Black Boxes: Making Ends Meet in Data Driven Networking

Sasu Tarkoma
Helsinki Institute for Information Technology
sasu.tarkoma@hiit.fi

Dirk Trossen
BT Research
dirk.trossen@bt.com

Mikko Särelä
Nomadiclab, Ericsson Research
mikko.sarela@nomadiclab.com

ABSTRACT

In this paper, we consider issues pertaining to end-to-end communications and mobility of network end-points in a network where the forwarding is based on data and interest. We view the network as a collection of black boxes, based on a set of recursive rendezvous functions. The boxes hide their internal topology and expose only labels and interest definitions to outside. We study the concept of completeness of network topology and find it useful for understanding and optimizing publish/subscribe based network architectures. We also show that many issues related to making network end points meet on multiple layers, such as mobility, become much simplified and some even trivial. Towards the end of the paper, we outline challenges and future work in this area.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Distributed networks; C.2.6 [**Internetworking**]: Routers

General Terms

Design, Security

1. INTRODUCTION

The current dominant networking stack, the Internet Protocol suite, suffers from a number of limitations, but works reasonably well for current demands [18]. Current challenges include mobility, efficient global multicast, and multihoming. The network structure also makes hosts and services especially vulnerable to unwanted traffic and Denial-of-Service (DoS) attacks. In addition, reconciliation of end-to-end reachability [10, 4] with other networking requirements that arise from scarcity of IP addresses, and untrustworthy environment (e.g. firewalls, NAT, and other middleboxes [31]) is a much studied, albeit hard to solve problem.

Many of the mobility and security related problems stem from the fact that the IP address specifies both the identity and location of a host. This means that a location management scheme is needed to update any changes in a mobile node's IP address to its peers. In addition, various intermediaries, such as NAT - boxes may filter and modify packets. Private networks created using NAT create yet another problem for the architecture, because typically only outbound connections are allowed.

We sketch a networking architecture based on data and interest driven packet delivery motivated by the publish / subscribe paradigm [14] in Section 2. In our architecture, the network presents itself as a black box to those utilizing it recursively, i.e. a local network is a black box, as is a domain, and the whole inter-domain network itself. A black box hides the internal topology to outsiders utilizing it and exposes only those labels and interest definitions necessary to outside users. Our aim is to define an internetworking architecture that can be run on top of both L2 and L3.

Our architecture has native support for multicast and anycast primitives; indeed, it is built upon them. In this kind of networking architecture many traditional problems, such as mobility and multihoming, become easier to solve, some even trivial. However, the scalability of data and interested-oriented networking is still a major challenge. We assume that edge clients indicate their willingness to receive data by subscribing, and it is the aim of the network to act as a substrate for this data delivery process from distributed data sources.

In our work, we approach the scalability challenge using a recursive set of rendezvous functions. We consider the properties of the network model with more detail in Section 3 and discuss the various problems pertaining to this model in Section 4. Section 5 gives a brief overview of related work and finally we conclude in Section 6.

2. BLACK BOXES

In contrast to traditional end point-oriented network concepts, information centric approach, revolving around particular data and interest for this data, has a significantly different structure and architectural principles. In the following, we outline the principles that lay the foundation for our architectural model. We then investigate the crucial role of rendezvous in this model before going into details on the functions for our architecture.

2.1 Principles of Design

Information is central in the architecture we intend to consider. While IP decomposes an information-centric view of the world by introducing topological constraints on the

MobiArch'08, August 22, 2008, Seattle, Washington, USA.
Copyright 2008 ACM 978-1-60558-178-1/08/08 ...$5.00.

delivery of information (i.e., the subnet-structure of IP networks), we intend to primarily build the system architecture around information from the viewpoints of semantics and scope.

In this information centrism, two concepts are fundamental, namely the recursive nature of *information semantics* and the notion of *information scope* as a concept of reachability.

