

06441 Abstracts Collection
Naming and Addressing for Next-Generation
Internetworks
— Dagstuhl Seminar —

Bengt Ahlgren¹, Lars Eggert², Anja Feldmann³, Andrei Gurtov⁴ and Tom R.
Henderson⁵

¹ SICS - Kista, SE

`Bengt.Ahlgren@sics.se`

² NEC Europe - Heidelberg, DE

`lars.eggert@netlab.nec.de`

³ TU München, DE

`anja.feldmann@telekom.de`

⁴ HIIT - Helsinki, FI

`gurtov@cs.helsinki.fi`

⁵ Boeing Phantom Works - Seattle, US

`thomas.r.henderson@boeing.com`

Abstract. From 29.10.06 to 01.11.06, the Dagstuhl Seminar 06441 “Naming and Addressing for Next-Generation Internetworks” was held in the International Conference and Research Center (IBFI), Schloss Dagstuhl. 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. Naming, addressing, network architecture, next-generation networks, security, privacy

06441 Summary – Naming and Addressing for Next Generation Internetworks

The design of naming and addressing for data networks is a fundamental architectural consideration, and several current or anticipated problems in the Internet - including mobility dynamics, forwarding table growth in the core routers, and security - point out possible limitations with naming and addressing schemes in use today. A seminar on the topic of naming and addressing for next generation internetworks was held at the Schloss Dagstuhl from October 29 to November 1, 2006. Researchers from different fields discussed their views and recent results pertaining to naming and addressing problems. Over twenty talks covered

topics such as routing, naming components, APIs, mobility, delay-tolerant architectures, flat routing and deployment issues. This article briefly summarizes the seminar presentations and discussions.

Keywords: Network architecture, scalability, mobility, heterogeneity, extensibility, naming, addressing

Joint work of: Henderson, Thomas, R.; Gurtov, Andrei; Eggert, Lars; Dannewitz, Christian

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2007/1129>

A naming and addressing tutorial based on Saltzer

Bengt Ahlgren (SICS - Kista, S)

Discussions on naming and addressing in computer and telecommunication networks often result in confusion because the participants have a different understanding of what names and addresses are. The discussion can be made clearer by separating the notions of names and the objects the names refers to, and then carefully distinguishing between properties of the name and properties of the object.

Furthermore, with the help of Saltzer, we argue that there is no fundamental difference between a name, an address and an identifier.

The difference between them lies in the properties of the namespaces and how the namespaces are used. Finally, we argue that it is useful to view "address of" as a relation between objects, implemented as a binding between the names of those objects.

Identity-Locator Merge

Jari Arkko (Ericsson - Jorvas, FIN)

We evaluate the original Internet Architecture in terms of its use of identity and locator in the same token, the IP address, and how our ability to use that identity has eroded over the years. The often suggested remedy for this is a so called identity-locator separation. However, this separation comes at a cost. In particular, to support referrals such an architecture requires the implementation of a mapping function from an identifier to locator. We argue that it is unclear whether the benefits of the split outweigh the costs. As an example of an architecture that does not have these costs we present an approach based on Cryptographically Generated Addresses (CGAs). This approach attempts to bring back the identity role of an address in a way that makes it use possible and secure.

Keywords: Identity-Locator Split, Cryptographically Generated Addresses, Internet Architecture

Joint work of: Arkko, Jari; Bagnulo, Marcelo

What's in a name? Is it possible to use better naming in APIs to solve existing architectural problems?

Saleem Bhatti (University of St Andrews, GB)

In IP-based networks, when deploying functions such as mobility, multi-homing and local addressing (NATs), there are some strong architectural challenges to end-to-end network operation leading to problems in practical use of these functions. Meanwhile, the use of traditional sockets based APIs exposes engineering detail at various levels in the stack to the applications programmer and can compound these problems. This may expose the applications programmer to levels of detail that "entangle" code associated with the application-level communication with some of the engineering detail of the communication stack. We would like to pose the question above in order to discuss whether it is possible to "disentangle" the stack engineering from the application level code and so fix some of the challenges currently being experienced with deployment of the functions listed above.

Joint work of: Atkinson, Ran; Bhatti, Saleem

On What to Name

Ken Calvert (University of Kentucky, USA)

It is well-known that the present Internet architecture conflates several separable concerns, viz., routing, forwarding and addressing.

We are developing a "Postmodern" Internet architecture that attempts to separate these concerns to a greater degree. The routing/forwarding component of our design is based on the assignment of identifiers to channels — as opposed to nodes, machines, or even programs. This talk will present an overview of our architecture and attempt to point out some of the benefits we see in naming channels instead of nodes.

