IEEE Communications MAGAZINE

July 2009, Vol. 47, No. 7

www.comsoc.org

Network and Service Management

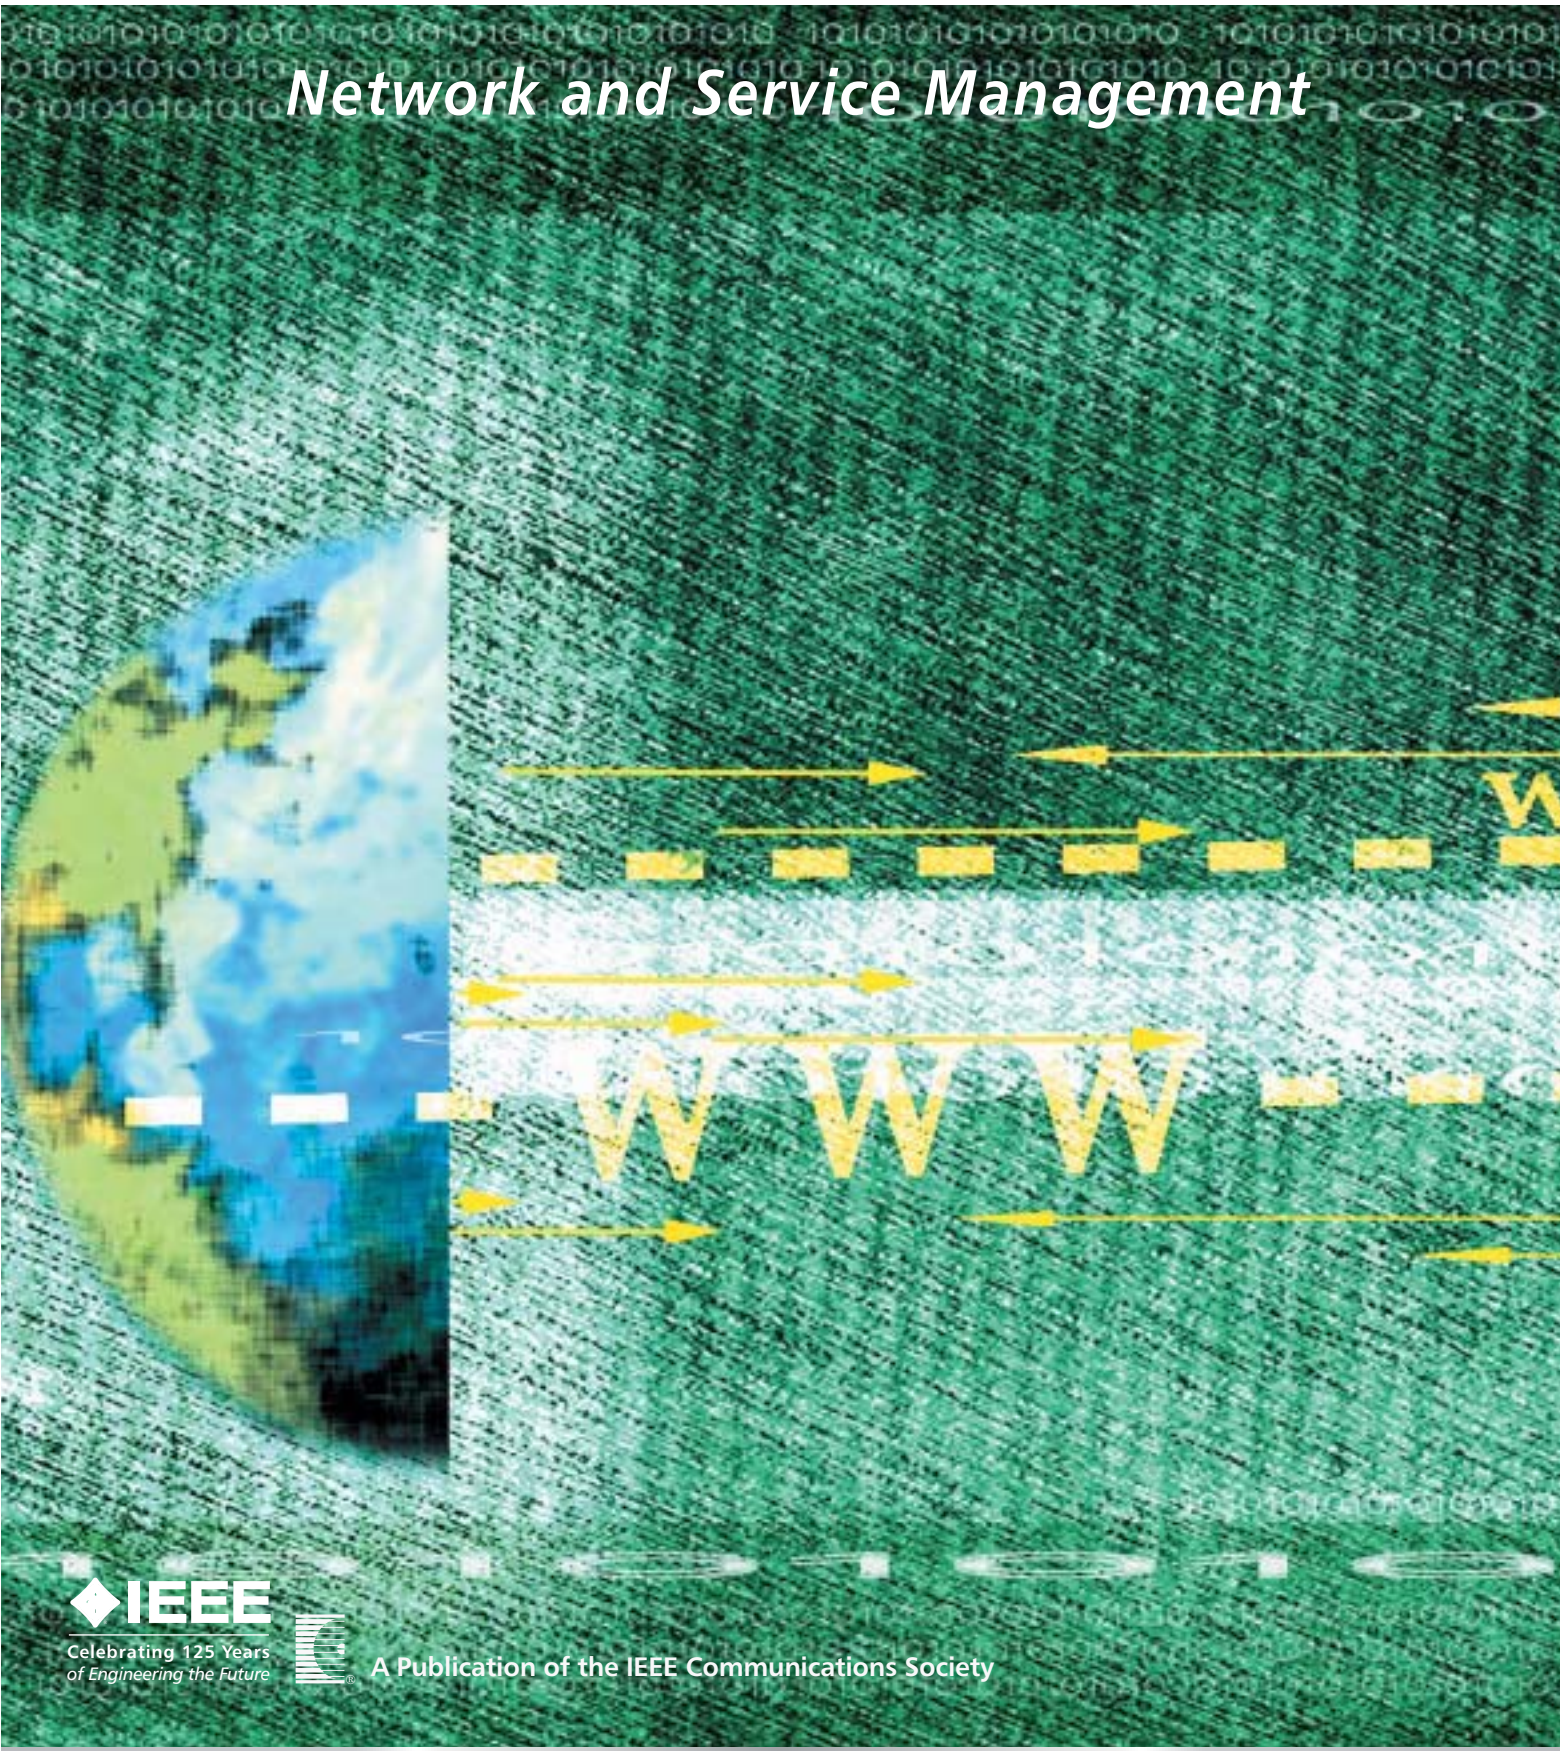

IEEE Celebrating 125 Years of Engineering the Future

A Publication of the IEEE Communications Society

Now Showing:
The Future of Video Entertainment

Video is driving a market transition in networking
that creates unique challenges for service providers.
Evolve your network with Cisco into a medianet
and go beyond being a service provider
to become an experience provider.
Cisco is the only company, which along with its partners,
can address video and rich media
from an end-to-end perspective
by leveraging deep expertise in IP, video
and customer premise solutions.
Move video more efficiently,
and energize customer satisfaction
and increase revenues.
Deliver the connected life that's more social,
more personal and more interactive.

Begin your transformation at www.cisco.com/go/sp/medianet

welcome to
the human network.

CISCO

FOR RENT

Telogy's Got It! *The Latest Telecom Test Equipment Available for Rent.*

Fluke | DTX-1800 Cable Analyzer

EXFO | FTB-7000 OTDRs

Fujikura | FSM-60S Fusion Splicer

Anritsu | S331D/S332D Site Master

Advantages of Renting Telecom Test Equipment from Telogy

- Availability of multiple products for various applications
- Be a fiscal hero and save your company money
- Rent test equipment that you use infrequently
- An array of equipment from various manufacturers
- Telogy takes care of the maintenance and ownership costs
- You focus on the job at hand not the equipment
- Weekly and Monthly rates
- 2 day FREE transit time allowed
- Competitive pricing, superior service
- Pre-paid shipping for easy rental return
- Rent-Use-Return, Telogy makes it easy!

Fujikura EXFO Anritsu FLUKE

RENT IT

We have flexible terms from 1 week to 48 months.

TELOGY™

WWW.TELOGYLLC.COM

1.800.835.6494

Director of Magazines
Steve Gorshe, PMC-Sierra, Inc. (USA)

Editor-in-Chief
Nim K. Cheung, ASTRI, (Hong Kong)

Associate Editor-in-Chief
Steve Gorshe, PMC-Sierra, Inc. (USA)

Senior Technical Editors
Nirwan Ansari, NJIT (USA)
Tom Chen, Swansea University (UK)
Roch H. Glitho, Ericsson Research (Canada)
Andrzej Jajszczyk, AGH U. of Sci. & Tech. (Poland)
Torleiv Maseng, Norwegian Def. Res. Est. (Norway)

Technical Editors
Koichi Asatani, Kogakuin University (Japan)
Mohammed Atiquzzaman, U. of Oklahoma (USA)
Tee-Hiang Cheng, Nanyang Tech. Univ. (Rep. of Singapore)
Jacek Chrostowski, Scheelite Techn. LLC (USA)
Sudhir S. Dixit, Nokia Siemens Networks (USA)
Nelson Fonseca, State U. of Campinas (Brazil)
Joan Garcia-Haro, Poly. U. of Cartagena (Spain)
Abbas Jamalipour, U. of Sydney (Australia)
Vimal Kumar Khanna (India)
Janusz Konrad, Boston U. (USA)
Nader Mir, San Jose State U. (USA)
Amitabh Mishra, Johns Hopkins University (USA)
Sean Moore, Avaya (USA)
Sedat Ölçer, IBM (Switzerland)
Algirdas Pakstas, London Met. U. (England)
Michal Pioro, Warsaw U. of Tech. (Poland)
Harry Rudin, IBM Zurich Res.Lab. (Switzerland)
Hady Salloum, Stevens Inst. of Tech. (USA)
Heinrich J. Stüttgen, NEC Europe Ltd. (Germany)
Dan Keun Sung, Korea Adv. Inst. Sci. & Tech. (Korea)
Naoaki Yamanaka, Keio Univ. (Japan)

Series Editors
Ad Hoc and Sensor Networks
Edoardo Biagioni, U. of Hawaii, Manoa (USA)
Silvia Giordano, Univ. of App. Sci. (Switzerland)
Automotive Networking and Applications
Wai Chen, Telcordia Technologies, Inc (USA)
Luca Delgrossi, Mercedes-Benz R&D N.A. (USA)
Timo Kosch, BMW Group (Germany)
Tadao Saito, University of Tokyo (Japan)
Design & Implementation
Sean Moore, Avaya (USA)
Integrated Circuits for Communications
Charles Chien (USA)
Zhiwei Xu, SST Communication Inc. (USA)
Stephen Molloy, Qualcomm (USA)
Network and Service Management Series
George Pavlou, U. of Surrey (UK)
Aiko Pras, U. of Twente (The Netherlands)
Topics in Optical Communications
Hideo Kuwahara, Fujitsu Laboratories, Ltd. (Japan)
Jim Theodoras, ADVA Optical Networking (USA)
Topics in Radio Communications
Joseph B. Evans, U. of Kansas (USA)
Zoran Zvonar, MediaTek (USA)
Standards
Yoichi Maeda, NTT Adv. Tech. Corp. (Japan)
Mostafa Hashem Sherif, AT&T (USA)

Columns
Book Reviews
Andrzej Jajszczyk, AGH U. of Sci. & Tech. (Poland)
Communications and the Law
Steve Moore, Heller Ehrman (USA)
History of Communications
Mischa Schwartz, Columbia U. (USA)
Regulatory and Policy Issues
J. Scott Marcus, WIK (Germany)
Jon M. Peha, Carnegie Mellon U. (USA)
Technology Leaders' Forum
Steve Weinstein (USA)
Very Large Projects
Ken Young, Telcordia Technologies (USA)
Your Internet Connection
Eddie Rabinovitch, ECI Technology (USA)

Publications Staff
Joseph Milizzo, Assistant Publisher
Eric Levine, Associate Publisher
Susan Lange, Digital Production Manager
Catherine Kemelmacher, Associate Editor
Jennifer Porcello, Publications Coordinator
Devika Mittra, Publications Assistant

IEEE Communications MAGAZINE

July 2009, Vol. 47, No. 7

www.comsoc.org/~ci

Topics in Network and Service Management

Series Editors: George Pavlou and Aiko Pras

18 **Series Editorial**

20 **Network Virtualization: State of the Art and Research Challenges**
Recently network virtualization has been pushed forward by its proponents as a long-term solution to the gradual ossification problem faced by the existing Internet and proposed to be an integral part of the next-generation networking paradigm. However, many technical issues stand in the way of its successful realization. This article investigates the past and the state of the art in network virtualization along with the future challenges that must be addressed to realize a viable network virtualization environment.
N. M. Mosharaf Kabir Chowdhury and Raouf Boutaba

27 **Future Internet = Content + Services + Management**
By taking the viewpoint of an Internet user who is interested in using Internet services and not so much in the protocols that move data around, the authors first describe some possible future Internet services. Then they derive some network and service management requirements, and discuss some of them in more detail.
Jürgen Schönwälder, Marc Fouquet, Gabi Dreo Rodosek, and Iris C. Hochstatter

34 **With Evolution for Revolution: Managing FEDERICA for Future Internet Research**
Traditional management solutions have followed an evolutionary path, although the scale of the Internet and emerging novel architectures such as peer-to-peer, ad hoc networks, as well as virtualization-capable network infrastructures require focused and possibly revolutionary changes in management approaches. The authors elaborate on challenges posed by the renaissance of virtualization as experienced in the planning, development, and operation of the FEDERICA infrastructure.
Peter Szegedi, Sergi Figuerola, Mauro Campanella, Vasilis Maglaris, and Cristina Cervello-Pastor

40 **Managing Interdomain Traffic in Latin America: A New Perspective Based on LISP**
The characteristics of Latin American network infrastructures have global consequences, particularly in the area of interdomain traffic engineering. The authors analyze the peculiarities of LA interdomain routing architecture, and provide up-to-date data about the combined effects of the multihoming and TE practices in the region. They outline one of the solutions under discussion at the IRTF, the Locator/Identifier Separation Protocol, and examine its potential in terms of interdomain traffic management in the context of LA.
Marcelo Yannuzzi, Xavi Masip-Bruin, Eduardo Grampin, Roque Gagliano, Alberto Castro, and Martin German

Accepted from Open Call

50 **Recent Advances in Amplify-and-Forward Two-Hop Relaying**
The authors review an important class of wireless cooperation protocols known as amplify-and-forward relaying. One or more low-complexity relay nodes assist the communication between sources and destinations without having to decode the signal. This makes AF relaying transparent to modulation and coding of the source/destination communication protocol.
Stefan Berger, Marc Kuhn, Armin Wittneben, Timo Unger, and Anja Klein

57 **Tuning Radio Resource in an Overlay Cognitive Radio Network for TCP: Greed Isn't Good**
The authors illustrate the performance of Transmission Control Protocol in an overlay cognitive radio network under dynamic spectrum access. They show that the performance of TCP in overlay CR networks that implement DSA is quite different from its performance in conventional networks, which do not allow DSA.
Teerawat Issariyakul, Laxminarayana S. Pillutla, and Vikram Krishnamurthy

The MathWorks
Accelerating the pace of engineering and science

क्या आप MATLAB बोलते हैं?

Over one million people around the world speak MATLAB.
Engineers and scientists in every field from aerospace and semiconductors to biotech, financial services, and earth and ocean sciences use it to express their ideas.
Do you speak MATLAB?

Saturn's northern latitudes and the moon Mimas. Image from the Cassini-Huygens mission.

Related article at mathworks.com/ltc

PHOTO: European Space Agency ©2008 The MathWorks, Inc.

MATLAB®

The language of technical computing

2009 Communications Society Elected Officers

Doug Zuckerman, *President*
Andrzej Jajszczyk, *VP–Technical Activities*
Mark Karol, *VP–Conferences*
Byeong Gi Lee, *VP–Member Relations*
Sergio Benedetto, *VP–Publications*
Byeong Gi Lee, *President-Elect*

Members-at-Large

Class of 2009

Thomas LaPorta, Theodore Rappaport
Catherine Rosenberg, Gordon Stuber

Class of 2010

Fred Bauer, Victor Frost
Stefano Galli, Lajos Hanzo

Class of 2011

Robert Fish, Joseph Evans
Nelson Fonseca, Michele Zorzi

2009 IEEE Officers

John R. Vig, *President*
Pedro A. Ray, *President-Elect*
Barry L. Shoop, *Secretary*
Peter Staecker, *Treasurer*
Lewis M. Terman, *Past-President*
E. James Prendergast, *Executive Director*
Curtis A. Siller, Jr., *Director, Division III*
Nim Cheung, *Director-Elect, Division III*

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/ci. 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. Non-member subscription: $400. Single copy price is $25.

EDITORIAL CORRESPONDENCE: Address to: Editor-in-Chief, Nim K. Cheung, Telcordia Tech., Inc., One Telcordia Drive, Room RRC-1B321, Piscataway, NJ 08854-4157; tel: +(732) 699-5252, e-mail: n.cheung@ieee.org.

COPYRIGHT AND REPRINT PERMISSIONS: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limits of U.S. Copyright law for private use of patrons: those post-1977 articles that carry a code on the bottom of the first page provided the per copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923. For other copying, reprint, or republication permission, write to Director, Publishing Services, at IEEE Headquarters. All rights reserved. Copyright © 2009 by The Institute of Electrical and Electronics Engineers, Inc.

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://commag-ieee.manuscriptcentral.com/). Instructions can be found at: http://www.comsoc.org/pubs/commag/sub_guidelines.html. For further information contact Steve Gorshe, Associate Editor-in-Chief (steve_gorshe@pmc-sierra.com). All submissions will be peer reviewed.

BPA INTERNATIONAL

64 **HANDOVER IN MULTIHOP CELLULAR NETWORKS**
The authors introduce various handover scenarios in multihop cellular networks. In addition, they present handover schemes where relay stations are located either inside a cell or on the boundary between two adjacent cells and investigate the effects of the deployment position of relay stations on handover performance. The simulation results show that multihop cellular networks for both deployment scenarios can achieve 90 percent throughput increase over single-hop cellular networks. The results also show that the overall throughput of the multihop cellular networks with relay stations inside a cell is higher than for those with relay stations on the boundary between two adjacent cells, whereas the opposite is observed for the throughput of cell boundary users.
SUNGHYUN CHO, EDWARD W. JANG, AND JOHN M. CIOFFI

74 **SECURING WIRELESS IMPLANTABLE DEVICES FOR HEALTHCARE: IDEAS AND CHALLENGES**
Implantable devices hold great potential for pervasive healthcare, enabling the identification, monitoring, and treatment of patients regardless of their location. To realize this goal, these devices must be able to communicate wirelessly with external devices. However, wireless communication presents many vulnerabilities: an attacker can eavesdrop on transmitted information, use implanted devices to track patients, or spoof an implanted device. The authors detail the threats that face wireless implantable devices, surveys the work addressing these threats, and identifies open issues for future research.
KRIANGSIRI MALASRI AND LAN WANG

81 **POLICY MANAGEMENT FOR ENUM SYSTEM ENABLING PRIVACY AND SECURITY**
ENUM (Telephone Number Mapping) is a key enabler in the convergence between IP-based networks and the traditional PSTNs that may result in additional complexity in commercial relationships and regulation of the telecommunications sector. In particular, the current ENUM may significantly increase the risk of unscrupulous use of the information managed. The aim of this article is to describe a new functional reference model for ENUM, and provide some new requirements that enable user privacy and security.
FRANCESCO SILLETTA AND PAOLO DE LUTIIS

86 **AN OVERVIEW OF DOWNLINK RADIO RESOURCE MANAGEMENT FOR UTRAN LONG-TERM EVOLUTION**
Radio resource management algorithms ranging from bearer admission control to semi-persistent and dynamic packet scheduling, fast link adaptation, and transmission control of multi-antenna transmission modes are addressed in this article for UTRAN long-term evolution. It is demonstrated how the collocation of the RRM algorithms at the base station with easy access to air interface measurements offers opportunities for efficient cross-functional optimization between layers 1, 2, and 3.
KLAUS INGEMANN PEDERSEN, TROELS EMIL KOLDING, FRANK FREDERIKSEN, ISTVAN ZSOLT KOVACS, DANIELA LASELVA, AND PREBEN ELGAARD MOGENSEN

94 **A NEW CLASS OF FLOATING-POINT DATA FORMATS WITH APPLICATIONS TO 16-BIT DIGITAL-SIGNAL PROCESSING SYSTEMS**
Sixteen-bit programmable digital signal processors suffer from inadequate dynamic range and noise performance due in part to the use of standard data formats with few bits available for numeric precision. A solution to this problem is developed that involves the use of irregular data formats. A new class of irregular floating-point formats is developed. A specific format is derived from this class that provides greater dynamic range and improved noise performance for 16-bit DSP applications.
MANUEL RICHEY AND HOSSEIN SAIEDIAN

102 **UNDERSTANDING INTERFERENCE AND CARRIER SENSING IN WIRELESS MESH NETWORKS**
Wireless mesh networks aim to provide high-speed Internet service without costly network infrastructure deployment and maintenance. The main obstacle in achieving high-capacity wireless mesh networks is interference between the mesh links. The authors analyze the carrier sensing and interference relations between two wireless links and measure the impact of these relations on link capacity on an indoor 802.11a mesh network testbed. They show that asymmetric carrier sensing and/or interference relations commonly exist in wireless mesh networks, and we study their impact on the link capacity and fair-channel access. In addition, they investigate the effect of traffic rate on link capacity in the presence of interference.
JEONGKEUN LEE, SUNG-JU LEE, WONHO KIM, DAEHYUNG JO, TAEKYOUNG KWON, AND YANGHEE CHOI

The President's Page 6
Conference Calendar 10
New Products 12
Global Communications Newsletter 13
Advertisers Index 112

FREE ONLINE TUTORIAL

from IEEE Communications Society

www.comsoc.org/freetutorials

Tutorials Now™

LTE TECHNOLOGY AND LTE TEST; a desk side chat

Long Term Evolution (LTE) technology is the next step in the evolution of UMTS cellular networks. "LTE technology and LTE test; a desk side chat" covers the motivation for introducing LTE, the technology basics including physical layer parameterization, radio procedures, MIMO, and network / protocol architecture as well as typical test and measurement challenges.

TAKE THIS VALUABLE TUTORIAL NOW!

FREE ACCESS FOR A LIMITED TIME ONLY

IEEE COMMUNICATIONS SOCIETY

FREE ACCESS SPONSORED BY ROHDE&SCHWARZ

For other tutorial sponsor opportunities, please contact Eric Levine, Associate Publisher: 212-705-8920, e.levine@comsoc.org

THE PRESIDENT'S PAGE

THE PATH TO BECOMING AN IEEE COMMUNICATIONS SOCIETY CONFERENCE

DOUG ZUCKERMAN

Each year the IEEE sponsors a wide variety of technical conferences around the world through its organizational units, particularly the 38 societies. The Communications Society (ComSoc) is the second largest IEEE society and sponsors approximately 85 events (conferences and workshops) either by direct involvement or by co-sponsorship. Conferences represent a substantial proportion of our efforts as a society. The 25 ComSoc conferences that we call portfolio events bring the society prestige, advance communications research and practice, and account for roughly half of our revenue each year. Frequently, some papers from those conferences become very popular IEEE digital library downloads. This column goes behind the scenes to describe what our society does to support those conferences of interest to our members.

I am pleased to share this month's column with Fred Bauer (Director-Conference Development), Merrily Hartmann (Director- Conference Operations), and Raouf Boutaba (Director-Conference Publications).

Fred Bauer works at Xorp Inc., a startup in Silicon Valley. Previously he has worked at companies such as Nokia, PacketHop, SRI, and Intel. Over the years, Fred has been an active participant in IEEE INFOCOM and SECON conferences. Most recently he has served as General Chair of INFOCOM in 2003 and as Standing Committee Chair for SECON. Fred has also served as Technical Activities Secretary (2004-2005), Distinguished Lecturer Selection Chair (2004-2005), Director of Meetings & Conferences (2006-2007), Director of Conference Development (2008-2009), and Board of Governor's Member-At-Large (2008-2010). Fred received his Ph.D. in computer engineering from the University of California at Santa Cruz in 1996.

Merrily retired from SBC Communications, Inc. (formerly Southwestern Bell and now AT&T) in 2000 with 25 years of service. She started her career with Bell Telephone Laboratories (Indian Hill Laboratory - Naperville, IL). She has served on numerous conference organizing committees: NTC '82, GLOBECOM '86, GLOBECOM '89, GLOBECOM '93, GLOBECOM 2001, GLOBECOM 2004, RFID 2007, RFID 2008, AND RFID 2009. She also served (2006-2007) on the IEEE Communications Society's GLOBECOM/ICC Management and Strategy (GIMS) Committee, which is responsible for providing strategic guidance and management oversight of the Society's flagship conferences ICC and GLOBECOM. Merrily received her Bachelor of Sciences degree in Mathematics and Computer Sciences from the University of Illinois, Chicago.

Raouf Boutaba is Professor of Computer Science at the University of Waterloo (Canada). He has published over 300 papers in refereed journals and conferences and has two US patents. He is the recipient of the Premier's Research Excellence Award, several Best Paper Awards, as well as other recognitions from academia and industry. He received the IEEE ComSoc Harold Sobol Award for Exemplary Service to Meetings & Conferences in 2007, and the IEEE ComSoc Fred W. Ellersick Prize in 2008. Raouf is active within ComSoc in many capacities: Chair of the Kitchener-Waterloo Chapter; Chair of the Information Infrastructure TC; founding Chair of the Autonomic Communications subcommittee; Director of Conference Publications; and member of the Education Board. Previously, he served as Director of Related Societies, the first Director of Standards, Vice Chair of the Information Infrastructure TC, and a member of the On-line Content Board. Raouf is founding Editor-in-Chief of *IEEE Transactions on Network and Service Management*, and editor for other journals. He chairs the NOMS/IM steering committee, has organized several conferences, and is a ComSoc Distinguished Lecturer.

ORGANIZATION AND RESPONSIBILITIES

Prior to 2008, conferences were organized under the VP of Membership Services and had just one director — the Director of Meetings and Conferences. In 2008 the Communications Society reorganized conferences under a new VP of Conferences (Mark Karol) and divided the work of the Director of Meetings and Conferences into three directors: Conference Development (Fred Bauer), Conference Publications (Raouf Boutaba), and Conference Operations (Merrily Hartmann). The three directors are assisted in their duties by ComSoc's staff led by Bruce Worthman, Director of Conferences, Finance & Administration.

Conference Development is the forward-looking part of ComSoc's conference management. Fred, the Director of Conference Development, looks for opportunities to serve conferences and attendees. In addition, he looks to integrate those conferences identified by Technical Committees or the Board of Governors as strategic conferences to include in ComSoc's portfolio. Examples of efforts to better support conferences include automated Student Travel Grant applications, less US-centric banking practices for non-US conferences, automated VISA applications for conference authors, and access to plagiarism detection tools. An example of a conference that has joined ComSoc's portfolio events is IEEE ANTS, a new regional event in networking in India. Often the best ideas for conference improvements come from the organizers themselves, so we encourage you to come up with new ideas for conference services.

Ultimately, events that are approved and sponsored by ComSoc publish their proceedings in IEEE Xplore(r). In this respect, Raouf, the Director of Conference Publications, provides guidance for all operational aspects. One such aspect concerns intellectual property matters, providing better access to IEEE copyright regulations for conference proceedings and working with IEEE for copyright acquisition of ComSoc co-sponsored conferences and their proceedings inclusion in IEEE Xplore(r). Other operational

aspects include contribution and implementation of IEEE conference publication policies such as author no-show policy, the policy of integrated workshop proceedings inclusion in Xplore(r), or plagiarism and self-plagiarism policies. Finally, it is the responsibility of Raouf's board to review and maintain the quality of ComSoc conference publications, which involves various factors such as the timeliness of posting of conference proceedings to IEEE Xplore(r), the coverage of ComSoc conference proceedings in publication databases and citation indices, and the establishment of quality metrics to evaluate ComSoc sponsored and co-sponsored conferences. We will be happy to receive your feedback on this or any of the above aspects of conference publications.

FRED BAUER RAOUF BOUTABA MERRILY HARTMANN

Merrily, Director of Conference Operations, is responsible for day-to-day conference management of the portfolio events (listed on our conferences website at www.ieee-comsoc.org/confs) and managing the technically co-sponsored event process. There are three forms of co-sponsorship, which are governed by IEEE policies:

• Sole Sponsorship: Sole-sponsored events are organized and financed entirely by ComSoc. They are led by society volunteers, endorsed by one or more technical committees, supported by ComSoc's staff, and meet the stringent technical criteria of the IEEE. This is the preferred sponsorship option for new conferences seeking ComSoc's financial involvement. Sole sponsored ComSoc events do not apply for sponsorship each year since, as part of the ComSoc conference portfolio, the relationship is in force until either party chooses to end it.

We have recently refined the definition of an important sub-category of a sole sponsored event: the "Technical Committee (TC) workshop." While ComSoc has had successful examples of TC workshops for years, this category has lacked a crisp definition. The definition of a TC workshop is as follows: "A TC workshop is founded or initiated by a ComSoc Technical Committee (TC), is managed by the TC's leadership, has ComSoc as its sole financial sponsor, ComSoc staff provides support functions as needed, and "IEEE" precedes the workshop name."

• Technical Co-Sponsorship: Technically co-sponsored events have clearly identified non-profit financial sponsors, and at least three members of a ComSoc Technical Committee are directly and substantially involved in the Technical Program Committee activities for the event. Technical co-sponsorship is valid only for the year granted. Subsequent events in a series must re-apply for co-sponsorship as described below.

• Financial Co-Sponsorship: ComSoc financially co-sponsored events have substantial organizational, financial, and technical involvement by ComSoc volunteers and staff. This form of sponsorship is reserved for those events that combine ComSoc's financial involvement with that of one or more IEEE organizational units or non-IEEE entities. While a few portfolio events such as OFC/NFOEC and MILCOM are financially co-sponsored, the preferred form of financial sponsorship is Sole Sponsorship.

TECHNICAL CO-SPONSORSHIP PROCESS

Technical co-sponsorship of outside conferences benefits the society in three ways: ComSoc's logo appears on conferences of high technical quality; ComSoc reaches potential new members by adding those conference attendees who opt-in to ComSoc email messages; and the event's proceedings become part of IEEE's digital libraries. However, technically cosponsored events carry the biggest risk for ComSoc in terms of content, quality, and branding. Every year over 100 non-profit conferences organized outside of the Communications Society (and many times outside of IEEE) apply for technical co-sponsorship. Many of the co-sponsorship applicants have preconceived notions about the process that may or may not be consistent with the actual process.

The application process for event technical co-sponsorship has the following steps.

1. Event organizers complete the co-sponsorship application form found on the ComSoc conferences website at least six months before the event. This form includes such event details as the organizing committee contacts; the names of three ComSoc TC members with direct and substantial involvement in the Technical Program Committee; the call for papers (CFP), the website URL, and the acceptance and attendance statistics for the previous year's event.

2. Merrily meets periodically with Fred, Bruce, and Bridget Erlikh, ComSoc's Meeting and Conferences Administrator, to review and screen pending applications. Applications found to be incomplete will require follow-up communications with the organizers. Events advertising pending ComSoc co-sponsorship or displaying the ComSoc logo prior to approval are asked to remove the offending advertisement. Events in direct conflict with existing ComSoc conferences or out of the scope of ComSoc's interests are denied.

3. Completed applications that pass this step will be sent to the Technical Committees (TCs) for review. Each TC independently reviews the application to determine if at least three of its members are directly and substantially involved in the Technical Program Committee activities and if it wishes to endorse the application. This review cycle typically takes 2-4 weeks, but on some occasions it can take longer.

4. Merrily reviews the results from the Technical Committees and accepts or denies the application for co-sponsorship based on the TC's recommendations.

5. Bridget notifies event organizers via email when a decision is reached. If accepted for co-sponsorship, an event may display the ComSoc logo and list ComSoc as a technical co-sponsor on their website and other materials. Note that events supported by either ComSoc local chapters or sister societies must also go through the same process each year for technical co-sponsorship.

In the vast majority of cases, the above steps complete the process. On occasion, an event that has been denied will request a rebuttal review. Organizers may resubmit their application providing relevant new information. Organizers who resort to the rebuttal process are urged to make every effort to ensure their event is worthy of co-sponsorship, particularly by directly and substantially involving TC members in their Technical Program Committee. This rebuttal application takes the same course described above. After the rebuttal review is complete, the Director of Conference Operations' decision is final.

CONTINUING TO SERVE

The members of ComSoc can be justifiably proud of the breadth and depth of our conference portfolio. Each represents a significant amount of effort by our members, maintains and enhances the society's reputation for technical excellence, and is a substantial investment in furthering our technical field. We look forward to continuing to serve you, the members of ComSoc, and the conferences you find worthy of our efforts.

Hot Topics
Global Experts
Online Tutorials at Your Own Pace
FREE 5-Minute Preview

www.comsoc.org/tutorialsnow

Tutorials Now™

Newly Posted Online Tutorials by Global Experts

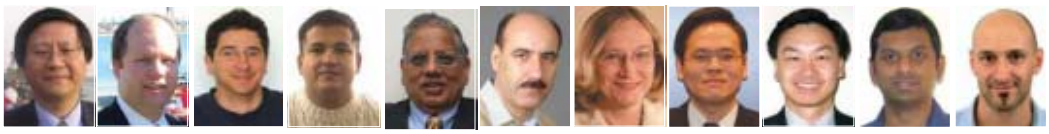

Hsiao Hwa Chen, Ilir Progri, Nicola Marchetti, Muhammad Rahman, Dharma Agrawal, Behrouz Farhang, Celia Desmond, Benny Bing, Thomas Chen, Harish Viswanathan, Mario Baldi

Wolfgang Kellerer, Gerald Kunzmann, Stefan Zöls, Leonid Kazovsky, Andreas Eberhart, Mischa Dohler, Hamid Aghvami, Eldad Perahia, Robert Stacey, Constantinos Papadias, Angel Lozano, Stefano Bregni

- **The Next Generation CDMA Technology**
  Hsiao-Hwa Chen, National Cheng Kung University
- **Indoor Geolocation Systems**
  Ilir F. Progri, Giftet
- **Future Gigabits Systems: Towards Real 4G and Cogninitve Radios**
  Nicola Marchetti and Muhammad Imadur Rahman
- **IP Multimedia Substem (IMS): Evolution to New Capabilities**
  Vijay K. Varma
- **Design and Performance Issues in Wireless Mesh Networks**
  Dharma P. Agrawal, University of Cincinnati
- **Signal Processing Techniques for Spectrum Sensing and Communications in Cognitive Radios**
  Behrouz Farhang-Boroujeny (University of Utah)
- **Project Management for Telecommunications Projects - Ensuring Success**
  Celia Desmond, World Class Telecommunications
- **Emerging Technologies in Wireless LANs: Theory, Design, Deployment**
  Benny Bing, Georgia Institute of Technology
- **Web Security**
  Thomas M. Chen, Southern Methodist University
- **Next Generation Cellular Networks: Novel Features & Algorithms**
  Harish Viswanathan, Alcatel-Lucent
- **MPLS - the importance of offering the right solution at the right moment**
  Mario Baldi, Politecnico di Torino
- **WiMax: Mobilizing the Internet**
  Benny Bing, Georgia Institute of Technology
- **Peer-to-Peer Technologies for Next Generation Communication Systems – Basic Principles and Advanced Issues**
  Wolfgang Kellerer, DoCoMo Communications Laboratories Europe, Gerald Kunzmann, Technische Universität München, & Stefan Zöls, Technische Universität München
- **Broadband Fiber Access**
  Prof. Leonid G. Kazovsky, David Gutierrez, Wei-Tao Shaw & Gordon Wong, Stanford University
- **Modern Web Applications with Ajax and Web 2.0**
  Andreas Eberhart, HP Germany
- **Wireless Cooperative Communication Networks**
  Mischa Dohler, France Telecom R&D & Hamid Aghvami, Kings College London
- **IEEE 802.11n: Throughput, Robustness, and Reliability Enhancements to WLANs**
  Eldad Perahia & Robert Stacey, Intel Corporation
- **MIMO Systems for Wireless Communications**
  Constantinos Papadias, Athens Information Technology; Angel Lozano, Bell Labs
- **Synchronization of Digital Telecommunications Networks**
  Stefano Bregni, Politecnico di Milano

Tutorial cost: Member US$200, Nonmember US$250
Tutorials that have been available online for over one year are priced at US$50 for Communications Society members.
Tutorials contain the original visuals and a voice-over by the presenters.
Length of Each Tutorial: 2.5 to 5 hours. **Number of Slides:** 78 to 477
High speed Internet connectivity suggested; PCs only. email: society@comsoc.org

69 hot topics to choose from!
Take a FREE 5-Minute Preview Now!
For complete tutorial titles, or to register
www.comsoc.org/tutorialsnow

IEEE COMMUNICATIONS SOCIETY
www.comsoc.org

CONFERENCE CALENDAR

2009

JULY

● ***NGI 2009 - 5th EURO-NGI Conference on Next Generation Internet Networks, 1-3 July***
Aveiro, Portugal. Info: http://www.ngi2009.eu

■ ***IEEE WiMAX 2009 - 2009 IEEE Mobile WiMAX Symposium, 9-11 July***
Napa, CA. Info: chenkc@cc.ee.ntu.edu.tw

■ ***IWQoS 2009 - Int'l. Workshop on Quality of Service 2009, 13-15 July***
Charleston, NC. Info: http://iwqos09.cse.sc.edu

● ***IEEE Policy 2009 - Int'l. Symposium on Policies for Distributed Systems and Networks, 20-22 July***
London, U.K. Info: http://IEEE-Policy.org

● ***NDT 2009 - 1st Int'l. Conference on Networked Digital Technologies, 28-31 July***
Ostrava, Czech Republic. Info: http://arg.vsb.cz/NDT2009/

AUGUST

● ***ICCCN 2009 - 18th Int'l. Conference on Computer Communications and Networks, 2-6 Aug.***
San Francisco, CA. Info: http://www.icccn.org/icccn09/

● ***ICASID 2009 - Int'l. Conference on Anti-Counterfeiting, Security and Identification in Communication, 20-22 Aug.***
Hong Kong, China. Info: http://www.ieee.org.hk/asid2009/

● ***ITU K-IDI 2009 - ITU-T Kaleidoscope 2009 — Innovations for Digital Inclusion, 31 Aug.-1 Sept.***
Mar Del Plata, Argentina. Info: http://www.itu.int/ITU-T/kaleidoscope2009/

■ ***IEEE EDOC 2009 - 13th IEEE Int'l. Enterprice Computing Conference, 31 Aug.-4 Sept.***
Auckland, New Zealand. Info: https://www.se.auckland.ac.nz/conferences/edoc2009/

SEPTEMBER

● ***ISWCS 2009 - Int'l. Symposium on Wireless Communication Systems, 7-10 Sept.***
Siena, Tuscany, Italy. Info: http://www.iswcs.org/iswcs2009/

● ***ICUWB 2009 - 2009 IEEE Int'l. Conference on Ultra Wideband, 9-11 Sept.***
Vancouver, BC, Canada. Info: http://www.ICUWB2009.org

● ***IEEE LATINCOM 2009 - IEEE Latin America Communications Conference, 10-11 Sept.***
Medellin, Antioquia, Colombia. Info: http://www.ieee.org.co/~comsoc/latincom

● ***APNOMS 2009 - 12th Asia-Pacific Network Operations and Management Symposium, 23-25 Sept.***
Jeju, Korea. Info: http://www.apnoms.org/2009/

● ***WiCOM 2009 - 2009 Int'l. Conference on Wireless Communications, Networking and Mobile Computing, 24-26 Sept.***
Beijing City, China. Info: http://www.wicom-meeting.org/

Softcom 2009 - Int'l. Conference on Software, Telecommunications and Computer Networks, 24-26 Sept.
Split-Hvar-Korcula, Hungary. Info: http://www.fesb.hr/SoftCOM

OCTOBER

APCC 2009 - 15th Asia-Pacific Conference on Communications, 8-10 Oct.
Shanghai, China. Info: http://www.apcc2009.sjtu.edu

■ ***ATC 2009 - 2009 Int'l. Conference on Advanced Technologies for Communications, 12-14 Oct.***
Hai Phong, Vietnam. Info: http://ww.atc09.org

● ***ICFIN 2009 - 1st Int'l. Conference on Future Information Networks, 14-17 Oct.***
Beijing, China. Info: http://conference.bjtu.edu.cn

■ ***MILCOM 2009 - 2009 IEEE Military Communications Conference, 16-21 Oct.***
Boston, MA. Info: http://www.milcom.org

■ Communications Society sponsored or co-sponsored conferences are indicated with a square before the listing; ● Communications Society technically co-sponsored or cooperating conferences are indicated with a circle before the listing. 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.curran@comsoc.org; fax: +1-212-705-8999. Items submitted for publication will be included on a space-available basis.

Call for Papers, Tutorials, Professional Studies & Workshops

ICCT 2010

The STS International Conference on Communication Technologies

Sponsored by the Saudi Telecommunications Society (STS)

"Toward Advanced Telecommunication Industries for Activating the Knowledge Economy"

King Saud University, Riyadh, Kingdom of Saudi Arabia, January 18-20, 2010

Contributories are solicited in but limited to the following:

- Telecom markets liberalization: *Competition and its influence on the local economy*
- Telecom basic and value-added services: *Enhancing performance and reducing cost*
- Future telecom technologies
- Frequency spectrum management plans
- The role of ICT n the knowledge-based economy

Papers of no more than 6 single-spaced (10 points font) pages including Figures & Tables following IEEE guidelines for preparing papers, should be submitted electronically to: stsconference2010@gmail.com before 20 / 8 / 2009. For more information visit the STS website.

CONFERENCE CALENDAR

● ***LANOMS 2009 - 6th Latin American Network Operations and Management Symposium, 19-21 Oct.***
Punta del Este, Uruguay. Info: http://www.lanoms.org/2009/

● ***DRCN 2009 - 7th Int'l. Worksho on the Design of Reliable Communications Networks, 26-29 Oct.***
Washington, DC. Info: http://www.drcn.us/

● ***ICIN 2009, 26-29 Oct.***
Bordeaux, France. Info: http://www.icin.biz/

NOVEMBER

● ***AH-ICI 2009 - First Asian Himalayas Int'l. Conference on Internet, 3-5 Nov.***
Kathmandu, Nepal. Info: http://www.ah-ici.org/ah-ici2009

● ***COMCAS 2009 - 2009 Int'l. Conference on Microwaves, Communications, Antennas and Electronic Systems, 9-11 Nov.***
Tel Aviv, Israel. Info: http://www.comcas.org

WCSP 2009 - Int'l. Conference on Wireless Communications and Signal Processing, 13-15 Nov.
Nanjing, Jiangsu, China. Info: http://www.icwcsp.org

● ***IEEE-RIVF 2009 - 2009 IEEE-RIVF Int'l. Conference on Computing and Communication, 13-17 Nov.***
Danang, Vietnam. Info: http://www.rivf.org

■ ***IEEE GLOBECOM 2009 - IEEE Global Communications Conference, 30 Nov.-4 Dec.***
Honolulu, HI. Info: http://www.ieee-globecom/2009

● ***GEOSS Workshop 2009 - GEOSS Workshop XXXII, 30 Nov.***
Honolulu, HI. Info: http://www.ieee-earth.org/Conferences/GEOSSWorkshops

DECEMBER

● ***ICICS 2009 - 7th Int'l. Conference on Information, Communications and Signal Processing, 7-10 Dec.***
Macau, China. Info: http://www.icics.org/2009/

■ ***ANTS 2009 - 2009 3rd Int'l. Symposium IEEE Advanced Nteworks and Telecommunications Systems, 14-16 Dec.***
New Delhi, India. Info: http://www.ieee-ants.org

● ***HEALTHCOM 2009 - Int'l. Conference on eHealth Networking, Applications and Services, 16-18 Dec.***
Sydney, NSW, Australia. Info: http://www.inext.uts.edu.au/healthcom09

2010

JANUARY

■ ***IEEE CCNC 2010 - IEEE Consumer Communications and Networking Conference, 9-12 Jan.***
Las Vegas, NV. Info: http://www.ieee-ccnc.org/

FEBRUARY

WSA 2010 - Int'l. ITG Workshop on Smart Antennas, 23-24 Feb.
Bremen, Germany. Info: http://www.ant.uni-bremen.de/wsa2010

MARCH

■ ***IEEE ISPLC 2020 - 14th Int'l. Symposium on Power Line Communications and Its Applications, 28-31 March***
Rio de Janeiro, Brazil. Info: http://www.ieee-isplc.org

APRIL

■ ***IEEE DYSPAN 2010 - IEEE Int'l. Symposium on Dynamic Spectrum Access Networks, 6-9 April***
Singapore. Info: http://www.ieee-dyspan.org

■ ***IEEE/IFIP NOMS 2010 - IEEE/IFIP Network Operations and Management Symposium, 19-23 April***
Osaka, Japan.

CONEC IP67 RATED CONNECTORS
When the Going Gets Tough.

Take advantage of a great choice of CONEC industrial interface connectors with IP67 protection. The ideal solution for rough environments.

USB 2.0 and RJ 45 Industrial Ethernet connection systems available as plastic or full metal versions for heavy duty applications.

CONEC offers a broad range of plug and receptacle housing kits, cable assemblies and accessories.

USB 2.0 and RJ 45 Connector Systems from CONEC: quality keeps connected!

CONEC®
TECHNOLOGY IN CONNECTORS™
Garner, NC, 27529
Tel. + 1 919 460 8800
Fax + 1 919 460 0141
E-mail info@conec.com
www.conec.com

NEW PRODUCTS

10/40/100 GBPS ENCRYPION IN FPGAS

Achronix Semiconductor Corporation

Achronix FPGAs are packed full of world-class technology, including the unmatched 1.5 GHz picoPIPE™ acceleration technology and world-class 10 Gbps SerDes. Achronix FPGAs consist of a traditional synchronous frame surrounding a picoPIPE logic fabric.

Leveraging picoPIPE technology for the FPGA fabric has two major benefits compared with a traditional fabric. First, when a design is mapped into the picoPIPE fabric, it can be pipelined automatically. Any circuit that will benefit from manual pipelining can automatically achieve increased performance when mapped into picoPIPE technology. Second, the need for global clocks has been removed. picoPIPE technology uses very fine grained local clocking to ensure data can travel at very high speeds.

The high-performance fabric is complemented by the Achronix FPGA frame, which includes configurable I/Os, SerDes, clocks, and PLLs. The frame provides all off-chip interfaces and forms the boundary between the picoPIPE core and these interfaces.

Containing up to 40 lanes of 10 Gbps SerDes, the Speedster™ family offers the highest SerDes bandwidth available in an FPGA today, and enables an aggregate SerDes bandwidth of 800 Gbps.

Speedster FPGAs include four embedded DDR1/2/3 controllers, each offering up to 72 bits of data at 1066 Mbps. Having embedded controllers saves valuable programmable resources and alleviates design challenges related to DDR implementation. The DDR controllers are fully by-passable so the pins can be used as general I/O if the DDR controllers are not needed.

Achronix non-SerDes I/Os are among the fastest in the industry. With performance of up to 1000 Mbps (or 1066 Mbps for DDR3) Achronix I/Os ensure data can be shifted into and out of the picoPIPE fabric fast enough to keep up with the 1.5 GHz internal throughput.

www.achronix.com

HIGH PERFORMANCE, INTEGRATED RF QUADRATURE/LOW IF TRANSCEIVER

CML Microcircuits

The CMX991 is a single-chip, high performance, RF transceiver that provides the core functions required to implement a full-featured radio transmitter and receiver. It operates from 100MHz to 1GHz and its I/Q architecture supports multiple modulation types and bandwidths with a single radio design.

The half-duplex CMX991 integrates Tx modulators, Rx demodulators, IF PLL and IF VCO subsystems to minimise the external circuits needed when implementing a complete transceiver.

User-selected modes suit different application requirements.

The Tx path includes an I/Q modulator to accurately generate modulation at the IF frequency, which may then be translated to the final RF frequency by an integrated image-reject up-converter system.

The I/Q modulator IF output is also made available for conversion to RF via external circuits, if desired.

The Rx path includes an integrated 1st Rx mixer having two outputs to support two external 1st IF filter choices, then an integrated 2:1 input mux followed by VGA and wideband signal level measurement functions, to support AGC implementation.

The 1st IF signal is then either I/Q demodulated to Zero-IF or mixed to a Low IF output.

