

Proposal for a book chapter in
" Media Networks; Architectures, Applications, and Standards"

Chapter: "Information-Centric Networking: Current Research Activities and Challenges"

B. Mathieu, P. Truong, J.F. Peltier, W. You - Orange Labs – Lannion – France

G. Simon – Telecom Bretagne – Brest - France

Introduction

The Internet usage over the past decade has shifted away from browsing to content dissemination. Host-to-host communications tend to disappear to make way for one-to-many or many-to-many distribution and retrieval of content objects with an increasing request for a better support of mobility. Users only want to know what content is available and how to get it rapidly anywhere, on demand, on any device and on the go. They do not care about the location of the content.

While imaginative solutions through incremental changes to the network stack or overlay networks have been successfully proposed to address these new usage patterns, it is widely admitted that these solutions have also limitations and drawbacks, so that they do not provide enough real benefits in terms of scalability, security, mobility, and manageability to encourage the creation or deployment of more ambitious and innovative services. They are for example confronted to the problem of intercompatibilities due to contentions or tussles between Internet stakeholders that generally have adverse interests and want to promote their solution. In addition, continuously adding overlays or patches to upgrade Internet often results in increased complexity that favours vulnerabilities, limits manageability and leads to lack of flexibility. All the current workarounds do not actually solve at the root of the problem: they are still host-centric, based on

the endpoint addressing and forwarding principles over which Internet was built historically, while the large majority of today's network communications are content-centric, i.e. concern production and consumption of pieces of data. The following picture 1 depicts this concept, which is known as "from hourglass to lovehandles" problem.

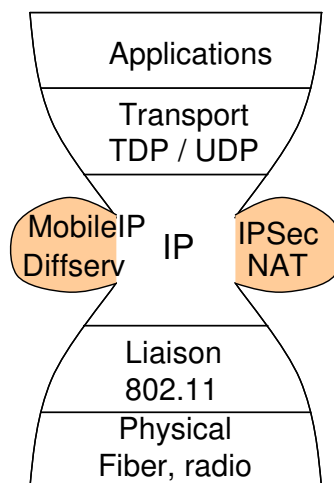


Figure 1: Lovehandles problem in IP

Recognizing that the host-oriented paradigm does not cope with the emergence of voracious content-consuming applications, several projects and initiatives have been started over last years to envision a clean slate foundation for the future Internet promoting information at the centre of the design considerations. We propose in this paper an overview on this rising Information-Centric Networking (ICN) paradigm. In section 2, we present the ICN model with the naming, addressing and forwarding issues. The section 3 highlights the most relevant results of the different research activities on the topic in the literature. In section 4, we outline challenges and requirements that need to be addressed for the development of the ICN paradigm. Section 5 is devoted to concluding remarks.

Information-Centric Networking Paradigm

Motivation for a Network of Information

Internet was originally designed to be an end-to-end connectivity substrate for the delivery of data. All the subsequent enhancements made for improving its architecture revolved around the conversation model that consists of communications between machines using the IP protocol. The current Internet architecture is now a constantly and rapidly evolving interconnection of thousand networks that act as simple carriers providing basic packet delivery services without guarantees, meaning that they make their best effort to try to deliver to receivers anything that senders want to send while only using IP addresses to identify end-hosts for data forwarding and unawarely considering what is being delivered.

Unfortunately, this endpoint-centric communication model does not cope with the overwhelming usage in current Internet anymore. The vast majority of today's traffic is indeed driven by information production and retrieval, ranging from simple RSS feed aggregators to advanced multimedia streaming services, including user-generated or dynamic Web content. The relationships between network entities are not restricted to the view of network topology, but also represent social or content-aware connections between users that can additionally share common interests (eg newsgroups, online photo and video sharing service such as Flickr or YouTube, social networks Facebook or Twitter, P2P-based file sharing). The resulting graph modelling today's Internet communications is then very complex. Nodes are mainly pieces of content, rather than endpoints (locators) addressed by the underlying IP protocol, and are connected by one or more types of interdependency based on the notion of intention, interest or policy-based membership. While still overlaying on top of the host-to-host conversation model,

all the current workarounds for Internet support of emerging information-centric applications increase the complexity and do not efficiently map all the relevant ties between the nodes in terms of:

