] J. Rosenberg and H. Schulzrinne, "Session Initiation Protocol (SIP): Locating SIP Servers," IETF RFC 3263, June 2002.
[15] ITU-T Rec. E.721, "Network Grade of Service Parameters and Target Values for Circuit-Switched Services in the Evolving ISDN," May 1999.

Biographies

Jaime Garcia-Reinoso (jgr@it.uc3m.es) received a telecommunications engineering degree in 2000 from the University of Vigo, Spain, and a Ph.D. in telecommunications in 2003 from the University Carlos III of Madrid, Spain. He is currently an associate professor at University Carlos III of Madrid, having joined in 2002, and has published over 50 papers in the field of broadband computer networks in top magazines and congresses. He has been involved in several international and national projects related to broadband access, peer-to-peer overlays, next generation networks, signaling protocols, and information-centric networking like EU IST MUSE, Trilogy 2, EU H2020 XHaul, and the Spanish MASSES project.

Ivan Vidal (ividal@it.uc3m.es) received a telecommunication engineering degree in 2001 from the University of Vigo, and a Ph.D. in telematics engineering in 2008 from the University Carlos III of Madrid. He is currently working as a visiting professor at Universidad Carlos III de Madrid. His research interests include multimedia networking, unmanned aircraft vehicles, and information-centric networking. He has been involved in several international and national research projects related to these topics, including EU IST MUSE, the Spanish MEDIANET project, and the Spanish Ministry of Defense project DRONE. He has published several papers in magazines and conferences, lately in the areas of mobility support in IMS networks, UAVs, and ICN management.

Paolo Bellavista (paolo.bellavista@unibo.it) is an associate professor at the University of Bologna, Italy. His research activities span from mobile computing to pervasive ubiquitous middleware, from vehicular sensor networks to big data adaptive stream processing and adaptive scalable multimedia, from IoT middleware for cloud integration to resource management in SDN. He has published more than 60 magazine/journal articles and more than 80 conference/workshop papers in those fields, reporting results from several national- and EU-funded projects. He serves on the Editorial Boards of *IEEE TNSM*, *IEEE TSC*, *Elsevier PMC*, and *Springer JNSM*. See also http://lia.deis.unibo.it/Staff/PaoloBellavista/ for additional details and a complete publications list.

Ignacio Soto (isoto@it.uc3m.es) received a telecommunication engineering degree in 1993, and a Ph.D. in telecommunications in 2000, both from the University of Vigo, Spain. He was a research and teaching assistant in telematic engineering at the University of Valladolid from 1993 to 1999. He was an associate professor at the Universidad Politcnica de Madrid from 2010 to 2012. In 1999 he joined University Carlos III of Madrid, where he currently works as an associate professor in telematic engineering. His research activities focus on mobility support in packet networks, heterogeneous wireless access networks, and vehicular networks.

Pedro A. Aranda Gutiérrez (pedroa.aranda@telefonica.com) obtained a Master's degree in telecommunications at the Universidad Politécnica de Madrid in 1989 and his Ph.D. in informatics at the Universität Paderborn in 2013. He currently works at the Technology Exploration unit at the Network Virtualisation and Innovation directorate at Telefónica, I+D/GCTO Unit. He has worked for Telefónica, I+D for the last 24 years and has participated in joint research projects of the EU starting at the RACE I Framework Programme until now. He is currently acting as technical manager of the NetIDE project of the 7th Framework Programme. His technical experience and interests range from IP networking with a foundation in IP routing and inter-domain routing architectures to the current SDN frameworks from a networking and programmer's point of view.

Self-Healing in Mobile Networks with Big Data

Emil J. Khatib, Raquel Barco, Pablo Muñoz, Isabel de la Bandera, and Immaculada Serrano

To maintain a certain quality of service, self-healing systems must complete their tasks in a reasonable time. The conjunction of a big volume of data and the limitation of time requires a big data approach to the problem of self-healing. This article reviews the data that self-healing uses as input and justifies its classification as big data. Big data techniques applied to mobile networks are examined, and some use cases along with their big data solutions are surveyed.

Abstract

Mobile networks have rapidly evolved in recent years due to the increase in multimedia traffic and offered services. This has led to a growth in the volume of control data and measurements that are used by self-healing systems. To maintain a certain quality of service, self-healing systems must complete their tasks in a reasonable time. The conjunction of a big volume of data and the limitation of time requires a big data approach to the problem of self-healing. This article reviews the data that self-healing uses as input and justifies its classification as big data. Big data techniques applied to mobile networks are examined, and some use cases along with their big data solutions are surveyed.

Introduction

With each new generation of communication technologies, the control data and network performance measurements in the operation and maintenance (O&M) subsystem grow in both volume and speed. A larger number of users and quantity of transferred data means a larger amount of measurements. In addition, the higher demand for quality and larger bandwidths pushes for a faster rate of generation and consumption of these data. Downtimes due to problems that are not solved quickly or incorrect optimization cause a high opportunity cost and increased overall O&M costs.

For these reasons, the Next Generation Mobile Network (NGMN) Alliance [1] and the Third Generation Partnership Project (3GPP) [2] defined *self-organizing networking* (SON) as a set of principles and concepts to add automation to mobile networks so that they require less maintenance than traditional networks while improving quality of service (QoS). SON functions take as input the measurements generated by mobile networks, and produce output that depends on the purpose of the function.

SON functionalities can be classified in three large categories:

- Self-configuration: functionalities that automate the planning and deployment of the network
- Self-optimization: functionalities that keep the configuration parameters always working at the optimal level to offer the best QoS
- Self-healing: functionalities that automate the solution of problems, reducing human intervention and minimizing downtime. Self-healing includes fault detection, diagnosis, compensation, and recovery.

Self-healing [3] aims to automate troubleshooting, which is one of the most important O&M tasks. The main objective of troubleshooting is finding and fixing malfunctions in the network. In the currently deployed LTE networks, this task is manually done by monitoring some variables that reflect the state of the network. Self-healing functions include data processing algorithms that analyze O&M data in order to identify and fix problems. Some attempts at defining an implementation for self-healing functions have been made using knowledge-based systems (KBSs), such as Bayesian networks [3], fuzzy logic controllers [4], or case-based reasoning [5] for automatic diagnosis. These systems have been proposed theoretically without acknowledging the problems inherent to large databases.

When data largely increase, traditional processing techniques have very poor performance. The big data paradigm deals with this type of dataset by applying new techniques that exploit the latest hardware and software innovations.

In mobile networks, performance measurement data have already passed that threshold; therefore, it must be studied through the paradigm of big data. In [6, 7], a general vision of big data for SON functions was given, without going into details on each SON functionality. This article goes one step further, focusing on self-healing functions. In particular, in this article, self-healing is reframed as a big data problem, and some specific requirements for the development of big-data-compliant self-healing functions are proposed. To further illustrate these principles, some use cases are shown, where modifications on previously existing algorithms that work with big-data-compliant inputs are proposed in order to use big data processing techniques.

In the following section, the data sources and their uses for troubleshooting in mobile networks are reviewed. Next, the big data concept is explained, and the characteristics of self-healing functionalities that are compliant with big data principles are described. After that some use cases are exposed, and big-data-compliant

Emir J. Khatib, Raquel Barco, Pablo Muñoz, and Isabel de la Bandera are with Universidad de Málaga; Immaculada Serrano is with Ericsson.v

0163-6804/16/$25.00 © 2016 IEEE

solutions are proposed. Finally, the conclusions of the study are reviewed.

Troubleshooting Data in Cellular Networks

Each element in a cellular network interacts with the rest of the network in order to establish communication, and each interaction produces several events (connection establishments, handovers, etc.). When a problem happens, the statistical data of these events can hold valuable information about the root cause. Therefore, in modern networks, all the elements save information about the events that they observe.

Data Sources

The troubleshooting process uses data sources that indicate the state and behavior of the network. These data sources are usually recorded on site in the base station (BS: the element that connects the network to the user terminals), which is connected to a data collection subsystem that regularly gathers all the information in a centralized database. The most used data sources are:

Performance Management Metrics: Each BS keeps an array of counters that increase with specific events, such as established or dropped connections. These counters are accumulated over a variable time period known as the report output period (ROP). Other measurements are also taken and averaged during this time interval, such as the received power.

Fault Management Alarms: Along with the counters, BSs monitor specific problematic events. The occurrence of these events is registered in a binary indicator (i.e., the alarm). The nature of alarms is varied, such as software errors or hardware integrity problems.

Configuration Management Parameters: The configuration management (CM) parameters of each BS are adjusted by the engineers or SON functions. These parameters regulate the network operation, so they are important information sources for better understanding how the events affect the network performance.

Call Traces: Measurements taken in the time interval and channel where a communication took place. Each call trace contains registers (e.g., counters and alarms) and measurements related to a specific connection between a user equipment (UE) and the network.

Others: Other information sources that are sometimes used in the troubleshooting process are trouble tickets (previously known problems of the affected sector or neighboring sectors), engineer actions (timestamped actions taken toward solving a previous problem or optimizing the performance), drive tests (on-site measurements of the radio signals received by UEs), customer complaints, and so on.

Performance management (PM) metrics, fault management (FM) alarms, and CM parameters are collected and stored in a raw format in the form of timestamped variables (time series) indexed by BS, but in order to make them more usable, they are usually transformed into more human-readable composite variables, performance indicators (PIs). PIs measure a specific magnitude, such as the proportion of established or dropped connections over a time interval. PIs are aggregated over more varied timescales, such as hours or days. A reduced set of the PIs is representative of the general behavior of the element, and therefore are used to better understand it. These are the key PIs (KPIs). In turn, KPIs are aggregated over groups of elements in order to represent the general behavior of the network or parts of it. Call traces are stored as files indexed by the session identifier, but they also have references to the BS and UE. Some PIs are calculated using the data of call traces. For this, the data in call traces referring to a specific BS are aggregated in *synthetic counters* and then used to calculate the new PIs. These PIs are saved and used along with the PIs calculated from CM/PM/FM.

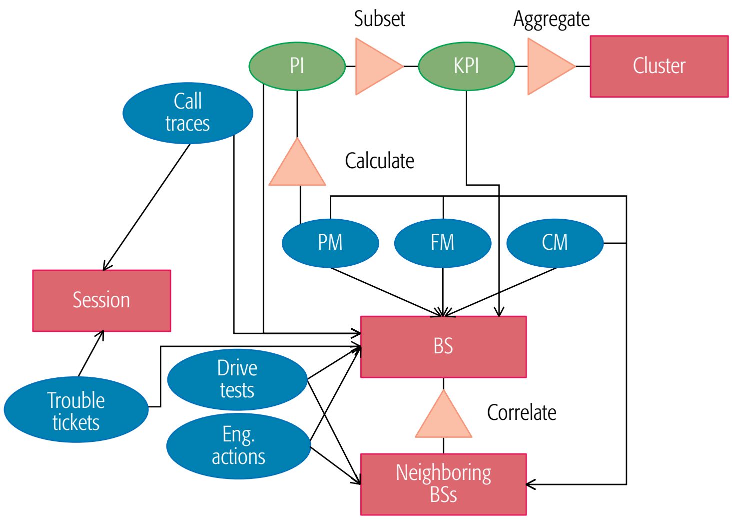

Figure 1. Data types and the relationships established by the information each type contains.

It is important to note that although the data sources are collected physically in individual network elements (usually the serving BSs), the information they contain may be relevant for other elements too, such as neighboring BSs. Figure 1 shows the complex relationships between the data types (marked in blue for the raw collected data and in green for the composite data) and the network elements on which they contain information (marked in red), as well as the processes that relate them (marked in orange).

Troubleshooting Procedure

The manual troubleshooting workflow has four main subtasks.

Detection: The process of determining that there is a problem in the network, and pinpointing the element or elements that are affected. To do this, troubleshooting engineers will usually monitor a very reduced set of KPIs. When one or more of these KPIs is degraded, a list of the worst offenders is extracted, that is, the elements (e.g., BSs) that are degrading the KPI averages most.

Diagnosis: Once the problematic elements have been determined, troubleshooting engineers must find out the root cause (i.e., why they are failing). The study of low-level indicators, as well as logs or event records may help in the deter-

To better understand the meaning of the data, experts often use statistical techniques to simplify the representation of the collected variables. For instance, to detect degradations on PIs or KPIs, thresholding is often used, that is, the value is considered degraded if it is above or below a certain threshold.

mination of the root cause. In some cases, active measurements are taken, such as drive tests.

Compensation: The troubleshooting process may take anywhere between minutes to several days. Therefore, it is important to redirect the resources of the network temporarily to give service to the users in the affected area. Since this temporary configuration is sub-optimal, the users may perceive reduced QoS.

Recovery: Once the root cause is known, the required actions to fix it are taken. These actions range from simple resets or configuration changes that can be ordered remotely, to hardware fixes or replacements that need on-site repairs. The recovery action may or may not solve a problem, so the results of the action are taken into account on subsequent repetitions of the diagnosis subtask.

Manual Data Processing

To better understand the meaning of the data, experts often use statistical techniques to simplify the representation of the collected variables. For instance, to detect degradations on PIs or KPIs, thresholding is often used, that is, the value is considered degraded if it is above or below a certain threshold. Thresholding is specially important in the detection stage, but it is also used to discretize the value of PIs, classifying them as either good or bad values. With discretized variables, heuristic rules of the "if … then …" form are often used in the diagnosis process. This pattern is generally used by experts in their observation of low-level PIs, and sometimes these rules are coded as programs and used in the database containing the PM, FM, and CM variables to test for known fault states. Another technique that is sometimes used to test the relations among variables is correlation. For instance, alarms or engineer actions are correlated with PIs to see to what extent they are responsible for the observed behavior.

There is no normalized course of action for applying these processes. It is usually up to the judgment and experience of the troubleshooting expert which techniques to use and when based on their observations.

Big Data in Self-Healing

In the last few years, big data solutions have been proposed for many applications. Generally, wherever a large user base is present, there are applications where large datasets are used, and therefore, processes will have to be optimized for them. Mobile networks are a perfect example of such a scenario, and in this article, the focus is set on their troubleshooting.

Introduction to Big Data

As a consequence of the decrease in the prices of storage hardware, as well as the increase in bandwidth in mobile networks and the growth in the number of connected electronic devices, the volume of data generated by our society is increasing exponentially. All these data contain information about a wide spectrum of aspects that may be interesting for all types of businesses. As the data produced grow, the information contained in them increases in both quantity and accuracy, and so does the demand for extracting that information.

The amount of available data is often too big and unstructured to be treated with traditional statistic methods to achieve the results and the speed to cater for the needs of the market, so new techniques must be used. Big data is the new paradigm that encompasses the principles and techniques for making sense of data in this new scenario. A data set is considered big-data-compliant when it follows a set of principles known as the 3 Vs of big data [8] as listed below.

Big Volume: The quantity of data that must be processed is large, because either there are many individual small information units (e.g., PM counters that are very simple and structured, but generated by a large number of BSs) or each information unit is large (e.g., call traces, which contain variable fields and many measurements, and the information is potentially referred to several different BSs). The exact boundary for the volume of data to be considered big data is largely dependent on the application, the time constraints, and the available hardware. In the case of mobile networks, each BS produces the CM, PM, and FM data in each ROP. The number of individual variables is on the order of hundreds to thousands for each BS, and the number of BSs in a network ranges in the tens of thousands. Additionally, each of the millions of individual connections generates a call trace.

High Velocity: The information is produced at a rate that requires special techniques in order to process it before new data is produced. Again, there is no exact boundary to consider a data source fast; it depends on the hardware resources that are available. In a cellular network, the data is generated in every ROP, which usually is 15 minutes long. Therefore, the total data must be collected, stored, and processed during this time interval. It is crucial to have quick results that help to prevent severe degradations in the performance of the network. In the case of troubleshooting, all the processes of detection, diagnosis, recovery, and compensation must be done before users perceive a severe loss in QoS. The time frames for this range between minutes and hours. However, currently, with manual troubleshooting, it is common that problems take days to resolve.

High Variety: Data sources have varied formats, which require preprocessing for homogenization or altogether separate processing pipes, often containing unstructured data, which require preprocessing in order to structure them so they can be processed, and are extracted from different physical/logical interfaces (requiring special equipment or software drivers). Also, the data units may contain information about different entities that may or may not be needed for a specific application. As shown previously, in cellular networks, data formats are very varied. For instance, PM and CM variables are given as numerical values, whereas FM variables are Boolean. Call traces register individual events, and for each event, a different set of measurements is provided. Other variables, such as trouble tickets or user complaints, are much more complex, and may contain important information for troubleshooting.

Big data makes heavy use of distributed computing for both data processing (cloud computing) and storage (data warehousing). Some de facto standards are the lambda system architecture [9] for the high-level design of big data systems, the MapReduce [10] programming architecture for low-level algorithm design, and NoSQL [11] databases for storage.

Big Data Techniques

In mobile networks, we propose to use the lambda architecture, which is a common architecture for cloud computing. This architecture is specially designed for data sources where information is being accumulated over time and live results are requested. The lambda architecture has two data pipelines, the batch layer and speed layer.

Batch Layer: In this layer, the algorithms are applied over datasets spanning a long period of time to extract detailed information. This pipeline is slow, since it deals with large volumes of data, so it does not offer immediate results when new data is added. Nevertheless, it must still be fast enough to process the data at the rate at which they are produced.

Speed Layer: It is often necessary to obtain an approximate result that is available immediately after new data is collected. This pipeline processes the latest data along with a data set spanning over a small period of time in the past. It can also use the output of the batch layer for previous periods of time. The results of the speed layer tend to be more inaccurate and prone to errors (due to missing data, input errors, etc.), since the focus is always on speed. In environments where proactive detection of problems is performed, the early approximate results of the speed layer are vital.

To reduce the processing time, parallelization is heavily used in big data. The time a process may take is reduced roughly by a factor of the number of independent processors on the computing cloud where the process is run. A common architecture for parallelization, which can be applied in cellular networks, is MapReduce (Fig. 2), where the data is processed using two functions. The *map* function is applied over each data unit that is independent (i.e., contains the required information without the need for other data units) in a separate process in the computing cloud. The output of the *map* function over each data unit is then aggregated with the *reduce* function.

Data sources do not always provide structured data that can easily be stored in a table without loss of information. The NoSQL paradigm refers to database implementations that do not use the traditional table-based architecture. In wireless networks, we propose to apply the *key-value* model, which is a common architecture, where each *record* (or *value*) is uniquely identified by a hashable variable (the key). In turn, each record can contain a variable number of fields, which simplifies the storage of unstructured data and reduces the required storage space.

Applying Big Data Techniques to Self-Healing

The main objective of self-healing is to automate the troubleshooting task described earlier by using programs that replicate the manual processes. Commercial requirements (the demand for QoS) create the need for a fast and reliable

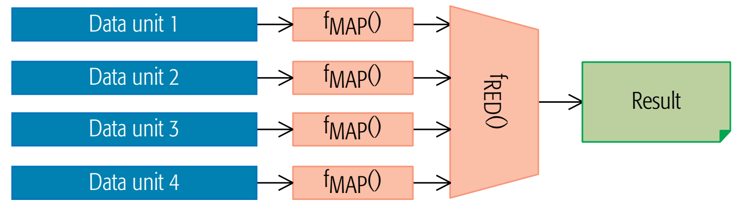

Figure 2. MapReduce architecture.

troubleshooting system that minimizes downtime. Therefore, automation in troubleshooting is required.

Self-healing algorithms are usually implemented using knowledge-based systems (KBSs) that imitate the process of human experts in order to accomplish a task. KBSs are algorithms composed of two parts:

- Knowledge base (KB): a codified representation of the field knowledge, that is, the knowledge that the experts need in order to complete the task. To generate and improve the KB, a continuous data mining (DM) process is run in the batch layer;
- Inference engine (IE): the procedures that use the KB in order to complete the task. The IE conforms to the speed layer.

Some KBSs that have been previously used in troubleshooting are fuzzy logic [4] or Bayesian networks [12]. It is important to follow the guidelines of big data when designing these implementations (i.e., creating parallelizable algorithms). For an algorithm to be parallelizable, its design must guarantee that the final result is the same when it is run as a single process and when the task is divided among multiple instances.

Use Cases

This section presents some self-healing examples and how they can be adressed by means of big data techniques.