The CMX991 provides differential and single-ended Rx output options and differential amplifiers for flexible signal conditioning.

The CMX991 can be used where highly linear modulations are being used, for example: applications such as TETRA, where a typical transmitter solution would include the CMX998 Cartesian Feedback Transmitter.

www.cmlmicro.com

LTC4266: HIGH EFFICIENCY QUAD PSE CONTROLLER ADDRESSES IEEE 802.3AT POE+ STANDARD

Linear Technology Corporation

Linear Technology Corporation introduces the LTC4266, a 4-port Power over Ethernet (PoE) controller for Power Sourcing Equipment (PSE) required to provide IEEE 802.3at (25.5W) or proprietary higher power levels. Next-generation PoE applications call for more power to support demanding features, while increasing power efficiency in an effort to be more "green" and reduce costs. The LTC4266 provides up to 100W over 4-pair Ethernet cabling and is fully compliant with the new IEEE 802.3at PoE+ standard and backward compatible with the prior IEEE 802.3af PoE standard. To help conserve power, the LTC4266 delivers the lowest-in-industry heat dissipation by using low RDS(ON) MOSFETs and 0.25Ohm sense resistors, eliminating the need for expensive heat sinks and providing a more robust PSE solution.

The LTC4266 is suitable for a wide variety of PSE applications, including next-generation switches, routers, hubs and midspans. Users will appreciate the extremely low power dissipation which makes thermal design significantly easier than when designing with PSEs that integrate more fragile, normally higher RDS(ON), MOSFETs. Up to four independent PSE channels are available. IEEE 802.3at powered devices (PDs) can request up to 25.5W of power by responding to 2-event classification from the LTC4266, confirming that the PD is indeed a high power device. PD discovery is accomplished using a proprietary dual-mode, four point detection mechanism that ensures the best immunity from false PD detection. Advanced power management also includes prioritized fast shutdown, 14.5-bit per-port voltage and current readback, 8-bit programmable current limits and 7-bit programmable overload current limits. A 1MHz I2C interface allows a host controller to digitally configure the IC or query port readings.

The LTC4266 is suitable for a wide variety of PSE applications, including next-generation switches, routers, hubs and midspans. Users will appreciate the extremely low power dissipation which makes thermal design significantly easier than when designing with PSEs that integrate more fragile, normally higher RDS(ON), MOSFETs. Up to four independent PSE channels are available. IEEE 802.3at powered devices (PDs) can request up to 25.5W of power by responding to 2-event classification from the LTC4266, confirming that the PD is indeed a high power device. PD discovery is accomplished using a proprietary dual-mode, four point detection mechanism that ensures the best immunity from false PD detection. Advanced power management also includes prioritized fast shutdown, 14.5-bit per-port voltage and current readback, 8-bit programmable current limits and 7-bit programmable overload current limits. A 1MHz I2C interface allows a host controller to digitally configure the IC or query port readings.

The LTC4266 is offered in commercial and industrial temperature grades, and is available in RoHS-compliant 5mm x 7mm QFN-38 or SSOP-36 packages. The LTC4266 is priced starting at $6.95 each in 1,000 piece quantities and is available today in production quantities. The LTC4266 provides an upgrade path from Linear's existing PSE controllers.

www.linear.com

Global
Communications
Newsletter
July 2009

Digital Romania: Broadband for All Romanians? Not Yet!

By Nicolae Oaca, Romania

Romania has one of the poorest broadband penetration levels in the European Union, and there is no "Internet for all Romanians" plan yet. In November 2008, just a few weeks before elections, the ministry of communications presented a plan aimed at covering 40 percent of households (the EU average in 2007!) by 2015 at 2 Mb/s, investing €1.8 billion from the state budget and attracting €383 million from the EU. In March 2009 the new government launched a public debate on a new plan aimed at providing 1 Mb/s access to 80 percent of homes by 2015 and spending only €1.25 billion from the state budget.

Forty-seven percent of Romanians live in rural areas, which does not compare favorably with the EU average of less than 10 percent, while the income in rural areas is very low; most families' incomes come from agriculture based on very small bits of land. In fact, there are two Romanias: an urban one with developed communications networks and a good IT culture, and a rural Romania with a very poor infrastructure and almost without IT culture. Romania's GDP or average wage is one of lowest in the EU; hence, huge efforts are needed to connect all Romanians to broadband.

Digital Romania's Costs

In March 2002, at the last census in Romania, there were 21.6 million citizens and 7.32 million households, out of which 3.36 million were in rural areas. By the end of 2008, 27 percent more households, or 2 million, were connected to the Internet, resulting in about 5.4 million homes connected. In rural areas only 8 percent of houses were connected. To make a cost evaluation, one could consider Greece's or Ireland's costs of €1000/house or Alcatel Lucent estimates for connecting rural houses of €6000/house by optical fiber (the Broadband World Forum Europe, September 30, 2008), resulting in €5.4 billion or €32.4 billion, respectively. The real cost of Digital Romania would lie between these two figures.

Financing Digital Romania

€32.4 or even €5.4 billion is a huge amount of money, if one considers Romania's GDP: about €140 billion in 2008. Financing Digital Romania would be a real problem, and one should take into consideration all the stakeholders: government, local communities, service providers, ministry of communications, the regulatory body, the EU, and so on.

The Romanian government should contribute, first, with telecom revenues: sale of telecom assets (RomTelecom's IPO, telcos privatization: Radiocomunicatii, Telecomunicatii CFR, Teletrans), resulting in about €1 billion, license sales (WiMAX, etc.), and the yearly fees charged by the regulatory body for usage of resources (frequencies, numbers, etc.). In addition, the state budget should provide financing, and Romania should borrow from financial institutions.

An important part of financing should come from the EC. European telecom commissioner Viviane Reding, who launched "Broadband for All Europeans," is fighting to finance it. Accordingly, in January 2009, the EC announced €1 billion investments in broadband for rural areas. Accessing this fund should take into consideration the population to be covered in rural areas rather than the current distribution key for the European Agricultural Fund for Rural Development, as the EC announced. The EC has often highlighted Romania's lack of Internet access, so they should recognize the Romanian reality by allotting an important amount of money to solve the problem.

Service providers should also contribute to Digital Romania: financing, deploying, and operating networks via public-private partnerships. Industry regulations should take into account that these new networks are under the universal service obligation and provide incentives.

In many rural villages there are projects aimed to build roads using EU financing, which alongside roads means ducts for water, residual water, and so on. All these projects should include the communications network too: ducts for/with cables (optical, copper) with output to every house. The extra cost is very low compared to the total cost and could solve the problem in these areas.

Project Length

Digital Romania is the *raison d'etre* of the ministry of communications and could be accomplished during its term (four years) in order for the new government not to pass it on to the next government. A first phase could take into consideration, as in Ireland, 100 percent coverage at 1 Mb/s by 2010 using wireless access building on the existing GSM and CDMA networks. In the longer term the plan should be to cover all households with broadband access (50 or 100 Mb/s) via NGNs.

The Romanian government is considering investments in infrastructures in is fight against the economic crisis, but is not considering telecom infrastructures yet. Digital Romania is a chance for Romania to accelerate economic growth, to fight the economic crisis, and to invest in the country's economic future.

The eBario Project: A Rural ICT Internet Access Initiative in Malaysia

By Alvin Yeo, Poline Bala, Peter Songan and Khairuddin Ab Hamid, Universiti Malaysia Sarawak, Malaysia

Malaysia aspires to be a knowledge-based society by 2020; therefore, the Malaysian government has placed much emphasis in providing access to information via the provisioning of information and communications technologies (ICTs) to both urban and remote communities. The eBario project is one such project that has successfully engaged a remote community to employ ICTs and is sustaining the telecentre through revenue generated from the telecentre's activities.

Bario is a remote rural community in Sarawak, East Malaysia, and is located close to the Malaysia-Indonesia (Kalimantan) border (Fig. 1). The only practical way to get to Bario is a one-hour flight on a 19-seater Twin Otter plane. There are no roads leading to Bario, and a land journey requires a river journey and a 14-day trek across forested mountains.

There are 12 longhouses in Bario that are homes to 1000 people, the majority of whom are Kelabits, one of the smallest ethnic groups in Sarawak. They are mainly farmers, growing the famous fragrant Bario rice. Bario was selected because of its extreme isolation. It has only basic infrastructure and no 24-hour electricity supply; water is available through gravity-fed systems. Communications are conducted using radio calls, and by passing messages to departing passengers and getting messages from passengers arriving at the airport. Hence, if one could successfully implement such a project in Bario, one can do so anywhere.

Figure 1. *Location of Bario.*

Figure 2. *Lodge owners using a computer at the telecentre.*

The eBario Project

The eBario project was conceived in 1999 as a research project by Universiti Malaysia Sarawak, with the goal of exploring how the deployment of ICTs can bring about social and economic development within remote communities in Sarawak.

A distinctive characteristic of this project is its emphasis on people, organization, and processes, rather than technology per se. A steering committee comprising representatives from different groups of the community (e.g., women and youth) was set up. Since major decisions were made in this committee, it captures very closely and addresses the actual needs of the community. In addition, a good rapport was established; the community was engaged and involved throughout the whole project. Computers were introduced at the school, and training was provided for teachers as well as the community. Access to the Internet was then provided through very small aperture terminals (VSATs) as this was (and still is) the most practical means given the region's geography. Lastly, computers and Internet access were made available to the community through the telecentre. Two VSATs were deployed, one for the school and another for the telecentre.

Solar panels were used to power the telecentre because of the prohibitive costs associated with diesel powered generators and the logistics of transporting the diesel to this remote hinterland.

A computer laboratory was established in 2001 in the secondary school with 16 computers and Internet access. A telecentre equipped with 10 computers and Internet access was established in 2002. Computer literacy training was conducted based on the training-of-the-trainer approach. Consequently, there is increased computer literacy among the students, teachers, and community.

The community is now able to communicate with the outside world using telephone and email. This is of immense importance, especially during emergencies. At the telecentre, the main users are the lodge operators who use email to liaise with their potential tourists directly to confirm bookings online (Fig. 2). Some use the Internet to promote their trekking services (http://www.kelabit.net). The senior citizens use emails to keep in touch with their family members living away, albeit with the help of staff members at the telecentre. The telecentre charges for their service as a first step toward sustainability. The telecentre also provides office services such as printing, scanning, and laminating.

A bigger objective of the project is to improve the livelihood and increase the quality of life of the community. More events are being organized to leverage on the strengths of this region. The climate is relatively cool (about 20°C, which can drop to 12–15°C), and the people are some of the most hospitable in Malaysia, if not the world. The Kelabit ethnic community is rich in culture, with rituals such as dances and ornate handicrafts. Thus, the telecentre organizes events like the Annual Slow Food Festival and the eBario Knowledge Fair, a conference whose aim is to provide a network platform for academicians and practitioners in the area of bridging the digital divide. The telecentre also works with government agencies on organizing training and seminars in Bario. As event organizers and providing tourism packages, the telecentre receives funding and a commission for their services (i.e., a consistent income contributing toward sustainability of the telecentre). An unexpected outcome of these events is that an even closer-knit community in Bario has been established, in that more people are now involved in the organization of these events, especially women.

In addition to the increased ICT awareness and usage by the community, another obvious impact has been an increase

(Continued on Newsletter page 4)

Telecommunications Engineering and the Bologna Declaration in Spain

By Pablo Pavón Mariño,
Vice Dean of Telecommunication Engineering, Polytechnic University of Cartagena, Spain

The realization of the European Higher Education Area (EHEA) is a common aim of 46 countries in Europe. The effort is usually referred to as the Bologna Process, after the Bologna Declaration of June 1999. Its main objective is to provide European universities with the tools for promoting citizens and knowledge mobility, breaking the obstacles caused by historical diversity among the European higher education systems. A central objective is the adoption of easily readable and comparable degrees in a scheme based on two main cycles, undergraduate and graduate. Uniformity is favored by the establishment of a common system of credits to encourage mobility among the European countries. European cooperation is also endorsed for quality assurance with a view to developing comparable criteria and methodologies.

Such noble and high principles must fight against a plethora of particular situations in the different countries. In Spain the Bologna Process is currently going through its most controversial step: the definition of the new catalog of degree plans inside the new Bologna-aware framework ordered by the government. With some exceptions like medicine and architecture, new graduate degrees in Spain are constrained to be four years long (240 ECTS), including a final degree project, and optional time spent in external practices. Postgraduate degrees last one or two years (60 to 120 ECTS). In contrast to the previous framework, no nation-wide closed catalogs of graduate and postgraduate studies exist. Therefore, novel degrees can be designed and proposed by Spanish universities. Then the proposals require the approval of a central evaluating authority (ANECA, Agencia Nacional de Evaluación de la Calidad y la Acreditación), which checks the academic adequacy of the degree plan and its feasibility according to university resources. Also, a control mechanism has been added to penalize unsuccessful ephemeral studies. ANECA is in charge of auditing the degree academic results every six years, with the potential of removing from the catalog any degree in any university.

However, this general ruling of the graduate and postgraduate degree catalog has a singular exception for a set of technical degrees, based on traditional engineering professions in Spain. Telecommunications engineering is one of them, together with others like naval, mining, industrial, aeronautics, and up to eight different traditional engineering academic qualifications. In the pre-Bologna framework, technical studies in those professions are separated into two levels: technical engineers (three years) and engineers (five years). Technical engineering degrees provide a faster introduction into the job market, while Engineering degrees comprise deeper and more generalized studies. The migration of technical studies to the Bologna framework has been designed in a peculiar manner. Each old technical engineering degree (with some exceptions) has been converted into a new regulated graduate degree. Regulated degrees are peculiar as their structure is defined by a degree template, which also enumerates a minimum set of student learning outcomes associated with the degree plan. Universities willing to incorporate a new technical degree "X" into its academic offer must propose a degree plan compliant with the degree template "X". Naturally, ANECA is responsible for validating this compliance. A similar strategy is in place for the old five-year engineering degrees, which are now mapped to the concatenation of a regulated graduate degree in the particular profession and a regulated postgraduate degree, which must also be compliant with a degree template.

The grounds for the existence of degree templates for the

(Continued on Newsletter page 4)

Highlights from ONDM 2009

By Admela Jukan, General Chair, ONDM 2009, Germany

The 13th IEEE/IFIP International Conference on Optical Network Design and Modelling (ONDM 2009) was held February 18–20, 2009 at the Technische Universität Carolo-Wilhelmina zu Braunschweig, Germany. The conference had the technical sponsorship of the IFIP Photonic Networking Working Group (WG 6.10) and IEEE Communications Society. Moreover, ONDM 2009 received an endorsement of two IEEE Technical Committees: Optical Networking (ONTC), and High Speed Networking (TC-HSN). As is established practice at ONDM events, the 13th instalment focused on cutting-edge state-of-the-art research in optical networking and newly emerging areas. ONDM 2009 not only encouraged submissions of research papers in the area of optical networking but also in areas relating optical networking to other areas and disciplines, such as the role of optical satellite communications and the optical network in future Internet design. Controversial ideas and approaches, and their open discussion were strongly encouraged. Moreover, ONDM 2009 contributed to community strengthening by inciting interactivity among senior and junior participants as well as among industrial and academic participants, in Europe and internationally.

The conference attracted submissions from Asia, Europe, and South and North America. All submitted papers went through a peer review process, and accepted papers were thematically distributed among six technical sessions, ranging from topics in routing and switching to integration and field trials. The keynote speech, "Optical Technology Evolution in Dynamic Market Environments," was given by Dr. Chris Glingener, CTO of ADVA Optical Networking. The Technical Program was further enriched by a number of exciting invited talks given by the most prominent optical networking academics and industrial representatives worldwide. Prof. Biswanath Mukherjee gave a talk on "Future Internet Design." Dr. Nikos Karafolas inspired his audience with a talk on optical communications in space. Andy Hougthon, an EU Officer, gave a talk on future European initiatives, entitled "Building the Future (Inter)Network in Europe: Work in Progress." Further invited speakers included Thomas Michael Bohnert (SAP), Ralf-Peter Braun (T-systems), and Dimitra Simenidou (University of Essex). The conference also presented two panels, one in the area of carrier-grade Ethernet and the other on interoperability of optical networks and the Internet. Most of the presentations and invited talks can be downloaded from the conference Web site, http://www.ida.ing.tu-bs.de/ondm2009/.

The conference social event took place in the beautiful hall of Altes Rathaus in Braunschweig (Dornse), built in the 13th

(Continued on Newsletter page 4)

ONDM 2009/*continued from page 3*

century. During the dinner event, the TPC co-chairs, Admela Jukan and Ken-ichi Kitayama, acknowledged the invaluable efforts of all TPC members. An honorable recognition was given to Biswanath Mukherjee (University of California, Davis) for his service relative to IEEE technical co-sponsorship. Special acknowledgments were extended to Josep Sole Pareta and Xavi Masip, both of the University of Catalonia, for their service to organization and coordination with IFIP. The conference awarded two papers and their authors with travel grants, one to Politecnico di Torino, Italy, and the other to Universidad Técnica Federico Santa María, Chile. The conference's best paper award was given to "New Dynamic Network Design and Provisioning Algorithms for Broadband Connection Services Considering Fairness," co-authored by Masahiro Nakagawa, Hiroshi Hasegawa, Ken-Ichi Sato, Ryuta Sugiyama, Tomonori Takeda, Eiji Oki, and Kohei Shiomoto, Japan.

The local organizers of the conference as well as TPC co-chairs sincerely hope that ONDM 2009 was a community event that will be remembered for the technical quality of its papers and the hospitality of the oldest German technical school, Technische Universität Carolo-Wilhelmina zu Braunschweig. The next ONDM event will be in Kyoto, Japan, in February 2010, hosted by the University of Osaka, under the chairmanship of Prof. Ken-ichi Kitayama. The optical networking community looks forward to the event in Japan, and says "See you in Kyoto in 2010!"

Global Communications Newsletter

www.comsoc.org/pubs/gcn

STEFANO BREGNI
Editor

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

IEEE COMMUNICATIONS SOCIETY

MARK KAROL, VICE-PRESIDENT CONFERENCES
BYEONG GI LEE, VICE-PRESIDENT MEMBER RELATIONS
NELSON FONSECA, DIRECTOR OF LA REGION
GABE JAKOBSON, DIRECTOR OF NA REGION
TARIQ DURRANI, DIRECTOR OF EAME REGION
ZHISHENG NIU, DIRECTOR OF AP REGION
ROBERTO SARACCO, DIRECTOR OF SISTER AND RELATED SOCIETIES

REGIONAL CORRESPONDENTS WHO CONTRIBUTED TO THIS ISSUE

THOMAS M. BOHNERT, SWITZERLAND (THOMAS.MICHAEL.BOHNERT@SAP.COM)
JOSÉ MARIA MALGOSA-SANAHUJA, SPAIN (JOSEM.MALGOSA@UPCT.ES)
BORHANUDIN MOHD ALI, MALAYSIA (BORHAN@ENG.UPM.EDU.MY)
NICOLAE OACA, ROMANIA (NICOLAE_OACA@YAHOO.COM)

A publication of the
IEEE Communications Society

BOLOGNA DECLARATION/*continued from page 3*

technical degrees is the manner in which professional habilitation is granted to engineers in Spain: the professional habilitation of an engineer in a specific field is automatically obtained after the successful completion of the associated engineering degree. This means that on one hand, no licensing exam is needed to obtain the professional habilitation; but on the other hand, the appropriate university degree is the only valid way to achieve it. In every profession the professional habilitation is different for three-year and five-year engineers. Therefore, it seems logical to think of different degree templates for regulated graduate and postgraduate degrees.

In the telecommunications engineering profession, four pre-Bologna three-year degrees existed in Spain: (i) Sistemas de Telecomunicación (RF engineering, signal processing, and signal propagation), (ii) Telemática (telematics and networking), (iii) Electrónico (electronics in the telecommunications field), and (iv) Imagen y Sonido (audio and video). Each one is now mapped into a four-year regulated degree, with its own degree template. One postgraduate degree in telecommunications engineering has also been defined. Currently, a small number of Spanish universities are offering the new degrees in telecommunications. Most of the universities are still preparing their degree proposals, or have their proposals under evaluation by ANECA (a process that can take about half a year). This process is expected to be completed during 2009 or the first half of 2010. In academic year 2010–2011 new students can only be enrolled in the Bologna-aware degrees. Then old and new degrees will coexist for some years, and each university will be involved in managing the required adaptation systems for students willing to move from old to new degrees. But that is another story.

THE EBARIO PROJECT/*continued from page 2*

in the number of tourists, contributing to enhanced employment opportunities and more revenue for the community. Also, more youths and their families are staying in Bario to operate the accommodation and tourist activities, thus stemming rural-urban migration. Today Bario has seven lodges and six homestays.

Acknowledgments

The eBario Project acknowledges funding received from IDRC Canada, the Demonstrators Application Grants Scheme, Ministry of Science Technology and Innovation, the cooperation of the Bario community, and the eBario Management Committee.

For more information

UNESCAP, Economic Planning Unit, Malaysia, National Institute of Public Administration. (2006). Guidebook on Developing Community e-Centres in Rural Areas: Based on the Malaysian Experience. United Nations, New York.
Repository: http://www.unimas.my/ebario, http://coeri.unimas.my
Tourism in Bario: http://www.ebario.com
Bario community: http://www.kelabit.net
Contact: Alvin Yeo alvin@fit.unimas.my

Advertise your products in IEEE Communications Magazine

Reach 55,000 Potential Customers!

Contact us today for more information or rates

Tel. +1 212 705 8920

Fax. +1 212 705 8999

Email: e.levine@comsoc.org

Website: www.comsoc.org/adv

SERIES EDITORIAL

TOPICS IN NETWORK AND SERVICE MANAGEMENT

George Pavlou *Aiko Pras*

This is the eighth issue of the series on Network and Service Management that is published twice a year, typically in April and October. This issue, however, is being published in July, having been delayed by three months for logistical reasons, and the next issue will be published in November. The series provides 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.

An important recent development in the community was the change of chairmanship in IFIP WG6.6, the working group on the Management of Networks and Distributed Systems. Prof. Raouf Boutaba of the University of Waterloo, Canada, stepped down as chair of IFIP WG6.6 after serving two terms of three years each. Under Raouf's active leadership, the flagship management conferences Integrated Management (IM) and Network Operations and Management Symposium (NOMS) became more valuable to the community, and previously self-standing events such as DSOM, MMNS, and IPOM got integrated into Manweek. In addition, the new *IEEE Transactions on Network and Service Management* (TNSM) was launched. After a call for nominations, Prof. Boutaba's tasks were collectively taken over by Dr. Aiko Pras of the University of Twente, Netherlands (Chair) and Dr. Olivier Festor of INRIA, France (Vice-Chair); you may note that Dr. Pras is also a co-editor of this series. The new chairs plan to continue the organization of successful events such as IM and NOMS, as well as smaller events such as DSOM, MMNS, and others. Key to the success of these events has been the good collaboration with the IEEE Communications Society sister organization the Committee on Network Operations and Management (CNOM). Under the leadership of the new chairs, this successful collaboration will be continued, as well as coordination and cooperation with other IFIP WGs.

Another important development for the community has been the completion of three and a half years of integration work of the European EMANICS project on Management of the Internet and Complex Services at the end of June 2009. This is a European Network of Excellence that brings together 13 research institutions active in the management of the future Internet. It encompasses work areas dealing with integration (long-term vision, virtual laboratory, and testbeds), dissemination (a new European conference, electronic dissemination, training and technology transfer, open source initiatives), and joint research activities (scalable, economic, and autonomic management). For more information, visit the project site: http://www.emanics.org/. You will also find there a newsletter with community news, events, and developments, which is published roughly three times a year. EMANICS has also established a European conference on Autonomous Infrastructure Security & Management (AIMS). The third AIMS conference is taking place in Twente, Netherlands, June 30–July 2, 2009, see http://www.aims-conference.org/2009/.

The key annual event in this area, which this year was the 11th IEEE/IFIP Integrated Network Management Symposium (IM 2009), was held June 1–5 on Long Island, New York; http://www.ieee-im.org/2009/. The second key annual event in this area is Manweek 2009, which brings together conferences and workshops such as DSOM, MMNS, and IPOM, and also MACE and NGNM. This year's Manweek will take place October 26–30 in beautiful Venice, Italy; http://www.manweek.org/2009/.

Finally, we should mention that this issue has invited articles addressing the management of the future Internet. Given that in the United States there are initiatives such as FIND and GENI, in Europe there is the Future Internet Assembly and the FIRE program, and in Japan there is NWGN, there is considerable current interest in

the future Internet and its management. Two of the four articles in this issue address future Internet management, and we expect more articles on this topic in the next issue.

We again experienced an overwhelming interest in the eighth issue, receiving 27 submissions in total (1 major revision and 26 new articles). For all the articles we got at least three and sometimes four independent reviews. We finally selected four articles, resulting in an acceptance rate of 14.8 percent. It should be mentioned that the acceptance rate for all the previous issues has ranged between 18 and 25 percent, making this series a highly competitive place to publish. We intend to maintain our rigorous review process in future issues, maintaining the high quality of the published articles.

The first article, "Network Virtualization: the Past, Present and Future" by Chowdhury and Boutaba, presents a survey of network virtualization techniques that allow multiple heterogeneous network architectures to coexist on the same physical substrate and examines relevant future challenges.

The second article, "Future Internet = Content + Services + Management" by Schoenwaelder, Fouquet, Dreo Rodosek, and Hochstatter, takes the view of a future Internet user who is more interested in services than the supporting protocols, derives relevant service and network management requirements, and discusses some of them in detail.

The third article, "With Evolution for Revolution: the FEDERICA Approach" by Szegedi, Figuerola, Campanella, Maglaris, and Cervello Pastor, discusses the FEDERICA infrastructure, which extends the virtualization capabilities of current software/hardware infrastructures in order to provide a flexible environment to host disruptive next-generation architectures.

Finally, the fourth article, "Managing Interdomain Traffic in Latin America: Facts and Perspectives" by Grampin, Gagliano, German, Castro, Masip-Bruin, and Yannuzzi, outlines one of the solutions for separation of the IP address space, the Locator/Identifier Separation Protocol (LISP), and examines its potential in terms of interdomain traffic management in Latin America.

We hope that readers of this issue 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@ee.ucl.ac.uk) is a professor of communication networks in the Department of Electronic and Electrical Engineering, University College London, United Kingdom, where he coordinates the activities of the Networks and Services Research Laboratory. 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. His research interests focus on network management, networking, and service engineering, including aspects such as traffic engineering, quality of service management, policy-based systems, autonomic networking, multimedia service control, and communications middleware. He has been instrumental in a number of European and U.K. research projects, and has contributed to standardization activities in ISO, ITU-T, and IETF. He was the technical program co-chair of the Seventh IFIP/IEEE Integrated Management Symposium (IM 2001) and the Tenth IFIP/IEEE International Conference on Management of Multimedia and Mobile Networks and Services (MMNS 2008).

AIKO PRAS (a.pras@utwente.nl) is an associate professor in the Departments of Electrical Engineering and Computer Science at the University of Twente, Netherlands, and a member of the Design and Analysis of Communication Systems Group. He received a Ph.D. degree from the same university for his thesis titled *Network Management Architectures*. His research interests include network management technologies, Web services, network measurements, and accounting. He is chairing IFIP Working Group 6.6, "Management of Networks and Distributed Systems," and is a research leader in the European Network of Excellence on "Management of the Internet and Complex Services" (EMANICS). He has also been contributing to research and standardization activities as a member of the Internet Research Task Force (IRTF) Network Management Research Group (NMRG). He was the technical program co-chair of the Ninth IFIP/IEEE Integrated Management Symposium (IM 2005) and is a Steering Committee member of the IFIP/IEEE NOMS and IM Symposia.

TOPICS IN NETWORK AND SERVICE MANAGEMENT

Network Virtualization: State of the Art and Research Challenges

N. M. Mosharaf Kabir Chowdhury and Raouf Boutaba, University of Waterloo

ABSTRACT

Recently network virtualization has been pushed forward by its proponents as a long-term solution to the gradual ossification problem faced by the existing Internet and proposed to be an integral part of the next-generation networking paradigm. By allowing multiple heterogeneous network architectures to cohabit on a shared physical substrate, network virtualization provides flexibility, promotes diversity, and promises security and increased manageability. However, many technical issues stand in the way of its successful realization. This article investigates the past and the state of the art in network virtualization along with the future challenges that must be addressed to realize a viable network virtualization environment.

INTRODUCTION

In recent years, the concept of network virtualization has attracted significant attention in the debate on how to model the next-generation networking paradigm that can replace the existing Internet. Architectural purists view network virtualization as a tool for evaluating new architectures, whereas pluralists conceive virtualization as a fundamental diversifying attribute of the next-generation architecture itself [1]. They believe that network virtualization can eradicate the so-called *ossifying forces* of the current Internet by introducing disruptive technologies [1, 2].

Network virtualization is defined by decoupling the roles of the traditional Internet service providers (ISPs) into two independent entities [2, 3]: infrastructure providers (InPs), who manage the physical infrastructure, and service providers (SPs), who create virtual networks (VNs) by aggregating resources from multiple InPs and offer end-to-end services. Such an environment will proliferate deployment of coexisting heterogeneous network architectures free of the inherent limitations of the existing Internet.

In this article we survey the past and the state of the art of network virtualization, and provide a better understanding of the key research challenges. The rest of this article is organized as follows. First, four somewhat similar ideas (virtual local area networks [VLANs], virtual private networks [VPNs], programmable networks, and overlay networks) are briefly reviewed. Next, a reference business model and a conceptual architecture of a network virtualization environment (NVE) are presented, identifying the characteristics and critical design factors to materialize it. Following this, a number of past and present research projects on network virtualization and related concepts are summarized. Finally, a detailed study of the key issues is presented emphasizing open research challenges with an objective to stoke wide interest among researchers in this field.

HISTORICAL PERSPECTIVE

The concept of multiple coexisting logical networks has appeared in the networking literature several times in the past, and can be categorized into four main classes: VLANs, VPNs, active and programmable networks, and overlay networks.

VIRTUAL LOCAL AREA NETWORK

A VLAN is a group of logically networked hosts with a single broadcast domain regardless of their physical connectivity. All frames in a VLAN bear a VLAN ID in the medium access control (MAC) header, and VLAN-enabled switches use both the destination MAC address and VLAN ID to forward frames. Since VLANs are based on logical instead of physical connections, network administration, management, and reconfiguration of VLANs are simpler than in their physical counterparts. In addition, VLANs provide elevated levels of isolation.

VIRTUAL PRIVATE NETWORK

A VPN is a dedicated network connecting multiple sites using private and secured tunnels over shared or public communication networks like the Internet. In most cases, VPNs connect geographically distributed sites of a single corporate enterprise. Each VPN site contains one or more customer edge (CE) devices that are attached to one or more provider edge (PE) routers.

Based on the protocols used in the data plane, VPNs can be classified into the following broad categories.

Layer 1 VPN — The layer 1 VPN (L1VPN) framework emerged in recent years from the need to extend layer 2/3 (L2/L3) packet switching VPN concepts to advanced circuit switching

0163-6804/09/$25.00 © 2009 IEEE

domains. It provides a multiservice backbone where customers can offer their own services, whose payloads can be of any layer (e.g., asynchronous transfer mode [ATM] and IP). This ensures that each service network has an independent address space, an independent L1 resource view, separate policies, and complete isolation from other VPNs.

Layer 2 VPN — Layer 2 VPNs (L2VPNs) transport L2 (typically Ethernet) frames between participating sites. The advantage is that they are agnostic about the higher-level protocols, and consequently more flexible than L3VPN. On the downside, there is no control plane to manage reachability across the VPN.

Layer 3 VPN — A layer 3 VPN (L3VPN) is characterized by its use of L3 protocols in the VPN backbone to carry data between the distributed CEs. There are two types of L3VPNs.

In the *CE-based* VPN approach, the provider network is completely unaware of the existence of a VPN. CE devices create, manage, and tear down the tunnels between themselves. Sender CE devices encapsulate the passenger packets and route them into carrier networks; when these encapsulated packets reach the end of the tunnel (i.e., receiver CE devices), they are extracted, and actual packets are injected into receiver networks.

In the *PE-based* approach, the provider network is responsible for VPN configuration and management. A connected CE device may behave as if it were connected to a private network.

Higher-Layer VPNs — VPNs using higher-layer (e.g., transport, session, or application) protocols also exist. SSL/TLS-based VPNs are popular for their inherent advantages in firewall and NAT traversals from remote locations. Such VPNs are lightweight, easy to install and use, and provide higher granularity of control to their users.

Active and Programmable Networks

Active and programmable networks research was motivated by the need to create, deploy, and manage novel services on the fly in response to user demands. In addition to programmability, they also promote concepts of *isolated environments* to allow multiple parties to run possibly conflicting codes on the same network elements without causing network instability.

Two separate schools of thought emerged on how to actually implement such concepts.

The Open Signaling Approach — Open signaling takes a telecommunication approach with a clear distinction between transport, control, and management planes that constitute programmable networks, and emphasizes quality of service (QoS) guarantees. An abstraction layer is proposed for physical network devices to act as distributed computing environments with well defined open programming interfaces allowing service providers to manipulate network states.

The Active Networks Approach — Active networks promote dynamic deployment of new services at runtime within the confinement of existing networks. Routers or switches in these networks can perform customized computations based on the contents of the *active* packets and can also modify them. Active networks allow the customization of network services at packet transport granularity and offer more flexibility than the open signaling approach at the expense of a more complex programming model.

Overlay Networks

An overlay network is a logical network built on top of one or more existing physical networks. The Internet itself started off as an overlay on top of the telecommunication network. Overlays in the existing Internet are typically implemented in the application layer; however, various implementations at lower layers of the network stack do exist.

Overlays do not require or cause any changes to the underlying network. Consequently, overlays have long been used as relatively easy and inexpensive means to deploy new features and fixes in the Internet. A multitude of application layer overlay designs have been proposed in recent years to address diverse issues, which include ensuring performance and availability of Internet routing, enabling multicasting, providing QoS guarantees, protecting from denial of service attacks, and content distribution and file sharing services. Overlays have also been used as testbeds (e.g., PlanetLab) to design and evaluate new architectures.

The authors in [1] point out that standard overlays falter as a deployment path for radical architectural innovations in at least two ways. First, overlays have largely been used as a means to deploy narrow fixes to specific problems without any holistic view. Second, most overlays have been designed in the application layer on top of IP; hence, they cannot go beyond the inherent limitations of the existing Internet.

Network Virtualization Environment

Unlike the existing all-IP Internet, a virtualized networking environment is a collection of multiple heterogeneous network architectures from different SPs. Each SP leases resources from one or more InPs to create VNs, and deploys customized protocols and services.

Business Model

The main distinction between the participants in the network virtualization model and the traditional model is the presence of two different roles, InPs and SPs, as opposed to the single role of the ISPs [2–4].

InP — InPs deploy and actually manage the underlying physical network resources. They offer their resources through programmable interfaces to different SPs. InPs distinguish themselves through the quality of resources they provide, the freedom they delegate to their customers, and the tools they provide to exploit that freedom.

Unlike the existing all-IP Internet, a virtualized networking environment is a collection of multiple heterogeneous network architectures from different SPs. Each SP leases resources from one or more InPs to create VNs and deploys customized protocols and services.

End users in the network virtualization model are similar to those of the existing Internet, except that the existence of multiple VNs from competing SPs provides them a wider range of choice. Any end user can connect to multiple VNs from different SPs for different services.

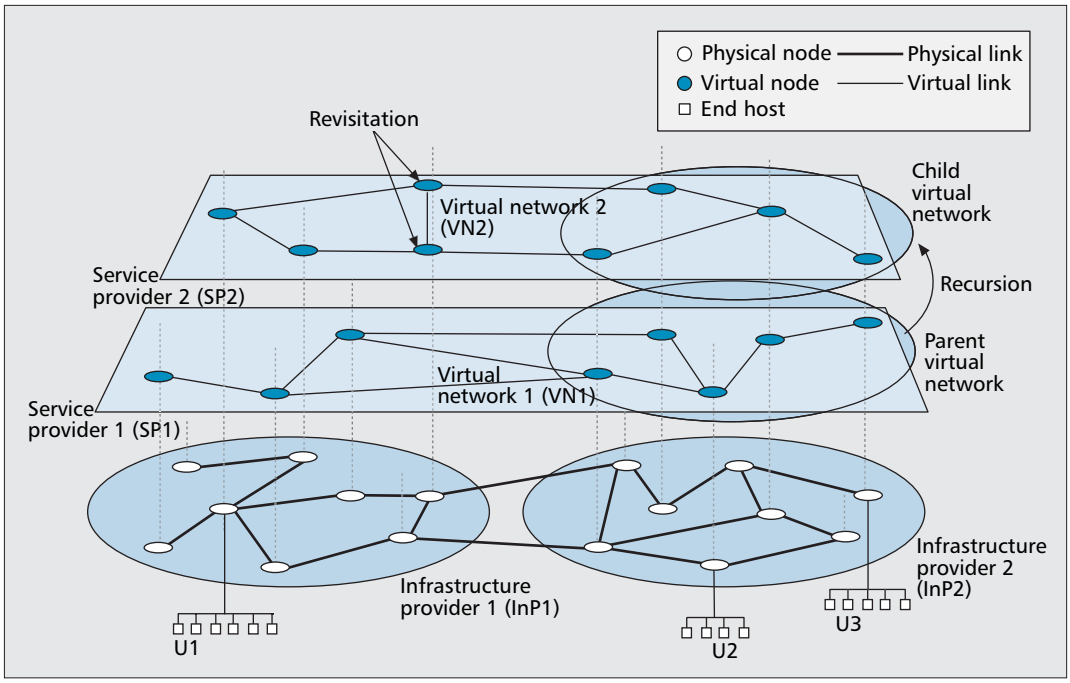

Figure 1. *Network virtualization environment.*

SP — SPs lease resources from multiple InPs to create and deploy VNs by programming allocated network resources to offer end-to-end services to end users. An SP can also provide network services to other SPs. It can also create child VNs by partitioning its resources and act as a virtual InP by leasing those child networks to other SPs (Fig. 1).

End User — End users in the network virtualization model are similar to those of the existing Internet, except that the existence of multiple VNs from competing SPs provides them a wider range of choice. Any end user can connect to multiple VNs from different SPs for different services.

Architecture

In an NVE the basic entity is a VN. A VN is a collection of virtual nodes connected together by a set of virtual links to form a virtual topology, which is essentially a subset of the underlying physical topology. Each virtual node is hosted on a particular physical node, whereas a virtual link spans over a path in the physical network and includes a portion of the network resources along the path.

Each VN is operated and managed by a single SP, even though the underlying physical resources might be aggregated from multiple InPs. Figure 1 depicts two VNs, VN1 and VN2, created by service providers SP1 and SP2, respectively. SP1 composed VN1 on top of the physical resources managed by two different InPs (InP1 and InP2), and provides end-to-end services to end users U2 and U3. SP2, on the other hand, deployed VN2 by combining resources from infrastructure provider InP1 with a child VN from service provider SP1. End users U1 and U3 are connected through VN2.

The owner of a VN is free to implement end-to-end services by deploying custom packet formats, routing protocols, forwarding mechanisms, as well as control and management planes. As mentioned earlier, end users have the choice to opt in to any VN. For example, U3 is subscribed to VN1 and VN2 managed by SP1 and SP2, respectively.

Architectural Principles

Network virtualization propounds the following principles for the next-generation networking paradigm.

Coexistence — Coexistence of multiple VNs is the defining characteristic of an NVE [1–3]. It refers to the fact that multiple VNs from different SPs can coexist together, spanning over part or full of the underlying physical networks provided by one or more InPs. In Fig. 1, VN1 and VN2 are two coexisting VNs.

Recursion — When one or more VNs are spawned from another VN creating a VN hierarchy with *parent-child* relationships, it is known as recursion as well as *nesting* of VNs [5]. Service provider SP1 in Fig. 1 leased away a portion of its allocated resources to SP2, to whom it appears simply as a virtual InP.

Inheritance — Child VNs in an NVE can inherit architectural attributes from their parents, which also means that the constraints on the parent VN automatically translate to similar constraints on its children [5]. For example, constraints imposed by InP2 will automatically be transferred to VN2 from VN1 through inheritance. Inheritance allows an SP to add value to the spawned child VNs before reselling them to other SPs [3].

Revisitation — Revisitation [6] allows a physical node to host multiple virtual nodes of a single VN. Use of multiple logical routers to handle

diverse functionalities in a large complex network allows an SP to logically rearrange its network structure and to simplify the management of a VN. Revisitation can also be useful for creating testbed networks. Figure 1 provides an example of revisitation in VN2.

Design Goals

The design goals for successfully realizing network virtualization have been addressed by different research groups. In order to materialize network virtualization, each of these design criteria should be fulfilled.

Flexibility — Network virtualization must provide freedom in every aspect of networking. Each SP should be free to implement arbitrary network topology, routing and forwarding functionalities, and customized control protocols independent of the underlying physical network and other coexisting VNs. For example, deploying source routing in today's network depends much on consensus among ISPs; in a virtualized environment, the owner of a VN should be able to offer source routing without having to coordinate with any other parties.

Manageability — By separating SPs from InPs, network virtualization will modularize network management tasks and introduce accountability at every layer of networking [3]. It must provide complete end-to-end control of the VNs to the SPs, obviating the requirement of coordination across administrative boundaries seen in the existing Internet.

Scalability — Coexistence of multiple networks is one of the fundamental principles of network virtualization. Scalability is an indispensable part of this equation. InPs in an NVE must scale to support an increasing number of coexisting VNs without affecting their performance.

Isolation — Network virtualization must ensure isolation between coexisting VNs to improve fault tolerance, security, and privacy. Network protocols are prone to misconfigurations and implementation errors. Virtualization must ensure that misconfigurations in one VN are contained within itself and do not affect other coexisting VNs.

Stability and Convergence — Isolation ensures that faults in one VN do not affect other coexisting VNs, but errors and misconfigurations in the underlying physical network can also destabilize an NVE. Moreover, instability in the InPs (e.g., routing oscillation) can lead to instability of all hosted VNs. Virtualization must ensure the stability of an NVE, and in case of any instability the affected VNs must be able to successfully converge to their stable states.

Programmability — To ensure flexibility and manageability, programmability of the network elements is an indispensable requirement. Only through programmability can SPs implement customized protocols and deploy diverse services. Two pressing questions in this respect must have satisfactory answers: "How much programmability should be allowed?" and "How should it be exposed?" A win-win situation must be found where programmability is easy, effective, and secure at the same time.

Heterogeneity — Heterogeneity in the context of network virtualization comes mainly from two fronts: first, heterogeneity of the underlying networking technologies (e.g., optical, wireless, and sensor); second, each end-to-end VN, created on top of that heterogeneous combination of underlying networks, can also be heterogeneous. SPs must be allowed to compose and run cross-domain end-to-end VNs without the need for any technology-specific solutions. Underlying infrastructures must also be capable of supporting heterogeneous protocols and algorithms implemented by different SPs. In addition, heterogeneity of end-user devices must also be taken into account.

Legacy Support — Legacy support or backward compatibility has always been a matter of deep concern while deploying any new technology. Conceptually, network virtualization can easily integrate legacy support by considering the existing Internet as just another VN in its collection of networks; but whether and how it can be done efficiently remains an open challenge.

Network virtualization must provide freedom at every aspect of networking. Each SP should be free to implement arbitrary network topology, routing and forwarding functionalities, and customized control protocols independent of the underlying physical network and other coexisting VNs.

Network Virtualization Projects

Over the years, the term *virtual network* has been used to describe different projects on VPNs, overlay networks, and active or programmable networks. But very few of them actually followed the pluralist view of network virtualization. In Table 1 we summarize the most significant past and on-going projects directly or indirectly related to network virtualization based on the following set of characteristics:

- ***Networking technology***: A handful of network virtualization prototypes have been developed for specific networking technologies with an aim to exploit unique characteristics of those networks to enable virtualization. Such projects include X-Bone for IP networks, Tempest targeting ATM networks, and the very recent GENI initiative that will be agnostic to any specific technology.
- ***Layer of virtualization***: Influenced by the existing Internet, researchers have naturally approached network virtualization in a layered manner. As a result, many projects have attempted to virtualize different layers of the network stack, starting from the physical layer (UCLP) and continuing up to the application layer (VIOLIN).
- ***Architectural domain***: Most projects have focused on particular architectural domains, which dictate the design choices made in the construction of architectures and services that can be offered on those platforms. Examples include network management (VNRMS), virtual active networks (NetScript), and spawning networks (Genesis).
- ***Level of virtualization***: To enable network virtualization, one must virtualize the nodes, links, and every other resource in the network. The level of virtualization refers to

Project	Architectural Domain	Networking Technology	Layer of Virtualization	Level of Virtualization
VNRMS [7]	Virtual network management	ATM/IP		Node/Link
Tempest [8]	Enabling alternate control architectures	ATM	Link	
NetScript [9]	Dynamic composition of services	IP	Network	Node
Genesis [5]	Spawning virtual network architectures		Network	Node/Link
VNET [10]	Virtual machine Grid computing		Link	Node
VIOLIN [11]	Deploying on-demand value-added services on IP overlays	IP	Application	Node
X-Bone [6]	Automating deployment of IP overlays	IP	Network	Node/Link
PlanetLab [12]	Deployment and management of overlay-based testbeds	IP	Application	Node
UCLP	Dynamic provisioning and reconfiguration of lightpaths	SONET	Physical	Link
AGAVE [4]	End-to-end QoS-aware service provisioning	IP	Network	
GENI	Creating customized virtual network testbeds	Heterogeneous		
VINI [13]	Evaluating protocols and services in a realistic environment		Link	
CABO [3]	Deploying value-added end-to-end services on shared infrastructure	Heterogeneous		Full

Table 1. *Characteristics of different network virtualization projects over the years.*

the granularity at which each VN can administer itself. At one end of this spectrum, node virtualization creates VNs by connecting virtual nodes on different physical nodes (e.g., PlanetLab). At the other end, CABO proposes the concept of true plurality where each VN has a semblance of the native network.

It can be noticed from Table 1 that over time, research on network virtualization has shifted focus toward creating a holistic and generalized NVE that features a completely virtualized (virtualization of all network elements), highly customizable (virtualization at lower layers), and technology-agnostic (creation of VNs over a heterogeneous combination of underlying networks) networking facility for the future Internet.

Research Challenges

Most of the existing research work related to network virtualization can at best be described as attempts to fix existing problems, rather than a conscious and focused push to build a complete NVE. As a result, several aspects of network virtualization remain unexplored, and many others require modification and improvement. In this section we summarize the key issues to be resolved for the realization of an NVE.

Interfacing

Every InP must provide an interface, following some standard, so that SPs can communicate with them and express their requirements. In addition, standard interfaces are also required to make programmability of network elements available to the SPs. On a similar note, appropriate interfaces between end users and SPs, as well as among multiple InPs and among SPs must also be identified and standardized.

Signaling and Bootstrapping

Before creating a VN, an SP must already have network connectivity to the InPs in order to issue its requests. This introduces circularity where network connectivity is a prerequisite to itself [3]. There must also be bootstrapping capabilities to allow SPs to customize the virtual nodes and virtual links allocated to them through appropriate interfaces. Both requirements call for at least another network that will always be present to provide connectivity to handle these issues, or an *out-of-band* mechanism to perform signaling and bootstrapping.

Resource and Topology Discovery

In order to allocate resources for requests from different SPs, InPs must be able to determine the topology of the networks they manage as well as the status of the corresponding network elements (physical nodes and interconnections between them, remaining capacities in nodes and links, etc.). Furthermore, two adjacent InPs must also be able to instantiate cross-domain virtual links to enable end-to-end VNs.

From an SP's point of view, a VN should be able to discover the presence and topologies of other coexisting VNs. This will allow VNs to communicate, interact, and collaborate between themselves to provide larger complex services.

Resource Allocation

Efficient allocation and scheduling of physical resources among multiple VN requests is extremely important in order to maximize the number of coexisting VNs, and increase the utilization and revenue of the InPs. The allocation of resources with constraints on virtual nodes and virtual links, also known as the VN embedding problem, can be represented using a mixed-integer program (MIP) [14]. Solving an MIP is known to be $\mathcal{NP}$-hard; so is finding optimal VN embedding.

Existing heuristic-based solutions focus on two major versions of the problem in the single-InP scenario: *offline*, where all the SPs' requests are known in advance, and *online*, the opposite. Even though various constraints and objectives make this problem computationally intractable, the presence of multifarious topologies and possible opportunities to exploit them still leave enough room for research on customized solutions and better approximation algorithms. In addition, VN embedding across multiple InPs is still a virtually untouched problem.

Admission Control and Usage Policing

When establishing a VN, an SP may require specific guarantees for its VNs' attributes as well as its virtual links' characteristics. InPs must perform accurate accounting, and implement admission control and distributed usage policing algorithms to ensure that they can deliver the guaranteed performance, and the hosted VNs do not exceed allocated resources either locally or globally. However, algorithms must be developed for complete VNs, instead of existing admission control or policing algorithms for individual nodes or links.

Virtual Nodes and Virtual Links

Virtual nodes allow multiple SPs to share the same set of physical resources and implement separate customized control protocols on them. Until now, router vendors have promoted virtual nodes as a tool for simplifying core network design, decreasing capital expenditure (CAPEX), and VPN purposes. A similar concept can be extended with programmability to create substrate routers that allow each SP to customize their virtual nodes. Scalability of an NVE is closely tied to scalability of the physical elements used by the InPs. Research in this direction should focus on increasing the number of virtual nodes any single physical router can hold.

To realize network virtualization, links between virtual nodes must also be virtualized. The ability to create tunnels over multiple physical links already exists in the context of VPNs. Similar tunneling mechanisms can also be used in VNs. The speed of transporting packets across a virtual link should be comparable to that of a native link, which translates into minimum encapsulation and multiplexing cost.

Naming and Addressing

Mapping between different address contexts is a well-known problem in the existing literature. But in the presence of different, often incompatible, addressing requirements in different VNs, it becomes even more complicated [15].

Naming and addressing should be decoupled in an NVE so that any end user can move from one SP to another with a single identity. Even though the concept of being simultaneously connected to multiple VNs from different SPs sounds similar to *multihoming*, the problem is exacerbated by the possible heterogeneity among different VNs [15].

Mobility Management

In an NVE, mobility of the devices must be supported congenitally, not using makeshift solutions as in the existing Internet. Mobility in this context does not just refer to its simplest form (i.e., geographic mobility of end user devices); routers in the core network can also move around using migration techniques. As a result, finding the exact location of any device at a particular moment and routing packets accordingly is a complex issue that needs simple solutions. In addition, end users can also move logically from one VN to another in order to access different services, which further complicates the problem.