The concept of recursive information semantics is expressed by rigorously basing our design on the notion of information, starting from very low level forwarding information up to complex semantics with certain contexts. For that, we assume flat labels that are used to identify communication relations between entities that are concerned with a particular class of information. Note that the particular topology for delivering this information is not of our concern (yet). We further assume that any higher layer semantics will be recursively built on top of this low level forwarding semantics. With this, we build the bridge between higher level concepts like ontologies and the lower level information that is used to deliver exactly these higher level concepts without particularly introducing the same topological constraint in reachability that we see in IP.

Given the concept of recursive information semantics, we turn towards the issue of reachability. Information scoping replaces the concept of topologies in IP networks, i.e., it represents the (limitation of) reachability of information. Given the outlined recursion on semantic level, scoping information and therefore the reachability of said information in itself is recursive. Hence, we assume that for certain levels of semantic, certain mechanisms exist for scoping the information related to the given semantic. There exists a variety of mechanisms to scope information on various levels. While we can see a GMPLS service manager as a means to scope forwarding labels in an all-optical network, rendezvous mechanisms are usually used for scoping information on routing level, while search or discovery are applied on higher levels of semantic. With the notion of scope defining the reachability of information, we assume that the architecture is neutral to the semantics and structure of the data sent.

With these concepts of recursive information semantics and information scopes for certain semantic levels, we further assume a communication model for our architecture that places the control over reception of data entirely in the hand of the receiver, provided that permission to receive has been granted. Hence, a model similar to publish/subscribe, in which data is only received once subscribed to it, is the foundation for our system architecture. However, the exact implementation of this communication model will vary depending on the level of semantics we are operating on.

With the above said in mind, we can formulate the following design principles for our system architecture:

P1 The architecture is information centric with recursively layering information semantics (information layering)

P2 Information is scoped with (recursively) layering regions of reachability (information scoping)

P3 Within each scope, the architecture is neutral towards the semantics and structure of data sent (scoped information neutrality)

P4 Information retrieval is receiver-controlled, provided access has been granted

In the remainder of our work, we focus on the impact of these principles on the design of a potential solution and the role of critical components like rendezvous. We restrict our considerations mainly to the equivalent of the internetworking (IP) layer although our principles are formulated as principles that apply for all layers throughout the system.

2.2 A Conceptual View: The Black Box Model

As outlined above, scoping communication is a central element in any networking architecture. In today's IP network, scoping the communication between individual endpoints is achieved by virtue of a network topology based identifier space, which is mapped onto a forwarding state within the network.

In an information-centric network however, an endpoint does not have any means of mapping the network topology other than through the publish and subscribe network primitives. This means that although it is possible to measure network properties edge-to-edge it is not possible to deduce how a data packet is delivered by the network. This is due to the difference of the forwarding paths in information-centric networks. These paths are created on-demand by active producers and consumers of data. There is no forwarding topology without these active entities. This means that the forwarding table is subject to frequent updates. Furthermore, scopes of information can change, e.g., due to the growth of the social network that scopes the reachability of some information. Hence, updates of forwarding paths are expected to be significantly more frequent than for IP forwarding tables (in typical IP networks, forwarding table updates are rather infrequent due to the typically less changing nature of relationships between networks and the relative stability of forwarding paths). The frequency of forwarding path updates is even more increased due to our aim to drop unwanted traffic as close to the source as possible. Hence, packets having labels that are not subscribed to are not forwarded in the network and therefore do not create specific forwarding state in the network. Typically, leases guarantee that any flows that have been setup are also eventually removed. This places a central requirement for the system to be able to support rapid updates in the routing and forwarding tables.

With this in mind, we develop a conceptual view of a potential information-centric system architecture that is based on a black box which is recursively built upon to assemble the higher level complexity in semantics and scope. In a sense, a black box represents a subnetwork which hides its structure from other subnetworks, implementing the recursive nature of reachability scope. As an analogy, we may consider a network that performs Network Address Translation (NAT) at each router thus hiding the network topology, as being such black box based system [12].

With this in mind, we can define another key design principle for this conceptual architecture, in addition to our general principles P1 through P4, namely the black box principle:

P5 Hiding of internal routing and forwarding state simplifies network management and external interfaces are subject to policies and compensation.

