10492 Abstracts Collection
Information-Centric Networking
— Dagstuhl Seminar —

Dirk Kutscher[1], Bengt Ahlgren[2], Holger Karl[3],
Börje Ohlman[4], Sara Oueslati[5] and Ignacio Solis[6]

[1] NEC Laboratories Europe – Heidelberg, DE, kutscher@neclab.eu
[2] SICS – Kista, SE, bengta@sics.se
[3] Universität Paderborn, DE, holger.karl@uni-paderborn.de
[4] Ericsson Research – Stockholm, SE, borje.ohlman@ericsson.com
[5] Orange Labs, FR, sara.oueslati@orange-ftgroup.com
[6] PARC – Palo Alto, US, isolis@parc.com

Zusammenfassung From December 5th to 8th 2010, the Dagstuhl Seminar 10492 on "Information-Centric Networking" was held in Schloss Dagstuhl – Leibniz Center for Informatics. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

Keywords. Information-Centric Networking, ICN, Content-Centric Networking, CCN, Data-Oriented Networking, DONA, NetInf, 4WARD, SAIL

1 Introduction

Information-Centric Networking (ICN) is one of the significant directions of current networking research. In ICN, the principal paradigm is not end-to-end communication between hosts - as it is in the current Internet architecture. Instead, the increasing amount of content that must be distributed requires alternatives: Architectures that work with information objects as a first-class abstraction; focusing on the properties of such objects and receivers' interests to achieve efficient and reliable distribution of such objects. Such architectures make in-network storage, multiparty communication through replication, and interaction models such as publish-subscribe generally available for all kinds of applications, without having to resort to dedicated systems such as peer-to-peer overlays and proprietary content-distribution networks.

The ICN approach is currently being explored by a number of research projects, both in Europe (4WARD, SAIL, PSIRP) and in the US (DONA, CCN).

Dagstuhl Seminar Proceedings 10492
Information-Centric Networking
http://drops.dagstuhl.de/opus/volltexte/2011/2943

The Delay Tolerant Networking (DTN) community has developed a message-oriented architecture that has been used along with ICN addressing and routing concepts. While these approaches differ with respect to their specific architecture, they share some assumptions, objectives and certain structuring architectural properties. In general, the aim is to develop network architectures that are better suited for content distribution, the currently prevailing usage of communication networks, and that better cope with disruptions in the communication service. The basic idea of ICN still leaves room for many variations. The Dagstuhl ICN seminar was intended as a catalyst for these variations and as a forum for discussing the following research topics:

- The relationship of networking architecture innovation vs. so-called over-the-top approaches in the application layer
- The support of an Internet of Things and Services by an ICN architecture
- How to migrate towards an information-centric architecture, and whether and how to use it as a migration enabler for, e.g., an IPv4/IPv6 technology step
- The role of and needs for naming and addressing and name resolution systems, along with the necessary security aspects of a naming scheme; a fundamental dichotomy between flat and hierarchical naming schemes needs to be resolved
- Efficiency and robustness of ICN data dissemination vs. specific content distribution overlay solutions
- The desirability of using specific transport protocols for ICN vs. the use of standard protocols like TCP or disruption tolerant protocols like the DTN Bundle protocol
- The integration and placement of caches inside a network
- Can the introduction of a new ICN architecture enable new types of applications that were too complex to create/operate/deploy/maintain in traditional networks?

The seminar delivered a comprehensive analysis of the state of the art in information-centric networking, progress on specific technical issues such as scalable addressing and content distribution, a better understanding of the legal requirements and application developer needs. It also touched upon possible next steps in research and helped to form an ICN community. The seminar has led to the organization of a SIGCOMM workshop[7] on the same topic that is co-organized by seminar organizers and participants.

2 Organization of the seminar

The seminar was organized as a 2.5 days seminar that provided room for presentation of approaches, results so far, as well as presentation and discussion of new ideas and selected specific topics.

