Pervasive Persistent Identification for Information Centric Networking

Karen R. Sollins
MIT CSAIL
32 Vassar St.
Cambridge, MA
+1 617 253 6006
sollins@csail.mit.edu

ABSTRACT

Identification is central to information or content centric networking, in order to enable referencing and access to the information objects. In this work we focus on identifiers and the identification system as a target of a design process, because without careful attention to the identifiers themselves and the approaches to selecting, assigning and using them, they may not meet their design goals. The paper begins with an examination of key issues central to the design of an identification system. With those in mind, we discuss the objectives of pervasiveness and persistence as requirements for identification in an information centric networking (ICN) approach. These lead to a set of design four goals: longevity, scalability, evolvability and security. We apply two key design principles, layering and modularity, to derive our design for the Pervasive Persistent Identification System or PPInS for information centric networking. The contributions of this work include (1) the design issues for identification systems, (2) analysis of goals and key design criteria for identification in an ICN approach, and (3) a principled design of PPInS.

Categories and Subject Descriptors

H.3.4 [**Systems and Software**]: Information Networks, C.2.1 [**Network Architecture and Design]**: Network Communications, C.2.2 [**Network protocols]**.

General Terms

Design

Keywords

Persistent pervasive identification, information centric networking

1 INTRODUCTION

At the core of an information or content centric networking design are identifiers, because, without some form of identification, referring to and accessing information is impossible. We focus on identifiers in this work as a target of the design process, because, without careful attention to the identifiers themselves or the approaches to selecting, assigning and using them, they may not meet a target set of goals. In turn, this suggests that we must begin any approach to using identifiers by clarifying the goals and any constraints on an identification scheme.

For demonstration we begin with a simplified example about medical information. Our medical record is a complex object, containing a history of illnesses, lab tests for which the original data is maintained by the lab, reports from a variety of doctors in a variety of practices, pointers to the medical records of parents and so forth, with constraints placed on it by the patients, their families and medical care staff, as well as regulators and the insurance industry. Different parts of it are the responsibility of different authorities, which may change with time as the patient moves, yet the record should have continuity. In addition, a medical record should survive not only for the lifetime of the patient, but perhaps the lifetimes of descendants. In this rather simplistic extended example we see requirements for identification assignment, management, access control, and reachability across space, time (e.g. persistence), scale, and changing conditions. In contrast, in other examples information may be both highly dynamic and continually of interest, such as quickly changing news stories, or the targets of flash mobs. The demands on an information identification system under these circumstances may be very different from the medical record situation.

We raise these issues here to illuminate the fact that goals for identification schemes can make the difference between an effective and useless or counterproductive scheme. For a global information layer as proposed in the current ICN projects, the examples only demonstrate a small part of the breadth of applicability. In order for these systems to be truly effective, identification of the information must also be designed to provide persistence and pervasiveness. We propose here an approach that takes advantage two design principles, layering and modularity. By using these approaches we demonstrate a reasoned and effective design of the Pervasive Persistent Identification[1] System or PPInS that meets our goals for identification in the context of ICNs. Although this work is proceeding in the context of PURSUIT[2], it is also well suited to the other current approaches to ICN.

In the remainder of the paper, we will begin with a brief analysis of the issues that are central to designing identification systems

ICN'12, August 17, 2012, Helsinki, Finland.

[1] We use the terms "identifier" and "identification" where others might use "names" and "naming" to avoid confusion.

[2] Throughout this paper we will refer to the combined PSIRP/PURSUIT project as "PURSUIT".

generally, as background to the later design. From there, we begin with our overall objectives of persistence and pervasiveness, which lead to our specific system goals. The combination of the goals and our design principles of layering and modularity direct the design the system as discussed in Section 3. Section 4 briefly reviews related research. The paper concludes with our contributions.

2 SYSTEMATIC THINKING ABOUT IDENTIFICATION SYSTEMS

Above we reflected on the fact that there are a number of different kinds of choices with respect to designing identification systems. In this section we organize the kinds of design choices, in order that later we can reflect on which ones are important in which ways. In other words, with respect to some designs a single factor in the design space may be critical and with respect to others the factor may be irrelevant or arbitrary. To make these choices in our design, we must understand what they are.