(The Postmodern Internet Architecture project is a collaboration among the Universities of Kentucky, Maryland and Kansas, funded as part of the NSF FIND program.)

Joint work of: Calvert, Ken; Bhattacharjee, B.; Griffioen, J.; Spring, N.; Sterbenz, J.

On What to Name

Ken Calvert (University of Kentucky, USA)

It is well-known that the present Internet architecture conflates several separable concerns, viz., routing, forwarding and addressing.

We are developing a "Postmodern" Internet architecture that attempts to separate these concerns to a greater degree. The routing/forwarding component of our design is based on the assignment of identifiers to channels — as opposed to nodes, machines, or even programs. This talk will present an overview of our architecture and attempt to point out some of the benefits we see in naming channels instead of nodes.

(The Postmodern Internet Architecture project is a collaboration among the Universities of Kentucky, Maryland and Kansas, funded as part of the NSF FIND program.)

Keywords: Network architecture, routing, forwarding

The Use of Key Based Routing for Next-Generation Internetworks

Thomas Fuhrmann (Universität Karlsruhe, D)

Structured overlay networks such as CAN, Chord, and Pastry have introduced a new kind of addressing scheme that evades many of the problems of the Internet addressing scheme. In particular, such networks offer key based routing which allows an application to work with (almost) arbitrary, location independent, fixed network addresses, for example, hashes of application level keys. Thereby, protocols like DNS or Mobile IP may become obsolete.

Recent work of the authors and independent researchers (cf. Fuhrmann, A Self-Organizing Routing Scheme for Random Networks, Proceedings of IFIP-TC6 Networking Conference '05; and Caesar et al., Virtual Ring Routing: Network Routing Inspired by DHTs, Proceedings of ACM SIGCOMM '06) has demonstrated how the ideas of structured routing overlays can be pushed down into the network layer, thereby potentially replacing IP with a new, key based routing protocol. Moreover, many of the requirements for future networks could be easily fulfilled by simple services that are built on top of a key based routing network (cf. Stoica et al., The Internet indirection infrastructure, SIGCOMM).

We present ongoing work along that direction.

Joint work of: Fuhrmann, Thomas; Kutzner, Kendy

Implementing Name & Address Virtualization

Richard Gold (University College London, GB)

The purpose of this talk is to provide an overview of existing techniques for implementing naming and addressing virtualization and to analyze the trade-offs between them

Keywords: Implementation, naming, addressing

Furthering the HIP Experiment

Tom R. Henderson (Boeing Phantom Works - Seattle, USA)

Many researchers are interested in the potential for the Host Identity Protocol (HIP) and related architectures that separate the role of identity and locator in the Internet. However, while there exists some early implementation experience and small demonstrations of the protocol, experiments to determine the impact of widespread adoption of HIP have been difficult to conduct to date. In this talk, we review the claimed benefits and costs of moving the IP stack towards a HIP architecture, and provide some sample experimental hypotheses to test these claims. We also discuss some of the barriers to widespread experimental adoption of an architectural extension such as, but not limited to, HIP. We call on the HIP experimental and research community to conduct or collaborate on these experiments.

References:

- [1] IRTF Host Identity Protocol research group,
<http://www.irtf.org/charter?gtype=rg&group=hip>
- [2] Email message to HIP research group mailing list:
<https://listserv.cybertrust.com/pipermail/hipsec-rg/2006-October/000360.html>

Keywords: HIP

An Axiomatic Basis for Communication

Martin Karsten (University of Waterloo, CA)

The de-facto service architecture of today's communication networks lacks a well-defined and coherent theoretical foundation. With layering as the only means for functional abstraction, the diversity of current technologies cannot be expressed consistently and analyzed properly. In this paper, we present an axiomatic formulation of fundamental mechanisms in communication networks. In particular, we reconcile the existing but somewhat fuzzy concepts of *naming* and *addressing* and present a consistent set of primitives that are sufficient to compose communication services. The long-term goal of this exercise is to better document, verify, evaluate, and eventually implement network services.

Joint work of: Karsten, Martin; Keshav, S.; Prasad, Sanjiva

Full Paper:

<http://www.cs.uwaterloo.ca/~mkarsten/papers/hotnets2006.html>

Is Mobility an Incompatible Architectural Challenge to IP Routing?