Monitoring, Configuration, and Failure Handling

To enable individual SPs to configure, monitor, and control their VNs irrespective of others, considerable changes are required from the level of network operations centers (NOCs) to intelligent agents at lower-level network elements. The concept of MIBlets [7] (partitioned management information bases [MIBs]) to gather and process performance statistics for each of the coexisting VNs instead of using a common MIB can be a good starting point. But a full-fledged robust monitoring framework needs more attention and effort.

Failures in the underlying physical network components can give rise to cascading series of failures in all the VNs directly hosted on those components. Detection, propagation, and isolation of such failures, as well as protection and restoration from them are all open research challenges.

Security and Privacy

Isolation between coexisting VNs can only provide a certain level of security and privacy through the use of secured tunnels, encryptions, and so on; but it does not obviate the prevalent threats, intrusions, and attacks to the physical layer and VNs. In addition to that, security and privacy issues specific to network virtualization must also be identified and explored. For example, programmability of the network elements can increase vulnerability if secure programming models and interfaces are unavailable. All these issues require close examination to create a realistic NVE.

Interoperability Issues

End-to-end VNs can span across multiple administrative domains, each using possibly heterogeneous networking technologies and management frameworks. Enabling virtualization in each of these technologies requires specific solutions for provisioning, operation, and maintenance. Interactions between such contrasting underlying infrastructures, while providing a generic and transparent management interface for SPs to

Efficient allocation and scheduling of physical resources among multiple VN requests is extremely important in order to maximize the number of coexisting VNs and to increase the utilization and revenue of the InPs.

Project	Originated	Link
4WARD	Europe	http://www.4ward-project.eu/
AKARI	Japan	http://akari-project.nict.go.jp/
CABO	United States	http://www.cs.princeton.edu/~jrex/virtual.html
Clean Slate	United States	http://cleanslate.stanford.edu/
GENI	United States	http://www.geni.net/
NouVeau	Canada	http://netlab.cs.uwaterloo.ca/virtual/
PlanetLab	United States	http://www.planet-lab.org/
Trilogy	Europe	http://www.trilogy-project.org/
UCLP	Canada	http://www.uclp.ca/
VINI	United States	http://www.vini-veritas.net/
X-Bone	United States	http://www.isi.edu/xbone/

Table 2. *Recent network virtualization related projects.*

easily compose and manage VNs, remains a daunting task. In addition, identification of the necessity, scope, and required interfaces for end-to-end communication across multiple VNs deserves close scrutiny.

Network Virtualization Economics

Unlike traditional networks where bandwidth is the chief commodity, virtual nodes are equally important as virtual links in an NVE. SPs are the buyers in this economy, whereas InPs are the sellers. There can also be brokers who act as mediators between the buyers and sellers. End users also participate as buyers of services from different SPs.

Traditionally, there are two general types of marketplaces: centralized and decentralized. Centralized marketplaces are efficient, but vulnerable and not scalable. On the other hand, fully decentralized marketplaces are extensible and fault-tolerant, but prone to malicious behavior and inefficiency. To find a trade-off between these two options, existing work on peer-to-peer (P2P) marketplaces can be extended to the domain of network virtualization.

Conclusion

Amid current trends of virtualizing practically every aspect of computing (e.g., operating systems, servers, and data centers), network virtualization stands at a unique point in the virtualization design space. On one hand, it is necessary to have a virtualized network to interconnect all other virtualized appliances to give each of the virtual entities a complete semblance of their native counterparts.

On the other hand, after enjoying years of rapid growth, the progress of the Internet and networking in general has come to a standstill. Most researchers now agree that a redesign is a bare necessity, not luxury. Network virtualization can take a leading role in this scenario to promote innovation through disruptive technologies. This realization has given birth to several projects all over the world that are directly or indirectly related to network virtualization (Table 2).

The materialization of an NVE needs to satisfy the requirements set by its characteristics and design goals, but fulfilling these requirements is not so easy. More insights into the challenges outlined in this article are needed for an open, flexible, and heterogeneous networking environment to become a reality.

References

[1] T. Anderson *et al.*, "Overcoming the Internet Impasse through Virtualization," *Computer*, vol. 38, no. 4, 2005, pp. 34–41.

[2] J. Turner and D. Taylor, "Diversifying the Internet," *Proc. GLOBECOM '05*, vol. 2, 2005.

[3] N. Feamster, L. Gao, and J. Rexford, "How to Lease the Internet in your Spare Time," *SIGCOMM Comp. Commun. Revi.*, vol. 37, no. 1, 2007, pp. 61–64.

[4] M. Boucadair *et al.*, "A Framework for End-to-End Service Differentiation: Network Planes and Parallel Internets," *IEEE Commun. Mag.*, vol. 45, no. 9, Sept. 2007, pp. 134–43.

[5] M. Kounavis *et al.*, "The Genesis Kernel: A Programming System for Spawning Network Architectures," *IEEE JSAC*, vol. 19, no. 3, 2001, pp. 511–26.

[6] J. Touch, "Dynamic Internet Overlay Deployment and Management using the X-Bone," *Comp. Networks*, vol. 36, no. 2–3, 2001, pp. 117–35.

[7] W. Ng *et al.*, "MIBlets: A Practical Approach to Virtual Network Management," *Proc. 6th IFIP/IEEE Int'l. Symp. Integrated Net. Mgmt.*, 1999, pp. 201–15.

[8] J. E. van der Merwe *et al.*, "The Tempest: A Practical Framework for Network Programmability," *IEEE Network*, vol. 12, no. 3, 1998, pp. 20–28.

[9] S. da Silva, Y. Yemini, and D. Florissi, "The NetScript Active Network System," *IEEE JSAC*, vol. 19, no. 3, 2001, pp. 538–51.

[10] A. Sundararaj and P. Dinda, "Towards Virtual Networks for Virtual Machine Grid Computing," *Proc. 3rd USENIX Virtual Machine Research Tech. Symp.*, 2004.

[11] P. Ruth *et al.*, "Virtual Distributed Environments in a Shared Infrastructure," *Computer*, vol. 38, no. 5, 2005, pp. 63–69.

[12] L. Peterson *et al.*, "A Blueprint for Introducing Disruptive Technology into the Internet," *SIGCOMM Comp. Commun. Rev.*, vol. 33, no. 1, 2003, pp. 59–64.

[13] A. Bavier *et al.*, "In VINI veritas: Realistic and Controlled Network Experimentation," *Proc. ACM SIGCOMM*, 2006, pp. 3–14.

[14] N. M. M. K. Chowdhury, M. R. Rahman, and R. Boutaba, "Virtual Network Embedding with Coordinated Node and Link Mapping," *IEEE INFOCOM*, 2009.

[15] N. M. M. K. Chowdhury, F. Zaheer, and R. Boutaba, "iMark: An Identity Management Framework for Network Virtualization Environment," *IFIP/IEEE Int'l. Symp. Integrated Net. Mgmt.*, 2009.

Biographies

N. M. Mosharaf Kabir Chowdhury (mosharafkabir@gmail.com) completed his Master's in computer science from the University of Waterloo, Canada, in 2009 and his Bachelor's in computer science and engineering from Bangladesh University of Engineering and Technology in 2006. His research interests include network virtualization, data center networking, and next-generation Internet architectures.

Raouf Boutaba (rboutaba@uwaterloo.ca) is a professor of computer science at the University of Waterloo. His research interests include network and service management. He is the founder and Editor-in-Chief of the *IEEE Transactions on Network and Service Management*. He is a distinguished lecturer of IEEE Communications Society, Chairman of the Technical Committee on Information Infrastructure, and Director of Conference Publications. He has received several best paper awards and other recognitions including the Premier's Research Excellence Award.

TOPICS IN NETWORK AND SERVICE MANAGEMENT

Future Internet = Content + Services + Management

Jürgen Schönwälder, Jacobs University Bremen
Marc Fouquet, Technical University of Munich
Gabi Dreo Rodosek and Iris C. Hochstatter, Universität der Bundeswehr München

ABSTRACT

While the term future Internet has gained a lot of interest recently, there is little agreement on what this term means or what the future Internet looks like. By taking the viewpoint of an Internet user who is interested in using Internet services and not so much in the protocols that move data around, we first describe some possible future Internet services. In a second step we derive some network and service management requirements, and discuss some of them in more detail.

INTRODUCTION

The term *future Internet* has gained a lot of interest recently, and since several research funding organizations have decided to support the development of the future Internet, a growing number of research projects are being established. There is, however, currently no agreement on what the technology of the future Internet will look like; nor is there agreement on what the goals of the various competing future Internet activities are. While many discussions in this context focus on technical aspects (e.g., how naming, addressing, or routing should be done in a future Internet), we believe that it is necessary to approach the design space from the service perspective.

We start our discussion with the observation that the current Internet provides services that only just work. As Mark Handley explains [1], the Internet's success has been that it basically provides 80 percent of the capability of an ideal network for 20 percent of the cost, and he points out that major changes were often last minute fixes to deal with an imminent problem; major architectural changes have become difficult in recent years. The Internet has reached a size where improvements to make the core Internet technology a bit nicer do not have much chance to succeed. This is one of the reasons some people believe it is time to work on a disruptive new Internet architecture and technology [2].

We approach the question of what a future Internet might be from a different, much less protocol-oriented perspective. For ordinary people, the Internet is defined by the set of services it provides. Most people associate the Internet with easy access to information and search engines, the availability of video, music, or general entertainment services, and access to online trading platforms, personal communication services, or social networking platforms. How the underlying network realizes these services is not important to the majority of Internet users. We therefore believe that it is necessary to approach the question of what a future Internet is by looking at the users of the Internet and the services they are likely to want to use in the future. Once we understand the driving factors, we can go back to the drawing board to identify the requirements the underlying communication network has to fulfill.

Our discussion in this article focuses on requirements impacting the network and service management plane of future networks. The remainder of this article is structured as follows. The next section discusses services we envision for the future Internet. We then identify network and service management requirements before we discuss some of them in more detail. We summarize our conclusions in the final section.

FUTURE INTERNET SERVICES

In this section we try to paint a picture of the future. To do so, we need to be aware of the colors and brushes we have at our disposal. In particular, we need to understand which technologies enable the creation of new services that were not possible before. Furthermore, we need to know how and why humans use online services, as this enables us to judge how likely it is for a scenario to become reality. To this end, we need to understand enabling technologies, but we should also take a look at which services are attracting early adopters.

KEY TECHNOLOGIES AND DRIVING FORCES

The following key technologies will enable the development of future services:

• *Ubiquitous good quality and cheap network access*: With the availability of high-quality optical wired networks and the increase of band-

 0163-6804/09/$25.00 © 2009 IEEE

The Internet needs to be treated as a critical infrastructure similar to power grids or fresh water supplies. However, the means to achieve a high level of dependability are likely very different: instead of hierarchical control structures, the Internet is much more likely to achieve dependability using massively distributed approaches.

width and quality of wireless networks, we can assume that connectivity is ubiquitous and cheap. In fact, we assume that network access is as widespread as access to electricity.

•*Novel human-computer interaction techniques*: The availability of cheap and compact sensor technology and new display technologies will revolutionize the way we interact with computers. Interaction with computers through gestures and displays embedded into everyday objects will change the way we experience computers [3]. Significant innovation can be expected in this area since novel human-computer interaction techniques help to distinguish products; hence, there is a clear incentive for companies to invest in this research area.

•*Sensors and availability of rich context information*: Sensor technology will not only enable new human-computer interaction styles but also drive new applications. We expect that sensors will be integrated in many everyday products, giving us access to rich contextual information [4]. Ubiquitous access to sensor data will enable applications we cannot imagine nowadays.

•*User-generated content and services (mashups)*: Technologies supporting user-generated content and services are developing quickly. The next big step is in mashup technologies, allowing users to combine existing content and services easily to provide novel integrated formats and services [5]. Semantic Web technologies might play an important role here once they become mature enough and accessible to a large user base.

From the perspective of ordinary Internet users, we believe that the following driving forces will shape the evolution of the Internet:

•*Personalized services*: Personalized services will become widespread on the Internet [6]. Starting from personalized and context-aware search results, we will see an increasing number of media productions targeting special interest groups, and entertainment services adapting to our preferred music or movie style and our current context and mood.

Closely related to personalized services are social networks that connect us with our friends and other people who share our interests. They depict a digital reflection of personal relationships and give us an important basis for questions of identity and trust in the Internet. Social networks will become mobile and also context aware in the near future.

•*Privacy and anonymity*: Privacy and anonymity will become increasingly important for Internet users [7]. While the awareness of these issues might be somewhat underdeveloped in today's Internet users, there are signs that this is going to change. There is already a market of specialized companies that offer Internet privacy services to (famous) people, and we expect that the sensitivity of ordinary Internet users to privacy and anonymity concerns will increase as well.

•*Computing and software as a network-centric service*: The era of personal computers installed with a large number of different applications is coming to an end. In the future we will see much simpler clients using network provided storage and software provided as a service. This frees users from dealing with nasty things such as backups and software updates. In addition, computing power can be accessed when needed through dedicated networked servers (sometimes called cloud computing [8]), reducing costs and system maintenance overhead.

•*Availability, reliability, and security*: Dependability and, in particular, availability, reliability, and security will be of major concern to Internet users as they increasingly depend on Internet services [9, 10]. The Internet needs to be treated as a critical infrastructure similar to power grids or fresh water supplies. However, the means to achieve a high level of dependability are likely very different — instead of hierarchical control structures, the Internet is much more likely to achieve dependability using massively distributed approaches.

With this understanding of key enabling technologies and driving forces in mind, we are ready to paint a picture of some future Internet services.

Examples of Future Internet Services

We expect that a large number of new Internet services will be created in the future. Due to space constraints, we limit our discussion to a small collection of future services. We hope our selection provides good coverage of the spectrum of novel services we envision.

Distributed Storage with Guarantees

Centralized data storage on media such as hard disks or DVDs has several disadvantages. The information is not available everywhere, old revisions of files are usually not archived, and the physical media have limited lifetimes. As many private computer users do not make regular backups, there is a high risk of losing data of high personal value (e.g., holiday photos, audio recordings, or personal videos).

In the future we expect that personal data will be saved in fully distributed storage networks. Storage services will offer virtually unlimited amounts of memory, they will keep a history of all changes to the user's files, they will make data accessible from anywhere on a wide range of different devices, they will provide privacy services, and they will offer storage guarantees (e.g., that the service maintains a certain number of redundant copies).

Several prototypes of fully distributed storage services based on peer-to-peer technologies already exist. However, these prototypes usually lack some of the mentioned features, the provided storage capacity is usually limited, systems often assume that contributing storage devices are located in more traditional computers, and they do not directly interface with devices such as cameras, multitouch displays embedded into furniture, or intelligent clothing.

Complete Life Replay

As enough (networked) storage becomes available, and microphones and cameras become smaller and smarter, it will be possible for people to record large portions of their life. Already today there are people who constantly carry a camera with them, which records everything they see. However, there is a lot more potential in this technology. The data collected by micro-

phones and cameras can be combined with information originating from our usage of other technical devices (e.g., cars or bikes) and from sensors embedded into everyday objects.

This will help us answer questions like "What was the name of that nice restaurant in Rome I visited last spring?" by simply replaying the relevant part of the recording of our life. Technically speaking, the major challenges are seamless integration of various data sources and automated indexing to support efficient retrieval of relevant information. For obvious reasons, effective privacy protection must be built into such a life replay service.

Health Services

Medical services provided by physicians will undergo major changes. The availability of cheap networked health monitoring devices and advanced sensors embedded into everyday objects will provide physicians with detailed data about our bodies. Persons with special risks will provide direct access to data about essential body functions to physicians, and it can be envisioned that online body monitoring services will become available, promising fast help in case of emergency.

With the availability of larger amounts of medical data, there will be a trend toward personalized medicine. Personalized medicine starts with the selection of combinations of pharmaceuticals optimized for the specific condition of the patient, and it might end with the production of personalized drugs combining chemicals in specific ways to optimize the positive effects for the patient while reducing unwanted side effects.

In the context of health services, we also envision new forms of online diagnostics. Questions such as "Am I pregnant?" might be answered by consulting a dedicated medical semantic search system and allowing it to access data collected by sensors installed in your bathroom. Many standard diagnostic tasks can potentially be provided as online services linked to medical mashups integrating pharmaceutical manufacturers, traders, online communities, and specialized semantic search systems.

Personalized and Mixed-Reality Entertainment

We envision further personalization of online entertainment services. Online offerings of music, movies, or games will not only take into account our rather static preferences in different music styles or types of games, but also our broader context, and in particular our current mood, abilities, and interests. Online radio stations will provide the music that best matches our current activities. Depending on whether you are currently cooking or cleaning the kitchen, you might receive different styles of music, and the songs played might even be selected based on the time available for carrying out an activity.

In general, we expect a growing mixture of real-world artifacts and artifacts that exist only in virtual worlds. So-called mixed-reality interfaces will enable new ways to participate in massively multiplayer online games (MMOGs). MMOGs are already played by millions of users, and the numbers are still growing. In these games social interaction is essential, as players form groups and solve quests together. With seamless integration of reality with virtual worlds, group forming and group interaction will lead to new experiences.

MMOGs are already played by millions of users and the numbers are still growing. In these games, social interaction is essential, as players form groups and solve quests together. With seamless integration of reality with virtual worlds, group forming and group interaction will lead to new experiences.

Requirements of Managing Future Services

As already indicated, the future Internet and the provided future services will change our lives, and especially the way we communicate and use our environment. Personalized services, data, and content will be accessible from everywhere by using heterogeneous equipment adapting automatically to the specifics of the underlying networks and the surrounding infrastructure. Future services will change the way users communicate and interact with each other (e.g., social networks, collaborative services). In addition, users will become more and more transparent as they will be constantly scanned by sensors, monitored by cameras, tracked (location), and profiled. If the underlying technologies and services are changing so dramatically, what are the management implications?

In contrast to traditional management approaches, where management is out-of-band, more or less centralized, and mostly added on later, the services of the future Internet require us to think of management issues already at the design phase of new services. As we all agree that management must not be an afterthought, a more detailed discussion is needed whether the management of the future Internet will be in-band or out-of-band, or even a mixture of both types. If we consider the millions of devices that need to be managed, we soon come to the conclusion that we need a new management paradigm to cope with those new challenges. Traditional out-of-band management is not designed to manage such a large number of managed objects. We have to think of new self-management approaches. Management needs to increasingly be part of the functionality of a managed object, not something that can be added afterwards. The authors believe that the strict differentiation between control and management planes (and even the user plane) will diminish.

The objective of this section is to identify requirements of managing future services, and classify them into those that can be fulfilled with traditional management approaches we are already familiar with, and those that are completely new. In more detail, the latter ones are:

- *Content and context as managed objects*: Traditionally, a managed object is the abstract representation of network resources and services that are managed. In the context of the future Internet, we have to consider new types of managed objects. We are not only dealing with managed objects representing resources associated with devices, networks, and services, but also with new managed objects: content and context. The modeling and integration of these new managed objects into management processes and supporting management platforms must be investigated.

One important aspect of future service management is moderation. Similar to today's web forums, any kind of service that involves creativity by the users or interaction between users will need some kind of moderation to ensure that all participants work together for mutual advantage.

•*Users acting as service providers*: Currently, we have a situation in which we have a large number of users using services and relatively few service providers. In the future Internet this will change since users themselves may act as providers. For example, a photographer can sell pictures of a specific event due to the network or even act as his/her own stream broadcaster. And considering car-to-car communication, a car may *inform* other cars about a problem on the street. Once users act as service providers, a moderation of services becomes an important management requirement. This by itself is not necessarily a novel aspect. However, with the distribution and increasing number of services, effective moderation mechanisms become of immense importance.

•*Personalization of services*: The current Internet provides mass services with no or little personalization. In the future services will be heavily personalized, either user-oriented or group-oriented. Personalization means not only defining parameters that can be configured, but services also need to be aware of the situation, context, and location. Thus, personalization includes a rather static component and a component that dynamically adapts to the specific situation, context, or location.

•*Seamless access to services and session mobility*: A new challenge is to access services, data, and content regardless of the underlying technology or used equipment. Besides roaming issues, aspects such as the heterogeneity and capabilities of end systems (smart phones, tablet PCs, gaming devices, cell phones, etc.) need to be managed. Session mobility addresses several management aspects such as tracking the set of active devices for a user and routing to the optimal device. Devices must register and authenticate themselves on behalf of a particular person whenever they detect an authorized user nearby.

•*Privacy of services and content*: Although privacy is not a new requirement, privacy aspects gain importance in the area of future services. They may become so important that users want to use *privacy services* to anonymize their *movements* in the ambient environment (not only their location or accessed Web site, but also other data collected by sensors or other devices, mostly not even known by the users).

•*Self-X management*: Due to the millions of objects that need to be managed and the associated scalability problems, it is necessary to rethink existing management approaches. A centralized approach (also hierarchical management) cannot address these challenges adequately. Self-X management is a new paradigm where a managed object reacts to its environment in an autonomous way, according to its rules. This paradigm seems to be a promising approach, and requires new concepts to be developed, such as coordination protocols and learning capabilities of managed objects.

•*Awareness*: Services need to be context-, situation-, location-, or even mood-aware. Thus, it is necessary to model concepts such as situation, context, and mood, and address these concepts from the management perspective. Challenges are caused by the many dependencies that can exist between services, situation, content, and context. For example, in a hospital certain services are automatically forbidden (e.g., cell phones are not allowed in intensive care rooms).

•*Security*: Since users are confronted with sensors, cameras, mobiles, networking devices, and so on all over, every device collecting sensitive data, it is more then ever important that security is ensured. The vision is that users have secure, trusted, and ubiquitous access to services. It is necessary to stress that traditional security approaches, which are more or less centralized, fail in such distributed environments, and that new approaches in terms of ambient security are required. The aspect of the quality of security needs to be approached as well.

•*Identity and trust management*: Management of identities and trust is closely related to security management, but also has privacy implications. Service providers need to trust the user's identity to ensure payment and prevent fraud, but at the same time the user is interested in limiting the amount of shared personal information and the ability of providers and third parties to collect information by combining the user's profiles from different services. Temporary identities provided by a trusted party may be a solution here. Equally important is proof of the identity of a service provider to the user.

Besides the identified new requirements, we also have to deal with known aspects such as service life cycle management, quality of service management, and the known general areas of management such as fault, configuration, accounting, performance, and security.

Implications for Network and Service Management

In this section we discuss selected network and service management aspects of the future Internet in more depth with a particular focus on aspects not yet widely addressed by the research community.

Moderation of Services

One important aspect of future service management is moderation. Similar to today's Web forums, any kind of service that involves creativity by users or interaction between users will need some kind of moderation to ensure that all participants work together for mutual advantage.

•*Copyright moderation*: The rights of content creators must be respected. On one hand, this implies that users do not publish illegal copies of movies and other copyrighted and protected works; on the other hand, users are also creative, so their rights have to be respected as well. A balance between these two aspects has to be found.

•*Peacekeeping*: In the anonymity of the Internet, some people tend to replace arguments with insults. Moderation is necessary here to keep the peace in each service's community.

•*Other illegal or offending content*: Of course, there is a lot more potential for illegal or offensive user-created content. For example, it must be ensured that inappropriate content is not made available to children.

It may be possible to automate parts of these

tasks (e.g., by filtering certain keywords). However, users may find ways to render such filters ineffective. Automated filters also have the problem of creating false positives and removing legitimate content. Therefore, human moderators will always be required as a final authority. However, it may be possible to place at least parts of the moderation authority into the hands of the service's community (e.g., by employing reputation systems).

As described earlier, there is an increasing trend of users taking the role of content and even service providers. Many users will not be able to properly moderate their own services, even though they are legally responsible for them. Therefore, moderation itself could also be offered as an external service by a third party service provider.

Privacy Management

In the future sensors will collect lots of personal data about every one of us. With the availability of cheap and reliable online storage, this sensitive data may be saved not on the user's hard disk but by some service provider that may even be located in a different country than the user. Therefore, the user must be sure that the saved data is not compromised.

Companies have strong incentives to use personal data (e.g., for personalized advertising) or even sell the data to third parties. Furthermore, employees of the companies may try to steal personal data for their own benefit. Recent years have also shown that many governments are interested in any data they can get for law enforcement and the prevention of terrorism.

There is still a lack of awareness regarding privacy issues today; people tend to give their addresses, phone numbers, and email addresses easily. However, life replays and health data are much more sensitive information. Most people will not want to store such data online or equip their house with sensors unless they can be sure that their privacy is protected.

When a service is used, it must be possible for the user to selectively reveal private information or give sensor access to the service provider based on the desired level of personalization. The user has to be allowed to control access to this data; that is, private information must not be given away to third parties without the user's approval. The user must be able to withdraw the usage rights at any time.

We regard privacy management as a major unsolved issue regarding future Internet services. Already today it is impossible for the average user to track what personal information has been given to whom. Therefore, we think that there will be novel privacy management services that help users control their private information.

Security Management

As we all know, there are many unsolved security threats on today's Internet, such as malware, phishing of (bank) account data, and denial of service attacks. Average Internet users have little knowledge of the threats, and even computer scientists are often surprised about new attack vectors.

In the future Internet the situation will become even more complicated. Web-based applications will require access to the user's data, and Web browsers will partially replace operating systems; therefore, security vulnerabilities in browsers or browser plugins could expose the user's locally and remotely stored data to attacks.

Composition of services might add an endless number of new threats. It will be very hard for a service provider to ensure that its own service does not contribute to a security problem when combined with some other services.

If users act as service providers, they will be responsible for the security of their customers. Few users will have the knowledge to make sure that their service is secure. Security management for this kind of situation may be offered by specialized companies.

There will be fewer centralized points where security can be enforced. Devices like sensors and actuators in home automation systems or car electronics are too small to run virus scanners, but as they may have wireless connectivity, it is also difficult to ensure that all of their traffic is routed through a central firewall that checks for security threats. In a distributed environment, security solutions have to be distributed as well.

In case of security vulnerabilities, all the new smart devices around us will need software updates, and it must be ensured that the update mechanism itself cannot be used for attacks.

Therefore, future services and network-capable devices need to be designed with security and security management in mind. As mentioned before, security management — as well as privacy management — can be offered as a service. Today's security suite software packages with firewalls and virus scanners may be seen as predecessors of such services.

Dynamic Service Management and Composition

As we see more and more user-generated content and services, the desire to compose new services out of existing ones to fulfill exactly one's own need increases. Together with widespread use of mobile devices, this aggravates the need for (semantic) service description and dynamic service composition.

Due to high user mobility and changing availability of services in ubiquitous environments, services have to be adapted and composed when they are actually used. This requires service discovery mechanisms to be carried out in an ongoing process, recalculating the matching between service offer and service demand, roaming between different service providers, and the transfer of sessions between different devices.

Awareness Management

Sensors anywhere will allow us to capture machine-readable information about the user's situation more or less anytime. This *context information* has to be captured from a plethora of different sources, processed to fit in quality and format, and then disseminated to any service that wants to use it.

Many users will not be able to properly moderate their own services, even though they are legally responsible for it. Therefore the moderation itself could also be offered as an external service by a third party service provider.

We believe that efficient content distribution and new services taking advantage of a networked world equipped with a large number of sensors will be the main forces driving the deployment of new technologies. Furthermore, we see great potential in the evolution of the network and service management plane.

So far, context information has mainly been used in restricted environments where everything from source to sink is controlled by one organization, which restricts the need for sophisticated context management. As more and more context sources and context consumers (context-aware services) are integrated in our everyday life, open infrastructures for managing context information will evolve.

To enable the interdomain exchange of context information, information models and exchange protocols have to be defined to describe the needs of context consumers as well as the offers of context producers including quality parameters. Context may be very sensitive information and be used in difficult, even dangerous, situations as a basis for decisions; thus, quality monitoring and trustworthiness of context sources may be needed.

One major challenge is the possible inconsistency and incompleteness of context information that arises, for example, as sensors announce different values for one distinct parameter. Specific sensors may be available in certain locations only, and other information is exclusively on hand for designated users. Finding, processing, and storing context information may overburden many networked nodes. Thus, distinct new actors like context providers or context brokers will relieve sensors and context-aware services of dealing with context management themselves.

User-Provided Services

As the user/service provider ratio changes significantly, and users offer their own content and services, many more transactions will take place between former end users. These new service providers do not have an infrastructure in place that guarantees a certain quality level and can deal with accounting. Instead of taking place inside one organization, the whole service life cycle (create, provision, manage, delete) will be distributed.

We are currently seeing the beginning of those changes. Technologies supporting mashups, defined service application programming interfaces (APIs), semantic description languages, and service creation platforms support the service creation process and will enable every Internet user to offer content and services. So-called cloud computing hides any technical details. Service hosting and content storage allow for professional service provisioning and easy management without the need to set up a single server. Many well-known management functions like quality of service, service level agreements, and accounting have to be adapted to much higher dynamics. More often than not, there will no longer exist established business connections where trust is gained over a long period; demands and offers will only last for certain situations. Future monetary transactions between users will also need much better support and reliability. Today's mostly central services will be offered in a distributed fashion; not only will users create services, but offer combined services and act as resellers as well.

Conclusion

This article argues that the term future Internet must be defined from a service perspective. We believe that taking a purely technical and protocol-centric viewpoint while defining the term runs the risk of optimizing or streamlining technology that will be very difficult to deploy on a global scale. The Internet has for a long time been just good enough, and a clean-slate design of better forwarding and control plane technology alone is not likely to provide sufficient incentives to get better technology deployed.

We believe that efficient content distribution and new services taking advantage of a networked world equipped with a large number of sensors will be the main forces driving the deployment of new technologies. Furthermore, we see great potential in the evolution of the network and service management plane. This is driven by stronger availability and reliability requirements, but also by the fact that users of future Internet services will request new moderation and management services. We believe that network and, in particular, new service management functions are becoming a distinguishing and revenue generating factor for many players in the future Internet.

Acknowledgments

The work on this article was supported in part by the EC ISTEMANICS Network of Excellence (#26854).We acknowledge the feedback received during the EMANICS "Future Internet" workshop that took place at Jacobs University in July 2008. Some ideas discussed here were created during the Dagstuhl seminar (# 07302) "Autonomic Management of Networks and Services" that took place in Dagstuhl in July 2007. Finally, we would like to thank the anonymous reviewers for their helpful comments.

References

[1] M. Handley, "Why the Internet Only Just Works," *BT Tech. J.*, vol. 24, no. 3, July 2006, pp. 119–29.
[2] A. Feldmann, "Internet Clean-Slate Design: What and Why?" *ACM SIGCOMM Comp. Commun. Rev.*, vol. 37, no. 3, July 2007, pp. 59–64.
[3] M. Wu and R. Balakrishnan, "Multi-Finger and Whole Hand Gestural InteractionTechniques for Multi-User Tabletop Displays," *Proc. 16th ACM Symp. User Interface Software Tech.*, Nov. 2003, pp. 193–202.
[4] M. Sichitiu and M. Kihl, "Inter-Vehicle Communication Systems: A Survey," *IEEE Commun. Surveys & Tutorials*, vol. 10, no. 2, 2008, pp. 88–105.
[5] D. Benslimane, S. Dustdar, and A. Sheth, "Services Mashups: The New Generation of Web Applications," *IEEE Internet Comp.*, vol. 12, no. 5, Sept. 2008, pp. 13–15.
[6] S. Ho and S. Kwok, "The Attraction of Personalized Service for Users in Mobile Commerce: An Empirical Study," *ACM SIGecom Exchanges*, vol. 3, no. 4, Jan. 2003, pp. 10–18.
[7] D. Wright *et al.* (Eds.), *Safeguards in a World of Ambient Intelligence*, Springer, 2008.
[8] A. Weiss, "Computing in the Clouds," *netWorker*, Dec. 2007, pp. 16–25.
[9] A. Avizienis, J.-C. Laprie, and B. Randell, "Fundamental Concepts of Dependability," *Proc. 3rd Info. Survivability Wksp.*, Oct. 2000.
[10] V. Igure and R. Williams, "Taxonomies of Attacks and Vulnerabilities in Computer Systems," *IEEE Commun. Surveys & Tutorials*, vol. 10, no. 1, 2008, pp. 6–19.

LTE, It's Business, Not Technology

ACADEMIC PRESS

Academic Press –
Information that you can trust.

NEW!

SAE and the Evolved Packet Core

Driving the Mobile Broadband Revolution

Pre-order your copy today!

System Architecture Evolution (SAE) is the core network architecture of 3GPP's future LTE wireless communication standard

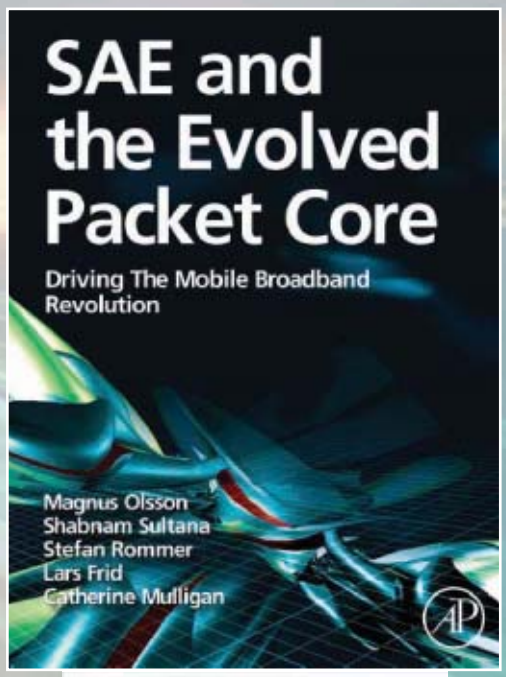

Evolved Packet Core (EPC) in Three Points

- Evolved Packet Core is a new mobile core for LTE – all-IP, end-to-end
- EPC must address a radically new set of network requirements to deliver true wireless broadband Quality of Experience (QoE)
- EPC must enable new business models and rapid introduction of new services

"The strengths of the book are that it is the first book dealing with a comprehensive description of the SAE architecture and that the authors are directly involved in SAE standardization and have a solid background."
—Dr. Kaiser

"I believe such a book will fill an empty knowledge place, making accessible the SAE concepts and capabilities to a wider spectrum of technological players in the market, beyond the small group of engineers currently working on the development of the SAE technical specifications in 3GPP. A well shaped book on SAE topic will be a good instrument in supporting the R&D work, product development work by the manufacturers, strategic planning of deployments by the network operators, as well as the development of the new services enabled by SAE."
– R. Babut

$125 USD | £64.99 | €81.95
August 2009 | 600 pages
ISBN: 9780123748263

By **Magnus Olsson; Shabnam Sultana; Stefan Rommer; Lars Frid and Catherine Mulligan**

ORDER ONLINE AT **Amazon.com** or **your favorite online retailer**
— Sign up for our **eNews** to receive the latest information you need at **www.elsevierdirect.com** —

e-news

FREE SAMPLE CHAPTER AND CONTENTS ONLINE at www.elsevierdirect.com

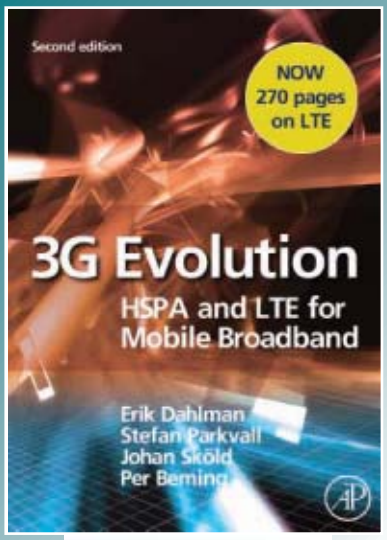

$119.95 USD | £69.99 | €89.95
October 2008 | 656 pages
ISBN: 9780123745385

3G Evolution, 2nd Edition
HSPA and LTE for Mobile Broadband

By **Erik Dahlman, Stefan Parkvall, Johan Skold** and **Per Beming**

The new edition has up-to-date coverage of the recently published LTE Release 8 radio-access standard, giving the reader insight into the ongoing and future process of LTE and LTE-Advanced standardization.

The authors of the book all work at Ericsson Research and are deeply involved in 3G development and standardization since the early days of 3G research. They are leading experts in the field and are today still actively contributing to the standardization of both HSPA and LTE within 3GPP. This includes details of the standards and technologies (160 new pages): LTE radio interface architecture, LTE physical layer and LTE access procedures.

"The best thing about this book is that unlike other books, the contents of this book are not directly picked up from specification. It seems the authors of the book have a good grasp of what they are writing and have spent enough time to make it easy to read. Language is very succinct. Each line makes sense, is relevant and helps increase understanding of LTE.

Overall this is probably the best book for LTE(and I have read several books). The book has been written in a fluid way and the flow of thought is connected. I have still not read the book completely, but I find myself going to the book again and again(in my free time or bedtime) just because it is a pleasant read."
—Jha Vidyesh Kumar

"This book offers the most comprehensive material concerning the LTE and HSPA standards that I know of. The material is well presented and is easy to follow without compromising the technical content. I really liked the book. For a person who does not attend the standards, I found the book to be an invaluable resource for my LTE work."
—Tony Rouphael

HOT FEATURE!

Log onto Elsevierdirect.com/9780123745385 and **listen to the authors speak** about their book –
Giving you insight into 3G Evolution!

ORDER ONLINE AT **Amazon.com** or **your favorite online retailer**
— Sign up for our **eNews** to receive the latest information you need at **www.elsevierdirect.com** —

FREE SAMPLE CHAPTER AND CONTENTS ONLINE at www.elsevierdirect.com

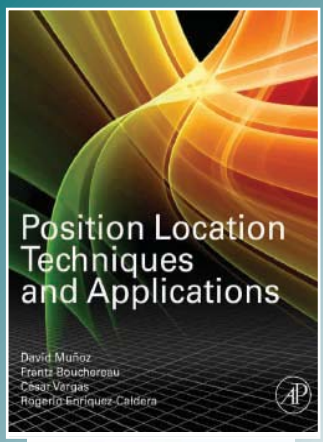

$119.95 USD | £59.99 | €87.95
April 2009 | 320 pages
ISBN: 9780123743534

Position Location Techniques and Applications

By **David Munoz, Frantz Bouchereau Lara, Cesar Vargas,** and **Rogerio Enriquez-Caldera**

This book covers the techniques and applications of position location, covering both terrestrial and satellite systems.

- Contains a complete guide to models, techniques and applications of position location
- Includes applications to wireless networks (WiFi, ZigBee, DVB networks), cognitive radio, sensor networks and reconfigurable and multi-hop networks, demonstrating the relevance of location positioning to these 'hot' areas in research and development
- Covers system design trade-offs, and the design of efficient positioning algorithms enables the creation of future location positioning systems
- Provides a theoretical underpinning for understanding current position location algorithms, giving researchers a foundation to develop future algorithms

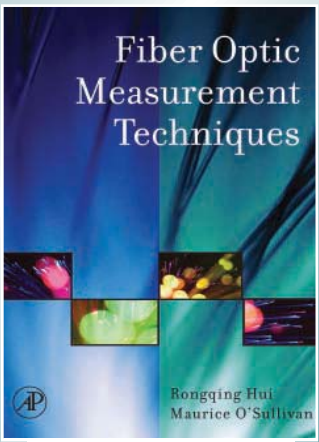

$99.95 USD | £49.99 | €73.95
January 2009 | 672 pages
ISBN: 9780123738653

Fiber Optic Measurement Techniques

By **Rongqing Hui** and **Maurice O'Sullivan**

Fiber Optic Measurement Techniques **is an indispensable collection of key optical measurement techniques essential for developing and characterizing today's photonic devices and fiber optic systems.**

The book provides you with descriptions of various fiber optic measurement methods with the emphasis on the understanding of optoelectronic signal processing methodologies.

"Optical fiber communication systems and networks constitute the core of the telecom infrastructure of the information society worldwide. Accurate knowledge of the properties of the contituent components, and of the performance of the subsystems and systems must be obtained in order to ensure reliable transmission, distribution, and delivery of information. This book is an authoritative and comprehensive treatment of fiber-optic measurement techniques, including not only fundamental principles and methodologies but also various instrumentations and practical implementations. It is an excellent up-to-date resource and reference for the academic and industrial researcher as well as the field engineer in manufacturing and network operations."
— **Dr. Tingye Li,** AT&T Labs (retired)

ORDER ONLINE AT **Amazon.com** or **your favorite online retailer**
— Sign up for our **eNews** to receive the latest information you need at **www.elsevierdirect.com** —

FREE SAMPLE CHAPTER AND CONTENTS ONLINE at www.elsevierdirect.com

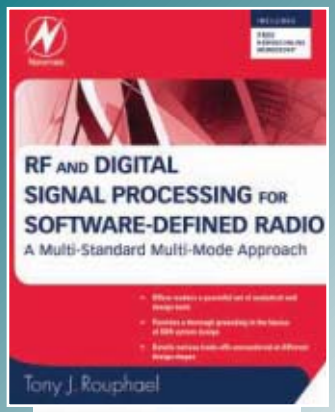

$69.95 USD | £39.99 | €57.95
November 2008 | 400 pages
ISBN: 9780750682107

RF and Digital Signal Processing for Software-Defined Radio

A Multi-Standard Multi-Mode Approach

By **Tony J. Rouphael,** RFMD, CA, USA

Understand the RF and Digital Signal Processing Principles Driving Software-defined Radios!

KEY TOPICS

- Modulation techniques
- RF system-design parameters
- Essentials of baseband and band-pass sampling and gain control
- Nyquist sampling converter architectures
- Oversampled data converter architectures
- Multi-rate signal processing

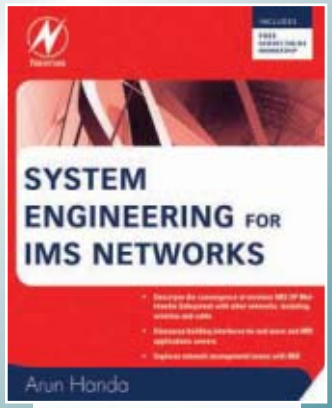

$69.95 USD | £39.99 | €57.95
February 2009 | 340 pages
ISBN: 9780750683883

System Engineering for IMS Networks

By **Arun Handa**

A Practical Guide to Designing Networks Supporting the IP Multimedia Subsystem!

This is the essential reference for anyone needing to quickly understand IMS and successfully create IMS networks.

- Describes the convergence of wireless IMS (IP Multimedia Subsystem) with other networks, including wireline and cable
- Discusses building interfaces for end users and IMS applications servers
- Explores network management issues with IMS

ORDER ONLINE AT **Amazon.com** or **your favorite online retailer**

— Sign up for our **eNews** to receive the latest information you need at **www.elsevierdirect.com** —

e-news

FREE SAMPLE CHAPTER AND CONTENTS ONLINE at www.elsevierdirect.com

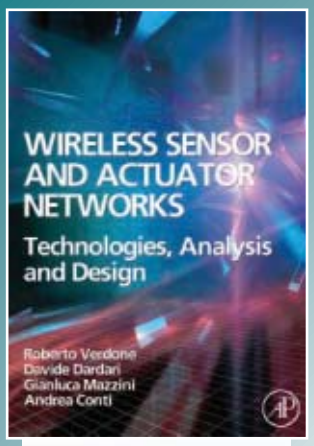

$99.95 USD | £59.99 | €86.95
March 2008 | 392 pages
ISBN: 9780123725394

Wireless Sensor and Actuator Networks
Technologies, Analysis and Design

By **R. Verdone, D. Dardari, G. Mazzini** and **A. Conti**

- A complete resource on how to evaluate the performance levels of wireless sensor network applications
- Describes the different technology options – MAC protocols, routing protocols, localization and data fusion techniques – and provides the means to numerically measure their performance

"This book is a thorough and accessible exposition on wireless sensor networks with a good balance between theory and practice: it is valuable for both students and practicing engineers, and is an essential addition for engineering libraries"
– Professor **Moe Win**, Associate Professor at the Laboratory for Information and Decision Systems (LIDS), Massachusetts Institute of Technology

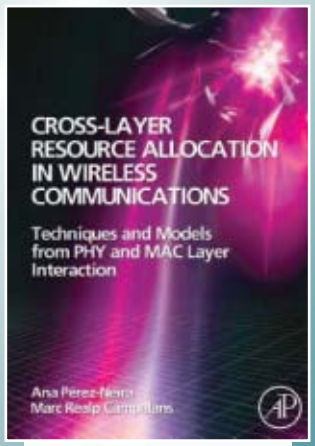

$99.95 USD | £49.99 | €69.95
October 2008 | 192 pages
ISBN: 9780123741417

Cross-Layer Resource Allocation in Wireless Communications
Techniques and Models from PHY and MAC Layer Interaction

By **A. Pérez-Neira** and **M. Realp Campalans**

The first book to give practical techniques and models for design and optimization of cross-layer resource allocation

- Unique cross-layer approach to optimizing network resources by linking the physical and network layers
- Presents resource allocation as a cross-layer design based on an optimization of MAC layer parameters within an accurate model of the PHY layer
- Provides cross-layer techniques, models and methodologies for improved quality of service

ORDER ONLINE AT **Amazon.com** or **your favorite online retailer**
— Sign up for our **eNews** to receive the latest information you need at **www.elsevierdirect.com** —

FREE SAMPLE CHAPTER AND CONTENTS ONLINE at www.elsevierdirect.com

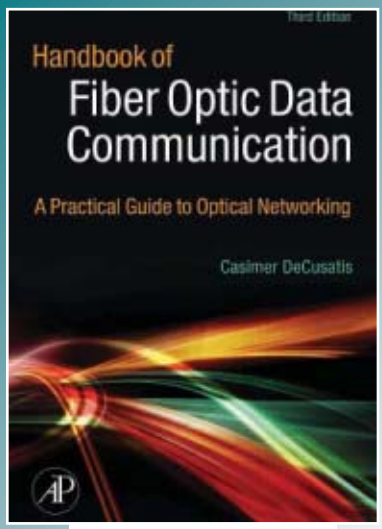

$99.95 USD | £59.99 | €86.95
January 2009 | 672 pages
ISBN: 9780123742162

Handbook of Fiber Optic Data Communication,
3rd Edition,
A Practical Guide to Optical Networking

Edited by **Casimer DeCusatis,** Distinguished Engineer, IBM Corporation

FEATURES INCLUDE:

- New Case Studies on Voice/Data Convergence, Redesigning Mainframe I/O, National LambdaRail, and optical peer-to-peer networks
- Includes an expanded listing of references on the World Wide Web, plus hard-to-find references for international, homologation, and type approval requirements
- Quick reference tables of all the key optical network parameters and a glossary that defines hundreds of technical terms and acronyms
- An accompanying website provides an online portal to the data communications community, containing podcasts on key topics, links to related material from technical societies, corporations and other professional organizations

Written for engineers by engineers, this Handbook will be an indispensable, hands-on reference for optical networks and equipment developers, designers, and installers, as well as for students studying optical fiber communications wanting an understanding of, and insight into, professional practice.

"One notable feature of this part is a chapter entitled Manufacturing, Environmental Laws, Directives and Challenges. Not typically found in any book on optical fiber communications, this section includes cpmarisons of environmental regulations in the USA, the European Union, Japan, China, and Korea. For any company contemplating manufacturing in any of these countries, this chapter provides a quick overview and comparison of the regulatory environments in them.

As the title states, this is a handbook and a useful reference for anyone working in optical fiber communications. The book has numerous references and provides an excellent introduction to many facets of optical fiber communications."
—**James Bechtel,** L-3 Photonics

ORDER ONLINE AT **Amazon.com** or **your favorite online retailer**
— Sign up for our **eNews** to receive the latest information you need at **www.elsevierdirect.com** —

FREE SAMPLE CHAPTER AND CONTENTS ONLINE at www.elsevierdirect.com

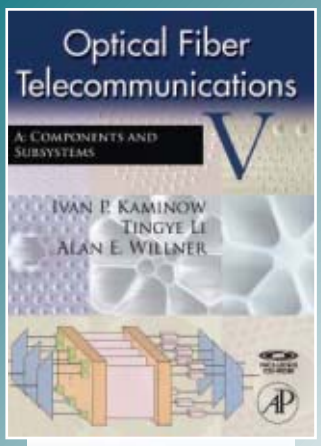

$150.00 | £79.99 | €116.00
February 2008 | 944 pages
ISBN 9780123741714

Optical Fiber Telecommunications V A – Components and Subsystems

By Ivan Kaminow; Tingye Li; and Alan E. Willner

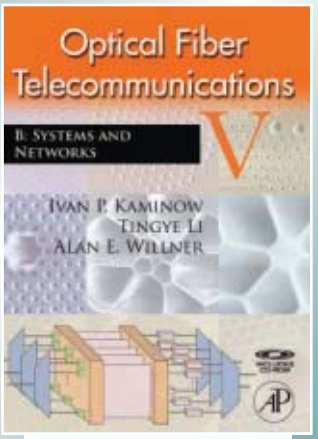

$150.00 | £79.99 | €116.00
February 2008 | 944 pages
ISBN 9780123741721

Optical Fiber Telecommunications V B – Systems and Networks

By Ivan Kaminow; Tingye Li; and Alan E. Willner

A classic for optical engineers, researchers, and graduate students for 29 years, now fully updated

For nearly three decades, the OFT series has served as the comprehensive primary resource covering progress in the science and technology of optical fiber telecom. It has been essential for the bookshelves of scientists and engineers active in the field. OFT V provides updates on considerable progress in established disciplines, as well as introductions to new topics. [OFT V]... generates a value that is even higher than that of the sum of its chapters.
— **Herwig Kogelnik,** Vice President Adjunct, Bell Labs, Alcatel-Lucent

... is a comprehensive and authoritative coverage of the latest research advances and development trends in the field, while upholding the highest standards of scholarly exposition and practical perspective. The wealth of material on innovative technologies and advanced applications will serve as an important and timely information resource ... for the advancement of telecommunications world-wide.
— **Leping Wei,** CTO, China Telecom

Lightwave systems constitute the nervous system of the industrial world and continue to evolve as innovations are introduced with enormous economic impact. The editors have very skillfully brought together authoritative chapters written by well known experts, encompassing new technologies that are enabling the rapid advances to their commercial deployment. This is a "must-have" book ...
— **Henry Kressel,** Managing Director, Warburg Pincus

ORDER ONLINE AT **Amazon.com** or **your favorite online retailer**
— Sign up for our **eNews** to receive the latest information you need at **www.elsevierdirect.com** —

FREE SAMPLE CHAPTER AND CONTENTS ONLINE at www.elsevierdirect.com

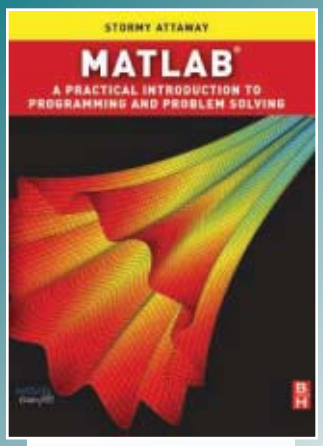

$44.95 USD | £22.99 | €32.95
March 2009 | 480 pages
ISBN: 9780750687621

MATLAB: A Practical Introduction to Programming and Problem Solving

By **Stormy Attaway**

"The text is straightforward and easy to read, with abundant examples. Sidebars offer quick opportunities to practice your understanding of each concept before moving on to the next one. The end of each chapter contains exercises that allow you to apply what you've learned to some practical sample problems, such as calculating wind chill factor. The author explains some of the differences in terminology between MATLAB and other languages, which will help people more easily apply the programming knowledge learned in this book to other languages (example: a vector in MATLAB is equivalent to what is called a one-dimensional array in other languages). User-defined functions (programs) are gradually described, starting with simple cases (calculating single values) and leading to more complex cases (functions that do tasks that do not return values, or return more than one value). The physical presentation of the book is very nice - the pages are printed in color on semi-glossy paper, and there is ample room in the margins for notes. I used this margin space to note a case where the MATLAB function in my version of the software differs slightly in name and usage from the one in the book (the author notes that this may be the case depending on the version of MATLAB you're running and the platform on which you're running it).
—**sandrift,** (see entire review on Amazon.com)

Go to Amazon.com to see more!