The seminar was structured in 4 main blocks:

[7] http://www.neclab.eu/icn-2011/

1. Presentation of on-going research activities
2. In-depth presentations and discussion of *naming*, *security*, and *routing and resolution* for ICN (Group Discussion 1)
3. In-depth presentations and discussion of *resource management and transport*, *ICN APIs and ICN hour glass waists*, and *deployments aspects, business models and incentives* for ICN (Group Discussion 2)
4. Discussion of seminar results and next steps

The seminar started, in the first block, with a set of presentations of on-going research activities (Section 4):

- Teemu Koponen: DONA (Data-Oriented Networking Architecture)
- Jim Thornton: NDN (Named-Data Networking)
- Bengt Ahlgren: NetInf (Network of Information) in the 4WARD project
- George Xylomenos: PURSUIT project

The seminar then addressed important specific ICN topics such as naming, security, routing and resolution. For that, a set of discussion starter presentation set the scene by summarizing important issues and by providing new ideas (Section 5):

- Christian Dannewitz: Naming and Security in Information-centric Networking
- Kevin Fall: Discussion on Information Centric Networking with a Security Focus
- Jarno Rajahalme: What's in a Data Name?
- Jussi Kangasharju: Naming and Search in Information-Centric Networks

These topics were then discussed in smaller groups (Group Discussion, part A), and the results of these discussions were presented and discuss in a plenary session (Section 6).

In the second block of specific ICN topics discussion, several discussion starter presentations on resource management, congestion control, and ICN in challenged networks have been given (Section 7):

- Van Jacobsen: Congestion Control and Transport in ICN
- Sara Oueslati: Ideas on Traffic Management in CCN
- Volker Hilt: Energy Consumption of Content-Centric Networks
- Jörg Ott: Delay-tolerant Networking: Elements of ICN
- Stephen Farrell: ICNing DTN
- Armando Caro: Content Based Networking in DTNs
- Christian Esteve Rothenberg: Compact Forwarding in Content-Oriented Networks
- Henrik Lundqvist: Deployment of Information Centric Networking from a Mobile Operator Perspective: Service Program Mobility
- Antonio Carzaniga: Content-Based Publish/Subscribe Networking and Information-Centric Networking

Aspects of these presentation were then discussed in *dedicated* groups on *resource management and transport*, *ICN APIs and the ICN hour glass waist*, and *deployment aspects, business models, and incentives* (Section 8).

The seminar was wrapped up by a discussion of common concepts, future research topics and next steps for the ICN community.

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2011/2942

3 A Survey of Information-Centric Networking

Bengt Ahlgren; Christian Dannewitz; Claudio Imbrenda; Dirk Kutscher; Börje Ohlman

In this paper we compare and discuss some of the features and design choices of the 4WARD Networking of Information architecture (NetInf), PARC's Content Centric Networking (CCN), the Publish-Subscribe Internet Routing Paradigm (PSIRP), and the Data Oriented Network Architecture (DONA). All four projects take an information-centric approach to designing a future network architecture, where the information objects themselves are the primary focus rather than the network nodes.

Keywords: ICN, CCN, NetInf, DONA, PSIRP

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2011/2941

4 On-going research activities

4.1 Architectural Commonalities and Implications

Teemu Koponen (NICIRA, SE)

This presentation reflects on experiences with DONA (Data-Oriented Networking Architecture) in recent years, presents identified commonalities of different approaches such as TRIAD, Haggle, DONA, CCN, PSIRP, NetInf, SCAFFOLD, S-GET and discusses architectural implications.

Keywords: DONA

4.2 CCN/NDN Overview

Jim Thornton (PARC – Palo Alto, US)

A brief, high-level overview of the CCN and NDN projects, focusing on generalizing the current communication network architecture into a distribution network architecture. The core goal of these projects is to create, refine and validate a 'narrow waist' protocol for information networking. There are a variety of hard problems raised by this goal, and we introduce the broad research agenda centered on trade-offs between the needs of applications and the challenges of routing, forwarding, and security.

4.3 4WARD Networking of Information overview presentation

Bengt Ahlgren (Swedish Institute of Computer Science – Kista, SE)

Networking of Information (NetInf) is the information-centric network architecture developed in the 4WARD EU project, and now being further worked on in the SAIL EU project. The NetInf information model handles information at different abstraction levels, including multiple encodings. Information object names have three fields: type, authenticator and label, similar to naming in DONA. The NetInf name resolution and routing framework allows multiple routing schemes in different administrative domains. The NetInf API borrows ideas from publish/subscribe. A proof-of-concept prototype is available as open source, implementing the naming scheme, security functions, NetInf core services, and some example applications, for instance, a Firefox web browser plugin.