James Kempf (DoCoMo USA Labs - Palo Alto, USA)

The mobility problem in IP networks derives from the basic use of IP addresses as locators and node identifiers. When a node moves into a new subnet, session continuity requires that it not change its IP address; yet in order for it to continue receiving traffic, it must obtain a new address that is topologically correct. The solution space for this problem is well known, and has been explored quite extensively both in standardization and experimental work. One part of the solution space consists of solutions that disaggregate routed traffic at some point and reaggregate it over an overlay network. Mobile IP is an example of solutions where the disaggregation point is located at some arbitrary point in the Internet, NETLMM and GPRS are solutions where the disaggregation point is in the local access network. These solutions either split the locator and identifier functions of the IP address into two addresses, like Mobile IP, or they use routing changes to update the location as the node moves so it can keep its same address.

HIP has explored a different part of the problem space, in which the locator and identifier function of the address are split. The address keeps the locator function and changes while the identifier is a new name space based on the hash of the public key. HIP does not require disaggregation and reaggregation of routed traffic; however, it does require some kind of infrastructure node (rendezvous server) to avoid problems with dropped sessions if two mobile nodes move at the same time.

Most mobility management solutions have poor interactions with routing. The disaggregation/reaggregation doesn't allow traffic to be traffic engineered using mobility as a criteria for traffic management. HIP also has this problem because the routing system doesn't have access to mobility information. A brief excursion into network coding shows that network coding improves forwarding performance in wireless networks with contention-based MACs, but does nothing to improve mobility's impact on routing.

Are there any other parts of the solution space that haven't been explored? Is there some way Internet routing in the large could be modified in some basic way to accommodate mobility, allowing routing decisions to be made based on mobility?

Keywords: Mobility, Mobile IP, HIP, GPRS, NETLMM, network coding

An Indirect Approach to Application Layer Identifiers

Miika Komu (Helsinki University of Technology, FIN)

One deployment obstacle for IPv6 has been that many applications are still not IPv6 enabled. We can argue that the reason for this is that the sockets API did not provide sufficient level of abstraction and indirection when it was originally designed. Later, improvements were made in RFC3493, especially on the client side APIs in terms of address-family-independent nodename and service name translation. On the server side, RFC4038 describes the use of IPv4 mapped IPv6 addresses in dual-stacks servers. In such a case, the server can bind to a single IPv6 socket, which can accept also IPv4 connections.

Further, the so called identifier/locator split, e.g. proposed by Host Identity Protocol (HIP), eases the transition because an IPv6 client and IPv4-only server application can talk to each other. However, new address spaces, such as IPv6 addresses and new identifiers introduced by HIP, maybe just the tip of the iceberg. We question whether the numerous networking applications should be burdened redundantly with the various types, lengths and presentations of identifiers. Instead, we propose some abstraction extensions to the existing sockets API. The abstraction is that we push some of the information from applications downwards into the networking stack and use constant format identifiers in applications. These identifiers act as indirection "handles" to the actual identifiers. We refer to these handles as "endpoint descriptors".

As an additional benefit, the binding from an endpoint descriptor to an endpoint identifier is dynamic, which allows the endpoint identifier to be changed transparently from the application. This may simplify e.g. opportunistic HIP implementations and provide some API flexibility in process migration systems. Alternatively, a single endpoint descriptor can be mapped to several underlying identifiers and might enable new types of group communication mechanisms. For example, it might allow implementing anycast or multicast on top of unicast, or even simultaneous use of several multicast addresses through a single socket.

Despite the endpoint descriptors hide the underlying identifier details, there may be a need for the application to pass an identifier from one host to another (referral). The endpoint descriptor is not suitable for this purpose as it is valid only in the local context of the host. In this case, the networking stack can expose the underlying identifiers to the application that may assist the stack by selecting the appropriate referral.

Finally, the endpoint descriptors might be treated equally to file descriptors. This way, endpoint descriptors inherit the access privilege properties of file descriptors, such read and write permissions for user ids and group ids.

Keywords: HIP, endpoint, descriptor, socket, API

Data-Oriented Network Architecture

Teemu Koponen (HIIT - Helsinki, FIN)

The current Internet architecture is built around a host-to-host communication model, and is perfectly suited for applications, such as file transfer and remote login, that focus on communication between pairs of well-known and stationary hosts. However, the vast majority of Internet usage today is data retrieval, where the user cares about content and is oblivious to its location. For data retrieval, we argue that the current Internet architecture is far from a comfortable fit owing to both naming- and protocol-level issues. In this work, we address these naming and protocol issues and present our resulting design, Data-Oriented Network Architecture (DONA). DONA does not change the underlying point-to-point IP layer, but provides two primitives that sit directly on IP: fetch, by which a client requests a piece of data by its name (not its location), and register, by which a host offers to serve a particular piece of data.