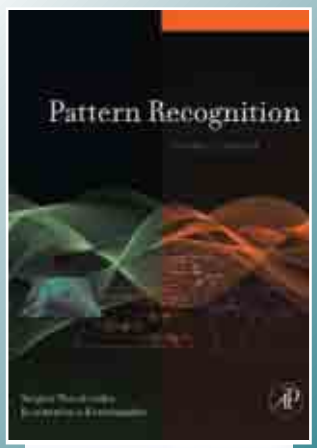

$99.95 USD | £49.99 | €73.95
October 2008 | 984 pages
ISBN: 9781597492720

Pattern Recognition, 4th Edition

By **Sergios Theodoridis** and **Konstantinos Koutroumbas**

This book considers classical and current theory and practice, of supervised, unsupervised and semi-supervised pattern recognition, to build a complete background for professionals and students of engineering. The authors, leading experts in the field of pattern recognition, have provided an up-to-date, self-contained volume encapsulating this wide spectrum of information. The very latest methods are incorporated in this edition: semi-supervised learning, combining clustering algorithms, and relevance feedback.

"The best textbook available today to introduce one to the important topic of Pattern Recognition. Comprehensive coverage; concise and clear presentation; provides the background and intuition needed to comprehend the material; and connects theory and applications throughout."
—**Harry Wechsler,** George Mason University

ORDER ONLINE AT **Amazon.com** or **your favorite online retailer**
— Sign up for our **eNews** to receive the latest information you need at **www.elsevierdirect.com** —

Biographies

Jürgen Schönwälder (j.schoenwaelder@jacobs-university.de) is an associate professor of computer science at Jacobs University Bremen, leading the research group on computer networks and distributed systems. He received his diploma in computer science in 1990 and his doctoral degree in 1996 from the Technical University of Braunschweig, Germany. His research interests are network management, distributed systems, wireless sensor networks, and network security. He is an active member of the Internet Engineering Task Force (IETF) where he has edited more than 20 specifications and standards, and chair of the Network Management Research Group (NMRG) of the Internet Research Task Force (IRTF).

Marc Fouquet received his Diploma in computer sciences from the University of Bonn, Germany, in 2004. After that he worked as a scientific staff member at the University of Tübingen, Germany, where he worked on future mobile networks in cooperation with Nokia Siemens Networks. In late 2008 he changed to Technische Universität München. His research interests are network security with a special focus on denial-of-service attacks, peer-to-peer networks, and the collection of management data in beyond 3G networks.

Gabi Dreo Rodosek received her M.Sc. degree from the University of Maribor, Slovenia, and a Ph.D. in computer science from Ludwig-Maximilians University of Munich, Germany. She holds the Chair of Communication Systems and Internet Services at the Universität der Bundeswehr München, Germany. Her research interests include integrated IT management, IT security of networks and operating systems, security management infrastructure, software defined radio, biometrics, and grid computing.

Iris C. Hochstatter received her Diploma in computer science from Ludwig Maximilians University Munich, Germany. She pursued graduate studies in technology management at the Center for Digital Technology and Management and the University of California at Berkeley. Since 2007 she has been a research assistant at the Department of Computer Science of the Universität der Bundeswehr München and a member of the European Network of Excellence EMANICS. Her research interests are ubiquitous computing, context-awareness, future Internet, and network management.

WCET CERTIFICATION

Wireless Communication Engineering Technologies (WCET) Certification Program

Vendor-neutral • Trans-national

Fall 2009 IEEE WCET Testing Window & Application Deadline

Application Period: 6 July - 4 September 2009
Application Deadline: 4 September 2009
Testing Window: 12 October - 31 October 2009

WWW.IEEE-WCET.ORG

• Free Candidate's Handbook • Free 1 hour Webinar
• IEEE Wireles Communications Professional Newsletter

IEEE Celebrating 125 Years of Engineering the Future

IEEE COMMUNICATIONS SOCIETY

IEEE Communications | Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page | qMags

With Evolution for Revolution: Managing FEDERICA for Future Internet Research

Peter Szegedi, Trans-European Research and Education Networking Association
Sergi Figuerola, i2CAT Foundation
Mauro Campanella, Consortium GARR
Vasilis Maglaris, National Technical University of Athens
Cristina Cervelló-Pastor, Universitat Politècnica de Catalunya

ABSTRACT

Over the last two decades the importance of data networking for human beings and systems has increased beyond any expectation in size, complexity, and impact on society. Today, technology offers the ubiquitous and constant possibility of being connected to the Internet at a wide range of speeds. Traditional management solutions have up to now followed an evolutionary path, although the scale of the Internet and emerging novel architectures such as peer-to-peer, ad hoc networks, as well as virtualization-capable network infrastructures require focused and possibly revolutionary changes in management approaches. This article elaborates on challenges posed by the renaissance of virtualization as experienced in the planning, development, and operation of the FEDERICA infrastructure. The European Community co-funded project FEDERICA, like other worldwide initiatives such as FIND/GENI in the United States, NWGN in Japan, and the FIRE program in Europe, is supporting the development of the future Internet. FEDERICA extends the virtualization capabilities of the current hardware and software to provide a flexible infrastructure to host disruptive testing by networking researchers.

INTRODUCTION

The number of hosts connected to the worldwide Internet is growing exponentially. Internet users number nearly 1.5 billion, and a wide variety of services are available along with a vast amount of information. However, it is now clear that the original Internet architecture needs to be improved, at least in the areas of security, management, monitoring, and mobility. The evolution of the Internet happened faster than the basic architectural concepts for a comprehensive design could possibly be formed, and management concepts have been developed just to keep pace with this evolution. The two main paths to solve the management challenges of the future Internet could be evolutionary (incremental) or revolutionary (clean slate). Either would be successful in the end; the bottom line is that we have to think about the manageability of the future Internet well in advance. It is also fundamentally necessary to plan a migration path that keeps the Internet operational at all stages.

Research activities on novel network architectures and protocols are anticipated all around the world to address a wide range of innovations. Pioneering clean slate initiatives such as Future Internet Design (FIND), Global Environment for Network Innovations (GENI), and Stanford University's "Clean Slate Design for the Internet" in the United States clearly indicate a number of possible areas for innovations that will be of significant importance for the Internet of the future. These areas include addressing and identification, cross-layer design, network virtualization, routing and traffic engineering, dynamic switching of optical circuits, decoupling of control and data, service discovery and composition, as well as management. The European Commission launched the Future Internet Research and Experimentation (FIRE) initiative in its 7th Framework Programme (FP7), and a key element is that new proposals for Internet architectures, protocols, and services should not be defined by paperwork, but should undergo early experimentation and testing in large-scale environments [1]. To provide efficient support for such innovative research activities, the FP7 project titled Federated E-Infrastructure Dedicated to European Researchers Innovating in Computing Network Architectures (FEDERICA) [2] is the enabling platform for all stakeholders involved in "Network of the Future" initiatives to develop highly innovative ideas, both evolutionary and clean

 0163-6804/09/$25.00 © 2009 IEEE

IEEE Communications | Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page | qMags

slate, tested in a production-like pan-European e-infrastructure used for proof of concepts.

To fulfill its goals, the management capabilities of the FEDERICA infrastructure are being developed to include virtualization, and actually control and monitor more than one network at the same time. The infrastructure itself will also provide a useful tool to support research in network management. The success of the project relies not only on the deployment of the infrastructure itself, but also on the results of the user projects dealing with revolutionary research.

This article is organized as follows. After a brief overview of the infrastructure projects supporting future Internet research, we provide a brief summary of the major network management strategies. Then we introduce the European project FEDERICA as a tangible use case of a virtualization-capable network infrastructure deployment. We discuss its design objectives, architecture, and basic processes, including challenges in virtualized network management. We describe the FEDERICA vision of the future Internet, and finally we conclude our article mentioning some future work.

Infrastructure Support for Future Internet Research

It has been globally realized that the novel concepts of the future Internet must be experimented and tested in large-scale environments, even in the early phase of the research. That is why such ambitious research initiatives as the National Science Foundation program GENI in the United States, New Generation Network (NWGN) in Japan, Future Internet Forum (FIF) in South Korea, the China Next Generation Internet (CNGI) project, and the G-Lab project in Germany together with the FIRE initiative at the European level are supported by infrastructure development activities.

To design and implement large-scale experimental facilities supporting revolutionary research on the future Internet, the introduction of network resource virtualization techniques is required. Virtualization in computer science is a mature and widely used technique, and PlanetLab [3], a global research network that supports the development of planetary-scale network services, was the first community testbed to use virtualization as the main principle, starting in 2002. The network slice has been defined as the collection of virtual machines running in a number of computer nodes of PlanetLab connected through the global Internet. All participating members are allowed to use and manage slices independently for their experiments, with site administrators able to manage slices through adding or removing virtual machines.

Although PlanetLab is a popular tool for networked services research, it may not be used for research on lower-layer network protocols and architectures. The GENI concept [4] addresses these shortcomings and extends the aforementioned virtualization principle to all layers of the infrastructure with the notions of substrate, slicing, and federation. GENI allows the creation of experiments that create a version of the network that could have its own management protocols and incompatibilities with the Internet.

FEDERICA follows an architecture similar to GENI, using substrate, slicing, and federation principles. Currently, finding the proper management concept for virtualization-capable networks and their services is a real challenge, as any solution should be simple and easy to use by clients (i.e., researchers) without wasting time managing things and concentrating only on the research objectives.

Finding the proper management concept for virtualization capable networks and its services is a real challenge, as any solution should be simple and easy-to-use by clients without wasting time on managing things and concentrating only on the research objectives.

Strategies in Network Management

Network and service management in future networks is becoming an important topic involving a whole set of evolutionary and novel approaches. Infrastructure support projects like FEDERICA provide a scalable and federated platform for ongoing research on various management approaches applying the infrastructure virtualization for network slicing as a main principle.

Evolutionary Approaches

The functional requirements of traditional network management are summarized by the well-known term fault, configuration, accounting, performance, and security (FCAPS). In the last few years, the importance of configuration management as well as network monitoring and measurement (i.e., fault and performance management) has increased. Accounting and authorization (i.e., security management) is also becoming more and more important in a virtual environment.

Classical management frameworks such as the open systems interconnection (OSI) network management framework, telecommunication management network (TMN), or Internet network management framework followed the traditional agent-manager centralized paradigm. As networks grow, though, management complexity and service requirements mean that such paradigms are no longer adequate and should be replaced with distributed management concepts. That is clearly illustrated by the evolution of the Simple Network Management Protocol (SNMP) for Internet Protocol (IP) networks. As an early effort in 1992, SNMPv2 introduced the concept of an intermediary manager for distributing management functions. Although SNMPv3 with embedded security features is widely deployed for monitoring network devices, it has some shortcomings in configuration management and managing complex systems [5], including virtualization-capable network devices, servers, and PCs.

Modern routers and switches are highly programmable devices and can accommodate management functions directly on the device. This design trend leads to new management approaches based on Extensible Markup Language (XML) and Web services, as well as some vendor-specific solutions. We do not call these trends revolutionary because they have been around for several years now and did not really change the network administrators' way of thinking. Since XML can be efficiently used as a tex-

The FEDERICA virtual infrastructure will be able to support experimental network research in data, control and management planes, as well as all protocol layers, over an existing pan-European substrate, but in isolation from the service environment.

tual encoding mechanism for application protocols, and object and interface specifications, the XML-based management technologies seemed the obvious choice for the next evolutionary step. In parallel with some proprietary developments like Juniper's JUNOScript [6] the benefits of the combined XML and Web-based approaches became clear. Such low-cost infrastructure with proven scalability and security features was developed by the Distributed Management Task Force (DMTF) and called Web-Based Enterprise Management (WBEM). The ideas behind JUNOScript and WBEM formed the basis of the Internet Engineering Task Force's (IETF's) standardization effort on the Network Configuration (NetConf) protocol [5].

Like SNMP, the WBEM and NetConf also follow the agent-manager paradigm, albeit with important differences regarding the capabilities of the management protocol and the associated information model. NetConf is qualitatively different from all the other approaches, given that it follows a document-based concept [5] as opposed to an individual managed object access approach.

Novel Approaches

Several novel ideas are emerging, and candidates for different network management approaches have been proposed based on highly distributed and sophisticated intelligence. Novel management approaches require revolutionary thinking *out of the box* and may drop the complete set of existing solutions. These approaches must be supported by large-scale experimental network facilities such as the FEDERICA virtual infrastructure for adequate proof of concept.

The policy-based network management concept started in the early 1990s [7]. The most significant benefit of policy-based network management is that it promotes the automation of establishing management-level objectives over various types of network devices. The network administrator would interact with the network by providing high-level abstract policies, and such device-independent and human-friendly policies are needed in a typically heterogeneous environment consisting of virtual resources.

The concept of management by delegation (MbD) [7] was introduced in parallel with the policy-based approach, pushing more management tasks to the agent side in a dynamic fashion. Code mobility can be considered the capability of a management procedure to distribute and relocate its components at runtime, and this means management tasks no longer have to be performed by managers. They simply generate management objectives and outline task procedures, with the execution of tasks delegated to the agents.

The concept of intelligent agents and active networks extends this initial idea [7]. An intelligent agent is an independent entity capable of performing complex actions and resolving management problems on its own. In contrast to code mobility, an intelligent agent does not need to be given task instructions to function, just the high-level objectives. However, as the entire network management system is autonomous, the problem of managing the intelligent agents also becomes increasingly important. Active networks, combined with code mobility, could present an effective enabling technology for distributing management tasks at the virtual network resource level.

These kinds of network management approaches are the subject of current research efforts all over the world and also supported by FEDERICA.

Use Case: The FP7 Project FEDERICA

We mentioned earlier the importance of infrastructure projects and early prototyping, so the following describes the European project FEDERICA. The project partners include a wide range of stakeholders on network research, national research and education networks (NRENs), DANTE, TERENA, academic and industrial research groups, and vendors.

Motivation and Design Objectives

FEDERICA is a two-and-a-half-year European Commission funded project to implement a versatile experimental network infrastructure for trialing new networking technologies [2]. It exploits the leading edge optical networking environment in campus LANs, NRENs, and GÉANT. The FEDERICA virtual infrastructure will be able to support experimental network research in the data, control, and management planes, as well as all protocol layers, over an existing pan-European substrate, but in isolation from the service environment.

FEDERICA is not directly about the future Internet, but provides key infrastructure for such research activities accomplished by hosting virtualized facilities in the substrate nodes. It should be stressed that one of the main FEDERICA design principles is experiment reproducibility, as we have to provide a controlled environment that is flexible, reliable, and allows reproducible running conditions in conjunction with simple and easy-to-use management solutions. Hence, the management of FEDERICA does not directly imply how the future Internet will be managed, but gives a good opportunity to understand it. FEDERICA infrastructure is not about performance, but rather about testing new ideas and principles.

FEDERICA Architecture and Procedures

The FEDERICA physical substrate is built on top of the GÉANT2 Pan-European backbone network interconnecting several NRENs around Europe. The topology is composed of 13 physical sites, four core and nine non-core nodes, connected by point-to-point links. FEDERICA is capable of creating parallel virtual slices of the physical substrate that can be assigned to multiple users simultaneously according to their needs (Fig. 1). FEDERICA node facilities include:

- Programmable high-end routers/switches (Juniper MX series) enabling logical router/switch instances only at the core nodes

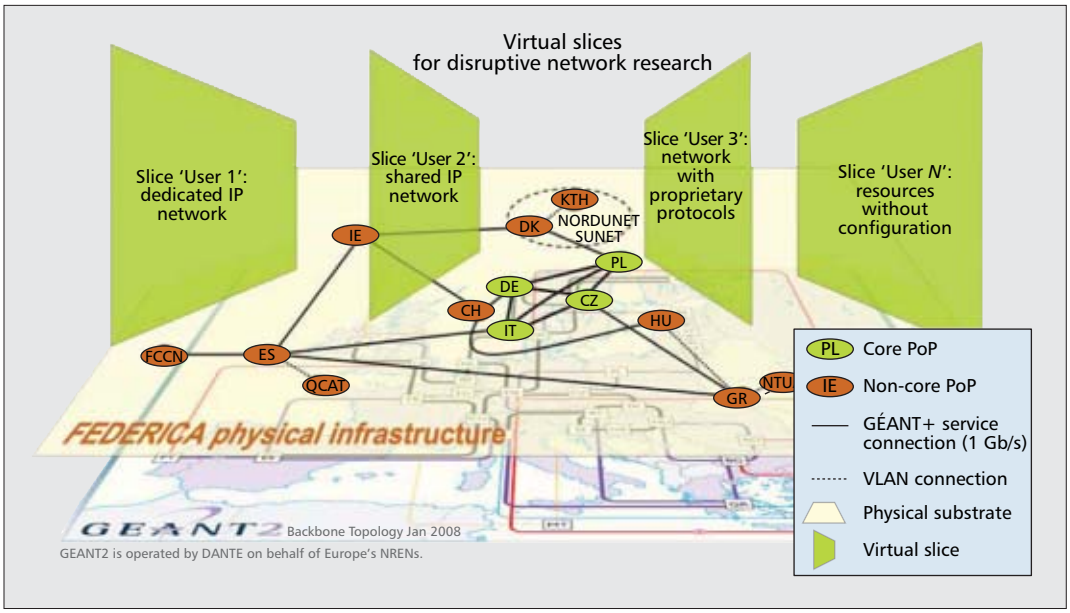

■ **Figure 1.** *The FEDERICA infrastructure.*

The FEDERICA infrastructure is engineered to enable reproducible experimentation of user scenarios when requested. This is accomplished by leveraging virtualization with appropriate hardware and software capabilities of the infrastructure.

- Multiprotocol switches (Juniper EX series) extending virtual links to non-core nodes
- PC-based virtualization-capable nodes (V-nodes on SUN servers) running virtualization software, implementing, for example, open source software routers and emulating end-user nodes

The FEDERICA infrastructure is engineered to enable reproducible experimentation of user scenarios when requested. This is accomplished by leveraging virtualization with appropriate hardware and software capabilities of the infrastructure, such as provisioning virtualized computing resources; binding separate physical network interfaces to virtual slices; assigning separate logical instances of switching/routing equipment; and partitioning physical communication links.

Compared to PlanetLab, FEDERICA's management concept is different as users are assigned virtual networks though a centralized decision making procedure. FEDERICA creates a private management network for each user to manage and configure network components inside the slice, down to the data link layer. Users subsequently access their slice resources via a dedicated proxy that safeguards these slices and the FEDERICA substrate from unauthorized and/or unscheduled demands, also acting as an access authentication engine. A FEDERICA node is not only considered an end host as in PlanetLab, but a network and computing element (i.e., programmable and physical router, switch, and server) that itself routes traffic and affects network reachability. In addition, the user is not forced to use a specific operating system or application but is free to install any software component inside the slice. While PlanetLab connections are based on the public Internet, FEDERICA allows the user to perform experiments on lower layers as well.

The main feature of FEDERICA is that users will be able to fully configure and manage the resources of their own slices, without affecting the physical infrastructure operation, and ensure the reproducibility of the experiment. The basic slice provisioning procedure performed by the FEDERICA network operation center (NOC) is illustrated in Fig. 2; note that initially, almost all of the actions performed by the NOC follow a manual procedure. However, the project envisages a dynamic system and a toolbench able to automatically create slices upon request.

Challenges in Virtualized Network Management

Within the scope of FEDERICA, it is very important to define a management strategy capable of dealing with virtualized resources. Currently there are some research projects in which some partners are involved, such as UCLPv2 [8] and MANTICORE [9], dealing with management approaches that could fit within the FEDERICA scope with some constraints. It is important, though, to consider the two novel design requirements FOR FEDERICA: isolation and federation. Isolation means that resources from one slice allocated to a specific user must act independent of resources allocated to another user when both are sharing the same physical infrastructure, while federation means that the resources in a slice can cooperate with outside resources.

The design process is considering the virtualization technology within a service oriented architecture (SOA) model. As an example, virtualization can be considered in the context of infrastructure as a service (IaaS) [10]. IaaS is a technique that represents a physical network device or network substrate as a software entity (e.g., aiming for quick and easy deployment of management approaches using Web services). The idea is that the users do not need to purchase physical hardware, but can instead pay the

FEDERICA may act as a bridge between the current Internet and the new one, acting also as a platform to test migration paths. Federating similar infrastructures is of paramount importance, and FEDERICA is already engaged in defining a service interface to federate resources with other initiatives.

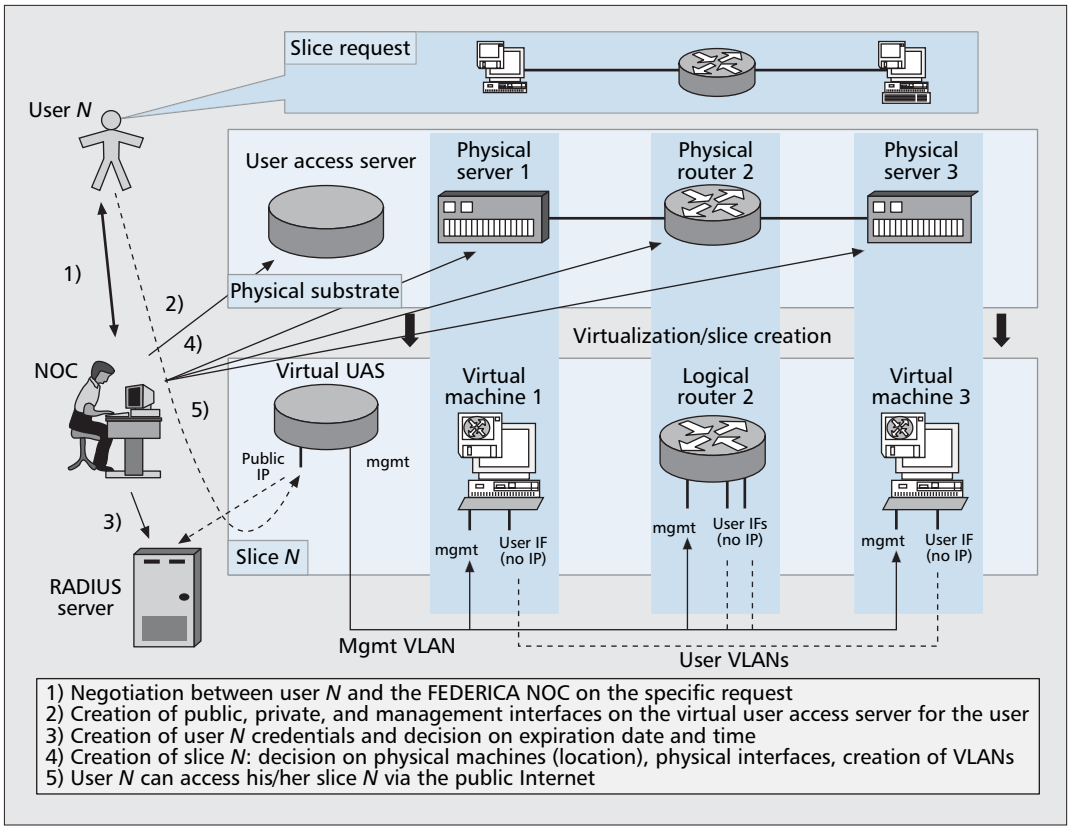

Figure 2. *Basic slice provisioning procedure.*

hardware owner for use of the resources as a service. During the lifetime of the service, the user can control and manage the requested infrastructure like the owner does.

The real challenge behind network virtualization is providing complete virtual infrastructures or slices as a service, which brings some advantages in respect to legacy management system architectures. In terms of scalability, the virtual infrastructure can follow business or project needs, since resources can scale efficiently, and when the requirement is over, network resources can be brought back to the provider. The virtualization system services can be exposed as Web services, which can be exported or connected directly to external SOA applications. SOA helps eliminate the network of layers, and allows network components and services to be accessed horizontally, vertically, and externally.

A particularly important challenge is the capability to federate different infrastructures. In that sense, infrastructure integrators can integrate resources from different management domains into a single management domain, thus federating disparate resources into a single infrastructure. Particularly for FEDERICA, there is some ongoing work to define a service interface in order to federate resources with other projects such as FP7 project OnaLab2 or even GENI in the future.

The IaaS concept is the basis of the MANTICORE project, where its virtual router configuration capabilities are going to be implemented within the FEDERICA toolbench to be delivered by the end of the project.

FEDERICA's Vision of the Future Internet

Trends and Requirements

The future Internet will develop according to two basic requirements: interconnecting a constantly growing number of elements and offering ubiquitous permanent connectivity to the majority of its users. The possibility for any user to connect to any other user is also considered fundamental to allow access to all kinds of information, as well as for sensors and objects. These requirements mandate a robust authorization and authentication infrastructure that must be available at all layers of the Internet.

The Internet will continue to be implemented as a set of interconnected autonomous domains due to scaling requirements, each with different technologies and capabilities, and without centralized management or control. Mobility, reachability, and security requirements will impose strong interdomain communication of various types of information [11]. Monitoring capabilities should then extend to mobile users traversing administrative boundaries and multiple parallel virtual networks on the same physical infrastructure. Such advances require ad hoc support in the hardware and development of new standards for virtual resource representations. In particular, the need is for a richer information system that tracks the relationships between the entities in each domain.

LIFE AFTER FEDERICA

FEDERICA may act as a bridge between the current Internet and the new one, as well as a platform to test migration paths. Federating similar infrastructures is of paramount importance, and FEDERICA is already engaged in defining a service interface to federate resources with other initiatives (e.g., the European project PanLabII and the U.S. GENI). The virtualization architecture proposed by FEDERICA allows the development of revolutionary schemes on top of the current Internet architecture, isolating the new disruptive protocols and services from the current protocol suite.

FEDERICA is planned to run until June 2010. Its results and research will be used in the European NREN environment and by the pan-European GÉANT backbone.

CONCLUSION

The Future Internet will need novel control and management capabilities. All related research approaches must be validated on large-scale experimental facilities without affecting production environments. The FEDERICA project uses simple evolutionary solutions to create an almost clean slate infrastructure for revolutionary research on existing NRENs infrastructures, including disruptive reproducible testing.

ACKNOWLEDGMENT

The FP7 project FEDERICA is partially supported by the European Commission under Grant Agreement no. RI-213107. The authors acknowledge the contribution of all project partners, and, specifically for their input, Radek Krzywania (PSNC), Björn Rhoads (KTH), Jean-Marc Uzé (Juniper Networks), and Jirí Navrátil (CESNET).

REFERENCES

[1] CORDIS, "Future Internet Research and Experimentation: An Overview of the European FIRE Initiative and Its Projects," Sept. 1, 2008; http://cordis.europa.eu/fp7/ict/fire/
[2] M. Campanella, "Federated E-Infrastructure Dedicated to European Researchers Innovating in Computing Network Architectures," *Future Internet Conf.*, Apr. 2, 2008; http://www.fp7-federica.eu/
[3] L. Peterson, "PlanetLab: A Blueprint for Introducing Disruptive Technology into the Internet," Jan. 2004; http://www.planet-lab.org/
[4] J. B. Evans and D. E. Ackers, "Overview of GENI Overview of GENI and Future Internet in the US," May 22, 2007; http://www.geni.net/
[5] G. Pavlou, *On the Evolution of Management Approaches, Frameworks and Protocols: A Historical Perspective*, Springer, Oct. 13, 2007.
[6] J. Schönwalder, A. Pras, and J-P. Martin-Flatin, "On the Future of Internet Management Technologies," *IEEE Commun. Mag.*, vol. 41, no. 10, Oct. 2003, pp. 90–97.
[7] R. Boutaba and J. Xiao, "Network Management: State of the Art," *IFIP World Comp. Cong.*, 2002.
[8] E. Grasa *et al.*, "UCLPv2: A Network Virtualization Framework Built on Web Services," *IEEE Commun. Mag.*, vol. 46, no. 3, Mar. 2008, pp. 126–34.
[9] V.Reijs and E. Grasa, "The MANTICORE Project: Providing Users with a Logical IP Network Service," *TERENA Net. Conf.*, May 2008.
[10] M. Lemay, "An Introduction to IaaS Framework," Aug. 5, 2008; http://www.iaasframework.com/
[11] N.M. M. K. Chowdhury, F.-E Zaheer, and R. Boutaba, "iMark: An Identity Management Framework for Network Virtualization Environment," *11th IFIP/IEEE Int'l. Symp. Integrated Net. Management*, Long Island, NY, June 2–4, 2009.

BIOGRAPHIES

PETER SZEGEDI (szegedi@terena.org) is currently the Joint Research Activity 2 leader of the FEDERICA project. He received his M.Sc. degree in electrical engineering from Budapest University of Technology and Economics, Hungary, 2002. He then worked toward a Ph.D. at the Department of Telecommunications of the same university. His main research interests include design and analysis of dynamic optical networks, especially optical Ethernet architectures, network control, and management processes. He worked for Magyar Telekom (2003–2007) and joined TERENA in January 2008.

SERGI FIGUEROLA (sergi.figuerola@i2cat.net) has been the coordi-nator of the Network Technologies Cluster of the i2CAT Foundation since 2004. He graduated in telecommunication engineering on the Technical University of Catalonia (UPC), Spain, in 2002, and holds a Master's in project management from La Salle (Universitat Ramón Llull, 2004). He is acting as Joint Research Activity 1 co-leader of the FEDERICA project and work package leader of the PHOSPHORUS project. He is currently pursuing his Ph.D. within UPC's Optical Communications Group.

MAURO CAMPANELLA (mauro.campanella@garr.it) graduated in physics in 1985. He is at present working for the Italian National Research and Education Network (GARR) and is responsible for research activities. He is one of the creators of the Premium IP QoS service of the European NREN network backbone GÉANT and created the architecture of the bandwidth on demand service of GÉANT2 (AutoBAHN). He is acting as the coordinator of the FEDERICA project.

VASILIS MAGLARIS (maglaris@netmode.ntua.gr) is the technical coordinator of FEDERICA. He is a professor of electrical and computer engineering at the National Technical University of Athens (NTUA), Greece, since 1989, directing the Network Management and Optimal Design Laboratory (NETMODE). He holds an Engineering degree from NTUA, 1974, and a Ph.D. degree from Columbia University, New York, 1979. He subsequently served in numerous positions in communications networking in the United States and Europe. Since October 2004 he has been chairman of the European NREN Policy Committee, responsible for the pan-European advanced network platform GÉANT.

CRISTINA CERVELLO-PASTOR (cristina@entel.upc.edu) received telecommunication engineering and Ph.D. degrees in 1989 and 1998, respectively, both from UPC. She is currently an associate professor in the Department of Telematics Engineering of UPC, which she joined in 1989. She also belongs to the Technological Committee of the i2CAT Foundation. Her research trajectory has been centered on the field of optical networks. She is currently the Joint Research Activity 1 co-leader of the FEDERICA project jointly with i2CAT.

The FEDERICA project uses simple evolutionary solutions to create an almost clean slate infrastructure for revolutionary research on existing NRENs infrastructures, including disruptive reproducible testing.

IEEE Communications | Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page | qMags

Managing Interdomain Traffic in Latin America: A New Perspective Based on LISP

Marcelo Yannuzzi and Xavi Masip-Bruin, Technical University of Catalonia
Eduardo Grampín, University of the Republic
Roque Gagliano, Latin American and Caribbean Internet Addresses Registry
Alberto Castro and Martín Germán, Technical University of Catalonia and University of the Republic

ABSTRACT

The characteristics of Latin American network infrastructures have global consequences, particularly in the area of interdomain traffic engineering. As an example, Latin America shows the largest de-aggregation factor of IP prefixes among all regional Internet registries, being proportionally the largest contributor to the growth and dynamics of the global BGP routing table. In this article we analyze the peculiarities of LA interdomain routing architecture, and provide up-to-date data about the combined effects of the multihoming and TE practices in the region. We observe that the Internet Research Task Force initiative on the separation of the address space into locators and identifiers can not only alleviate the growth and dynamics of the global routing table, but can also offer appealing TE opportunities for LA. We outline one of the solutions under discussion at the IRTF, the Locator/Identifier Separation Protocol, and examine its potential in terms of interdomain traffic management in the context of LA. The key advantage of LISP is its nondisruptive nature, but the existing proposals for its control plane have some problems that may hinder its possible deployment. In light of this, we introduce a promising control plane for LISP that can solve these issues, and at the same time has the potential to bridge the gap between intradomain and interdomain traffic management.

INTRODUCTION

Latin America (LA) is characterized by a low degree of peering connections among local domains, where most of the regional and transit traffic is exchanged outside LA, in U.S. network access points (NAPs). Internet service providers (ISPs) in LA typically purchase expensive connectivity from multiple transit providers, so one of their most important traffic engineering (TE) objectives is to load as much as possible their transit links. They also seek to improve both the performance and reliability of their transit traffic while minimizing costs. Unfortunately, the lack of investment in new network infrastructures increases the complexity, since transit traffic needs to be engineered over a bundle of links with heterogeneous capacities that are usually connected to different transit providers.

Currently, the process of finding the best trade-off for this challenging optimization problem is essentially manual, and tuned on a trial and error basis. The tools available today for Latin American ISPs are coarse-grained, and basically consist of the combination of Border Gateway Protocol (BGP)-based techniques, such as the utilization of variable prefix lengths together with AS-path prepending. The idea behind these techniques is to de-aggregate Internet Protocol (IP) prefixes, thus increasing the granularity of the BGP advertisements, allowing the distribution of traffic to be better controlled. The problem of this practice is that it fuels both the growth and dynamics of the global BGP routing table. Indeed, the average de-aggregation of IP prefixes in LA is twice as large as the global average. Issues like multihoming, suboptimal address allocation, and CIDR block de-aggregation for TE purposes, among others, are impacting on the scalability of the global BGP routing table, which soon will reach 300,000 entries [1]. This is a major concern for the Internet community [2].

To address this issue, the Internet Research Task Force (IRTF) is considering the separation of the address space into end-system identifiers (EIDs) and routing locators (RLOCs). The basic idea is that an EID represents an end-host IP address, while RLOCs represent the IP addresses where end hosts are located. The scaling ben-

This work was partially funded by the Spanish Ministry of Foreign Affairs and Cooperation through AECI-PCI A/019977/08.

 0163-6804/09/$25.00 © 2009 IEEE

IEEE Communications | Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page | qMags

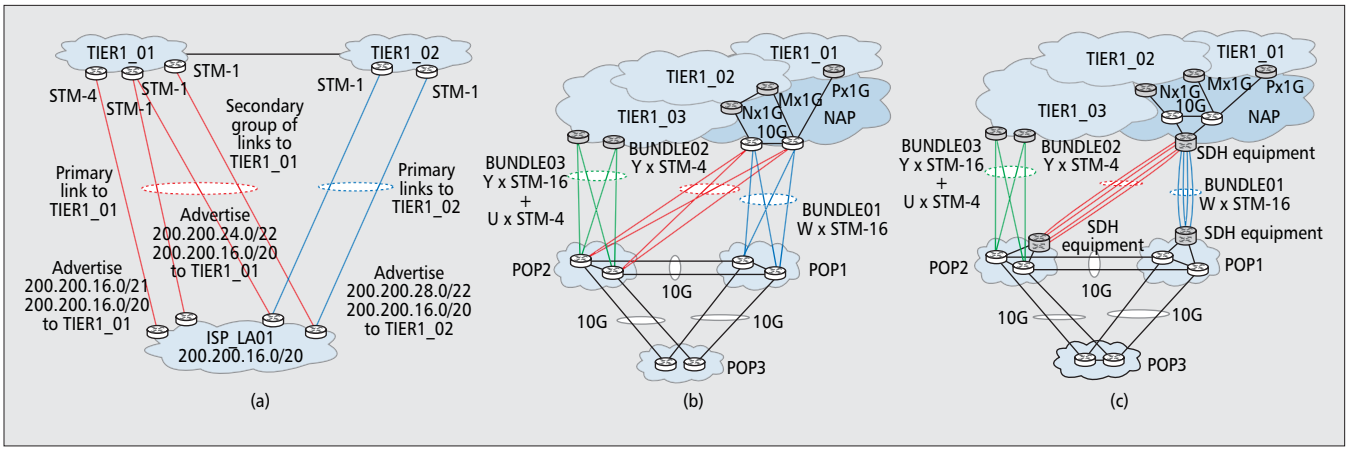

■ **Figure 1.** *a) Scenario #1: multihomed ISP with links of different capacity, load sharing, and backup routing policy; b) scenario #2: multihomed ISP with NAP presence; c) scenario #3: multihomed ISP with NAP presence and SDH multiplexers.*

efits arise when EID addresses are not routable through the Internet — only RLOC addresses are globally routable, allowing efficient aggregation of the RLOC address space.

Recent studies show that one of the solutions under discussion at the IRTF, the Locator/Identifier Separation Protocol (LISP) [3, 4], offers some key advantages. For instance, Quoitin *et al.* [5] show that the size of the global routing table can be reduced by roughly two orders of magnitude with LISP. That work also shows that LISP provides improved interdomain TE capabilities using a nondisruptive approach.

Despite these strengths, the proposals for the LISP control plane present some major challenges that are exposed and addressed in this article. These challenges lie in the fact that since EIDs are not globally routable through the Internet, a mapping system is necessary between EIDs and RLOCs.

In this article we provide up-to-date data about the characteristics and central issues of interdomain traffic management in LA. We outline the strengths of LISP, highlighting a set of important TE opportunities for LA, and we also introduce a novel control plane for LISP that deals with the challenges mentioned above. We discuss the importance of architecting the control plane so as to concurrently provide a highly efficient coupling to the DNS system, the path computation element (PCE) — if present — and an EID-to-RLOC mapping engine that borrows concepts from intelligent route control (IRC) techniques [6]. We conclude with directions for future research, especially in a promising field for LA we refer to here as intelligent route mapping (IRM).

Characterization of Interdomain Traffic Management in LA

In this section we describe some relevant characteristics of LA connectivity infrastructure, and examine their influence in aspects like interdomain routing and IP prefix de-aggregation policies.

LA Interdomain Routing Scenarios

Latin American Internet traffic has been growing at annual rates higher than 70 percent, and it will continue to grow steadily; therefore, capacity upgrade is of crucial importance to regional ISPs. In developed countries it is possible to plan capacity growth according to the observed usage and estimated traffic demands. Unfortunately, this is not the case for LA, where ISPs basically get the bandwidth they can buy, constrained to the available capacity along transit circuits. As a result, interdomain connectivity in LA is frequently composed of bundles of circuits of variable bit rates that are aggregated and used as primary as well as backup links. We proceed to illustrate some of these peculiarities by sketching three reference scenarios.

The first scenario, depicted in Fig. 1a, shows a small multihomed ISP with several links connected to two transit providers. In order to cope with the traffic growth, ISP LA01 has upgraded one of its primary links from an STM-1 to an STM-4. This is a common situation in LA, where economical and infrastructure constraints prevent ISPs from upgrading all its links at the same time.

The policy of provider ISP LA01 for managing its inbound traffic through links with different capacities is to use a combination of two BGP-based TE mechanisms, specifically, the de-aggregation of its prefix 200.200.16.0/20 and the utilization of AS-path prepending. More precisely, ISP LA01 splits the prefix 200.200.16.0/20 into three more specific prefixes: 200.200.16.0/21, 200.200.24.0/22, and 200.200.28.0/22. Prefixes 200.200.16.0/21 and 200.200.24.0/22 are advertised to TIER1 01 through the STM-4 and the secondary group of links, respectively, whereas prefix 200.200.28.0/22 is advertised to TIER1 02. In order to provide backup paths, ISP LA01 advertises the less specific prefix 200.200.16.0/20 to both TIER1 01 and TIER1 02. The effect of ISP LA01's inbound TE policy is that instead of simply advertising the prefix 200.200.16.0/20 to each transit provider, a total of six prefixes are advertised upstream. This de-aggregation of prefixes increases the size of the global BGP routing table. In addition, ISP LA01 may further

Region	IXPs	# of prefixes	De-aggregation factor (DF)
Africa	21	5K	3.46
Asia & Pacific	73	66K	2.81
Europe & Mid. East	123	67K	1.74
LA & Caribbean	24	26K	4.38
North America	88	124K	1.87
		Global BGP table	Global average
		288K	2.12

Table 1. *Statistics by region (data of April 2009, extracted from [1] and APNIC [7]).*

split the prefixes and will often tune its advertisements (on a trial and error basis), so as to adapt to variations in the inbound traffic. This, in turn, increases the dynamics of the global routing table.

In LA outbound traffic typically represents a small fraction of inbound traffic. However, the policy for managing outbound traffic requires special care, since the utilization of links with different capacities might degrade the performance of end users' applications, especially in frequent cases where the egress links are chosen using a round-robin scheme.

The second scenario considered here represents another frequent situation in LA, which comes up when the intercontinental connectivity from LA to the NAPs is part of the ISP network, meaning that the ISP owns termination equipment at both ends. This scenario is depicted in Fig. 1b, where the ISP backbone is composed of local (inexpensive) 10G bundles, and intercontinental (expensive) STM-x bundles to different transit providers. Note that traffic among POP1, POP2, and POP3 should be kept local, and clearly the choice of uplink/downlink for interdomain traffic might considerably affect the ISP's economics. The main complexity of this scenario lies in the utilization of bundles of links. This means that both inbound and outbound interdomain traffic policies should be further refined, and coordinated with intradomain routing, to optimize the usage of the expensive intercontinental links. The optimization of the distribution of traffic in this scenario usually involves a mixture of manual BGP-based TE techniques, TE based on tweaking the IGP metrics, and/or multiprotocol label switching (MPLS)-TE tunneling. These kinds of settings usually aggravate even more the de-aggregation of IP prefixes as well as the dynamics of the global BGP routing table.

The third scenario consists of a variation on the previous one, where the intercontinental connectivity is supported by aggregation at the synchronous digital hierarchy (SDH) layer (Fig. 1c). The strength of this solution is that it simplifies the layer 3 topology, since multiple links that were formerly connected at layer 3 are now managed as a single trunk. This approach has the potential to reduce the de-aggregation of IP prefixes. The weakness, on the other hand, is that the trunks are frequently composed of heterogeneous circuits, so the loss of granularity at layer 3 makes it extremely hard to fine-tune the distribution of traffic within the trunks.

Overall, these scenarios show the complexity of traffic management for ISPs in LA, dominated by largely manual trial-and-error procedures.

LA Peering Infrastructures and De-aggregation Factor

Table 1 shows the number of Internet exchange points (IXPs) per regional Internet registry (RIR). It is worth highlighting that 12 of the 24 Latin American IXPs are located in Brazil, since the latter gathers a large amount of LA's interdomain traffic. This suggests that TE solutions for LA may differ depending on the geographical area, since the network infrastructure in Brazil is substantially different from the rest of LA. Another characteristic of LA is the lack of significant regional content providers. LA consumes traffic mainly from the United States and Europe, showing only a small exchange of traffic among countries in the region. This limits the interest of large providers in deploying IXPs in LA, and is why most of the transit links are terminated in U.S. NAPs. However, peer-to-peer applications are changing the region's traffic profile, increasing the amount of regional traffic as countries share languages and cultural habits.

As described in scenario #1 in Fig. 1a, the de-aggregation of IP prefixes occurs when a domain advertises CIDR blocks with longer prefixes than those allocated by its RIR. To quantify this, the Internet community has defined the de-aggregation factor (DF), which represents a measure of the current routing table size vs. its aggregated size, and is formally defined as

$$DF = \left(\frac{\text{Prefixes in the Global Routing Table}}{\text{Aggregatable Prefixes}} \right). \quad (1)$$

Up-to-date values for the global routing table size and the DF are shown in Table 1 (columns three and four, respectively). The distribution of the DF in LA is shown in Fig. 2. The figure shows the DF vs. the number of upstream autonomous systems (ASs) for all the ASs registered in the LACNIC region. We observe that the larger DFs come from ASs with a small number of upstream ASs — many of them even with a single upstream AS. This suggests that several of the ASs that are connected to a few transit providers, leak their internal partitioning of prefixes in their BGP advertisements.

The large DF in LA is a consequence of:

- The problem of managing traffic over heterogeneous and complex infrastructures like the ones described in Fig. 1
- The intrinsic limitations of BGP-based TE techniques [8]
- The overloading of IP address semantics [2]

The adverse effects of the permanent increase in the DF are, fundamentally, the processing capacity needed by the routers supporting the global routing table (e.g., stringent memory,

CPU, and power requirements) and the impact on the BGP convergence time, since the larger the routing table, the slower the convergence. Moreover, the practice of adjusting the distribution of interdomain traffic based on the de-aggregation of IP prefixes adversely affects the dynamics of the global routing table.

In the next section we overview a novel approach that can dramatically reduce the size and churn rate of the global routing table, and at the same time offer a promising perspective for dealing with the peculiarities and complexities of interdomain traffic management in LA.

Overview of LISP

In April 2009 the Internet Engineering Task Force (IETF) chartered the LISP Working Group, and experimental Requests for Comments (RFCs) are expected by 2010. LISP [3, 4] uses IP-over-IP tunnels deployed between border routers located at different domains. The IP addresses configured on the external interfaces of the border routers act as RLOC addresses for the end systems in the local domain. Since an AS usually has several border routers, the local EID addresses can be reached through multiple RLOC addresses. LISP separates the address space into two parts, where only addresses from the RLOC address space are assigned to the transit Internet. Therefore, only RLOC addresses are routable through the Internet; EID addresses are considered routable only within their local domain. To illustrate the basics of LISP we use Fig. 3 (extracted from [3]). For a comprehensive understanding of LISP, the reader is referred to [3, 4].

LISP Data Plane

When local end host S with EID address 1.0.0.1 (Fig. 3) wants to communicate with end host D in a different domain whose EID address is 2.0.0.2, the following sequence of events occur in LISP.

The first step is the usual lookup of the destination address E_D in the DNS (E_D corresponds to 2.0.0.2 in the example in Fig. 3). Once E_D is obtained, the packets sourced from E_S traverse the domain and reach one of the local border routers. In LISP the latter are referred to as ingress tunnel routers (ITRs). Since only RLOC addresses are globally routable, when an ITR receives packets toward E_D, it queries the control plane to retrieve the E_D-to-RLOC mapping. After the E_D-to-RLOC mapping resolution, the ITR encapsulates and tunnels packets between the local RLOC address (ITR address 11.0.0.1 in the example) and the RLOC address retrieved from the mapping system, the egress tunnel router (ETR) address in LISP terminology (either 12.0.0.2 or 13.0.0.2 to E_D depending on the mapping). At the destination domain, the ETR decapsulates the packets received through the tunnel and forwards them to E_D — which, as mentioned above, is locally routable within the domain. From the first packet received, the ETR caches a new entry, solving in this way the reverse mapping for the packets to be tunneled back from E_D to E_S.