Data Reduction

Big datasets are often collected to be inspected by human operators or for unspecified future uses. Therefore, it is a common problem that the relevant information for a specific application is hidden among redundant or even unrelated information. Hence, for many SON algorithms, a necessary prior step is extracting the subset of data where the information is contained. In troubleshooting, some algorithms need to extract information from degraded indicators. These values can be extracted from historical O&M databases that store the timestamped indicator values, that is, all the data types described earlier, but in such databases, both normal and degraded values are indistinctly saved, with no indication of the status (i.e., normal/degraded). Therefore, it is important to devise methods to isolate the time intervals where the behavior of each element is degraded.

Since O&M databases are very large, it is important to perform this process with a parallelizable algorithm. Each parallel process must be independent of the result of other processes. An easy way to divide the work in this case is to analyze each BS individually. The algorithm described in [13] detects degraded intervals (DIs, time intervals where the behavior of a BS is degraded) by analyzing KPIs. For each BS, the time series of one KPI is sequentially analyzed and compared against a *good* threshold and a *bad* threshold. A DI starts when the value of the KPI falls below (or rises above) the *bad* value, and ends when it crosses the *bad* threshold again and then rises above (or falls below) the *good* threshold. This algorithm can easily be parallelized since each BS can be processed by a separate instance (by implementing the detection algorithm as a *map* function). The input database is filtered by the BS, and each instance of the *map* function reads the KPI time series of the filtered table. On the output, the result is aggregated, adding a field identifying the BS and finally applying a *reduce* function that saves all the detected DIs in the same table. An example of a result of the algorithm is shown in Fig. 3, where the *retainability* KPI is used to determine the DIs. This specific use case comes from a dataset containing 47 samples, each with two weeks of hourly data, making a grand total of 15,792 entries.

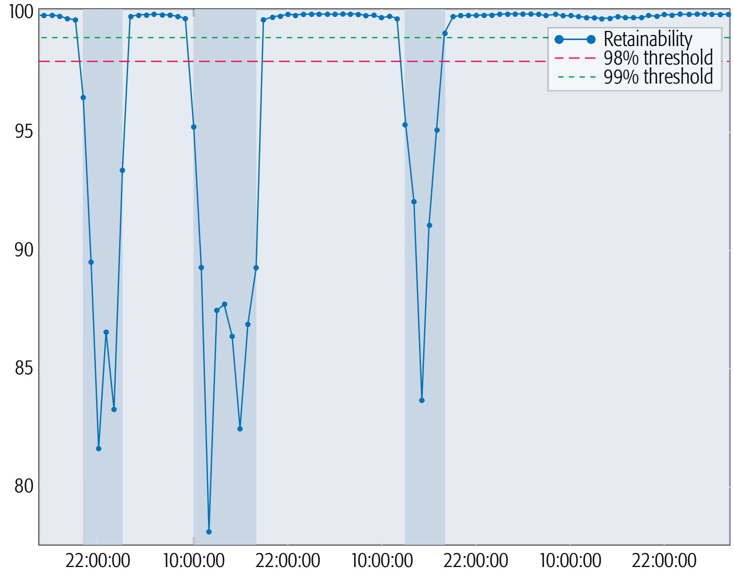

Figure 3. Detected degraded intervals.

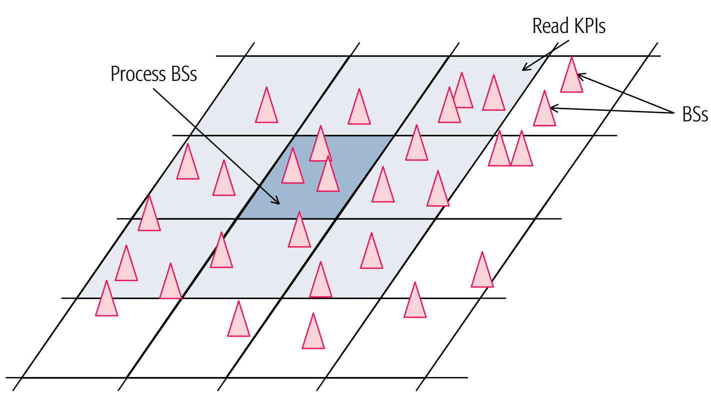

Figure 4. Terrain division for the detection of sleeping cells.

The total number of detected DIs is 359, and since the KPIs are aggregated over the duration of the DI, the reduced database contains 359 entries, representing a reduction of 97.73 percent of data volume.

Sleeping Cell Detection

A common problem in mobile networks is cell outage or sleeping cells, that is, cells that should be providing service but are not doing it at all for some reason. In scenarios where the density of cells is high, this problem is especially hard to detect, since the users are redirected to neighboring cells. Sleeping cells produce a low QoS, since the optimal cell for the affected users is not being used. This is usually done using some availability KPIs and alarms which indicate that the cell is down. But in a major outage where the whole BS is down and not responding to status queries (e.g., a major software fault or power outage), the network operator will not be able to detect the fault quickly. An alternate approach to outage detection is using the neighboring BS measurements to detect the outage by calculating its impact. Since in a network the number of BSs is normally large, and for each one all the data from its neighbors is used, it is easily determined that this is a big data problem. Moreover, the timeframe to detect and correct an outage is usually low, since the service for the affected users is degraded, leading to a bad user experience.

In [14], an algorithm for detecting sleeping cells based on decreased handovers with neighboring cells is described. To find sleeping cells, for each BS, the number of incoming handovers for the current and previous ROP is aggregated from the neighbor BSs' outgoing handovers. If the handovers have suddenly dropped to zero and the readings of other PIs (or the lack of PIs) of the cell indicate a malfunction, the cell is marked as a sleeping cell. To apply this algorithm under the big data principles, it should be considered that the full network has to be analyzed over a limited time (a ROP before new data is received). Thus, it is essential that multiple instances of the algorithm analyze separate parts of the network. In order to achieve this, the terrain can be divided into partitions that are the size of the maximum distance between neighbors, as shown in Fig. 4. Each instance of the algorithm sequentially tests each of the BSs contained in one partition by looking into the data of its neighbors, which are contained in the adjacent partitions. Therefore, each parallel instance only works with a reduced database containing only the data of the current and adjacent partitions.

Figure 5 shows the results obtained in [14] for the algorithm compared to other methods when applied to a simulated LTE network: availability PIs (the detection is made by monitoring certain PIs of the cell) and lack of PIs (a cell is selected as a sleeping cell if there are no PIs available). For each method, the results show the false positive rate (i.e., the percentage of wrong detected cases among the total outage cases) and the false negative rate (i.e., the percentage of undetected cases among the total of normal cases simulated). The results show that the proposed method is able to detect most simulated outages, leading to a low percentage of false negatives (5.9 percent), while availability PIs and lack of PIs methods present a high percentage of false negatives. These results show that the increase in the volume of processed data improves the detection capacity.

Diagnosis Based on KPI Correlation

In troubleshooting, once a problem is detected in a BS, the diagnosis task is launched. Normally, detection is done by observing very high-level KPIs, whereas diagnosis takes into account all the available information to find the root cause. A very important source of information is to find which PIs correlate the most with the occurrence of the problem. Since the KPIs are an indication of the general behavior of the BS, a list of the most correlated indicators will give an important clue to the root cause analysis. In [15], a method that performs this analysis is described. This algorithm takes into account the PIs of the affected BS and the neighboring sectors in order to simplify the task of diagnosis. The process of calculating the correlation of two time series is a computationally heavy operation, and the number of correlations that must be done is high (all the PIs of the analyzed sector, plus all the PIs of each neighboring sector), qualifying it as a big data problem; but since each PI can be processed independently, the algorithm is easily parallelizable. Again, the correlation process is implemented as a map function, and a reduce function creates a list of PIs ordered by correlation. In Fig. 6, a time series of a KPI in the diagnosed BS (the number of radio resource control connections) is shown, along with a highly correlated PI of a neighboring BS (a counter of bad coverage reports) and their correlation.

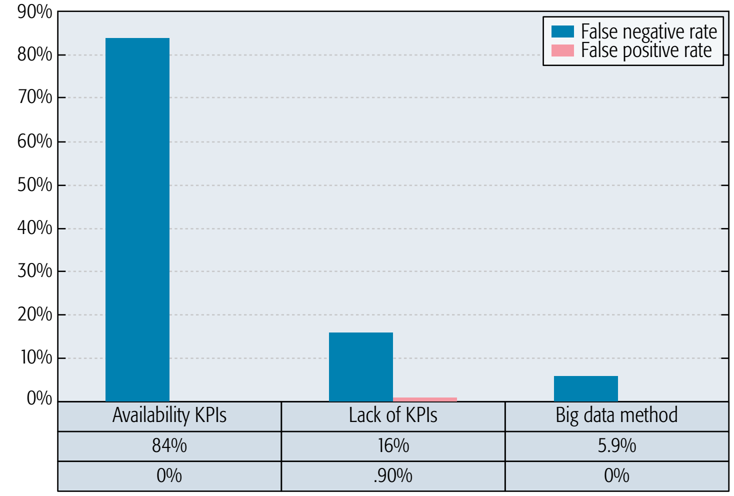

Figure 5. Comparison between the proposed big data method and other common techniques.

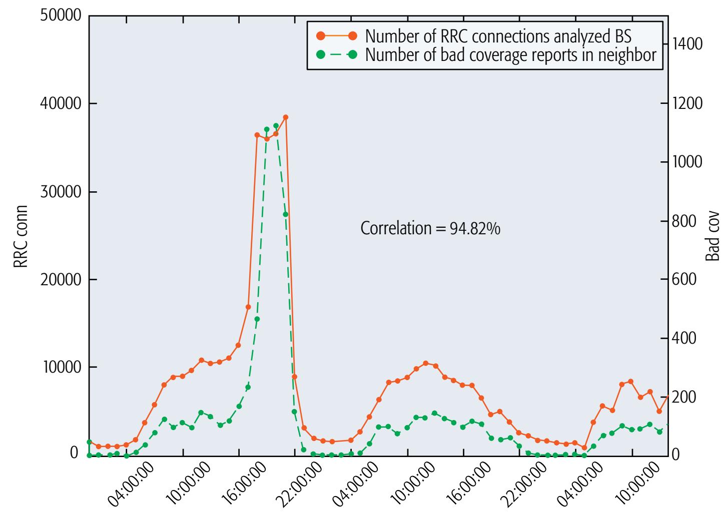

Figure 6. Correlation between the KPI and a PI of a neighboring sector.

Conclusions

Automation of the O&M of mobile networks is a need for mobile operators due to several forces driving costs up: increased traffic volume, variety of services, number of subscribers, demand for QoS, and competitiveness among operators. Modern networks (especially LTE and the future 5G networks) produce large volumes of O&M data, which contain the information needed for the self-healing functions that operators need. In fact, the volume of this data, along with the variety and time restrictions, calls for the use of the big data paradigm. Big data is a set of new techniques that exploit the ever increasing processing power of computer networks. Specifically, in this article, the main focus is parallelization. Several use cases have been displayed as examples where data processing algorithms can extract information contained in big-data-compliant data sources. In order for these algorithms to work properly, small modifications have been proposed in order to run the algorithm in parallel processes that do not interfere with each other. Parallelization reduces the processing time (or increases the processing capability) of these algorithms, so SON functions can be applied in a manner that is transparent for the users of the network. The proposals of this article can be extended to any SON algorithm that can be extended in a way which guarantees that the results of the parallelized version are the same as the original. In order to further test the proposed solutions of the use cases, they can be implemented using the commercially available big data processing solutions.

Acknowledgment

This work has been partially funded by Optimi-Ericsson, Junta de Andalucía (Agencia IDEA, Consejería de Ciencia, Innovación y Empresa, ref. 59288; and Proyecto de Investigación de Excelencia P12-TIC-2905) and ERDF.

References

[1] NGMN, *Use Cases Related to Self-Organising Network, Overall Description*, Apr. 2007.

[2] 3GPP, "Telecommunication Management; Self-Organizing Networks (SON); Self-Healing Concepts and Requirements," Rel. 11, TS 32.541.

[3] R. Barco, P. Lázaro, and P. Muñoz, "A Unified Framework for Self-Healing in Wireless Networks," *IEEE Commun. Mag.*, vol. 50, no. 12, Dec. 2012, pp. 134–42.

[4] E. J. Khatib *et al.*, "Data Mining for Fuzzy Diagnosis Systems in LTE Networks," *Expert Systems with Applications*, vol. 42, no. 21, 2015, pp. 7549–59.

[5] P. Szilagyi and S. Novaczki, "An Automatic Detection and Diagnosis Framework for Mobile Communication Systems," *IEEE Trans. Network and Service Management*, vol. 9, no. 2, June 2012, pp. 184–97.

[6] N. Baldo, L. Giupponi, and J. Mangues-Bafalluy, "Big Data Empowered Self Organized Networks," *Proc. Euro. Wireless 2014*, VDE, 2014, pp. 1–8.
[7] A. Imran and A. Zoha, "Challenges in 5G: How to Empower SON with Big Data for Enabling 5G," *IEEE Network*, vol. 28, no. 6, Nov. 2014, pp. 27–33.
[8] P. Russom, "Big Data Analytics," *TDWI Best Practices Report*, 4th qtr., 2011.
[9] N. Marz and J. Warren, *Big Data: Principles and Best Practices of Scalable Realtime Data Systems*, 1st ed., Manning, 2015.
[10] J. Dean and S. Ghemawat, "Mapreduce: Simplified Data Processing on Large Clusters," *Commun. ACM*, vol. 51, no. 1, Jan. 2008, pp. 107–13.
[11] R. Cattell, "Scalable SQL and NoSQL Data Stores," *ACM SIGMOD Record*, vol. 39, no. 4, 2011, pp. 12–27.
[12] R. Barco *et al*., "Continuous Versus Discrete Model in Autodiagnosis Systems for Wireless Networks," *IEEE Trans. Mobile Computing*, vol. 7, no. 6, June 2008, pp. 673–81.
[13] E. J. Khatib *et al.*, "LTE Performance Data Reduction for Knowledge Acquisition," *Proc. GLOBECOM Wksps.*, Dec. 2014, pp. 270–74.
[14] I. de la Bandera *et al*., "Cell Outage Detection Based on Handover Statistics," *IEEE Commun. Lett*., vol. 19, no. 7, 2015, pp. 1189–92.
[15] P. Muñoz *et al*., "Fault Diagnosis in Networks," Patent 4 24, 2015.

Biographies

Emil J. Khatib (emil@uma.es) received an M.Sc. degree in telecommunications engineering from the University of Málaga, Spain, in 2010. Until 2012 he worked as a research assistant in the Department of Systems Engineering and Automatics at the University of Malaga. Since spring 2012, he has been pursuing his Ph.D. studies in the Department of Communications of the University of Malaga. His research interests include communication protocols and artificial intelligence applied to telecommunications.

Raquel Barco holds an M.Sc. and a Ph.D. in telecommunication engineering. She has worked at Telefónica in Madrid, Spain, and at the European Space Agency (ESA) in Darmstadt, Germany. From 2000 to 2003, she worked part-time for Nokia Networks. In 2000 she joined the University of Málaga, where she is currently an associate professor. Her research interests include satellite and mobile communications, mainly focusing on self-organizing networks.

Pablo Muñoz received his M.Sc. and Ph.D. degrees in telecommunication engineering from the University of Málaga in 2008 and 2013, respectively. From 2009 to 2013, he was a Ph.D. fellow in self-optimization of mobile radio access networks and radio resource management. Upon completing his Ph.D, he has continued his career at the University of Málaga as a research assistant within an R&D contract with Optimi-Ericsson focusing on self-organizing networks. In 2014 he held a postdoctoral position granted by the Andalusian Government in support of research and teaching.

Isabel de la Bandera received her M.Sc. degree in telecommunication engineering in 2009 from the University of Málaga, where she is currently working toward her Ph.D. degree in telecommunications engineering. In 2008, she worked on radio-frequency identification projects with the Communications Engineering Department, University of Málaga. Since February 2010, she has been working on projects concerning radio resource management in next-generation mobile networks.

Immaculada Serrano obtained her M.Sc. degree in Universidad Politécnica Valencia. She collaborated with the Communication Department investigating arrays of antennas and completed her research at the University of Sheffield in 2003. She specialized further in radio after complementing her education with a Master's in mobile communications from Universidad Politecnica Madrid. In 2004 she joined Optimi and started a broad career in optimization and troubleshooting of mobile networks for different operators around the world, including a variety of consulting, training, and technical project management roles. In 2012, she moved to the Advanced Research Department with a strong focus on automatic troubleshooting and self-healing.

AD HOC AND SENSOR NETWORKS

Edoardo Biagioni

Silvia Giordano

Research in ad hoc and sensor networks has produced a plethora of papers, and very interesting work is still being published in this area, evolving in parallel with technological innovation, and intersecting with novel areas, particularly the Internet of Things, and more established areas such as wildlife monitoring.

One key aspect of wireless research has been and still is performance evaluation. In the past there has been deserved criticism of the way many wireless solutions have been evaluated, and the article from Papadopoulos, Kritsis, Gallais, Chatzimisios, and Noel, "Performance Evaluation Methods in Ad-Hoc and Wireless Sensor Networks: A Literature Study," deals with this unsolved problem. The authors analyze a large set of statistics on articles (674 papers in total) published in top representative conferences (i.e. ACM/IEEE IPSN, ACM MobiCom, ACM MobiHoc, and ACM SenSys) related to ad hoc and wireless sensor networks (WSNs) over the 2008-2013 period, to see how perfomance has been evaluated and whether this evaluation is sound. Interestingly, the authors find that there is a tendency to rely, more and more, on custom or open testbeds. The article also discusses simulation and experimental approaches as possible alternative to real testbeds.

From the beginning sensor technology has been the technology used for animal monitoring, and in particular for wildlife. The next article, "From Radio Telemetry to Ultra-Low Power Sensor Networks: Tracking Bats in the Wild" by Dressler, Ripperger, Hierold, Nowak, Eibel, Cassens, Mayer, Meyer-Wegener, and Kölpin, demonstrates how ad hoc and sensor networks can progress with technological innovation, and presents a new approach for monitoring based on ultra-low power systems. The authors then specifically apply this approach to bat monitoring, and validate the results from the first field experiments.

As stated above, ad hoc and sensor technology has proved to be a fundamental communication technology for the Internet of Things. The next article, "CheepSync: A Time Synchronization Service for Resource Constrained Bluetooth LE Advertisers" by Sridhar, Misra, Gurinder, and Warrior, discusses the problem of time synchronization, which is a key service for current and future applications in sensor networks and IoT. In this article the authors propose a time synchronization solution for Bluetooth Low Energy (BLE), and demonstrate its validity in both empirical and experimental settings.

These three articles are good examples of how research in ad hoc and sensor networks can still be fascinating and effective, and applicable to real life and technological innovation.

BIOGRAPHIES

EDOARDO BIAGIONI (esb@hawaii.edu) is an associate professor in the Department of Information and Computer Sciences at the University of Hawaii at Manoa. His research interests focus on networking, with an emphasis on ubiquitous wireless networking, but they have over time ranged widely from security to high-performance computing, programming languages, and human-computer interfaces. He received his Ph.D. degree from the University of North Carolina at Chapel Hill, and has been a series co-editor for *IEEE Communications Magazine* since 2006.

SILVIA GIORDANO [M] (silvia.giordano@supsi.ch) received her Ph.D. from EPFL, and she is a full professor at SUPSI, Switzerland, and an associate researcher at CNR. She directs the NetworkingLab. She has published extensively in the areas of QoS, traffic control, and wireless and mobile networking. She is a co-editor of the books *Mobile Ad Hoc Networking* (IEEE-Wiley, 2004) and *Mobile Ad Hoc Networking: The Cutting Edge Directions* (Wiley 2013). She is an ACM Distinguished Scientist, ACM Distinguished Speaker, and on the Board of ACM N2Women. She has been a series co-editor for *IEEE Communications Magazine* since 2004.

Performance Evaluation Methods in Ad Hoc and Wireless Sensor Networks: A Literature Study

Georgios Z. Papadopoulos, Kosmas Kritsis, Antoine Gallais, Periklis Chatzimisios, and Thomas Noël

The authors analyze a large set of statistics in articles published in the top representative ad hoc and wireless sensor networks related conferences during the period 2008–2013. We focus on the evaluation methodologies provided by researchers. More specifically, our goal is to explore the role of simulators and testbeds in the theoretical analysis of a model throughout the protocol development procedure.