Keywords: Network architecture, flat names, anycast

Joint work of: Koponen, Teemu; Chawla, Mohit; Lakshminarayanan, Karthik; Ramachandran, Anirudh; Tavakoli, Arsalan; Vasu, Atul; Shenker, Scott; Stoica, Ion

Privacy Aspects in Naming and Addressing

Janne Lindqvist (Helsinki University of Technology, FIN)

Names, identifiers or addresses are used in the Internet protocol stack for various purposes. We investigate how the uses of identifiers on different layers affect the privacy properties of protocols. Also, there are vast amounts of legacy authentication systems that do not take the privacy of the users into consideration. Many networks use, for example, MAC address or IP address based authentication, despite of their limited security properties. These authentication systems hinder the possibility to use e.g. pseudorandom MAC and IPv6 addresses for providing unlinkability and location privacy protection. Effectively this means that in some cases, many host-based privacy protection architectures are unusable. We have implemented a host-based privacy management system that addresses these problems. The implementation uses the Host Identity Protocol to provide authenticated and secure handovers for mobile nodes, while using pseudorandom identifiers below the transport layer. Transport and above layers are protected with IPsec ESP.

Keywords: Privacy, security, authentication, pseudonymity, unlinkability, location privacy, Host Identity Protocol

Naming and Addressing in Highly Mobile Ad Hoc Networks

Raquel Morera (Telcordia - Piscataway, USA)

A well-known problem is the coupling of names with the IP addresses used for routing. Often IP addresses identify hosts in the transport and application layers, which limits their ability to support multi-homing, allow Network Address Translation and rapidly reflect topological changes. Ideally, names should identify entities within a purely logical structure, with no correlation to topology, while addresses dynamically reflect the global network topology. Therefore, decoupling names from IP addresses increases flexibility, modularity and scalability. Decoupling names and IP addresses however requires mechanisms to perform the name to IP translation. In dynamic ad hoc networks, current name to IP translation mechanisms such as DNS fail to provide an adequate level of service when nodes, while in reach of each other, are disconnected from the Internet to which DNS servers connect to. M-DNS can partially solve the problem; however this is not a scalable solution. Autoconfigured location servers are a possible solution. In this approach the Logic Name Server (LNS) routes messages to the Location Server associated with a name, which finally performs the name to IP address translation. This approach delays the name to IP address translation to servers that are local to the node. While solutions based on autoconfigured servers are good for a certain degree of mobility, other solutions that violate the name to IP address decoupling principle may be more appropriate for small, very dynamic networks. We will present our current thoughts on this matter.

Joint work of: Morera, Raquel; McAuley, Anthony

See also: A. McAuley, R. Morera, "Name and Address Decoupling in Support of Dynamic Networks", IEEE MILCOM 2002, Anaheim, CA, USA, October 2002

What is an end-point anyway?

Börje Ohlman (Ericsson Research - Stockholm, S)

There is currently a lot of talk about separating identities from locators. The locators are seen as the addresses of attachment points of a layer 3 network. The identities are often referred to as some layer 3.5 objects sitting between the network and the transport layer, e.g. Host identities (HI) in HIP or NodeIDs in the Ambient Networks NodeID proposal. A relevant question is of course what objects these identities really refers to. What is clear is that a HI in HIP is not a unique identifier of a host (a host can host multiple HI) in the same way a node can have multiple NodeIDs. Nor is a HI identifying an endpoint of a communication flow, as multiple flows can be multiplexed on the same HI using ports. You can perhaps say that a HI identifies a logical host within a physical terminal.

Why does this matter? If we would like the network to support session and application mobility it would be nice to have endpoints that identify communication flows rather than parts of boxes. Unfortunately HI does not seem to be the right thing for this.

There are several reasons why this does not work. A major show stopper is that these identities are based on public/private key pairs. And as there, as I understand it, are no good way to move a private key between devices (unless you physically move a SIM card or so) using e.g. HI as identifiers for identifying end-points of flows that you want to move between devices does not seem like a good idea.

But if we look at the problem from the user perspective, moving flows between devices and applications seems like a neat thing to do. Yes, you can do this at the application level today, but only in predefined ways and basically only within one application. It would be nice if I could just take an audio flow and move it to a device displaying text. The problem of routing the flow via a transcoding device in the network I think is better understood than how we provide the user with a "handle" to grab the flow coming out of his audio device and moving it to the text device, in a generic way. That is, without having these two application being tailored to interwork with each other in this way.

This presentation will explore what different type of endpoints there are at different layers in the protocol stack, how they relate to each other and how they can be made useful from an enduser/application developer perspective.