4.4 Publish/Subscribe Internetworking: From PSIRP to PURSUIT

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

The goal of the PSIRP and PURSUIT projects is to design, prototype and evaluate an internetwork architecture based on the publish/subscribe model not only in the protocols, but also in the host implementations and programming interfaces. While the PSIRP project focused on the waist of the network, PURSUIT will also cover higher and lower layer issues, as well as continue evolving the PSIRP architecture and implementation. In this talk I will provide a brief overview of the goals, working methods and achievements of PSIRP, focusing on its three main components: rendezvous, topology and forwarding. I will also talk about the goals of PURSUIT, focusing on the issues raised during PSIRP and the additional challenges that PURSUIT takes on, such as transport and link specific issues.

Keywords: PSIRP, PURSUIT, ICN, Publish/Subscribe

5 Discussion starter presentations, part A

5.1 Naming and Security in Information-centric Networking

Christian Dannewitz (Universität Paderborn, DE)

There are several different naming scheme proposals in the information-centric networking (ICN) community. This presentation investigates the core differences between different approaches and investigates which are the core naming-related questions to answer within the community. In the presentation, a collection of general properties potentially desired for an ICN network are introduced, also looking at tradeoffs between different properties. As a conclusion, I look at possible approaches to integrate several different ICN naming schemes into a combined naming framework.

Keywords: Naming, security

5.2 Discussion on Information Centric Networking with a security focus

Kevin Fall (Intel Berkeley Labs, US)

This presentation discusses Information-Centric Networking with a security focus including ICN security model issues, data acces control, and network access security.

Keywords: ICN, security, access control

5.3 What's in a Data Name: Human-readable/Semantic-free/ self-certified: DoS Implications?

Jarno Rajahalme (Nokia Siemens Networks – Espoo, FI)

Choices in namespace security properties, like use of self-certification, have impact on denial-of-service resistance properties of the network.

Keywords: Self-certified names, infrastructure protection, certificate authorities

Joint work of: Jarno Rajahalme; Pasi Sarolahti

5.4 Naming and Search in Information-Centric Networks

Jussi Kangasharju (University of Helsinki – FI)

Naming of information in information-centric networks is an important and active topic. We can identify three types of names being used or proposed. First are machine-readable names, e.g., self-certifying names, which are efficient for routing and forwarding. Second are human-readable names which could be memorized and easily typed in by human users. Third are names that could be used in advertising. These are likely to be a sub-type of human-readable names. All three types of names appear to have a need, but most of the work has focused on machine-readable names. In this talk, we argue that ädvertizable"names serve a vital role in the success of information-centric networks and they must not be forgotten.

Searching in information-centric networks has several possibilities. On the one hand, we can integrate (some degree of) search functionality directly into the network infrastructure. On the other hand, normal web searching solutions based on crawling and indexing (e.g., Google) will be popular and important, hence their requirements need to be taken into account in the design of information-centric networks.

Keywords: Naming, Search, Self-Certification

6 Group discussions, part A

6.1 Group and topic setup

The discussion starter talks above had led to a discussion on the topics *Naming for Information-Centric Networking*, *Security for Information-Centric Networking*, and *Routing and Name Resolution for Information-Centric Networking*, which was prepared by the organizers by providing an initial list of discussion topics and specific questions for each topic.

Five group were formed to discuss all of these details, and the results from these discussions have been presented (see discussion results abstracts in Sections 6.2 to 6.5).

Naming for Information-Centric Networking

- What do we name at all?
- Names tied to topology / organization?
- Hierarchical vs. flat names?
- Name assignment?
- Unique names?
- Anonymous names?
- How many layers of naming?
- Meta data
- URIs, Search?

Specific questions on naming

- How would the ICN WWW look like?
- How would user interfaces look like? How to enter a name?
- How would you find the printer next room / the temperature sensor on the Mars rover?