Abstract

Verification of theoretical analysis is a vital step in the development of an application or a protocol for wireless networks. Most proposals are evaluated through mathematical analysis followed by either simulation or experimental validation campaigns. In this article, we analyze a large set of statistics in articles published (674 papers in total) in the top representative ad hoc and wireless sensor networks related conferences (i.e., ACM/IEEE IPSN, ACM MobiCom, ACM MobiHoc, and ACM SenSys) during the period 2008–2013. We focus on the evaluation methodologies provided by researchers. More specifically, our goal is to explore the role of simulators and testbeds in the theoretical analysis of a model throughout the protocol development procedure. We show that there is a tendency for more and more researchers to rely on custom or open testbeds in order to evaluate the performance of their proposals. Simulators indeed fail to reproduce actual environmental conditions of deployed systems. Experimentation with real hardware allows our research community to mind the gaps between simulation and real deployment. Still, as the experimental apFproach through custom testbeds results in a low reproducibility level (i.e., 16.5 percent), we investigate to what extent such performance evaluation methods will be able to bridge those gaps. We finally discuss experimental testbeds and their potential to replace simulators as the cornerstone of performance evaluation procedures.

Introduction

Ad hoc and wireless sensor networks (WSNs) have enabled a large variety of applications. Environmental and wildlife monitoring, clinical medical and home-care monitoring, monitoring and control of industrial processes including agriculture, and smart houses or cities are just some of the examples of ad hoc and WSN applications, where low-cost and easily deployed multi-functional sensor nodes are the ideal solution [1]. As a result, during the past few years we have experienced the emergence of a new paradigm called the Internet of Things (IoT) in which smart and connected objects cooperatively construct a (wireless) network of things [2]. However, the unique features of ad hoc and WSN technologies can pose significant challenges. Hence, envisioned solutions must be verified before being deployed in a real-world WSN deployment, either by utilizing simulators or emulators or through experimentations by employing testbeds.

Simulation evaluation is an essential phase during the design and development of an ad hoc or WSN infrastructure. However, environments in which ad hoc or sensor networks evolve are often application-specific and too complex to be reproduced precisely. More specifically, simulators allow users to implement some basic assumptions (e.g., link quality, radio propagation, medium interferences, topologies) [3]. Although the majority of the simulation models cannot capture real world complexity [4, 5], they are often utilized as a first step. Our purpose is to show that this step is not sufficient to demonstrate the consistency of a solution, as well as that low-cost devices have led researchers and engineers to enrich performance evaluation with testbeds.

Experimental evaluation is performed over either custom or open testbeds, and exhibits unexpected failures and problems that could possibly be faced during real-world deployments. Even though performing well over testbeds, they still remain in vitro deployments with more or less controlled environment conditions. Such a proof of concept must then be transposed into the real world. Designing and setting up a complete ad hoc or WSN system under real conditions that can support robust applications is a very complex task [6]. Researchers and production system architects first need an appropriate plan of deployment and later a number of tools, simulators/emulators, and testing facilities for real experiments in order to initially validate their concept or model and then develop the appropriate infrastructure.

Throughout this study, we compile a large set of statistics from a literature review of 674 articles published in top conferences on ad hoc and WSNs over the 2008–2013 period. We focus on the evaluation provided by authors, and especially to what extent experiments on testbeds have become a must for performance evaluation of new communication algorithms and protocols. Hence, we exhibit the tendency for performance evaluation procedures to rely on experiments with real hardware and environments, to the detriment of

Georgios Z. Papadopoulos, Antoine Gallais, and Thomas Noël are with the University of Strasbourg. Kosmas Kritsis and Periklis Chatzimisios are with Alexander TEI of Thessaloniki.

 0163-6804/16/$25.00 © 2016 IEEE

simulations. The question of scientific results vs. proof of concepts therefore arises. Indeed, we discuss the meaning of reproducibility and of a proof of concept as a prototype being designed to determine feasibility. In this article, we also analyze the selection of the evaluation methodology (e.g., simulator, testbed), and simplicity of the overall design that should be provided for validation, understanding, and explanation. Finally, this work aims to investigate and gather the pros from both simulation and experiments so that real-world experiments could lead to reproducible scientific results for our research community.

Performance Evaluation Procedures

In a typical research process cycle, once the modeling phase is done, network researchers and developers continue with the validation procedure in which they evaluate their concept by using either a simulator or an emulator. Later, network engineers and developers may proceed with experimentation to further cross-verify their proposal [7]. Thus, once both the simulation performance and experimental measurements are satisfactory, real deployments can be initiated.

Simulating Protocols or Experimenting Algorithms

Since we face complex environments that are very difficult to theoretically analyze, and we also take into account the difficulties of setting up a real-world (e.g., large-scale) deployment, simulations are often considered as the optimal approach for studying the performance of ad hoc and WSNs. Many open source and freely available simulators allow users to have better control of the nodes by often employing a graphical user interface (GUI), and to retain or simplify some assumptions in order to evaluate their solutions. Simulation evaluation is a provisioning procedure during protocol development. However, even if the simulation performance presents coherent results with mathematical analysis, past real-world deployments show that it is not recommended to proceed directly with real deployment since engineers may face unpredictable phenomena such as node crashing or network disconnection [5, 8]. Intermediate experimentation platforms can therefore be considered to bridge the gap between simulations and real-world deployments. Nevertheless, while simulations can offer wider sets of assumptions to test and therefore potentially more complete evaluations, testbeds impose many characteristics (e.g., physical environment, hardware, network topology). Such facilities offer the opportunity to have their solutions face real conditions, thus being more realistic than those modeled under software simulators. However, numerous parameters (e.g., radio dynamics, link stability and symmetry, impact of the weather on communications [9]) appear so unpredictable that they may lead to results that cannot be reproduced with sufficiently tight confidence intervals. The ambition to obtain scientific results should then lead researchers to allow for further repeatability of the presented results. As a result, during the simulation evaluation the environmental conditions should not affect the behavior of the nodes. Hence, it would be ideal if the authors first verify their model by employing experimental tests in order to reflect the reality that their proposals would face during real deployment.

A Thorough Literature Study

In this article, we carry out a thorough study over top representative conferences that are strongly related to ad hoc and WSN research fields. In particular, we have studied all articles that have been published at the ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN), ACM Annual International Conference on Mobile Computing and Networking (MobiCom), ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc), and ACM Conference on Embedded Networked Sensor Systems (SenSys) in order to derive the current tendency of the validation methodology that authors follow with respect to previously reported issues. Hence, we go through and study 674 articles in total, published in the conference proceedings for the last six years from 2008 to 2013, of which 596 are related to ad hoc and WSN (Fig. 1a). Indeed, we identified 78 articles that deal with other wireless technologies, such as WiFi and WiMAX, that are studied in the context of cellular networks. All of these papers have been found in MobiCom (i.e., 140 out of 185) and MobiHoc (i.e., 142 articles out of 175) conferences (Fig. 1b), which are not entirely dedicated to ad hoc and WSN but have a broad scope on mobility and wireless communications. In the rest of the article, we further emphasize our investigation over these 596 articles. During our investigation, we obtain a plethora of information for each work and then categorize the considered articles based on their common features.

Figure 1c provides detailed information about the total number as well as the ad hoc and WSN related published articles per proceeding year. We actually observe that there is a decreasing tendency of published articles in the proceedings; indeed, we identified 43 articles less from 2008 to 2013. More specifically, MobiHoc and IPSN reduced the total accepted articles from 44 to 24 (MobiHoc) and from 41 to 24 (IPSN), respectively, while MobiCom and SenSys kept a steady flow.

Modern technologies have introduced the feature of mobility. Consequently, the research community focuses on developing and testing such aspects and scenarios. Our study results justify this trend, due to the 148 articles (52.1 percent) that simulated mobile scenarios. Still, our statistical results for MobiHoc and MobiCom, the mobile oriented conferences, show that not all of their articles implement mobility scenarios. For instance, during the 2008 MobiHoc conference we determined only 13 out of 28 simulation-based articles that introduced mobility in their tests. As shown in Fig. 2a, 57 percent of articles involving mobility are induced less by our conference sample (half of the conferences, MobiCom and MobiHoc, being theoretically focused on mobility-related topics) than by the global enthusiasm for mobile scenarios.

Results of Analysis

Evaluation Procedures

In this subsection, we expose our analysis on the validation procedures that the authors followed. As a first step, we aimed to categorize

During the simulation evaluation the environmental conditions should not affect the behavior of the nodes. Hence, it would be ideal if the authors first verify their model by employing experimental tests in order to reflect the reality that their proposals would face during real deployment.

the reviewed articles according to the employed evaluation method. In particular, we examine the proportion of simulation, experimental, and mathematical (i.e., modeling or analysis) evaluated works. Our primary analysis exposes interesting results. More specifically, our investigation shows that the majority (i.e., 561) of the articles provide an analytical representation of their solution. The remaining 35 have only simulation or experimentation results. Furthermore, 284 verify their proposal by employing simulation evaluation, while on the other hand, 392 of the articles include experimental evaluation for their validation. Finally, only one out of five (i.e., 20.3 percent) articles examines all three phases of the research process cycle (i.e., analysis, simulation, and experimentation). The number of articles with the previously stated properties (with respect to the 596 studied papers) is illustrated in Fig. 2b.

We now present the characteristics of the articles that we studied. The percentage of simulation vs. experiment-based studies (with respect to 596 studied articles) is illustrated in Fig. 2c. As can be observed, while simulations and experiments used to be equally deployed until 2009, the usage of simulations is decreasing every year (except in 2011), while experimentations still remain present at a relatively stable rate.

Over the 2008–2013 period, 284 studies followed a simulation evaluation to test their proposal. We noted the simulator usage, the scales of simulated networks, and the programming languages used for custom simulators (Fig. 3). Only 43.3 percent of articles are validated through a known simulator, while 42.3 percent did not even provide any information about the tool that their authors utilized (Fig. 3a). Finally, 14.4 percent (with respect to 284 studied simulation-evaluated articles) developed a homemade simulator (Fig. 3a) by utilizing programming languages such as Python and Java (Fig. 3c for the distribution of the most popular programming languages).

We are next interested in determining the usage of the simulators. As can be observed in Fig. 3b, MATLAB is the first choice in our community, counting more than 35 articles, followed by TOSSIM, which was reported in almost 20 articles. Furthermore, Network Simulator 2 (ns-2) comes in third with 13 articles.

Nowadays, the research community is able to evaluate proposed protocols, models, and even new technologies over open testbeds at a very large scale [2]. Increasingly, network researchers are using experimentations to enlarge the scope of their performance evaluation (Fig. 2c). Moreover, as can be observed in Fig. 4a, our investigation shows that the majority of researchers, 91.3 percent, chose to set up their own testbeds. Even though to the current day there are a number of open facilities providing developers with the infrastructure needed for experimental ad hoc, WSN, or IoT studies, only 10.7 percent of the articles use open platforms. Our compiled statistics tend to show that researchers would rather favor their own setups for small-scale deployments. In fact, among the 392 articles exposing experimental results, 78 percent of them do not exceed 40 nodes for their experimental setup (Fig. 4a). Hence, the increased difficulty of apprehending a remote open testbed (e.g., specific hardware and software, network topology, booking procedure) may have induced researchers to set up their own relatively small-scale networks.

Finally, we evaluated the popularity of the devices in homemade experiments. In Figs. 4b and 4c the utility of the open testbeds and motes is presented. Even though a small number of articles experimented over open testbeds, we point out the popular open platforms. As observed in Fig. 4b, Harvard's Motelab comes first (11 articles), followed by TWIST (10 studies). That can be explained simply as those facilities were the first to open up to the scientific community. Regarding the Indriya testbed, even though

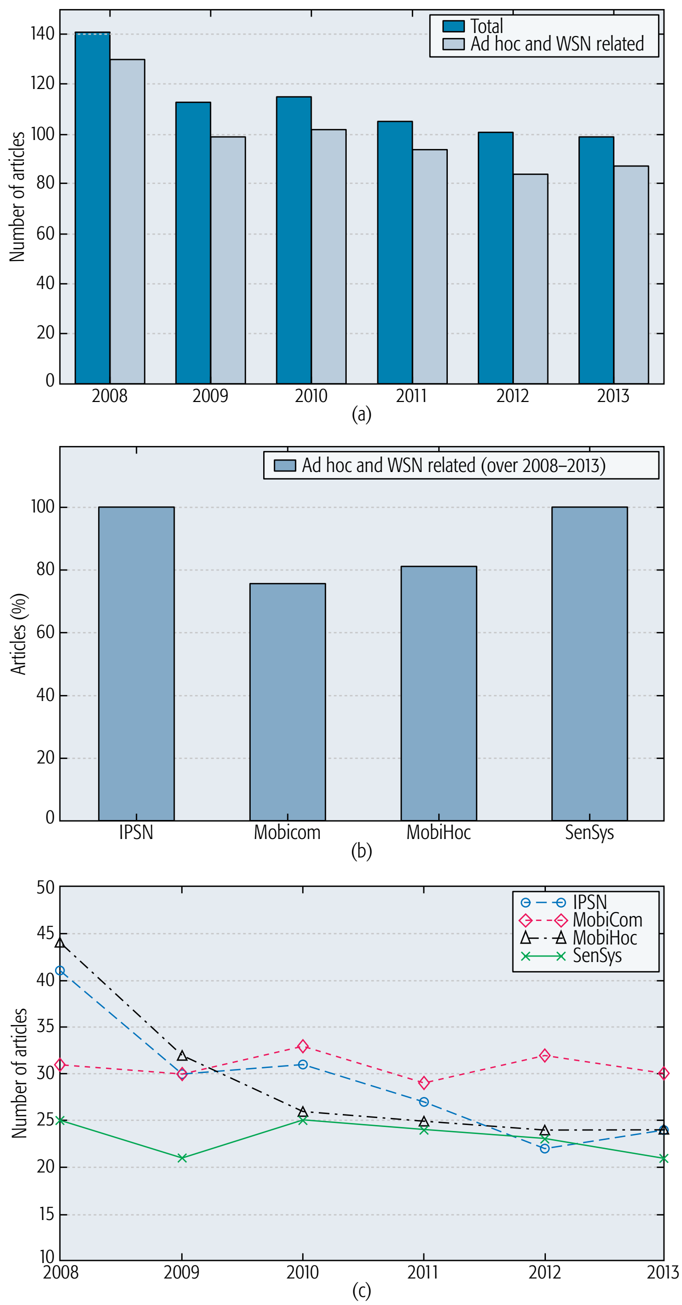

Figure 1. Published articles in ACM/IEEE IPSN, ACM MobiCom,, ACM MobiHoc, and ACM SenSys from 2008 to 2013: a) number of articles per year (all conferences are considered); b) appropriateness of our conference sample.; c) publication flows over the 2008–2013 period.

it was made available only in 2011, it was used in eight articles. The fact that users can interact with the testbed through the same intuitive web-based interface as MoteLab's could explain this success among the community.

Reproducibility

We continue our study by investigating the feasibility of reproducing results that are presented in the reviewed articles, for both simulation and experimental campaigns. To proceed, we looked for some critical information (e.g., simulation setup, simulator indication, simulator details such as version or library, number of nodes) that should be provided by the studied articles. In order to reproduce the proposed solution, we assumed that the authors should provide a complete simulation or experiment settings subsection.

Regarding the simulation-based evaluations, while only 43.3 percent of the articles indicate the simulator, 78.5 percent of those do provide some details about simulation setups. Among those, 72.5 percent give the precise number of involved nodes. Finally, we decided on incomplete setups as soon as there was a lack of critical details regarding the tools used during simulations. For instance, as discussed earlier, MATLAB stands as the most popular software for simulations. In order to use it as a network simulator, researchers must import external libraries (e.g., as developed by the WISLAB[1] team). It is difficult, if not impossible, to reproduce a simulation study when the version of a publicly available simulator is unknown, and only 21.5 percent provide us with the employed version or the utilized library of the simulator, which essentially concludes our outcome about the reproducibility of the simulation-evaluation articles.

We followed a similar methodology for the experimental-based validations. Taking into account the nature of open platforms (42 articles), we consider that these articles overall are reproducible. However, we counted eight papers where the authors tested their ideas over both custom and open testbeds, with only three of them providing enough information to be assumed reproducible. On the other hand, the experimental results that are retrieved through homemade testbeds can be considered as difficult or even impossible to reproduce. This is because most of them are deployed in offices, houses, or even outdoor installations where the environmental radio activity varies, due to the interpolation of external features including mobile phones, wireless routers, access points, and so on. Nevertheless, due to the nature (e.g., application layer) of the tested solution, we detected 31 homemade-based studies that may be reproduced. Finally, by summarizing the previous statements, we calculated that only 16.5 percent (65) of the experiment-based papers present reproducible results.

Further Discussions

Scientific Results or Proofs of Concepts?

Scientific results are expected to be repeatable, while a proof of concept is a realization of an idea that demonstrates its feasibility. Our initial investigation shows that most of the authors chose to validate their proposals over experimental evaluation. Our investigation highlights some interesting tendencies in the networking scientific community, especially around ad hoc and WSNs. As previously presented, an increasing number of

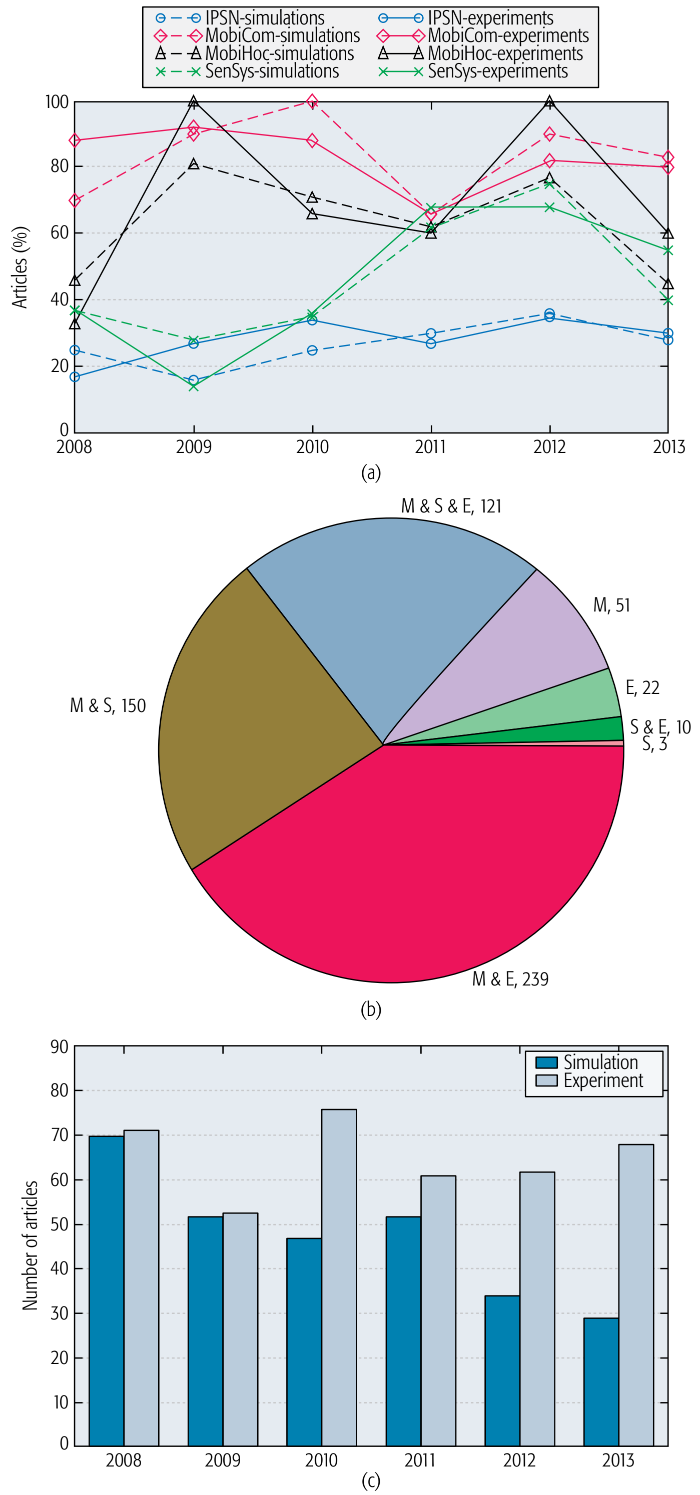

Figure 2. Mobility scenarios in performance evaluation procedures: a) mobility scenarios in performance evaluation procedures; b) use of mathematics (M), simulations (S), experiments (E), and their combinations in validation procedures of 596 ad hoc and WSN related articles; c) total simulation vs. experimentation evaluated articles.

[1] http://wislab.cz/

papers validate their proposals by using experimental evaluations.

