A Survey of Naming and Routing in Information-Centric Networks

Md. Faizul Bari, Shihabur Rahman Chowdhury, and Reaz Ahmed, University of Waterloo
Raouf Boutaba, University of Waterloo and Pohang University of Science and Technology
Bertrand Mathieu, Orange Labs

ABSTRACT

The concept of information-centric networking (ICN) defines a new communication model that focuses on what is being exchanged rather than which network entities are exchanging information. From the ICN perspective, contents are first class network citizens instead of hosts. ICN's primary objective is to shift the current host-oriented communication model toward a content-centric model for effective distribution of content over the network. In recent years this paradigm shift has generated much interest in the research community and sprung several research projects around the globe to investigate and advance this stream of thought. Content naming and content-based routing are core research challenges in this research community. In this survey, we analyze, compare, and contrast the naming and routing mechanisms proposed by some of the most prominent ICN research projects.

INTRODUCTION

Today's Internet architecture was designed in the 1960s and '70s with the intention to interconnect a few computing resources across a geographically distributed user base. The solution brought us the TCP/IP protocol stack, a communication model where Internet hosts can speak to each other by establishing communication pipes between them. It was an excellent match for client-server applications like HTTP, FTP, telnet, and SMTP. However, host-centric Internet architecture is becoming inadequate for the modern bandwidth-intensive Internet usage patterns, demanding a paradigm shift in contemporary content distribution mechanisms.

As a matter of fact, Internet usage patterns have shifted in the last several years from a host-oriented model to a content-oriented model. Between 2000 and 2006, peer-to-peer applications (e.g., BitTorrent) were the major contributors of Internet traffic. HTTP video has been dominating Internet traffic since 2007, due to the emergence of video streaming portals for user generated content (e.g., YouTube, Google-Video), video on demand (e.g., NetFlix, Hulu), IPTV services). According to the *Cisco Visual Networking Index 2012*, approximately 26 exabytes of Internet traffic is generated per month and approximately 55 percent of Internet traffic is going to be videos by the year 2014.

Historically, the Internet has evolved in an ad-hoc manner. Incremental patches were added to the Internet to handle new requirements as they arose. For example, the Domain Name System (DNS) offers the core functionality of name resolution, yet it was designed and implemented long after the Internet was deployed. Since DNS aims to resolve a name to an Internet host address, it lacks support for content replication, movement, and location awareness [1]. To overcome these problems content distribution networks (CDNs) were introduced. Similarly, peer-to-peer (P2P) file-sharing systems (like Gnutella and BitTorrent) were designed for multi-source content retrieval, inherent replication, and rapid content dissemination. These mechanisms have clearly contributed to improved content access over the Internet. However, they generally operate as overlays and do not leverage the knowledge of the underlying network topology to achieve optimal performance.

These shortcomings motivated the research community to look for alternate architectures for the future Internet (see [2] for a survey on future Internet architectures). Information-centric networking (ICN) is one such alternative. ICN's primary objective is to shift the current host-oriented communication model toward a content-centric model. It relies on location-independent naming, in-network caching, and name-based routing for effective distribution of content over the network. Although ICN has drawn a lot of attention from the research community, research in this area is still in its infancy. Numerous research challenges have to be addressed to bring ICN to life. Some of these challenges include secure and persistent naming, name-based routing, name resolution, in-network caching, on-demand replication, security, privacy, content dissemination, backward compatibility, and incremental deployment capability.

Naming and routing lie at the core of any ICN architecture. ICN projects have proposed

0163-6804/12/$25.00 © 2012 IEEE

diverse solutions for naming and routing. In this survey, we present a focused and in-depth discussion of the naming and routing mechanisms in major ICN projects. More general surveys on ICN can be found in [3, 4].

The rest of this article is organized as follows. We describe the content naming and routing mechanisms of five representative ICN projects. A comparative analysis of different aspects of naming and routing in these ICN proposals are presented. Our perspective on the requirements of an ideal content naming and routing model is introduced, and finally we conclude the article.

Project Name	Reference	Year
Combined Broadcast and Content Based (CBCB)	[6]	2004
Data Oriented Network Architecture (DONA)	[7]	2007
Network of Information (NetInf)	[8]	2009
Named Data Networking (NDN)	[9]	2009
Publish Subscribe Internet Technology (PURSUIT)	[10, 11]	2010

Table 1. *List of surveyed research projects.*

Major Research Projects

The concept of ICN dates back to 2000, when Cheriton *et al.* introduced the concept of name-based routing in TRIAD [5]. Subsequently, a number of research efforts have been dedicated to ICN. In this survey, we analyze, compare, and contrast the ICN projects listed in Table 1. The selected research projects provide a reasonable coverage of the diverse research efforts toward naming and routing in ICN.

Combined Broadcast and Content Based Routing

Combined Broadcast and Content Based Routing (CBCB) [6] is an application level overlay, which superposes a content based communication service over a generic point-to-point network. It has a publish/subscribe architecture, where publishers publish their contents using *messages* and subscribers advertise their interests using *predicates*. A *message* is a set of attribute-value pairs, while a *predicate* is a disjunction of conjunctions of constraints on individual attributes. The published messages are propagated over a broadcast tree from their sources. Nodes use the predicates to prune the branches of the broadcast tree to ensure delivery of the message to interested nodes only.

Naming — CBCB uses a set of attribute-value pairs to name a content in the network. An attribute has a name, a type, and a set of possible values. For example, in CBCB, the name of the content located at uwaterloo.ca/mfbari/srv_naming.pdf will take a form similar to the one in Fig. 1a.

CBCB's naming paradigm is unique in the sense that it differs from traditional URL-based naming as well as from flat naming schemes used by other content-oriented network architectures. But this scheme ensures neither name uniqueness nor secure content names.

Routing — CBCB has a publish/subscribe architecture. It performs routing of "publish" messages by content names, i.e., on attribute-value pairs. A CBCB router implements two protocols:

- Broadcast routing protocol
- Content-based routing protocol

The broadcast routing protocol uses the network topology information to ensure loop-free routing paths. Publishers publish their content using messages and broadcast the messages over the broadcast tree rooted at them. The content-based routing protocol prunes branches in the broadcast tree according to the predicates (interest) declared by the nodes, to ensure delivery of a published message to only those hosts that have expressed interest in that message. A router maintains a content-based forwarding table, where each interface i_k is mapped with a predicate p_k. A router forwards a message to interface i_k if the message's set of attribute-value pairs satisfy predicate p_k. The predicates in the routing table are constructed and updated using two mechanisms: receiver advertisements (RAs) and sender requests (SRs)/update replies (URs).

