

Information-centric networking – Ready for the real world?

Edited by

Ali Ghodsi¹, Börje Ohlman², Jörg Ott³, Ignacio Solis³, and
Matthias Wählisch⁵

- 1 University of California – Berkeley, US
- 2 Ericsson Research – Stockholm, SE, borje.ohlman@ericsson.com
- 3 Aalto University, FI
- 3 PARC – Palo Alto, US, Ignacio.Solis@parc.com
- 5 Freie Universität Berlin, DE, waehlich@ieee.org

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 12361 “Information-centric networking – Ready for the real world?”. The outcome of this seminar is based on individual talks, group work, and significant discussions among all participants. The topics range from application and performance aspects up to business, legal, and deployment questions. Even though significant progress is visible from the last Dagstuhl Seminar about ICN, there are still thrilling open research questions in all topic areas.

Seminar 02.–09. August, 2012 – www.dagstuhl.de/12361

1998 ACM Subject Classification C.2.1 Network Architecture and Design, C.2.5 Local and Wide-Area Networks—Internet

Keywords and phrases Information-centric, Network architecture, Application structure, Internet business models

Digital Object Identifier 10.4230/DagRep.2.9.1

1 Executive Summary


Ali Ghodsi

Börje Ohlman

Jörg Ott

Ignacio Solis

Matthias Wählisch

License  Creative Commons BY-NC-ND 3.0 Unported license
© Ali Ghodsi, Börje Ohlman, Jörg Ott, Ignacio Solis, and Matthias Wählisch

Information-centric networking (ICN) defines a communication paradigm that recognizes the dominant usage of the Internet as a substrate to disseminate and access content of all kinds: from traditional web pages to online social networks to file distribution to live and on-demand media feeds. With ICN, the focus shifts from the location at which a content object is stored (typically some server) to the object itself so that scale, efficiency, and robustness of content publication and retrieval can be improved beyond what current Content Distribution Networks (CDNs) can deliver.

Diverse instances of ICN networking architectures were developed, including CCN/NDN, NetInf, DONA, and LIPSIN, among others, and see experimentation at different scale in both academia and industry. The fundamental concepts of ICN have gained popularity in the



Except where otherwise noted, content of this report is licensed
under a Creative Commons BY-NC-ND 3.0 Unported license

Information-centric networking – Ready for the real world?, *Dagstuhl Reports*, Vol. 2, Issue 9, pp. 1–14
Editors: Ali Ghodsi, Börje Ohlman, and Ignacio Solis



Dagstuhl Reports

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

research community and have been taken up by several research activities that are addressing the topic from different angles.

Numerous research problems remain open, some of which (such as naming content) may find different (optimal) solutions in different deployments while others are more fundamental in nature and could affect the performance of all deployments. The latter include the performance benefits achievable through (cooperative) caching and caching at different points in the network, parallel content retrieval from multiple sources, and tradeoffs between native network layer and overlay-based ICNs. This second Dagstuhl Seminar on information-centric networking is intended to operate as a catalyst for these activities and provide a forum for discussing a selected subset of important research topics that have been identified so far. It will bring together researchers from different ICN backgrounds to discuss fundamentals that matter across the various platforms with the meta goal of identifying obstacles to be overcome, solutions, and paths towards real-world deployments.

In this seminar, we discussed the following core topics: (1) ICN applications and services, (2) ICN performance and comparison of alternative technologies, (3) business, legal, and deployment aspects.

2 Table of Contents

Executive Summary

<i>Ali Ghodsi, Börje Ohlman, Jörg Ott, Ignacio Solis, and Matthias Wählisch</i>	1
---	---

Overview of Talks

SAIL NetInf global connectivity routing and forwarding <i>Bengt Ahlgren</i>	5
ICN from a Service Provider Perspective <i>Marcus Brunner</i>	5
Testbeds and mounting large-scale demonstrations: Experiences with NDN <i>Patrick Crowley</i>	5
Authorization and ICNs – Beyond Open Content <i>Elwyn Brian Davies</i>	6
ICN Service Model Issues. Which are the ICN services, who provides them, and what does an ICN API look like? <i>Anders Eriksson</i>	6
ICN use cases and deployment (from device manufacturer perspective) <i>Myeong-Wuk Jang</i>	6
Deploying ICN at the network edge <i>Vikas Kawadia</i>	7
Resource management in ICN: Business relations and interconnection <i>Luca Muscariello</i>	7
Video Streaming In NDN Architectures <i>Ashok Narayanan</i>	7
Blackadder: Node Design for an Information-Centric Network Architecture <i>George Parisis</i>	8
Backscatter from the Data Plane – Threats to Stability and Security in Information-Centric Networking <i>Thomas C. Schmidt</i>	8
ICN evaluation: Why and How? <i>George Xylomenos</i>	9
What will be Inter-domain Policy in Content Centric Networks? <i>Eiko Yoneki</i>	9
Named Data Networking: Experimentation with the new architecture <i>Lixia Zhang</i>	10