We begin by considering the basic functions or purposes of identifiers: equality, access or reference, and meaning or mnemonics. The *equality* function answers the question of whether two objects are the same or different under a particular definition and set of conditions for equality. The *access* function produces some form of representation or path to an instance of the object. *Meaning,* as reflected in an identifier, produces some meta-information about the object derived solely from the identifier itself. The assignment of identifiers to objects enables one or more of these functions to be performed on the identifiers, depending on the implementation of the identification system The basic functions of an identification system can only be implemented in the context of a variety of design choices for its structure. We can define three key structural aspects of such a system: (1) the nature of the idspace itself, (2) the approach to assignment of identifiers to objects, and (3) the resolution of identifiers in the system. Although security is also critical, our design supports factoring it out. This will discussed under Goals in Section 3.1

2.1 The structure of the idspace

We begin with the structure or nature of the idspace itself. There are two there are two key facets to the structure, the scope or breadth of applicability of the space and the syntax or structure of the ids that can be defined in it. With respect to scope, if an idspace is ***global*,** it provides a single, shared pool of ids from which to select. In contrast, if an idspace is ***local***, it provides ids that may also be provided by another local idspace. The choice between local and global may be reflected in questions such as whether there is a need for a single overarching domain of discourse, whether selection and assignment of ids is centralized or distributed, and the degree to which operating in a sparsely or densely populated idspace is important.

There are several aspects to the syntax of the ids in an idspace: **size**, **internal structure**, and **character set**. The **size** of ids may be *fixed* or *variable* and within this latter category may be either *bounded* or *unbounded*. Separately, the **internal structure** of ids generally falls into one of three categories. At one extreme there is a *flat idspace*, in which ids have no structure and are simply selected from a large space. At the other extreme are very organized idspaces, usually *hierarchies*, in which an identifier is composed of elements selected from successive subspaces; such idspaces may be singly or multiply rooted.[3] In the middle there are *composite identifiers***,** which are less organized. Attribute based identification is such an example. Finally, the **character set** available to an idspace often is important, again because it may simplify processing or make human expression in the ids easier. The choices with respect to syntax are likely to be driven by the goals or requirements placed on the basic functions.

2.2 Identifier assignment

We next turn to the business of assignment of ids to objects and we observe three issues: who has authority to assign ids, the persistence of those assignments, and the uniqueness of the assignments. Again decisions with respect to each will depend on the design goals for the three basic functions. To begin, with respect to authority for assignment of ids, the two key choices are the degree of centralization in selecting an id, and whether the assignment authority is reflected in the ids themselves or not. The DNS and telephone number allocation/assignment systems demonstrate different points on this spectrum. The allocation authority is explicit DNS names, and not present in phone numbers.

The second design choice is the question of the persistence or lifetime of those assignments, *limited* or *permanent*. DNS assignment is time limited, and MAC address assignment permanent.

Our third issue with respect to id assignment is uniqueness. In some idspaces, it is important that each object have exactly one id. In others, it may be valuable or simply unimportant if an object has more than one id, or if more than one object share an id.

2.3 Resolution of identifiers

Resolution is the action of translating an id into something else. With respect to resolution of ids, we find four key design questions worth noting here:

Scope of resolution: This issue has to do with whether the resolution of a particular id is the same from any location (universal) or different from different locations (local). For example, if a service is replicated in several places, the id may be resolved to different instances for performance.

Existence and persistence of resolution: The issue here is whether or not it is required that an id be resolvable at all times. If an id may be viewed as a placeholder for an entity that does not exist at all times, the identifier may not be resolvable at all times.

Timing of resolution: One must also decide whether resolution occurs early, late, or piecemeal. As an example, CCN ([9], [17])[4] performs repeated partial resolutions.

Target domain for translation: Based on two distinct kinds of target domains, we identify two translation operations, ***aliasing*** and ***resolution***. Aliasing is translating an identifier into another identifier from the same idspace and abstraction space. Thus, a nickname for a human might be considered an alias. Providing aliasing usually is optional and therefore considered at design time. In contrast, resolution, the operation of translating from one idspace to another, is a central to an identification service.

The reason for this apparent diversion into design alternatives is to illuminate the fact that there are many design choices for identification schemes. In any id scheme, those choices should reflect the goals and other design criteria for the system in order to achieve

[3] Hierarchies are often simpler for humans to understand, but can be limiting, if it becomes necessary to move or rename nodes in the hierarchy.