The routers in the network periodically issue predicates as RAs to push their interest into all potential senders in the network using the broadcast tree rooted at the issuer. For instance, router 6 in Fig. 1b broadcasts its interest using an RA with predicate p_6. When router 4 receives the RA via interface i_6 it updates the interface's associated predicate in the routing table from *false* to $false \vee p_6 = p_6$ and forwards it to the other interfaces. If the received RA is a specialization of an existing predicate mapped with the receiving interface, the router prunes the propagation of the received RA. Figure 1b also depicts this scenario, where router 3 receives an RA with predicate p_2 at interface i_4, and stops forwarding it since it is a generalization of the predicate mapped with its receiving interface.

SRs/URs are used by the routers to pull content-based addresses from the other routers and update their routing tables. A router broadcasts an SR (router 5 in Fig. 1c), and each router on the broadcast tree that receives the SR sends a UR back to the issuing router. The leaf routers of the broadcast tree include their content based address in the UR (Fig. 1c). Other non-leaf routers accumulate all the URs they receive, add their content-based address to the set, perform logical OR operation on them to construct their UR, and send it to the interface where the SR originally arrived. The original issuer of the SR updates its routing table entries using the URs it receives through its interfaces.

Data Oriented Network Architecture

The Data-Oriented Network Architecture (DONA) [7] project argues for a clean slate redesign of the Internet name resolution system. It proposes to use a flat and self-certifying naming scheme with a hierarchically organized name

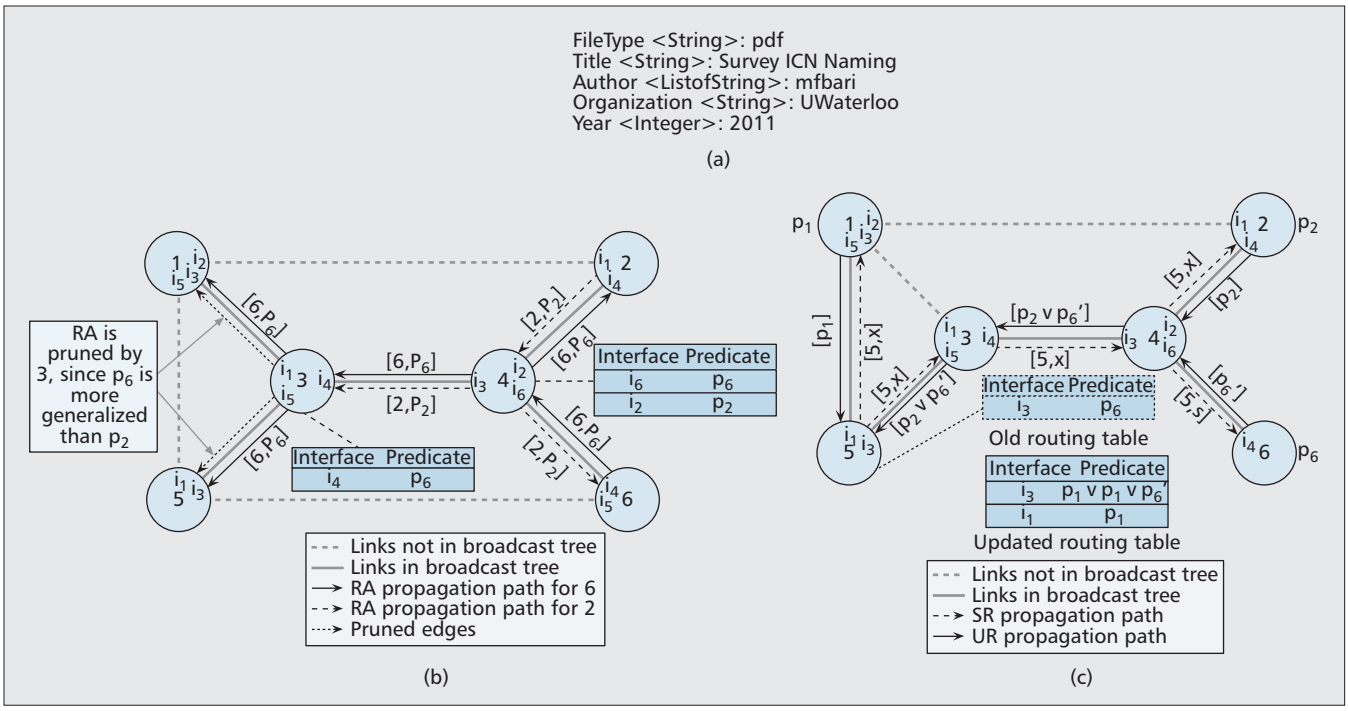

Figure 1. *CBCB routing table: a) naming scheme; b) RA propagation (p_6) and pruning (p_2); c) Routing table update of router 5 using SR/UR.*

resolution infrastructure to achieve three primary objectives:

- Availability in terms of reliability and low-latency
- Name persistence
- Content provenance

Naming — Every content in DONA is associated with a publishing entity called a Principle (owner). Names in DONA are of the form *P:L*, where *P* is the cryptographic hash of the owner's public key and *L* is an owner assigned label. The owner is responsible for the uniqueness and granularity of *L*. Names are globally unique, persistent, and not bounded to any organizational boundaries.

Meta-data associated with each content contains the full public key and digital digest signed by the owner. The *P* part of the name ensures provenance and the signature in the meta-data ensures content integrity. This mechanism opens up new opportunities for in-network caching. Any Internet host with a valid copy can serve as a source. CDNs can leverage this feature for providing access to contents from multiple owners, where clients retrieving those contents only need to trust the CDN. The task of verifying that the public key actually belongs to the owner is