An important consideration has to be placed on the point where functions, in particular rendezvous, are implemented in our design. In Clark's refinement to the original E2E principle, he outlines the importance of trustworthiness of execution rather than particular placement of functions in endpoints [8]. This revised E2E principle, called Trust-to-Trust (T2T), is considered in our design through another principle, stating

P6 Rendezvous should happen where the network is trusted to operate correctly in terms of communal, economical, and functional requirements.

We argue that this kind of black box network design has favourable properties, namely support for flexible network deployment, and enhanced security and privacy. On the other hand, additional network management tools are needed. We note that certain diagnostic packets can be implemented that tag through which routers a packet was delivered.

2.3 Rendezvous

We propose *rendezvous* as a central function of a black box. This function is responsible for associating data subscribers and publishers to some *scope* [15], implementing our design principle P2. According to P6, rendezvous should happen where the network is trusted to operate correctly in terms of communal, economical, and functional requirements. It is therefore a policy enforcement point and a mechanism for supporting the freedom of choice for end points. A networking architecture should offer as many *degrees of freedom* as possible. This RP constitutes a well-defined point where tussle is likely to commence [11] due to its importance in defining the information scope in particular scenarios.

Rendezvous functionality has been used in many distributed systems, for example HIP [19, 13], IP multicast (RFC 2362), i3 [29],FARA [9], PASTRY [27], and HERMES [23]. A rendezvous point can be seen as a fixed or non-fixed indirection point for communications.

Figure 1 considers the many faces of rendezvous. In addition to basic packet delivery capability, rendezvous plays a role in higher level activities, including inter-domain policies, and also communal aspects such as observed in services like Facebook and Flickr.

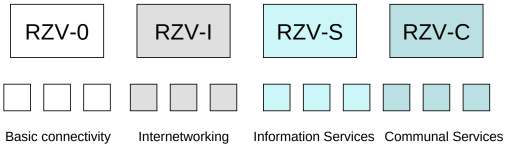

Figure 1: Many faces of rendezvous.

2.4 Rendezvous Function

As stated in Section 2.3, rendezvous is about offering the flexibility of defining (information) reachability. This is achieved through a cascade of rendezvous functions, each for a specific (semantic) abstraction level. Each rendezvous function resolves into an RP, which provides a convenient way to be able to support extensibility of the system, for example in terms of networking policies.

Our definition of the rendezvous function (RZV) is as follows:

Algorithm 1 The RZV function.

```
$RZV(X,D,P)$
    Let set $X \leftarrow resolve(X)$
    Let set $D \leftarrow rewrite(D)$
    $sendToRP(X, D, P, RZV)$
```

Here X is the network address of the next RP, and D is the data label or description which is the basis of routing and forwarding decisions. Both X and D are sets, and P is the packet payload. A packet is a triplet (X, D, P). The *resolve* function resolves to the next rendezvous points. The *rewrite* function performs necessary rewriting of the data D for sending the packet comprising of the triplet (X, D, P). The *sendToRP* primitive is a unicast, anycast, or multicast function that transfers the packet to the next logical RP X, where it is processed by an RZV function.

Since our work assumes data and interest centric operation (according to principle P4), we assume that the data packet description D is a label or a semi-structured descriptor. Our aim is to be able to support various data centric communication mechanisms using the same rendezvous function approach. The system must have a base rendezvous layer, which we denote RZV_0 that ensures basic connectivity for data packets. Then building on top of this, the idea is to introduce label-based forwarding, and ultimately content-based forwarding [5].

2.5 Flat Labels and Scopes

Flat self-certifying [17] labels seem to be a natural choice for a data oriented architecture. Recently, many systems based on globally unique flat namespaces have been proposed, for example UIP [16], TRIAD [7], Nimrod, i3 [29], DOA [2], and ROFL [3].

We propose to utilize flat labels in such a way that each packet identifier has two parts, namely *scope* and *data description*, both being flat labels. With this, the scope allows for the definition of subscription and publication context, effectively implementing our design principle P2. It also enables better aggregation of labels for recursive scoping. As an analogy, a scope could be equivalent to a domain name, and the data description an identifier of a particular web page.