Working Groups

Applications & API (1)	10
Applications & API (2)	11
Metrics and Evaluation (1)	11
Metrics and Evaluation (2)	12
ICN vs. DTN	12

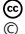

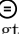
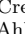
4 12361 – Information-centric networking – Ready for the real world?

Interdomain routing, forwarding	13
ICN Deployment	13
Participants	14

3 Overview of Talks

3.1 SAIL NetInf global connectivity routing and forwarding




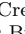
Bengt Ahlgren (SICS, SE)

License     Creative Commons BY-NC-ND 3.0 Unported license
© Bengt Ahlgren

The SAIL NetInf approach to global routing is a hybrid scheme using name-based routing assisted by routing hints. There are two levels of aggregation of routing information in order to make the system scalable, and to improve forwarding performance. First named data objects are grouped into aggregates using the authority part of the name. These aggregates are then associated with a set of routing hints through a lookup service. Requests for a named data object are forwarded using the highest priority hint in the set a router has a forwarding entry for leading towards a network location where the publisher makes the object available. The hint priority implements the second level of aggregation, as only the lowest priority hints need to be announced in the global routing system.

3.2 ICN from a Service Provider Perspective

Marcus Brunner (Swisscom AG – Bern, CH)

License     Creative Commons BY-NC-ND 3.0 Unported license
© Marcus Brunner

The slides are assessing the ICN approach from a service provider or operator perspective. This is done from a technical as well as a first try on the business aspects of an operator. Thought the system designs are different and not finally specified this is just a first approximation of some of the operator thinking behind the abstract approach of ICN. Many of the assumptions made are today's assumptions and might not be true anymore in the future depending on the future development of operator infrastructure.

3.3 Testbeds and mounting large-scale demonstrations: Experiences with NDN

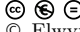
Patrick Crowley (Washington University, US)

License     Creative Commons BY-NC-ND 3.0 Unported license
© Patrick Crowley
Joint work of NDN team

In this talk, we will share both our motivations and methods for conducting large-scale demonstrations of the NDN architecture. We will also introduce open-source tools that we believe will enable others to mount similar demonstrations.

3.4 Authorization and ICNs – Beyond Open Content

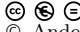
Elwyn Brian Davies (Trinity College Dublin, IE)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Elwyn Brian Davies

ICNs to date are primarily focussed on open access information. This presentation is intended to promote a discussion on applying authorization in order to restrict access to particular content that is published into an ICN.

3.5 ICN Service Model Issues. Which are the ICN services, who provides them, and what does an ICN API look like?

Anders Eriksson (Ericsson Research – Stockholm, SE)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Anders Eriksson

Joint work of Eriksson, Anders; Ohlman, Börje; Persson, Karl-Åke

Main reference A. Eriksson, B. Ohlman, K.-A. Persson, “What are the Services of an Information-centric Network, and Who Provides Them?”, in: Proc. of the Fourth Int’l Conf. on Advances in P2P Systems (AP2PS 2012), IARIA, 2012.

URL http://www.thinkmind.org/index.php?view=article&articleid=ap2ps_2012_1_20_30008

This presentation highlights the need to define ICN services in addition to a basic content retrieval service. The following services are proposed: advertise, publish, subscribe, event notification, and search. The event notification service can alert the search service when new data is advertised or published, so that the search service can index the data in a timely fashion. Finally, an example of an ICN API is described.

3.6 ICN use cases and deployment (from device manufacturer perspective)


Myeong-Wuk Jang (Samsung, KR)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Myeong-Wuk Jang

This talk includes ICN use cases and deployment issues from device manufacturer perspective. Although fundamental researches for ICN are very important, we should seriously consider CCN products that will be launched in the near future. The first product of ICN may not be networking devices, such as routers, but consumer devices, such as phones, computers, or TV. ICN should run over IP for a while. When we consider near future scenarios with ICN, there are still unresolved problems. These problems should be the target of our research and development.