[4] Throughout this paper we will refer to the combined CCN and NDN projects as "CCN", for simplicity.

both function and performance as required. Hence the identification system design we present here is driven by our objectives of supporting persistence and pervasiveness

3 GOALS AND DESIGN OF PPInS

For any system design it is necessary to clarify the goals of the system to both identify constraints and also expose where they do not exist. By designing to meet the goals, the design is constrained, suggesting that minimizing goals may maximize flexibility and evolution. In Section 3.1 we refine our objectives into four goals and follow that in Section 3.2 with their design implications.

3.1 Goals of PPInS

Throughout the Internet we find increasing agreement on the centrality of information objects, including accessing them, comparing them, sharing them, producing them, all at global scale. Increasing investment in the creation and management of information suggests an increased value in persistence and pervasiveness. These objectives for identification for ICN architectures lead us to four central goals that drive our system design: longevity, scalability, extensibility, and security.

Longevity: The goal of longevity is to make it possible for a long-lived information object to have a distinguishing identifier assigned to it for that same lifetime. In other words, in some deep sense the "meaning" of the identifier remains constant. A valuable identification system will support testing for that kind of equality through testing for the equality of identifiers. Thus, we can define equality of objects by equality of ids. One feasible approach is to define more manageable sub-idspaces for which equality is defined for both underlying objects and among the ids themselves.

Scalability: Although scalability is assumed in many information centric architectures, from the Web to the growing list of self—identified ICN projects, it is important to be explicit about it. We have gone from the scalability of the DNS to the Web to the ICN projects such as PURSUIT and CCN. PURSUIT ([20],[27]), with an identifier per object for each scope within which it is published and CCN, which identifies each packet or element of delivery (fragment of an object) distinctly, make significantly larger scaling demands than the Web.

Extensibility: The goal of extensibility is what makes this proposed architecture unique in the ICN community. If id equality reflects object equality and the kinds of objects grow with time, our system must support new definitions of equality. Hence we set a goal of supporting extensibility of identification schemes within a single framework. Each such scheme will need to define for itself an identifier syntax, the generation of ids, equality of ids within the syntax and assignment of ids, equality of objects assigned ids, resolution schemes for mapping an identifier onto some means of access to the object and semantics exposed in the ids. By incorporating multiple approaches into the framework, different design and efficiency choices can be tuned to more particular needs. The intention of this goal of extensibility is that within a single framework, all these existing and new choices should be able to co-exist.

Security: Without security a system such as this will not be viable. Each of the lead ICN approaches includes an approach to security, over which we can lay a PPInS end-to-end approach. That said, we are proposing a framework within which different id schemes can co-exist. Each one of those in turn will be the point at which a decision should be made about the degree and nature of security it provides. Thus, each one can provide its own definition of the trustworthiness of its security approach.

3.2 Design of PPInS

We describe the design of PPInS in the following three subsections on the overall design and organizing principles, the Pervasive Persistent Object Id (PPOID) design, and the PPInS layer structure.

3.2.1 Design principles: layering and modularity

A number of designs for PPInS could meet the goals above. Ours is intended to separate abstraction layers and cleanly enable research on aspects of the architectural framework. Below the PPInS layer sits an ICN that supports a pull model of delivery of a set of bits that represent the intended object. In the set of ICN projects, there are different approaches taken here, but each one is intended to deliver bits when requested.

We highlight the design in Figure 1. At the top in the user and application layer, we postulate a diverse and evolving set of mechanisms for supporting human or application friendly identification, such as global search tools, an attribute based identification, or long-lived references embedded in books, journals, and so forth. In turn, within the PPInS layer, a PPInS identifier is mapped into a particular idspace, and then resolved within that idspace. This allows for multiple resolutions services for any of the individual idspaces, as well as idspace-specific security solutions. This resolution maps to a supporting ICN identifier. In the case of PURSUIT this will be one or more fully qualified RIDs, because an object can be published in more than one scope. In CCN it will be a CCN hierarchical name, which may identify a single CCN