As shown at the bottom of Fig. 3, the mapping system can return multiple RLOC addresses for the same destination. Each of the entries returned has a *priority* and a *weight* attribute. The priority determines the order in which the ETRs must be selected, while the weight tells how to distribute the traffic among ETRs with the same priority. In the example the priorities

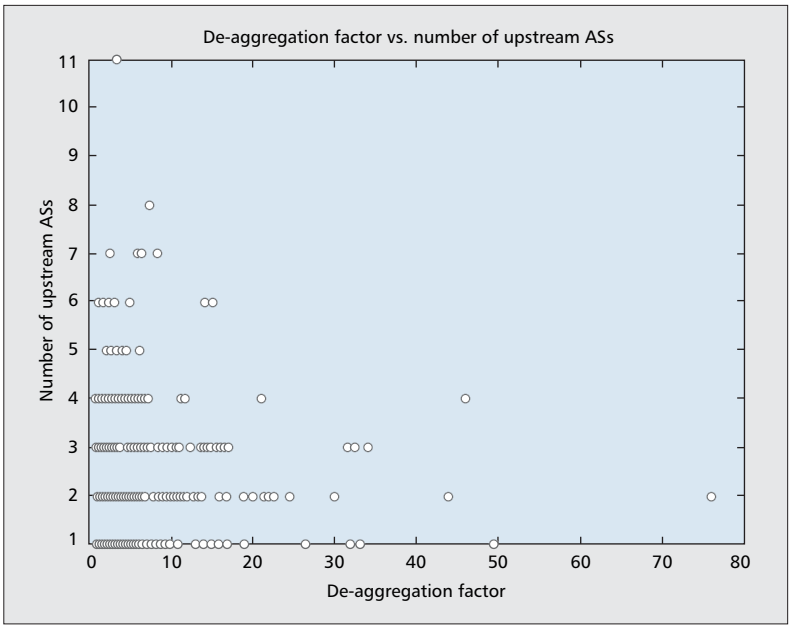

■ **Figure 2.** *Distribution of the de-aggregation factor as a function of the number of upstream providers in Latin America (data of April 2009, extracted from [1] and APNIC [7]).*

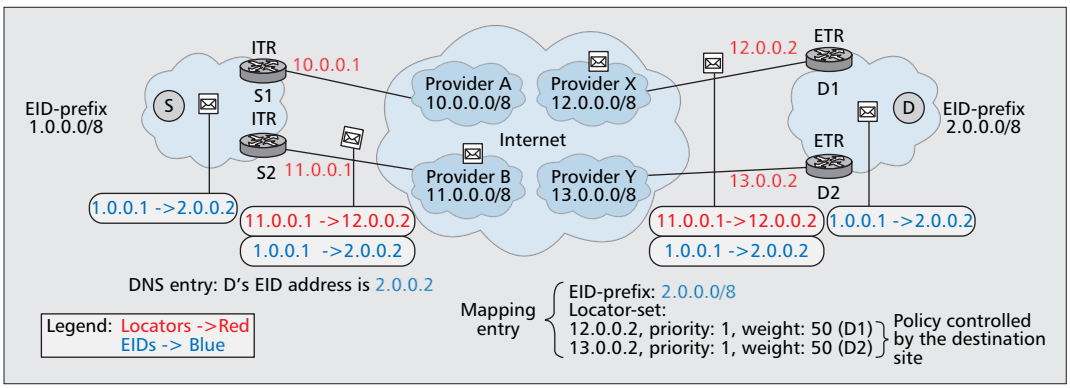

■ **Figure 3.** *The basics of LISP.*

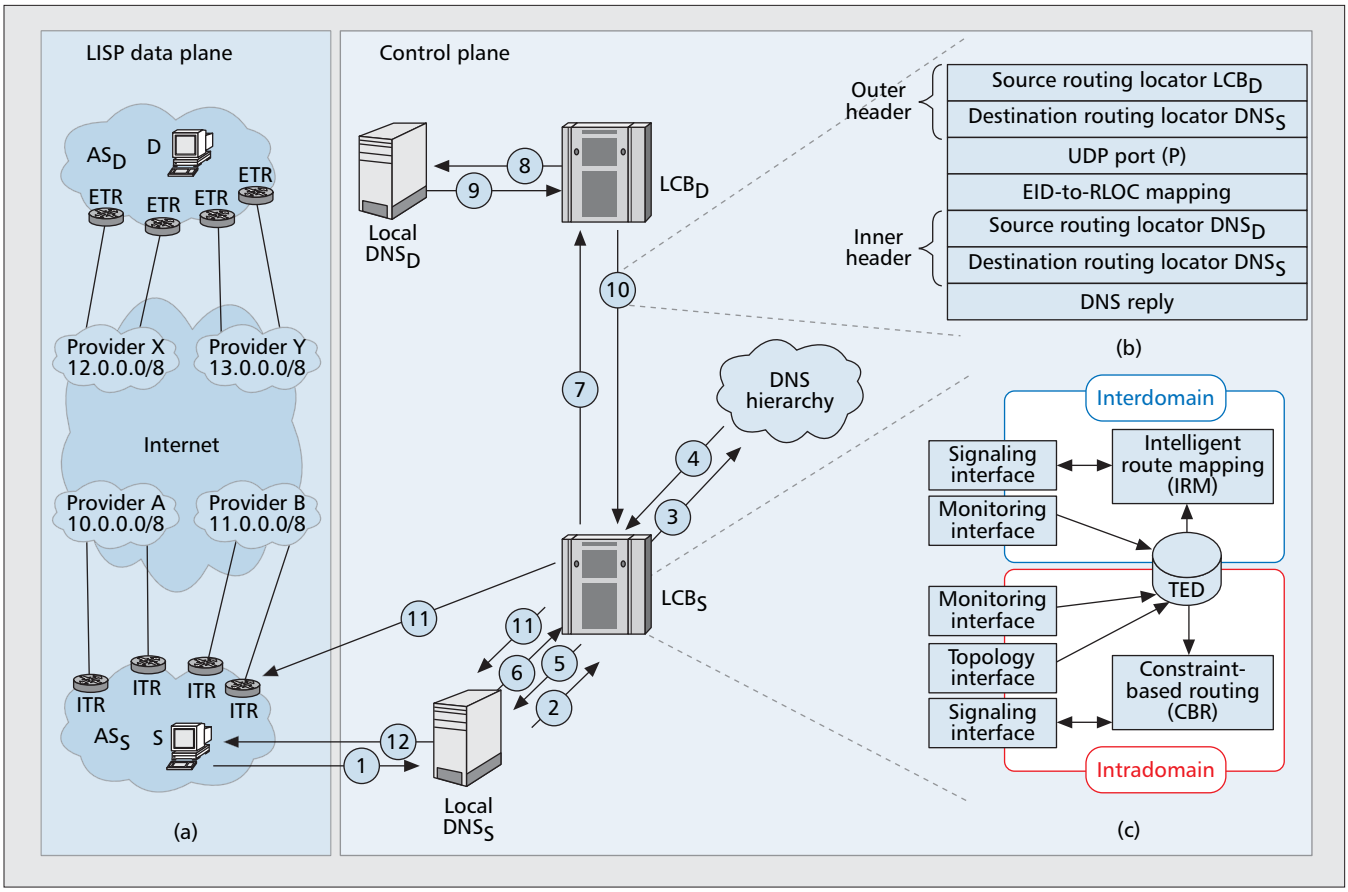

■ **Figure 4.** *Proposed control plane architecture.*

and weights of ETRs D1 and D2 are equal, so the traffic from *S* to *D* will be balanced between D1 and D2.

Overall, LISP has three major advantages. First, it does not introduce major changes to the routing system, and therefore it might be feasible to implement and deploy in the near future. Second, it has the potential to significantly reduce the size of the global routing table [5]. Third, the mapping system brings a wide set of TE opportunities, which in principle, can reach a granularity of a /32 prefix without impacting on the size or dynamics of the global routing table. Nevertheless, it is important to notice that special care must be taken, since LISP might end up moving the scalability issues from the global routing table to the global mapping system.

LISP-Based TE Opportunities for LA

Let us consider now the application of LISP to the three reference scenarios described in Fig. 1.

Scenario #1 — LISP enables multihomed sites to completely avoid running BGP. The only IP prefixes that need to be advertised to the global routing system are those of the wide area network (WAN) interfaces of the border routers (i.e., the ITRs'/ETRs' external addresses). In this framework the load sharing and backup policies of ISP LA01 in Fig. 1a can be managed by intelligent EID-to-RLOC mapping functions, which may be dynamically tuned using the LISP weight and priority attributes, respectively. The different capacities of the links can be used during the mapping resolution process, by appropriately unbalancing the weight attribute of the different links (e.g., the weights advertised could keep the same proportion as the link capacities). The advantages of applying LISP in this scenario are evident, since LA ISPs are released from the burden of handcrafting their BGP configurations. In addition, these advantages can be achieved without de-aggregating prefixes, and therefore without adversely impacting on the size or the dynamics of the global routing table.

Scenario #2 — In this case the TE problem addressed is twofold, since not only must the interdomain mapping function be solved, but also the traffic should be load balanced over heterogeneous links among the upstream providers and the ISP's regional points of presence (POPs). We argue that tools like constraint-based routing (CBR) (for the intradomain part) and an IRM engine borrowing concepts from IRC techniques [6] (for the interdomain part) can work together to accomplish these objectives. At least two appealing solutions can be sketched to drive the intradomain traffic according to the ISP policies. One option is to set up MPLS label switched paths (LSPs) on the fly for different flows, in order to *stitch* the inbound interdomain traffic to the appropriate *pipe* downstream (i.e., to the corresponding ISP POP). This is the usual task of the PCE [9]. The second option is to introduce the idea of interior

ITR/ETRs, which can perform internal mapping from/onto border routers.

Scenario #3 — Similar to the previous scenario, a combination of CBR and IRM may apply here as well. Note that the utilization of SDH bundles suggests that cross-layer TE techniques might also be needed.

LISP Control Plane

The basic role of the control plane is to provide the mapping system. Among the proposals under discussion are LISP-ALT, LISP-CONS, and LISP-DHT [4, references therein]. At present, the more developed proposals for a LISP control plane present at least three problems. First, the initial packets sent from a source EID address (E_S) to a destination EID address (E_D) can be dropped at the ITR during the EID-to-RLOC mapping resolution. Although caching techniques are being proposed to store the mappings at ITRs, a hit might not necessarily be found, because the mapping either has aged out or simply was never requested before.

Second, LISP might considerably increase the latency to start up the communication between end systems. Considering the one-way delay (OWD), a TCP connection between two end systems is established roughly around

$$T_{\mathrm{DNS}} + 2OWD(E_S,E_D) + OWD(E_D,E_S), \quad (2)$$

whereas with LISP (Fig. 3) it would roughly demand

$$T_{\mathrm{DNS}} + T_{\mathrm{resol}}^{\mathrm{map}} + 2OWD(E_S,E_D) + OWD(E_D,E_S) \quad (3)$$

under the usual assumption that ITRs and ETRs can encapsulate/decapsulate at line rate.

Third, for each traffic flow from S to D, the egress ITR is also used as the local ETR for the packets sent from D to S. This is to avoid a two-way mapping resolution, which would increase even more the latency shown in Expression 3. Clearly, this introduces a limitation in terms of inbound TE, which is particularly important in the case of LA, given that outbound and inbound traffic management policies typically do not match.

To cope with these issues, alternative solutions are being discussed. These alternatives require either some major changes to the DNS system, the addition of some debatable features to border routers, or using the control plane to transport data while the mapping is being resolved.

Control Plane Proposal

Our goal is threefold. First, we aim to prevent the potential dropping of packets at the ITRs while the EID-to-RLOC mapping resolution is being computed. Second, we aim to obtain and configurie the corresponding mapping during the normal DNS resolution process for destination E_D, and we want to achieve this goal without introducing changes to the DNS system. More precisely, we seek

$$T_{\mathrm{DNS}} + T_{\mathrm{resol}}^{\mathrm{map}} \approx T_{\mathrm{DNS}}. \quad (4)$$

Third, we aim to have the TE flexibility to choose different local ITR and ETR LISP routers for any given flow sourced at the domain.

Architecture

To pursue these objectives, we propose the scheme and set of steps depicted in Fig. 4. We introduce an entity we call a LISP control box (LCB), which might be a standalone device or run as an instance of a PCE, implemented as described in [10] and tailored for this purpose.

Step 1 — S queries DNS_S to obtain the EID address of destination D. The LCB (LCB_S) obtains the EID address of S (E_S) by interprocess communication (IPC) with DNS_S (e.g., using sockets) and computes the local RLOC to be used for the reverse mapping (i.e., for the incoming traffic from E_D to E_S) based on TE constraints. The algorithms used to determine the ingress RLOC (RLOC_S) are inherently the same used today by IRC techniques [6]. The E_S-to-RLOC_S mapping computation is performed by the IRM module of LCB_S.

Steps 2–9 — The LCBs are in the data path of the DNS servers; therefore, they can transparently analyze the exchange of DNS messages. Steps 2–9 represent the usual flow of iterative DNS queries performed by DNS_S, and the corresponding replies received from the DNS servers in the hierarchy (root server, top-level domain server, etc.). It is important to observe that these steps require the use of iterative queries between DNS servers. This approach is in line with the current trend of avoiding recursion between DNS servers. Indeed, supporting iteration is mandatory, whereas recursion is strictly optional (by default BIND9 performs iterative queries).

Step 10 — When LCB_D detects that the reply issued from DNS_D carries the address E_D, it encapsulates the reply into a UDP message; clearly variants of this approach can work as well-with source address LCB_D, destination address DNS_S, and a special transport port P that will be listened to by LCB_S at the source domain S (Fig. 4b). The payload of the outer packet contains the EID-to-RLOC mapping for E_D. It is worth highlighting that the mapping selection performed at LCB_D is precomputed by an online IRC engine running in the background (as mentioned in step 1, this computation is performed by the IRM module shown in Fig. 4c), so the mapping is always known beforehand. This means that LCB_D can encapsulate the answer from DNS_D roughly at line rate.

Step 11 — LCB_S is in the data path of DNS_S, so when LCB_S detects a packet toward DNS_S using port number P, it intercepts and decapsulates the packet and forwards the usual DNS answer to DNS_S. From the outer packet LCB_S discovers the address of LCB_D, retrieves the mapping for E_D, and configures through the signaling interface of the LCB all the ITRs according to that mapping. The advantage of pushing the mapping to all ITRs is that AS_S can carry out local TE actions and move part of its internal traffic without caring whether a mapping will

To cope with these issues, alternative solutions are being discussed. These alternatives either require some major changes to the DNS system, the addition of some debatable features to border routers, or using the control plane to transport data while the mapping is being resolved.

To have a preliminary evaluation of the performance of our control plane, we have prototyped the LCB in a standard Linux PC. For simplicity, our prototype is built into BIND9, meaning that the results shown here were obtained by integrating the LCB functionality into the DNS.

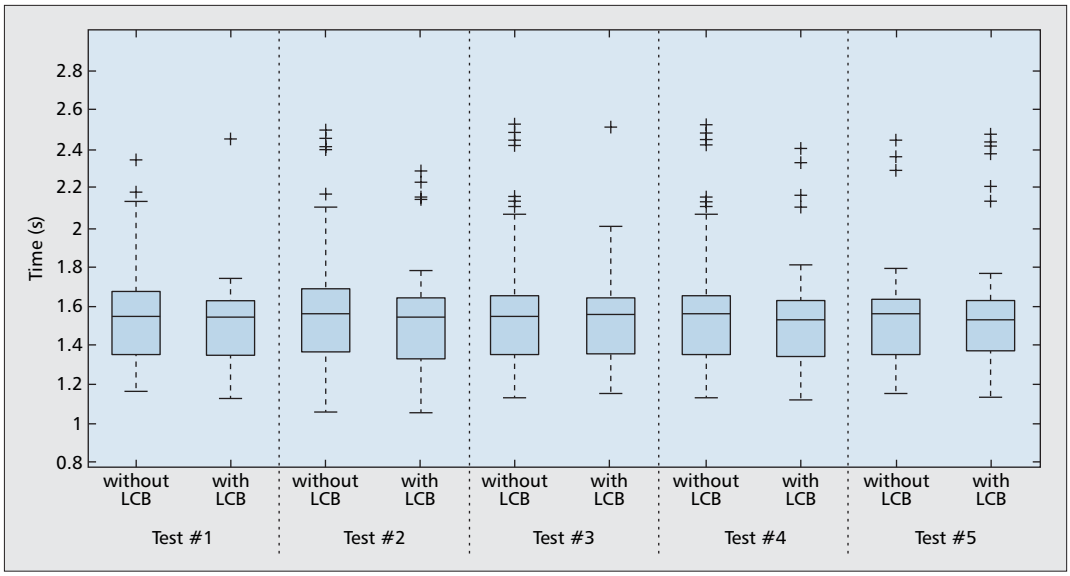

Figure 5. *Five tests showing the time distribution of a set of 1000 DNS lookups over the Internet. Each test corresponds to a round of 200 DNS lookups, 100 without LCBs, and 100 with LCBs.*

be in place in the relevant ITRs after TE optimization. The mapping information pushed to the ITRs consists of the tuple (E_S,E_D, $RLOC_S$, $RLOC_D$), supporting the utilization of two independent one-way tunnels depending on the reverse mapping computed by LCB_S during step 1. In other words, an ITR is capable of forwarding traffic to E_D, using as source address in the encapsulation an RLOC address that might be different from the addresses configured on its WAN interfaces (e.g., using source NAT) — although restricted to the pool of addresses assigned by the ISP to which the ITR is connected. Like in Fig. 4a, this approach helps AS_S exploiting the usual multiconnectivity to each of its providers.

Step 12 — DNS_S responds the usual DNS query to S.

After the usual DNS resolution process, E_S starts sending packets toward E_D, with the advantage that the mapping has already been configured at the ITRs, hence avoiding the potential dropping of packets. As we shall show in the next section, the overall process (steps 1–12) can be completed in approximately T_{DNS}, which we claim should be used as the upper bound for solving the mapping.

When the first data packet reaches the corresponding ETR in AS_D, the ETR:

1 Decapsulates the packet and forwards the inner packet to E_D
2 Obtains the reverse mapping (i.e., the E_S-to-$RLOC_S$ mapping)
3 Pushes this mapping to the rest of the ETRs and also updates the LCB_D mapping system via multicast or another mechanism

This action completes the two-way mapping resolution process. An interesting point is that our control plane allows each domain to achieve its TE policies congruently, since each domain has the freedom to independently decide its ingress and egress mappings.

In order to achieve the aforementioned goals, the IRM engine of the LCB needs to know the local topology and the availability of local resources. As shown in Fig. 4c, a promising approach is to provide a highly efficient coupling between the IRM engine of the LCB and the PCE — the basic components of a PCE are represented in the intradomain portion of Fig. 4c. Indeed, both can share and feed the TE database (TED), which is crucial to accomplish both intradomain and interdomain TE objectives concurrently. The LA scenarios shown in Figs. 1b and 1c can benefit considerably from this approach, since traffic can be coordinatedly balanced by the LCB and PCE, making possible the selection of a upstream/downstream forwarding path with high granularity (e.g., for prefixes or even end systems).

Preliminary Results

In order to have a preliminary evaluation of the performance of our control plane, we have prototyped the LCB in a standard Linux PC. For simplicity, our prototype is built into BIND9, meaning that the results shown here were obtained by integrating the LCB functionality into the DNS. More precisely, in Fig. 4 DNS_S and LCB_S are integrated in a single Linux PC, and the same applies to DNS_D and LCB_D. It is worth emphasizing that the control plane architecture proposed in this article is more general in scope; thus, more advanced implementations of the LCB prototype may not be embedded in BIND9 and may run as standalone devices.

Our initial goal is to verify that Eq. 4 holds when the LCBs are used. To this end, we performed a set of 1000 rounds of DNS lookups on the Internet. These rounds were clustered in five different tests (Fig. 5), each corresponding to a sequential set of 200 DNS lookups, half of them without the intervention of the proposed control plane and the other half utilizing our prototype. The five tests were carried out at different moments in time, where we measured the total time of the DNS lookups from the source E_S.

In our tests E_S and DNS_S/LCB_S were located at the University of the Republic premises in Montevideo, Uruguay, and E_D as well as DNS_D/LCB_D were located at the Technical University of Catalonia premises in Barcelona, Spain.

In this setting steps 2–9 in Fig. 4 are solved iteratively using the usual DNS hierarchy. As mentioned above, our goal is to assess if the control plane is able to obtain and configure the mapping during the usual DNS resolution process. Therefore, prior to each of the 1000 tests (i.e., both with and without LCBs), we flushed the DNS_S cache, enforcing in this way the lookup of E_D through the DNS system. This is only to ensure that during the trials, the validation of Eq. 4 is not biased by the DNS_S cache. It is important to observe that in an operational scenario, if E_D is cached in DNS_S, a mapping entry will already be configured at the ITRs in AS_S, so Eq. 4 will trivially hold; analysis of the potential interdependencies between the DNS and LISP caches and their aging policies is part of our future work.

After the steps 2–9 in Fig. 4, and upon detection of authoritative responses for A, AAAA, and MX record types in DNS_D, we obtain the EID-to-RLOC mapping for E_D (as mentioned before this is computed aforehand), and then proceed with steps 10–12.

The results for 1000 experiments are shown in Fig. 5. The horizontal mark inside each box is the median, the edges of the boxes are the 25th and 75th percentiles, and the whiskers cover the rest of the data gathered, excluding outliers. The outliers are represented as crosses and they are shown individually. Our preliminary results are promising, since Fig. 5 confirms that a simple prototype implementation of the control plane is able to accomplish the goal targeted in Eq. 4.

Discussion and Concluding Remarks

In this article we have presented the peculiarities of LA network infrastructures, and we have analyzed the TE opportunities that LISP may bring to the region, especially by developing customized intelligent route mapping techniques.

Our main contribution is the proposal of a control plane that tackles the challenges exposed in LISP. Among the strengths of this control plane are the following. First, the mappings can be configured during the normal DNS resolution process, and this can be implemented without introducing changes to the DNS system; the zone files, resource records, and so on remain unchanged. Second, by placing the LCBs in the data path of the DNS system, it is possible to transparently use the latter as a discovery mechanism of LCBs. Once a remote LCB is discovered, direct communication and even cooperation among LCBs might be exploited, adding therefore, extra capabilities to the control plane. Indeed, our control plane can support an overlay of LCBs, enabling on-the-fly refinement of the mappings toward (from) a set of popular destinations (sources) of a domain. Third, our control plane is technically supported by an architecture that blends two ongoing initiatives in the IETF, the PCE and LISP, which offers a promising perspective for LA, especially for the scenarios in Fig. 1.

Although the proposed control plane offers a promising approach, several aspects of the architecture need to be further explored. For example, both the DNS servers and the LCBs are end systems whose IP addresses are, in principle, EIDs (i.e., they are non-globally routable). This states the obvious problem that a *resolver of resolvers* is required. An option is to assign globally routable addresses to both of them, but arguments can be found against this too. In any case, the debate about the assignment of the address space to DNS servers is present in any locator-identifier separation scheme and must be appropriately addressed.

Issues such as how to solve mappings that, in principle, do not require a DNS resolution (e.g., ping 10.10.10.1), and the potential impact of these solutions on the overhead and dynamics of the DNS system need to be analyzed and evaluated. Other important issues to be explored are the security aspects of this control plane.

Issues such as how to solve mappings that, in principle, do not require a DNS resolution (e.g., ping 10.10.10.1), and the potential impact of these solutions on the overhead and dynamics of the DNS system, need to be analyzed and evaluated.

References

[1] CIDR Report; http://www.cidr-report.org/as2.0/
[2] D. Meyer, L. Zhang, and K. Fall, "Report from the IAB Workshop on Routing and Addressing," IETF RFC 4984, Sept. 2007.
[3] D. Meyer, "The Locator Identifier Separation Protocol (LISP)," *Cisco Internet Protocol J.*, Mar. 2008, pp. 23–36.
[4] D. Farinacci *et al.*, "Locator/ID Separation Protocol (LISP)," work in progress, May. 2009; IETF draft-ietf-lisp-01.txt.
[5] B. Quoitin *et al.*, "Evaluating the Benefits of the Locator/Identifier Separation," *Proc. ACM SIGCOMM MobiArch*, Kyoto, Japan, Aug. 2007.
[6] M. Yannuzzi *et al.*, "Improving the Performance of Route Control Middleboxes in a Competitive Environment," *IEEE Network*, vol. 22, no. 5, Sept./Oct. 2008, pp. 56–64.
[7] BGP Routing Table Analysis; http://thyme.apnic.net/
[8] R. Gagliano, "IPv4 De-aggregation in LACNIC Region-LACNIC XI," May 2008; http://www.lacnic.net
[9] A. Farrel, J. P. Vasseur, and J. Ash, "A Path Computation Element (PCE)-Based Architecture," IETF RFC 4655, Aug. 2006.
[10] E. Grampín and J. Serrat, "Cooperation of Control and Management Plane for Provisioning in MPLS Networks," *9th IFIP/IEEE Int'l. Symp. Integrated Net. Mgmt.*, Nice, France, May 2005.

Biographies

MARCELO YANNUZZI (yannuzzi@ac.upc.edu) received a degree in electrical engineering from the University of the Republic (UdelaR), Uruguay, in 2001, and D.E.A. (M.Sc.) and Ph.D. degrees in computer science from the Department of Computer Architecture, Technical University of Catalonia (UPC), Spain, in 2005 and 2007, respectively. He is with the Advanced Network Architectures Laboratory (CRAAX) at UPC, where he is an associate professor. He held previous positions with the Physics Department of the School of Engineering, UdelaR (1997–2003), and the Electrical Engineering Department of the same university (2003–2006). He worked in the industry for 10 years at the national telco in Uruguay (1993–2003).

EDUARDO GRAMPIN (grampin@fing.edu.uy) holds a degree in electrical engineering from UdelaR, and he got his Ph.D. from UPC in 2005. Currently, he is an associate professor at the Computer Science Department of the School of Engineering, UdelaR. Between 1993 and 2006 he worked at the National Telecommunications Administration (ANTEL) in Uruguay. He participated in the deployment of commercial Internet services in Uruguay in the mid–1990s,

and has worked on different aspects of planning and management of IP networks in ANTEL, mainly linked to projects toward a convergent architecture based on IP/MPLS, with an integrated management system.

ROQUE GAGLIANO (rgaglian@fing.edu.uy) received his electrical engineering degree from UdelaR in 2001, and his M.S. in electrical engineering from the University of Kansas in 2005. He worked as an area engineer in ANTEL between 1999–2003 and 2006–2008. He held previous positions with the Physics Department, UdelaR (1997–2002), Sprint Nextel Corp, Kansas City, Missouri (2005–2006), and South-East and ITTC Engineering, University of Kansas (2004–2005). He is now with LACNIC in Uruguay, and since 2002 he is also an assistant professor in the Electrical Engineering Department, UdelaR.

ALBERTO CASTRO (acastro@ac.upc.edu) got his engineering degree in computer science from UdelaR in 2007. Since 2008 he is a Ph.D. student working at CRAAX, UPC. He is an assistant professor in the Computer Science Department of the School of Engineering, UdelaR. His research interests are in IP/MPLS interdomain routing, peer-to-peer networks, network performance, and interdomain traffic engineering.

MARTIN GERMAN (mgerman@ac.upc.edu) received his computer science degree from UdelaR in 2007. He is a Ph.D. student at CRAAX, UPC. He is also an assistant professor in the Computer Science Department of the School of Engineering, UdelaR. His research interests are in the design of scalable routing paradigms for the Internet, novel routing/switching paradigms, network performance, and traffic engineering.

XAVI MASIP-BRUIN (xmasip@ac.upc.edu) received M.S. and Ph.D. degrees from UPC, both in telecommunications engineering, in 1997 and 2003, respectively. He is with CRAAX at UPC, where he is currently an associate professor in computer science. Since 2000 he has participated in many international and Spanish research projects. His publications include around 60 papers in national and international refereed journals and conferences. His current research interests lie in broadband communications, QoS management and provision, and traffic engineering.

MARKET FOCUS

Austria Centre, Vienna
Exhibition 21-23 September

In conjunction with

ECOC Exhibition 2009
Austria Centre
21-23 September 2009

Market Focus is back for 2009, and this year we are working to make this the most informative and prestigious Market Focus yet!

To create the highest quality, most valuable and informative content for attendees at this year's event, a high profile line-up of industry experts have been appointed to form an Advisory Committee representing the breadth of technologies in the optical communications space. The Committee includes representatives from telecom operators, system vendors, component manufacturers, industry analyst groups, and independent industry bodies.

The seven committee members have been confirmed as **Stephan Neidlinger** from *ADVA Optical Networking*, **Chris Clarke** from *Oclaro*, **Philippe Chanclou** from *France Telecom*, **Joeri Van Bogaert** *from Leoni*, **Michael Lebby** from the *OIDA*, **Daryl Inniss** from *Ovum*, and **Glenn Wellbrock** from *Verizon*.

2008 Speakers included

Magnus Breidne, CEO, Ignis Photonyx

Alain Couder, President and CEO, Bookham, Inc.

Dana Cooperson, VP Optical Networks, Ovum RHK

Kim Papakos, Senior Product Manager, Tellabs

For further information visit
www.ecocexhibition.com

Market Focus 2009 Topics

- **40 and 100 Gbit/s Ethernet Systems**
- **Optical Solutions for 3 and 4G Mobile Networks**
- **Mass-Market Broadband Fibre Access**
- **Role of Optical Comms in Green Telecoms**
- **Next-Generation Services**
- **Optical Network Agility**
- **Packet Optical Transport**

ACCEPTED FROM OPEN CALL

Recent Advances in Amplify-and-Forward Two-Hop Relaying

Stefan Berger, Marc Kuhn, and Armin Wittneben, ETH Zurich
Timo Unger and Anja Klein, Technische Universität Darmstadt

ABSTRACT

In this article we review an important class of wireless cooperation protocols known as amplify-and-forward relaying. One or more low-complexity relay nodes assist the communication between sources and destinations without having to decode the signal. This makes AF relaying transparent to modulation and coding of the source/destination communication protocol. It is therefore a highly flexible technology that also qualifies for application in heterogeneous networks comprising many nodes of different complexity or even standards. Recently, two-way relaying was introduced, which is readily combined with AF relaying. It is a spectrally efficient protocol that allows for bidirectional communication between sources and destinations. In order to investigate the potential of wireless AF relaying, we introduce three distributed network scenarios that differ in the amount of cooperation between nodes. New challenges arise in those networks, and we discuss approaches to overcome them. For the most general case of a completely distributed system, we present coherent relaying solutions that offer a distributed spatial multiplexing gain even for single-antenna nodes. Based on real-world experiments, we validate the feasibility of all schemes in our laboratory.

INTRODUCTION

Future wireless communication systems will be built for cooperation rather than mere coexistence. Cooperative communication is a hot topic of current research, and many people believe it to be the next big step after multiple-input multiple-output (MIMO) systems. The basic idea is that multiple nodes cooperate in order to increase the link quality, reliability, and data rate of the system. In the future, the density of active nodes competing for a common wireless channel in cellular as well as access or ad hoc networks will increase significantly. Therefore, node cooperation is an efficient means of achieving these gains. An overview of several cooperative diversity protocols for wireless networks can be found, for example, in [1]. Either user nodes or dedicated terminals may assist the communication between sources and destinations. We refer to both types of assisting nodes as relays. In the literature they are classified as either *full-duplex* or *half-duplex*. Full-duplex relays can simultaneously transmit and receive, whereas half-duplex relays cannot. As an example, the nodes in full-duplex frequency-division duplex (FDD) systems transmit and receive signals at the same time but use different frequency channels.

In the following we focus on half-duplex relays. However, the relaying concepts discussed in this article can also be used in full-duplex FDD systems. The typical approach to orthogonalize channel resources in a half-duplex system is to use time-division duplex (TDD) where reception and transmission at the relay are orthogonalized by using two orthogonal time slots:

- Transmission from source to destination and relay
- Transmission from relay to destination

Hence, a bidirectional transmission between two nodes via a relay requires four time slots to exchange only two messages. To increase the spectral efficiency of such bidirectional communication, the authors in [2] proposed a scheme — known as *two-way relaying* — that reduces the number of time slots to two: In the first time slot both nodes transmit their messages simultaneously to the relay. In the second time slot the relay transmits a combined version of the received signals to both nodes. Since each node knows its own transmitted signal, it can subtract the back-propagated *self-interference* prior to decoding. In a full-duplex FDD system, the transmission from the nodes to the relay and from the relay back to the nodes would take place at the same time, but on different frequency channels.

The general relay channel of conventional one-way relaying was first investigated in [3]. One source transmits to one destination while a single relay assists. The capacity of the general relay channel is still not known. Moreover, there is not even a cooperation strategy known that works best for this general case. At least two different basic signal processing strategies are distinguished at the relay. The first strategy, known as decode-and forward (DF) protocol, involves decoding of the source transmission at the relay. The re-encoded and possibly compressed signal

0163-6804/09/$25.00 © 2009 IEEE

is then forwarded to the destination. In terms of sum rate, this protocol is close to optimal when the source-relay channel is excellent, which is practically the case when source and relay are physically close or when dedicated relays are placed intentionally in such a way that a good connection to the source is ensured.

For the second strategy — known as compress-and-forward (CF) or *quantize-and-forward* protocol — the relay does not decode the source signal, but uses its observations in a different way. The received signal is quantized (and possibly compressed) and then forwarded to the destination. This protocol is most efficient in cases where the source-relay and source-destination channels are of comparable quality, and the relay-destination link is good. In this situation the relay may not be able to decode the source signal, but nevertheless has an independent signal observation that can be used to assist the decoding at the destination. Finally, the amplify-and-forward (AF) protocol is a special case of the CF strategy where the signal processing at the relay is only linear. For a multi-antenna relay this means that each transmit antenna may forward a linear combination of the received signals of all receive antennas. In [4] the authors develop capacity scaling laws for such networks and propose protocols that achieve this capacity scaling in the limit for large numbers of relays.

In this article we focus on the use of the AF protocol for applications where the complexity of all involved nodes differs significantly. Based on three scenarios described later, we discuss the achievable gains as well as challenges arising for different signal processing schemes at the relays. Another section also contains results obtained by a real-world demonstrator that show the feasibility of AF relaying in all three scenarios. Finally, we conclude this article by naming open issues and giving an outlook on future research.

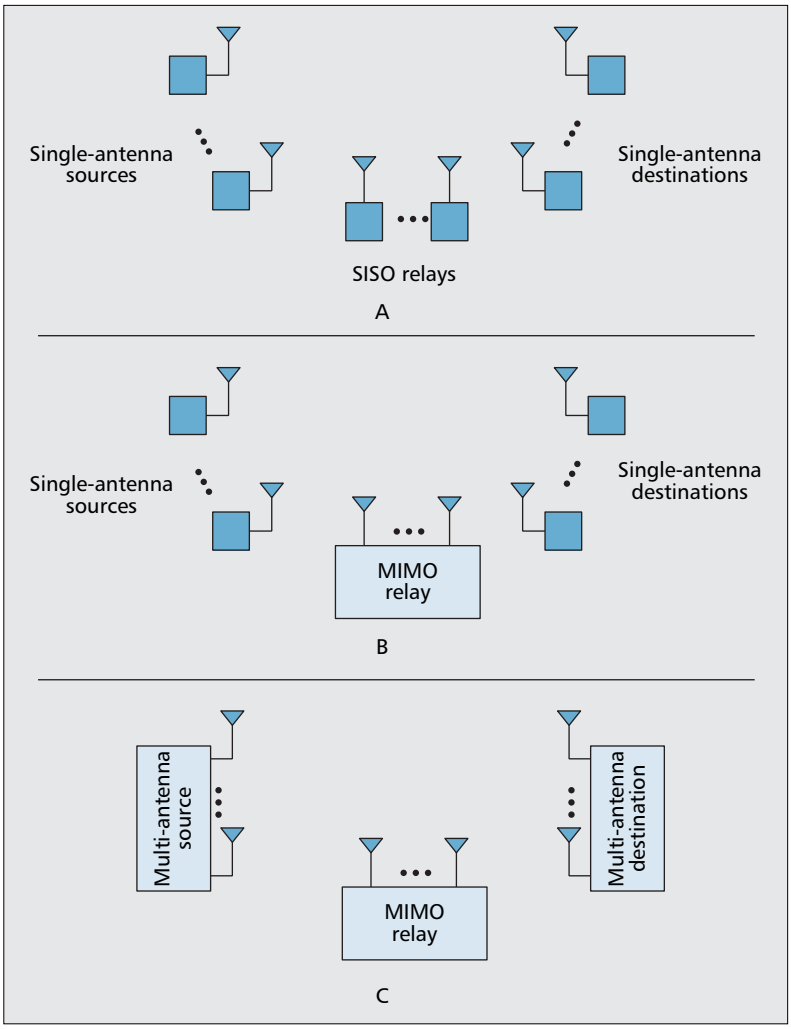

Figure 1. *Amplify-and-forward relaying scenarios with an amount of cooperation that increases from A to C: from distributed nodes and relays in scenario A to multi-antenna nodes and relays in scenario C.*

Scenarios and System Model

We investigate three scenarios, as depicted in Fig. 1, that differ in the complexity of the participating nodes. In scenario A we consider a distributed wireless network with single-antenna nodes and single-antenna AF relays. N sources and N destinations communicate with the help of M relays (multiuser relaying). Each source communicates with a single dedicated destination terminal, together forming a source/destination pair. A distributed spatial multiplexing gain can be achieved if the AF relay gains are chosen appropriately. This essentially allows multiple source/destination pairs to communicate concurrently over the same physical channel. Note that DF relaying is not applicable in this scenario because due to the strong multiuser interference, the single-antenna relays are not able to efficiently decode the data streams of all sources. In scenario B a single multi-antenna relay assists the communication between N single-antenna source/destination pairs. Because all relay antennas can fully cooperate, the number of available degrees of freedom increases compared to scenario A, and efficient precoding schemes can be implemented at the relay. A typical application could be a sensor network with one central node of high complexity. Finally, scenario C comprises one multi-antenna source/destination pair and a multi-antenna relay. This configuration could correspond to a single link in a cellular system with high-data-rate bidirectional communication.

With increasing cooperation between the relay antennas (scenario B) and between the source and destination antennas (scenario C), increasing performance gains are available. In this sense, scenario C constitutes an upper bound on the performance of scenario B. Note that for scenarios A and B, the extension to the case where some of the sources and destinations employ multiple antennas is straightforward.

Two-way relaying can basically be used in all three scenarios. In [5] it has been extended from a single-antenna configuration to the MIMO case for DF relaying, in [6] to MIMO AF relaying. Scenario C allows very efficient implementation of two-way relaying, because there is no interuser interference, and the self-interference can be completely subtracted from the data stream at both nodes. In contrast to that, the relay has to suppress $2N - 1$ interferers in scenarios A and B.

The mutual information is a measure of qual-

The maximum achievable distributed spatial multiplexing gain, which enters the sum rate expression 1 as pre-log factor, lies in the order of the number of source/destination pairs. If the relay gain factors are chosen properly, this can even be realized with low-complexity single-antenna nodes.

ity for the communication link. For a single channel realization and source/destination pair, it is given by [7]

$$I = \frac{1}{2}\log_2 \det(\mathbf{I} + (\mathbf{R}_\mathrm{I} + \mathbf{R}_\mathrm{N})^{-1}\mathbf{R}_\mathrm{S}), \qquad (1)$$

where $\mathbf{R}_\mathrm{I}$, $\mathbf{R}_\mathrm{N}$, and $\mathbf{R}_\mathrm{S}$ are the covariance matrices of the interference (if present), noise, and signal 1 at the destination, respectively. The pre-log factor 1/2 comes from the fact that a two-hop transmission cycle is assumed to require two time slots of equal length. Averaging the mutual information over many channel realizations delivers the average rate.

Scenario A

In classical wireless multiple access scenarios, the amount of resources available for each individual user reduces with the number of terminals sharing the channel. Coherent multiuser relaying is a promising approach to break this paradigm by allowing multiple users to communicate concurrently on the same physical channel (distributed spatial multiplexing gain). Although we lose spectral efficiency by requiring two time slots for each transmission cycle, large gains are feasible for more than two source/destination pairs. The maximum achievable distributed spatial multiplexing gain, which enters the sum rate expression 1 as pre-log factor, lies in the order of the number of source/destination pairs. If the relay gain factors are chosen properly, this can even be realized with low-complexity single-antenna nodes.

Relay Gain Allocation

In [8, references therein] coherent gain allocation schemes that achieve a distributed spatial multiplexing gain are discussed. One approach to allow multiple users to access the channel simultaneously is to compute the relay gain factors such that the source/destination streams are completely orthogonalized in space (multiuser zero-forcing [MUZF] relaying). The complex baseband signals are scaled and rotated at the relays in such a way that interuser interference is suppressed completely. A necessary but not sufficient condition for this to be possible is $M > N(N-1)$. Note that if we allow for relays with multiple antennas, significantly fewer relays are required. We refer to the case that $M = N(N-1)+1$ as minimum *relay configuration*. The gain vector is then (apart from a complex-valued scaling factor) uniquely determined. If the number of relays is larger than the minimum relay configuration, there is room for further optimization (e.g., to achieve additional diversity gain).

Another coherent gain allocation scheme that achieves a distributed spatial multiplexing gain is based on a minimum mean squared error (MMSE) criterion. This approach achieves even better performance than MUZF but requires additional knowledge of the noise covariance matrices in order to compute the relay gain factors. Furthermore, since a small amount of interuser interference is tolerated, no minimum number of relays is required. Performance instead degrades smoothly for decreasing M, which makes this scheme also suitable for smaller system configurations.

Major Challenges in Distributed Wireless Networks

Compared to classical point-to-point communications links, new challenges (e.g., distributed time synchronization) arise if multiple distributed nodes want to cooperate. However, since AF relaying is transparent to modulation and coding, the relays need only to be synchronized on a time slot — not a sample — basis. This is a great alleviation of a major demand. Furthermore, in order to coherently forward the signals, the relays require channel knowledge. For the gain allocation schemes described above (MUZF and MMSE relaying) they have to know the complete first-hop and second-hop channel matrix. This involves channel estimation and dissemination of the estimates between all relays so that the gain factors can be computed locally at the relays. Although this is a major effort, in low-mobility scenarios the gain largely outweighs the cost. Finally, it was widely believed that all relays require a global phase reference in order to be able to coherently forward the signal. We showed that this is not true if a basic principle is followed when the channels are estimated. Our results have been proven on a real-world demonstrator, where we successfully implemented the MUZF and MMSE gain allocation schemes [9].

Channel Estimation — In order to provide all relay nodes with knowledge of the first-hop and second-hop channels (we speak of "global channel state information"), a *channel estimation* and a *dissemination phase* are required. If the relays are not phase synchronous, it is important to estimate all channels in the same "direction." Consider a point-to-point communication situation where node A wants to transmit data to node B. The propagation channel between the two nodes can be estimated in two directions: either node A transmits a training sequence to node B that estimates the channel ("forward direction"), or B transmits the training sequence and the estimation is done by A ("backward direction"). Although the wireless propagation channel between the two nodes is reciprocal, the equivalent baseband channels are generally not if A and B are not phase-synchronous. This is because the transmission involves shifting the signal to and from passband. The current (random and unknown) local oscillator (LO) phases of the nodes enter the signal as a positive (shifting from baseband to passband) or negative (shifting from passband to baseband) phase rotation. Consequently, in order to obtain knowledge of the forward (backward) channel at node A (B), it has to be estimated at B (A) and fed back to A (B). Likewise, the first-hop and second-hop channel coefficients in the two-hop scenario have to be estimated all in forward direction (from sources to relays and from relays to destinations) or all in backward direction (from destinations to relays and relays to sources). The estimated product channels (i.e. the product of estimated first-hop and second-hop channel matrices) are then independent of the LO phase offsets of the relays. For MUZF and MMSE relaying this means that the computed gain factors are also independent of the LO phases.

Once the channel coefficients are estimated, they have to be disseminated to the relays. The number of required channel uses is $2NM$. In order to broadcast a channel coefficient, the respective node has to transmit the signal at a rate that all relays are able to decode. A short-range secondary system (e.g., Bluetooth or ultra-wideband (UWB)) could provide a very efficient means of exchanging the required channel state information (CSI).

LO Imperfections and Phase Noise — Another crucial issue that arises in distributed networks with multiple relays is LO accuracy. In contrast to point-to-point communication systems or single-relay networks (e.g., scenarios B and C), the unknown and random LO phases at the relay terminals pose a potential problem for networks where multiple relays want to coherently forward the signal at the same time. The individual relays' LO phase enters the signal in the mixer, where the received passband signal is shifted to baseband. Fortunately, this phase error is compensated if the signal is mixed to passband again *given that the LO phase has not changed in the meantime*. Thus, phase stability is a crucial point. In practical systems, LO phase noise introduces a random change of the LO phases during the time between reception and retransmission, thus destroying coherency between the relays. In general, the phase uncertainty increases with time (e.g., Wiener phase noise model, which is widely used in orthogonal frequency-division multiplexing [OFDM] literature). Therefore, the larger the time difference between reception and retransmission, the larger the potential phase error.

In practical systems the LO of each node also exhibits a frequency offset. Fortunately, the error introduced by a carrier frequency offset when mixing the signal to baseband at a relay is compensated for when mixing it to passband again.

Performance Evaluation and Demonstration

In this section we assess the performance gain of two-hop multiuser relaying in scenario A over a classic single-hop (point-to-point) reference scenario, where N single-antenna source/destination pairs communicate in a time-division duplexing (TDD) manner. In both cases the total transmit power for each transmission cycle (i.e., two time slots for the relaying case and N time slots for the point-to-point reference case) is the same. In the two-hop case, the sum transmit power of all sources is equal to the sum transmit power of all relays. For the Matlab simulation results, all channels are frequency-flat and subject to i.i.d. Rayleigh block fading of unit variance.

In the simulation results, the same transmit power is used for both the relaying and reference scenarios. The resulting average sum rates are then plotted vs. the receive signal-to-noise ratio (SNR) of the reference case (simply denoted SNR in Fig. 2). First, we show results of Monte Carlo simulations in MATLAB, where the simulation environment is perfectly defined. Then we prove that our findings regarding phase noise and channel estimation hold in practical systems by providing performance results of the real-world demonstrator.

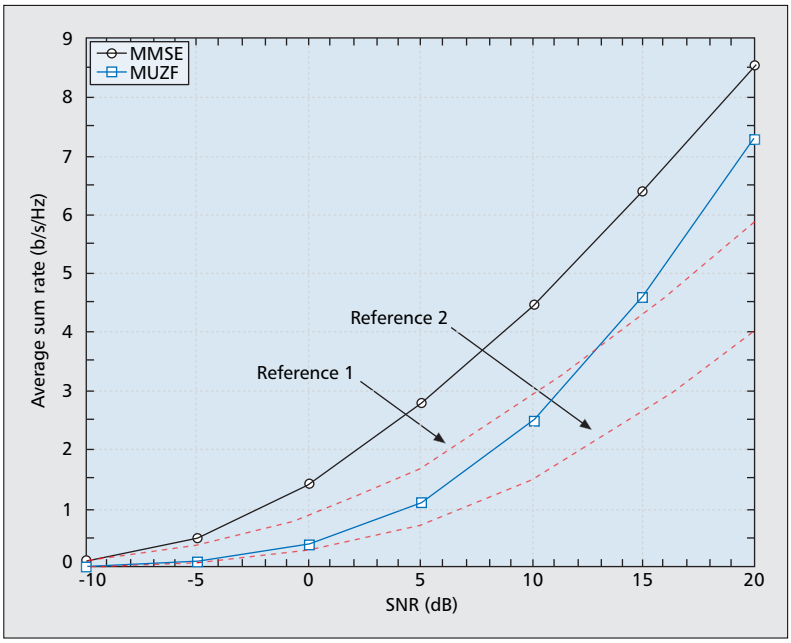

■ **Figure 2.** *Average sum rate versus SNR for 4 single-antenna source/destination pairs and 13 relays.*

Simulation Results — We performed Monte Carlo simulations in Matlab, where all channel estimates are perfect and there is no phase noise. Two cases for the reference scenario are considered:

- All channel coefficients exhibit unit variance (denoted reference 1).
- The variance of the channel coefficients in the reference scenario is 1/4, while it is one for the first-hop and second-hop channel coefficients in the relaying case (reference 2). In an environment where the path loss is similar to free-space propagation, this corresponds to a situation where the relays are halfway between sources and destinations.

Figure 2 shows the average sum rate for MUZF and MMSE relaying as well as the two reference cases. Obviously, the relaying schemes are able to offer performance gains over the reference cases. MUZF and MMSE relaying will converge for high SNR and offer a distributed spatial multiplexing gain in the order of the number of source/destination pairs.

Measurement Results — A custom-built radio testbed called Radio Access with Cooperating Nodes (RACooN), which can be used to demonstrate multiuser relaying scenarios, is available at the Communications Technology Laboratory at ETH. Ten identical single-antenna nodes can each act as either source, relay, or destination node in a wireless network.

We set up a network comprising two source/destination pairs and three relays operating at a carrier frequency of $f_c = 5.5$ GHz. This is the minimum relay configuration for MUZF. The propagation environment is a typical laboratory room of about 8 m × 4 m containing cupboards

and tables with electronic equipment. In Fig. 3 we plot the average signal-to-interference ratio (SIR) at destination 1 for MUZF relaying and noncoherent forwarding (mere amplitude scaling at the relays). Compared to noncoherent forwarding, the ZF gain allocation is able to suppress the interference on average by about 20 dB in this system configuration. Correlation of the source-relay-destination channels has a large impact on the performance. The less correlated the channels, the better the gain allocation is able to suppress interference. The results prove that coherent multiuser relaying indeed works in the real world even if there is no LO phase reference at the relays.

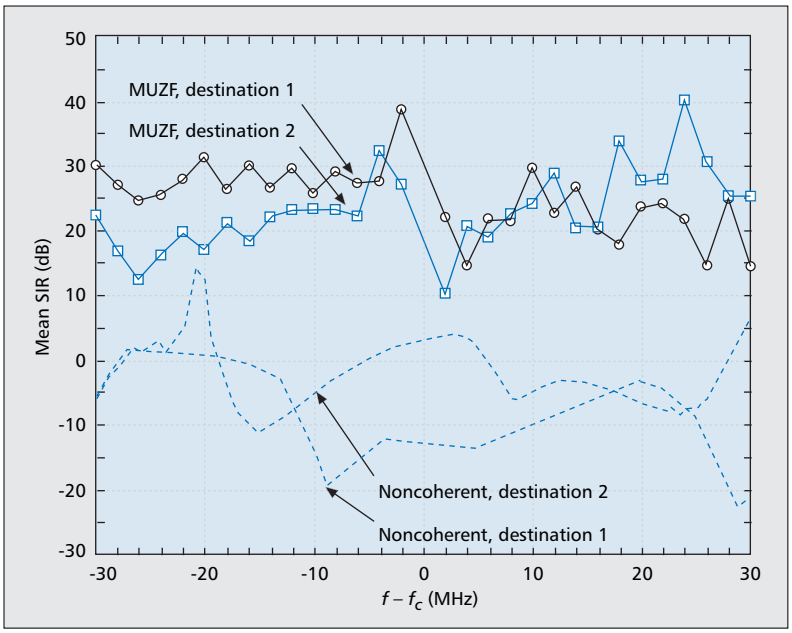

■ **Figure 3.** *Measured mean SIR for MUZF relaying and noncoherent forwarding.*

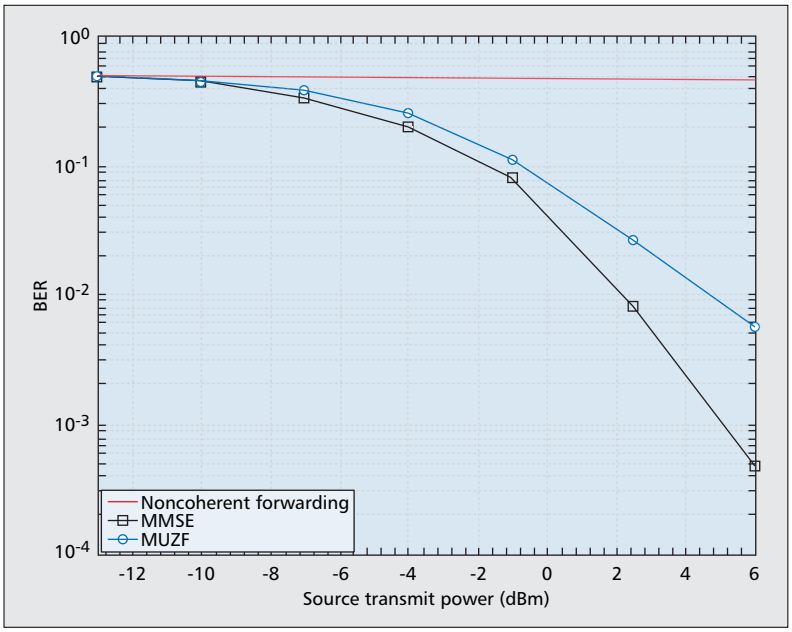

■ **Figure 4.** *Bit error rate at destination 1 versus the transmit power of each source.*

Finally, we use OFDM to obtain frequency-flat subchannels and compute the gain factors for each of them. Bit error curves are then generated for a single subcarrier. In Fig. 4 we show the results at destination 1 for MUZF relaying, MMSE relaying, and noncoherent forwarding as reference. We see that with the coherent gain allocation schemes, communication between both source/destination pairs indeed becomes possible. Note that these experiments not only validate the feasibility of scenario A, but also of scenarios B and C because increasing the amount of cooperation among the nodes actually reduces the challenges.

Scenario B

Compared to scenario A, in scenario B it is much easier to provide the relay with global CSI because the antennas are collocated and the dissemination phase is thus not required. Furthermore, since a single LO reference is used for all relay antennas, phase noise does not destroy coherence. Finally, each relay transmit antenna may forward a linear combination of the received signals of all receive antennas. In the following we discuss how these simplifications can be exploited in one-way and two-way relaying.