3.7 Deploying ICN at the network edge

Vikas Kawadia (BBN Technologies – Cambridge MA, US)


License  Creative Commons BY-NC-ND 3.0 Unported license
© Vikas Kawadia

Joint work of Kawadia, Vikas; Hoon, Jeremy; Parkes, David C.

We discuss the issues in deploying ICN at the mobile edge of the network. We consider incentives for users to participate in an ICN, in particular saving broadband quota by obtaining content locally from other users. We consider a simplified case where multiple users are competing for wireless bandwidth. We design a mechanism for incentive compatible dynamic prioritization of user data on shared routers, such as 3G/4G devices. We adopt a recent idea proposed by Babaioff et al. (2010) in a multi-armed bandits context to create a mechanism that is truthful for buyers, meaning that users or user devices can bid straightforwardly, and design a revenue pooling method for incentive alignment that makes the scheme faithful for sellers, meaning that sellers maximize revenue by using prioritization algorithms that preserve buyer truthfulness. The mechanism works for constant per-byte buyer values, private to each buyer, a wide range of stochastic demand models, and is detail-free, in that the rules operate without knowledge of demand models or network conditions. We show simulation results demonstrating the efficiency gains from dynamic prioritization, as well as the effectiveness of revenue pooling.

3.8 Resource management in ICN: Business relations and interconnection


Luca Muscariello (Orange, FR)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Luca Muscariello

In the talk we review the content-delivery value chain and how new business models, heavily driven by video delivery, are changing the interconnection relations between the different entities in the Internet. The talk is based on the viewpoint of a large network operator but also considers ISPs in general and their potential strategies in the long term with respect to the content-delivery value chain.

3.9 Video Streaming In NDN Architectures

Ashok Narayanan (Cisco Systems – Lexington, US)


License  Creative Commons BY-NC-ND 3.0 Unported license
© Ashok Narayanan

This talk describes some of the issues seen with video streaming when implemented as an NDN application. It touches on the various service-provider video streaming schemes (MPEG/UTP, progressive/RTMP and HTTP/adaptive), and describes some of the high-level issues when trying to re-implement HTTP adaptive streaming as an NDN application. It goes in-depth into bandwidth estimation which is a key feature of adaptive video streaming but is complicated by ICN/NDN architectures. One of the general lessons to be learned is

that applications which try to adapt to "network conditions" require significant changes for NDN/ICN because the representation of network conditions is quite different.

3.10 Blackadder: Node Design for an Information-Centric Network Architecture

George Parisis (University of Cambridge, GB)

License  Creative Commons BY-NC-ND 3.0 Unported license
© George Parisis

Joint work of Parisis, George; Trossen, Dirk


Main reference D. Trossen, G. Parisis, "Designing and Realizing an Information-Centric Internet," in IEEE Communications Magazine, Volume 50, Issue 7, pp. 60–67, July 2012.

URL <http://dx.doi.org/10.1109/MCOM.2012.6231280>

Information-centric networking has been touted as an alternative to the current Internet architecture by several research groups. Our work addresses a crucial part of such a proposal, namely the design of a network node within an information-centric networking architecture. We describe the service model exposed to applications and other network nodes and we demonstrate a video streaming application that runs natively on top of a VPN network consisting of VMs across all European partners of the PURSUIT EU FP7 project.

3.11 Backscatter from the Data Plane – Threats to Stability and Security in Information-Centric Networking

Thomas C. Schmidt (HAW – Hamburg, DE)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Thomas C. Schmidt

Joint work of Schmidt, Thomas C.; Wählisch, Matthias; Vahlenkamp, Markus


Main reference M. Wählisch, T.C. Schmidt, M. Vahlenkamp, "Backscatter from the Data Plane – Threats to Stability and Security in Information-Centric Networking," arXiv:1205.4778v2 [cs.NI].

URL <http://arxiv.org/abs/1205.4778v2>

Information-centric networking proposals attract much attention in the ongoing search for a future communication paradigm of the Internet. Replacing the host-to-host connectivity by a data-oriented publish/subscribe service eases content distribution and authentication by concept, while eliminating threats from unwanted traffic at an end host as are common in today's Internet. However, current approaches to content routing heavily rely on data-driven protocol events and thereby introduce a strong coupling of the control to the data plane in the underlying routing infrastructure. In this paper, threats to the stability and security of the content distribution system are analyzed in theory and practical experiments. We derive relations between state resources and the performance of routers and demonstrate how this coupling can be misused in practice. We discuss new attack vectors present in its current state of development, as well as possibilities and limitations to mitigate them.