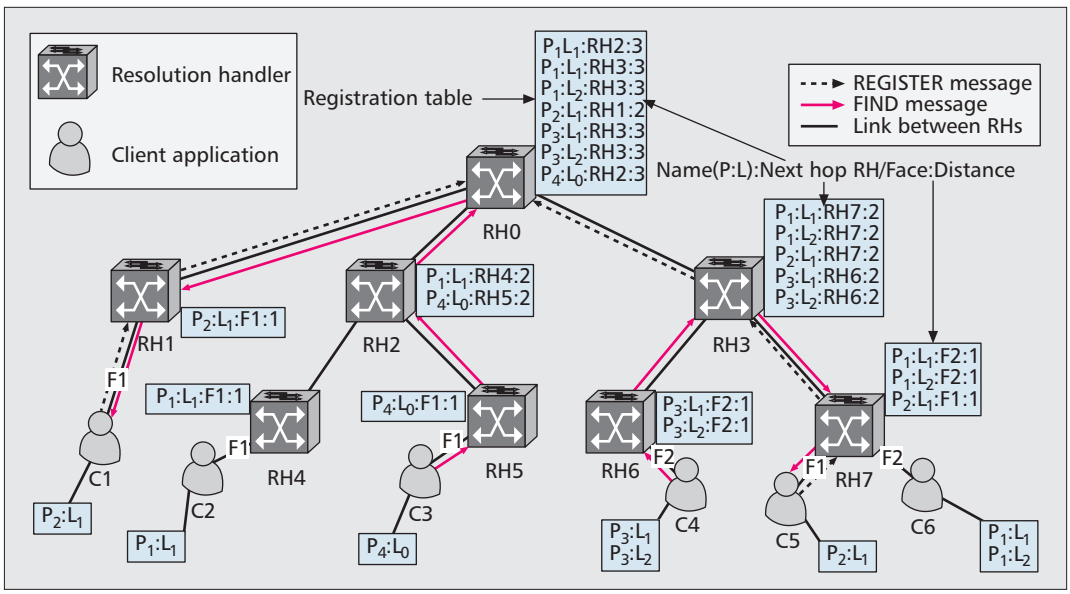

Figure 2. *DONA RH hierarchy with* REGISTER *and* FIND *message propagation.*

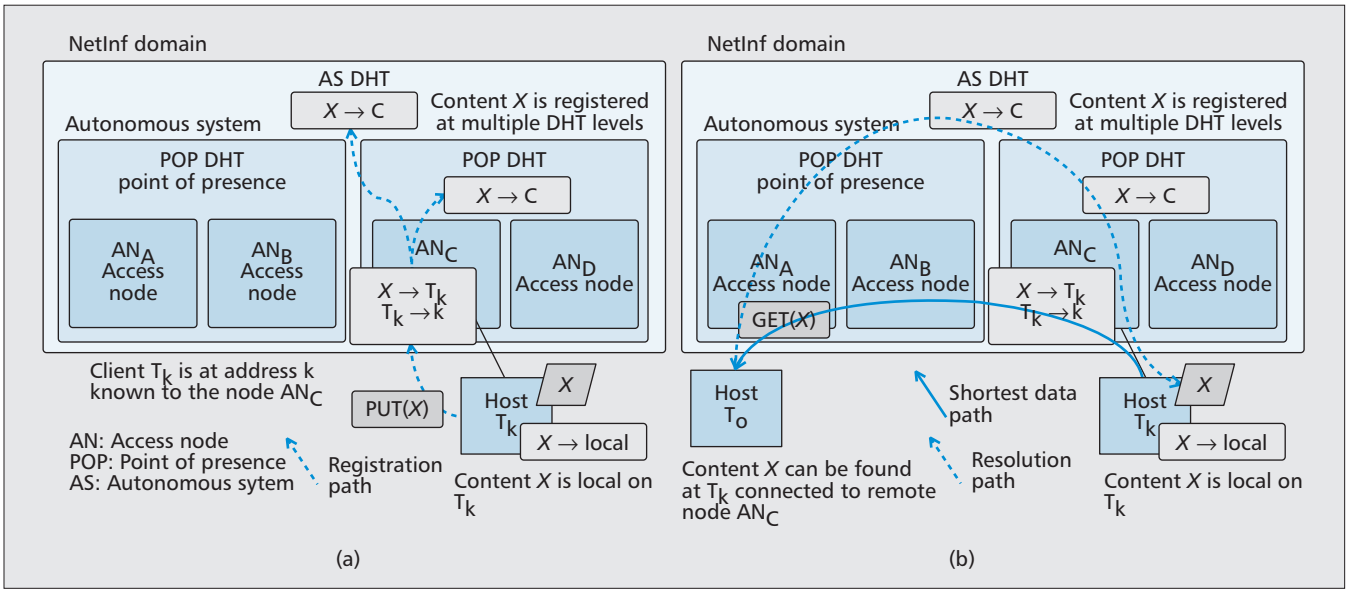

Figure 3. *MDHT content registration, resolution, and retrieval in NetInf: a) register content X at three levels: AN, POP, AS; b) resolve and retrieve content X.*

left for the receiver. This problem can be eliminated in two ways:

- Store the public key of an owner under a special label that is known to all.
- Depend on a PKI or Web of Trust (WoT) infrastructure for key verification.

Routing — Name resolution is performed using the route-by-name paradigm deployed above the IP layer. Name resolution entities are referred to as resolution handlers (RHs). Each domain or administrative entity owns one logical RH. These RHs are organized hierarchically following the organizational and social structure of the Internet (Fig. 2). The resolution infrastructure offers a very simple interface, providing only two operations: *FIND*(*P*:*L*) and *REGISTER*(*P*:*L*). *FIND*(*P*:*L*) locates the object named *P*:*L*, while *REGISTER*(*P*:*L*) sets up necessary states in the RHs to route subsequent *FIND* messages effectively. Routing between the RHs is performed directly on the name, and network operators can define global and local routing policies similar to Border Gateway Protocol (BGP). As a *FIND* message is forwarded, the hop-by-hop domain-level address can be appended to it. Once a *FIND* message is resolved, content can be delivered to the client by sending it over the reverse of the appended path. Alternatively, DONA can use IP routing to return discovered content to the client.

Figure 2 shows the registration tables of different RHs and paths traversed by *REGISTER* and *FIND* messages for a content called P_2: L_1, which is registered by two client applications *C*1 and *C*5 attached to *RH*1 and *RH*7, respectively. Register tables store 3-tuples consisting of ⟨*P*: *L*, next hop *RH*, distance⟩. Each RH forwards the *REGISTER* message to its provider (parent) RH when no record with that name exists in its register table or the new *REGISTER* comes from a copy closer than the previous one. *REGISTER* messages from *C*1 and *C*5 traverse the paths *C*1 → *RH*1 → *RH*0 and *C*5 → *RH*7 → *RH*3 → *RH*0, respectively, and set up necessary states for effectively discovering the closest copy of the content from any location in the network. After receiving a *FIND* message, if there is a matching entry in the registration table, the *FIND* is forwarded toward the next hop RH; otherwise, the *FIND* is forwarded toward its parent RH. This mechanism is guaranteed to discover the closest registered copy, which is evident from the *FIND* message forwarding example shown in Fig. 2. In terms of scalability, DONA imposes name resolution load on the RHs according to their position in the hierarchy and therefore is not scalable as the Tier-1 RHs need to store all names in the network.