In one-way relaying N sources share the physical channel, which means that $N - 1$ interferers are received at each destination. The gain allocation at the relay is required to suppress this interference because it cannot be canceled by the destinations themselves. However, there are additional degrees of freedom available for the gain allocation compared to scenario A because the relay antennas can cooperate. The number of required antennas is thus reduced significantly. In general, $M = N$ relay antennas are sufficient in order to separate the data streams of all N sources. The separation can be achieved by multiuser beamforming approaches at the relay (e.g., according to the ZF or MMSE criterion). This beamforming is a combination of receive beamforming in the first-hop multiple access channel and transmit beamforming in the second-hop broadcast channel.

In two-way relaying, $2N$ sources access the channel simultaneously. Since each destination can eliminate the self-interference coming from its own transmitted data stream, $2(N - 1)$ effective interferers are present. This essentially means that the gain allocation at the relay has to suppress about double the amount of interfering data streams as in one-way relaying. If the relay has at least $2N$ antennas, it can also suppress the self-interference at all destinations, allowing to sources and destinations of low complexity to be implemented. In this case multiuser beamforming approaches similar to the one-way relaying case can be applied [6].

Scenario C

With respect to channel estimation/dissemination and the impact of relay phase noise, scenario C behaves as scenario B. Since all source and destination antennas are also collocated, joint processing of their transmit and receive signals now becomes possible.

For one-way relaying, N data streams can be transmitted simultaneously if $M \geq N$. Since the destination node can apply joint decoding, there is no interference in scenario C. The relay antennas can then be used exclusively to maximize the desired figure of merit (e.g., sum rate). Note that for $M < N$ the rank of the equivalent two-hop channel collapses. In [10] the authors show how to maximize the instantaneous sum rate for one-way relaying in scenario C by beamforming at the relay. The spatial substreams can be separated by beamforming matched to the eigenmodes of the first-hop and second-hop channels. By appropriate pairing of the first- and second-hop subchannels, and using water filling to weight the subchannels, the instantaneous sum rate can be maximized.

In two-way relaying $2N$ data streams can be transmitted simultaneously if $M \geq 2N$. Since the self-interference can be completely suppressed at the destinations, the beamforming at the relay can be used exclusively to maximize the instantaneous sum rate as in one-way relaying. The relay cannot adapt simultaneously to the eigenmodes of both first-hop channels and both second-hop channels. However, it can adapt to the eigenmodes of the joint first-hop channel and joint second-hop channel, where the joint channels are concatenations of the two matrix channels between the nodes and the relay. With this adaption, the number of optimization variables for the beamforming at the relay is always given by $(2N)^2$ even for the case $M > 2N$. Up to now, the beamformers at the relay are determined by numerical methods; analytical solutions achieving the maximum sum rate are subject to future research.

In this article a comparison of the sum rates of one-way and two-way relaying in scenarios B and C is presented. Figure 5 shows a comparison of the sum rates in one-way and two-way relaying for $M = 4$ and $N = 2$. The SNR is given by the ratio between the transmit power of each transmitter and the noise variance at the respective receiver. It is assumed that the SNR for the transmission from each source to the relay is the same as the SNR from the relay to each destination. Obviously, the two-way relaying scheme significantly outperforms the one-way relaying scheme due to the more efficient utilization of the available channel resources.

As a reference, the average sum rate of a point-to-point one-hop MIMO system with $N = 2$ antennas is depicted for $\mathrm{SNR}^{(0)} = 8$ dB and $\mathrm{SNR}^{(0)} = 16$ dB, where $\mathrm{SNR}^{(0)}$ is the SNR of the point-to-point link between source and destination. The plot shows that two-way relaying outperforms the reference case if the relay is placed such that the source/relay and relay/destination links exhibit an SNR of at least 12 dB and 19 dB, respectively.

Figure 6 finally shows the average sum rate vs. M for SNR = 12 dB and $N = 2$. Obviously, for $M = 4$ the spatial multiplexing gain is already fully exploited in two-way relaying. By increasing M, the array gain as well as the spatial diversity increase. The relative increase in average sum rate is the same for one-way and two-way relaying. For large M, the average sum rate of two-way relaying in scenario B converges to the sum rate of two-way relaying in scenario C because the interference in scenario B can be suppressed more efficiently by the relay if M increases.

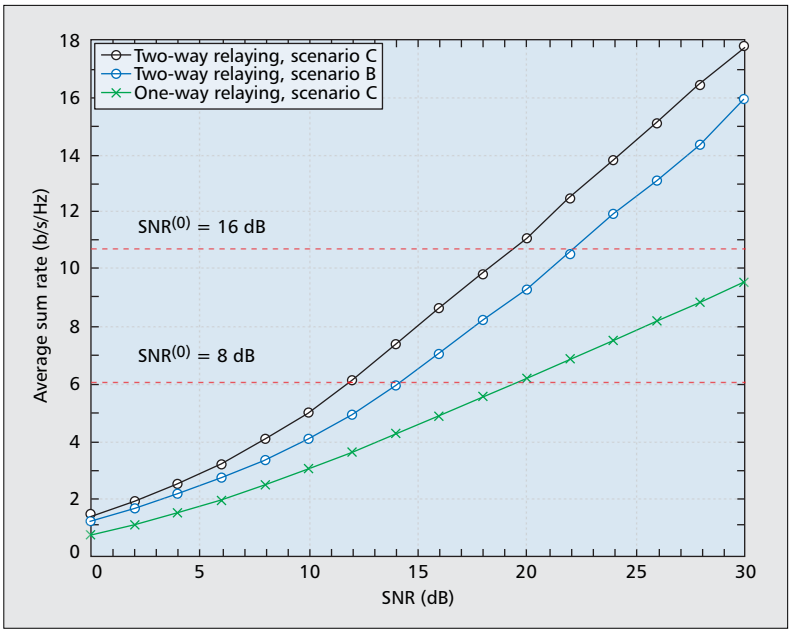

Figure 5. *Average sum rate for one-way and two-way relaying,* M = *4 antennas at the relay,* N = *2 antennas at S1 and S2.*

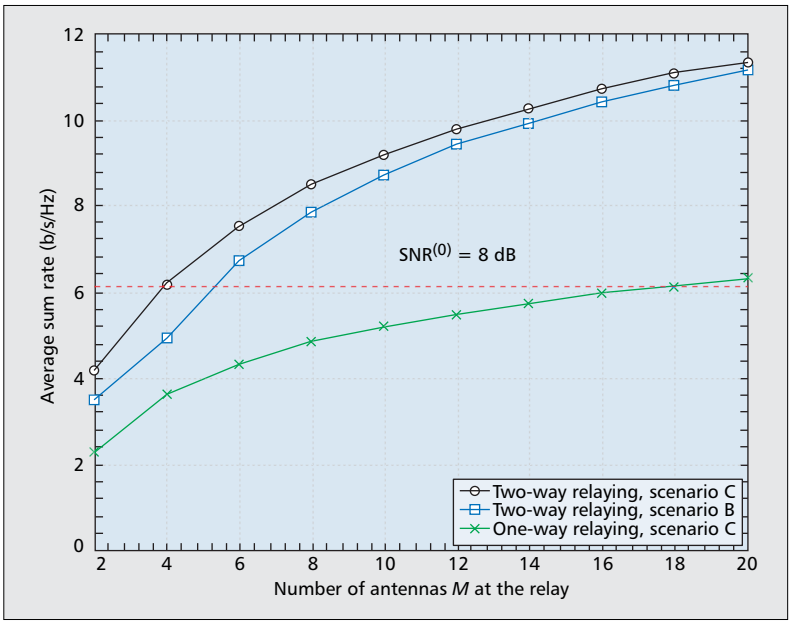

Figure 6. *Average sum rate for one-way and two-way relaying, SNR = 12 dB,* N = *2 antennas at S1 and S2 average sum rate (b/s/Hz).*

Open Issues and Future Research

A lot of research is currently being done on the topic of wireless multiuser communications. AF relaying is a promising candidate technology for future wireless cooperative systems. Although some interesting and very promising results have been obtained, there are still many issues that have not received much attention yet. These are,

Efficient protocols for the dissemination of system coordination data or CSI using for example a secondary communication system like Bluetooth or UWB will be key issues for future cooperative communication systems.

for example, two-way relaying schemes in the presence of multiple distributed relays or gain allocation schemes for broadband systems, the assignment of sources, relays, and destinations among a set of nodes, efficient strategies to cope with mobile nodes entering or leaving the system, or robust signal processing schemes taking system imperfections like noisy or outdated CSI, timing errors, phase noise, or hardware imperfections into account. Furthermore, efficient protocols for the dissemination of system coordination data or CSI, using, for example, a secondary communication system like Bluetooth or UWB, will be key issues for future cooperative communication systems. And finally, fundamental performance bounds of distributed communication protocols remain to be found.

References

[1] J. N. Laneman, D. N. C. Tse, and G. W. Wornell, "Cooperative Diversity in Wireless Networks: Efficient Protocols and Outage Behaviour," *IEEE Trans. Info. Theory*, vol. 50, no. 12, Dec. 2004, pp. 3062–80.

[2] B. Rankov and A. Wittneben, "Spectral Efficient Signaling for Half-Duplex Relay Channels," *Asilomar Conf. Signals, Systems, and Computers 2005*, Pacific Grove, CA, Nov. 2005.

[3] E. C. van der Meulen, "Three-Terminal Communication Channels," *Adv. Appl. Prob.*, vol. 3, 1971, pp. 120–54.

[4] H. Bölcskei *et al.*, "Capacity Scaling Laws in MIMO Relay Networks," *IEEE Trans. Wireless Commun.*, vol. 5, no. 6, June 2006.

[5] I. Hammerström *et al.*, "MIMO Two-Way Relaying with Transmit CSI at the Relay," *Proc. IEEE Sig. Proc. Advances in Wireless Commun.*, Helsinki, Finland, June 2007.

[6] T. Unger and A. Klein, "On the Performance of Two-Way Relaying with Multiple Antenna Relay Stations," *Proc. IST Mobile and Wireless Commun. Summit*, Budapest, Hungary, July 2007.

[7] A. Lozano and A. M. Tulino, "Capacity of Multiple-Transmit Multiple-Receive Antenna Architectures," *IEEE Trans. Info. Theory*, vol. 48, no. 12, Dec. 2002, pp. 3117–28.

[8] S. Berger and A. Wittneben, "Cooperative Distributed Multiuser MMSE Relaying in Wireless Ad-Hoc Networks," *Asilomar Conf. Signals, Sys., and Computers*, Pacific Grove, CA, Oct. 2005.

[9] —, "A Coherent Amplify-and-Forward Relaying Demonstrator Without Global Phase Reference," *IEEE Int'l. Symp. Pers., Indoor and Mobile Radio Commun.*, Cannes, France, Sept. 2008.

[10] O. Munoz, J. Vidal, and A. Agustin, "Linear Transceiver Design in Nonregenerative Relays With Channel State Information," *IEEE Trans. Sig. Proc.*, vol. 55, no. 6, June 2007, pp. 2593–2604.

Biographies

Stefan Berger [M] (berger@nari.ee.ethz.ch) received his Dipl.-Ing. degree in electrical engineering from Munich University of Technology, Germany, in 2003. Since 2003 he has been working toward his Ph.D. degree in electrical engineering at the Communication Technology Laboratory (CTL), ETH Zurich, Switzerland. His research interests are in the field of distributed wireless ad hoc and sensor networks with special focus on cooperative wireless relaying.

Marc Kuhn [M] (kuhn@nari.ee.ethz.ch) received his Diploma degree in electrical engineering and Dr.-Ing. degree frm Saarland University, Saarbrücken, Germany, in 1998 and 2002, respectively. Since 2002 he has been with CTL at ETH Zurich. His main research interests are in wireless communications, including space-time coding, MIMO communication, cooperative wireless communications, wireless relaying, and heterogeneous wireless access networks.

Armin Wittneben [M] (wittneben@nari.ee.ethz.ch) received his M.S. and Ph.D. degrees in electrical engineering from the Technical University of Darmstadt, Germany, in 1983 and 1989, respectively. In 1989 he joined the central research unit, ASCOM Tech, of ASCOM Lt., Switzerland, as a postdoctoral researcher. In 1992 he became head of the Wireless Technology and Security Department. In 1996 he also served as co-director of ASCOM Tech. In 1997 he received the Venia Legendi (Habilitation) in communication technology from the Technical University of Darmstadt. in 1998 he accepted a position as full professor of telecommunication at Saarland University. Since 2002 he has been a full professor of wireless technology at ETH Zurich and is serving as head of CTL. His research interests include all aspects of wireless communication with special emphasis on communication theory, signal processing, cooperative wireless signaling and protocols, distributed MIMO systems, and ultra-wideband/ultra-low-power communication.

Timo Unger [S] (t.unger@nt.tu-darmstadt.de) received his Dipl.-Ing. degree in electrical engineering from the Technical University of Darmstadt in 2005. He is currently working toward his Ph.D. degree in elecrical engineering in the Communications Engineering Laboratory of the same university. His research interests include future mobile radio systems with focus on two-hop relaying and multiple antenna communications.

Anja Klein [M] (a.klein@nt.tu-darmstadt.de) received her diploma and Dr.-Ing. degrees in electrical engineering from the University of Kaiserslautern, Germany, in 1991 and 1996, respectively. From 1991 to 1996 she was a member of staff with the Research Group for RF Communications at the university. In 1996 she joined Siemens AG, Mobile Networks Division, Munich and Berlin. She was active in the standardization of third-generation mobile radio in ETSI and 3GPP, for instance, leading the TDD group in RAN1 of 3GPP. She was a vice president, headinig a development department and a systems engineering department. In May 200 she joined Darmstadt University of Technology as a full professor, heading he Communications Engineering Laboratory. Her main research interests are in mobile radio, including multiple access, single- and multicarrier transmission schemes, multi-antenna systems, radio resource management, network planning and dimensioning, relaying, and cooperative communication. She has published over 150 refereed papers and has contributed to five books. She is an inventor and a co-inventor of more than 45 patents in he field of mobile radio. In 1999 she was inventor of the year at Siemens AG. She is a member of Verband Deutscher Electrotechniker-Inforationstechnische Gesellschaft (VDE-ITG).

ACCEPTED FROM OPEN CALL

Tuning Radio Resource in an Overlay Cognitive Radio Network for TCP: Greed Isn't Good

Teerawat Issariyakul, Laxminarayana S. Pillutla, and Vikram Krishnamurthy,
The University of British Columbia

ABSTRACT

In this article we illustrate the performance of Transmission Control Protocol in an overlay cognitive radio network under dynamic spectrum access. We show that the performance of TCP in overlay CR networks that implement DSA to be quite different from its performance in conventional networks, which do not allow DSA. The key difference is that secondary users in an overlay CR network have to cope with a new type of loss called service interruption loss, due to the existence of primary users. We demonstrate on an NS2 simulation testbed the surprising result: Excessive radio resource usage leads to a decrease in aggregate TCP throughput. This behavior is in contrast to the behavior of TCP in conventional networks, where throughput increases monotonically with the available radio resource.

INTRODUCTION

Cognitive radio (CR) is an emerging communications paradigm, wherein a wireless transceiver unit can sense the surrounding environment and adapt itself accordingly. CR technology, along with dynamic spectrum access (DSA), has the potential to alleviate the shortage of radio resource. For many years, it was believed that the spectrum shortage was due to an increasing number of wireless applications and their substantial bandwidth usage. However, spectrum measurements in Washington, DC, New Orleans, San Diego, Atlanta, Chicago, and other metropolitan areas show that vast portions of licensed spectrum are not in use (i.e., they are underutilized) [1]. This finding suggests that an improvement in spectrum utilization can alleviate the spectrum shortage problem. A transceiver equipped with CR capability can access and release spectrum with more agility. Such agility can allow secondary users (who are unlicensed) to access the spectrum when primary users (i.e., spectrum licensees) are inactive. Referred to in the literature as DSA, this approach clearly reduces spectrum underutilization, hence alleviating the spectrum shortage problem.

Most of the CR research focused on the lower layers (i.e., physical and medium access control [MAC] layers) of the network protocol stack. Despite sporadic works that study routing issues in CR networks, no work to date has considered transport layer issues in CR networks [2]. In this article we study the performance of Transmission Control Protocol (TCP) in a CR network.

TCP has become a ubiquitous transport layer protocol for reliable non-real-time data transmission over the Internet. TCP was originally designed to control network congestion in wired networks, and was later enhanced to cope with channel errors in wireless networks [3]. The advent of DSA introduces a new type of packet loss called *service interruption loss* (see the definition given later), in addition to the existing packet losses resulting from network congestion and channel errors. Consequently, there is a strong motivation to study the performance of TCP under service interruption loss in CR networks that employ DSA.

This article provides a detailed survey of CR networks. We present their salient features, access techniques, and architectures. We also discuss a surprising result regarding TCP performance in overlay CR networks (see the definition given in the next section): the TCP throughput of a secondary user does not always increase with the available radio resource. Due to service interruption from the primary user, TCP performance of a secondary user could degrade when it tries to acquire too much radio resource. This finding poses the need to optimally control how secondary users obtain the radio resource in order to fully utilize the available radio resource.

The rest of this article is organized as follows. The next section discusses CR evolution and provides a survey of various regulatory issues, access techniques, and architectures proposed in the literature. We then describe the basic operation of TCP in conventional networks, and emphasize the need to study TCP performance over a CR network. The following section shows our simulation study of TCP performance in CR networks allowing DSA. Conclusions are given in the final section.

This work was supported in part by the Canadian Natural Sciences and Engineering Research Council and by the Asia-Pacific Telecommunity.

 0163-6804/09/$25.00 © 2009 IEEE

A general centralized network architecture consists of two main entities. One is a base station, which schedules the data transmission of users in the network. The other entity is the spectrum broker, which is responsible for allocating the radio resource to users.

CR: Evolution, Access Techniques, and Architectures

Evolution and Salient Features of CR

CR is a key enabling technology to achieve DSA of licensed spectrum bands. The current interest in CR technologies is due to a report published by the Federal Communications Commission (FCC) pointing out that vast portions of spectrum are underutilized [1]. This finding suggested that granting access to unlicensed users could lead to significant improvement in spectrum utilization.

The term CR was coined by Mitola in [4], although its genesis can be traced to the concept of software defined radio (SDR). According to the FCC, "an SDR is one in which operating parameters such as frequency range, modulation type, or output power can be altered by software without making any changes to the hardware components." Taking the definition of SDR a step further, Mitola defined CR as "an SDR that senses its environment, tracks changes and reacts accordingly" [4].

CR is characterized by two main features:

- *Cognitive capability*: This refers to the ability of the radio to identify spectrum opportunities and adapt to the situation accordingly. The typical steps involved in this adaptive operation are *spectrum sensing*, *spectrum analysis*, and *spectrum decision*. Spectrum sensing involves detection of white spaces (i.e., the spectrum portion not in use). Spectrum analysis involves analyzing the characteristics of the detected white space (e.g., characterizing the white space based on its time-varying radio characteristics and the primary user activity). The spectrum decision specifies the action (e.g., adjusting transmission rate, power levels, and/or bandwidth) that should be taken after spectrum sensing and analysis [2, 5].
- *Reconfigurability*: This refers to the ability of the radio to be configured dynamically according to the environment. The various parameters that can be reconfigured are, for example, operating frequency, modulation type, and transmission power [5].

A general CR network consists of two subnetworks: a *primary network* and a *secondary network*. The primary network consists of primary users and the primary base station, whose function is to control and regulate the users of the network, similar to a cellular system. The secondary network does not own a license and tries to utilize the spectrum in an opportunistic manner. A secondary network may or may not have a base station, depending on the architecture.

We note that CRs can be conceived for even the unlicensed bands of the spectrum, as noted in [2]. However, in this article we focus our attention on CR networks for access to the licensed bands only.

Spectrum Access Policy: Own It or Share It?

The inception of CR has raised spectrum regulatory issues. Two schools of thought — *property rights* and *spectrum commons* — are being debated extensively [6]. Advocates of property rights argue that spectrum owners should have an absolute right over their spectrum. Proponents of spectrum commons, on the other hand, believe that sharing spectrum would be more efficient and greatly alleviate the problem of spectrum underutilization. Hence, there are constant regulatory and legal policy discussions on these two conflicting ideas, the result of which is still inconclusive.

Access Techniques for CR Networks

Access techniques for CR networks can be classified into two types:

- *Overlay approach (or interference-free approach)*: In this approach the secondary users access the portion of the spectrum that is not used by primary users. As a result, there is virtually no interference to the primary users. We refer to CR networks that employ overlay access techniques as *overlay CR networks*.
- *Underlay approach (or interference-tolerant approach)*: In this approach the secondary users access the network by spreading their signals over a wide frequency band. The underlay approach imposes severe constraints on the transmission power of secondary users. Operating below the noise floor of primary users, secondary users are allowed to interfere with primary users up to a certain tolerable level. We refer toCR networks that employ underlay access techniques as *underlay CR networks*.

The current spectrum management policy adopted by the FCC, based on the property rights model, does not allow secondary user operation in the licensed spectrum, irrespective of the access techniques. Hence, there may be regulatory and legal discussions on how secondary users should operate in the licensed spectrum.

Architectures for CR Networks

The architecture for CR networks can be either *centralized* or *distributed*. Table 1 summarizes various CR architectures proposed in the literature.

Centralized CR Network Architectures — A general centralized network architecture consists of two main entities. One is a base station, which schedules the data transmission of users in the network. The other entity is the spectrum broker, which is responsible for allocating the radio resource to users (primary, secondary, or both). The spectrum broker can be a primary or secondary base station [2], or a dedicated entity dealing with spectrum allocation [7, 9]. The sensing functionality in a centralized CR network can be performed by secondary users or the spectrum broker. The sensed spectrum information is used by the spectrum broker to create a spectrum allocation map for radio resource allocation.

The following are examples of centralized CR architectures proposed in the literature. A more comprehensive survey of CR architectures is given in [2].

	Centralized	Distributed
Sensing entity	Secondary users/spectrum broker	Secondary users
Spectrum broker	Separate entity/primary base station/secondary base station	None
Access granted to	Primary and secondary users	Secondary users
Examples	Spectrum pooling [7] DIMSUMnet [9]	CORVUS [8] Nautilus [10, 11]

Table 1. *Summary of cognitive radio architectures.*

In general, the suggested solutions to improve performance of TCP in wireless networks are classified into three basic groups based on their fundamental philosophies: end-to-end solutions, split-connection solutions, and link-layer solutions.

Spectrum Pooling — In this architecture spectrum from different owners is put together in a common pool [7]. The architecture is based on the well-known orthogonal frequency-division multiplexing (OFDM). The advantage of an architecture based on OFDM is that interference to primary users can be minimized by allocating zero power to the subcarriers occupied by them.

Dynamic Intelligent Management of Spectrum for a Ubiquitous Mobile Network — In dynamic intelligent management of spectrum for a ubiquitous mobile network (DIMSUMnet) [9], a spectrum broker owns a contiguous chunk of spectrum called the coordinated access band (CAB).[1] The CAB is then leased according to requests from different users for a specified time.

Distributed CR Network Architectures — This type of architecture has no centralized agent like a spectrum broker or base station to coordinate spectrum access of secondary users. It is similar to that used in wireless ad hoc networks. The main difference from wireless ad hoc networks is the presence of primary users and service interruption loss, which will be defined in the next section.

Distributed CR networks can be classified into *cooperative* or *non-cooperative* networks. In a cooperative network the users share the interference information and determine spectrum allocation based on this shared information. In a non-cooperative network the users access the spectrum based on local policies, since there is no communication regarding interference information among users of the network. In general, the cooperative approach achieves better throughput than the non-cooperative approach. However, a cooperative approach may incur extra overhead due to the communication among users.

We now briefly summarize the existing distributed CR network architectures.

Cognitive Radio Approach for Usage of Virtual Unlicensed Spectrum — In the CR approach for usage of virtual unlicensed spectrum (CORVUS) [8], secondary users form a communication group called a secondary user group (SUG). The secondary users in an SUG help each other in sensing the channel for primary user activity. If primary users are inactive on a channel, the SUG takes that channel for communication within the group.

We note that CORVUS can be used even in the centralized mode. In each SUG a secondary user could be selected to control spectrum access of other secondary users.

Nautilus — This project devises distributed, efficient, and scalable algorithms for spectrum sharing in CR ad hoc networks [10, 11]. The project proposes spectrum access algorithms based on two approaches: one based on graph coloring [10] and another based on local bargaining [11]. In the former approach the spectrum allocation is transformed into a graph-coloring problem, where a secondary user corresponds to a vertex, while a radio channel corresponds to a color. Since a graph-coloring problem is NP-hard, the authors propose heuristic approximation algorithms to solve the problem [10]. In the latter approach the secondary users affected by topological changes form local bargaining groups and adapt their spectrum assignment to a new optimal conflict free assignment [11].

TCP in Conventional vs. CR Networks

TCP in Wired and Conventional Wireless Networks

TCP was originally conceived for wired networks as a means to avoid and control network congestion, and to provide reliable end-to-end delivery of user data. Whenever the sender identifies loss of data (either in the form of several duplicated acknowledgments or an absence of acknowledgment for a timeout interval), TCP reacts to the loss by reducing its transmission window size before retransmitting lost packets. This reduction in window size eases the load on intermediate links, hence controlling congestion in the network. The subsequent increase in window size depends on whether TCP is in *slow start phase* or *congestion avoidance phase*. In the slow start phase the window size grows linearly in response to every acknowledged packet, while the window size grows sublinearly in the congestion avoidance phase. In general, the essential functionality of TCP remains the same, despite the modifications to improve its performance. For example, TCP variants (TCP-Tahoe, TCP-Reno, and TCP-Vegas) implement the same window-based adaptation mechanism, but have minor differences.

[1] *CABs can be designated frequency bands in cellular, PCS, or TV broadcasting systems.*

TCP performance under the congestion losses, channel error losses, and/or data collision losses has been studied extensively. However, due to the service interruption loss, TCP performance in an overlay CR network could be significantly different from that in a conventional network.

TCP was originally developed for wired networks and later enhanced for wireless networks. Unlike in a wired network where packet loss is mainly caused by congestion, the majority of packet loss in a wireless network is due to channel errors (caused by fading, interference, and shadowing) or packet collision (i.e., when more than one mobile accesses the channel simultaneously). While two consecutive packet losses due to network congestion are highly correlated, those due to channel errors or collision may not be related at all.

In general, the suggested solutions to improve performance of TCP in wireless networks are classified into three basic groups based on their fundamental philosophies: *end-to-end solutions*, *split-connection solutions*, and *link-layer solutions* [3]. Among these three solution approaches, link-layer solutions are the most popular, since they do not require any modification of the network architecture or TCP operation.

Service Interruption Loss in Overlay CR Networks

The concept of DSA introduces a new type of loss called service interruption loss for the secondary users of an overlay CR network. *Service interruption loss refers to the loss experienced by secondary users due to the intervention of primary users while transmitting data*. Service interruption loss is absent in underlay CR networks, since in an underlay approach the secondary users operate far below the power level of the primary users, hence ruling out the possibility of collisions with the primary users of the spectrum.

The amount of service interruption loss for the secondary users of the overlay CR network is determined by the amount of primary user activity in the network. To capture the amount of primary user activity, we define load factor (A) as the amount of traffic *all* the primary users generate on a *single* channel. Mathematically, it is stated as

$$A = \frac{T_{ON}}{T_{ON} + T_{OFF}} \times \frac{N_p}{L}, \qquad (1)$$

where T_{ON} and T_{OFF} denote the mean ON time and OFF time of the primary users, N_p denotes the number of primary users, and L denotes the total number of channels in the spectrum.

The rationale behind Eq. 1 is as follows. First, the average traffic generated by each primary user is $T_{ON}/(T_{ON} + T_{OFF})$. Hence, the total traffic generated by all the primary users is $N_p \cdot T_{ON}/(T_{ON} + T_{OFF})$. Since this traffic is generated over L radio channels, the total traffic (or load) generated by the primary users on a single channel is given by (A) defined as in Eq. 1. In other words, if the traffic generated by all the primary users is distributed equally among all the radio channels, A in Eq. 1 denotes the amount of primary user traffic on one radio channel. For example, $A = 0.5$ implies that every radio channel is used by all primary users on average 50 percent of the time. In general, a higher value of the load factor (A) implies lower aggregate TCP throughput of the secondary users.

The service interruption loss is different from the following three major losses that occur in a conventional network. First, losses due to network congestion are usually caused by buffer overflow and are assumed to be correlated. Second, losses due to channel errors depend on channel characteristics. Depending on whether the user is moving slowly or rapidly, the losses due to channel errors can be assumed correlated or independent. The third and final type of losses are due to collision between various users of the network that depend on the underlying medium access control (MAC) mechanism. Service interruption loss due to DSA, on the other hand, does not depend on network conditions, channel characteristics, or the underlying MAC mechanism. It depends on the amount of primary user activity in the spectrum, which in turn depends on extraneous factors such as geographical location and time of the day. The report in [1] notes that in a given geographical area, for the most part the spectrum is devoid of primary user activity. In such cases the duration of service interruption loss may be small, since it is possible for the secondary user to find an alternate channel for transmission almost immediately. However, if a given geographical region has heavy primary user activity, it may not be possible for a secondary user to find an alternate channel immediately, hence leading to a longer duration of service interruption loss.

TCP performance under congestion losses, channel error losses, and/or data collision losses has been studied extensively. However, due to service interruption loss, TCP performance in an overlay CR network could be significantly different from that in a conventional network. This poses a clear need to study TCP performance in an overlay CR network. As we shall see from our simulation study presented in the next section, a definite number of channels must be carefully chosen in order to optimize TCP performance in an overlay CR network.

Simulation Study of TCP Performance in an Overlay CR Network

System Model of an Overlay CR Network

In this article we consider an overlay CR network with N_p primary users and N_s secondary users (Fig. 1). Primary users send their data to a primary network via its base station. On the other hand, each secondary user transfers a file with infinite size to the server connecting to its base station. Both primary and secondary users share L radio channels to communicate with their base stations. We assume that the data rate and propagation delay of each radio channel are b_r b/s and d_r s, respectively. Also, the link that connects the base station and server has a data rate of b_l b/s and a propagation delay of d_l s. Each secondary user is allowed to access the free channels only. While using a channel, a secondary user may be asked (by a primary user) to give up its ongoing transmission and return the channel to the primary user. In this case the

packet transmitted by the secondary user will be lost due to service interruption.

We model the primary user activity according to a continuous-time Markov chain with two states, ON and OFF. This modeling assumption implies each primary user generates a data packet that holds the channel for an exponentially distributed time with mean T_{ON}. After transmitting a packet, a primary user turns off for an exponentially distributed time with mean T_{OFF}. When switching to an ON state, a primary user first looks for a free channel to transmit data. If there is no free channel, a primary user will force one of the secondary users (if any) to give up transmission on the radio channel. If all the channels are occupied by primary users, the new primary user will not be able to transmit its data.

Each secondary user transfers a file using TCP as its transport layer protocol. Again, based on a perceived packet loss pattern, TCP feeds and stops feeding packets to the underlying link layer. A secondary user attempts to transmit each link layer packet according to the following spectrum access mechanism. At any time, a secondary user can access a maximum of L_s radio channels, which implies there are always $(L - L_s)$ channels reserved for the primary users. Among the L_s channels, a secondary user chooses to transmit (or) not to transmit on a free channel with probability p and $1 - p$, respectively.[2] If more than one secondary user chooses to transmit on the same channel simultaneously, all the transmitted packets will be lost. The above access mechanism is invoked repeatedly until all the link layer packets are transmitted or the number of occupied channels is greater than L_s.

Simulation Parameters

We run the simulation using an open source network simulation tool, NS2. Unless specified otherwise, we assume $N_p = 30$, $L = 20$, $T_{OFF} = 0.1$ s, $b_r = 100$ kb/s, $d_r = 2$ ms, $b_l = 100$ Mb/s, and $d_l = 2$ ms. The value of T_{ON} is determined by the choice of the load factor value (A). The value of b_r chosen above is typical for wireless data applications. For example, the code-division multiple access (CDMA) 2000 1xEV-DO standard (which is one of the six radio interfaces under IMT-2000) supports a maximum data rate of 154 kb/s [12]. We run the simulation for a duration of 500 s, which we found to be sufficient for TCP to reach its steady state. Each secondary user employs TCP New Reno[3] to upload the file, and chooses transmission probability value (p) equal to $1/N_s$, where N_s denotes the number of secondary users.[4] Finally, we assume there are no losses due to channel errors (or) buffer overflow on all the TCP flows of secondary users.

Simulation Results

Figure 2 shows the plot of aggregate TCP throughput of secondary users vs. the maximum channels for secondary users (L_s) for the case when there are no primary users (i.e., $N_p = 0$). With no primary users, this case corresponds to a conventional random access ad hoc network. As can be seen from Fig. 2, without service interruption from primary users, the aggregate TCP throughput of secondary users increases monotonically as L_s increases. Using random access, the secondary user experiences no collision for $N_s = 1$. Therefore, this case has higher aggregate TCP throughput than the multi-user case (i.e., $N_s = 5$ and $N_s = 10$). Note that the aggregate TCP throughput for $N_s = 1$ saturates at $L_s = 18$ since the TCP window size has reached its maximum value. However, for the multi-user case (i.e., $N_s = 5$ and 10), the aggregate TCP throughput increases continuously since in these

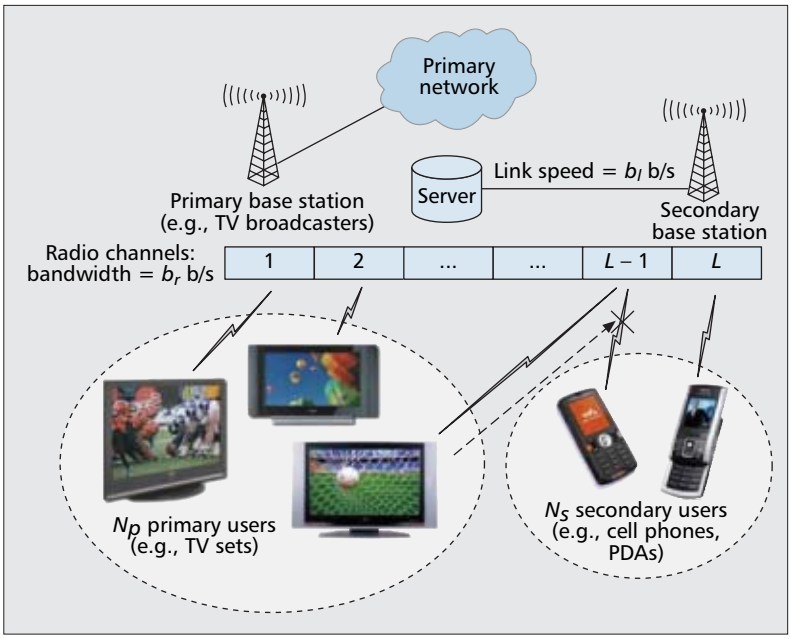

Figure 1. *System model of an overlay CR network: N_p primary users and N_s secondary users access their base stations via L shared radio channels. Each radio channel has a bandwidth of b_r b/s. The primary base station connects to the primary network, while the secondary base station connects to a server via a wired link with a bandwidth of b_l b/s. When needed, a primary user can force a secondary user to give up transmission so that it can use the released channel.*

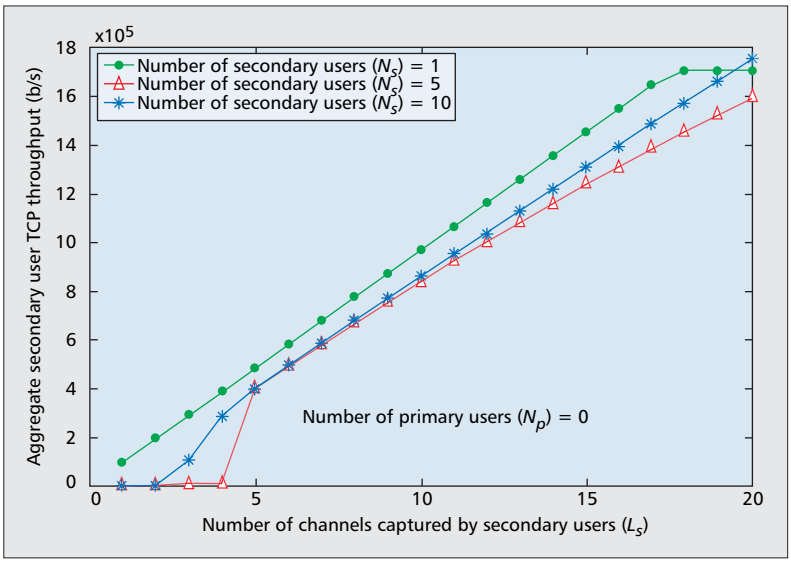

Figure 2. *Aggregate secondary user TCP throughput vs. the maximum channels for secondary users (L_s) when $N_p = 0$. Due to the absence of service interruption from the primary users, this case is similar to an ad hoc network with random access. The figure shows that the aggregate TCP throughput of secondary users increases monotonically with L_s.*

[2] *This type of random access is widely used for ad hoc wireless networks, where there is no base station to control and coordinate the transmission of secondary users.*

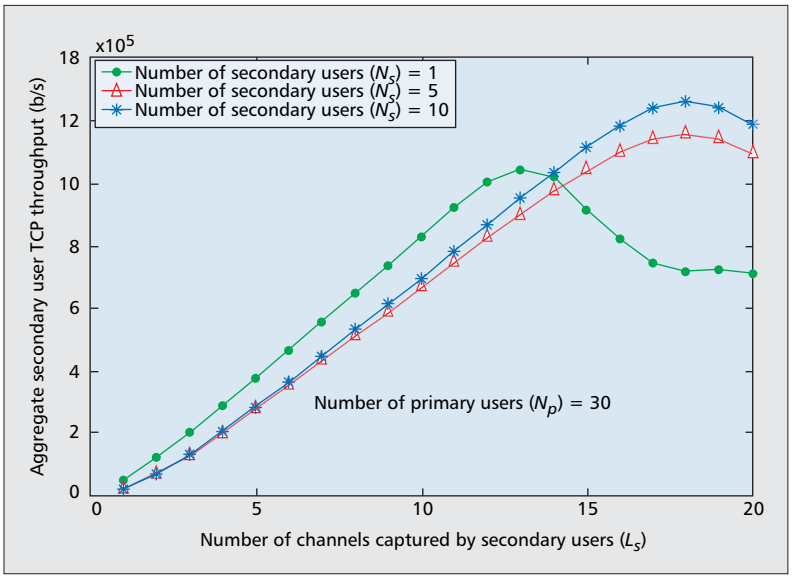

Figure 3. *Aggregate secondary user TCP throughput vs. the maximum channels for secondary users* (L_s) *in the presence of service interruption due to 30 primary users. The figure shows that the throughput is no longer monotone with respect to* L_s. *There exists a well defined optimal value of* L_s *that maximizes the aggregate TCP throughput.*

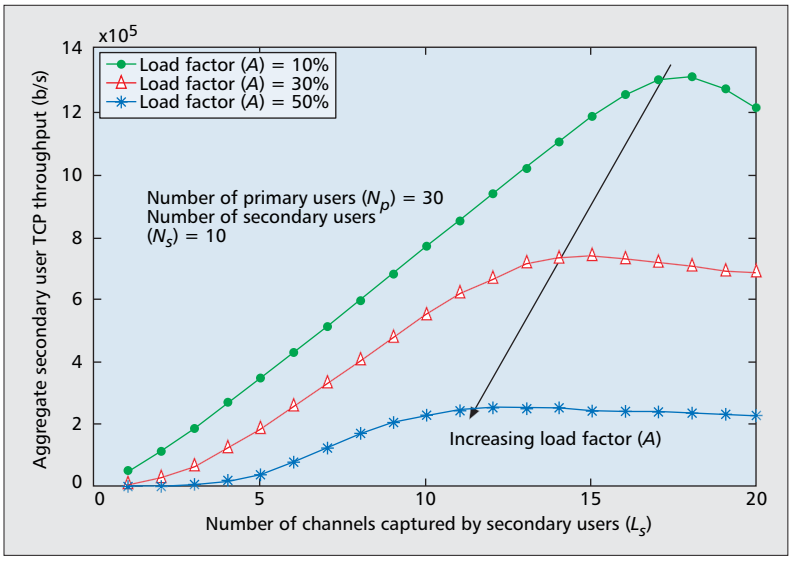

Figure 4. *Aggregate secondary user TCP throughput vs. the maximum channels for secondary users* (L_s) *for various values of load factor* (A) *with service interruption due to 30 primary users. The figure shows that the aggregate TCP throughput of secondary users decreases as the load factor increases.*

[3] TCP New Reno is an improved version of Reno that avoids multiple reductions of congestion window size when several TCP packets from the same window of data are lost [13].

[4] This choice of $p = 1/N_s$ *minimizes the collision probability among* N_s *secondary users.*

cases the TCP window size has not reached its maximum.

Figure 3 shows the impact of service interruption on the aggregate TCP throughput of secondary users. Here, we set the number of primary users (N_p) equal to 30. In this case the throughput gradually increases as the maximum channels for secondary users (L_s) increase up to a certain point beyond which the throughput decreases. Hence, there exists a well defined optimal value of L_s (denoted L_s^*) that maximizes the aggregate TCP throughput.

This non-monotonic behavior of TCP in overlay CR networks that implement DSA can be explained as follows. Intuitively, as the maximum channels for secondary users (L_s) increase, we expect an increase in aggregate TCP throughput. However, an increase in L_s also increases the probability of service interruption from a primary user, leading to increased packet loss. This packet loss (detected via timeout or duplicate acknowledgment) leads to a decrease in TCP window size, and hence significant degradation of aggregate TCP throughput of secondary users. When adopting DSA in overlay CR networks, secondary users need to be judicious in choosing the number of channels for data transmission.

We also observe from Fig. 3 that the optimal value of L_s (L_s^*) also increases as the number of secondary users increases. Since every secondary user accesses the channel with the same probability $p = 1/N_s$, on average, all the free channels are distributed equally among the secondary users. For example, suppose the number of free channels are 15. In a single-user case all 15 free channels are used by only one user. When N_s = 5, each secondary user employs, on average, 15/5 = 3 channels to transmit TCP packets. With a small number of channels per user, a success in single transmission could gain a significant increase in TCP throughput compared to the packet loss due to service interruption. Therefore, the peaks (i.e., L_s^*) in Fig. 3 shift to the right as the number of secondary users increases.

Figure 4 shows the variation of aggregate TCP throughput with respect to the maximum number of channels L_s for different values of the load factor (A). In this case we set N_p = 30 and N_s = 10, and the load factor A is chosen to be 10, 30, and 50 percent, respectively. A higher value of the load factor implies an increase in the probability of service interruption from the primary user and also a decrease in transmission opportunities for a secondary user. Hence, the aggregate TCP throughput of the secondary users decreases as the load factor increases.

Figure 5 plots the optimal value of L_s (L_s^*) vs. the load factor (A). We set the number of primary and secondary users to be 30 and 1, respectively. The non-monotone behavior of L_s^* with respect to the load factor A in Fig. 5 can be explained as follows. First, a small load factor implies that primary users rarely need the channel; service interruption is rare in this case. With small service interruption probability, a secondary user should take a chance and transmit on as many channels as possible to increase TCP throughput. Second, as the load factor increases, the service interruption probability increases. In this case we should decrease L_s^* to avoid service interruption. The decrease of L_s^* (with respect to A) stops when the traffic load is extremely large. In this case primary users tend to occupy all the channels for most of the time, and secondary users may not be able to acquire a channel. Therefore, secondary users should take a chance and transmit packets on as many channels as possible, even though they are prone to increased service interruption from the primary users. Due to the above non-monotone behavior, a secondary user needs to be judicious in setting the maximum number of channels for the secondary users viz. L_s.

Figure 5 also shows the effect of radio channel bandwidth on L_s^*. An increase in radio channel bandwidth decreases packet transmission time. In this case packet transmission could be complete before a primary user service interruption. Increasing radio channel bandwidth therefore decreases the service interruption probability. In a high bandwidth case we can allow a secondary user to acquire more radio channels without increasing the risk of service interruption. Hence, we observe in Fig. 5 that increasing the radio channel bandwidth can lead to a higher optimal value of L_s (i.e., L_s^*).

Conclusions

In this article we review the evolution of CR and the various architectures envisaged for its realization. The DSA technique introduces a new type of loss called service interruption loss for the secondary users in an overlay CR network. We show via simulations that *TCP performance under service interruption loss in an overlay CR network is significantly different to its performance in a conventional network*. In particular, we observe that in an overlay CR network that adopts DSA, there is an optimal number of channels the secondary users need to capture to maximize their aggregate TCP throughput. After all, a secondary user cannot blindly utilize all the available channels and expect the best result.

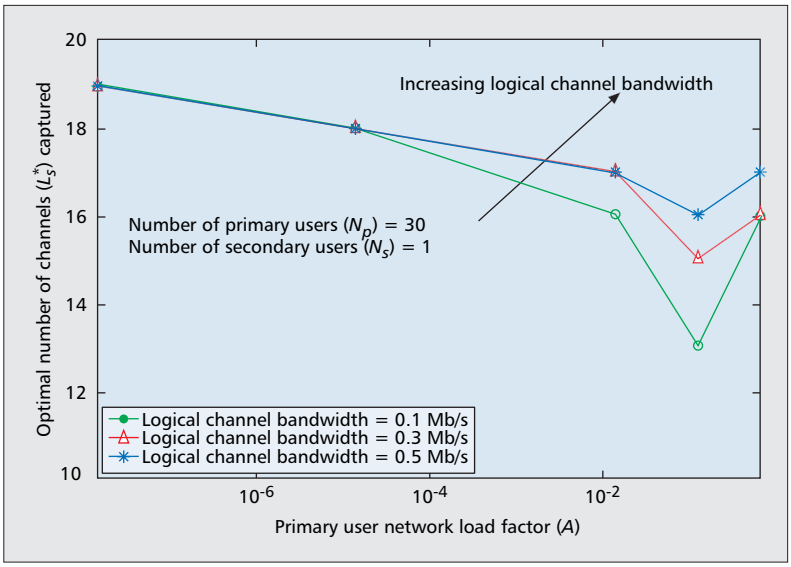

Figure 5. *Optimum number of channels (L_s^*) a secondary user captures to maximize its TCP throughput vs. load factor (A) for different values of bandwidth. The figure shows that for a given bandwidth, the optimal number of channels (L_s^*) captured is high for small and large load factor values, and low for moderate load factor values.*

References

[1] FCC, "Spectrum Policy Task Force Report," ET Docket no. 02-135, Nov. 15, 2002.

[2] F. Akyildiz *et al.*, "NeXt Generation/Dynamic Spectrum Access/Cognitive Radio Wireless Networks: A Survey," *Elsevier Comp. Net.*, vol. 15, no. 13, Sept. 2006, pp. 2127–59.

[3] H. Balakrishnan *et al.*, "A Comparison of Mechanisms for Improving TCP Performance over Wireless Links," *IEEE/ACM Trans. Net.*, vol. 5, no. 6, Dec. 1997, pp. 756–69.

[4] J. Mitola III, *Cognitive Radio: An Integrated Agent Architecture for Software Defined Radio*, Ph.D. thesis, KTH Royal Inst. Technology, Sweden, 2000.

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

[6] J. M. Peha, "Approaches to Spectrum Sharing," *IEEE Commun. Mag.*, vol. 43, no. 2, Feb. 2005, pp. 10–12.

[7] T. Weiss and F. K. Jondral, "Spectrum Pooling: An Innovative Strategy for the Enhancement of Spectrum Efficiency," *IEEE Commun. Mag.*, vol. 42, no. 3, Mar. 2004, pp. 8–14.

[8] R. W. Brodersen *et al.*, "CORVUS: A Cognitive Radio Approach for Usage of Virtual Unlicensed Spectrum," White Paper, Berkeley Wireless Research Center, 2004, pp. 1–21.

[9] M. M. Buddhikot *et al.*, "DIMSUMNet: Spectrum Management in Coordinated Dynamic Spectrum Access Based Cellular Networks," *Proc. IEEE DySPAN '05*, Nov. 2005, pp. 299–307.

[10] H. Zheng and L. Cao, "Device-Centric Spectrum Management, " *Proc. IEEE DySPAN '05*, Nov. 2005, pp. 56–65.

[11] L. Cao and H. Zheng, "Distributed Spectrum Allocation via Local Bargaining," *Proc. IEEE SECON '05*, Sept. 2005, pp. 475–86.

[12] 3GPP2, " CDMA 2000 High Rate Packet Data Air Interface Specification," tech. rep. no. C.S20024 v. 2.0, Oct. 2000.

[13] J. F. Kurose and K. W. Ross, *Computer Networking: A Top-Down Approach*, Addison-Wesley, 2008.

Biographies

TEERAWAT ISSARIYAKUL [M] (teerawat@ece.ubc.ca) received his B.Eng and M. Eng. from Thammasat University and the Asian Institute of Technology (AIT), Thailand, in 1997 and 1999, respectively. In 2005 he completed his Ph.D. at the Department of Electrical and Computer Engineering, University of Manitoba. In 2001–2005 he was also affiliated with Telecommunication Research Laboratory (TRLabs). In 2005-2006 he worked as a postdoctoral fellow at the Department of Electrical and Computer Engineering, University of British Columbia. Currently, he is a system engineer at TOT Public Company Limited, Thailand. He also holds an adjunct faculty position at the Asian Institute of Technology. His current research area is modeling and optimization of cognitive radio networks and intelligent transportation networks.

LAXMINARAYANA S. PILLUTLA [S'04] (laxp@ece.ubc.ca) obtained his B.Eng. from the University of Madras, India, and his M.S. from Wichita State University, Kansas, in electrical engineering. He received a Ph.D. from the University of British Columbia in electrical engineering under the supervision of Dr. Krishnamurthy in 2008. His research interests are in general communication theory (CDMA, OFDM, and MIMO systems), cross-layer design of wireless networks, wireless sensor networks, cognitive radio, and vehicular telematics.

VIKRAM KRISHNAMURTHY [S'90, M'91, SM'99, F'05] (vikramk@ece.ubc.ca) received a B.E. degree from the University of Auckland, New Zealand, in 1988 and a Ph.D. degree from the Australian National University, Canberra, in 1992. Since 2002 he has been a professor and Canada Research Chair of the Department of Electrical Engineering, University of British Columbia, Vancouver, Canada. Prior to 2002 he was a chaired professor with the Department of Electrical and Electronic Engineering, University of Melbourne, Australia. His current research interests include stochastic control and computational game theory, stochastic optimization and scheduling, and statistical signal processing. He has served as Associate Editor for several journals including *IEEE Transactions on Signal Processing*, *IEEE Transactions on Aerospace and Electronic Systems*, *IEEE Transactions on Automatic Control*, *IEEE Transactions on Nanobioscience*, and *Systems and Control Letters*. From 1009 to 2010 he is serving as a Distinguished Lecturer for the IEEE Signal Processing Society.