Security for Information-Centric Networking

- What are the security goals?
- What needs to be authenticated in ICN?
- What are the new threats?
- Requirements (PKI?, Requirements on routers/nodes)
- Trust chains

Specific questions on security

- How to design / control Wiki Leaks in ICN?
- Will ICN better protect privacy than today's networks?
- Trusted Computing Platforms for ICN?

Routing and Name Resolution in Information-Centric Networking

- Names and resolution architectures
- Name spaces (on different layers, in different domains)
- Name-based routing vs. resolution (could use a good definition)
- Routing on non-aggregatablenames
- Dynamicity in topology changes, mobility
- Resolution approaches (Distributed / step-wise resolution / forwarding, late binding)

- Route symmetry

Specific questions on routing and name resolution

- How to use ICN with a uni-directional satellite link?
- How would "DNS for ICN" look like and how many do we need?

Keywords: Naming, Security, Routing, Resolution

6.2 Group A1

The following two issues were discussed:

1. Should signatures be mandatory?
   Signatures in names are useless by themselves. Signatures with a trust model are useful, but trust model varies by application. Some application trust models will require 2 signatures. Also, signature transition would also be difficult to handle. Signing with bad ECDSA parameters might be a nice DoS vector if routers try verify. So no reason to make signatures mandatory, i.e., part of the ICN thin waist. Hashes are quite good enough. This position was expressed by S. Farrel, it is, however, not a consensus position.
2. Geolocation and ICN?
   ICN by separating identity and location, eliminates a rather simple way of mapping to geography. So what do we replace it with? It was proposed to rely on encryption. Not in the thin waist, but we expect ICN to support. Then, access control boils down to key distribution.
   One option is that access network provider certifiy consumers; ISP asserting something about their customers (e.g. hardwire line connection). CP will rely on this certification. This may involve explicit agreements between CPs and ISPs (Otherwise, there might be a trust problem: Do I believe this is an ISP?). The consumer could also claim a certificate from its ISP (I am in your geo!).
   An alternative option is to rely on "content firewall", in the form of a "Border router" (enterprise, ISP) preventing a collection of content from transiting out of the border.

Keywords: Security, Signatures, Geo-Location

Joint work of: Group A1 participants

6.3 Group A3

Context dependent naming can simplify for users. Names should be syntactically constant but take their semantic meaning from the context, e.g. `THE_PRINTER_IN_THIS_ROOM`. In PSIRP the resolution service is a global service, how is it getting local context? Isn't that what scopes are used for? Using flat names (i.e. non human readable names) means you need name resolution service.

Mapping between application names and ICN naming, how should this be done? This relates to Joseph Halpern's work on naming, http://www.cs.cornell.edu/home/halpern/node8.html. What is needed to build a trust structure? You need the triple name, key and data. That triple is both sufficient and necessary. Self-certifying names only have two of the three. You need contextual data NetInf etc. does not have that. There is a problem with delegation of trust. Only the resolution system has the triple, this means that you need to trust the resolution system.

When developing the Arpanet they initially had an integrated protocol providing both transitivity and reliability (IP and TCP integrated). When Jon Postel was working on a voice packet service he realised that he only needed transitivity and thus proposed splitting up TCP and IP. What can be learnt from this is that, when creating a new communication architecture, it is very useful to try it out for other applications than what was intended in the early design phase to see how general the architecture is. This should be kept in mind, and practised when we are building a new ICN architecture.

Today we have a security model where we create a secure channel to a box we trust and then we trust the information that the box delivers to us via that channel. With ICN, where you can trust the data by itself, will the need for trusted computing platforms be reduced? How will it be different? You still will need to trust your local box which is doing your data verification and the rendering of the received information. Assumably there should be a reduced need for trust in remote hosts as you can verify the data received. This can ease the requirement on e.g. cloud computing environments. Some remote host you will still need to trust, e.g. name resolution servers, and those being part of an infrastructure used for delivering the names that you'll put your trust in.