3. NETWORK PROPERTIES

3.1 Completeness

The completeness of data subscriptions and data advertisements is given by Definition 1. This formulation is flexible enough to be useful for various routing protocols. Completeness may be used to characterize the whole routing system. In addition, it may also be used to characterize a part of the routing system, such as a *path*. If a graph, subgraph, or path is not complete, then it is *incomplete* [30].

Definition 1 *A data advertisement A is complete in a network system PS if there does not exist a router r with a data subscription that has not processed a matching A. Similarly, a data subscription S is complete in PS if there does not exist a router r such that r has a data advertisement that matches with S and S is not active on r.*

3.2 Rendezvous and Completeness

We now extend the rendezvous function in such a way that it is easy to check whether or not part of the rendezvous function has successfully completed. This is accomplished by assuming that the rendezvous function maps to a tree of rendezvous points (RPs). This is given by Definition 2. Figure 2 illustrates packet and message exchanges in the rendezvous process.

Definition 2 *A distributed rendezvous function is complete if the signalling path to the RPs identified by the function are complete according to Definition 1.*

Our definition of the rendezvous function with completeness checking (RZVC) is as follows:

Algorithm 2 The *RZVC* function.

```
$RZVC(X,D,P)$
  1. Let set $X \leftarrow resolve(X)$
  2. Let set $D \leftarrow rewrite(D)$
  3. Let set $I \leftarrow incomplete(X,D)$
  4. $\forall x \in I$: $sendToRP(\{x\}, D, P, RZVC)$
```

If a part of the rendezvous process is complete with respect to the input parameters (and any modified parameters during the process), it is not necessary to continue the rendezvous process further.

We observe that the changes in complete parts do not require any interaction with the client, the incomplete parts are the most interesting from the network topology viewpoint.

Moreover, the number of nested RPs gives a natural metric for the topology update process for a particular data item D. This number, denoted by $h(D)$, can be used to gauge the level of completeness of a logical space of the information network. The localhost is defined as $h(0)$ and the local area network as $h(1)$. In addition to completeness, it is possible to determine the Round-Trip-Time (RTT) of $h(x)$, for some x. This can be seen as a basis for network diagnostic functions.

Algorithm 3 The *CMPL* function that checks completeness.

```
$CMPL(X,D)$
  1. Let set $X \leftarrow resolve(X)$
  2. Let set $D \leftarrow rewrite(D)$
  3. Let set $I$ be initially empty
  4. Forall $x \in X$: If $D$ is active on $x$ and a completeness ack
     has not been received then $I \leftarrow \{x\}$
  5. If $I \neq \oslash$ then return INCOMPLETE
  6. return COMPLETE
```

3.3 Impact of Mobility

Mobility support can be realized on different levels in the network protocol stack and for different types of endpoints. In this section, we consider the physical mobility of publishers, subscribers, and networks. Such mobility support, implemented in a handoff mechanism, needs to be supported by the network. This requires that the routing and forwarding topology is updated accordingly during and after handoffs.

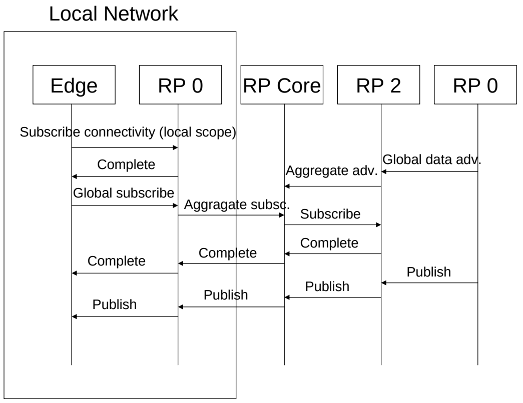

Figure 2: Example interactions.

A simple way to cope with a handoff on the subscriber and publisher level is for subscribers to re-subscribe to those publications they are subscribing to, and for publishers to re-establish their publications at the new location. The rendezvous service is responsible for maintaining the proper set of rendezvous points and multicast forests.