We focused on the simulation and experimentation setups in order to determine if they were sufficiently described to allow for repetition of the evaluation procedure. While Kurkowski *et al.* focused on mobile ad hoc networking (MANET), hence looking for simulation parameters specific to mobility (e.g., speed of nodes, speed delta, pause time, pause delta), we aimed at a larger scope by gathering various sets of setup parameters. This is especially true for all observed experimentations among which setups are highly different (e.g., hardware, physical topologies, radio environment). The reproducibility level of experimental studies is lower than the simulation one. This is even more dramatic as this latest has not varied much since the study of Kurkowski *et al*. More specifically, the authors identified 29.8 percent of the simulation-based articles that did not identify the simulator used in the research. As mentioned earlier, regarding the four conferences we observed over the 2008–2013 period, this proportion rose to 42.3 percent. In addition, they calculated only 12.1 percent of the articles where the simulator version was mentioned. Furthermore, the authors were concerned that more than 90 percent of the published results might include bias. As result, they conclude that approximately 12 percent of the MobiHoc simulation-based results appear to be repeatable. In [10], numerous pitfalls throughout the simulation life cycle had already been observed. Those tendencies, as already highlighted by Kurkowski *et al.*, take away from the goals of making the research repeatable, unbiased, realistic, and statistically sound.

As previously observed, over the last six years, fewer and fewer papers have actually considered simulations during their performance evaluation process. Still, the simulation phase allows researchers to demonstrate that the main principles of their proposal are indeed effective before implementing them over a testbed [7].

However, in order for users to be able to continue their proof of concept validation, we can avoid the necessity of getting familiar with various simulators and testbed platforms. Emulators such as TOSSIM[2] and COOJA[3] were developed to bridge the gap between simulation and experimentation, by being very close to real embedded systems in terms of architecture compilation targets. In fact, by utilizing these simulators, the very same code remains unchanged over the transfer from simulation to experimental campaign.

We are coming to a trade-off between realism and reproducibility. More specifically, on one hand there are more published articles that are (closer to) real deployment, while on the other hand the reproducibility level of the studies decreases. So far, the proportion of papers using experimentations that allow the conditions of an experiment to be reproduced remains very low (< 11 percent). Moreover, all those testbeds are highly different (e.g., hardware, physical topologies, radio environment), and each would require specific guidance to allow for scientific results to be obtained.

In [10], the authors proposed simulation study guidance. If the enthusiasm for experiments in networking scientific papers is to be confirmed, we should also be able to establish such mandatory steps to ensure statistically sound results. The significant number of open access and large-scale testbeds that have been deployed in recent years [2] provide appropriate tools and experimental facilities for researchers and engineers to perform real experiments in order to further analyze their protocols. Open testbeds allow users to easily

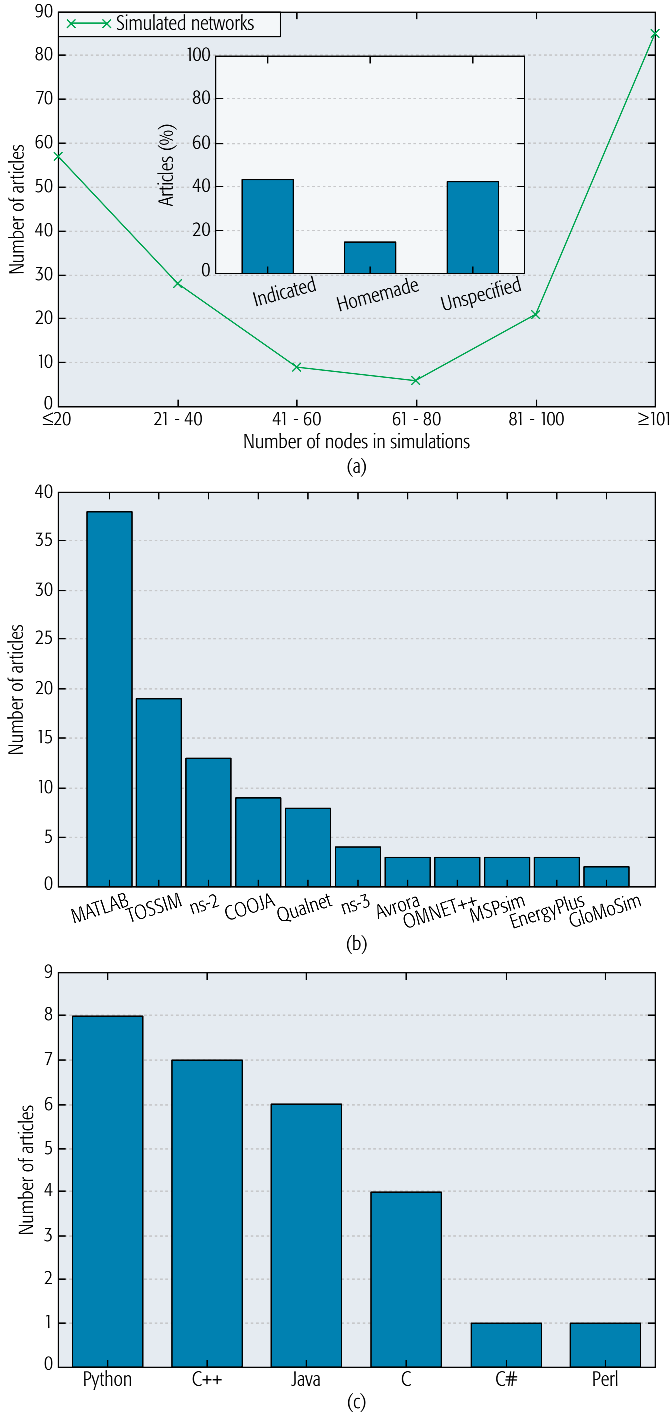

Figure 3. Use of mathematics (M), simulations (S), experiments (E), and their combinations in validation procedures of 596 ad hoc and WSN related articles: a) simulator usage and scales of simulated networks; b) popularity of simulators; c) programming language popularity for custom simulators.

[2] http://tinyos.stanford.edu/tinyos-wiki/index.php/TOSSIM

[3] http://www.contiki-os.org

deploy source code (which could be the same as the that of the simulator) on a sensor node and to flash it with no delay. Those open platforms thus allow for more rigorous, transparent, and replicable testing of proposed protocols and models.

Researchers, by connecting remotely (e.g., via ssh) to one open platform, may set up and initiate an experiment by using the terminal. Hence, the previously reported simulators along with open testbeds allow the research community to get a flavor of real deployments while maintaining a unique programming code. More importantly, obtaining performance evaluation measurements over a large-scale network (both for simulation and experiments) can be at no cost at all.

Finally, after following all the previously presented steps and by obtaining coherent results, researchers may consider initiating a real deployment by utilizing their verified and refined protocol.

Applications

While studying the 674 papers, we observed that the vast majority of papers actually mention some classical envisioned applications (e.g., defense, environment monitoring) but then focus on networking solutions that are application-independent. Regarding the type and nature of the problems that were addressed, we collected data about the correspondence of the studied articles to the layers of the open systems interconnection (OSI) model. We also identified papers that took into account some cross-layer design methodology.

As observed in Fig. 5, the most common approaches were at the application layer and with cross-layer design. While papers related to the former investigated new kinds of information that could be collected by ad hoc and WSNs, contributions related to the latter were concerned with the high constraints imposed by low-cost sensor and mobile devices that impose consideration of cross-layer approaches.

Mobility

Mobility is a key aspect of future designs. While the majority of existing and used simulators allow use and creation of mobility models, testing and executing such scenarios during an experimentation procedure require the involvement and combination of advanced and intelligent technologies such as robots. Consequently, very few of the widely popular open platforms support mobility [11]. Actually, there are a number of challenges that need to be addressed having mobile robots in a testbed: charging, remote administration, and maintenance of the robots. Indeed, robots must be able to reach their docking stations automatically. Conversely, remote users must be able to interact with robots over reliable links (e.g., WiFi). Even though these challenges can be addressed, testbed administrators then face the issue of localizing mobile devices in order to allow for repeatable trajectories. Indoor deployments cannot rely on GPS solutions and thus impose the computation of distance approximations being based on other available inputs (e.g., received signal strength intensity) or costly technologies (e.g., 3D camera with range detector sensors for the mapping of the environment). Furthermore, even with perfect localization of all robots, trajectories would be very difficult to replay, especially due to the odometry drift. Some 3D cameras using range detector sensors aim to handle this drift. Still, they lack the ability to compute the path where not enough landmarks exist in open space and large-scale environments.

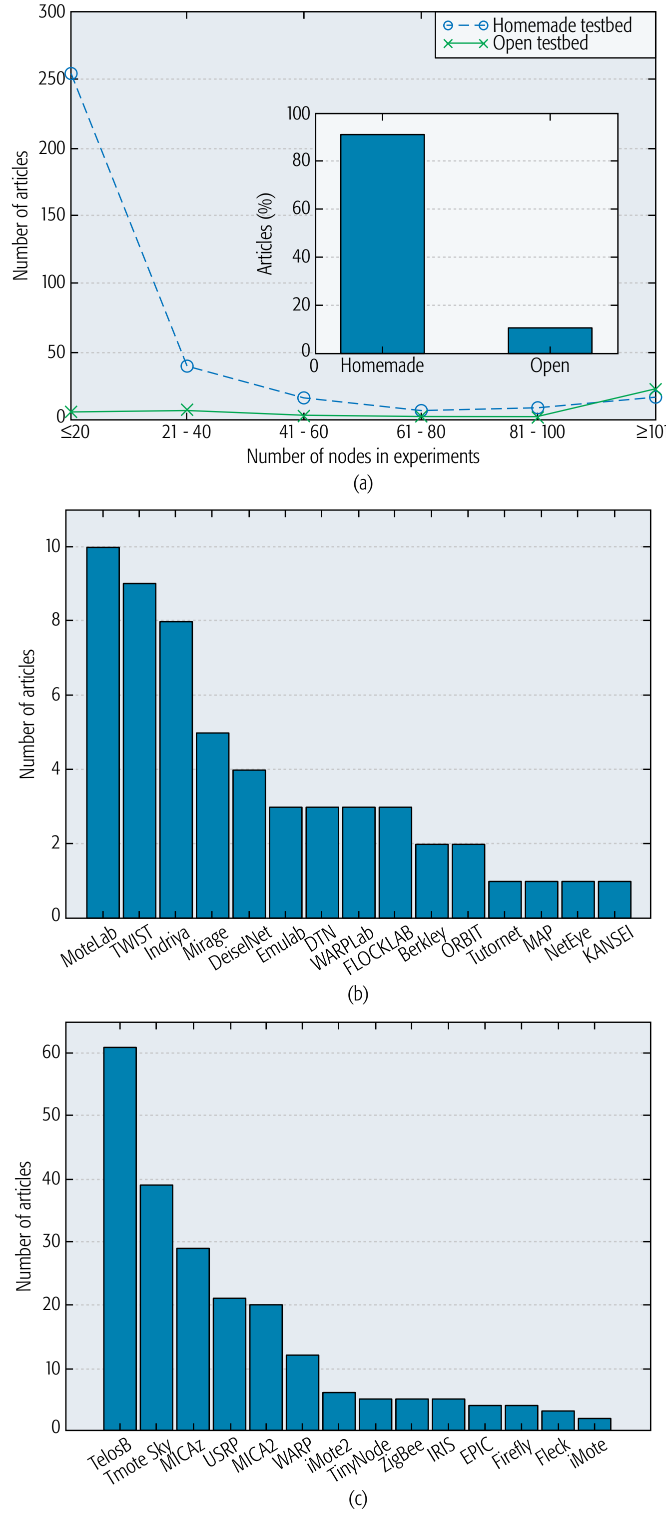

Figure 4. Total simulation vs. experimentation evaluated articles: a) testbed usage and scales of experimented networks; b) popularity of open testbeds; c) motes popularity.

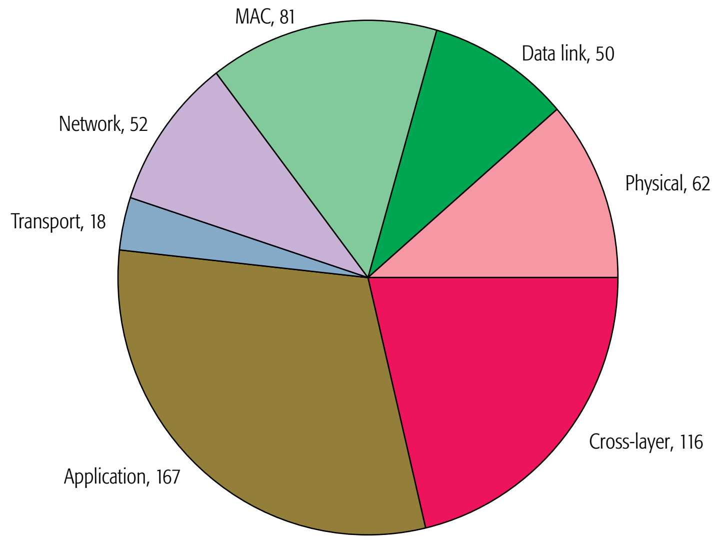

Figure 5. Published articles in ACM/IEEE IPSN, ACM MobiCom, ACM MobiHoc, and ACM SenSys from 2008 to 2013.

Conclusions

In this article, we review 674 papers that were published in four major and representative conferences in ad hoc and wireless sensor networks over the 2008–2013 period. We especially focus on the performance evaluation procedures in order to raise the question of whether simulations and experiments lead to scientific results or proofs of concepts. It is undeniable that simulators make the whole process of validation easier, faster, and less expensive. On the other hand, with the growing development of open and realistic testbeds, researchers may overcome the technical challenges and economical barriers of real-world deployment to perform a thorough experimental evaluation of their ideas in wide-scale platforms. Simulators and open testbeds are two crucial and complementary design and validation tools; ideally, the development process should start from theoretical analysis by providing bounds and indication of its performance, be validated and verified by simulations, and finally confirmed in open testbeds. Hence, once the entire procedure is successfully done and the performance results show coherence, researchers can push their solution to engineers in order to proceed with real deployment.

Simulation evaluations should allow for reproducible setups, thus producing scientific results that can be reproduced and verified by anyone in the community. In the context of experiments, our future work will focus on allowing researchers to get guidance for conducting experimentations over different testbeds in order to cover larger sets of assumptions. Finally, as far as the specific issues studied by our research community are concerned, they can be considered a different approach from the one we followed, and binding those findings with our study, will also be a straightforward extension to our current work.

References

[1] J. Yick, B. Mukherjee, and D. Ghosal, "Wireless Sensor Network Survey," *Computer Networks*, vol. 52, no. 12, 2008, pp. 2292–330.

[2] A. Gluhak *et al.*, "A Survey on Facilities for Experimental Internet of Things Research," *IEEE Commun. Mag.*, vol. 49, no. 11, Nov. 2011, pp. 58–67.

[3] G. Z. Papadopoulos *et al.*, "Adding Value to WSN Simulation Using the IoT-LAB Experimental Platform," *Proc. 9th IEEE Int'l. Conf. Wireless and Mobile Computing, Networking and Commun.*, 2013.

[4] D. Hiranandani, K. Obraczka, and J. Garcia-Luna-Aceves, "MANET Protocol Simulations Considered Harmful: The Case for Benchmarking," *IEEE Wireless Commun.*, vol. 20, no. 4, Aug. 2013, pp. 82–90.

[5] G. Barrenetxea *et al.*, "The Hitchhiker's Guide to Successful Wireless Sensor Network Deployments," *Proc. 6th ACM Conf. Embedded Networked Sensor Systems*, 2008.

[6] H. Kdouh *et al.*, "Performance Analysis of a Hierarchical Shipboard Wireless Sensor Network," *Proc. 23rd IEEE PIMRC*, 2012.

[7] I. Stojmenovic, "Simulations in Wireless Sensor and Ad Hoc Networks: Matching and Advancing Models, Metrics, and Solutions," *IEEE Commun. Mag.*, vol. 46, no. 12, Dec. 2008, pp. 102–07.

[8] K. Langendoen, A. Baggio, and O. Visser, "Murphy Loves Potatoes: Experiences from a Pilot Sensor Network Deployment in Precision Agriculture," *Proc. 20th IEEE Int'l. Symp. Parallel and Distrib. Processing*, 2006, pp. 174–74.

[9] C. Boano *et al.*, "The Impact of Temperature on Outdoor Industrial Sensornet Applications," *IEEE Trans. Industrial Info.*, vol. 6, no. 3, 2010, pp. 451–59.

[10] S. Kurkowski, T. Camp, and M. Colagrosso, "MANET Simulation Studies: The Incredibles," *ACM Mobile Comp. and Commun. Rev.*, vol. 9, no. 4, 2005, pp. 50–61.

[11] A.-S. Tonneau, N. Mitton, and J. Vandaele, "A Survey on (Mobile) Wireless Sensor Network Experimentation Testbeds," *Proc. IEEE Int'l. Conf. Distrib. Computing in Sensor Sys.*, 2014.

Biographies

Georgios Z. Papadopoulos (gpapadopoulos@unistra.fr) is a research assistant at the ICube Laboratory at the University of Strasbourg, France. He received his Ph.D. degree in computer science from the University of Strasbourg in 2015, his M.Sc. in telematics engineering from University Carlos III of Madrid, Spain, in 2012, and his B.Sc. in informatics from Alexander Technological Educational Institute (TEI). of Thessaloniki, Greece, in 2011. His research interests primarily span the area of wireless sensor networks and the Internet of Things. He is mainly working on the design and development of energy-efficient and low-delay algorithms and protocols for the MAC and routing layers under mobility and dynamic traffic assumptions. He received a Best Paper Award at IFIP Med-Hoc-Net '14 and was nominated for Best Student Paper Award at IEEE Sensors '14. He has already participated in various national (e.g., FIT equipex) and international (e.g., FLAVIA, PHC EXPRESS) projects.

Kosmas Kritsis received his B.Sc in informatics from the Alexander TEI of Thessaloniki in 2015. During academic year 2011–2012 he actively participated in the Erasmus program as a student at the Computer Science Department of University Carlos III of Madrid. He participated for a two-month period, starting in September 2013, in the Erasmus Internship program as a student researcher at the ICube Laboratory, University of Strasbourg. In November 2013 he was employed for a six-month period as a CISCO network and telecommunication administrator at the Network Infrastructure Department of Aristotle University of Thessaloniki. During summer 2014 he joined the Cypriot National Guard and served as a soldier. Currently he is a post-graduate student at the Sound and Music Computing M.Sc. program of the Pompeu Fabra University of Barcelona His research interests focus on wireless communications and multimedia.

Antoine Gallais obtained a Ph.D. in computer science in 2007 from the University of Lille 1, France. He then joined the University of Strasbourg where he has been an associate professor since 2008. His research interests include sensor and mobile ad hoc and mesh networks, especially focused on medium access control and routing protocols. He has been involved in the organization of several international events (e.g., Program Co-Chair of ICST AdHocNets 2014 and 2015, local Co-Chair of IEEE Wimob '13, Publicity Co-Chair of WiSARN '11–Fall, Web Chair of ICST Adhocnets '11) and has served on the TPC of several national and international conferences (e.g., IEEE GIIS 2013 and 2014, Adhocnets '11–15). He has already participated in various national (e.g., CNRS RECAP, ARC Inria IRAMUS, and currently FIT equipex and ANR IRIS) and international (e.g., CNRS-WIDE collaboration, PHC EXPRESS) projects.

Periklis Chatzimisios (SMIEEE) serves as an associate professor at Alexander TEI of Thessaloniki. Recently, he has been a visiting academic/researcher at the University of Toronto and Massachusetts Institute of Technology. He is involved in several standardization activities serving as a member of the Standards Development Board for IEEE ComSoc (2010–present) and Secretary of the IEEE 1907.1 Working Group. He holds Editorial Board positions for several IEEE/non-IEEE journals, and he is the author/editor of 8 books and more than 100 peer-reviewed papers with more than 1500 citations. He received his Ph.D. from Bournemouth University, United Kingdom, in 2005, and his B.Sc. from Alexander TEI of Thessaloniki in 2000.

Thomas Noël is a professor at the University of Strasbourg, France. His research activities include several aspects of wireless communications networks and telecommunications systems. He is particularly interested in network mobility, self-organized mobile networks, mobile network architecture and protocols, wireless sensor networks, ubiquitous computing, and multicast and group communications.