Being a root-CA is exorbitantly expensive. There are not good incentives for them to do their job properly. One example is when a person called up four root-CAs and asked for a signing certificate for live.com, which is Microsofts site for Windows updates, the only question three of the four asked was if he wanted to pay with VISA or Mastercard. Trust needs to be based on evidence; you trust those that have proved in the past that they are trustworthy.

An alternative could be use of Simple Distributed Security Infrastructure (SDSI), http://groups.csail.mit.edu/cis/sdsi.html, which create local namespaces distinguished by the unique public key of the entity defining the names, instead of trying to create a globalized namespace. A major issue with SDSI is as there is no central authority involved there is no clear business case for how to deploy it. Another alternative to investigate is if each publisher could provide its own resolution system.

Keywords: Naming, name resolution, security, trust

Joint work of: Lixia Zhang, Van Jacobsen, Börje Ohlman, Eiko Yoneki

6.4 Group A4

The group started to discuss naming of information, for instance, structured or unstructured (flat) namespaces, names for entering in browsers, names that are globally unique or not, and the granularity of names.

In networking we quickly conclude that we need globally unique names, but in the real world this is the exception. Most, if not all, names are instead as local as possible within some context.

The group discussed naming granularity. If there are names for small objects (packets or even smaller), there is more overhead per byte. Larger objects mean less relative overhead for security, routing, etc. There is however a difference if the namespace is hierarchical and support aggregation. The group agreed that it is "natural"to name objects that we store in a single file in our filesystems. For complex objects, like web pages consisting of many sub-objects, it may be more natural for the user if the whole page can be considered an object with its own name.

The group then turned to security and privacy. *Will ICN prevent spam?* Probably not. The persistence of information was discussed. *How can you remove information that you previously published?* If published information has a stated lifetime in its metadata, you cannot technically make everyone abide to it. Legal measures are also needed. Access control is another problem. If content is distributed encrypted, access control turns into a key management problem, and may be relying on a trusted platform module (TPM) wherein the decrypted content can be handled.

Are we with ICN creating the perfect tool for repressing free speech? We concluded that this to some degree depends on whether you can create your own publisher keys or not. If you can, the situation improves compared to today, else it gets worse.

We discussed whether ICN will be more secure or not compared to current networks. On the one hand, ICN provides integrity and authentication checks for every information object. On the other hand many security issues are more due to that users don't really understand how the security works, or they are simply fooled, so more security mechanisms won't help. Perhaps more user education will?

How will ICN look like for the user? Not much different at all compared to today. The group was a little dissapointed that we could not find clear benefits with ICN for the end user.

Keywords: Naming, Security

Joint work of: Armando Caro; Jussi Kangasharju; Pan Hui; Henrik Lundqvist; Bengt Ahlgren

6.5 Group A5

In this group, there was an extensive dicussion on the concrete semantics of different information-centric networking approaches (in particular, about CCN, NetInf, and PSIRP). A lag of a formally described interface semantics was identified for all these existing approaches. As a complementing approach to document-oriented information-centric networking, the notion of identifying events by content-based filtering and providing a true event-notification semantic was discussed. There was, however, no clear consensus whether this is indeed a complement, a competing approach, or already covered by at least some existing approaches (NetInf, in particular, claims to have such functionality already included). This discussion again highlighted the need for a formal definition of what information-centric networking actually implies.

On the topic of whether human-readable names are necessary (and what that would imply for the user interface), consensus was quickly reached that such names are in fact not necessary – however, this discussion has to be put into context whether a human enduser or an application-level programmer is considered (the first clearly needed human-readable names much in the form of today's URLs, the latter likely not interested in the concrete representation of such names). An even stronger point was debated whether any kind of names at all are required or whether it would not be preferable to use predicates defining matching content instead (concerns about accuracy and false positives of such predicates were voiced). The particular structure of the names has obvious repercussions on the network's efficiency; if names can be chosen by an adversary and placed at arbitrary points in the network, no efficient routing/forwarding scheme is possible (from theoretic considerations).

The particular security challenges for information-centric networking seem not to be well understood or agreed upon at this point in the discussion.