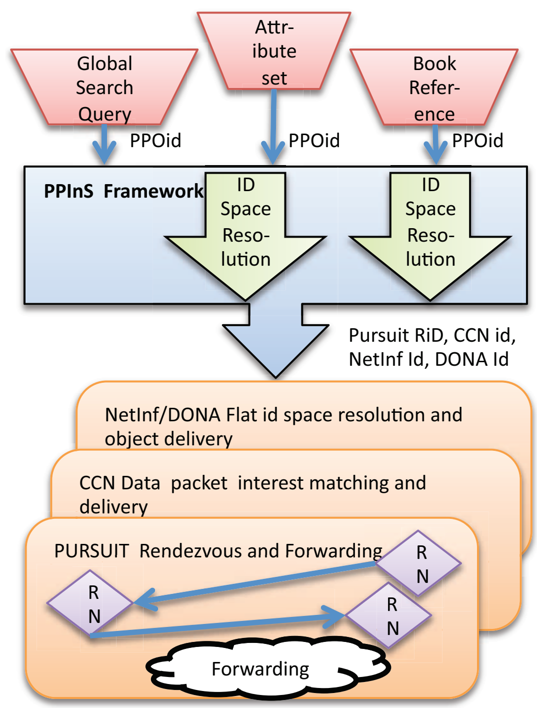

Figure 1: Identification architecture: User/application layer, PPInS layer, and supporting ICN layer options

data packet or a set of them that share their higher order parts of the id hierarchy. In NetInf or DONA the identifier will be for a single object. The point is that resolution in the PPInS layer will result in an id that in turn will be used by the supporting ICN to deliver the requested information to the requestor, without the supporting ICN id system having to support our goals of true persistence, scalability, evolvability of the id spaces or more than the core security that is currently proposed by these supporting ICNs. Our design for the PPInS layer is guided by our two design principles of layering and modularity.

Layering: The layering in this design allows for isolation of functionality that has distinct identification requirements into the distinct layers. Our design approach reflects three layers, because the identification requirements at each layer are distinct. User-friendly identification at the top should be simple, expressive, but not necessarily global, persistent, nor unique, in contrast with what is required of the PPInS layer. Below that identifiers must be designed for realtime resolution and delivery. With these different goals, we have argued ourselves into a three-layer design.

Modularity: Our organizing principle of modularity allows for the provision of independent parallel functional elements within a layer, each of which provides essentially the same functionality, but with different design constraints. These may have to do with scale, flexibility, efficiency, locally shared definitions of equality or any number of other design criteria. At the PPInS layer, different idspaces can be based on different design choices, as well as supported by more than one resolution scheme. The PPInS layer translates from PPInS ids (PPOIDs) to PURSUIT RIDs, hierarchical CCN Ids, or the flat names of NetInf or DONA. The PPInS layer is modularized, containing a dispatch capability that determines the PPInS idspace of the PPOID and dispatches to the Subspace PPOID Resolution Service (S-PPOID RS) of that idspace.

In the case where PURSUIT is the supporting ICN, PURSUIT *scopes* provide another form of modularity. The act of publishing a set of RIDs in a scope can be reflective of a declaration of association among the set of objects, defined by being in the same scope. In addition, each scope supports a dissemination methodology shared by all the objects published in it. So, the modularity provided by scopes may be either declarative or methodological and is reflective of a commonality within the scope and differentiation between scopes, as with the commonality supported within an idspace and the differentiation between idspaces.

3.2.2 At the center: the PPOID

At the heart of the PPInS approach is the persistent, pervasive identifier. Because standards track Uniform Resource Names (URNs) ([13], [22], [24]) match our requirements, we adopt that approach, depicted in Figure 2. These ids or "names" were originally designed as part of a suite of identifiers that included URLs, so the prefix of "urn:" was chosen, followed by the namespace id (NID) and the namespace specific string (NSS) in the NID space.

As with URNs, there must be agreement on the *meanings* of the NID names. This therefore is a small, but global idspace requiring global resolution. For organizational reasons, assignment control of these name components for URNs was given to the Internet Assigned Numbers Authority, IANA. It is required that there be at least one approach to mapping from the NID identifier to a service that can resolve the NSS into underlying ICN ids, but as in the URN case, an evolving set is possible (and perhaps likely).

3.2.3 Design of the PPInS layer

As discussed above, the PPInS layer is composed of a two-part structure. Overarching the layer is the global idspace service or G-PPOID RS. That service hands off resolution of the PPOID to the appropriate distinct sub idspace resolution service, or S-PPOID RS.

The Global PPOID Resolution Service (G-PPOID RS)

The G-PPOID RS must be both reachable from anywhere and efficient. There are two key functions that must occur within this service: creation of new schemes (i.e. assignment of scheme ids) and resolution from a scheme id to a service that embodies that scheme, specifically an S-PPOID RS. As with URNs, the expectation is that creation of schemes will be infrequent and low volume. In order to insure that such a service is viable, IANA or some other similar organization will vet them and verify that there is at least one viable S-PPOID RS deployment available. This organization will not be responsible for guaranteeing that such a service exists at all times, but will provide some degree of confidence in it. If a particular appropriate S-PPOID RS is inoperable, PPOIDs from that idspace may be unresolvable and hence inaccessible, but this will not affect the overall system.