From Radio Telemetry to Ultra-Low-Power Sensor Networks: Tracking Bats in the Wild

Falko Dressler, Simon Ripperger, Martin Hierold, Thorsten Nowak, Christopher Eibel, Björn Cassens, Frieder Mayer, Klaus Meyer-Wegener, and Alexander Kölpin

Abstract

Sensor networks have successfully been used for wildlife monitoring and tracking of different species. When it comes to small animals such as smaller birds, mammals, or even insects, the current approach is to use extremely lightweight RF tags located using radio telemetry. A new quantum leap in technology is needed to overcome this limitation and enable new ways to observe larger numbers of small animals. In an interdisciplinary team, we are working on the different aspects of such a new technology. In particular, we report on our findings on a sensor-network-based tracking solution for bats. Our system is based on integrated localization and wireless communication protocols for ultra-low-power systems. This requires coding techniques for improved reliability as well as ranging solutions for tracking hunting bats. We address the technological and methodical problems related to system design, software support, and protocol design. First field experiments have been conducted that showcase the capabilities of our system.

Introduction

Wildlife tracking was one of the early applications in wireless sensor networks (WSNs) and remained, among a few others, one of the most successful [1, 2]. Sensor-networking-based wildlife monitoring provides more sophisticated methods for biologists to study individuals of a specific species in terms of gathering a huge amount of data by long-term observations. When it comes to tracking smaller animals like very small mammals or birds, radio telemetry is still considered state of the art [3]. This means that for localizing a single individual, at least two radio receivers operated by biologists in the field are needed to obtain a single position sample per triangulation. Here, both the rate of localization samples is very low — too low for continuous tracking — and the observation is usually restricted to very few individuals.

In the scope of the BATS[1] project, we are developing a new sensor-network-based system for monitoring group dynamics of bats in their natural habitat. In particular, we go one step further than related activities and investigate potentials of ultra-low-power sensor systems carried by the bats to monitor contacts or encounters between individuals and to track their routes at high spatial and temporal resolution. Tracking bats is especially challenging because direct observations on flying bats are almost impossible due to their nocturnal activity and high mobility. Mouse-eared bats (*Myotis myotis*), one of the most protected species in the European Union, are the main study target. The key challenge is that the animals, with an average body weight of about 20 g, can carry sensors of at most 2 g (including a 1 g battery), which is even less weight than a sheet of paper in A5 format. Comparable sensors published in the literature typically weigh more than 100 g, with Encounternet tags being outstanding with a weight of only 10 g — which is one order of magnitude higher compared to our requirements [2]. This poses a series of fundamental research questions that need to be addressed from a multi-disciplinary perspective ranging from hardware design, system support, and communication protocol engineering.

Our bat tracking scenario involves a variety of functions that need to be integrated with a main focus on energy efficiency. In a first stage, we aim at tracking flying bats using a stationary ground network. The concept is outlined in Fig. 1. The ground stations localize and track the bats using a combination of received signal strength (RSS) and phase-based localization techniques. Given the energy constraints of the mobile nodes, special modulation schemes are needed. When talking about the mobile nodes, we need to emphasize that a 1 g battery cannot directly power the microcontroller and the radio transceiver for a longer time period — in particular, we have to charge a capacitor first to provide peak operating current for the system. We investigated a combination of wake-up receiver (WuRx) and carefully controlled duty cycling to recharge the capacitor.

In a second stage, the system is also used to

Sensor networks have successfully been used for wildlife monitoring and tracking of different species. When it comes to small animals such as smaller birds, mammals, or even insects, the current approach is to use extremely lightweight RF tags located using radio telemetry. A new quantum leap in technology is needed to overcome this limitation and enable new ways to observe larger numbers of small animals.

Falko Dressler is with the University of Paderborn; Simon Ripperger and Frieder Mayer are with the Museum für Naturkunde Berlin, Leibniz Institute for Evolution and Biodiversity Science; Martin Hierold, Thorsten Nowak, Christopher Eibel, Klaus Meyer-Wegener, and Alexander Kölpin are with the Universität Erlangen-Nürnberg; and Björn Cassens is with TU Braunschweig.

[1] Dynamically adaptive applications for bat localization using embedded communicating sensor systems, http://www.for-bats.org/

 0163-6804/16/$25.00

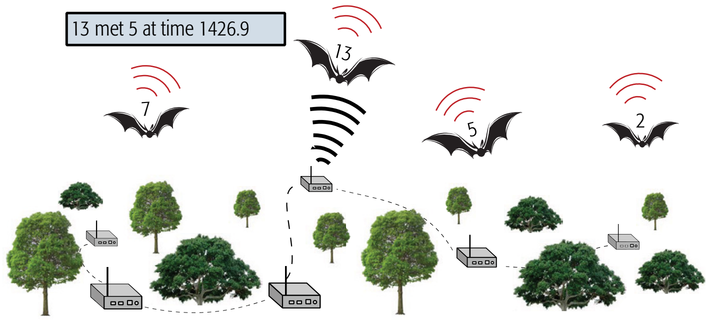

Figure 1. Conceptual scenario for bat tracking in the wild: stationary ground nodes are used to localize and track bats based on emitted beacon messages. The mobile nodes are also used to monitor encounters and download this data to the ground network to increase the observation range beyond the ground station area.

exchange contact information among the mobile nodes to increase the observation range beyond the ground station area. The encounter information is then downloaded to the ground network when in communication range, as determined by the WuRx. In the following, we refer to this data as *chunks* representing sets of contact information collected by a mobile node on the bat. The download follows the same power constraints and needs to be integrated with the localization signal. As the channel quality varies quickly in the given environment, additional error control mechanisms need to be integrated. We explored the capabilities of erasure codes (ECs) and were able to show that substantial improvements are possible.

Right now, the project is going beyond initial fundamental research on all the aspects mentioned. First field experiments have been completed in Germany as well as Panama showcasing the capabilities of the envisioned architecture. In this article, we summarize all the related inherently inter-disciplinary research challenges and outline conceptual approaches solving the problems, toward a novel ultra-low-power sensor networking solution for tracking bats in the wild.

Our main contributions can be summarized as follows:
- We report on first field tests using our 2 g sensor platform enabling, for the first time, tracking of animals weighing as few as 20 g;
- We developed a novel system architecture supporting precise localization and tracking;
- We summarize system design aspects from the operating system to wake-up receiver design;
- We finally outline the novel integrated wireless communication and localization protocols from PHY to forward error control using erasure codes;
- We briefly report about first successful outdoor experiments in Germany and Panama.

Related Work: Studying Bats in the Wild

The first projects relied on typical sensor platforms as used in academic research labs (e.g., the Great Duck Island project) or on special hardware that is even robust enough to be carried by larger animals (e.g., the ZebraNet project). More recently, tracking has become a major application besides the collection of several sensor readings. Also, technological advances have enabled new generations of sensor nodes that can be used to track much smaller animals such as the Iberian lynx [1]. Wireless digital transceiver technology has even rendered the automated mapping of social networks in wild birds possible (e.g., in the Encounternet project) [2]. From these successful approaches to wildlife monitoring using sensor networks, we learned about hardware design issues, network management, and data collection techniques.

The state-of-the-art technology for bat tracking is still radio telemetry. However, this method requires high labor costs since two or more persons must manually observe one or a few individuals at a time. The reward for this great effort is a minimal number of animal positions that are separated by several minutes and usually contain localization errors of tens of meters.

Observing the movements of individual animals in their natural habitat is one of the most difficult tasks in the field of behavioral biology; however, it is key to understanding complex biological processes such as foraging, social interactions, migration, and gene flow. Recent technological advances in satellite-based animal localization and automatized data acquisition are restricted to medium-sized to large mammals and birds due to the considerable weight of available transmitters [4].

The most promising approach for tracking large-scale movements of small animals from space is represented by the ICARUS initiative. It is expected to start in 2016, and tags will initially weigh 5 g but should become considerably lighter within a couple of years [5]. Until ICARUS finally launches and tags undergo further miniaturization, traditional radio telemetry still represents the state of the art of bat tracking as radio transmitters are available with a weight down to 0.2 g. However, this technique can only provide a rough estimate of foraging movements based on animal positions that are separated by several minutes and contain localization errors of at least tens of meters. However, the number of individuals as well as the observable area are limited.

The study of social interactions among individually identifiable bats is especially challenging in the wild. To date, the only option to automate this is to use extremely light passive integrated transponder tags (PIT-tags), which can be identified within known roosting sites that are equipped with PIT-tag readers [6]. The only possibility to document group dynamics while foraging is again radio telemetry [7].

The BATS sensor network aims at implementing both automated high-precision positioning of many individuals at a time and documenting interactions of the observed bats. The high temporal and spatial resolution of data will render the reconstruction of individual flight trajectories possible. Communication among mobile nodes will shed light on interactions among bats at the individual level during the nightly activity phase, which were impossible to study until now. The advances of the BATS system hold the potential to gain deeper understanding of bat behavior, for example, habitat use, analysis of flight maneuvers, and group dynamics.

The BATS Ground Network

Experiment Management

Many projects on wildlife tracking mentioned in the introduction have used a data-stream management system (DSMS) to collect the data. This has the advantage of processing the data early (i.e., selecting and aggregating them), which has proven to reduce power requirements for WSNs. This in general is favorable, as battery lifespans increase and maintenance costs are reduced.

Queries are deployed to the ground network and later even to the mobile nodes on the bats to define the early processing. The stream operators invoked by the queries can be adjusted or even replaced if the biologists want to tweak resolution to support their experiments. It is still a research issue to conduct dynamic stream operator replacement efficiently in sensor networks.

Due to technical restrictions, especially the communication data rate between mobile nodes and ground nodes, not all acquired data can be sent. For example, providing a timestamp and duration of a meeting, which is mandatory for further analysis, uses almost 70 percent of the available data rate. Other data like different received signal strength indicator (RSSI) values cannot be sent without any violation against data rate requirements. Therefore, only a subset of acquired data can be sent to the base station network. Thus, a decision has to be made to select the most suitable data to send. Depending on the current research focus, the system will be adapted to increase the quality of the relevant data while balancing the other goals: providing additional data for later offline analysis and increasing node lifespans.

Ground Nodes

The ground nodes consist of MicroZed boards[2] equipped with a custom software defined radio (SDR) RF frontend based on the Analog Devices AD9361 transceiver chip as we assume no tight energy constraints for the ground nodes. This versatile SDR platform allows exploitation of the localization methods described above. Comprising a field programmable gate array (FPGA) and a dual ARM core, the ground nodes feature enough processing power to enable complex range estimation and direction finding algorithms.

The ground nodes also have to wake up the mobile nodes in order to initiate and coordinate transmissions when the mobile nodes are in range of the network. Periodic beacon signals allow the mobile nodes to be loosely synchronized with the ground stations. Using these beacon signals, mobile nodes can decide whether to transmit localization signals or to save energy.

Ground Network

In order to support both localization as well as data collection from the mobile nodes, an efficient and decentralized distributed data storage and lookup is needed. A promising concept is a distributed hash table (DHT) integrated with ground network routing capabilities. There are various protocols for sensor networks providing standard DHT functionalities [8]; however, most of them rely on globally valid topology information, need geographic location information, or do not take into consideration the physical position of nodes, which leads to increased routing paths. We selected the Virtual Cord Protocol (VCP) [8], which overcomes many of these shortcomings. VCP supports routing in the ground network topology using a virtual cord based on neighborhood information. For data management, each node is then responsible for data with a hash value that matches its virtual node identifier.

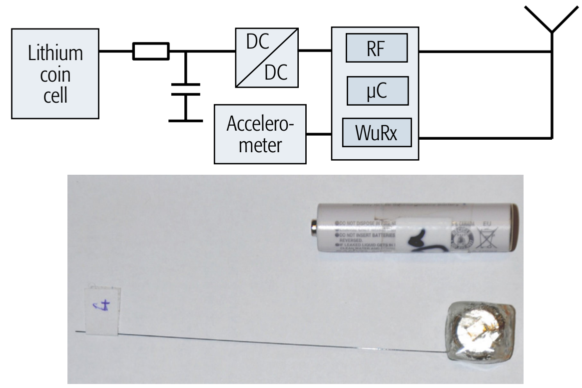

Figure 2. Hardware architecture and assembly of the mobile node.

Mobile Node Design

Hardware Design

The building blocks of the mobile node and a successfully used prototype are depicted in Fig. 2. Due to its high energy density, a lithium primary coin cell battery is used to power the tag. Caused by the maximum current of lithium cell supply capability, a buffer capacitance is applied. A DC/DC switching converter down converts the variable capacitor voltage to a constant system-on-chip input voltage of 1.8 V. The system on chip (SoC), which is the key component of the tag, comprises a microcontroller, a dual-band front-end for transmission and reception in the 868 MHz and 2.45 GHz bands, and a WuRx operating in the 868 MHz band. The first prototypes containing a Cortex-M0+, an Si4460, and an AT86RF233 have been built to set up a system demonstrator. Besides the SoC, an accelerometer is placed on the tag to facilitate motion detection. A dual-band antenna is shared by the regular transceiver and the WuRx. The whole hardware assembly is protected against physical influences (e.g., humidity or the attempt of the bat to scratch it off) by an epoxy sealing.

The WuRx must be suitable for two different operating conditions: the communication between mobile nodes and the communication between ground node and mobile node. Defined by the spacing of ground nodes, the maximum distance for ground node to mobile node communication is approximately 50 m. According to the communication channel, this corresponds to an attenuation of 65 dB (free space path loss)

[2] http://www.zedboard.org/

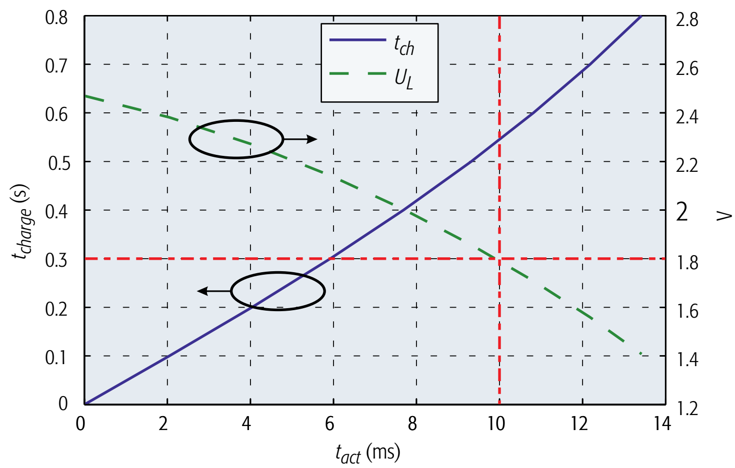

Figure 3. Dependency of the active period and the recharging time.

to 78 dB (free space plus linear fading with 0.25 dB/m) at 868 MHz. Given a transmission power of 10 dBm, this leads to a minimum required sensitivity of –55 dBm and –68 dBm, respectively. We rely on the concept described in [9], which is suitable for both operating conditions.

Power Management

The lithium primary battery offers high capacitance, but still the current that can be drawn is limited to ~0.5 mA. However, the SoC, especially the transceiver, demands several milli-amperes when active. To satisfy this demand a buffer capacitance is integrated, which has to be charged by the battery before it is discharged by the SoC. Hence, continuous operation of the transceiver is not feasible. Furthermore, a trade-off between active period and recharging time exists: the longer the transceiver is active, the longer the capacitor has to be recharged. With the assumption that the current drain during the sleeping phase of the tag and the recharging during the active period is negligible, the dependency between the maximum active period t_{act} and the minimum recharging time t_{ch} is visualized in Fig. 3. The fact that the capacitor's voltage decreases during its discharge leads to the existence of an absolute maximum time limit of the active period; that is, when the voltage of the capacitor at the end of the active period U_L falls below the 1.8 V input voltage of the SoC. The value of U_L is also given in Fig. 3 and determines the limit of the active period to be about 10 ms The resulting protocol design challenges are discussed in the following section.

System Support

Hardware limitations inevitably have an effect on the resources that are available to the software components running on top. In the BATS project, this limitation is even more difficult to manage as we have to assume a deeply embedded system with a microcontroller and available memory that are considerably more limited than, for example, common embedded devices such as smart phones or lab-style sensor nodes. This demands a common software infrastructure to efficiently use underlying hardware components and render resource-efficient applications possible. However, these applications still need to meet certain other criteria apart from pure energy efficiency (e.g., complying to certain deadlines such as soft and hard real time). For tiny embedded systems, a multitude of different operating systems exist, each possibly having divergent approaches since they are geared toward different use cases. However, most of these operating systems still aim to be general purpose, which results in a certain degree of potential efficiency loss. Alternatives are real-time operating systems such as sloth [10]. This particular system allows full use to be made of the available hardware support, which leads to smaller program code and space that the operating system allocates in memory.

To address the requirement of having an operating system and applications that also need to be optimized regarding energy efficiency, we rely on both the SEEP energy estimation framework developed by Hönig *et al.* [11] and energy measurements from our self-developed measurement device. Furthermore, this gives us the possibility to write predictable code in terms of its expected lifetime on the actual hardware. By following the framework's profile-driven approach, we have created an energy profile for our target platform, and therefore are able to retrieve estimated expected energy consumption values for a sloth application under test (e.g., with CPU-intensive code such as erasure encoding processing, described below). This enables us to proactively develop program code and make energy optimizations before the actual system is deployed and field tests are conducted.

We further extended sloth with support of dynamic program code reconfiguration; that is, the possibility to efficiently switch between different code parts that already reside on the platform at runtime. With this mechanism, we can react to either a query sent by the front-end user (i.e., a biologist) or predictions automatically derived from ongoing data stream analyses. By way of example, we are able to alter the bit fields of the meeting, which consist of a meeting's duration, starting time, and RSSI values.

Communication Protocols

Protocol Design

When it comes to the design of low-energy communication protocols in sensor networks, three main approaches have been identified in the literature:

- Duty cycling, that is, periodically switching between active and passive state to power off the main components in the passive stage, with synchronization explicitly required
- Low-power listening, that is, "waking up" the receiver node using multiple transmission attempts (either full messages or wake-up preambles) to dismiss the synchronization requirement
- Wake-up medium access control (MAC) protocols, that is, using dedicated hardware to wake up the node in case of an upcoming transmission (e.g., PW-MAC [12])

Given the energy constraints discussed, the communication protocol for mobile to ground communication needs to be designed in a completely novel way. First of all, duty cycling is an inherent feature given the recharging cycles of the capacitor that powers the radio transceiver. Furthermore, unnecessary transmissions need to be prevented when the bat is not in communication range of at least one ground node. Here, a multi-stage WuRx is used to completely power off both the radio transceiver and the microcontroller.

These two concepts can be combined to benefit from both advantages. Duty cycling helps reduce the energy consumption (and supports recharging the capacitor) when the bat is in range of a ground node, and the multi-stage WuRx triggers initiation of this duty cycling and turns off all digital components if not needed. The entire cycle is controlled by the ground network (i.e., all ground nodes are assumed to be synchronized). We assume a frequency of wake-up pulses of up to 10 Hz for trajectory estimation.

In each cycle, a wake-up pulse disseminated by the ground nodes wakes up the mobile nodes. In a second phase of the project, we aim to even encode data on the wake-up signal. This information can be used to coordinate channel access in order to avoid collisions if multiple bats are within range of a ground node.

Obviously, more than a single mobile node needs to be supported by the mobile to ground communication protocol and track multiple individuals at the same time. Timing is controlled by the base station. We picked this approach to increase the robustness of the protocol even though (without an uplink signal) this limits the number of mobile nodes to the number of available time slots. Furthermore, guard intervals have been introduced because the sensor nodes are not synchronized perfectly and the oscillators might drift considerably.

Encoding and Modulation for Combined Localization and Data Communication

As energy awareness is one of the most crucial aspects of our system, the energy spent on RF activity for localization and communication has to be minimized. We recently proposed a signaling scheme that combines localization and data signals [13]. In the presented approach, binary offset carrier (BOC) modulation is used to simultaneously transmit data and provide accurate range measurements.

BOC modulation is well known in the field of global navigation satellite systems (GNSSs). In contrast to GNSS and due to the limited energy, very short burst signals are used for communication and localization instead of continuous signals. A further motivation for burst signals is to avoid near-far effects in local real-time locating systems (RTLSs) by time-division multiplexing.

Due to the limited observation area, pure subcarrier tracking is applied in the BATS system as this approach maximizes root mean square (RMS) bandwidth and thus leads to a minimum range estimation variance. Data transmission is realized by modulation of the subcarriers. This modulation broadens the subcarriers and also decreases the RMS bandwidth, but still has only a negligible impact on the range estimation accuracy. However, data decoding errors have a rather substantial influence on the distance estimation as they result in a mismatch of the correlated sequences, which then leads to signal-to-noise ratio (SNR) degradation. This significant increase in the range estimation variance is shown in Fig. 4.