3.12 ICN evaluation: Why and How?


George Xylomenos (Athens University of Economics and Business, GR)

License  Creative Commons BY-NC-ND 3.0 Unported license
© George Xylomenos

ICN evaluation seems to be very tricky: common networking metrics seem insufficient, existing traffic and network models are likely inappropriate and there qualitative arguments both for and against ICN. The real question though is what are we trying to prove? If existing metrics seem wrong, maybe this is because we expect ICN to offer something other than improved throughput or reduced delay. Whatever that is, it must be something to do with data naming.

3.13 What will be Inter-domain Policy in Content Centric Networks?

Eiko Yoneki (University of Cambridge, UK)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Eiko Yoneki

Inter-domain routing with BGP is not shortest paths, which involves policies on routing. In Information Centric Networking (ICN), it requires a policy based routing protocol with a finer granularity of policies (i.e. content names level rather than hosts). DiBenedetto et al. [1] explored routing policies in ICN. However, we have not seen much research in Inter-domain policy in ICN.


In ICN, the policies could be driven by economic incentives by deployment of sharing cache among peers, rebating routing, and multi-payment in multi-path, which could be complex. Rajahalme et al worked on incentive-compatible caching and peering in [2]. Agyapong et al also described implication for protocol design and public policy in economic incentives [3]. An economic model of ICN, where various stake holders play a complex game based on their incentives is an essential issue in ICN research. For example, content cache placement could depend on the inter-domain routing policy. This problem is important and interesting since no longer can the policy be operated on top of network protocols (i.e. TCP/IP), and it needs to be embedded within ICN.

References

- 1 S. DiBenedetto, C. Papadopoulos, and D. Massey, *Routing policies in named data networking*, in Proc. of the ACM SIGCOMM workshop on Information-centric networking (ICN), 2011
- 2 J. Rajahalme, M. Sarela, P. Nikander, and Sasu Tarkoma, *Incentive-compatible caching and peering in data-oriented networks*, ReArch 2008.
- 3 P. Agyapong and M. Sirbu, *Economic Incentives in Content-Centric Networking: Implications for Protocol Design and Public Policy*, 39th Research Conference on Communication, Information and Internet Policy, 2011.

3.14 Named Data Networking: Experimentation with the new architecture

Lixia Zhang (Univ. California – Los Angeles, US)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Lixia Zhang

Joint work of the NDN project team

Main reference The Named Data Networking (NDN) Project, NDN Technical Report NDN-0001, October 2010.

URL <http://named-data.net/>

This talk reports on the progress of our NDN testbed development. Since its launch two years ago, the Named-Data Networking project has focused the research effort on the design, development, and actual usage of a variety of applications running over NDN networks. The NDN testbed is used not only to support application experimentations but also to drive research into network routing, monitoring, and management under the NDN architecture. Our experience suggests that architecture design research must include experimental components in order to verify whether, and how well, the proposed design can solve real problems. Our experience also shows that solving real problems not only forces architectural details to be filled in, but also validates and shapes the direction of the architectural development.

4 Working Groups

In this seminar, we organized two sessions for discussions in smaller groups. The first session covered two topics. Each topic has been analyzed by two individual groups. The second session focused on three different topics. In the following, we present the main results.

4.1 Applications & API (1)

Participants

Carsten Bormann, Antonio Carzaniga, Lars Eggert, Anders Eriksson, Xiaoming Fu, Volker Hilt, Pan Hui, Anders Lindgren, Ignacio Solis, George Xylomenos

Discussion and Open Questions

Starting from the key question what an API is, the group discussed where the API is located. In general, one may distinguish between a library API (running at the edge) and a core API (running everywhere). The group further discussed functions of a common API. This includes (a) advertise a name, (b) make data available, and (c) get data based on name. Requesting data might be split into `get` and `subscribe`. `Get` asks the network to retrieve a piece or a set of existing data based on a name. In contrast to this, `subscribe` asks the network to retrieve the data for a give amount of time.

Other topics that need consideration in the future are:

- Is there an Object Oriented API?
- Should CCN add an official way to advertise content?
- Are (CCN) conventions protocols?
- What about streams/flows?