- **Security:** today's IP-based network security requires trust on end-hosts and on connections over which content transfer occurs. As a result, the main flaw in IP addressing is that the network accepts anything from senders regardless of information contained in packets, provided that the senders appear legitimate. This situation leads to unsolicited and malicious messages sent to receivers.
- **Mobility and multi-homing:** the dual role of IP addresses as both network layer locators and transport layer identifiers limits the flexibility of the Internet architecture for a more efficient support of mobility and multihoming. In particular, transport protocols are bound to IP addresses to identify communication sessions, which are interrupted when an address changes.
- **Multicast delivering:** IP Multicast protocol was designed as an after-thought add-on around the original point-to-point communication model to offer the ability to send information to a group of receivers. However IP Multicast has never taken off outside of LAN environments because of its scalability shortcomings, so that complex overlay solutions are rather preferred to deploy multicast services at large scales.
- **Scalability and QoS guarantees:** with the rapid proliferation of content distribution services, costly solutions such as overlay networks (e.g. CDN or P2P) have been proposed to alleviate the huge demand of bandwidth and to improve user experience by pushing and caching content to the network's edge, but performance bottlenecks still persist in the last mile and the inability for network operators to control traffic traversing

their networks often results in conflicts of business interests or an inefficient network resource optimization.

To overcome these issues from the incompatibility between the usage and communication models on Internet, the Information-Centric Networking, also called content-centric or data-centric networking paradigm has been proposed over few years to target a clean slate architecture redesign by placing content in the foreground, at the heart of network transactions. In the remainder of our paper, we will use the terms information or data as well as the term content interchangeably.

Concepts of Information-Centric Networking

The ICN paradigm consists of communications that revolve around the production and consumption of information matching user interest. The principal concern of the network is to expose, find and deliver information rather than the reachability of end-hosts and the maintenance of conversations between them. As a global view, the paradigm can be divided into two functional parts: information dissemination or exposure, and information retrieval. On a rather low level from the ICN perspective, a network is a set of interconnected pieces of information, also called as content, information or data objects, which are addressed by names for routing and managed by applications or services at the higher middleware level. The naming scheme for identifying content objects in ICN is intended to replace the current IP naming scheme, which mixes host locations and content identifiers (like an URL for instance). In particular, the name of a content object is globally unique and independent of its location (i.e. the host holding the data). Content objects is an abstract notion and can be any type, including for

example web applications (e.g. a piece of mail generated by online webmail services), static or user-generated content (photos, videos, documents, etc), real time media streams such as VoIP, VoD, Web TV, or online videos and music (a stream being considered as series of chunks of data), more complex interactive multimedia communications, or even devices (e.g. routers, data servers, etc) for network management. Objects are sometimes organized into clusters to define social relationships or some ontology between them (as illustrated by the notion of scope in PSIRP); they can also be mutable, combined or aggregated to form new objects.

In ICN, senders do not send content directly to receivers, and any data object delivery is controlled by receivers (cf Figure 2). A sender (or content provider) that has objects to distribute does not actually transmit them in the network, but it rather sends advertisement messages to inform the network that it has content to diffuse, without knowledge of receivers that may be interested in it. A receiver or consumer declares its interest for some content, without knowledge of potential senders that may have queried content. Only when receiver's intention matches a published information object, the network initiates a delivery path from the sender to the receiver so that content retrieval can now start for the receiver. The match of interest rather than the findability of endpoint that provides content dictates thus the establishment of a communication in ICN. The information-centric network connects producers to consumers, disregarding the underlying hosts involved in communications. The focus is only on content, not on the hosts storing content. Only queried content is delivered to receivers that have asked for it beforehand. Another key principle that comes with the interest-oriented networking in the ICN paradigm is the use of dynamic content caching to enable fast, reliable and scalable content delivery with maximized bandwidth to avoid congestion. A router (in the delivery path from the sender to the receiver) can for instance cache content objects that traverse it so that subsequent queries for the

same objects can be satisfied rapidly by the router. This means that routing in ICN consists in finding and delivering copies of data objects to consenting receivers from the most efficient location in the network.




Figure 2: Content dissemination and retrieval in ICN (with caching capability)

When using caching for content retrieval, it is important to give users every guarantee that cached content did indeed come from the original source. ICN affixes trust directly on content itself, rather than the end-to-end connexion carrying it. This allows users to assert immediately the safeness of content and therefore evading host-related threats and vulnerabilities.