Network of Information

Network of Information (NetInf) [8] is a part of the EU FP7 projects 4WARD and SAIL. The 4WARD project is more focused on naming and content searching, while the SAIL project focuses on network transport issues. NetInf proposes using flat and self-certifying names similar to DONA.

Naming — Similar to DONA, NetInf names have two parts, *P*: *L*, where *P* is the hash of owner's public key and *L* is a label chosen by the owner. For a static content, *L* is the hash of the content itself. But for dynamic content, a fixed ID is used as *L* and a digital signature (stored in meta-data) ensures content integrity. NetInf proposes binding using the public/private key pair to the content instead of the owner, so a single owner may use multiple public/private key pairs. Owner authenticity and identification is determined from public key chaining information stored in meta-data. This feature enables anonymous but secure content publishing.

Routing — NetInf uses a multilevel DHT-based name resolution service called MDHT [12] that provides name-based anycast routing. As shown in Fig. 3, MDHT is a topologically embedded

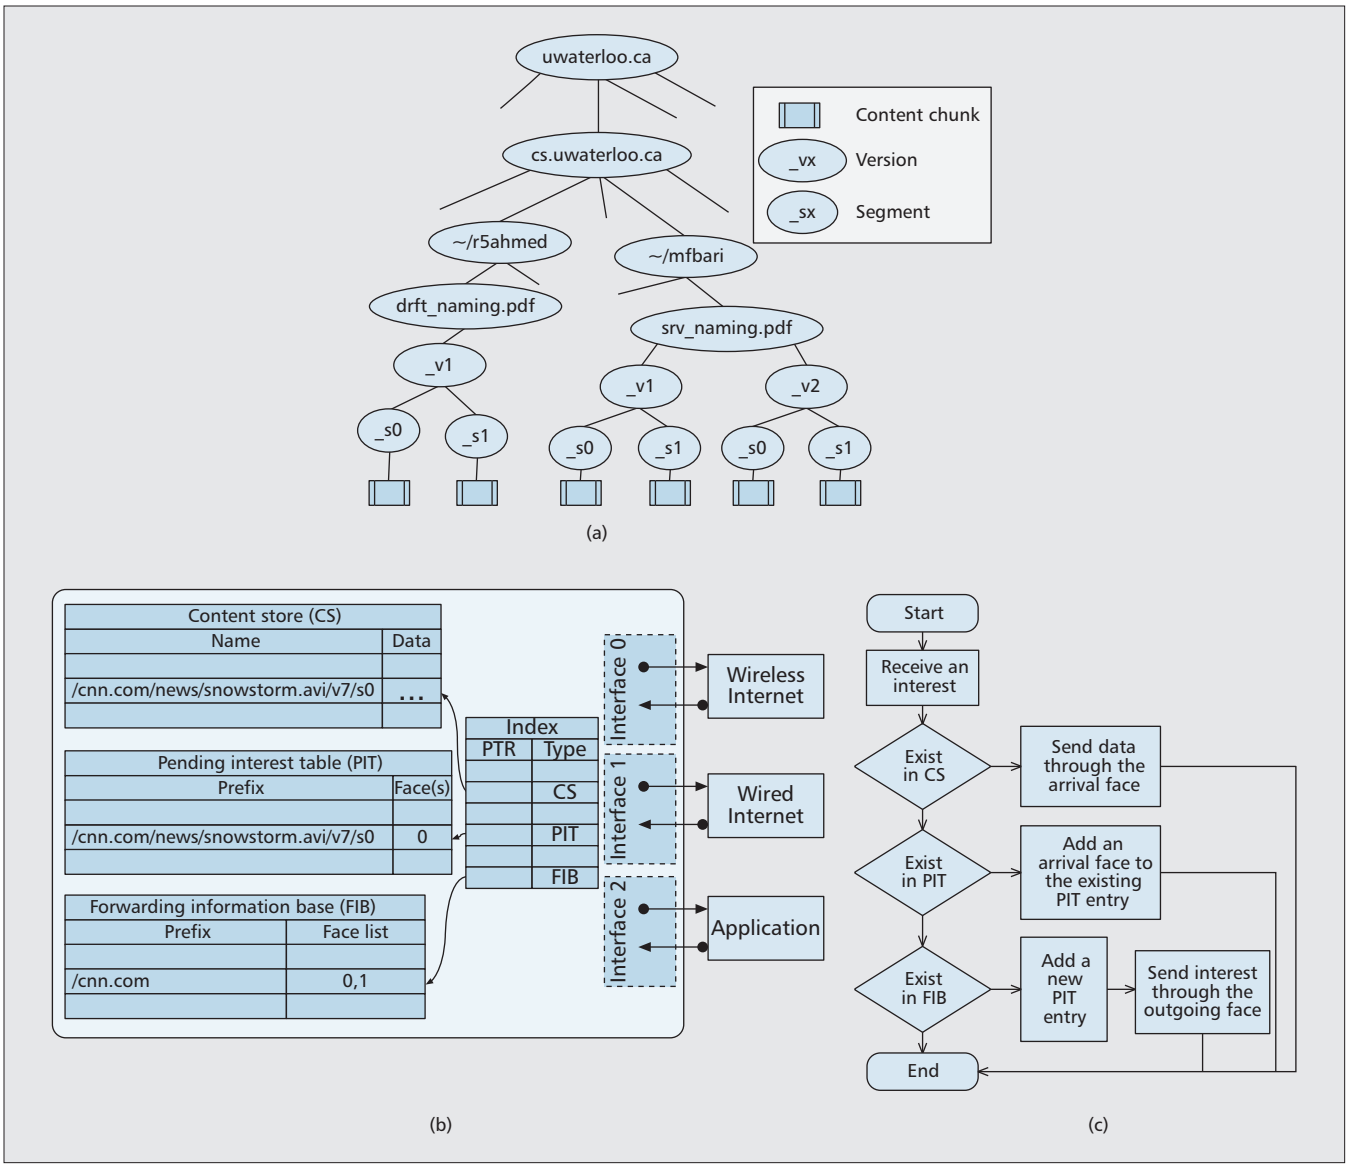

Figure 4. *NDN forwarding engine and interest processing: a) naming scheme; b) forwarding engine; c) interest processing.*

multilevel, nested, hierarchical DHT that utilizes locality in request patterns to minimize intra-AS routing latency. In Fig. 3 three DHTs are nested in the access node (AN), point of presence (POP), and autonomous system (AS) levels, respectively. Each of these DHTs (DHT areas) can run its own DHT algorithm, and any node can take part in multiple DHTs. Intra-area routing and forwarding is done according to the rules of the local DHT algorithms. Inter-area routing is done by finding a node in the local DHT that also takes part in the next higher level DHT.