4.2 Applications & API (2)

Participants

Marcus Brunner, Antonio Carzaniga, Elwyn Davies, Myeong-Wuk Jang, Gunnar Karlsson, Ashok Narayanan, Börje Ohlman, Lixia Zhang

Discussion and Open Questions

Any API discussion needs to clarify at which level the API will be deployed (application versus host stack API; host versus network API).

Raising the basic question whether we can define a common API across different ICN technologies, the group concludes that this seems not possible. Even basic verbs are not common across different approaches (e.g., *publish* is not relevant in NDN, *advertises* are not relevant in NetInf). Even an API for each individual ICN scheme is challenging as a number of subtleties are not fully defined. At the current state, we do not understand enough about our applications to be able to define these APIs yet. The consensus of the group on this topic is that a *complete* API definition may be premature, but we need to keep working at our current model to harden it.

In the subsequent discussion, the group analyzed (a) *search* as a network primitive, (b) ICN as a network storage, (c) content revocation, and (d) interactive services. Only little consensus was found. Regarding the question: “Can ICN offer reliable, persistent network storage?”, there are two observations. First, there is no guaranteed reliability without explicit agreement. Second, ICN offers ability to delegate storage. All that is missing is some protocol behaviour to bridge the gap between agreement and publication. This can be implemented by PubSub and with NDN.

Regarding content revocation they argue that this is not likely in any case. It is also not apparently a good use of anybody’s time. Do something simple (like do-not-cache bits) and the rest of the problem belongs to the application.

In addition to this, the following open questions remain:

- Are networks sender driven or receiver driven?
- How do you implement a push model in NDN?
- API for “the network cannot find this”?
- Security?

4.3 Metrics and Evaluation (1)

Participants

Giovanna Carofiglio, Patrick Crowley, Van Jacobson, Vikas Kawadia, Dirk Kutscher, Luca Muscariello, George Parisis, Christian Esteve Rothenberg, Thomas C. Schmidt, Eiko Yoneki

Discussion and Open Questions

The main topics of this group includes (a) project testbeds, (b) metrics, and (c) economic aspects. The group noted the importance of controllable testbeds, i.e., dedicated hardware resources for repeatable experiments.

For an exhaustive list of metrics, the group pointed to the presentation given by Giovanna Carofiglio. They highlighted that the benefits of low-level caching, which helps in case of

packet loss (retransmission) and mobility scenarios. Memory can be used to implement three things: rate adaptation, repairing, and re-use (sharing).

An additional focus in the context of metric discussion was forwarding properties and congestion control. A major insight lies in the consideration of RTT variation. A common assumption is that spectrum of RTTs in CCN is lower. Other issues regard fairness with locally cached content (low RTT) and far-away content (high RTT). In particular, a CCN application cannot know where named pieces come from.

Regarding economic aspects, there was no agreement on the revenue model (i.e., which stakeholder gets which value). Metrics for network management and OPEX profiles should also be part of future work.

4.4 Metrics and Evaluation (2)

Participants

Bengt Ahlgren, Somaya Arianfar, Ken Calvert, Kevin Fall, Holger Karl, Pasi Sarolahti, Matthias Wählisch

Discussion and Open Questions

The group agreed that concrete testbeds should be preferred in contrast to generic testbeds. Then, the objectives of a testbed have been discussed. It is impossible to *verify* that an architecture is good. It is doable to *falsify* that an architecture is good. However, the most reasonable intention behind a testbed is to get concepts *running*.

A major open question is: How to quantify “betterness”?

The group came up with only small additions to Giovanna’s list of metrics. This includes ease of programmability, accountability, ease of introducing net business models. The group noted that performance evaluation will only give arguments to dismiss ICN. Hence, we need elaborate arguments beside performance evaluation to push ICN.

Discussions have emphasized that limiting ICN to caching will miss 90% of today’s applications.

4.5 ICN vs. DTN

Participants

Bengt Ahlgren, Somaya Arianfar, Patrick Crowley, Elwyn Brian Davies, Kevin Fall, Pan Hui, Gunnar Karlsson, Vikas Kawadia, Dirk Kutscher, Anders Lindgren, Pasi Sarolahti, Ignacio Solis, Eiko Yoneki

Discussion and Open Questions

After clarifying basic terms Delay Tolerant Networking (DTN) and ICN, the group presented a nice table comparing both approaches. The group concluded that ICN and DTN are not redundant but similar. Both schemes generally can provide disruption tolerant operations for link outages. There are two essential requirements for DTN support, location-independent naming and node with storage, which is shared with ICN.