ACCEPTED FROM OPEN CALL

Handover in Multihop Cellular Networks

Sunghyun Cho, Hansei University
Edward W. Jang and John M. Cioffi, Stanford University

ABSTRACT

This article introduces various handover scenarios in multihop cellular networks. In addition, this article presents handover schemes where relay stations are located either inside a cell or on the boundary between two adjacent cells and investigates the effects of the deployment position of relay stations to handover performance. The simulation results show that multihop cellular networks for both deployment scenarios can achieve 90 percent throughput increase over single-hop cellular networks. The results also show that the overall throughput of the multihop cellular networks with relay stations inside a cell is higher than for those with relay stations on the boundary between two adjacent cells, whereas the opposite is observed for the throughput of cell-boundary users. The intercell handover latency in multihop cellular networks is increased by 20 ~ 56 percent compared with that in single-hop cellular networks because of the increased number of handovers and signaling overhead. However, by deploying relay stations on the boundary between two adjacent cells, the service-interruption time caused by inter-cell handover is reduced by 80 percent compared with that of single-hop cellular networks.

INTRODUCTION

Recently, multihop or relay networks have been widely considered as supplementary technology in next-generation wireless systems such as fourth generation (4G), Third Generation Partnership Project Long-Term Evolution (3GPP LTE), and IEEE 802.16m to increase the cell radius or combat the shadowing effect, which is mainly caused by large obstacles between transceivers [1–5]. The application of the multihop concept to cellular networks, however, raises many technical issues, such as the best positions for the base station (BS) and relay stations (RSs); the number of RSs; spectrum allocation and multiplexing between the BS and RSs; scheduling; and handover [1, 2]. Especially, the introduction of RS in cellular networks creates additional handover scenarios and increases the number of handovers. In conventional cellular systems, handover occurs only when a mobile station (MS) moves to different cells or different sectors of the same cell, whereas additional handovers occur in a multihop cellular network (MCN) between the BS and the RSs or between two different RSs. These additional MCN handovers can cause serious ping-pong problems and increase signaling overhead.

Unfortunately, only a few research articles have studied handover issues in relay networks [6–8]. Reference [6] proposes a relay-assisted vertical handover scheme for hybrid cellular and wireless local area network (WLAN) systems. In this scheme, an active MS, which has an active connection, uses non-active MSs as relay stations in handover regions when a handover request is rejected or delayed for an unacceptable period of time. Reference [7] categorizes and evaluates the performance of different types of handovers in multihop radio-access networks (MRANs). In [7], the multihop handover schemes are categorized into forced handover and route optimization-based handover, according to the handover initiation method. In addition, the signaling mechanisms for these handover scenarios are proposed, and handover delay and signaling overhead are also investigated. Reference [8] proposes relay-assisted handover with geo-locating information in a hybrid ad hoc cellular system. An MS is assumed to be able to establish direct connections with nearby MSs to form a temporary wireless relay network, and the MS-to-MS interface is based on a WLAN protocol like the integrated cellular and ad hoc relaying (i-CAR) system proposed in [2]. The relay-assisted handover in [8] uses relaying technology to avoid unnecessary handovers and call drop caused by abrupt channel degradation.

Although these works studied relay-handover problems, the main interest of using the multihop concepts is not in cellular networks but in ad hoc networks. Recently, the handover schemes in MCNs have been defined in the baseline document for the draft standard of IEEE 802.16j [9]. Media-access-control (MAC)-layer handover procedures in IEEE 802.16j are almost the same as handover procedures in IEEE 802.16e, except that an RS relays handover-associated messages between an MS and a BS. However, scanning procedures in IEEE 802.16j are more complicated than those in IEEE 802.16e. IEEE 802.16j defines two different types of RS such as transparent-mode RS and non-transparent-mode RS, which support the one-hop relay and the two-hop relay, respectively. Only a non-transparent-mode RS broadcasts its own preamble, time-aligned with its

This work was supported by the Korea Research Foundation Grant funded by the Korean Government (MOEHRD, Basic Research Promotion Fund) (KRF-2006-352-D00138).

 0163-6804/09/$25.00 © 2009 IEEE

serving BS preamble. Therefore, a non-transparent-mode RS is included in a neighbor list for handover scanning and broadcasts a different neighbor advertisement message that is suitable for its service area [9]. On the other hand, a transparent-mode RS is excluded from a neighbor list, and a neighbor advertisement message is broadcasted by a BS. IEEE 802.16j also defines a mobile RS handover to support an RS migration from a serving BS to another BS. However, few technical contributions in IEEE 802.16j have studied the effect of relay-deployment structures to the design of handover algorithms and the handover performance of the cellular systems because IEEE 802.16j must be backward-compatible with the IEEE 802.16e systems, and consequently, the main interesting points are to design handover-signaling procedures and amend the current mobile-broadband wireless access (MBWA) specification [10].

This article investigates handover schemes for two different RS deployment structures, where an RS is located either inside a cell or on the boundary between two adjacent cells. Detailed cell planning, spectrum allocation, frame structure, and handover signaling for MCNs are also explained. The performance of the handover schemes in an MCN is compared with a single-hop cellular network (SCN) in terms of the average throughput of all users and the cell-boundary users, the handover latency, and the signaling overhead.

Handover Scenarios in MCNs

This section describes all possible handover scenarios in MCNs regardless of the RS deployment structure. Handovers in MCNs can occur when an MS moves between different BSs, between different RSs, or between a BS and an RS. Figure 1 depicts diverse handover scenarios in MCNs. This article classifies each handover scenario into intracell handover and intercell handover with the inter-sector or softer handover excluded. The detailed handover scenarios are explained in the following sections.

Intracell Handover

Scenario 1 (Intracell RS-RS Handover) — In this scenario, an MS performs handover between two different RSs in the same cell. In Fig. 1, MS1 moving from RS1b to RS1a in cell 1 illustrates handover scenario 1. The BS easily can control the handover process because both the serving RS and the target RS are under their own control, and inter-BS information or signaling is not required. To prevent packet loss during the handover process, automatic retransmission request (ARQ) status should be consistent between a serving RS and a target RS. In this scenario, transfer of the ARQ status to the target RS is not required during handover if the ARQ function is located in the BS. In addition, the current layer 3 (L3) address can be used after the handover.

Scenario 2 (Intracell BS-RS Handover) — In scenario 2, an MS changes its communication node from a BS to the RS of the same cell, or vice versa, which is the handover of MS2 in Fig. 1. In this case, the BS easily can control the handover process as in scenario 1, and the ARQ status transfer and the L3 address renewal are not required.

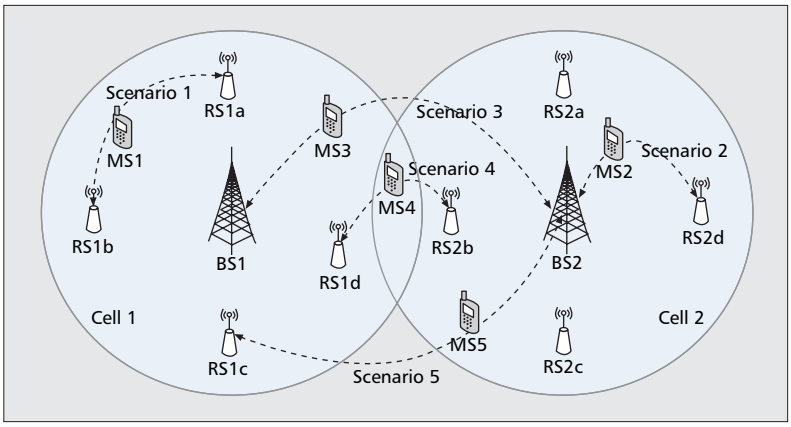

■ **Figure 1.** *Handover scenarios in multihop cellular networks.*

Intercell Handover

Scenario 3 (Intercell BS-BS Handover) — Scenario 3 is exactly the same as the inter-BS handover in the conventional cellular systems. The ARQ status should be transferred during the handover process if the ARQ is controlled by the BS, and the L3 address is also reassigned when the subnet is changed by the handover.

Scenario 4 (Intercell RS-RS Handover) — In scenario 4, an MS performs handover from an RS to the RS of different cells, which is the handover of MS4 in Fig. 1. This scenario can cause relatively larger signaling overhead than other scenarios because it requires inter-BS signaling and RS-BS signaling in both cells. In addition, the channel quality of the MS in this handover region can be seriously attenuated by the intercell interference from the adjacent BSs and RSs. The ARQ status and the L3 address management are similar to those of scenario 3.

Scenario 5 (Intercell BS-RS Handover) — In scenario 5, an MS moves from a BS to the RS of different cells, or vice versa, which is the handover of MS5 in Fig. 1. Inter-BS signaling is also required in this scenario, and the ARQ status and the L3 address management are the same as scenario 3 except for additional signaling between the RS and the BS.

Handovers for the Different MCN Structures

The design principle and performance of handovers depend highly on the RS deployment structure in the MCN. This section describes the detailed handover operation and signaling according to the RS deployment structures. This section focuses especially on the intercell RS-RS or BS-RS handover schemes for different RS deployment structures because these handovers are the most complicated scenarios among the available handover scenarios in MCNs and occur more frequently than intercell BS-BS handovers

The frequency-reuse factor (FRF) of the MCN is basically assumed to be one, which means that all BSs and RSs share the same spectrum band. However, fractional frequency reuse schemes can be used to mitigate inter-cell interference.

if the RSs are densely deployed on the cell-boundary region. This article considers the fixed multihop systems where RSs are fixed like a BS. In addition, it is assumed that the handover decision is made by BSs to increase the handover efficiency and reduce the cost of RSs. The following sections describe two different RS deployment structures and the detailed intercell handover schemes under the described RS deployment structures.

Relay-Deployment Structures

Cell Structure — The cell structure of an MCN can be varied according to the deployment method of RSs. This article considers two kinds of RS deployment methods as shown in Fig. 2a and Fig. 2b. As shown in Fig. 2a and Fig. 2b, RSs are located inside a cell in MCN structure 1 (MCN 1), while RSs are located on the boundary between two adjacent cells in MCN structure 2 (MCN 2). This article considers these two structures because MCN 1 is most widely used in the previous works [1–3], whereas MCN 2 is much preferable to intercell RS-RS handover. Most of the previous work [1–3] assumes that RSs are located inside a cell and controlled by the BS to increase the cell coverage or high-data-rate region, as in MCN 1. Meanwhile, RSs are supposed to be synchronized with the adjacent two BSs and controlled by the BSs in MCN 2; for example, RS1 is synchronized with BS0 and BS1 in Fig. 2b and shares the radio resources and control information with BS0 and BS1. The frequency-reuse factor (FRF) of the MCN is basically assumed to be one, which means that all BSs and RSs share the same spectrum band. However, fractional frequency reuse schemes [10, 11] can be used to mitigate intercell interference.

Multiplexing and Frame Structure — Because BSs and RSs are assumed to use the same frequency bands, the resources of BSs and RSs should be coordinated for multiplexing. Time-division multiplexing (TDM) and frequency-division multiplexing (FDM) mainly were used in the latest works [1, 2]. This article mainly considers TDM between BS and RSs; whereas, FDM also can be used without significant modification of the considered frame structure. Figure 2c depicts an example of TDM-based frame structure for an MCN, which is compatible with the non-transparent-mode frame structure of IEEE 802.16j [9] except that an RS can transmit its preamble. In this frame structure, the duplexing between downlink (DL) and uplink (UL) is based on time-division duplexing (TDD). Each frame in the DL transmission begins with a preamble followed by a frame-control header (FCH), a DL map (DL-MAP), and possibly a UL map (UL-MAP). BSs and RSs transmit their own preamble signal in a preamble region. An FCH is a control header to inform the transmission scheme and length of MAP. The DL-MAP and UL-MAP include DL and UL channel-usage information, respectively. The R-MAP, which indicates the DL and UL relay channel usage, is located following the UL-MAP or defined as an extension of the MAP. For the frame synchronization among a BS, RSs, and MSs in a cell, RSs and MSs must receive control information such as the FCH, DL-MAP, and UL-MAP regardless of channel quality. Therefore, the FCH, DL-MAP, and UL-MAP would be transmitted using low-order modulation with lower code rate. The frame structure following the MAP consists of a DL sub-frame period and UL sub-frame period. The DL sub-frame includes at least one DL access zone for the BS to its subordinate MS or RS transmissions, and optionally, includes a transparent zone for the RS to its subordinate MS transmissions. The UL sub-frame includes the UL access zone and UL relay zone. The MS can transmit UL control or data information to the BS or RS during the UL access zone; meanwhile the RS can transmit UL control or data information to the BS or another RS during the UL relay zone. The ranging channel in the UL access zone is shared by the RS and MS. In each frame, a relay receive/transmit transition gap (R-RTG) is inserted between the DL access zone and the DL transparent zone and between the UL access zone and the UL relay zone. In MCN 2, the same DL and UL sub-frame for an RS are allocated by the adjacent two BSs.

Intercell Handover Schemes in MCN

Handover Scenarios — Although intracell handover scenarios are almost the same for both deployment schemes, intercell handover scenarios can be quite different. Figure 3a shows the point of handover executions and the related handover scenarios of two different RS deployment structures. As seen in this figure, MS1 located in MCN 1, moves from cell 1 to cell 2, and accordingly, performs three handovers, namely, intracell BS-RS handover, intercell RS-RS handover, and intracell BS-RS handover. However, MS2 located in MCN 2 moves from cell 2 to cell 3 and performs only two handovers, namely, intracell BS-RS handover and intercell BS-RS handover. In this case, MS2 can perform the intercell handover without switching the current serving RS, which is quite different from the previously defined handover scenario 4 or 5. Thus, this type of handover is newly defined as the advanced handover scenario 5 or advanced intercell BS-RS handover, whereas the MS handover procedure in IEEE 802.16j falls under handover scenario 4. The detailed intercell handover procedures and signaling in MCN 1 and 2 are compared in the following section.

Received-Signal-Strength Measurement and Handover Decision — The received-signal-strength (RSS) measurement process can cause packet-loss or packet-transmission delay because an MS cannot receive data packets from the serving BS during the RSS measurement for the neighbor BSs. To minimize the packet-loss or the transmission delay during measurement, puncturing has been used in conventional code-division multiple-access systems. In mobile WiMAX systems based on IEEE 802.16e standardization [10], an MS requests a measurement period during which a BS does not send any data packets to the MS to prevent the packet loss. In an MCN, the RSS measurement overhead

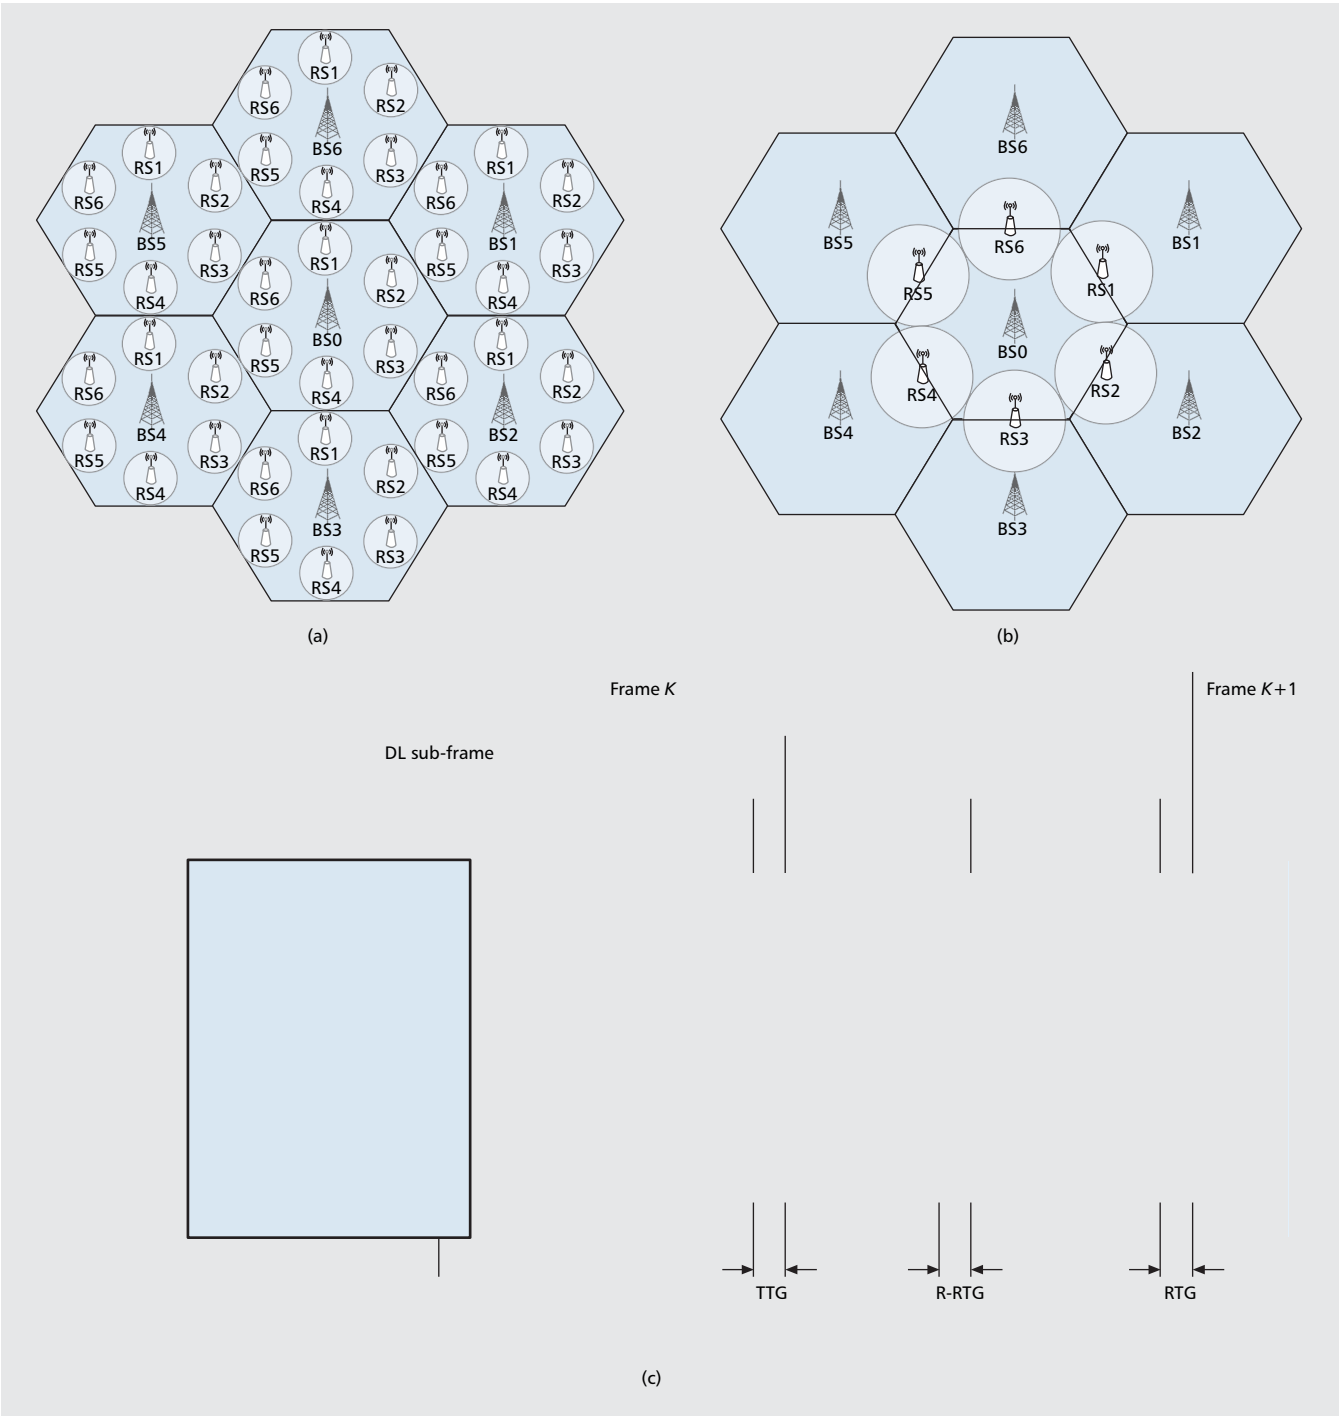

■ **Figure 2.** *The cell and frame structures for multihop cellular networks: a) multihop cellular network structure 1; b) multihop cellular network structure 2; c) an example of frame structure for multihop cellular networks.*

increases because an MS should measure the RSS of the neighbor RSs, as well as the neighbor BSs. Therefore, packet-loss or packet-transmission delay caused by handover can be increased in an MCN compared to an SCN. In IEEE 802.16j, a non-transparent-mode RS can make its own neighbor list based on its service area; accordingly, an MS should receive neighbor lists from a serving BS and an RS for precise measurement. In the advanced handover scenario 5, an RS can generate its own neighbor list, or a serving BS can make a neighbor list for RSs and forward the list to each RS. Regardless of the signaling procedures for RSS measurement, the MCN 2 can reduce the measurement overhead compared with the MCN 1 because the number of RSs in the measurement list is smaller than it is in the MCN 1. In Fig. 3a, MS2 measures the RSS of BS2, RS3, and BS3, while MS1 measures the RSS of BS1, RS1, RS2, and BS2. The handover decision is based on the RSS measurement result. When the serving BS receives the measurement report from the MS or the RS, it determines the handover execution and direc-

Numeric range	Format S = sign bit M = mantissa 0 = exponent 1 = separator	Significant binary digits
1.0–0.5	SMMMMMMMMMMMMMM1	15
0.5–0.25	SMMMMMMMMMMMMM10	14
0.25–0.125	SMMMMMMMMMMMM100	13
0.125–0.0625	SMMMMMMMMMMM1000	12
0.0625–0.03125	SMMMMMMMMMM10000	11
0.03125–0.015625	SMMMMMMMMM100000	10
0.015625–0.0078125	SMMMMMMMM1000000	9
0.0078125–0.00390625	SMMMMMMM10000000	8
0.00390625–0.001953125	SMMMMMM100000000	7
0.001953125–0.0009765625	SMMMMM1000000000	6
0.0009765625- 0.00048828125	SMMMM10000000000	5
0.00048828125–0.000244140625	SMMM100000000000	4
0.000244140625–0.0001220703125	SMM1000000000000	3
0.0001220703125–0.00006103515625	SM10000000000000	2
0.000030517578125	s100000000000000	1
0	s000000000000000	0

Table 2. *Fractional fixed point data format recast with trailing zeros.*

This format provides almost double the dynamic range of standard fixed point at the cost of decreased precision at the very top of the numeric range (1.0–0.5). It also provides improved precision over fixed point throughout the remainder of the numeric range.

article we examine only one new data format, a new fractional format. The main reasons for introducing a new fractional format are that this format has a path into a C data type with the embedded-C standard (ISO, 2003a), and it has a proven record for use in signal-processing applications. If we can develop a fractional format with improved dynamic range that doesn't sacrifice noise performance, we will have accomplished something worthwhile.

Any format optimized for fractional DSP should take a clue from the noise performance of fixed point and have the largest mantissas at the top of the range. The mantissa should then fall off gradually to the smallest mantissa. Such a format derived from our new class is shown in Table 4 and illustrated in Fig. 1b. This format provides a significant dynamic-range improvement over the standard fixed-point format and the prevailing 16-bit floating-point format (s10e5 or binary16). It should also provide significantly improved noise performance over the binary16 floating-point format. Its only weakness is that a single bit of precision was sacrificed in the top two ranges when compared to fixed point. For a fractional format, this is significant, and we can do better by sacrificing some dynamic range.

By examining the binary fixed-point format, we learn that a single term from each of the ranges (A1 through M13 in Table 4) is combined to form a single 15-bit magnitude. This procedure is repeated to form a single 14-bit magnitude and on down the line. The result is that the binary fixed-point format can be considered a single instance of the new class of floating-point formats we developed in Table 3. To develop a new fractional format, we follow the same approach, but skip the creation of the first 15-bit magnitude. This leaves us with two terms of each exponent size, and we order them in decreasing precision as shown in Table 5 and illustrated in Fig. 1c.

This format provides almost double the dynamic range of standard fixed point at the cost of decreased precision at the very top of the numeric range (1.0–0.5). It also provides improved precision over fixed point throughout the remainder of the numeric range. As we discover, the increased dynamic range and additional overall precision actually improve the round-off noise performance of this format when compared to 16-bit fixed point.

Simulation and Results

A simulation was performed to validate improved noise performance for the second new fractional format. Here we provide a summary

Format S = sign bit M = mantissa 0 = exponent 1 = separator	Significant binary digits (assuming normalized Mantissa and implied leading 1)	Identifier (for later reference)
S1MMMMMMMMMMMMM1	14	A1
S1MMMMMMMMMMMM10, S01MMMMMMMMMMMM1	13	B1 B2
S1MMMMMMMMMMM100, S01MMMMMMMMMMM10, S001MMMMMMMMMMM1	12	C1 C2 C3
S1MMMMMMMMMM1000, S01MMMMMMMMMM100, S001MMMMMMMMMM10, S0001MMMMMMMMMM1	11	D1 D2 D3 D4
5 values	10	E1–E5
…	…	
S1100…–S0…011	1	M1–M14
S1000…–S0…001	1	N1–N15
S000000000000000	0 & special use for S = 1	O1, O2

Table 3. *A new class of floating point formats.*

of the simulation and results, which are described in greater detail elsewhere [8].

An amplitude modulation (AM) receiver simulation was used to compare the noise performance for various data formats. The AM format was selected because the problem is well understood, and even if in decline, the format is still widely used. This simulation contained several typical DSP algorithms, which include the following: quantization to simulate 16-bit analog-to-digital (A/D) conversion, finite impulse response (FIR), and infinite impulse response (IIR) filters, demodulation, AGC, Hanning window, fast Fourier transform (FFT), and signal-to-noise ratio (SNR) measurement through Parseval's Theorem. AGC techniques were used following several stages to improve the native performance of the fractional format.

The following formats were simulated: IEEE 32-bit floating point (as a comparison baseline), s10e5 16-bit floating point, a 16-bit logarithmic format, fractional 16-bit fixed point, and the second new fractional format presented here. The simulation was performed for both weak-signal and strong-signal cases and both with and without the use of a single, large, post-multiply accumulator. Noise was not added to the simulation, so the resulting noise is a consequence of round-off errors during calculation, quantization to simulate A/D conversion, and out-of-band filter rejection (just over 50 dB). The simulation results are shown in Table 6.

As can be seen from Table 6, the new fractional format significantly outperformed the other 16-bit formats in terms of noise performance, and it approaches the performance of 32-bit floating point for this simulation. It also provides almost twice the dynamic range of traditional fixed point. The first fractional format was not simulated, but would drastically increase dynamic range with only a minor degradation of noise performance when compared to the second fractional format. It is still very useful for problems that require a larger dynamic range.

Better digital filters can also be achieved using the new fractional format than are possible using either fixed- or floating-point data types of equivalent word width. As shown in Fig. 2, a FIR-notch filter was created using the IEEE 32-bit floating-point format. The coefficients were then converted (with rounding) to the 16-bit fixed-point fractional format, the s10e5 16-bit floating-point format, and the 16-bit new fractional format. The coefficients were then converted back to floating point, and their magnitude response was plotted. As seen in Fig. 2, when the digital coefficients were expressed in the new fractional format, the rejection band improved by approximately 10 dB compared to the other two 16-bit formats. This improvement may become even more pronounced in IIR filters, which employ feedback.

Implementation Considerations

Having derived a data format that actually improves numeric performance for 16-bit DSP applications, we are left with several implementation issues that must be addressed before the format is truly useful. For non-real-time software applications such as simulation or data storage, the new fractional format has been coded into a C++ class with overloaded arithmetic operators that has provided excellent results. Unfortunately, this class performs arithmetic operations at a much slower rate than a traditional fractional format would. Furthermore, a data format with two exponent fields and a table-lookup mechanism may be difficult to implement in hardware and may execute more slowly than a traditional fractional format. The next logical step in our format development process is to render the second new fractional format into a form that is easily implemented in hardware. We accomplish this with a reorganization of the second new fractional format that allows easy encoding and decoding of the format into a 32-bit two's complement number.

Assume a 16-bit number partitioned into two fields. The first field is a single-bit field representing a shift flag. The second field is a 15-bit compressed two's complement number. The number of sign bits (leading ones or zeros) in the two's complement field is reduced by half (rounding up). This organization of bits proves equivalent to the new fractional format illustrated in Fig. 1c and defined in Table 5. Here we provide a mechanism for expanding the 15-bit field into a 32-bit two's complement number, and the shift bit indicates if the number is to be shifted to the left by one bit or not.

To decode the format, the leading digits of the 15-bit two's complement field that are all of the same binary value are to be doubled. In

Numeric range	Format S = sign bit M = mantissa 0 = exponent 1 = separator Red = actual digit Black = implied digit	*N*	*F*
1.0–0.5	S0.1MMMMMMMMMMMMM1	14	A1
0.5–0.25	S0.01MMMMMMMMMMMM1	13	B2
0.25–0.125	S0.001MMMMMMMMMMMM10	13	B1
0.125–0.0625	S0.0001MMMMMMMMMMM10	12	C2
0.0625–0.03125	S0.00001MMMMMMMMMMM100	12	C1
0.03125–0.015625	S0.000001MMMMMMMMMMM1	12	C3
0.015625–0.0078125	S0.0000001MMMMMMMMMM1	11	D4
0.0078125–0.00390625	S0.00000001MMMMMMMMMM1000	11	D1
0.00390625–0.001953125	S0.000000001MMMMMMMMMM100	11	D2
0.001953125–0.0009765625	S0.0000000001MMMMMMMMMM10	11	D3
…	…	…	
	S0. "104 0s" 000000000000001	1	N15
0.0	0000000000000000	0	O1

Notes: *N* = number of significant binary digits. *F* = format from Table 3.

Table 4. *A first attempt at a new fractional format.*

Having derived a data format that actually improves numeric performance for 16-bit DSP applications, we are left with several implementation issues that must be addressed before the format is truly useful.

other words, if there are five ones or five zeros at the beginning of this field, you replace them with ten ones or ten zeros. If the shift bit is a zero, then you remove one of the leading sign bits at the front of the field. This number is then left justified into a 32-bit two's complement format, and the least significant bits are all set to zero. You now have expanded your 16-bit number into a 32-bit number with maximum precision for the very largest fractions.

The data can now be operated upon with a traditional 32-bit two's complement arithmetic unit. However, we should point out that a 15-bit DSP multiplier may be more appropriate because there are not more than 15 bits of precision in either multiplicand. The resulting product can be shifted into a 32-bit format after the 15×15 bit multiplication is complete.

Encoding from a 32-bit (or larger) two's complement word to a 16-bit new fractional format is accomplished by following the reverse operation. The number of leading zeros is halved (rounding up). The shift bit is then set to zero if the original number of leading sign bits is odd. Once the number of leading sign bits has been halved and the shift bit set, then rounding must be performed to compress the value into the resulting 15-bit two's complement format.

Although the number format itself has not increased in size (16-bits) with this approach, the computational element (multiplier, arithmetic logic unit [ALU], accumulator) may now have increased in size from 16 to 32 bits (at least for addition). This format should prove easy to implement in hardware but may result in a drop in throughput and an increase in power consumption due to the use of a 32-bit arithmetic unit together with encoding and decoding logic. Though this approach may slightly reduce the throughput of a 16-bit DSP device, it should not significantly affect the size or cost of the device because these attributes are driven primarily by the size of the main data and address buses. The use of a smaller 15×15-bit multiplier may also help with the size issue.

Some might argue that a hardware implementation of this new format would not be worthwhile because it does not directly improve important hardware performance parameters such as circuit complexity, delay, and power consumption. This point is granted. However, implementing this new format in hardware directly improves algorithm performance for many DSP applications. Obviously, not all DSP applications require additional dynamic range or improved noise performance, but other DSP applications go to great lengths algorithmically to obtain additional performance in these two categories. For these types of systems, improved algorithm

to be useful in a real-time application, the new format should be directly implemented in hardware, most conveniently in the arithmetic element of a programmable 16-bit DSP processor. Almost every application that uses a 16-bit DSP could be improved by using the new fractional format.

performance may indirectly improve power consumption, delay, algorithm complexity, cost, and size.

For many applications, the slight trade-off in throughput and power is well worth the improved dynamic range and noise performance achieved with the new format. Many applications that now require 32-bit floating point could be realized at less cost if this new data format were natively provided in a 16-bit programmable DSP device. Furthermore, many applications that currently utilize programmable 16-bit DSPs could experience improved performance with this new format.

Conclusions

In this article we introduced a new class of floating-point formats with uneven numeric precision. We illustrated how to derive specific formats from this class by creating two new fractional formats. We then presented simulation results to verify that the second new fractional format outperformed traditional 16-bit fixed- and floating-point formats in terms of noise performance and also dynamic range (for fixed point). Finally we remapped the new derived format into a form that is more convenient for hardware implementation.

Numeric range	Format S = sign bit M = mantissa 0 = exponent 1 = separator Red = actual digit Black = implied digit	*N*	*F*
1.0–0.5	S0.1MMMMMMMMMMMMM1	14	A1
0.5–0.25	S0.01MMMMMMMMMMMMM10	14	B1, C2, D3,..., M13
0.25–0.125	S0.001MMMMMMMMMMMM1	13	B2
0.125–0.0625	S0.0001MMMMMMMMMMMM100	13	C1, D2, ...
0.0625–0.03125	S0.00001MMMMMMMMMMM1	12	C3
0.03125–0.015625	S0.000001MMMMMMMMMMM1000	12	D1, E2, ...
0.015625–0.0078125	S0.0000001MMMMMMMMMM1	11	D4
0.0078125–0.00390625	S0.00000001MMMMMMMMMM10000	11	E1, F2, ...
0.00390625–0.001953125	S0.0000000001MMMMMMMMM1	10	E5
...	...	...	
	S0. "16 0s" 000000000000001	1	M13
0.0	0000000000000000	0	O1

Notes: *N* = number of significant binary digits. *F* = format from Table 3.

Table 5. *A second attempt at a new fractional format.*

Format	Weak signal SNR without accum	Weak signal SNR with accum	Strong signal SNR	Dynamic range
IEEE 754 32-bit floating point	31.97	31.97	50.53	1530
s10e5	8.44	7.90	42.06	252
16-bit logarithmic	8.23	8.32	38.61	385
16-bit fixed point fractional	13.30	24.93	44.42	96
New fractional (#2)	21.91	27.16	50.13	181

Note: Among the 16-bit formats blue indicates the best performer and red the second best performer.

Table 6. *Summary of simulation results (dB).*

We emphasized the fact that to be useful in a real-time application, the new format should be directly implemented in hardware, most conveniently in the arithmetic element of a programmable 16-bit DSP processor. Almost every application that uses a 16-bit DSP could be improved by using the new fractional format. Some obvious potential applications are:

- CD audio and other digital audio data formats
- Satellite TV and HDTV radio frequency (RF) data formats for signal processing
- Modems and other telephony applications
- Toys that employ speech generation or recognition
- Cell phone and software radio applications
- Implanted DSP devices (e.g., cochlear ear implants)

Some important aspects of this article that must be emphasized are:

- Both the new class of floating-point formats and the new fractional formats are applicable to data sizes other than 16-bits.
- The new class of formats may have applications outside the realm of DSP.
- Additional new data formats and more detailed development is provided in the authors' previous work [8].

For many years, the signal-processing industry has used primarily either fixed-point or binary floating-point number formats to represent digital signals. If the field of data formatting is opened up to include irregular formats, actual performance in terms of dynamic range and round-off noise can improve.

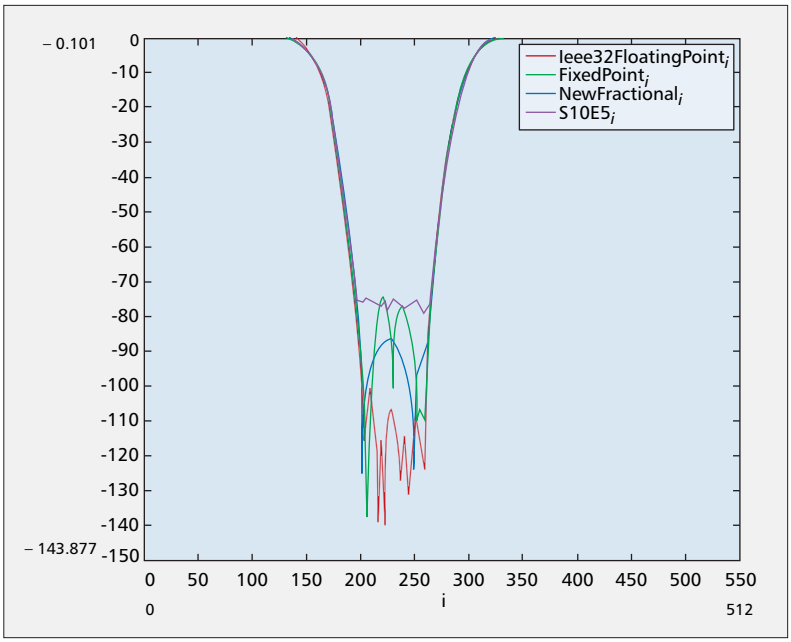

Figure 2. *A comparison of digital filter coefficient performance.*

References

[1] ISO/IEC JTC1/SC22 WG14/N854, "DSP-C: An Extension to IOS/IEC IS 9899:1990," 1998; http://www.open-std.org/jtc1/sc22/wg14/www/docs/n854.pdf (accessed Nov. 13, 2006).

[2] A. V. Oppenheim and R. W. Shafer, *Discrete-Time Signal Processing*, 2nd ed., Prentice-Hall, 1999.

[3] ISO/IEC JTC1/SC22 WG14/N1021, "Extensions for the Programming Language C to Support Embedded Processors: ISO/IEC DTR 18037," 2003; http://www.open-std.org/jtc1/sc22/wg14/www/docs/n1021.pdf (accessed Nov. 13, 2006).

[4] IEEE Std P754 Draft 1.2.5, "Draft Standard for Floating-Point Arithmetic"; math.berkeley.edu/~scanon/754 (accessed Nov. 13, 2006).

[5] ISO, "Rationale for International Standard-Programming Languages — C," revision 5.1, 2003; http://www.open-std.org/jtc1/sc22/wg14/www/C99RationaleV5.10.pdf (accessed Oct. 18, 2006).

[6] S. W. Smith, *The Scientist and Engineer's Guide to Digital Signal Processing*, California Technical Publishing, 1999; http://www.analog.com/processors/learning/training/dsp_book_index.html (accessed Aug. 18, 2006).

[7] B. Parhami, *Computer Arithmetic: Algorithms and Hardware Designs*, Oxford Univ. Press, 1999.

[8] M. Richey, "The Application of Irregular Data Formats for Improved Performance in 16-bit Digital Signal Processing Systems," Master's thesis, Univ. of Kansas, Dec. 2006.

Biographies

HOSSEIN SAIEDIAN [SM] (saiedian@eecs.ku.edu) received his Ph.D. from Kansas State University in 1989. He is currently a professor and an associate chair in the Department of Electrical Engineering and Computer Science at the University of Kansas (KU) and a member of the KU Information and Telecommunication Technology Center (ITTC). His primary area of research is software engineering and information security. He has over 100 publications on a variety of topics in software engineering and computer science.

MANUEL RICHEY (manuel.richey@honeywell.com) graduated from the University of California at Davis in 1983 with a Bachelor's degree in electrical engineering. In 2007 he earned his Master's degree in computer science from the University of Kansas. He is a principal engineer at Honeywell International Inc. where he has worked since 1985. At Honeywell, he has worked primarily in the fields of digital signal processing and software radio. He holds seven U.S. patents and has four additional patents pending.

IEEE Communications | Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page | qMags

ACCEPTED FROM OPEN CALL

Understanding Interference and Carrier Sensing in Wireless Mesh Networks

Jeongkeun Lee and Sung-Ju Lee, Hewlett-Packard Laboratorie
Wonho Kim, Princeton University
Daehyung Jo, Samsung Electronics
Taekyoung Kwon and Yanghee Choi, Seoul National University

ABSTRACT

Wireless mesh networks aim to provide high-speed Internet service without costly network infrastructure deployment and maintenance. The main obstacle in achieving high-capacity wireless mesh networks is interference between the mesh links. In this article, we analyze the carrier sensing and interference relations between two wireless links and measure the impact of these relations on link capacity on an indoor 802.11a mesh network testbed. We show that asymmetric carrier sensing and/or interference relations commonly exist in wireless mesh networks, and we study their impact on the link capacity and fair-channel access. In addition, we investigate the effect of traffic rate on link capacity in the presence of interference.

INTRODUCTION

Most research in wireless mesh networks (WMNs) focuses on the capacity improvement in the mesh backbone. As wireless mesh nodes typically have no mobility and employ carrier-sense multiple-access (CSMA) protocols in a multihop environment, interference is the main obstacle in achieving high capacity. Understanding the wireless link behavior in the face of interference helps to manage the network performance by various measures: traffic engineering, route selection, layout planning, and so on. Because of the multihop nature of WMNs, WMNs suffer from hidden interference and self interference. Hence, the effect of interference is more severe than in single-hop, wireless local-area networks (WLANs). Recent work proposes to use multiple channels and/or multiple radios [1] to overcome interference and increase capacity. Nevertheless, we concentrate our effort on studying interference together with carrier sensing in a single-channel WMN because:

- It is a good first step in investigating the wireless-channel behavior with interference.
- The findings still apply to networks with multiple channels and radios.
- Many deployed WMNs use only a single radio and channel because of low cost and complexity [2].

We study interference in a WMN by investigating its relation with carrier sensing (CS). We categorize their relations on two wireless links and perform testbed experiments to evaluate their effect on the link performance. Using our indoor 802.11a testbed measurements, we study the degree of interference on each relation category and also the fairness of channel access between the two links. We use both saturated and rate-controlled network traffic to investigate the impact of traffic load on link performance.

The IEEE 802.11 standard suggests the use of request-to-send/clear-to-send (RTS/CTS) to solve the hidden interference problem. However, the RTS/CTS mechanism is shown to be ineffective in eliminating the hidden node problem [3] and fails to increase the multihop capacity [4]. Therefore, we investigate the CS and interference relations without considering RTC/CTS.

We believe the testbed experimental analysis helps us better understand interference and its impact on capacity—an aspect that is critical for the protocol design and wireless capacity improvement. We characterize non-endpoint- sharing, pairwise 802.11 links, based on their CS and interference relation. This method of link-pair categorization was proposed by Garetto *et al*. [5] using modeling and simulation. But there is a requirement to validate models in a real system deployment. We describe the behavior of 802.11 links based on measurements and present new findings that modeling and simulation cannot reveal.

The rest of the article is organized as follows. We provide a background study in the next section. The following two sections present testbed results and analyses; the relation of CS and interference with saturated traffic is investigated, and then the rate-controlled traffic is discussed. We conclude in the final section.

BACKGROUND

To define CS and interference relations, we consider two links, L_1 and L_2, that do not share end-points. On L_1 and L_2, nodes S_1 and S_2 trans-

This work was supported in part by the Hewlett-Packard Laboratories University Relations Program, the IT R&D program of MKE/IITA [2007-F-038-03, Fundamental Technologies for the Future Internet], and the IT R&D program of MKE/IITA [2008-F-034-02, Development of Security-Quality Guarantee Technology in Resilient Networks].

0163-6804/09/$25.00 © 2009 IEEE

IEEE Communications | Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page | qMags

mit data to nodes R_1 and R_2, respectively, as illustrated in Fig. 1. Arrows in the figure indicate the directions of data/acknowledgment (ACK) packets, interference signals, and sensed signals. The CS relation between the two links is decided by the two senders S_1 and S_2, whereas the interference relation depends on how one sender affects the other link: that is, S_1's effect on L_2 and S_2's effect on L_1.

CS Relation

In the 802.11 CS mechanism, a wireless station withholds its transmission when it senses an ongoing transmission on the medium. In Fig. 1, when S_1 senses S_2's transmission, a CS relation exists between the two links, and we say L_2 is carrier sensed by L_1. We estimate whether S_1 senses the transmission of S_2 by observing S_1's broadcast transmission throughput when S_2 is transmitting simultaneously. When S_1 senses S_2's transmission due to the 802.11 CSMA/CA mechanism, its transmission throughput is about one-half of the throughput measured when S_1 does not sense S_2's transmission. Figure 1 illustrates an asymmetric CS relation. In our testbed network deployed at Hewlett-Packard Laboratories (HP Labs) in Palo Alto, 17 percent of the node pairs exhibit an asymmetric CS relation, 14 percent exhibit a mutual CS relation, and 69 percent of node pairs do not sense each other.

Interference Relation

We define the interference relation independently of the CS relation. In Fig. 1, if S_1 and S_2 simultaneously transmit and S_2's transmission hinders R_1's reception of S_1's packet, we say L_1 is interfered by L_2 (or S_2) regardless of whether the interferer S_2 is carrier sensed by S_1. In reality, if both S_1 and S_2 sense each other, they do not transmit simultaneously, assuming the CS and collision-avoidance mechanisms. However, we define the interference relation independently of the CS relation for the simplicity of case enumeration.

Interference becomes more effective as the signal-to-interference ratio (SIR) at R_1 becomes smaller due to S_2's transmission. The SIR is defined as "received signal strength (RSS) from S_1 *minus* RSS from S_2" in dB scale. If the SIR at R_1 goes below a certain threshold (which depends on the PHY bit rate of S_1's transmission), S_2 effectively interferes R_1's reception of packets from S_1. If the SIR is above the threshold, the physical layer capture (PLC) in 802.11 occurs, and S_1's signal is successfully received at R_1 despite S_2's simultaneous transmission. The arrival-timing order and the RSS of two packets at a receiver determines whether PLC occurs [6].

Note that R_1's ability to receive S_2's packets does not necessarily mean that R_1 is interfered by S_2. Its converse is not true, either. For example, if S_1 is much closer to R_1 than S_2 is, R_1 can still decode the packets from S_1 due to the high SIR despite the interferer S_2. In general, interference is determined by the SIR relation between the two senders and one receiver. Thus, estimating interference, based solely on the existence of a communication link, which is popular in the literature, is not correct, and hence, recent research tends to use the SIR-based interference models.

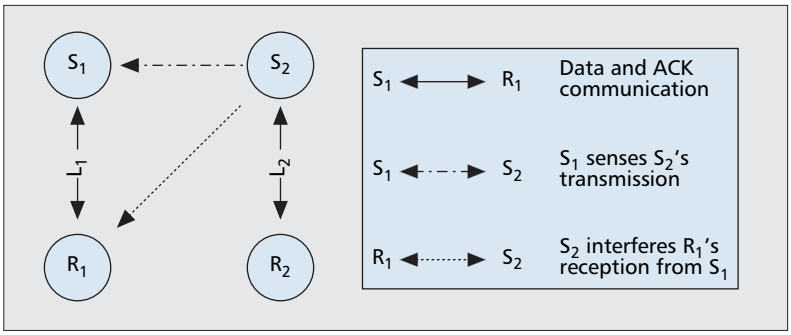

■ **Figure 1.** *Example of carrier sensing and interference between* L_1 *and* L_2.

As we consider large data packets (i.e., a medium access control [MAC] frame size of 1064 bytes), we ignore the interference caused by ACK packets because the frame size is 14 bytes. Supporting observations have been made in [7]. This simplification helps us to focus on the interaction between CS and inter-data packet interference.

State Representation of CS and Interference

Although CS and interference states can be continuous, for the sake of simplicity, we consider a binary state for both CS and interference. If S_1 senses S_2's transmission, the CS state of L_1 by L_2 is expressed as C_1:1 and otherwise, C_1:0. Similarly, if L_1 suffers interference from L_2 (or S_2), the interference state of L_1 by L_2 is F_1:1 and otherwise, F_1:0. Thus, the CS and interference relation of a link to the other link is defined as one of four states: 11, 10, 01, or 00. In the example in Fig. 1, L_1's state by L_2 is C_1:1/F_1:1, and the state of L_2 by L_1 is C_2:0/F_2:0. We use a subscript X for a *don't care* state; X can be either 1 or 0.

CS and Interference with Saturated Traffic

To help in understanding, we distinguish between the throughput at a transmitter (TX) and the goodput at a receiver (RX). Throughput is defined as the number of bits-per-second transmitted at the transmitter's application; whereas, goodput is defined as the number of bits-per-second received at the receiver's application. In the presence of interference, RX goodput may be smaller than TX throughput.

In this section, to understand the effect of CS and interference on the capacity of two links, we first consider the case-by-case, link throughput/goodput when both senders always have packets to send (i.e., traffic is saturated). Because each of the CS and interference relations between two links can have one of four states (11, 10, 01, or 00) as shown in the previous section, the number of combinations is 16 for the CS and interference states of the two links.

We first exclude the four cases when two senders sense each other (C_1C_2:11, F_1F_2:XX) because our measurement shows that there is little interference in these cases. Both links have a fair share of broadcast/unicast throughput/goodput.

In general, interference is determined by the SIR relation between the two senders and one receiver. Thus, estimating interference, based solely on the existence of a communication link, which is popular in the literature, is not correct, and hence, recent research tends to use the SIR-based interference models.