Keywords: ICN semantics, Content-based Networking, Security

Joint work of: Group A5 participants

7 Discussion starter presentations, part B

7.1 Congestion Control and Transport in ICN

Van Jacobsen (Palo Alto Research Center – CA, USA)

Based on the concept of *flow balance* for packet-based data communications, this talk disucussed resource management issues from a content-centric networking perspective.

Keywords: CCN

7.2 Ideas on Traffic Management in CCN

Sara Oueslati (Orange Labs, FR)

Jacobson has argued convincingly that the Internet should be re-designed to facilitate content dissemination. His proposed content-centric networking (CCN) paradigm would bring significant advantages, notably with respect to security, mobility and effciency. The CCN architecture is still incom- plete, however, notably in respect to the way bandwidth sharing should be controlled in CCN to ensure applications experience acceptable quality. The aim of this presentation is to discuss this issue.

It is necessary, for instance, that packets of voice calls should experience negligible delay when they compete for bandwidth with high speed docu- ment downloads. Similarly, it should not be possible for users to unduly impact the quality experienced by others by greedily or maliciously request- ing downloads at a rate that is too high. These functions are performed, imperfectly, in IP through various QoS mechanisms in the network and TCP congestion control implemented in end systems. It is necessary to carefully examine whether IP QoS can be transposed and to and a CCN replacement for TCP.

Our proposal is that CCN should implement flow-aware controls where a flow would be identified by the object name included in Interest and Data packets. We argue that fair bandwidth sharing on network links is sufficient to meet performance requirements as long as additional controls are in place to limit the impact of overloads. More elaborate, user controlled sharing is advocated for the "last mile"resources between user and CCN access node.

Bandwidth sharing controls have strong economic implications since they determine what service level agreements are feasible. For present purposes we assume a simple business model where transport is unbundled from any value added service. This highlights the importance of traffic controls and emphasizes the incentive to use caching which is an important feature of CCN.

The presentation is structured as follows. We first recall salient features of CCN that are necessary for our discussion. In the next sections, we proceed successively to consider traffic management in the network, where the only key for sharing is the packet name, and on the link from access node to user, where the specific requirements of individual ows can be taken into account.

Keywords: Traffic management, CCN, fairness, QoS, business model, transport

Joint work of: Sara Oueslati; Jim Roberts; Nada Sbihi

7.3 Energy Consumption of Content-Centric Networks

Volker Hilt (Alcatel-Lucent Bell Labs – Holmdel, US)

To meet the ever-increasing demand for content, content and network providers are rapidly expanding their server and network infrastructure. Even today, the servers and network devices used for content distribution consume a substantial amount of energy. In this talk, I will introduce an energy efficiency analysis of various content dissemination strategies. A key result of this study is that a change in network architecture from host-oriented to content-centric networking (CCN) can open new possibilities for energy-efficient content dissemination. I will present an analysis of the energy-efficiency of a CCN architecture and present trace driven results. Our results show that CCN is more energy efficient than conventional CDNs and P2P networks, even under incremental deployment of CCN-enabled routers.

7.4 Delay-tolerant Networking: Elements of ICN

Jörg Ott (Aalto University, FI)

DTN operation and paradigms lend themselves quite nicely to support information-centric networking ideas.

Self-contained messages with identifiable content and operations are the basic building blocks.

Suitable application and protocol design will allow for a smooth transition between reliable infrastructure and probabilistic ad-hoc operation.

Keywords: ICN, DTN

7.5 ICN for challenged networks

Stephen Farrell (Trinity College Dublin, IE)

Report on 2010 arctic DTN trial from the N4C project as a way of presenting requirements for ICN in challenged networks, and outline of a planned bundle protocol query (BPQ) extension to the bundle protocol to support ICN.

Keywords: ICN, DTN

7.6 Content Based Networking in DTNs

Armando Caro (Raytheon BBN Technologies – Cambridge, US)

We have investigated some aspects of Content Based Networking (CBN) within the context of Delay/Disruption Tolerant Networks (DTNs). This talk begins by presenting our perspective on DTNs and its fluid relationship with more well-connected Mobile Ad hoc Networks (MANETs) and even fixed infrastructure networks like the Internet. This perspective motivates our vision and our approach to CBN in DTNs. We present the high level concepts for a infrastructureless content distribution mechanism that takes into account network topology dynamics, regional content demand differences, resource constraints, and user perceived latency.