The following sections detail the above-mentioned concepts of ICN and show how they can be achieved through the information object model that involves content dissemination and retrieval.

Naming Information Objects

Retrieving content in ICN can be decoupled into two phases: content discovery and content delivery. Content discovery includes naming and addressing content objects while content

delivery defines the rules of routing and forwarding content objects in the network.

Naming content consists in identifying each content object living in the network by a globally unique name, so that users can address objects by their name to determine the location host or the nearest cache holding a copy of the queried object before the network creates a delivery path from this host to the receiver asking for content. The naming to addressing mapping is a one-to-many relation because of the use of content caching and replication in ICN.

In this section, we first focus on the naming scheme that allows to define the conceptual principles of ICN. We delay our discussion about addressing and forwarding for content delivery in the next section.

Properties of naming

As everything is information in ICN, an emphasis has been put on naming to define the information object model. The objective of naming is not only to uniquely identify content objects in the network, but also to include important properties such as pertinence, usability, scalability and security.

- **Uniqueness:** Information objects have to be named in a globally unique way in the network. This uniqueness is necessary to route content by their name.
- **Persistence and location-independence:** The name is invariant and independent of the location of the host that stores data. Content can then be replicated or moved from one hosting location to another in the network topology without service disruption, provided that the provider continues to serve it.
- **Usability and scalability:** As a network of information, we can expect to have a huge number of data objects in the network, which are interconnected and can have multiple

interdependencies between them. Moreover, objects are not necessarily static and can be mutable, e.g. being fragmented in small segments or evolving from one version to another at some time in the future such as a weather RSS feed. In addition to be scalable, the naming scheme should then be usable for dynamic objects and should also allow deletion of objects (e.g. by authorizing them to be invalidated themselves with a TTL mechanism).

- Security: ICN embeds security directly in content rather than assuming trust on users or securing the communication channels that deliver it. Content-based security means that there is a cryptographic binding between content and its name to ensure that information objects are self-certifiable in the sense that when a user receives its queried object, it can rapidly verify that the object was truly originated by the genuine provider. This binding generally consists of a hash of the content provider's private key in the name of the content object. This allows the authentication of the content provider by creating a signature of the data object using the provider's private key, which can be verified with its self-certificated public key sent as metadata to receivers along with the data object. Confidentiality and integrity of content are also guaranteed by a public key encryption. An important challenge for self-certification of content is to handle revocation of objects when the public key is compromised or data is updated.

While the name structure is often made up of several parts to reflect these properties, there are actually two types of namespace in the ICN literature: hierarchical and flat. Hierarchical names, as proposed by Van Jacobson, make it possible to use the concepts of IP routing lookups based on longest prefix matching and to organize content in a way similar to DNS, i.e. aggregating

content at different levels over trusted domains to improve security and the authenticity of the provenance of content which is useful for Digital Rights Management (e.g. `parc.com/videos/widgetA.mpg/`). However, the binding between hierarchical names and administrative domains compromises the persistence of content objects. In this case, a flat namespace may be preferred.

The presented properties of naming allow us to assert the following principles for ICN:

- Resiliency to service disruption and network failure by replicating data in various points in the network.
- Increased performance by enabling caching for content retrieval throughout the network (the nearest copy of a data object is returned to users).
- Native support of multicast, mobility and multihoming, rather than add-on solutions. This is achieved with the decoupling between naming and location of data.

Metadata

In addition to naming, content objects also have attributes represented by means of metadata. The concept of metadata is important in ICN to provide information or description about a content object or its relations with other objects. For example, metadata associated to a photo can give indications about the image resolution, the author, the date or any other data inserted by software.

Metadata attributes are used for several purposes in ICN. The semantic nature of these attributes is firstly profitable for applications to manage content objects and to understand how they can be

used. For example, in the Van Jacobson's implementation of VoIP using the content-centric paradigm [3], metadata contain descriptive information about users involved in the conversation. As object names are numerical identifiers, it is usually necessary in ICN to let users to make keyword or description based searches for content. This means that a resolution often exists at the application level to map human understandable attributes into object names. Most search engines use metadata associated to objects to implement this resolution in a distributed way.