The registration process for content X is shown in Fig. 3a. Host T_k registers content X at three different levels: AN, POP, and AS. The AN stores two mappings: the first one says that content X belongs to host T_k, and the second one says that host T_k can be found at address k, which can be an IP address or a private address to access node C. POP and AS level DHTs map the content X to the access node C. Figure 3b shows the name resolution and data transmission path for content X. Host T_o looking for content X, first looks it up at its local AN, if it is not found then at its local POP, and after that at the AS level DHT. If the lookup is unsuccessful at the AS level, T_o looks up the name in the Resolution Exchange (REX) system, which is an independent entity responsible for managing registration, updates, and aggregation of names on a global level. Aggregated bindings generated by the REX system are cached by the AS level DHTs to reduce load on the REX system.

Named Data Networking

Named Data Networking (NDN) [9] is one of the four NSF FIA projects to explore and design future Internet architectures. The objective of NDN is to completely redesign the Internet by replacing IP with content chunks as a universal component of transport.

Naming — Names in NDN are composed of multiple components arranged in a hierarchy (Fig. 4a). A component can be any string of arbi-

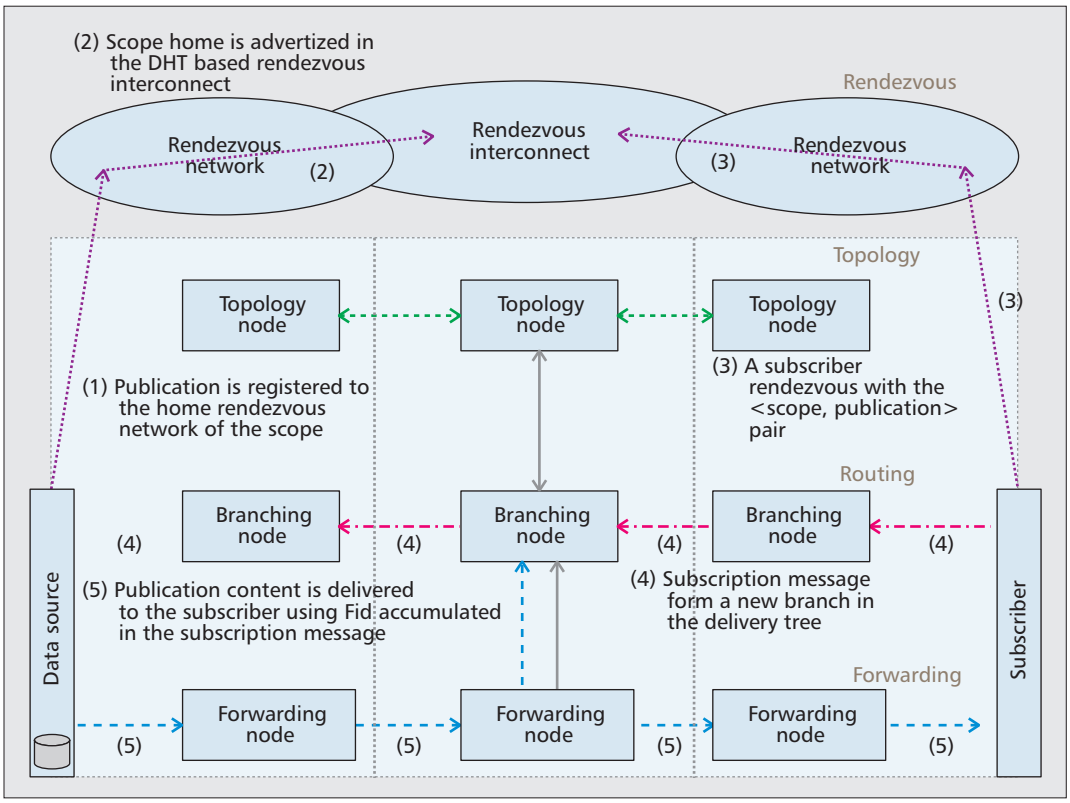

Figure 5. *Rendezvous, topology, routing, and forwarding layers of PSIRP.*

NDN names are human friendly, and non-persistent due to the hierarchical structure. Though the full name of a content with SHA256 digest is unique, the user assigned part of a name may not be unique.

trary length. The NDN transport layer imposes no restriction on names except the component structure. Names are generated and assigned by users. NDN only proposes the structure and anticipates that naming standards will emerge and become standardized through the development of different types of applications. Names also contain information like version and segment numbers. All NDN names implicitly contain a SHA256 digest of the content for resolving ambiguity. To provide content authenticity and integrity, name to content mappings are digitally signed and delivered with the content. NDN names are human friendly, and non-persistent due to the hierarchical structure. Though the full name of a content with SHA256 digest is unique, the user assigned part of a name may not be unique.

Routing — Two key messages are exchanged in NDN: Interest and Data. These messages are routed using a *route-by-name* paradigm. A client issues Interest for a content by broadcasting over all available connections. Any NDN node having the original or replicated copy of that content or caching it, can satisfy the Interest by responding with the corresponding Data. Both Interest and Data identify the content being exchanged by name. Interest and Data have a one-to-one mapping to maintain a strict flow balance. So Data is sent only in response to an Interest and that Interest is consumed by the Data.

An NDN node maintains the following tables: Forwarding Information Base (FIB), Pending Interest Table (PIT), and Content Store (CS). The NDN forwarding engine is shown in Fig. 4b. The FIB is used to store probable source(s) for Data. The PIT stores the return path state for possible Data messages in response to the Interests forwarded upstream. CS is used for caching Data. The interest processing mechanism in an NDN node is depicted in Fig. 4c. NDN can perform efficient routing depending on the component structure of names. It performs prefix aggregation and loop free forwarding for routing table compression and reduced messaging overhead yet the achievable scalability is far below that of IP routing. The NDN project is developing an Open Shortest Path First (OSPF) like routing protocol for named data, called OSPF-N, for populating and updating the routing tables [13].

Publish Subscribe Internet Technologies

PURSUIT [11] is an EU FP7 project recently launched as a follow up to a former EU FP7 project, Publish Subscribe Internet Routing Paradigm (PSIRP) [10]. PSIRP proposed a clean slate routing architecture for ICN shifting the current *send-receive* based Internet toward the *publish-subscribe* paradigm. One of the main goals of the PURSUIT project is to develop Internet-scale deployable components of the PSIRP architecture.