A similar mechanism can be used for router and network mobility, although in this case function of re-subscription and re-publishing becomes a function of aggregating a set of publications and subscriptions. With this, the mobile router that receives a subscription from its mobile network, becomes a rendezvous point for the nodes it serves. When the mobile router moves, it simply updates the publish and subscribe requests to the rendezvous system, which updates the forwarding tables in the network to deliver data to the new location.

To optimize this, the mobile router can attach to a static anchor point in the fixed network and route all subscriptions and publications through it. When it moves, it first re-subscribes and publishes to that anchor point to allow the data flow to continue immediately. It then finds another anchor point that is closer to it, and asks it to be its anchor point. The anchor points then do ordinary mobility signaling in the network. This is in principle similar to Hierarchical Mobile IPv6 Mobility [6]. A similar rendezvous-point-based mobility management scheme has been proposed in conjunction with the HIP architecture [13].

The main problem then becomes how to optimize the rendezvous system for this. We observe that if the data labels are already subscribed or published in the new network, the handoff is complete.

4. DISCUSSION

The novelty of the proposal stems from the combination of black box based network model and the use of recursive rendezvous processes on multiple logical layers in combination with publish/subscribe primitives. The aim of this mechanism is to better support current data intensive network applications, such as YouTube, Flickr, many mashups, and many web pages that are frequently updated. The approach is motivated by the need to support many different kinds of network technologies, including DHT-based overlays and new kinds of networking solutions, such as all-optical GMPLS core networks that are based on label switching. The label-based approach might yield performance improvements also in broadcast networks such as wireless access networks.

There are currently on the order of 10^5 IP prefixes and 10^4 autonomous system (AS) numbers. Each black box can be seen as an AS, and scopes can be seen as equivalent to IP prefixes albeit with more flexibility due to open and multicast nature. Therefore we would need to store on the order of 10^9 entries in the backbone to be able to connect data subscribers and publishers between networks. Given that there is a large identifier space associated with each scope, say 10^{10}, simply using the scope to subscribe and publish will result in many false positives. Further work is needed to understand the scalability properties of data-centric networks with recursive rendezvous. We note that various probabilistic techniques can be used to aggregate flat labels, which typically increase false positive rate. A number of different routing algorithms for scale-free topologies have been developed [21]. It is an open issue how compact routing technique can be combined with the rendezvous-based approach presented in this paper.

The proposed rendezvous-based mechanism introduces implicit policy-based routing on multiple layers. Related work has demonstrated that policy-based routing does not seem to exacerbate the maximum congestion when compared to shortest-path routing [1].

The black box model of the network enhances privacy and accountability simultaneously. It enhances privacy of network actors by making the internal structure of a network private, i.e. a black box, and by making routing based on the data identifiers rather than destination identifiers. At the same time, explicit rendezvous system that matches the wishes of publishers and subscribers is a natural control point for controlling access to resources and, if so wanted, creating a transaction that the subscriber or publisher cannot repudiate afterwards. Such transaction protocols are out of the scope of this work and are left for future work.

5. RELATED WORK

The presented work has been influenced by a number of existing systems that extend core network features by introducing more indirection in the way data is delivered. Most research on data and interest oriented networking has focused on application layer overlay networks. In this context, the Siena system can be considered to be a classic example of a distributed content-based routing system that was implemented in the application layer [26].

The Session Initiation Protocol (SIP) [25, 28] implements a session-based communication mechanism with a server-based rendezvous mechanism, i.e., the so-called SIP proxy maps the URI-based names of participants onto their IP addresses from which they have registered. Another service model is provided through the SIP event framework [24], implementing a generic publish-subscribe framework on top of the SIP rendezvous mechanism. Extensions to the server-based rendezvous mechanism in SIP have been widely investigated. For instance, the IETF is currently investigating P2P-based rendezvous replacements. Other work investigated the usage of local multicast as a rendezvous mechanism [22].