Keywords: Locator-id split, end-points, session mobility, application mobility

Full Paper:

http://www.ambient-networks.org/docs/Host_Identity_Indirection_Infrastructure_Hi3.pdf

Delay-tolerant Networking: Some Aspects concerning Naming and Addressing

Jörg Ott (Helsinki University of Technology, FIN)

Delay-tolerant Networking (DTN) accepts disconnections and (potentially long) delays and the resulting lack of an end-to-end path as fundamental communication characteristics to enable information exchange in challenged environments. This includes (sparse) sensor networks and interplanetary communications but also mobile communications which may or may not use the Internet Protocols as underlying communication substrate. To enable communication in such environments, DTN relies on asynchronous exchange of potentially arbitrary size messages ("bundles"). From a naming and address perspective, DTNs are special as they cannot rely on an omnipresent infrastructure to take care of translations of names and addresses.

The DTN architecture defined in the DTNRG of the IRTF uses URIs as sole means of naming/addressing: for endpoints, applications, and possibly (data)

objects. URIs drawn from the same name space—Endpoint Identifiers (EIDs)—are used to uniquely identify individual entities ("singletons") but may also refer to groups. No restrictions are placed on the URI schemes so that multiple distinguishable namespaces are used. Routing takes place based upon URIs so that potentially all context information (destinations, applications, etc.) are easily available at every node, thus enabling different routing protocols for different name spaces. DTN allows deferring the ultimate resolution to a concrete device address to other nodes (late binding), where intermediate forwarding may occur via default routes or (loose) source routes.

Keywords: Delay-tolerant Networking, Naming, Addressing

Labels and Names taking over Addresses

Christian Tschudin (Universität Basel, CH)

Labels and Names taking over Addresses Christophe Jelger and Christian Tschudin, University of Basel

When the Internet started, addresses ruled the net. Although domain names and path labels were added at a later stage, these "additions" are now taking over many tasks formerly performed by addresses. Ultimately, addresses will become an auxiliary construct rather than a core network ingredient.

In this talk we describe this ongoing transformation, pointing out how formerly orthogonal activities like name lookup and routing are now merging (Caesar et al's Routing on Flat Labels, SIGCOMM06), as well as the trend towards more and more shortlived IP address allocation coupled with repeated address translation. In the ANA project (Autonomic Network Architecture), we anticipate this evolution and envisage a basic forwarding layer with scope-restricted labels at the bottom, and attribute sets as extended names at the top. Inbetween, multiple (and potentially competing) resolution schemes exist inside network instances, as well across network compartments, which integrate resource discovery and search with routing. We report on the current state of the discussions.

Joint work of: Tschudin, Christian; Jelger, Christophe

Labels and Names taking over Addresses

Christian Tschudin (Universität Basel, CH)

When the Internet started, addresses ruled the net. Although domain names and path labels were added at a later stage, these "additions" are now taking over many tasks formerly performed by addresses. Ultimately, addresses will become an auxiliary construct rather than a core network ingredient.

In this talk we describe this ongoing transformation, pointing out how formerly orthogonal activities like name lookup and routing are now merging (Caesar et al's Routing on Flat Labels, SIGCOMM06), as well as the trend towards more and more shortlived IP address allocation coupled with repeated address translation. In the ANA project (Autonomic Network Architecture), we anticipate this evolution and envisage a basic forwarding layer with scope-restricted labels at the bottom, and attribute sets as extended names at the top. Inbetween, multiple (and potentially competing) resolution schemes exist inside network instances, as well across network compartments, which integrate resource discovery and search with routing. We report on the current state of the discussions.

Keywords: Network architecture, name resolution, transient addresses, Selnat, Lunar.

Joint work of: Jelger, Christophe; Tschudin, Christian

Ambient Networks Internetworking Architecture

Rolf Winter (NEC Europe - Heidelberg, D)

The Internet consists of independent networks that belong to different administrative domains and vary in scope from personal area networks, private home networks, corporate networks to ISP and global operator networks. These networks may employ different technologies, communications mediums, addressing realms and may have widely different capabilities. The coming years will add a significant level of dynamic behavior, such as mobile nodes and moving networks, which the Internet must support. At the same time, there is a need to address the increasing levels of harmful traffic and denial-of-service attacks. The existing Internet architecture does not support dynamic behavior or secure communication to a sufficient degree. This talk outlines the Ambient Networks internetworking architecture that allows heterogeneous networks to work together without loss of functionality. Some of techniques employed in this architecture include reliance on cryptographic node identifiers, identity routers and localized addressing realms.

Keywords: Inter-domain routing, Ambient Networks