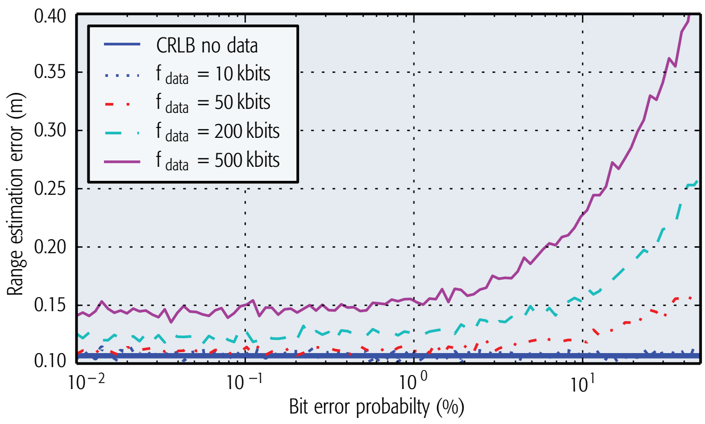

Figure 4. Impact of data transmissions on the maximum achievable range estimation performance.

Improving Communication Reliability

The channel quality may vary quickly due to the continuous movements of bats and the heterogeneous forest environment; thus, the communication is in general assumed to be highly unreliable, and error control techniques (ECs) have to be applied. We consider ECs as a specific class of forward error correction codes and a promising approach in our BATS scenario. ECs are widely employed to improve the reliability in wireless transmissions [14]. Compared to the simplistic approach of sending chunk replica together with the original data as well as to the classic automatic repeat request (ARQ) mechanism, ECs offer better performance with reduced costs in terms of energy consumption. Likewise, ECs show better efficiency than on demand chunk retransmissions realized by acknowledging successfully transmitted chunks.

The significant difference between the various ECs is the mathematical background of the encoding and decoding algorithms. Reed-Solomon (RS) codes such as Cauchy and Vandermonde share the same algorithms; however, they work on different kinds of matrices, whereas codes like Tornado vary significantly in the algorithm itself. We investigated the mentioned codes for their applicability in our scenario [15]. These codes support different code rates that essentially define the possible error correction vs. the overhead for additional coding data. To the best of our knowledge, there is no study on the feasibility of ECs for scenarios with spontaneous connectivity such as the scenario we are investigating with its specific channel properties.

The usage of ECs and replicated sending inevitably increases energy consumption. Primarily, the sending of redundant chunks drains energy; however, in the former case the execution of the encoding algorithm has to be taken into consideration as well. This trade-off between

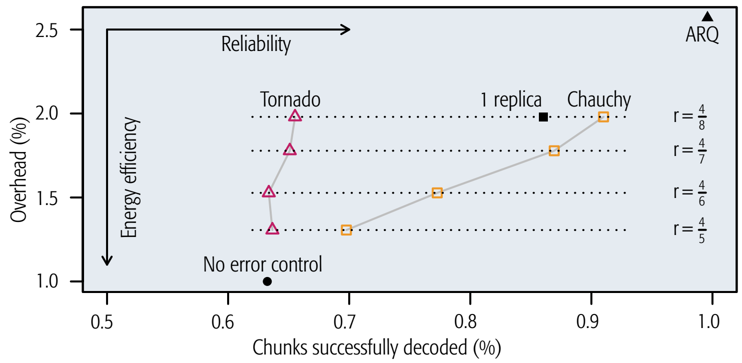

Figure 5. Reliability gain achieved by of the different error control strategies and code rates r vs. their energy efficiency in terms of the necessary number of packet transmissions.

improved reliability and the overhead caused by redundant chunks is outlined in Fig. 5. The graph illustrates the gain in reliability for the different error control techniques in comparison to the energy efficiency. The plotted results have been collected in a series of simulations based on the discussed mobile to ground communication protocol and assuming a typical packet error rate of about 20 percent in addition to the used two-ray path loss model to resemble multipath fading effects in the simulation.

As we move from left to right in the graphs, reliability measured against the amount of recovered data increases, whereas moving from bottom to top the energy efficiency decreases with increasing overhead. For reference, the non-replicated sending is also indicated, obviously not inducing any overhead but at the cost of very low reliability.

As we can see, ARQ as well as Tornado-based ECs either significantly increase the overhead or lead to only marginal reliability improvements. However, combining the wireless communication with a chunk-based RS code, we observe substantial improvements at acceptable energy costs. This especially holds for code rates of $r = (4/7)$ to $r = (4/8)$.

First Experiments on Wild Bats

The basic functionality of the BATS concept has been validated on the target species *Myotis myotis*, the greater mouse-eared bat, and on the tropical fringe-lipped bat, *Trachops cirrhosus*, in Panama. In order to demonstrate the technical feasibility of building an energy-efficient proximity sensor node of less than 2 g with a theoretical battery life of at least one week, we performed a field test on four individual bats in a maternity colony of mouse-eared bats in Upper Franconia. The presence and absence of the tagged individuals in the colony has been documented, and interactions among the tagged individuals have been surveyed. Communication of mobile nodes with the base station served as an indicator of presence in the colony, while RSSI measurements were used to estimate the distance between two bats. During a second field experiment conducted in Gamboa, Panama, we successfully documented encounters among members of a social group of the fringe-lipped bat outside the roost while hunting. Furthermore, we tracked foraging movements of individual bats in a small area of about 20 m × 25 m based on field strength measurements. Tagged animals of the focus species *Myotis myotis* and *Trachops cirrhosus* are shown in Fig. 6.

Conclusion and Future Work

We have reported on our findings toward a new era of ultra-low-power sensor systems used for tracking bats in the wild. Even though the BATS sensor network has been designed to observe bats (i.e., small animals that are moving in three dimensions at high speed), it will also be applicable to a wide range of vertebrates, including mammals, birds, and reptiles, and even certain invertebrates (e.g., large beetles). In our first field tests, we succeeded in collecting contact information of bats in their natural environment and documented foraging movements. These very promising early results encourage further investigations and research in this inherently interdisciplinary project. Many of the scientific findings can be adapted to other application fields — we believe that our technical solutions will substantially impact research on ultra-low-power sensor networks in general.

Certainly, there are still many open research questions and big challenges to be addressed. This particularly includes the option to provide even more reliable wireless communication without increasing the energy budget. Such

Figure 6. Individuals of the focus species *Myotis myotis* (left) and *Trachops cirrhosus* (right) carrying prototypes of mobile nodes that are attached between the shoulder blades with surgical cement.

functionality is needed, for example, for online reconfiguration and even software updates of the mobile nodes. We are also thinking about integrating sensors to combine physiological and environmental data with tracking data. The applicability to a wide species spectrum across taxa (not only bats) may even be increased by further miniaturization.

Acknowledgments

This work has been supported in part by the German Research Foundation (DFG) under grant no. FOR 1508.

This project is inherently multi-disciplinary, and many faculty members and students have been involved in making the presented system possible. We acknowledge support from the following people: Bastian Bloessl, Markus Hartmann, Sebastian Herbst, Rüdiger Kapitza, Margit Mutschlechner, Niels Hadaschik, Lucila Patino-Studencka, Wolfgang Schröder-Preikschat, Jörn Thielecke, Thomas Ussmüller, and Robert Weigel.

References

[1] A.-J. Garcia-Sanchez *et al.*, "Wireless Sensor Network Deployment for Monitoring Wildlife Passages," *Sensors*, vol. 10, no. 8, 2010, pp. 7236–62.

[2] C. Rutz *et al.*, "Automated Mapping of Social Networks in Wild Birds," *Current Biology*, vol. 22, no. 17, 2012, pp. R669–71.

[3] S. K. Amelon *et al.*, "Ecological and Behavioral Methods for the Study of Bats," T. H. Kunz and S. Parson, Eds. Johns Hopkins Univ. Press, Oct. 2009, pp. 57–71.

[4] M. Hebblewhite and D. T. Haydon, "Distinguishing Technology from Biology: A Critical Review of the use of GPS Telemetry Data in Ecology," *Philosophical Trans. Royal Soc. B: Biological Sciences*, vol. 365, no. 1550, June 2010, pp. 2303–12.

[5] E. Pennisi, "Global Tracking of Small Animals Gains Momentum," *Science*, vol. 334, no. 6059, Nov. 2011, p. 1042.

[6] G. Kerth, N. Perony, and F. Schweitzer, "Bats Are Able to Maintain Long-Term Social Relationships Despite the High Fission-Fusion Dynamics of Their Groups," *Proc. Royal Soc. London B: Biological Sciences*, vol. 278, no. 1719, Sept. 2011, pp. 2761–67.

[7] D. K. N. Dechmann *et al.*, "Group Hunting – A Reason for Sociality in Molossid Bats?" *PLoS ONE*, vol. 5, no. 2, Feb. 2010, p. e9012.

[8] A. Awad, R. German, and F. Dressler, "Exploiting Virtual Coordinates for Improved Routing Performance in Sensor Networks," *IEEE Trans. Mobile Computing*, vol. 10, no. 9, Sept. 2011, pp. 1214–26.

[9] N. Pletcher, S. Gambini, and J. Rabaey, "A 2GHz 52 mW Wake-Up Receiver with –72dBm Sensitivity Using Uncertain-IF Architecture," *IEEE J. Solid-State Circuits*, vol. 44, no. 1, Jan. 2009, pp. 269–80.

[10] W. Hofer *et al.*, "Sloth: Threads as Interrupts," *Proc. 30th IEEE Real-Time Sys. Symp.*, Washington, DC, Dec. 2009, pp. 204–13.

[11] T. Hönig, *et al.*, "SEEP: Exploiting Symbolic Execution for Energy-Aware Programming," *Proc. 23nd ACM Symp. Operating Systems Principles*, 4th Wksp. Power-Aware Computing and Systems, Cascais, Portugal, Oct. 2011, pp. 4:1–4:5.

[12] L. Tang *et al.*, "PW-MAC: An Energy-Efficient Predictive-Wakeup MAC Protocol for Wireless Sensor Networks," *Proc. IEEE INFOCOM 2011*, Shanghai, China, Apr. 2011, pp. 1305–13.

[13] T. Nowak *et al.*, "Combined Localization and Data Transmission in Energy-Constrained Wireless Sensor Networks," *Proc. IEEE Radio Wireless Week*, IEEE Topical Conf. Wireless Sensors and Sensor Networks, San Diego, CA, Jan. 2015.

[14] C. R. Berger *et al.*, "Optimizing Joint Erasure- and Error-Correction Coding for Wireless Packet Transmissions," *IEEE Trans. Wireless Commun.*, vol. 7, no. 11, Nov. 2008, pp. 4586–95.

[15] M. Mutschlechner, B. Li, R. Kapitza, and F. Dressler, "Using Erasure Codes to Overcome Reliability Issues in Energy-Constrained Sensor Networks," *Proc. 11th IEEE/IFIP Conf. Wireless On Demand Network Systems and Services*, Obergurgl, Austria, Apr. 2014, pp. 41–48.

Biographies

FALKO DRESSLER [SM] (dressler@ccs-labs.org) is a full professor of computer science and head of the Distributed Embedded Systems Group at the Department of Computer Science, University of Paderborn. He received his M.Sc. and Ph.D. degrees from the Department of Computer Science, University of Erlangen in 1998 and 2003, respectively. He is a Senior Member of ACM. His research activities are focused on adaptive wireless networking and self-organization methods with applications in wireless ad hoc and sensor networks, vehicular networks, and nano-networking.

SIMON RIPPERGER is a biologist and is working as a postdoctoral researcher at the Museum für Naturkunde Berlin, Leibniz Institute for Evolution and Biodiversity Science. He graduated from the University of Ulm in conservation genetics and movement ecology in neotropical bats. His present research focus is on ecology and behavior of tropical and temperate bat species.

MARTIN HIEROLD received his Diploma in electrical engineering in 2010 from the University of Erlangen-Nuremberg, Germany. Currently he is working toward his doctoral degree. After graduation he worked as a development engineer in a microwave test laboratory. Since 2013 he is with the Institute for Electronics Engineering, Friedrich-Alexander University of Erlangen-Nuremberg, Germany. His research interests are in the areas of wireless communication, local positioning, and low-power wireless sensor networks.

THORSTEN NOWAK received his Diploma in engineering from the University of Ulm, Germany, in 2009. In 2008, he joined the Fraunhofer Institute for Integrated Circuits, where he was involed in the develepoment of localization systems and sensor data fusion techniques. Since 2013, he has been with the Insitute of Information Technology, University Erlangen-Nuremberg, working toward his Ph.D. His focus is on multipath mitigation techniques, multisensor data fusion, localization systems, and RFID.

CHRISTOPHER EIBEL received his diploma degree in computer science from the Friedrich-Alexander-Universit√§t Erlangen-Nurernberg (FAU) in 2012. He is currently a Ph.D. candidate at the Department of Computer Science 4 (Distributed Systems and Operating Systems) at the same university. His research interests include energy proportionality and energy awareness in distributed systems.

BJÖRN CASSENS received his B.Sc. degree from the University of Applied Science of Osnabrueck in 2010 and his M.Sc. degree in 2013 from the Technical University of Braunschweig, both in electrical engineering. His current research interest is decreasing the energy demand of embedded systems by utilizing software and hardware induced energy-saving capabilities.

FRIEDER MAYER is an evolutionary biologist and curator at the Museum für Naturkunde Berlin, the Leibniz Institute for Evolution and Biodiversity Science. His research interests are in the fields of behavioral ecology, population dynamics, and speciation. Since his graduation in biology he has worked on various aspects of behavior, ecology, and evolution of bats.

KLAUS MEYER-WEGENER has been a full professor of computer science (data management) at the University of Erlangen and Nuremberg since October 2001. From 1975 to 1980, he studied computer science at the Darmstadt Institute of Technology and finished with the degree of Diplom-Informatiker. After that, he became a research assistant in the Department of Computer Science at the University of Kaiserslautern. In 1986, he received his Ph.D. From October 1987 to December 1988 he was granted a leave to work as an adjunct research professor at the Naval Postgraduate School in Monterey, California.

ALEXANDER KÖLPIN received his diploma in electrical engineering in 2005, his doctoral degree in 2010, and his Venia Legendi in 2015, all from the University of Erlangen-Nuremberg. Currently he works with the Institute for Electronics Engineering, University of Erlangen-Nuremberg as a group leader of his research group, Circuits, Systems and Hardware Test. His research interests are in the areas of microwave circuits and systems, local positioning, and wireless sensor networks.

There are still many open research questions and big challenges to be addressed. This particularly includes the option to provide even more reliable wireless communication without increasing the energy budget. Such functionality is needed, for example, for online reconfiguration and even software updates of the mobile nodes.

CheepSync: A Time Synchronization Service for Resource Constrained Bluetooth LE Advertisers

Sabarish Sridhar, Prasant Misra, Gurinder Singh Gill, and Jay Warrior

This article presents and describes *Cheep*Sync, a time synchronization service for BLE advertisers, especially tailored for applications requiring high time precision on resource constrained BLE platforms. Designed on top of the existing Bluetooth v4.0 standard, the *Cheep*Sync framework utilizes low-level timestamping and comprehensive error compensation mechanisms for overcoming uncertainties in message transmission, clock drift, and other system-specific constraints.

Abstract

Clock synchronization is highly desirable in distributed systems, including many applications in the Internet of Things and Humans. It improves the efficiency, modularity, and scalability of the system, and optimizes use of event triggers. For IoTH, BLE — a subset of the recent Bluetooth v4.0 stack — provides a low-power and loosely coupled mechanism for sensor data collection with ubiquitous units (e.g., smartphones and tablets) carried by humans. This fundamental design paradigm of BLE is enabled by a range of broadcast advertising modes. While its operational benefits are numerous, the lack of a common time reference in the broadcast mode of BLE has been a fundamental limitation. This article presents and describes *Cheep*Sync, a time synchronization service for BLE advertisers, especially tailored for applications requiring high time precision on resource constrained BLE platforms. Designed on top of the existing Bluetooth v4.0 standard, the *Cheep*Sync framework utilizes low-level timestamping and comprehensive error compensation mechanisms for overcoming uncertainties in message transmission, clock drift, and other system-specific constraints. *Cheep*Sync was implemented on custom designed nRF24Cheep beacon platforms (as broadcasters) and commercial off-the-shelf Android ported smartphones (as passive listeners). We demonstrate the efficacy of *Cheep*Sync by numerous empirical evaluations in a variety of experimental setups, and show that its average (single-hop) time synchronization accuracy is in the 10 μs range.

Introduction

A common time reference is an important requirement for distributed systems [1], and developing such a service on *constrained* devices is particularly challenging [2–10]. In the recent past, a new class of constrained platforms based on Bluetooth Low Energy (BLE)[1] have emerged. BLE is different from other wireless technologies because it combines a standardized communication technology designed for low-power systems, and a new sensor-based data collection framework. It also offers easy integration with most handheld devices (e.g., smartphones and tablets), something toward which traditional wireless sensor networks (WSNs) are still working. BLE has inherited several technical features from classic Bluetooth that provide for robust reliable connections. However, the most significant difference is its asymmetric design. While the communication foundation is based on a master-slave architecture, it offers a new feature in the form of an *advertisement* (configurable through the broadcast/peripheral mode of BLE). This new mode offers unidirectional correspondence between two or more LE devices using advertising events, thereby achieving a communication solution without entering into a bonded connection (as required by classic Bluetooth devices). Such a loosely coupled manner of data transfer is undoubtedly more energy efficient, but also unearths other limitations. For example, the broadcast mode of communication *does not* provision for a time synchronization service; even though this feature is included in the solution stack, it can only be availed upon pairing.

Challenges

BLE provides a range of broadcast advertising modes, of which the *most* energy efficient is the non-connectable undirected advertising mode (`ADV_NONCONN_IND`): a transmit-only broadcaster mode without any listen window. Establishing time synchronization in this mode of BLE operation is challenging due to many reasons. First, the traditional techniques of message passing among different elements is not supported by this network architecture (i.e., Bluetooth v4.0), and as a result, timing uncertainties cannot be compensated by exchanging time stamped packets, or "pings" between nodes. Second, devices receiving advertisement data (e.g., smartphones) have high functional asymmetry compared to the BLE broadcast units. They typically run on a multithreaded and multitasking operating system (OS) (e.g., Android) where the measure of system latencies and their associated uncertainties can be

[1] BLE Specification Adopted Documents. https://www.bluetooth.org/en-us/specification/adopted-specifications

Sabarish Sridhar is with M.S. Ramaiah Institute of Technology, Bangalore; Prasant Misra and Gurinder Singh are with the Robert Bosch Center for Cyber Physical Systems, Indian Institute of Science, Bangalore; Jay Warrior is with Mobiatrics LLC. Sabarish Sridhar was an intern and Jay Warrior was the chief technologist at the Robert Bosch Center for Cyber Physical Systems, Indian Institute of Science, Bangalore during the course of this work.

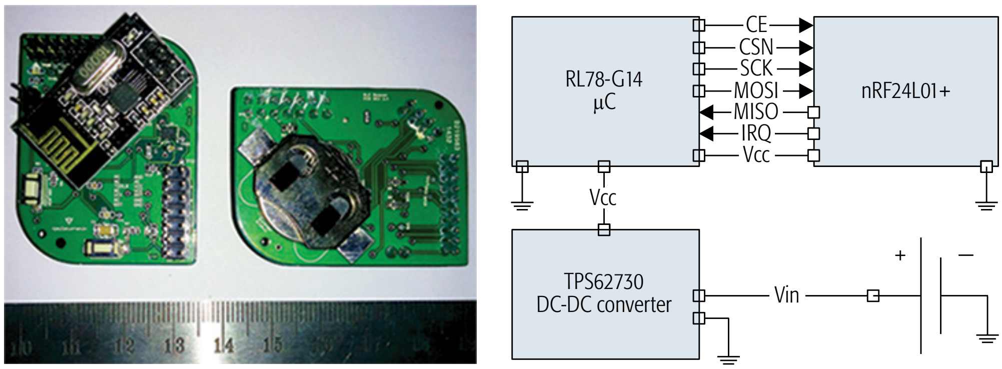

(a) (b)

Figure 1. nRF24*Cheep*: custom designed BLE Beacon platform. The major components include: *RL78-G14* microcontroller, *NRF24L01+ 2.4 GHz* RF transceiver, *TPS62730* synchronous step-down DC-DC converter, and *CR2032* battery holder. a) platform overview; b) functional representation of the major platform components.