Naming — PSIRP uses the same naming scheme as DONA and content names are called resource identifiers (RIds). PURSUIT continues to use the same naming scheme. Content persistence is ensured by using special network entities called data sources that reside at the network

Existing ICN naming schemes belong to three categories: hierarchical, flat, and attribute-value pair based. All three categories of names can be aggregated to some level, to improve routing table scalability. NDN performs prefix aggregation on hierarchical names for scalability.

edge (Fig. 5). Data sources periodically refresh content publication states in the network.

A PSIRP network is based on the concept of *scopes*, where scopes are identified by Scope identifiers (SIds.) Scopes control access rights, authorization, reachability, availability, replication, persistence, and upstream resources of a content. Content publication (*publish*) and content request (*subscribe*) are based on ⟨*SId*, *RId*⟩ pairs. An implicit assumption of this mechanism is that content publishers will publish and subscribers will subscribe to contents in a scope that they trust.

Routing —The routing infrastructure of PSIRP comprises four components: *Rendezvous*, *Topology*, *Routing*, and *Forwarding* (Fig. 5). PSIRP assumes that the network consists of Autonomous Systems (Domains) similar to the current Internet.

The Rendezvous component consists of one rendezvous network (RN) per domain. They are interconnected by a global (Internet scale) hierarchical DHT-based rendezvous interconnect (RI). RI acts as middleman between the publishers and subscribers by matching data sources of certain publications with subscribers' interests. An RN is responsible for locating the publications and scopes of its network. An RN also advertises its scope to the RI so that the scope becomes globally reachable. Each individual RN can be constructed using a hierarchical name resolution system like DONA.

The Topology component consists of one topology node (TN) per domain. Each TN is responsible for managing the intradomain topology and load balancing. TNs also exchange interdomain path vectors between themselves, similar to BGP, for policy compliant interdomain routing.

A number of branching nodes (BN) build up the Routing component. A BN uses the topology information maintained by a TN to route subscription messages from subscribers toward data sources and cache popular content. The intradomain routing guarantees that subscription for a certain data source will always go through the same BN in the domain.

The Forwarding component has a number of forwarding nodes (FNs). FNs use Bloom filter based forwarding to implement a simple and fast forwarding algorithm to send back a content to the subscriber.

Subscribers send subscription requests to their home RN to get content locators. Then the BN of a subscriber's domain uses the network topology information obtained from the TN to forward the subscription request. Content request packets accumulate a return path in a Bloom filter called the forwarding identifier (FId). The FNs use this FId to send back a content to the subscriber. Typical phases of publish/subscribe process are shown in Fig. 5.

Comparative Analysis

In this section we analyze, compare, and contrast the content naming and routing mechanisms of the above presented research projects based on several competing design choices.

Hierarchical vs. Flat vs. Attribute-Value Based Naming

Existing ICN naming schemes belong to three categories: hierarchical, flat, and attribute-value pair based. All three categories of names can be aggregated to some level, to improve routing table scalability. NDN performs prefix aggregation on hierarchical names for scalability. CBCB uses attribute-value pairs to prune unnecessary branches of the broadcast tree. Predicates with common attributes at a CBCB router can be combined to reduce the number of routing table entries. Flat names used by DONA, NetInf/MDHT, and PURSUIT, in the form P: L can be aggregated at the publisher level. But this aggregation is not very effective in DONA. That's why DONA suffers from the problem of increased load at higher level RHs. Flat names are more suitable for DHT based lookup services like NetInf/MDHT and PURSUIT, where storage load is distributed uniformly between the resolution nodes. PURSUIT also performs another level of aggregation on names at the scope (sID) level.

Flat vs. Human Friendly Names

The use of a cryptographic hash in a content name hides the underlying content's semantics from human users and makes the names difficult to remember. This fact makes the self-certifying flat names in DONA, NetInf and PURSUIT less human friendly. On the other hand, the naming schemes of CBCB and NDN are more human friendly because of their hierarchical structure or attribute-based partition, which make them easier to remember, and they provide more information about the content's semantics. But this human friendliness comes at the cost of some challenges: ensuring global uniqueness, security binding [14], and authenticity. This raises a fundamental question — *Should we prefer human friendly names with possible security vulnerability OR flat, location independent and self certifying names that can ensure content integrity?*

Recent usage trends can help in answering this question. Nowadays people like to find content by doing a search in a search engine using some keywords. Using search keywords is easier than remembering a full URL. This idea is also supported by the recent trend of merging the address bar and search box into one unified bar in web browsers. One can argue in favor of bookmarks, but doing a search on Google is faster than going through hundreds of bookmarks. This trend has an implication: non-human-friendly names are not a barrier in accessing content as long as the contents have some keywords associated with them. What users care more about is getting the authentic content. Therefore, in the ICN context, we argue that content names can be less human friendly to incorporate features such as location independence and self-certifying capability.

Name Resolution vs. Name-Based Routing

Most ICN naming schemes are based on the concept of identifier/locator split. Existing mechanisms for discovering a content from its location-independent identifier can be categorized under two major approaches: name resolution

and name-based routing. The process of name resolution involves two steps: in the first step the content name is resolved to a single or a set of locators (e.g., IP address) and in the second step, the request is routed to one of these locators using topology based shortest path routing (e.g., ISIS, OSPF). On the other hand, in name based routing, request forwarding is performed directly based on the identifier (name) alone and some sort of state information is set-up along the way so that the content can travel back to the requester. Among the presented research projects, NetInf/MDHT and PURSUIT follow the name resolution approach, while CBCB, DONA, and NDN follow the name based routing approach.

NetInf/MDHT has topologically embedded, nested, and hierarchically organized DHTs for Intra-AS routing and global name registration system called REX for Inter-AS routing. On the other hand, PURSUIT uses multiple rendezvous networks (not related to any AS) interconnected by an Internet-scale DHT. A rendezvous network uses a name resolution system similar to DONA's.

Although CBCB, DONA, and NDN follow the name-based routing approach, their naming schemes are quite different. DONA uses flat names, NDN uses hierarchical names, and CBCB uses a list of attribute-value pairs to name a content. DONA has a hierarchical name resolution infrastructure (closely matching the Internet AS hierarchy) for storing indexes. NDN's approach to routing is similar to traditional topology based IP routing, where IP addresses are replaced by content names. However, the routing protocol of CBCB is quite special, as explained earlier, in that it uses a mix of flooding and publish/subscribe-based forwarding in its approach to routing.