It may be valuable for an S-PPOID idspace to be served by more than one, but a small number of resolution services, in order to support geographic variation, evolution, different performance criteria, etc., suggesting the need for reasoned choice here. Because of both the relatively small number of idspace ids and low turn over, this service is an excellent candidate for widespread replication of its underlying information. The publish/subscribe paradigm allows for efficient and timely updates of new information.

A Subspace PPOID Resolution Service (S-PPOID RS)

Each S-PPOID RS will operate independently. Each one's task is to translate a PPOID of its own idspace into ids in the supporting

"urn:"	<NID> ":"	<NSS>

Figure 2: The PPOID as a URN

ICN; in the case of PURSUIT, this may be a set of RIDs, because an object might be published in more than one scope. We expect to see wide variation in the designs of these services, to meet a variety of performance, policy, and domain requirements. That said, there are some common design requirements. First, because there is no concept of location at this level of the overall ICN, each S-PPOID RS must be prepared to serve the needs of clients that may be anywhere, and where their location is not knowable. Hence, the span of each service must be global. Second, the set of PPOIDs that each S-PPOID RS is supporting is likely to be significantly more dynamic and larger than the G-PPOID RS. In most of these ICNs, the identifiers are announcements of

In addition, we must remember that in most of these ICN systems, the publication or advertisement of an object by means of an ID, may be separate from whether or not the underlying object exists. an RID is an announcement of publication (in the present). Although a PPOID and the underlying object may exist, there may be times when S-PPOID RS may not be able to map the PPOID to a currently viable ICN id. It is worth noting here that the intention is that the PPOID of an object be its identifier for its lifetime. Separately, within the ICN supporting layer, it may or may not have a currently assigned identifier. For example, PURSUIT RIDs announce current accessibility, and can be withdrawn, without the PPOID becoming invalid. Thus, the S-PPOID may not be able to translate a valid PPOID from its idspace at all times, if there is no currently valid RID published for the object. This is completely acceptable behavior to be supported by S-PPOID RSs. This partitioning of resolution also allows for tunable, localized security.

To return to our original design goals for PPInS, the primary objectives of persistence and pervasiveness have been the key drivers of this design. PPOIDS provide persistent ids and the tiered resolution structure provides pervasive, scalable, and modular resolution.

4 RELATED WORK

Because our focus is the design an identification system for a novel, ICN architecture, we review two topics here, the requirements and choices in identification, and the design choices in such information-based systems. [5]

Identification requirements and choices: There is a long, but somewhat intermittent history of designing name or idspaces to meet specific requirements. Very briefly, the DNS ([15], [16]) and the World Wide Web built its URL scheme ([3]) on top of it were designed to be global in reach, hierarchical in assignment, management and resolution, and long-lived, but not persistent. Later, Sollins' Information Mesh project ([23]) and related work on URNs in the IETF made persistence a primary goal, as did the Handle System ([26]), based on a global handle registry.

We now turn to contemporary projects. The NetInf project ([2], [18]) and the Data-Oriented Network and Architecture (DONA) ([11]) provide both uniform persistence and security through slightly different flat designs. The uniformity might allow for, but does not actively support, the variation in approaches we seek. Of the NSF FIA projects, CCN is most closely related to the work proposed here, with three significant differences. First their hierarchical names are human-meaningful, a challenge in a constantly changing environment, because the changes must be reflected in the hierarchy; the identifiers can be either persistent or reflective of change, but not both. Second, in CCN the named entities are packets, the units of transmission for an object, not the object itself. Third, CCN takes on a more significant challenge with respect to security, specifically validity, provenance and relevance, than the other ICN approaches in the work of Smetters and Jacobson ([21]). Last, we reflect on the PURSUIT. At its core is an identifier (a Rendezvous identifier or RID) to make rendezvous and forwarding scalable and efficient. An RID is published in a scope, within which it is unique and which itself is published, so rendezvous occurs through a DAG of scopes. Thus, the identifiers are intentionally not required to be unique across all scopes nor necessarily persistent, although they do, as with NetInf and DONA, support self-certification. Because PURSUIT is not overlaying persistence and ubiquity on its underlying rendezvous and delivery, it supports our choice of a separate identification layer most effectively.