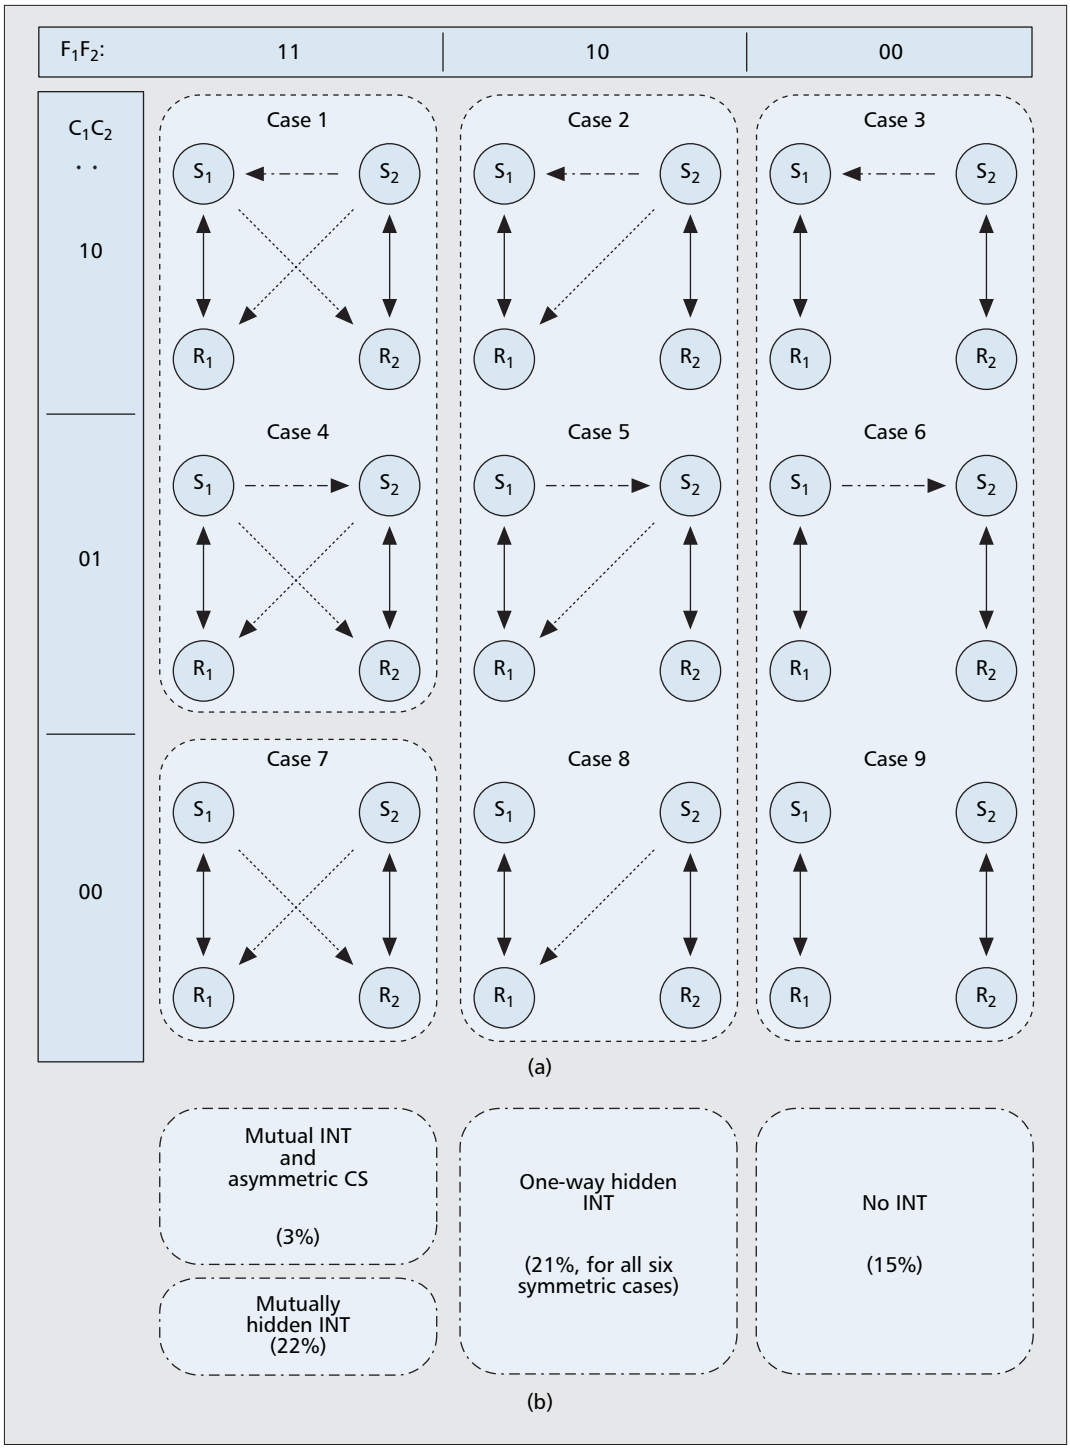

Figure 2. *Nine distinct cases and four groups. In the group names, INT denotes interference.*

After excluding the above four cases, we have 12 remaining cases. The cases where only L_1 is interfered and the cases where only L_2 is interfered result in symmetric throughput/goodput scenarios. We thus ignore one of two symmetric interference states and exclude three cases when only L_2 suffers interference (F_1F_2:01). We consider three cases where only L_1 suffers interference (F_1F_2:10).

We now have nine remaining cases as shown in Fig. 2a. The row index represents the CS relation, and the column index represents the interference relation. Although we see two pairs of symmetric cases (cases 1, 4 and cases 3, 6), we do not remove them for the sake of presentation. Based on the distinctive performance characteristics of the nine cases, we classify them into four groups as shown in Fig. 2b.

Testbed Experiments

We embodied nine network topologies corresponding to Fig. 2a in our testbed, which consists of small-form factor, single-board computers with mini-PCI 802.11a cards using the Atheros

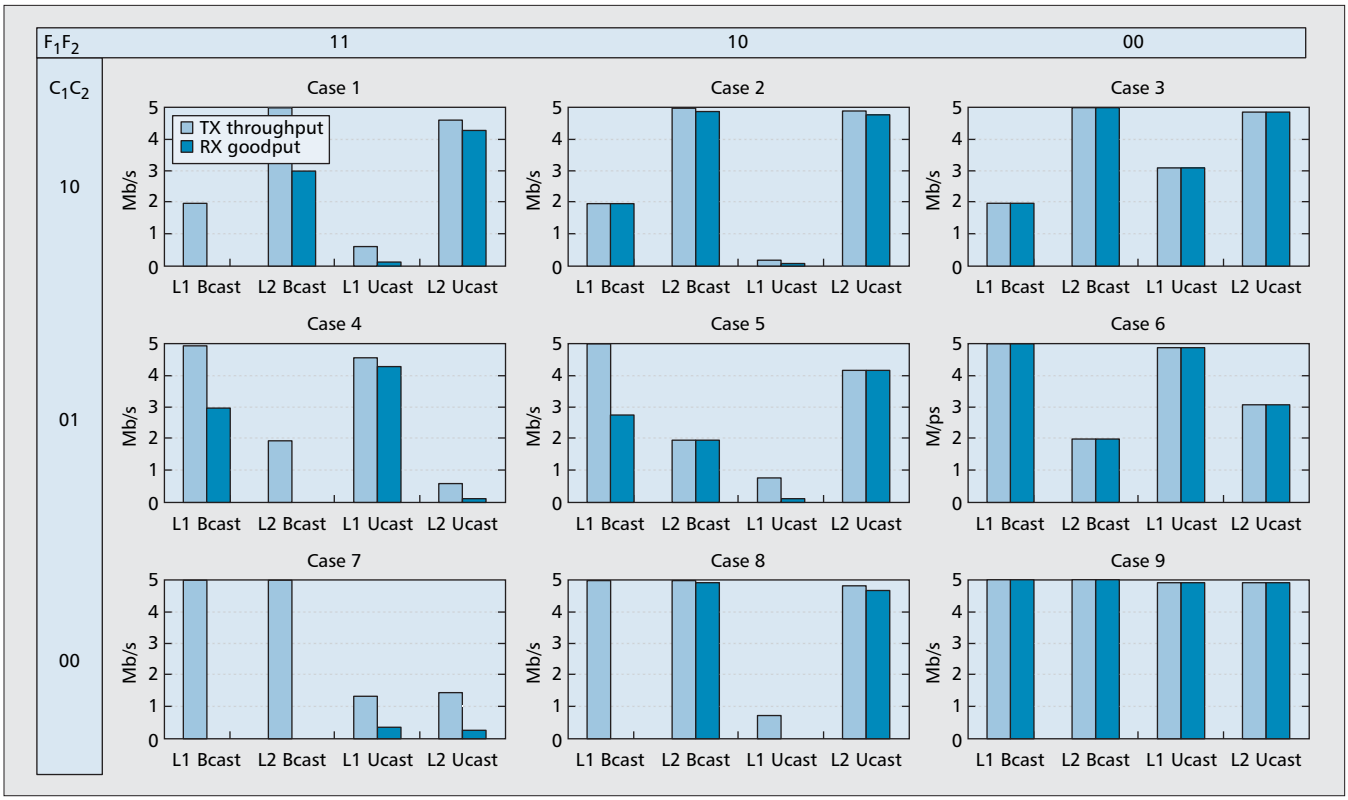

Figure 3. *Link throughput/goodput for nine cases in two saturated links.*

chipset. To embody each topology case with different CS and interference relations, we controlled the node placement, transmission power, and transmission/receiving antennas. The applications on S_1 and S_2 generate User Datagram Protocol (UDP) packets of 1000 bytes and continuously send them to R_1 and R_2, respectively, to make their output queue always backlogged (saturated). Each broadcast/unicast transmission period was 30 seconds. During the entire experiment, the PHY rate is fixed at 6 Mb/s, which is the lowest and the most robust bit rate in IEEE 802.11a. We used a clear 802.11a channel to avoid interference from other networks.

We first measure the maximum achievable throughput/goodput of a link, which is used as a reference value. The maximum TX throughput (sending-traffic rate at a transmitter) and RX goodput (traffic rate successfully delivered to a receiver) of each link are measured by activating one link while deactivating the other link. The measured maximum throughput/goodput of broadcast/unicast is about 5 Mb/s in our experiments. Because of the PHY/MAC overhead, the maximum throughput/goodput is approximately 5 Mb/s, although the PHY bit rate is 6 Mb/s. We designate this 5 Mb/s, *interference-free* throughput/goodput as the *channel capacity* of the 6 Mb/s PHY rate.

For each nine topology cases, we measured the throughput and goodput of UDP broadcast and unicast of L_1 and L_2 when both links are simultaneously active and the input traffic is always backlogged. Before testing the simultaneous transmissions, we verified that the interference-free throughput and goodput of each link reaches the channel capacity. Figure 3 shows the throughput/goodput of each link for the nine cases with saturated traffic. We now analyze the unique performance characteristics of the four groups classified in Fig. 2b, based on the observations from Fig. 3.

No Interference

Because both links in this group (cases 3, 6, and 9) are free from interference, goodput is the same as throughput for every case. Analyzing this group helps us understand the effects of (especially, asymmetric) CS on TX throughput.

Case 9 shows the expected results: when two links operate independently without sensing each other, both links fully utilize the channel capacity.

When the CS relation is asymmetric, the TX throughput of the sender, which does not sense the other sender, is the channel capacity as expected (L_2 in case 3 and L_1 in case 6). The throughput of the CS sender, however, differs from the expectation that it is half of the channel capacity. For broadcast, TX throughput is 2 Mb/s (smaller than 2.5) while it is 3.1 Mb/s for unicast. For broadcast, we identify the extended inter-frame spacing (EIFS) of 802.11 as the cause of smaller throughput. In 802.11 MAC, a node that received a packet that it cannot decode may go into the EIFS mode and waits until either receiving an error-free frame or the expiration of the EIFS time, which is longer than the double of distributed inter-frame-spacing (DIFS) time in 802.11.

For example, consider case 6. Suppose the CS sender S_2 attempts to transmit broadcast data and chooses a back-off counter that is smaller than S_1's, as illustrated by "S_2 wins" in Fig. 4a. Because S_1 does not sense S_2's transmission, S_1 also starts its transmission. As S_1's transmission starts later then S_2's, S_2 receives only the latter

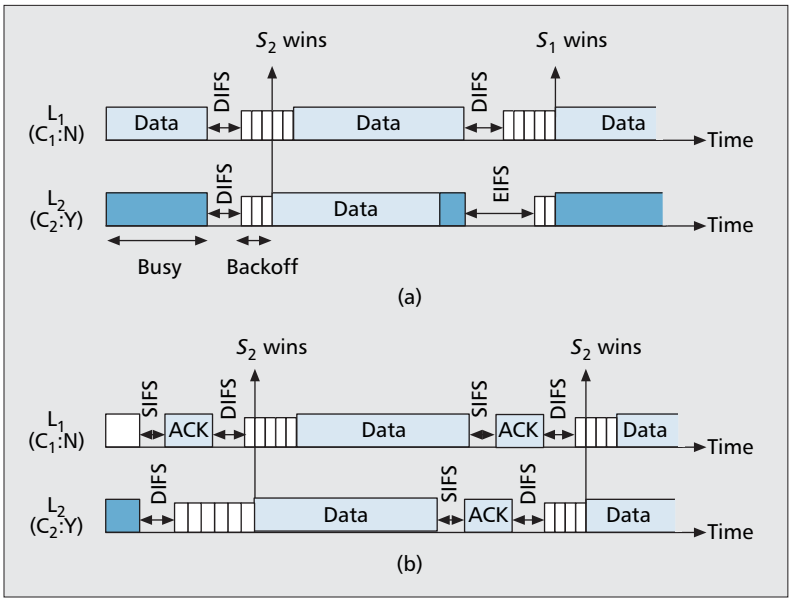

■ **Figure 4.** *Channel access scenario of asymmetric CS links: a) broadcast; b) unicast. The dark blue area indicates CS busy state sensed by* S_2.

part of S_1's transmission right after finishing its own transmission. Because S_2 missed the preamble and PHY header of S_1's packet, S_2 is unable to decode S_1's packet and hence goes into the EIFS mode. While S_2 waits for EIFS duration, S_1 begins its back-off counter earlier than S_2, which increases the possibility that S_1 will win the channel contention and start the next transmission before S_2's back-off counter expires ("S_1 wins" in Fig. 4a). In this manner, S_2, the sender that carrier senses, has transmission probability of less than 50 percent after its previous transmission. According to the 802.11 MAC standard, EIFS occurs only when the PHY indicates to the MAC that an 802.11 frame transmission has begun but did not result in the correct reception of a complete MAC frame. In Fig. 4a, where S_2 cannot recognize the beginning of S_1's frame transmission which occurs while S_2 is transmitting, based on the standard, EIFS should not occur. Our back-off time measurement and analysis, however, prove that EIFS occurs in the case of Fig. 4a with our Atheros 802.11 cards. We believe the card can recognize a 802.11 frame transmission with its preamble missed, although the card cannot decode the frame correctly.

When unicast traffic is examined, we observe a different scenario, as illustrated in Fig. 4b. In this case, the CS sender, S_2, receives an ACK from its receiver R_2 after its data transmission; thus, S_2 is less likely to sense the latter part of S_1's transmission and go into the EIFS mode than the broadcast scenario. Furthermore, because S_2 does not sense R_1's ACK transmission, S_2 will start a decrement of its back-off counter while S_1 is receiving an ACK from R_1. Thus, S_2, with unicast transmission, has a transmission probability greater than 50 percent not only after its previous transmission (second "S_2 wins" in Fig. 4b), but also after yielding the channel by sensing S_1's data transmission (first "S_2 wins" after busy state).

We confirmed the above arguments by observing the transmission order and inter-transmission time of S_1 and S_2 by using a 802.11 packet sniffer.

One-Way Hidden Interference

In this group of cases 2, 5, and 8, where the interference relation is asymmetric and at least one of two senders cannot sense the other, the goodput exhibits severe unfairness between the two links. In particular, with unicast transmission, the link whose interference state is F:1 has nearly zero goodput, while the other link's goodput is close to the channel capacity. We can learn the following:

The Hidden Node Problem Occurs No Matter Which Sender is Hidden — We observe that a link is interfered when at least one of the senders of the link pair is hidden; either the sender of the link or the interfering sender of the other link (i.e., C_1C_2 is 10, 01, or 00). The term hidden interference commonly refers to only the case, C_1C_2:00. However, two more cases should be addressed:

- The sender is hidden from the interferer (e.g., in case 2, S_1 is hidden from its interferer S_2, and L_1 has almost zero unicast goodput).
- The interferer is hidden from the sender (e.g., the interferer S_2 in case 5 is hidden, and L_1 has almost zero unicast goodput).

The Winner Takes All — In this group, the interference-free (F:0) link takes most, if not all, of the channel capacity with unicast traffic. This is evident in case 5, where the broadcast goodput result shows a nearly fair share between the two links; whereas, the unicast goodput share is extremely unfair. This is due to the exponential back-off mechanism, which forms a vicious cycle in case 5 as follows:

1. S_1 of the interfered link L_1 increases its contention window when it does not receive an ACK from R_1.
2. S_1's packet transmission rate decreases.
3. S_2 has more chances to transmit while S_1 refrains its transmission due to the increased contention window. Hence, S_2's TX throughput increases.
4. As the interferer S_2's TX throughput increases, the packet-drop rate of L_1 increases, and S_1's contention window is doubled again.

As the above cycle repeats, L_1's unicast RX goodput drastically drops and reaches near zero, while L_2 takes most of the channel capacity. Note that this phenomenon in unicast contrasts with that of broadcast. Because there is no explicit ACK mechanism for broadcast packets, the sender does not know whether its broadcast packet collided, and no backoff occurs. That is why the broadcast TX throughput of the interfered links in this group does not decrease despite packet collisions resulting from interference.

Mutual Interference and Asymmetric CS

In this group of cases 1 and 4, we find that if both links suffer from interference and the CS relation is asymmetric, the link that does not sense takes most of the channel capacity. The

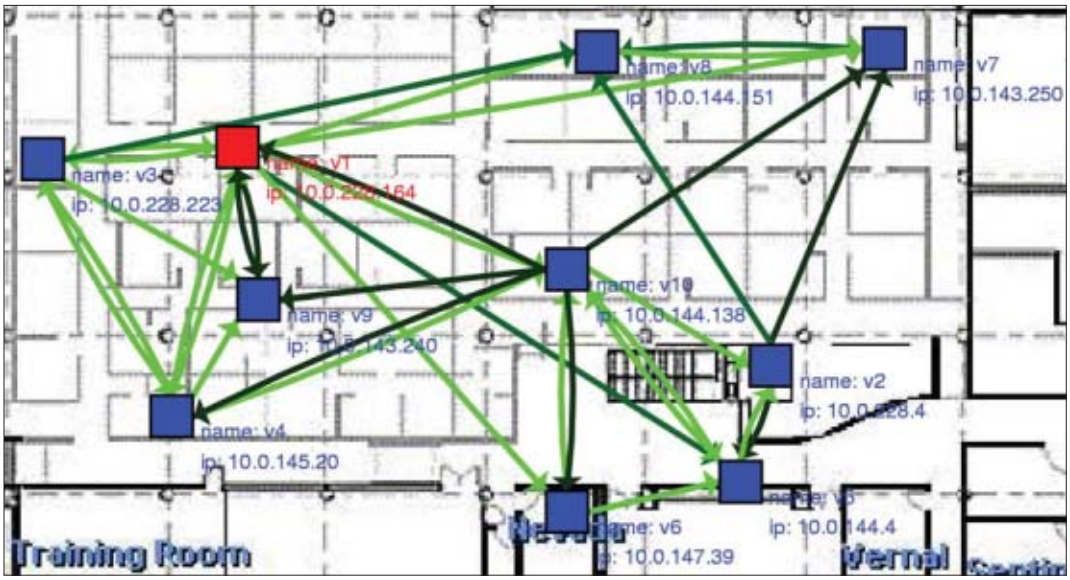

■ **Figure 5.** *Snapshot of a testbed deployment.*

In 802.11 MAC, a node that received a packet that it cannot decode may go into the EIFS mode and waits until either receiving an error-free frame or the expiration of the EIFS time, which is longer than the double of distributed inter-frame-spacing (DIFS) time in 802.11.

sender sensing the other sender's transmissions, for example S_1 in case 1, begins to yield the channel which in turn, decreases the other link's (L_2) collision probability at R_2. However, L_1's collision probability at R_1 remains high as L_2 does not sense L_1 and hence, does not yield the channel. Thus, S_1's contention window increases faster than S_2's. Therefore, L_2 takes most of the channel capacity, and L_1's goodput reaches almost zero.

In this group, the aggregate broadcast goodput (3 Mb/s) is much less than the aggregate unicast goodput (4.4 Mb/s). In broadcast, the CS sender (e.g., S_1 in case 1) is not aware of the interference due to the lack of ACK feedback and does not back off or decrease its transmission rate. Thus, both links suffer from interference and broadcast goodput degradation.

Mutually Hidden Interference

In this group of case 7, both links interfere and are hidden from each other, which results in the poor performance of both links. The goodput degradation becomes intensified because of saturated traffic from both senders. This problematic topology scenario commonly exists; 22 percent of link pairs in our 10-node testbed fall into this category. Note that in this case, the payload size, as well as the PHY bit rate can yield a somewhat different throughput/goodput as reported in [8]. We found that our measurement results follow the model in [8]. For example, the performance degradation becomes intensified as the larger payload size is used.

When we observe the aggregate goodput with unicast for all four groups, there are important findings. The aggregate unicast goodput reaches almost the channel capacity in all groups except this *mutually hidden interference (INT)* group. In the other groups, if any one of two senders senses the other, the aggregate unicast goodput reaches almost the channel capacity regardless of the interference relation. This is true even in cases when both links interfere with each other, as shown in cases 1 and 4. In the example of case 1, due to the back-off mechanism, S_1 decreases its transmission rate, which helps the other to survive interference.

Occurrence Frequency

We measured how frequently each of four groups occurs in our 10-node testbed deployed at HP Labs. Figure 5 shows the layout of the deployed testbed from our topology management program. Note that in this experiment, we did not control the antenna/node positions and transmission power to observe the natural topological characteristics. We chose a pair of links where each link has greater than 2.5 Mb/s-interference-free goodput. We found a total of 152 such pairs. In Fig. 5, we can observe that many links are asymmetric in the natural deployment. Figure 2b shows the occurrence percentage of each group from the 152 pairs in our testbed. Note that here we examine link pairs and not node pairs. The occurrence frequency of the *one-way hidden-interference* group contains all the six cases including the three excluded symmetric cases. As for the cases when both senders sense each other (which are omitted in Fig. 2), the occurrence frequency is 39 percent.

Summary

The major findings from our measurement study are listed as follows. First, when the CS relation is asymmetric, the CS sender has more TX throughput with unicast transmission than with broadcast transmission due to 802.11 MAC mechanisms (EIFS and unicast ACK). Second, hidden interference should include two asymmetric cases: when the interferer is hidden from the sender and when the sender is hidden from the interferer. Third, in the *one-way hidden-interference* group and the *mutual-interference and asymmetric-CS* group, the unicast goodput share is extremely unfair. In our testbed results, about 24 percent of link pairs suffer from this extreme unfairness. Fourth, in the mutually hidden INT group, which occurs frequently (22 percent), both links suffer from poor performance.

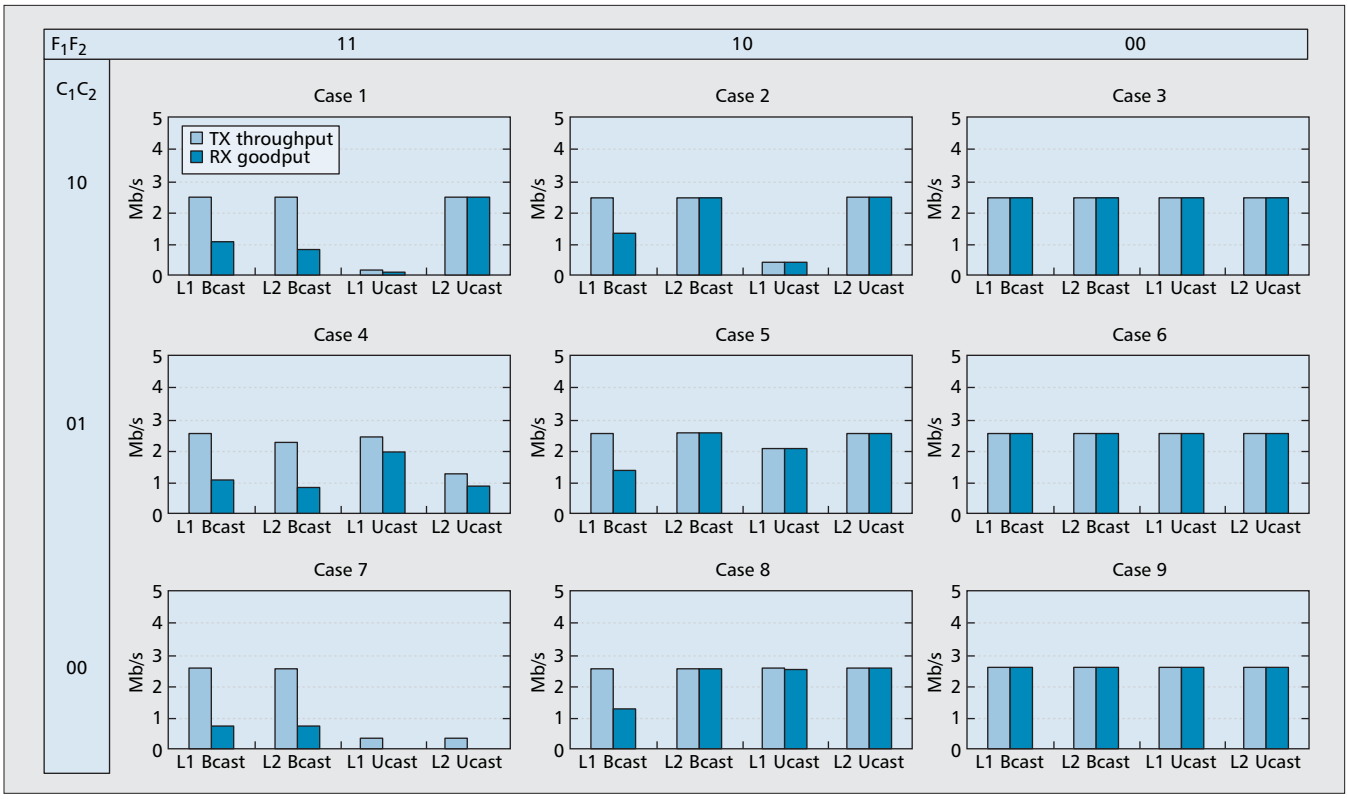

■ **Figure 6.** *Link throughput/goodput for nine cases with 2.5 Mb/s traffic.*

CS and Interference with Rate-Controlled Traffic

We now investigate the link behavior with rate-controlled traffic. A sender's application program periodically transmits packets by the given traffic rate. The UDP transmission rate (TX throughput) cannot exceed the given rate in any of the cases. We observed that the broadcast/unicast throughput/goodput results for the 1-Mb/s traffic-rate range is between 0.8 and 1 Mb/s in all cases. Two links have a fair share of goodput regardless of the relation between interference and CS when the traffic load is light (1 Mb/s).

We present the 2.5-Mb/s traffic-rate results in Fig. 6. Below are the findings from this experiment.

The links that had low goodput in the saturated channel scenario have greater goodput (up to 2.5 Mb/s) when we test with the rate-controlled traffic. Furthermore, the aggregate unicast goodput with rate-controlled traffic becomes almost as high as the aggregate goodput from the saturated scenarios in some cases; compare cases 5 and 8 in Figs. 3 and 6. This observation gives an important clue for traffic engineering and network management. As most off-the-shelf wireless LAN and mesh devices do not support time-division multiple-access (TDMA)-like link scheduling, we cannot avoid link interference in the CSMA MAC systems even if the link interference information is available. Mesh backbone networks are usually stationary and operated/managed by a single operator. Thus, after we obtain the CS and interference information (by measurement or prediction), we can control the traffic intensity on each link by selecting the appropriate routing paths and shaping the traffic at the mesh nodes. By doing so, although it is not possible to completely avoid link interference and packet collisions in CSMA systems, we can control the effect of link interference and eventually utilize the network capacity more efficiently, as the given CS and interference relations allow.

Packet-drop rate is significantly reduced with rate-controlled traffic. In case 1 in Fig. 6, L_1 suffers (hidden) interference from L_2, and its unicast TX throughput is smaller than broadcast TX throughput. In unicast mode, how frequently S_1 performs backoff depends on L_2's traffic rate. However, note that its unicast RX goodput is equal to the unicast TX throughput; whereas, broadcast RX goodput is less than half of the broadcast TX throughput. With rate-controlled interfering traffic, we observe that the transmission failure rate is not so high as to cause packet drops. By default in 802.11 MAC, the failure of four consecutive retransmissions indicates a packet drop. Because the traffic rate of the interfering L_2 is controlled, L_2 does not exploit all the channel-idle time during L_1's backoff. This allows L_1 to use some portion of the channel.

This observation of reduced packet-drop rate indicates another implication of traffic engineering in multihop wireless networks. Multihop routing protocols consider a wireless link to be broken when consecutive packets drop and use an alternate path or attempt to rediscover a route even though the link is active and could be suffering from only temporary interference. This re-routing instability [9] causes drastic perfor-

mance degradation. Thus, rate control can alleviate the MAC-level fairness problem and also mitigate upper-layer performance variation in wireless mesh networks.

Conclusion

Understanding the link behavior in the presence of interference is the key to achieving high-capacity wireless mesh networks. We address this important task by studying the interaction of interference and CS. We experimentally characterize the different couplings of CS and interference relations that are possible between two 802.11 links. We investigate the impact of these relations on the capacity of the link pairs, with both saturated and rate-controlled traffic. We show that the asymmetric CS and interference relations, together with the traffic intensity, determine the link capacity and the channel access fairness.

We value this experimental study as a first step toward a more comprehensive understanding of interference in wireless mesh networks. First, we are now studying a measurement-based approach to estimate the topology of a given (deployed) wireless network — CS relation, interference relation, and link performance. The RSS-based prediction (RBP) methodology [10] shows that we can estimate link throughput/goodput by using the measurement data of this article. Second, although we study only one-hop links in this article, we can apply our scheme to multihop routes. In fact, because our scheme identifies hidden interferers and their impact on link capacity, we can estimate the self interference (or intra-flow interference) of multihop routes by evaluating the relation between each pair of non-adjacent links on the route (e.g., first and third links of the route, first and fourth links, second and fourth links, etc.). Extending our work to multiple interferers is not as straightforward. A simple solution is to generate additional CS/interference relations that consider multiple links. This approach, however, will result in a large number of relation cases, and generating each case topology is a difficult task. We believe a measurement-based modeling is a more efficient and effective approach for multi-link interference. Applying the major findings of this study, such as the impact of asymmetric CS relations on hidden interference and link capacity, into multi-link interference modeling is part of our ongoing work.

References

[1] K. Ramachandran *et al.*, "Interference-Aware Channel Assignment in Multi-Radio Wireless Mesh Networks," *Proc. IEEE INFOCOM*, Barcelona, Spain, Apr. 2006.

[2] Tropos Networks; http://www.tropos.com/

[3] K. Xu, M. Gerla, and S. Bae, "How Effective is the IEEE 802.11 RTS/CTS Handshake in Ad Hoc Networks?" *Proc. IEEE GLOBECOM*, Taipei, Nov. 2002.

[4] J. Bicket *et al.*, "Architecture and Evaluation of an Unplanned 802.11b Mesh Network," *Proc. ACM MobiCom*, Cologne, Germany, Aug. 2005.

[5] M. Garetto, J. Shi, and E. W. Knightly, "Modeling Media Access in Embedded Two-Flow Topologies of Multi-hop Wireless Networks," *Proc. ACM MobiCom*, Cologne, Germany, Aug. 2005.

[6] A. Kochut *et al.*, "Sniffing Out the Correct Physical Layer Capture Model in 802.11b," *Proc. IEEE ICNP*, Berlin, Germany, Oct. 2004.

[7] J. Padhye *et al.*, "Estimation of Link Interference in Static Multi-hop Wireless Networks," *Proc. Internet Measurement Conf.*, Berkeley, CA, Oct. 2005.

[8] H. Chang and V. Misra, "802.11 Link Interference: A Simple Model and a Performance Enhancement," *Proc. NETWORKING*, Waterloo, Canada, May 2005.

[9] P. C. Ng and S. C. Liew, "Re-Routing Instability in IEEE 802.11 Multi-Hop Ad-Hoc Networks," *Proc. IEEE LCN*, Tampa, FL, Nov. 2004.

[10] J. Lee *et al.*, "RSS-Based Carrier Sensing and Interference Estimation in 802.11 Wireless Networks," *Proc. IEEE SECON*, San Diego, CA, 2007.

Biographies

JEONGKEUN LEE (jklee@hp.com) received his B.S. and Ph.D. degrees from the School of Computer Science and Engineering, Seoul National University, in 2001 and 2007, respectively. During his graduate program, he was a visiting student at the University of California at Santa Cruz, Fraunhofer FOKUS in Germany, and HP Labs in the United States. He started his post-doctoral research at HP Labs in October 2007. His research focuses on mobile wireless networking systems, specifically mesh networks and localization/tracking of mobile nodes.

SUNG-JU LEE [SM] (sjlee@hp.com) received his Ph.D. degree in computer science from the University of California, Los Angeles, in 2000. Since completing his Ph.D. research, he has been with Hewlett-Packard Laboratories, working as a research scientist for the Mobile and Media Systems Laboratory. His recent research interests include wireless mesh networks, network management, and network interference. He has published more than 60 papers in the field of computer networks and distributed systems. He is a member of ACM.

WONHO KIM (wonhokim@princeton.edu) received his B.S. degree in computer science and engineering from Seoul National University in 2006. In fall 2007 he began his Ph.D. studies at Princeton University, New Jersey. Since February 2006 he has been with the Multimedia & Mobile Communications Laboratory at Seoul National University as an intern researcher. His research interests include wireless and mobile networking and distributed systems. In 2003 he joined KEBT and worked on development of smart-card operating systems.

DAEHYUNG JO (daehyung.jo@samsung.com) received his B.S. and M.S. degrees in computer science and engineering from Seoul National University in 2006 and 2008, respectively. He has been working in the Digital Media and Communication R&D center of Samsung Electronics since 2008. His research intrests include mesh and ad hoc networks, node tracking, and wireless/wired network connectivity.

TAEKYOUNG KWON (tkkwon@snu.ac.kr) obtained his B.S., M.S., and Ph.D. degrees from the Department of Computer Engineering, Seoul National University, in 1993, 1995, and 2000, respectively. He has been an assistant professor in the School of Computer Science and Engineering at the same university since 2004. Before that, he was a post-doctoral research associate at the University of California at Los Angeles and City University of New York. During his graduate studies he was a visiting student at IBM T. J. Watson Research Center and the University of North Texas. His research interest lies in sensor networks, wireless networks, IP mobility, and ubiquitous computing.

YANGHEE CHOI (yhchoi@snu.ac.kr) received a B.S. from Seoul National University, an M.S. from the Korea Advanced Institute of Science, and a D.Eng. in computer science from the École Nationale Superieure des Telecommunications Paris, in 1975, 1977, and 1984, respectively. Since 1991 he has been a professor at the School of Computer Science and Engineering, Seoul National University. He was with the Electronics and Telecommunications Research Institute from 1977 to 1979 and spent 1984 to 1991 as the director of the Data Communication Laboratory and the Protocol Engineering Center. He also worked at the Centre National D'Étude des Telecommunications, France, from 1981 to 1984, and at the IBM Thomas J. Watson Research Center from 1988 to1989 as a visiting scientist. He is the director of the Future Internet Forum of Korea and president-elect of the Korea Institute of Information Scientists and Engineers.

We believe a measurement-based modeling is a more efficient and effective approach for multi-link interference. Applying the major findings of this study into multi-link interference modeling is part of our ongoing work.

■ FIVE DAYS ■ FOUR CONFERENCES ■ ONE EXHIBITION

EUMW2009 PUTS CUTTING EDGE TECHNOLOGY ON A GLOBAL STAGE

EuMW 2009

EUROPEAN MICROWAVE WEEK

Rome

28 September - 2 October 2009

www.eumweek.com

The historical, cultural and stylish city of Rome will be the backdrop for the 12th European Microwave Week when it plays host to Europe's premier Microwave, RF, Wireless Technology and Radar event for the very first time. The vibrant Italian capital complements the focussed and challenging Week, which covers FIVE days, encompassing FOUR cutting edge conferences and ONE dynamic trade and technology exhibition featuring leading players from across the globe

THE EXHIBITION

29 September - 1 October 2009

• 7,500 sqm of gross exhibition space •
• 5,000 key visitors from around the globe •
• 1,700 - 2,000 conference delegates •
• In excess of 250 exhibitors •

VISIT & SEE

• International Companies - meet the industry's biggest names and network on a global scale
• Cutting-edge Technology - exhibitors showcase the latest product innovations, offer hands-on demonstrations and provide the opportunity to talk technical with the experts
• Technical Workshops - get first hand technical advice and guidance from some of the industry's leading innovators
• Four Conferences
- European Microwave Integrated Circuits Conference (EuMIC)
- European Microwave Conference (EuMC)
- European Wireless Technology Conference (EuWiT)
- European Radar Conference (EuRAD)

SIGN UP!

To apply to exhibit, register for FREE as a visitor to the exhibition or to discover prices related to conference attendance, log onto www.eumweek.com

Co-sponsored by: MTT-S, GAAS

Supported by: IET

Endorsed by: IEEE, AES, Electron Devices Society, AES

Organised by: horizon house, EuMA European Microwave Association

Official Publication: Microwave Journal

39th European Microwave Conference 2009

The 39th European Microwave Conference

EuWiT 2009 — The 2nd European Wireless Technology Conference

EuRAD 2009 — The 6th European Radar Conference

EuMIC 2009 — The 4th European Microwave Integrated Circuits Conference

CALL FOR PAPERS AND PROPOSALS

IEEE International Conference on Communications (ICC 2010)

will be held in Cape Town, South Africa, from 23-27 May 2010, prior to the Soccer World Cup also being held in South Africa. The conference is aimed at addressing key themes on "Communications: Accelerating Growth and Development." The program will feature major Symposia, Tutorials, Panel Discussions, and Workshops. Full details of submission procedures are available on the IEEE ICC 2010 website, www.ieee-icc.org/2010/.

The organizers of IEEE ICC 2010 as well as our attendees expect accepted papers to be presented at the conference. IEEE reserves the right to exclude a paper from distribution after the conference (e.g., removal from IEEE Xplore) if the paper is not presented at the conference.

EXECUTIVE COMMITTEE

Minister Naledi Pandor (General Chair) (Department of S&T, South Africa)

Khaled B. Letaief (Vice General Chair) (eekhaled@ust.edu.hk)

Reuben September (Vice General Chair) (Telkom, South Africa)

Fambirai Takawira (Executive Co-Chair) (ftakaw@ukzn.ac.za)

Wilhelm Leuschner (Executive Co-Chair) (Leuschner@up.ac.za)

Jan C. Olivier (TPC Co-Chair) (corne.olivier@eng.up.ac.za)

Chengshan Xiao (TPC Co-chair) (xiaoc@mst.edu)

Sherman Shen (Symposia Chair) (xshen@bbcr.uwaterloo.ca)

Jerry D. Gibson (Panel Co-Chair) (gibson@ece.ucsb.edu)

Dumisa W. Ngwenya (Panel Co-Chair) (dwngwenya@icasa.org.za)

Abbas Jamalipour (Workshop Co-Chair) (a.jamalipour@ieee.org)

Andreas Molisch (Workshop Co-Chair) (molisch@usc.edu)

Geoffrey Li (Tutorial Co-Chair) (liye@ece.gatech.edu)

Mischa Dohler (Tutorial Co-Chair) (mischa.dohler@cttc.es)

Ross Murch (Keynote Chair) (eermurch@ee.ust.hk)

Ted Rappaport (GITC Advisor) (wireless@mail.utexas.edu)

Merrily Hartmann (GIMS Advisor)

Roger Peplow (Finance Chair) (roger@ukzn.ac.za)

Sunil Maharaj (Conf. Operations Chair) (sunil.maharaj@up.ac.za)

Marius Mostert (Business Chair) (mostert@telkom.co.za)

Steve Alves (Exhibition Chair) (steve_alves@concilium.co.za)

Viv Crone (Patronage Chair) (vcrone@za.spescom.com)

Bea Lacquet (Publication Chair) (beatrys.lacquet@wits.ac.za)

Saurabh Sinha (Registration Chair) (saurabh.sinha@up.ac.za)

Jingxian Wu (Student Travel Grant Chair) wuj@uark.edu

PLANNED TECHNICAL SYMPOSIA

> Paper Submission: 4 Sep 2009

Authors are invited to submit original technical papers for presentation at symposia and publication in the conference proceedings from the following list: oriented information in a broad range of communication areas from cognitive radio networks, ultrawide bandwidth systems and automotive networking and telematics to cognitive and cooperative wireless networks and peer-to-peer communications.

General Symposium on Selected Areas in Communications

Tarik Taleb, talebtarik@ieee.org
Sumit Roy, sroy@u.washington.edu
Chunming Qiao, qiao@computer.org
Yacine Ghamri-Doudane, ghamri@ensiie.fr
Riccardo Raheli, raheli@unipr.it

Communication Theory Symposium

Daniela Tuninetti, danielat@uic.edu
Holger Boche, holger.boche@mk.tu-berlin.de
Guosen Yue, yueg@nec-labs.com

Signal Processing for Communications Symposium (SPS)

S. Boussakta, s.boussakta@ncl.ac.uk
Donglai Xu, d.xu@tees.ac.uk
T.J. Lim, limtj@comm.utoronto.ca

Wireless Communications Symposium

Yi Qian, yqian@nist.gov
Ranjan K. Mallik, rkmallik@ee.iitd.ernet.in
Sennur Ulukus, ulukus@umd.edu
Wei Chen, wchen@tsinghua.edu.cn

Wireless & Mobile Networking Symposium

Thomas Hou, thou@vt.edu
Randall Berry, rberry@ece.northwestern.edu
Hai Jiang, hai.jiang@ece.ualberta.ca
Farid Naït-Abdesselam, farid.nait-abdesselam@lifl.fr

Optical Networks and Systems Symposium

Mounir Hamdi, hamdi@cse.ust.hk
Andrea Bianco, andrea.bianco@polito.it
Srinivasan Ramasubramanian, srini@ece.arizona.edu

Next-Generation Networking and Internet Symposium

Wanjiun Liao, wjliao@cc.ee.ntu.edu.tw
Achille Pattavina, pattavina@elet.polimi.it
Hussein Mouftah, mouftah@site.uottawa.ca

Communication QoS, Reliability and Modeling Symposium

Nelson Fonseca, nfonseca@ic.unicamp.br
Charalabos Skianis, cskianis@aegean.gr
Toshinori Tsuboi, tsuboi@cs.teu.ac.jp

Ad Hoc, Sensor and Mesh Networking Symposium

Nirwan Ansari, Nirwan.Ansari@njit.edu
Walaa Hamouda, hamouda@ece.concordia.ca
Nei Kato, kato@it.ecei.tohoku.ac.jp
Hongchi Shi, hs15@txstate.edu

Multimedia Communications, Communication Software and Services Symposium

Mohammed Atiquzzaman, atiq@ou.edu
Zhu Li, zhu.li@ieee.org
Marcus Brunner, brunner@nw.neclab.eu

Communication and Information System Security

Wenjing Lou, wjlou@ece.wpi.edu
Pramode K Verma, pverma@ou.edu
Abderrahim Benslimane, abderrahim.benslimane@univ-avignon.fr

TUTORIALS

> Tutorial Proposal Due: 4 Sep 2009

Proposals are invited for half- or full-day tutorials in communications and networking topics. Proposals should be submitted to liye@ece.gatech.edu and mischa.dohler@cttc.es for review.

PANEL DISCUSSIONS

> Panel Proposal Due: 28 Aug 2009

Submissions are sought for panel discussions the latest technical and business issues in telecommunications topics. Proposals should be submitted to gibson@ece.ucsb.edu and dwngwenya@icasa.org.za for review.

WORKSHOPS

> Workshops Proposal Due: 30 June 2009

Submissions are sought for workshops on the latest technical and business issues in communications and networking topics. Proposals should be submitted to a.jamalipour@ieee.org and molisch@usc.edu for review.

Advertisers' Index

Company	Page
Cisco	Cover 2
Conec Corporation	11
ECOC 2009	49
Elsevier	Insert
European Microwave Week	110
GL Communications	19
ICC 2010	111
John Wiley & Sons, Ltd.	Cover 3
Omicron Lab	7
Samsung	Cover 4
STS-ICCT 2009	10
Telogy	1
The Main Works	3

ADVERTISING SALES OFFICES

Closing date for space reservation: 1st of the month prior to issue date

NATIONAL SALES OFFICE
Eric L. Levine
Advertising Sales Manager
IEEE Communications Magazine
3 Park Avenue, 17th Floor
New York, NY 10016
Tel: 212-705-8920
Fax: 212-705-8999
e-mail: e.levine@comsoc.org

SOUTHERN CALIFORNIA
Patrick Jagendorf
7202 S. Marina Pacifica Drive
Long Beach, CA 90803
Tel: 562-795-9134
Fax: 562-598-8242
patrick@roymcdonald.com

NORTHERN CALIFORNIA
George Roman
4779 Luna Ridge Court
Las Vegas, NV 89129
Tel: 702-515-7247
Fax: 702-515-7248
Cell: 702-280-1158
george@roymcdonald.com

SOUTHEAST
Scott Rickles
560 Jacaranda Court
Alpharetta, GA 30022
Tel: 770-664-4567
Fax: 770-740-1399
srickles@aol.com

EUROPE
Rachel DiSanto
Huson International Media
Cambridge House, Gogmore Lane
Chertsey, Surrey, KT16 9AP
ENGLAND
Tel: +44 1428608150
Fax: +44 1 1932564998
Email: r.disanto@husonmedia.com

Convergent Technologies

Don't Forget Your IEEE Member Discount

WILEY
www.wiley.com

Advanced Wireless Networks

Cognitive, Cooperative & Opportunistic 4G Technology, 2nd Edition

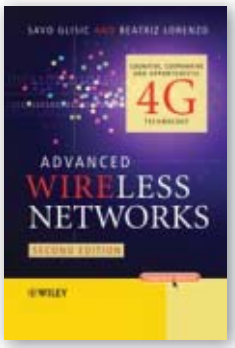

SAVO GLISIC and BEATRIZ LORENZO

Presents a comprehensive integral presentation of the main issues in 4G wireless networks, showing the wide scope and inter-relation between different elements of the networks. It provides systematic coverage of the latest global research results in the field of advanced wireless technology networks at a time when the industry is paving the way for 4G.

978-0-470-74250-1 • June 2009 • Hbk • 896pp
$170.00

Networking Fundamentals

Wide, Local and Personal Area Communications

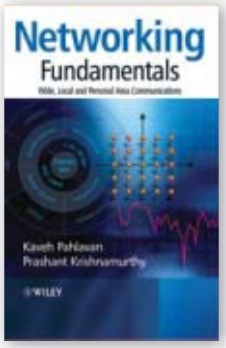

KAVEH PAHLAVAN and PRASHANT KRISHNAMURTHY

Provides essential information on networking technologies that are used in both wired and wireless networks designed for local area networks (LANs) and wide-area networks (WANs). Contains numerous illustrations, case studies and tables to supplement the text.

978-0-470-99289-0 • May 2009 • Hbk • 656pp
$180.00
978-0-470-99290-6 • May 2009 • Pbk • 656pp
$90.00

Femtocells

Opportunities and Challenges for Business and Technology

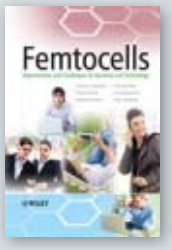

SIMON SAUNDERS; STUART CARLAW; ANDREA GIUSTINA; RAVI RAI BHAT; SRINIVASA RAO and RASA SIEGBERG

Outlines how low-power wireless access points can be used by mobile operators to provide high-speed wireless access to existing mobile handsets in home and offices, enhancing coverage and capacity and delivering new services, while maximising the benefits of licensed spectrum.

978-0-470-74816-9 • June 2009 • Pbk • 248pp
$69.95

Storage Networks Explained

Basics and Application of Fibre Channel SAN, NAS, iSCSI,InfiniBand and FCoE, 2nd Edition

ULF TROPPENS; RAINER ERKENS; WOLFGANG MUELLE-FRIEDT; RAINER WOLAFKA and NILS HAUSTEIN

covers the fundaments, techniques and functions of storage networks such as disk subsystems, Fibre Channel SAN, Internet SCSI (iSCSI), Fibre Channel over Ethernet (FCoE), Network Attached Storage (NAS), file systems, and storage virtualization.

978-0-470-74143-6 • July 2009 • Hbk • 568pp
$110.00

Mobile Peer to Peer

A Tutorial Guide

FRANK H. P. FITZEK and HASSAN CHARAF

This book provides the basic programming skills required to set up wireless communication links between mobile devices, offering a guide to the development process of mobile peer to peer networks.

Wiley series on Communications Networking & Distributed Systems

978-0-470-69992-8 • July 2009 • Hbk • 264pp
$50.00

Space-Time Layered Information Processing for Wireless Communications

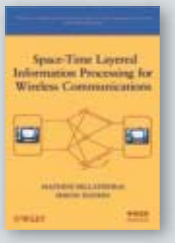

MATHINI SELLATHURAI and SIMON HAYKIN

This comprehensive book provides coverage of cutting-edge research in the field, with an emphasis on Multiple-Input Multiple-Output (MIMO) wireless communication systems. Examples illustrate theory wherever possible, while end-of-chapter problems challenge readers to reinforce their understanding of the subject matter.

978-0-471-20921-8 • April 2009 • Hbk 204pp
$95.95

Near-Capacity Multi-Functional MIMO Systems

Sphere-Packing, Iterative Detection and Cooperation

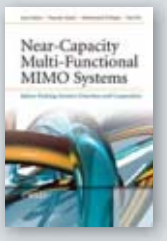

LAJOS HANZO; OSAMAH ALAMRI; MOHAMMED EL-HAJJAR and NAN WU

This book consolidates a large body of knowledge in one place, providing a unified treatment of differentially encoded space-time codes. Material is logically presented allowing advanced readers to turn directly to any relevant chapter of interest.

978-0-470-77965-1 • May 2009 • Hbk • 712pp
$200.00

A Guide to the Wireless Engineering Body of Knowledge (WEBOK)

IEEE Communications Society

Included in the WEBOK is a large list of additional publications and documents that candidates should review to become completely prepared for the WCET exam. The book is also a valuable reference for any wireless professional.

978-0-470-43366-9 • April 2009 • Pbk • 253pp
$69.95

HOW TO ORDER

Wiley books are available through your Bookseller. Alternatively send your order to:
John Wiley & Sons Inc., 111 River Street, Hoboken, NJ 07030, USA

For telephone orders:
Call 1-800-225-5945 (1-800-US-WILEY)
Fax: (212) 850-8888
E-mail: custserv@wiley.com

ORDER WILEY ONLINE...

- Visit the Wiley Electrical Engineering homepage: www.wiley.com/electrical
- Wiley Communications Technology Website: www.wiley.com/go/commstech/
- Amazon Wiley Communications Technology bookstore: www.amazon.com/professional/

Please note all prices correct at time of going to press but subject to change.

04553 04/09

Why did Samsung develop WiMAX?

We like speed.

Samsung has taken a leading role in WiMAX since the beginning. We were on the board of the WiMAX Forum and helped develop the iEEE 802.16e standard. We've been involved with everything from infrastructure to chip design and devices. We launched the first commercial Mobile WiMAX network in the world, then in the United States. Today, we're working hard to bring WiMAX technology to everyone. We're working fast, too.

For the latest information on Samsung and WiMAX, go to www.samsung.com/wss.

SAMSUNG

©2009 Samsung Telecommunications America, LLC ("Samsung"). All rights reserved.
All product and brand names are trademarks or registered trademarks of their respective companies.