The system is composed of two units: beacon and control. The beacon unit consists of resource constrained sensing tags that are deployed in the region of interest. The control unit consists of a resourceful gateway device capable of listening and receiving broadcast data contained in BLE advertisements.

many orders of magnitude higher than those at the transmitter end. Third, low-level timestamping on such multifunctional receiver devices can be performed to a certain limit and is subject to system restrictions of the underlying firmware. Therefore, motivated by the need to overcome the above limitations but be able to establish a common time reference across resourced constrained BLE devices operating in the `ADV_NONCONN_IND` mode, we propose *Cheep*Sync.

Contributions and Road Map

Due to the architectural constraints, the key ideas of *Cheep*Sync are:

- Not to synchronize the nodes in the network, but to make the devices that use these nodes to synchronize
- Piggyback on the device mobility aspect (as they are inevitably carried by people) and use them as "synchronization mules" for the "broadcast" system

Thus, it offers a flexible piggyback design wherein running the time service does not require data transactions to be temporarily suspended. Therefore, time synchronization with *Cheep*Sync is highly implicit rather than explicit. Since *Cheep*Sync rides on the BLE broadcast framework, it is scalable to the point that the framework has to offer.

In this article, we describe our experiences in building a custom BLE beacon platform that uses BLE fakery over a general-purpose radio — a tool for conducting research in this direction — in the next section. It is followed by a detailed design and analysis of the *Cheep*Syc architecture and performance that is able to achieve an average time synchronization accuracy in the range of 10 μs. The final sections provide a concise background of existing work in this related field, and we conclude with a summary of the areas covered in the article.

System Overview

The system is composed of two units: beacon and control. The *beacon unit* consists of resource constrained sensing tags that are deployed in the region of interest. They are responsible for measurement of simple physical parameters, and disseminating that information through undirected BLE broadcasts. The *control unit* consists of a resourceful gateway device capable of listening and receiving broadcast data contained in BLE advertisements. In our case, the beacon unit is a custom designed nRF24*Cheep* BLE platform, and the control unit is an Android v4.4.4 smartphone that uses Bluedroid (the default Android Bluetooth stack).

nRF24*Cheep*

The custom designed beacon platform, nRF24*Cheep* (Fig. 1), consists of an RL78-G14 microcontroller, an nRF24L01+ 2.4 GHz RF transceiver with an embedded baseband protocol engine (Enhanced ShockBurst™), a TPS62730 synchronous step-down DC-DC converter, and a CR2032 battery holder. The microcontroller has an RL78 core with 16–512 kB flash memory, and operates at a maximum clock speed of 64 MHz with a high precision (±1 percent) on-chip oscillator. It also provides a range of *interval timers* for different application requirements. The nRF24L01+, although not strictly compatible with BLE, can be made to operate as a BLE transmitter by configuring the transceiver settings.[2] This allows BLE listeners/scanners to view the nRF24L01+ as a BLE device and decode its advertisement data. The nRF24L01+ interfaces with the application controller (RL78-G14) over a high-speed Serial Peripheral Interface (SPI) bus. Enhanced ShockBurst™, designed to handle all the high-speed link layer operations, is based on packet communication and supports 1 to 32 bytes of dynamic payload length. Data flow between the radio front-end and the microcontroller is through internal first-in first-out buffers (FIFOs).

The nRF24L01+ module has the following eight interfacing pins, of which four are SPI related: CSN, SCK, MISO, MOSI; the remaining ones are Vcc, GND, IRQ, and CE. CE is used to control data transmission and reception in TX and RX modes, respectively. In TX mode, CE is always kept low except when the packet has to be transmitted, and is done by loading the TX FIFO and then toggling the CE pin. IRQ is the interrupt pin, and can be used to assert three internal interrupts: data received, data transmitted, and maximum number of transmit retries reached.

[2] The necessary RF configurations for BLE compliance can be obtained from Bit-Banging Bluetooth Low Energy: http://goo.gl/JDpflR

™ Enhanced ShockBurst is a trademark of Nordic Semiconductor.

Send time is the time spent by the transmitter to assemble the message and trigger the send request to the MAC layer. It is, therefore, a function of the processor load and the system call overhead of the respective OS ported on the transmitter platform.

*CHEEP*SYNC: TIME SYNCHRONIZATION PROTOCOL

The basic *Cheep*Sync mechanism uses the broadcaster mode of BLE to transmit a single advertisement packet from the beacon (transmitter) unit to the control (receiver) unit. The broadcasted message contains:

- The transmitter's current timestamp value, which is a counter field that increments every *interval* millisecond, and is the estimated local time at the transmission of the advertisement packet;
- The aggregate delay incurred during the transmission of the previous packet

On message reception, the receiver obtains the corresponding "wall'" clock time expressing time (in nanoseconds) since the epoch. In principle, two broadcast packets provide a *synchronization* point between the transmitter and the receiver.[3] The difference between the local and "wall" clock time of a synchronization point estimates the clock offset of the transmitter.

The counter value *ts_counter* is needed to estimate the time elapsed on the beacon unit. If the underlying platform uses a b bit field to store the timestamp value, and the timer fires every *interval* ms, the total time that can be represented by the timestamp field is ($2^b * interval$) ms. Therefore, the configuration of $b = 24$ bits and $interval = 100$ ms can be used to make the timestamp value work for approximately 19 days without rolling over.[4]

In the next subsection, we discuss the general uncertainties associated with RF message delivery and then converge to our specific use case. It is then followed by a detailed explanation of the structure of the BLE advertisement packet with specifics into the packet restructuring as per the nRF24L01+ Enhanced ShockBurst protocol engine.

SOURCES OF TIME SYNCHRONIZATION ERROR

We shall use the following error decomposition model [2–4, 9, 11] to better understand the sources of latency, and modify it according to the specifics of the platform and radio of interest.

Send Time: It is the time spent by the transmitter to assemble the message and trigger the send request to the medium access control (MAC) layer. Therefore, it is a function of the processor load and the system call overhead of the respective OS ported on the transmitter platform. nRF24*Cheep* does not have an OS port, and makes direct system calls to the underlying hardware without any (potential) soft routing. This enables more user control over different system modules (albeit with increased complexity), and thus helps to reduce delays that are typically nondeterministic and were previously difficult to calibrate.

Access Time: It is the delay incurred waiting for access to the transmit channel up to the point when transmission begins, and is specific to the MAC protocol in use. It is considered the least deterministic part of the message delivery system. BLE *does* provision for a (time/frequency-division multiple access, TDMA/FDMA) MAC, but it is only operational in connection mode. Access control rules have not been defined for the BLE broadcast mode of communication, so packets get pushed on to the physical channel as and when they are flagged for transmission.

Transmission Time: A function of the length of the message and the radio speed, it is the time taken by the transmitter to transmit the message and is a deterministic component.

Propagation Time: Once the message has left the transmitter, it is the time needed to transit to the receiver. For many application requirements (wherein the channel length is under 300 m), this delay is highly deterministic and less than 1 μs.

Reception Time: It is the time taken by the receiver to receive the message. In our case, it is the most nondeterministic part of the message delivery mechanism as the receivers are Android ported smartphones that run multiple tasks and process threads at the same time.

Receive Time: It is the time required to process and notify the incoming message to the receiver application.

We perform low-level timestamping, at both the transmitter and receiver ends, to overcome the above stated uncertainties in a message transaction, the details of which are discussed subsequent to the BLE advertisement packet format that follows below.

ADVERTISEMENT PACKET FORMAT

There is a single format for a BLE (advertisement or data) `Packet`, and it consists of the following four fields:

1. Preamble (1 octet)
2. Access address (4 octets)
3. Protocol data unit (PDU: conventionally[5] 2–39 octets, but limited to 2–32 octets in nRF Enhanced ShockBurst™ packet format)
4. Cyclic redundancy check (CRC, 3 octets)

As per the core specifications of an advertisement packet, the 8-bit preamble and 32-bit access address were set to `10101010b` and `0x8E89BED6`, respectively. The preamble is used in the receiver to perform frequency synchronization, symbol timing estimation, and automatic gain control training. A 24-bit CRC is appended to the end of every packet, and is calculated over the PDU. It is important to note that some fields in the packet definition, marked RFU, are reserved for future use, and are set to zero at transmission and ignored upon receipt. Depending on the PDU size, a BLE advertisement packet length could vary from 10 to 40 octets in the nRF Enhanced ShockBurst mode (as opposed to the standard 10–47 octets payload). The *advertisement* channel PDU has a 16-bit header and a variable-size payload.

Header: The header consists of the following six fields spanning over 2 octets:

1 `PDU type` (4 bits);
2 `RFU` (2 bits);
3 `TxAdd` (1 bit);
4 `RxAdd` (1 bit);
5 `Length` (6 bit);
6 `RFU` (2 bit).

The `PDU type` was set to `ADV_NONCONN_IND (0010b)` for *transmitting* non-connectable undirected advertising events. The following `RFU`, `TxAdd`, and `RxAdd` fields were not used, and hence were set to zero. The payload size is indicated by the `Length` field, and can vary between 6 to 30 octets (instead of the standard 37 octets).

Payload: The `Payload` for the `ADV_NON-`

[3] We explain how the determinism of time delays on the nRF24*Cheep* beacon platform benefits our implementation. Therefore, for our platform of choice, only a single broadcast packet (instead of two) is apt for reliably synchronizing the transmitter and the receiver below.

[4] Our system design is based on the assumption that there must be a control unit that passes by the beacon units at least once over a 19-day period. However, under overriding circumstances, the system can determine the number of counter rollovers.

[5] By conventional/standard, we mean the guidelines provided in the Bluetooth core specification v4.

CONN_IND PDU consists of the following two fields:

1. AdvA (6 octets)
2. AdvData (0–24 octets, instead of the standard 0–31 octets)

The AdvA field holds the device address of the advertiser, which can be either public (if TxAdd = 0) or random (if TxAdd = 1). The AdvData field contains the advertisement data, and it consists of two logical parts: significant and non-significant. The significant part contains a sequence of AD structures. Each AD structure consists of the following two fields to populate a separate item of user data:

1 Length (1 octet)
2 Data (Length octets)
- AD type (1 octet)
- AD data (Length - 1 octets)

The non-significant part extends the advertising data to the remaining octets and contains all zeroes. For our implementation, we use the AD structures (defined in the core specification) presented in Table 1.

The AD type Flags sets the discoverability preference of the device, and the General Discoverable Mode makes it detectable unconditionally. The Local Name (Shortened) AD type sets the short user-readable name of the device. The Manufacturer Specific Data AD Type is a generic, freely formattable data field, and includes the 24-bit timestamp value and an 8-bit transmit time delay (of the previous packet).

Timestamping at the Transmitter

On the nRF24*Cheep* platform, an advertisement packet is transmitted by loading the TX FIFO and pulling the CE pin to a high state. One of the fields that is pushed into the SPI buffer is the current timestamp value. A second timestamp is also recorded once the TX_DS interrupt is seen by the microcontroller (i.e., when the radio's IRQ pin is pulled down low), signaling the success of the transmit event. The difference between these two timestamps provides an accurate estimate of the time delays incurred due to send, access, and transmission; and this information is encapsulated into the next advertisement packet. As discussed above, due to the high determinism of send, access, and transmission time delays on the nRF24*Cheep* beacon platform, our implementation uses a *single* broadcast packet (instead of two) to reliably synchronize the transmitter and the receiver. The timestamping characteristics at the transmitter end are shown in Fig. 2, depicted as a histogram showing the distribution of the transmitting time interval recorded for 35,000 broadcast packets. The distribution appears Gaussian with best fit parameters of $\mu = 0.201829$ μs and a minuscule $\sigma^2 = 5.19537e{-}07$ μs. This latency characterization supports the determinism of our approach on the transmitter end.

Timestamping at the Receiver

The reception time on the Android phone can be further divided into the following delivery delays:

- Time taken by the Broadcom BCM radio to receive the message and raise an interrupt
- Time taken by the standard UART platform driver to hook into the BCM radio and register a RX event

Field	Value
Length AD type AD data	2 octets Flags General discoverable mode
Length AD type AD data	6 octets Local name (shortened) –
Length AD type AD data	4 octets Manufacturer-specific data Timestamp, transmit time delay (previous packet)

Table 1. Utilized AD types.

- Time taken by the Bluedroid stack to poll the UART driver, waiting to check if there are any bytes to be read (using the userial_read_thread() method in userial.c)

Therefore, there is a significant nondeterministic time lag between the instants when an RF message is actually received by the radio to when it is processed by the Bluedroid stack. Moving to software layers beyond the UART will limit portability across different Android phones. Taking these facts into consideration, the lowest layer accessible entity on the Android Bluedroid stack is userial.c.[6]

userial.c interfaces with the standard UART driver, and provides a common interface for the time-sync to work on multiple phones without introducing many hardware changes. However, timestamping at the level of userial.c introduces two complications. First, it is a generic container for catching all types of events (notification and other messages) sent out by the BCM radio to the underlying driver; hence, there is a need to correctly identify the Bluetooth event. Second, this problem can be overcome by waiting for userial.c to process the event list, but would introduce finite variable delays. Therefore, timestamps are recorded at the instant when an event is received at userial.c, but before any processing is performed on the list.

The clock_gettime(CLOCK_REALTIME) method, which returns the real-time clock of the system in nanoseconds since the Epoch (00:00 1 January, 1970 UTC), is utilized to perform low-level timestamping on the receiver (phone) end. This representative value is passed on to the higher-layer application, and the receive timestamp corresponding to the received BLE packet is subsequently determined. It is possible that a notification was received by userial.c from the radio, and the application layer was called following that particular notification being recorded as a timestamp. This introduces an error since the timestamp of the notification and not that of the message is being used by the program.

To overcome this problem, we adopt an algorithmic approach that determines the timestamp corresponding to a received packet. The algorithm first iteratively records the timestamps for the previous and current packets, and then subtracts each current packet's timestamp value from every previous packet's timestamp value. Upon completion, the value with the least devi-

The algorithm first iteratively records the timestamps for the previous and current packets, and then subtracts each current packet's timestamp value from every previous packet's timestamp value. Upon completion, the value with the least deviation is taken as the best fit timestamp for the particular packet.

[6] In the Android Bluedroid stack, userial.c is located at: /external/bluetooth/bluedroid/hci/src/

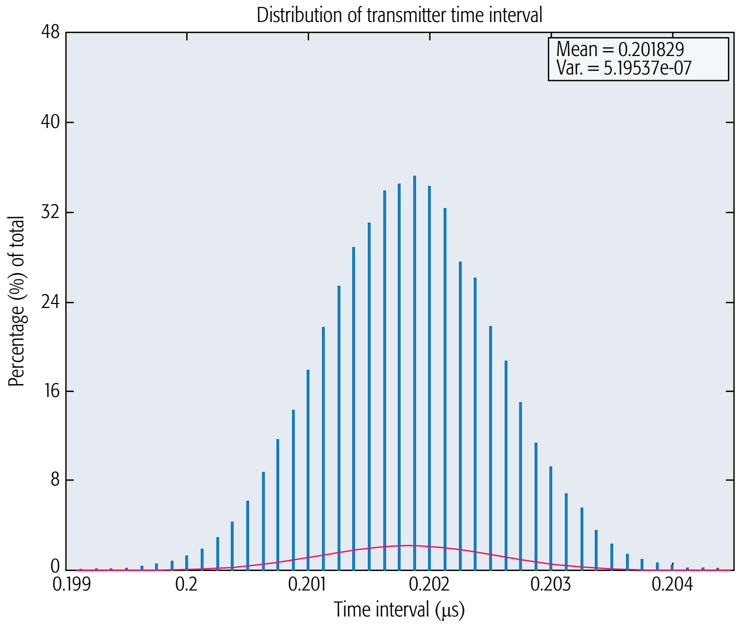

Figure 2. Transmitter side timestamping characteristics. Histogram showing the distribution of transmitter time interval recorded for 35,000 broadcast packets, grouped into 1 µs buckets. The curve is a plot of the best fit Gaussian parameters with $\mu = 0.201828$ µs and $\sigma^2 = 5.19537e{-}07$ µs.

ation is taken as the best fit timestamp for the particular packet.

Clock Drift Management

The clock quality and speed on the transmitter (nRF24*Cheep*) and the receiver (Nexus~5 phone) are vastly different. *Cheep*Sync makes continuous skew adjustments over a measurement window on the phone unit as: $k = (o_i - \overline{o})/(t_i^r - \overline{t^r})$; where, $\overline{t^r}$ is the average elapsed time at the reference clock and $\overline{o}$ is the average offset up to the i^{th} sample point. This linear regression method offers a fast mechanism to find the frequency and phase errors over time.

Figure 3 shows the performance of time synchronization with and without clock drift compensation for an interval value of 100 ms over a period of 13 h. Without drift compensation, the mean error in synchronizing the transmitter and receiver clock is in the millisecond range, but can be brought down to a few microseconds with correction (including both low-level timestamping and clock drift management). For the specific case of *interval* = 100 ms, the error before drift correction was $\mu = 2.62$ ms, $\sigma^2 = 6.37312$ ms, and 95 percent cumulative probability of 8.2 ms; while the respective estimates after correction were recorded as $\mu = 0.0667$ ms, $s^2 = 3.89512e{-}02$ µs, and 0.64 µs at the 95 percent cumulative probability level. We also experimented with an interval value of 200 µs [12], but the interval timer of 100 ms was chosen due to its high stability and lower energy cost.

Multi-Device Time Synchronization

Considering the system dynamics, there are two forms of time synchronization across multiple devices.

Many-Tx-to-One-Rx Synchronization: The scenario of synchronizing multiple Tx's to a single Rx is a simple extension of the case of deriving an estimate of time for a single (transmitter, receiver) pair, wherein the control unit becomes the reference point to perform synchronization. A reference point contains a pair of local and global timestamps where both of them refer to the same time instant. The control unit receives periodic broadcasts from beacon units within their coverage zone; otherwise, their records are not entered. When the control unit collects the required measurement points, it estimates the skew and offset of the observed beacons and derives their coordinated time measure with respect to the global time.

Many-Tx-to-Many-Rx Synchronization: The scenario of synchronizing multiple Tx's to multiple Rx's combines the above system infrastructure with a mechanism for different control units to share the skew and offset information of already visited beacon units. This could be achieved by peer-to-peer interaction between the control units or through the cloud infrastructure. For ease of implementation, we choose to take the latter alternative with the Google Cloud platform. In this case, the absolute timestamp of the control unit along with the timestamp of the beacon unit is inserted into the cloud. This data is then made accessible to other control units using the Google App Engine database. Whenever the timestamp of a new beacon unit is recorded by a control unit, a notification is sent to all other control units using Google cloud messaging. Once a control unit receives this message, it observes the last time it obtained the timestamp of the same beacon and computes the time difference. For instance, let the first control unit record timestamp t_a from beacon i at time instance τ_a. Let the second control unit record another timestamp from the same beacon i with value t_b at time instance τ_b. Using this information, the offset (i.e., time difference) between the phones is measured as $(\tau_a - \tau_b - (t_b - t_a) * k)$.

The relative time on the control units (which are Android phones) may vary from one to another by up to several seconds. It is therefore necessary to constantly synchronize with the nearest Network Time Protocol (NTP) server repeatedly in order to maintain very high precision for our time synchronization. For our implementation, we used an Android NTP application,[7] and was configured to synchronize the control units at an interval of 30 s.

Porting CheepSync into Mobile Applications

The reference design of nRF24*Cheep*, source code of *Cheep*Sync, and build instructions for a custom Android ROM are publicly available at https://github.com/prasantmisra/cheep-sync, with some preliminary results reported in [13]. The accuracy levels reported in the following section can only be obtained with the nRF24*Cheep* platform, while other BLE platforms will/may lead to higher time sync error values. As for the mobile phone platform, a custom Android ROM that includes the modified `userial.c` file needs to be installed. The same link also contains an example to capture the steps for incorporating *Cheep*Sync into a mobile application.

Experimental Evaluation

In this section, we evaluate the accuracy of *Cheep*Sync in a variety of controlled and uncontrolled experiments.

*Cheep*Sync in Controlled Experiments

Study 1: The aim of this study was to obtain the baseline performance level of *Cheep*Sync for a *many-Tx-to-one-Rx* synchronization scenario. Thus, it was designed with a *static* setup of 8 beacon units and 1 control unit that was always covered by all the beacons. The beacons were configured to broadcast advertisement packets at an interval of 100 ms and at their lowest transmit power of –20 dB. This experiment was conducted for about 13 hours wherein an average of 10000 packets were received by the control unit from each beacon.