Keywords: Content based network, delay tolerant, disruption tolerant, mobile ad hoc, infrastructure content distribution

Joint work of: Armando Caro; Vikas Kawadia; Niky Riga

7.7 Compact forwarding in content-oriented networks

Christian Esteve Rothenberg (University of Campinas, BR)

Advances in efficient packet forwarding techniques have been central to continuously moving traffic smoothly through the Internet at increasing rates. Much work has been invested in data structures and algorithms for packet forwarding and classification. Research in the design of forwarding table compacting techniques has been a continuum since the early 90s, and still goes on, yielding novel compact representations for structured graphs such as tries, new algorithms and data structures for IP lookups, packet classification, and advances in high-speed memory technologies among others.

Content-oriented network architectures are characterized by introducing new namespaces for content objects. A common property of the proposed naming schemes is relying on flat identifiers (e.g., 256-bit hash values) and/or long, non-fixed size URL-like names (e.g., TRIAD, CCN) to uniquely identify single pieces of content. Other network architectures that separate identifiers from locators or aiming at scalable Ethernet designs, face similar challenges when handling packets carrying flat identifiers. A flat naming scheme simplifies address administration or content identification but is hard to scale due to the lack of aggregation capabilities. Structured identifiers (e.g., NDN) are also hard to handle at wire speed due to the challenges of performing LPM-like lookup operations on arbitrary long identifiers resulting from non-fixed size components.

Similar to the advances in algorithms and data structures that enabled the feasibility of high-performance IP routers, we surmise that new enablers in the forwarding plane may be fundamental to the realization of content-oriented networks. More specifically, we expect probabilistic techniques to play a key role to guide the construction of data structures well-suited for the requirements of packet forwarding in content-oriented networks.

Motivated by the needs of content-oriented networking proposals, we have explored new approaches to the fundamental trade-offs of packet routing to provide forwarding services with scalability, multicast-friendliness and security in mind. The main idea behind compact forwarding is taking a probabilistic approach to the problem of packet forwarding in networks centered on content identifiers rather than traditional host addresses.

Due to the lack of aggregation capabilities of flat labels and the compact forwarding goal of seeking the minimal information base to deliver packets at scale, we have dived into solutions based on error-prone probabilistic data structures providing lossy compression functionality.

A fundamental question explored is where to place the packet forwarding state, in network nodes or in packet headers? Solutions for both extremes are proposed. In the SPSwitch, approximate forwarding state is kept in network nodes. In LIPSIN, the state is carried in the packets themselves. Both approaches are based on probabilistic packet forwarding functions inspired by variations of the Bloom filter data structure. The approximate forwarding state comes at the cost of additional considerations due to the effects of one-sided error-prone data structures. By exchanging correctness (traduced in forwarding efficiency penalties) for space/memory time requirements (traduced in reduced forwarding information base in packet headers and network nodes), new spots in the design space can be explored.

Keywords: Forwarding, routing, state, multicast, probabilistic, data structure, Bloom filter

7.8 Deployment of Information Centric Networking from a Mobile Operator Perspective: Service Program Mobility

Henrik Lundqvist (DoCoMo Euro-Labs – München, DE)

Among mobile operators there is a strong interest in delivering value added services composed by multiple service components. Information centric concepts can be applied to future service delivery platforms, for example name resolution and routing based on names. However, there is also a case for adding extensions to support service delivery, in particular software processes can be considered as a generalization of information objects. For example, this requires placing processing nodes in the network in addition to caches, it creates new security challenges, and interconnection of multiple objects has implications on service placement.

7.9 Deployment of Information Centric Networking from a Mobile Operator Perspective: Service Program Mobility

Antonio Carzaniga (University of Lugano – IT)

I argue that content-based publish/subscribe communication is an essential form of communication for several important applications, and that such communi- cation could and should be realized as a network service. I also argue that the notion of content-centric networking proposed by Van Jacobson et al is complementary to content-based publish/subscribe networking, and that both are important for the design of a more general information-centric network layer.