Three further papers address design choices. First, similar to our approach, Ghodsi, Koponen and colleagues in their two recent papers on an evolvable architecture ([8], [12]) suggest both a layer of indirection and identification in identifiers of modularized parallel approaches within their framework. Also, Ghodsi et al. ([6]) argue about particular design criteria for namespaces. Although the arguments are incomplete, the approach is valuable. The approach taken to trust is well stated. On the other hand, the definition of denial of service extremely limited, there is no discussion of revocation, and while the suggestion that self-certification is not user-friendly is true, that is exactly the reason for a layered approach, in which different capabilities and characteristics can be provided in different layers.

Scalability: The PURSUIT concept of *scopes* (see [27]) as a mechanism to achieve scalability was derived from Sollins on the concept of *regions* ([25]), although it was preceded by the nested Internet routing structure of Autonomous Systems. Both the flat and strictly hierarchical approaches of the systems discussed above assume an Internet like approach to resolution, based on longest prefix matching, which in turn suggests that prefixes can be consolidated efficiently. Unfortunately, neither the unconstrained hierarchy of CCN nor the flat idspaces of DONA and NetInf are likely to allow for effective consolidation, in contrast with PURSUIT.

Persistence: A number of the efforts include persistence as a key goal. NetInf is explicit in its commitment to persistence. MobilityFirst ([4], [14]) inherits it from the Handle system, although it is not called out as such. The Information Mesh project and URN work were explicit and up front about their commitment to persistence as an underlying goal.

Extensibility and evolvability: This is an area that has become increasingly popular. Not only do some of the FIA proposals call it out explicitly (XIA ([5], [28]) especially), but recently in the papers by Ghodsi et al. ([7], [8]), first in proposing an architecture for "innovation", followed by a broader analysis of requirements for evolution in architecture. A related set of authors, Popa et. al ([19]) proposed HTTP as the protocol at the narrow waist of an architecture allowing for evolution above and below it using URLs for identification. XIA specifically focuses on the ability to define new *principals*, distinguished by name and protocol.

As mentioned above in some of the specifics, there is a small but growing trend in considering sets of architectural criteria. Examples of this are Trossen et al. ([27]), Ahlgren et al. ([1]), and Ghodsi et al. ([7], [8]). In addition, we note here that a number of the ICN systems also target one or another form of security as a goal or design criterion.

5 CONTRIBUTIONS

The contributions of this paper fall into three categories. The first is a framework for making architectural and design choices for identification systems. The second translates the objectives of pervasiveness and persistence into the four architectural goals of longevity, scalability, and evolvability, and security. The third is a principled design, based on layering and modularity, for PPInS

that enables it to overlay any of the existing ICN approaches, thus relieving them of some of their more difficult problems and allowing them to demonstrate their strengths more effectively.

6 ACKNOWLEDGMENT

This work was funded in part by The National Science Foundation under Grant 0915629 and the Intel Corporation.

7 REFERENCES

[1] Ahlgren, B., Dannewitz, C., Imbrenda, C., Kutshcer, D., Ohlman, B. 2011 *A Survey of Information-Centric Networking (Draft)*, Dagustuhl Proceedings 10492, http://drops/dagstuhl.de/opus/volltexte/2011/2941 (Feb. 2011).

[2] Ahlgren, B., D'Ambrosio, Dannewitz, Eriksson, Golic, Gronvall,Home, Lindgren, Mammela, Marchisio, Makela, Nechifor, Ohlman, Pentikousis, Randriamasy, Rautio, Renaul, Seittenranta, Strandberg, Tanauca, Vercellone, Zeghlache, 2010 *Second NetInf Architecture Description,* 4WARD Deliverable, D 6.2 (Jan. 2010).

[5] For each project we have selected only a small representative subset of the documentation.

[3] Berners-Lee, T., Masinter, L., McCahill, M., 1994 *Uniform Resource Locators*, IETF RFC 1738 (Dec. 1994).

[4] Dong, L, Liu, H., Zhang, Y., Paul, S., Raychaudhuri, D., On the Cache-and-Forward Network Architecture, *Proc. IEEE ICC 2009*, (Dresden, Germany, June, 2009).