Name resolution approaches can guarantee discovery of any content in the network in a bounded (on network size) number of hops. On the other hand, name-based routing approaches do not guarantee discovery of content. Instead they promise a high probability of content discovery, which is most of the time proportional to the number of visited nodes. Update message overhead in name resolution approaches is lower than that of name-based routing approaches, as the later requires flooding the whole network for update propagation. However, node failures in a name resolution system may render a portion of the index inaccessible even though the content is available. This problem does not exist in name-based routing approaches as they can discover alternate routing paths due to message flooding (NDN) or broadcasting (CBCB). Moreover, the accumulated storage requirement for a name resolution system is much more than that of a name-based routing approach. Name resolution approaches typically need to maintain two databases: name to IP mapping in the resolution system and IP reachability information in the routing system, whereas name-based routing approaches only need to maintain a mapping between names and network locations.

Clean Slate vs. Incremental Approach

ICN research projects can be divided in two categories in this regard. Some projects require clean slate deployment, while others can be incrementally deployed or coexist with current Internet technologies. CBCB, NDN, and PURSUIT follow a clean slate design which does not require IP routing. CBCB routing tables index next hop routers against logical predicates. Routing tables in NDN index next hop routers against content names instead of IP addresses. PURSUIT requires the deployment of a global DHT to interconnect rendezvous networks. On the other hand, DONA and NetInf (MDHT) propose a new name resolution infrastructure (replacing DNS), while still using topological IP routing protocols (e.g. BGP, OSPF, ISIS, RIP). Both can be deployed incrementally, one ISP at a time. As incremental deployment is a preferable feature for any practical solution, approaches like DONA and NefInf seem more suitable than others in this respect.

Scalability

According to the BGP Routing Table Analysis Reports,[1] the biggest Internet routing table contains around 4×10^5 BGP routes for covering about 3.8×10^9 IPv4 addresses and 6×10^8 hosts. This 10^4 scaling factor between IPv4 addresses and BGP routes is achieved by prefix based routing and route aggregation. However, the number of addressable ICN contents is expected to be several orders of magnitude higher. For example, Google has indexed approximately 10^{12} URLs,[2] which would impose 7 orders of magnitude scalability requirement on a name-based routing scheme analogous to BGP. In DONA, Tier-1 RHs need to store all names in the network. The same is true for the REX system in NetInf and the global interconnecting DHT in PURSUIT. NDN needs to flood each new message over the whole network, whereas CBCB needs to broadcast publish/subscribe messages over a large number of network domains. Currently, there are $\sim 10^8$ registered *second level domains*[3] in the Internet, which are analogous to the content publishers in the context of ICN. Even if the names can be aggregated at the publisher level as discussed earlier, the name based routing tables need to handle around 10^8 routes, which is 4 orders of magnitude larger than the biggest BGP routing table size. Evidently, all of the considered ICN projects will face severe scalability problems for an Internet-scale deployment.

Name-based routing approaches do not guarantee discovery of content. Instead, they promise a high probability of content discovery, which is most of the time proportional to the number of visited nodes.

Requirements of an Ideal Content Naming and Routing Model: Our Perspective

Based on our analysis, we now present our perspective on suitable design choices for naming and routing models in ICN.

[1] *http://bgp.potaroo.net*

[2] *http://googleblog.blogspot.com/2008/07/we-knew-web-was-bit.html*

[3] *www.registrarstats.com/TLDDomainCounts.aspx*

contents in ICN need globally unique, secure, location independent and human friendly names. But it is difficult to find one single naming scheme that satisfies all of these properties. Rather, a multilayer naming scheme that combines self certifying names with human friendly keywords will be more suitable in practice.

NAMING

Name structure: We believe that instead of having human friendly names for web content, it is more important that a content has a set of owner selected keywords assigned to it. These keywords can be later used by search engines to index them and enable end users to search for their desired content.

Self-certification: For content naming, self-certifying flat names can be adopted. There are two main reasons for that:

- They intrinsically provide name persistence, security binding, authenticity, and global uniqueness.
- Scalability analysis provided by DONA [7] suggests that flat name based routing is within a grasp of today's technology.

In short, contents in ICN need globally unique, secure, location-independent, and -friendly names. But it is difficult to find one single naming scheme that satisfies all of these properties. Rather, a multilayer naming scheme that combines self-certifying names with human-friendly keywords will be more suitable in practice.

ROUTING

For content routing we identify a list of desirable properties for any ICN routing mechanism. However, achieving all of these features in one routing scheme may be difficult.

Content state: Whether it is a name resolution or name-based routing approach, it should provide low-latency network-level primitive operations for content (original, replica, or cached) registration, meta-data update, and deletion. None of the presented research projects explicitly mention meta-data update or content deletion. Whether content deletion should be explicit or expiry-time based or some hybrid is an interesting question in this context.

Discovery of closest copy: The routing mechanism should be able to route a content request to the closest (based on some network metric) copy. This feature ensures the reduction of inter-domain traffic.

Resolution and retrieval locality: Message propagation for name resolution and retrieval should not leave the network domain that contains both the source and the content.

Discovery guarantee: The routing mechanism should provide guarantees on discovery of any existing content, regardless of the content's popularity and replication level.

Scalability: The number of contents for ICN is on the order of trillions. Any ICN routing/name resolution scheme needs to scale to at least this many contents and possibly beyond to accommodate future growth. The trade-off between routing stretch (ratio between routing path length and minimum length path) and routing table size needs to be analyzed, while keeping in mind the huge number of names and physical limitations imposed by memory technologies.

Network-level deployment: Ideally the content retrieval process in ICN should be a one-step process, either by combining name resolution and routing in a single step or by completely eliminating the name resolution part. Except for NDN, all other presented mechanisms require an overlay-based name resolution step. However, NDN, being an unstructured-flooding based routing protocol, neither guarantees content discovery nor ensures scalable routing table size and manageable update message overhead. An effective routing mechanism for ICN may require combining the advantages of structured and unstructured routing mechanisms, while still operating at the network layer without requiring any overlays.

Security infrastructure: Most proposed naming schemes use public/private key pairs for ensuring content integrity and provenance. However, the task of linking a public key to a real-world entity is left for a trusted third party (e.g., PKI). But there are various security risks associated with such schemes [15]. This has been a largely unexplored area in the context of ICN and needs further investigation to determine a better security model.

CONCLUSION

In this article, we have surveyed content naming and routing mechanisms of five ICN research projects. We have distilled a set of design alternatives and presented a qualitative comparison of the competing design decisions made in these projects. This is by no means an exhaustive list of projects in this area; instead, our target was to focus on a representative set. We have also presented our perspective on the requirements of an ideal naming and routing model for ICN. The results of our survey indicate a clear lack of a de facto naming and routing model, and provide a starting point for the reader interested in this exciting research area.