The experimentation results are shown in Fig. 4. The system-level performance of *Cheep*Sync is depicted in Fig. 4a, which shows an average error of 8 μs and a 95 percent error probability of less than 0.04 ms. Figure 4b disaggregates the combined (drift compensated) measurements into respective representations for each of the eight beacons in the system. Here, it is evident that a large percentage of the beacon units show consistent behavior with an average error level of approximately 10 μs and a worst case value of 0.04 ms; except beacon 7, which shows 2 times better average performance but shoots over in the worst case performance level.

Study 2: The aim of this study was to obtain the baseline performance level of *Cheep*Sync for the scenarios of *one-tx-to-many-rx* and *many-tx-to-many-rx* synchronization. For the prior scenario, the static setup consisted of a single beacon and two control units, while for the latter scenario, the setup consisted of eight static beacons and two static control units. The control units were always kept within the coverage range of the beacon(s), while all other beacon/control unit configuration parameters were kept consistent with study 1. Figures 4c and 4d show the cumulative time-sync error (i.e., the time difference between control units as explained earlier) distributions for the prior and latter setup, respectively, and it is conditioned on contributing error factors including clock drift, NTP drift, and so on. The result obtained from Fig. 4c suggests that there is a 95 percent probability of the time synchronization error to be less than 5 ms with mean < 10 ms. A similar observation is also noted from Fig. 4d where the result suggests that there is a 90 percent probability of the time sync error to be less than 10 μs, although with the possibility of errors as large as 10 ms for 5 percent higher confidence levels.

*Cheep*Sync in Uncontrolled Experiments

Study 3 — This experiment was designed to evaluate the performance of *Cheep*Sync in an uncontrolled setup, and hence quantify its deviation from the respective benchmark results obtained from studies 1 and 2. The experimentation space was in a [20 × 20] m portion of our office floor. For our evaluation, we took the services of two people who were each handed an Android smartphone running the *Cheep*Sync time service and kept it for about an hour. The respective office space was instrumented with a set of five beacon units in a manner that all beacons did not cover every end of the experiment zone. All other system configurations, such as the beaconing rate and transmit power levels, were kept the same as in study 1/2.

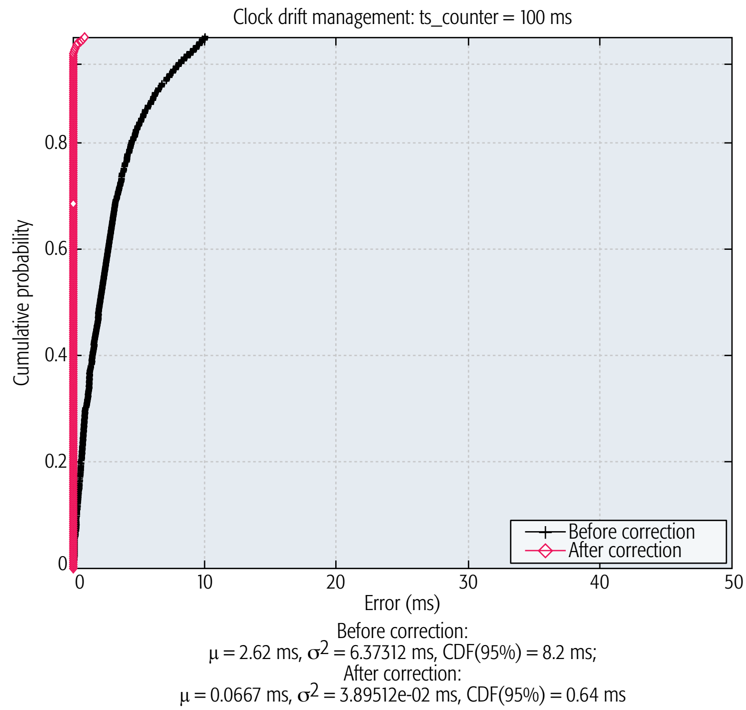

Figure 3. Clock drift management. There is one order of improvement in synchronization accuracy after compensating for clock drift, at both the average and 95 percent probability levels.

For the *many-tx-to-one-rx* scenario (Fig. 4e), a mean error of 12 μs was observed, while its 95 percent error probability was less than 0.04 ms. As for the *many-tx-to-many-rx* scenario (Fig. 4f), the 95 percent error probability was less than 22 ms with a mean error of 10 μs. Both of these results are in good agreement with the baseline observations recorded in Figs. 4a and 4c, thereby establishing the efficacy of *Cheep*Sync.

Discussion

Time synchronization solutions have some basic features in common: a messaging protocol to exchange timestamped packets between nodes (some acting as clients and others as time servers) in a network, techniques for overcoming nondeterministic delays, and an adjustment mechanism to update the local clock. However, they differ in aspects including whether the physical clocks of the network are kept consistent internally or are synchronized to external standards; if the server is an arbiter of the client clock or is considered as a canonical clock; and so on.

For synchronization in WSNs, there are multiple algorithms such as Reference Broadcast Synchronization (RBS) [3], Flooding Time Synchronization Protocol (FTSP) [2], Timesync Protocol for Sensor Networks (TPSN) [4], Glossy [5], and hardware assisted clock synchronization (HACS) [10] that use message passing as the basic mechanism to compensate for delays. *Cheep*Sync explores a form of synchronization that differs from traditional WSNs, and is specifical-

[7] The Android NTP application is available at https://play.google.com/store/apps/details?id=ru.org.amip.ClockSync&hl=en. While alternatives such as GPS exist, NTP was chosen due to its wide availability and benefits of scalability, self-configuration, robustness to network failures, and energy efficiency.

CheepSync explores a form of synchronization that differs from traditional WSNs, and is specifically designed for the technology segment that uses BLE undirected broadcasts and smartphones. Its fundamental design philosophy is to provision for synchronization between transmitters and receivers, as opposed to traditional WSN protocols.

ly designed for the technology segment that uses BLE undirected broadcasts and smartphones. Its fundamental design philosophy is to provision for synchronization between transmitters and receivers, as opposed to traditional WSN protocols that synchronize a set of receivers with one another. Performing receiver-receiver synchronization removes uncertainties associated with send and access delays, the biggest contributors to non-determinism in latency, from the critical time path. System-level complexities are further reduced by using symmetric WSN platforms that run on simple OSs. All of these factors conglomerate toward higher levels of synchronization accuracy than is typically evident in traditional WSNs. *Cheep*Sync, in contrast, operates on highly asymmetric devices with vast differences in system complexity. Its synchronization methodology, system architecture, and wireless communication standard utilized for delivering the respective solutions are completely different. *Cheep*Sync achieves a high synchronization accuracy that is better than the time-sync levels of RBS, TPSN, and HACS, but is less precise compared to FTSP or Glossy (Table 2). We refer

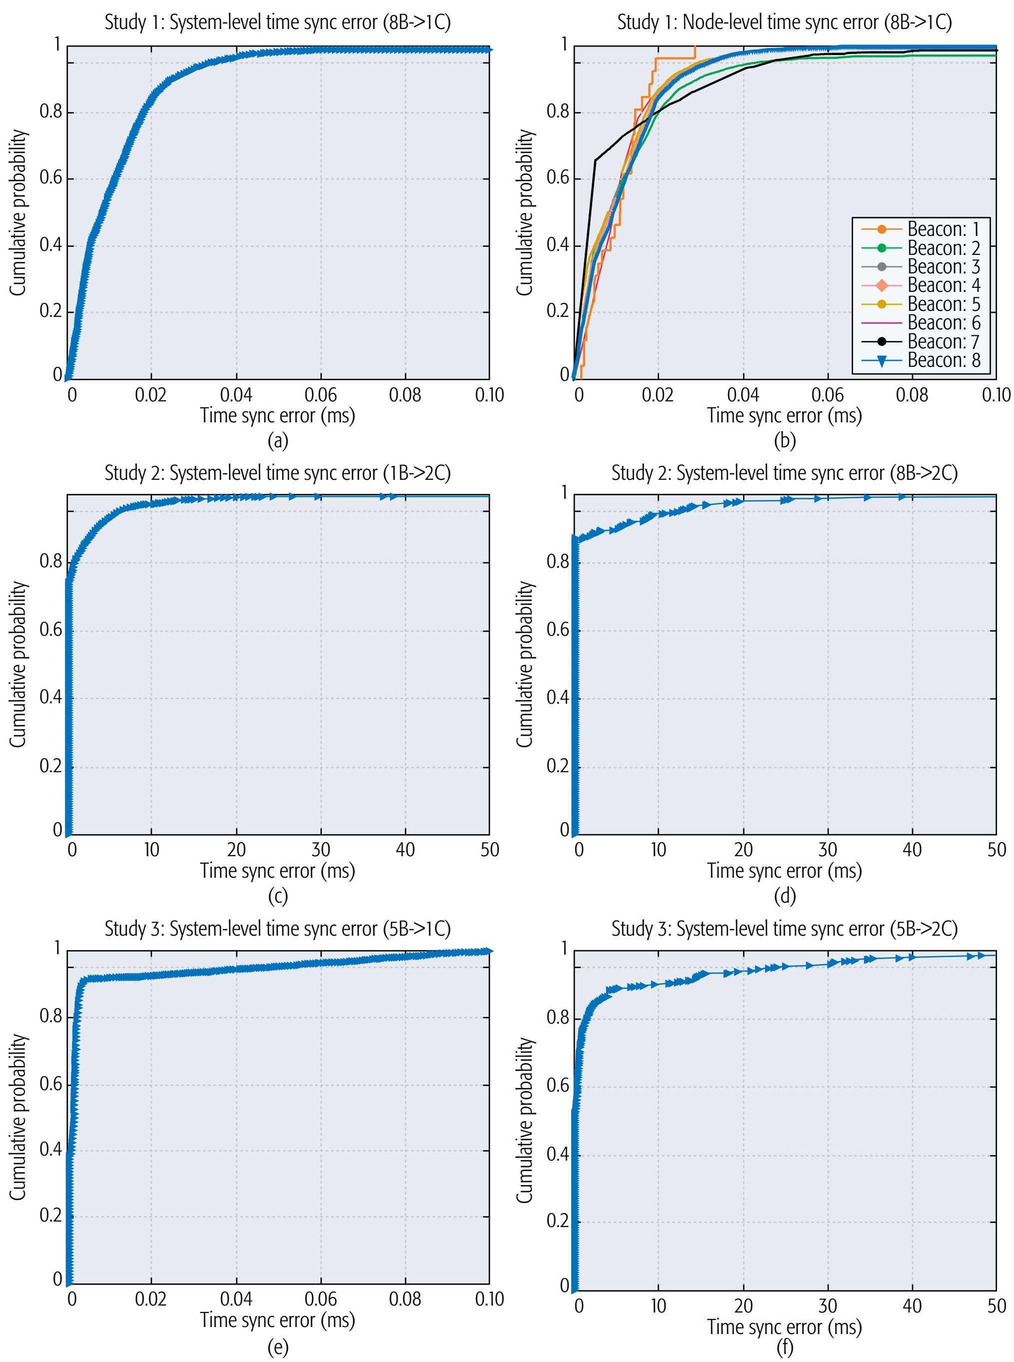

Figure 4. Performance of *Cheep*Sync. a) many-tx-to-one-rx: $\mu = 8\ \mu s$, CDF(95 percent) = 0.04 ms; b) many-tx-to-one-rx; c) one-tx-to-many-rx scenario: $\mu = 10\ \mu s$, CDF(95 percent) = 5 ms; d) many-tx-to-many-rx scenario: $\mu = 10\ \mu s$, CDF(95 percent) = 10 ms; e) many-tx-to-one-rx scenario: $\mu = 12\ \mu s$, CDF(95 percent) = 0.04 ms; f) many-tx-to-many-rx scenario: $\mu = 10\ \mu s$, CDF(95 percent) = 22 ms.

our reader to Serpedin *et al.* [14] for an extensive survey and analysis of various solutions in this space.

Application

The loosely coupled BLE data collection framework enables new ways of architecting Internet of Things and Humans (IoTH) systems. An example application is to monitor the *hand sanitization* status of healthcare personnel (HCP) in an intensive care unit (ICU) to control and prevent the spread of hospital acquired infections (HAIs), where unclean hands are the most common factor contributing to this cause. CleanHands [15] is a recently proposed system that uses a combination of BLE beacon tags and HCP's mobile phones for detecting occurrences of noncompliance with hand hygiene. This application requires different levels of time accuracy. For instance, a few milliseconds may be sufficient to disambiguate the handwashing event of HCPs, but a significantly higher accuracy level of a few microseconds would be required to secure the same system against possible gaming.

Conclusion

*Cheep*Sync is a time synchronization service for BLE advertisers, and therefore is a key enabler for a variety of IoTH applications where mobile crowdsourcing, using existing infrastructure and ubiquitously used platforms along with humans as the data mules, has emerged as an alternate architecture. These applications typically require different levels of time accuracy. By empirical evaluations, we show that CheepSync is capable of gracefully handling timing requirements as low as 10 μs. It is built on the existing Bluetooth v4.0 standard, and hence is generic to all devices using the low energy profile of Bluetooth.

Acknowledgments

This work has received funding from the Robert Bosch Centre for Cyber Physical Systems, IISc Bangalore under grant agreement RBCCPS/PC-0043.

Time sync. protocol	Avg. accuracy (μs) (single hop)
RBS [3]	29.10
TPSN [4]	16.90
FTSP [2]	01.48
Glossy [5]	0.50
HACS [10]	1000.00
*Cheep*Sync	10.00

Table 2. Accuracy measure of time synchronization algorithms.

By empirical evaluations, we showed that CheepSync is capable of gracefully handling timing requirements as low as 10 μs. It is built on the existing Bluetooth v4.0 standard, and hence is, generic to all devices using the low energy profile of Bluetooth.

References

[1] L. Lamport, "Time, Clocks, and the Ordering of Events in a Distributed System," *Commun. ACM*, vol. 21, no. 7, July 1978, pp. 558–65.

[2] M. Maróti *et al.*, "The Flooding Time Synchronization Protocol," *Proc. 2nd Int'l. Conf. Embedded Networked Sensor Systems*, New York, NY, 2004, pp. 39–49.

[3] J. Elson, L. Girod, and D. Estrin, "Fine-Grained Network Time Synchronization Using Reference Broadcasts," *Proc. 5th Symp. Operating Systems Design and Implementation*, New York, NY, 2002, pp. 147–63.

[4] S. Ganeriwal, R. Kumar, and M. B. Srivastava, "Timing-Sync Protocol for Sensor Networks," *Proc. 1st Int'l. Conf. Embedded Networked Sensor Systems*, New York, NY, 2003, pp. 138–49.

[5] F. Ferrari *et al.*, "Efficient Network Flooding and Time Synchronization with Glossy," *Proc. 10th Int'l. Conf. Information Processing in Sensor Networks*, Apr. 2011, pp. 73–84.

[6] D. L. Mills, "Internet Time Synchronization: The Network Time Protocol," *IEEE Trans. Commun.*, vol. 39, no. 10, Oct 1991, pp. 1482–93.

[7] R. Gusella and S. Zatti, "The Accuracy of the Clock Synchronization Achieved by Tempo In Berkeley Unix 4.3bsd," *IEEE Trans. Software Engineering*, vol. 15, no. 7, Jul 1989, pp. 847–53.

[8] F. Cristian, "Probabilistic Clock Synchronization," *Distributed Computing*, vol. 3, no. 3, 1989, pp. 146–58.

[9] H. Kopetz and W. Ochsenreiter, "Clock Synchronization in Distributed Real-Time Systems," *IEEE Trans. Computers*, vol. 36, no. 8, Aug. 1987, pp. 933–40.

[10] A. Rowe, V. Gupta, and R. Rajkumar, "Low-Power Clock Synchronization Using Electromagnetic Energy Radiating from AC Power Lines," *Proc. 7th ACM Conf. Embedded Networked Sensor Systems*, New York, NY, 2009, pp. 211–24

[11] M. Horauer *et al.*, "Psynutcevaluation of a High Precision Time Synchronization Prototype System for Ethernet LANs," *Proc. 34th Annual Precise Time and Time Interval Meeting*, 2002, pp. 263–79.

[12] S. Sridhar *et al.*, "CheepSync: A Time Synchronization Service for Resource Constrained Bluetooth Low Energy Advertisers," tech. rep., 2015.

[13] S. Sridhar, P. Misra, and J. Warrior, Poster Abstract: "Cheepsync: A Time Synchronization Service for Resource Constrained Bluetooth Low Energy Advertisers," *Proc. 14th Int'l. Conf. Information Processing in Sensor Networks, IPSN '15*, New York, NY, 2015, pp. 364–65.

[14] E. Serpedin and Q. M. Chaudhari, *Synchronization in Wireless Sensor Networks: Parameter Estimation, Performance Benchmarks, and Protocols*, Cambridge Univ. Press, 2009.

[15] P. Misra *et al.*, poster abstract, "Cleanhands: An Integrated Monitoring System for Control of Hospital Acquired Infections," *Proc. 14th Int'l. Conf. Information Processing in Sensor Networks*, New York, NY, 2015, pp. 348–49.

Biographies

Sabarish Sridhar is currently pursuing a B.Tech. in electronics and communications at M.S. Ramaiah Institute of Technology, Bangalore. He is passionate about healthcare, sensor networks, and smart systems. His current research interests include automating analog design using geometric programming, image processing with field programmable gate arrays, and development of embedded frameworks for sensor systems that are economically viable.

Prasant Misra [SM] is a senior MTS at the Robert Bosch Centre for Cyber Physical Systems in the Indian Institute of Science, Bangalore. He received his Ph.D. from the University of New South Wales, Sydney, Australia, in 2012. His current research interests include low-power sensing/communication, signal processing, and energy-efficient computing with a focus on system design and implementation within the general framework of cyber physical systems and the Internet of Things. He has many years of experience in technology development, and has worked in different roles and capabilities for Keane Inc. (now a unit of NTT Data Corporation), India; CSIRO ICT Centre, Australia; Red Lotus Technologies, United States; and SICS Swedish ICT, Sweden, the outcomes of which have resulted in either commercial products or publications in premier sensornet forums such as *ACM/IEEE IPSN* and *ACM TOSN*. His professional and research contributions have been recognized by numerous awards, of which it is noteworthy to mention the ERCIM Alain Bensoussan/Marie Curie Fellowship (2012) and the AusAID Australia Awards Leadership Program (2008). He has also served on the organizing/technical committees of a number of international conferences. He is a member of ACM, Secretary of the IEEE Computer Society (Bangalore Section), and an Associate Technical Editor of *IEEE Communications Magazine*.

Gurinder Singh Gill did not have a biography available at the time of publication.

Jay Warrior has over 20 years of creating new high-technology-based business opportunities for Agilent Technologies, Hewlett-Packard, Emerson, Fisher-Rosemount, and Honeywell in the United States and Asia. He is driven by a systems-wide perspective to problems, integrating long-term trends with strategy, technology development, and user-centric design techniques in his work. He puts these into practice at Mobiatrics LLC, a healthcare focused startup. He also works extensively with TiE and IESA on developing the IoT ecosystem for India. His previous experiences include his roles as chief technologist at the Robert Bosch Centre for Cyber Physical Systems at IISc, Bangalore, a technology transfer organization, chief technologist and strategist for the Network Solutions business at Agilent Technologies, and managing director of New Business Creation at Agilent. He has led the the development of multiple large-scale distributed systems, including the HART protocol, a de facto standard in process automation. He holds a Ph.D. in control and dynamical systems, and currently has over 25 patents covering key inventions in networking and diagnostics technology.

COMPANY .... PAGE

CURRENTLY SCHEDULED TOPICS

TOPIC	ISSUE DATE	MANUSCRIPT DUE DATE
SOCIAL AND MOBILE SOLUTIONS IN AD HOC AND SENSOR NETWORKING	JULY 2016	JANUARY 11, 2016
SDN USE CASES FOR SERVICE PROVIDER NETWORKS	OCTOBER 2016	JANUARY 31, 2016
5G RADIO ACCESS NETWORK ARCHITECTURE AND TECHNOLOGIES	NOVEMBER 2016	FEBRUARY 1, 2016
ENABLING MOBILE AND WIRELESS TECHNOLOGIES FOR SMART CITIES	DECEMBER 2016	FEBRUARY 29, 2016
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 FEBRUARY ISSUE

UNDERWATER WIRELESS COMMUNICATIONS

OPTICAL COMMUNICATIONS

AUTOMOTIVE NETWORKING