[5] Han, D, Anand, Dogar, Li, Lim, Machado, Mukundan, Wu, Akella, Andersen, Byers, Seshan, Steenkiste, 2012 XIA: Efficient Support for Evolvable Internetworking, *Proc. USENIX NSDI '12* (Apr., 2012, San Jose, CA).

[6] Ghodsi, A., Koponen, T., Rajahalme, J., Sarolahti, P., Shenker, S., 2011 Naming in Content-Oriented Architectures, *Proc. SIGCOMM ICN '11* (Toronto, Aug. 2011).

[7] Ghodsi, A., Koponen, T., Raghavan, B., Shenker, S., Singla, A., Wilcox, J. 2011 Information-Centric Networking: Seeing the Forest for the Trees, *Proc. HotNets '11* (Cambridge, MA, Nov. 2011).

[8] Ghodsi, A., Koponen, T., Raghavan, B., Shenker, S., Gingla, A., Wilcox, J., 2011 Intelligent Design Enables Architectural Evolution, *Proc. Hotnets '11* (Cambridge, MA, Nov. 2011).

[9] Guo, N., 2011 *Scalable Information-Sharing Network Management*, MIT-CSAIL-TR 2011-06-07 (June, 2011).

[10] Jacobson, V, Smetters, D., Thornton, J., Plass, M., Briggs, N., Braynard, R., 2009 Networking Named Content, *Proc. CoNEXT 2009* (Rome, Dec. 2009).

[11] Koponen, T. Chawla, M., Chun, B-G., Ermolinskiy, A., Kim, K., Shenker, S., Stoica, I., 2007 A Data-Oriented (and Beyond) Network Architecture, *Proc SIGCOMM '07* (Kyoto, Aug. 2007).

[12] Koponen, T., Shenker, Balakrishnam, Feamster, Ganichev, Ghodsi, Godfrey, McKeown, Parulkar, Raghavan, Rexford, Arianfar, Kuptsov, 2011 Architecting for Innovation, *Computer Communications Review* (July 2011).

[13] Moats, R., 1997 *URN Syntax*, IETF RFC 2141 (Standards Track) (May, 1997).

[14] MobilityFirst, 2012 http://mobilityfirst.winlab.rutgers.edu/.

[15] Mockapetris, P., 1987 *Domain Names – concepts and facilities*, IETF RFC 1034, Std 0013 (Nov. 1987).

[16] Mockapetris, P., 1987 *Domain Names – implementation and specification*, IETF RFC 1035, Std 0013 (Nov. 1987).

[17] NDN 2012, http://www.named-data.net/

[18] Ohlman, B., Ahlgren, Brunner, D'Ambrosio, Dannewitz, Eriksson, Gronvall, Horne, Marchisio, Marsh, Nechifor, Pentikousis, Randriamasy, Rembarz, Renault, Strandberg, Talaba, Ubillos, Vercellone, Zeghlache, 2009 *First NetInf Architecture Description*, 4WARD Deliverable D 6.1 (April 2009).

[19] Popa, L., Ghodsi, A., Stoica, I. 2010 HTTP as the Narrow Waist of the Future Internet, *Proc. Hotnets '10* (Monterey, CA, Oct. 2010).

[20] PURSUIT 2012 http://www.fp7-pursuit.ed/PursuitWeb/: published papers, technical reports, deliverable documents.

[21] Smetters, D., Jacobson, V. 2009 *Securing Network Content*, PARC Technical Report (Oct. 2009).

[22] Sollins, K., Masinter, L.1994 *Functional Requirements for Uniform Resource Names*, IETF RFC 1737 (Dec. 1994).

[23] Sollins, K. Van Dyke, J. 1995 Linking in a Global Information Architecture, *Proc. Fourth WWW Conference* (Boston, 1995).

[24] Sollins, K., 1998 *Architectural Principles of Uniform Resource Name Resolution*, IETF RFC 2276 (Jan. 1998).

[25] Sollins, K. 2003 Designing for scale and differentiation, *Proc. SIGCOMM FDNA '03* (Karlsruhe, Aug. 2003).

[26] Sun, S., Lannom, L., Boesch, B. 2003 *Handle System Overview*, IETF RFC 3650 (Informational) (Nov. 2003).

[27] Trossen, D., Sarela, M., Sollins, K. 2010 Arguments for an Information-Centric Internetworking Architecture, *Computer Communications Review* (April, 2010).

[28] XIA 2012, http://www.cs.cmu.edu/~xia/