ACKNOWLEDGMENT

This work was supported in part by the Natural Science and Engineering Council of Canada (NSERC) under its Strategic program, in part by Orange Labs France, and in part by the WCU program through the Korea Science and Engineering Foundation funded by the Ministry of Education, Science and Technology (Project No. R31-2008-000-10100-0).

REFERENCES

[1] M. Walfish, H. Balakrishnan, and S. Shenker, "Untangling the Web from DNS," *Proc. 1st Conf. Symp. Networked Systems Design and Implementation — Volume 1*, ser. NSDI'04, Berkeley, CA, 2004, pp. 17–17.
[2] J. Pan, S. Paul, and R. Jain, "A Survey of the Research on Future Internet Architectures," *IEEE Commun. Mag.*, vol. 49, no. 7, July 2011, pp. 26–36.
[3] B. Ahlgren *et al.*, "A Survey of Information-Centric Networking," *IEEE Commun. Mag.*, vol. 50, no. 7, pp. 26–36, July 2012.
[4] J. Choi *et al.*, "A Survey on Content-Oriented Networking for Efficient Content Delivery," *IEEE Commun. Mag.*, vol. 49, no. 3, Mar. 2011, pp. 121–27.
[5] D. Cheriton and M. Gritter, "TRIAD: a Scalable Deployable NATbased Internet Architecture," Technical Report, Jan. 2000, available: http://www-dsg.stanford.edu/triad/
[6] A. Carzaniga, M. Rutherford, and A. Wolf, "A Routing Scheme for Content-Based Networking," *INFOCOM 2004, 3rd Annual Joint Conf. IEEE Computer and Commun. Societies*, vol. 2, Mar, 2004,vol. 2, pp. 918–28.
[7] T. Koponen *et al.*, "A Data-Oriented (and Beyond) Network Architecture," *SIGCOMM Comput. Commun. Rev.*, vol. 37, Aug. 2007, pp. 181–192.

[8] C. Dannewitz *et al.*, "Secure Naming for a Network of Information," *INFOCOM IEEE Conf. Comp. Commun. Wksp.*, 2010, Mar. 2010, pp. 1 –6.
[9] V. Jacobson *et al.*, "Networking Named Content," *CoNEXT*, J. Liebeherr et al., Eds. ACM, 2009, pp. 1–12.
[10] D. Lagutin, K. Visala, and S. Tarkoma, "Publish/Subscribe for Internet: PSIRP Perspective," *Towards the Future Internet Emerging Trends from European Research*, vol. 4, 2010, pp. 75–84.
[11] N. Fotiou *et al.*, "Developing Information Networking Further: From PSIRP to PURSUIT," *Int'l. ICST Conf. Broadband Communications, Networks, and Systems (BROADNETS)*, 2010 (invited paper), Oct. 2010.
[12] M. D'Ambrosio *et al.*, "MDHT: A Hierarchical Name Resolution Service for Information-Centric Networks," *Proc. ACM SIGCOMM Wksp. Information-Centric Networking, ACM*, 2011, pp. 7–12.
[13] L. Wang *et al.*, "OSPFN: An OSPF Based Routing Protocol for Named Data Networking," Technical Report NDN Technical Report NDN-2012-13, July 2012.
[14] A. Ghodsi *et al.*, "Naming in Content-Oriented Architectures," *Proc. ACM SIGCOMM Wksp. Information-Centric Networking*, ser. ICN '11, New York, NY, USA: ACM, August 2011, pp. 1–6.
[15] C. Ellison and B. Schneier, "Ten Risks of PKI: What You're Not Being Told About Public Key Infrastructure," *Computer Security Journal*, vol. 16, no. 1, 2000, pp. 1–7.

Biographies

Md. Faizul Bari (mfbari@uwaterloo.ca) is a Ph.D. student at the School of Computer Science at the University of Waterloo. He received M.Sc. and B.Sc. degrees in computer science and engineering from Bangladesh University of Engineering and Technology (BUET) in 2009 and 2007, respectively. He also holds the position of assistant professor at BUET. He has served as a reviewer for many international conferences and journals. His research interests include future Internet architecture, network virtualization, and cloud computing. He is the recipient of the Ontario Graduate Scholarship, Presidents Graduate Scholarship, and David R. Cheriton Graduate Scholarship at the University of Waterloo. He also received a Merit Scholarship and Dean's award at BUET.

Shihabur Rahman Chowdhury (sr2chowdhury@uwaterloo.ca) received his B.Sc. degree in computer science and engineering from Bangladesh University of Engineering and Technology. He is currently working toward his Masters degree at the David R. Cheriton School of Computer Science, University of Waterloo. His research interest include future Internet architecture, peer-to-peer systems, and algorithms.

Reaz Ahmed (r5ahmed@uwaterloo.ca) is working as an associate professor at the Department of Computer Science and Engineering, BUET. He received his Ph.D. degree in computer science from the University of Waterloo, in 2007. He received M.Sc. and BSc. degrees in computer science from BUET in 2002 and 2000, respectively. He received the IEEE Fred W. Ellersick award in 2008. His research interests include future Internet architectures, wide area service discovery and content sharing peer-to-peer networks with focus on search flexibility, efficiency and robustness.

Raouf Boutaba [F] (rboutaba@uwaterloo.ca) received M.Sc. and Ph.D. degrees in computer science from the University Pierre and Marie Curie, Paris, France, in 1990 and 1994, respectively. He is currently a professor of computer science at the University of Waterloo and a distinguished visiting professor at the Division of IT Convergence Engineering at POSTECH. His research interests include network, resource and service management in wired and wireless networks. He is the founding editor in chief of *IEEE Transactions on Network and Service Management* (2007–2010) and on the editorial boards of other journals. He has received several best paper awards and other recognitions such as the Premier's Research Excellence Award, the IEEE Hal Sobol Award in 2007, the Fred W. Ellersick Prize in 2008, the Joe LoCicero and Dan Stokesbury awards in 2009, and the Salah Aidarous Award in 2012.

Bertrand Mathieu [SM] (bertrand2.mathieu@orange.com) is a senior researcher at France Telecom, Orange Labs since 1994. He received a Diploma of Engineering in Toulon, an M.Sc. degree from the University of Marseille, and a Ph.D. degree from the University Pierre and Marie Curie, Paris. His research activities are related to dynamic overlay networks, P2P networks, and Information centric networking. He has contributed to several national and European projects, and published more than 30 papers in international conferences, journals, and books. He is a member of several conferences, Technical Program Committees, and an SEE Senior Member.

An effective routing mechanism for ICN may require the combining of the advantages of structured and unstructured routing mechanisms, while still operating at the network layer without requiring any overlays.