The TRIAD architecture considered how to use NATs in the network architecture [7]. The main idea is that NATs are too valuable not to be included in the future Internet architecture. FARA (Forwarding directive, Association, and Rendezvous Architecture) defines an abstract naming model that decouples end-systems from their network addresses. They focus on abstract definition of the architecture and do not rely on a global namespace, but rather a rendezvous function that routes the first packet form a source to a destination [9].

A discussion item within the IRTF End-to-End mailing list, originally sent out by Jon Crowcroft [12], proposed a 100% NAT solution for IP networks, which would result in a flat label architecture on the routing and forwarding level that is similar to the one discussed in this paper. Although being an interesting thought experiment, the proposal did not address the problem from the viewpoint of changing the underlying communication paradigm rather than avoiding denial-of-service attacks; a claim that was not resolved in the discussion following the proposal. Furthermore, the question around solving the rendezvous in such system remained unanswered (originally DHT was proposed).

DONA (Data-Oriented Network Architecture, ICSI) aims to introduce data-centric operations to the networking architecture. DONA inserts a data-handling shim layer right above the network layer and resolves names by directly routing to data. This does not involve DNS and a lookup, but just routing on names [20].

6. CONCLUSIONS

In this paper, we considered issues pertaining to end-to-end communications and mobility of network end-points in a publish/subscribe motivated network architecture. This architecture is based on data-driven communications, in which the network connects publishers and subscribers of data. The connection is done using flat self-certifying labels with scoping, and using a recursive rendezvous function. The aim of the rendezvous mechanism is to support policy-based routing, and meet the demands of various network applications. We considered how the rendezvous functions can be used to realize network management operations and to ensure that topology updates are performed properly. We also briefly discussed scalability issues pertaining to data-driven architecture. The purpose of this paper is to consider rendezvous as a general networking primitive. Future work includes investigation of scalability and performance issues pertaining to black boxed and rendezvous based networking.

Acknowledgments

This work has been funded by the EU FP7 PSIRP project (Contract INFSO-ICT-216173). We would like to thank PSIRP partners for insightful discussions.

7. REFERENCES

[1] A. Akella, S. Chawla, A. Kannan, and S. Seshan. On the scaling of congestion in the internet graph. *SIGCOMM Comput. Commun. Rev.*, 34(3):43–56, 2004.

[2] H. Balakrishnan, K. Lakshminarayanan, S. Ratnasamy, S. Shenker, I. Stoica, and M. Walfish. A layered naming architecture for the internet. In *ACM SIGCOMM 2004, Portland, OR*, Sept. 2004.

[3] M. Caesar, T. Condie, J. Kannan, K. Lakshminarayanan, and I. Stoica. Rofl: routing on

flat labels. In L. Rizzo, T. Anderson, and N. McKeown, editors, *SIGCOMM*, pages 363–374. ACM, 2006.

[4] B. Carpenter. Architectural Principles of the Internet. Internet Engineering Task Force: RFC 1958, June 1996.

[5] A. Carzaniga and A. L. Wolf. Forwarding in a content-based network. In *Proceedings of ACM SIGCOMM 2003*, pages 163–174, Karlsruhe, Germany, Aug. 2003.

[6] C. Castelluccia. HMIPv6: A hierarchical mobile IPv6 proposal. *ACM SIGMOBILE Mobile Computing and Communications Review*, 4(1):48–59, 2000.

[7] D. R. Cheriton and M. Gritter. TRIAD: A New Next-Generation Internet Architecture. `http://www-dsg.stanford.edu/triad/`, July 2000.

[8] D. Clark and M. Blumenthal. The end-to-end argument and application design: the role of trust. In *Proceedings of the Conference on Communication, Information and Internet Policy (TPRC)*, Sept. 2007.

[9] D. Clark, R. Braden, A. Falk, and V. Pingali. FARA: Reorganizing the Addressing Architecture. *ACM SIGCOMM Computer Communication Review*, pages 313–321, 2003.

[10] D. D. Clark. The design philosophy of the DARPA internet protocols. In *SIGCOMM*, pages 106–114, Stanford, CA, Aug. 1988. ACM.