4.6 Interdomain routing, forwarding

Participants

Ken Calvert, Giovanna Carofiglio, Antonio Carzaniga, Xiaoming Fu, Volker Hilt, Holger Karl, Luca Muscariello, Ashok Narayanan, Borje Ohlman, George Parisi, Christian E. Rothenberg, Thomas C. Schmidt, Matthias Wählisch, George Xylomenos

Discussion and Open Questions

The intention of this group meeting was to discuss both, interdomain routing and forwarding. However, The time was too short to discuss forwarding aspects.

The routing topic splits basically into two parts: (a) routing on topology-independent labels and (b) routing on topology-dependent labels.

For the first, there are two broad options. Distributing routing tables based on topology-independent (TI) names, and converting topology-independent names to topology-dependent (TD) addresses via lookup schemes. Routing tables and forwarding are then based on TD addresses. Obviously, hybrid concepts may also exist. Still an open question: Can interdomain routing scale based on topology-independent labels? Even if we currently cannot imagine routing schemes that can solve it, it does not mean it cannot be solved.

Regarding routing on topology-dependent labels there was also no real consensus. Apart from mobility aspects people did not seem interested.

Overall, the participants concluded that a tutorial on network theory may help improve the debate about TI or TD labels.

After this, the extension of interdomain routing policy has been discussed. It is quite likely that this will happen in the future as CDNs already exchange much more sophisticated policies. It has been identified as a useful area of future research.

Further open questions are

- Can incomplete-table type routing schemes work?
- Can proposed lookup service store # of states

4.7 ICN Deployment

Participants

Carsten Bormann, Marcus Brunner, Lars Eggert, Anders Eriksson, Myeong-Wuk Jang, Lixia Zhang

Discussion and Open Questions

Starting from the observation that ICN is easy to deploy as an overlay or in the infrastructure, it is still not fully clear why people would want to do this. In other words, we still miss a killer application.

As a disruptive technology arising that we can piggyback ICN on, the following has been noted. First, a big player that could make ICN actually happen is Google with Android. However, why would they want to do this. Elaborating this could be part of future work.

Another interesting application field for ICN may also machine-to-machine communication and the Internet of Things (IoT). If applying ICN to IoT, it will give us some sane routing (compared to the current protocols pushed in the IETF), which could be a killer app.

Participants

- Bengt Ahlgren
Swedish Institute of Computer Science – Kista, SE
- Somaya Arianfar
Aalto University, FI
- Carsten Bormann
Universität Bremen, DE
- Marcus Brunner
Swisscom AG – Bern, CH
- Ken Calvert
University of Kentucky, US
- Giovanna Carofiglio
ALCATEL – Marcoussis, FR
- Antonio Carzaniga
University of Lugano, CH
- Patrick Crowley
Washington University, US
- Elwyn Brian Davies
Trinity College Dublin, IE
- Lars Eggert
NetApp Deutschland GmbH – Kirchheim, DE
- Anders Eriksson
Ericsson Res. – Stockholm, SE
- Christian Esteve Rothenberg
CPqD – Campinas, BR
- Kevin Fall
Qualcomm Corp. – Berkeley, US
- Xiaoming Fu
Universität Göttingen, DE
- Volker Hilt
Alcatel-Lucent, DE
- Pan Hui
TU Berlin, DE
- Van Jacobson
PARC – Palo Alto, US
- Myeong-Wuk Jang
Samsung, KR
- Holger Karl
Universität Paderborn, DE
- Gunnar Karlsson
KTH – Stockholm, SE
- Vikas Kawadia
BBN Technologies – Cambridge MA, US
- Dirk Kutscher
NEC Laboratories Europe – Heidelberg, DE
- Anders Lindgren
Swedish Institute of Computer Science – Kista, SE
- Luca Muscariello
Orange Labs, FR
- Ashok Narayanan
Cisco Systems – Lexington, US
- Börje Ohlman
Ericsson Res. – Stockholm, SE
- George Parisis
University of Cambridge, GB
- Pasi Sarolahti
Aalto University, FI
- Thomas C. Schmidt
HAW – Hamburg, DE
- Ignacio Solis
PARC – Palo Alto, US
- Matthias Wählisch
FU Berlin, DE
- George Xylomenos
Athens University of Economics and Business, GR
- Eiko Yoneki
University of Cambridge, GB
- Lixia Zhang
Univ. California – Los Angeles, US