I use the term *content-based publish/subscribe communication* to refer to the immediate transmission of short ephemeral messages (e.g., event notifications) from producers to all interested consumers. An essential feature of this form of communication is that the information flow is instigated by the producer. This "push" communication mode is in contrast with the traditional "pull" communication model of the Web, in which producers only respond to explicit consumer requests. Needless to say, these two modes are at some level equivalent, in that they can implement each other. However, it is still important to distinguish primitives that are *designed* to implement one over the other.

Another essential feature of content-based publish/subscribe networking is that the flow of information is determined by consumer interests predicated upon message content. More specifically, a consumer declares a selection criterion. In general, such a criterion (or *predicate*) may apply to the content of an individual message, but also to other properties of the environment, or even to the whole flow of messages to that consumer. For example, a consumer might be interested in receiving “sport news” and “network management events,"but at the same time it might want to limit the number of messages received per time unit, possibly according to the time of day, and it might also require that all messages be authentic with respect to a given set of credentials. This way of delivering information (by content) distinguishes content-based publish/subscribe networking from content-centric networking (although it might be applicable there, too) and other traditional network services such as IP multicast.

The thesis I put forth is that (1) several applications motivate content-based publish/subscribe communication, (2) the nature of this communication model motivates its development as a network service, (3) content-based publish/subscribe networking differs significantly from content-centric networking, both in its purpose and in the nature of the communication, and yet (4) content- based publish/subscribe networking and content-centric networking embody compatible services that admit to a common architecture and that might well be realized on the basis of synergistic protocols.

Keywords: Content-based communication, publish-subscribe

8 Group discussions, part B

A detailed discussion of aspects of the presentations described above were then dicussed in three dedicated groups on:

- resource management and transport;
- ICN APIs and the ICN hour glass waist; and
- deployment aspects, business models, and incentives.

Keywords: Resource management, transport, API, deployment, business models, incentives

Joint work of: Bengt Ahlgren; Holger Karl; Dirk Kutscher; Börje Ohlman; Sara Oueslati; Ignacio Solis

8.1 Resource Management

An extensive discussion about details of resource management took place, concentrating on CCN as a case example. Questions about how to pace interests in various scenarios were discussed (e.g., differences between DSL and PON downlinks, how to reuse interest from several interested parties, and how this relates to different fairness constraints imposed on ill-behaving users). Another issue was buffer memory and whether it is useful or necessary to distinguish between memory used for keeping the pending interests and the actual content, and how to behave if either of these memory types fills up. Several open issues were identified as well: how to gauge timer settings, prioritizing interests (does it make sense to keep state about the originating user per interest), how much explicit signaling to include (e.g., does a node receiving and interest acknowledge that it could store the interest locally or that is has been dropped?). On a conceptual level, interests behave differently than standard flows since both the interests themselves as well as the corresponding data flows can be merged, deviating from the standard models of flow theory. The discussion also briefly touched upon possibilities to run denial-of-service attacks against CCN (e.g., random request attacks to fill up the pending interest memory).

Generalizing from CCN as a case study was considered to be difficult as, for example, NetInf pursues quite different resolution mechanisms than CCN and likely would need very different resource management solutions. On the one hand, resource management in NetInf might be considerably simpler (less or even no state to be maintained per flow); on the other hand, it is not obvious how to enforce desirable properties like symmetric traffic.

Keywords: ICN semantics, Content-based Networking, Security

Joint work of: Group B1 participants

8.2 ICN API and ICN Hour Glass Waist

8.3 Deployment

Moving CDNs into the network is one of the key drivers for ICN technology. It is essential that any new caching model allows for generating revenue from advertisements intertwined with content, at least as well as is possible in today's networks. Legal frameworks strongly influence how caching can be used for copyrighted content. ICNs could help democratizing the use of CDN business, allowing pay/earn as you go business models. ICN can offer new opportunities by combining cloud processing resources with caching.

Keywords: ICN semantics, Content-based Networking, Security

Joint work of: Börje Ohlman; Björn Groenvall; Henrik Lundqvist; Jussi Kangasharju; Jarno Rajahalme; Volker Hilt; Armando Caro