[11] D. D. Clark, J. Wroclawski, K. R. Sollins, and R. Braden. Tussle in cyberspace: defining tomorrow's internet. In *SIGCOMM '02: Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 347–356, New York, NY, USA, 2002. ACM.

[12] J. Crowcroft. *100% NAT - a DoS proof Internet.* IRTF, Feb. 2006. IRTF End-to-End (E2E) mailing list.

[13] L. Eggert and J. Laganier. *Host Identity Protocol (HIP) Rendezvous Mechanisms.* IETF, Oct. 2004. Internet draft, work in progress.

[14] P. T. Eugster, P. A. Felber, R. Guerraoui, and A.-M. Kermarrec. The many faces of publish/subscribe. *ACM Comput. Surv.*, 35(2):114–131, 2003.

[15] L. Fiege, A. Zeidler, A. P. Buchmann, R. Kilian-Kehr, and G. Mühl. Security aspects in publish/subscribe systems. In *Third Intl. Workshop on Distributed Event-based Systems (DEBS'04)*, Edinburgh, Scotland, UK, May 2004.

[16] B. Ford. Unmanaged internet protocol: taming the edge network management crisis. *SIGCOMM Comput. Commun. Rev.*, 34(1):93–98, 2004.

[17] M. Girault. Self-Certified Public Keys. *Advances in Cryptology (EUROCRYPT)*, pages 490–497, 1991.

[18] M. Handley. Why the internet only just works. *BT Technology Journal*, 24(3):119–129, July 2006.

[19] M. Komu, S. Tarkoma, J. Kangasharju, and A. Gurtov. Applying a cryptographic namespace to applications. In *DIN '05: Proceedings of the 1st ACM workshop on Dynamic interconnection of networks*, pages 23–27, New York, NY, USA, 2005. ACM Press.

[20] T. Koponen, M. Chawla, B.-G. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker, and I. Stoica. A data-oriented (and beyond) network architecture. In J. Murai and K. Cho, editors, *SIGCOMM*, pages 181–192. ACM, 2007.

[21] D. Krioukov, k c claffy, K. Fall, and A. Brady. On compact routing for the internet. *SIGCOMM Comput. Commun. Rev.*, 37(3):41–52, 2007.

[22] S. Leggio, J. Manner, A. Hulkkonen, and K. Raatikainen. Session Initiation Protocol Deployment in Ad-Hoc Networks: a Decentralized Approach. In *Proceedings of the International Workshop on Wireless Ad-Hoc Networks (IWWAN2005)*, London, UK, May 2005.

[23] P. R. Pietzuch. *Hermes: A Scalable Event-Based Middleware.* PhD thesis, Computer Laboratory, Queens' College, University of Cambridge, Feb. 2004.

[24] J. Rosenberg. A session initiation protocol (sip) event package for registrations, 2004.

[25] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. Johnston, J. Peterson, R. Sparks, M. Handley, and E. Schooler. Sip: Session initiation protocol, 2002. [RFC 3261].

[26] D. Rosenblum. A tour of siena, an interoperability infrastructure for internet-scale distributed architectures. In *In Ground System Architectures Workshop*, Feb. 2001.

[27] A. I. T. Rowstron and P. Druschel. Pastry: Scalable, decentralized object location, and routing for large-scale peer-to-peer systems. In *Middleware '01*, pages 329–350, London, UK, 2001. Springer-Verlag.

[28] H. Schulzrinne and E. Wedlund. Application-layer mobility using sip. *SIGMOBILE Mob. Comput. Commun. Rev.*, 4(3):47–57, 2000.

[29] I. Stoica, D. Adkins, S. Zhuang, S. Shenker, and S. Surana. Internet Indirection Infrastructure. In *Proceedings of ACM SIGCOMM*, August 2002.

[30] S. Tarkoma and J. Kangasharju. On the cost and safety of handoffs in content-based routing systems. *Computer Networks*, 51(6), Apr. 2007.

[31] M. Walfish, J. Stribling, M. Krohn, H. Balakrishnan, R. Morris, and S. Shenker. Middleboxes no longer considered harmful. *USENIX Symposium on Operating Systems Design and Implementation*, 2004.