Beyond semantic meanings, metadata can provide cryptographic inputs to perform more complex security checks on content objects. A simple example consists in including in data objects metadata carrying the content provider's public key combined with a digital signature so that receivers can assert that content did exactly come from the provider. The network can also rely on metadata to perform QoS guarantees on information objects, especially coming from real-time streaming multimedia applications (e.g. prioritizing content objects from a VoIP session). In particular, we can define network metadata to perform some network access control or to collect statistics on traffic usage to monitor network health. Generally speaking, metadata provide an efficient tool for ICN to supervise the network.

There are different ways for implementing metadata in ICN. As metadata are by definition nothing other than data about data, any content object can be metadata for other content objects, and in other words, metadata are simply content objects as any other ones in the network [6]. This design consideration involves an additional class of identifiers to use in conjunction with the name of content so that we can know which metadata items are related to a content object. However, for simplicity, metadata are often considered as labels at the same level as the name.

Addressing and Forwarding for Content Delivery

Addressing defines the reachability of information objects in the network by mapping names to hosting locations. In ICN literature, it is also referred to as name resolution to make an analogy with the current DNS resolution that resolves host names to their corresponding IP.

Name resolution is equivalent to the network layer of today's Internet and includes content location (routing queries) and forwarding for content dissemination and content delivery. As a network of information revolves around content producers and consumers, the publish/subscribe communication paradigm [1] which decouples the sender from the receiver appears as the most relevant networking concept from which ICN takes its inspiration. Most of the information-centric architectures proposed in the research community fairly reuse the principles of the pub/sub networking communication to implement content forwarding, which is thus based on two functional steps we can mention as REGISTER (or PUBLISH) and FIND (or SUBSCRIBE). It is not clear if the ICN paradigm can (or should) replace or not the IP layer completely (clean slate approach for the future Internet). ICN can be actually overlaid over any forwarding layer, including IP itself. While the debate is ongoing, current technologies proposed to implement routing in ICN are overlaid over IP facilities and for the time being, can be classified into two categories:

- One-step name resolution: name-based routing [4], FIB-based forwarding [2];
- Two-step name resolution: Name Resolver Service for translating an object ID into one or multiple source locators [9], Rendezvous-based communication [6] or probabilistic routing (e.g. using Bloom Filters) [11];

In an one-step name resolution, the request message for an information object is directly

forwarded from the requester to the source or any cache which can serve the query. Instead of exchanging route information based on IP prefix advertisements as does current BGP routing, routers in the one-step name resolution advertise names of information objects, so that each router can determine efficient routes to objects. Forwarding then consists in finding the better match between the name of the requested information object and the entries of the routing tables.

The second two-step approach, based on the publish/subscribe model, relies on an intermediate mapping proxy to resolve object names into network identifiers (or locators) which are used for routing queries to content sources. Users send a subscription message to the proxy with the name of the desired information object. The proxy is then responsible for getting the requested object back to users. Routing protocols are generally based on Distributed Hash Tables (DHTs) to distribute this name to locator mapping function over the global network.

Examples of these two different routing approaches will be detailed in the when describing the architectures from research projects related to ICN.

Transport Layer for Information Objects

In the current Internet architecture, transport layer functions such as error detection, lost data retransmission, bandwidth management, flow control or congestion avoidance are implemented at the endhost level as end-to-end communication processes. This is inconsistent with the ICN paradigm in which the role of endhosts is very different compared to traditional IP networks since communication sessions are only information-centric, disregarding the involved endpoints. As a consequence, if naming and addressing content as discussed in the previous sections allows

us to implement the functions of the network and lower layers required for content-centric networking, the transport layer also need to be considered to completely remove the dependence on endpoints. Although some proposals exist in the literature to translate transport functions in ICN, discussions are still opened. For example, senders and receivers are decoupled in ICN, and because of caching, a requester can receive its stream of queried content from multiple sources in an unpredictable way. The challenge is then how to perform transport control per data source under uncertainty that there is no way to know these sources in advance. One solution could be to let the receiver to control congestion avoidance, which is estimated from the source feedback, as described in [13].

The IP layer implements datagram fragmentation, so that fragmented packets can pass through a link with a maximum transmission unit (MTU) smaller than the original datagram size. In ICN, the equivalent notion is chunking which means that a source can serve content in a series of chunks. A chunk is typically the smallest identifiable piece of a content object, but chunking is more than fragmentation in the sense that a chunk can be split up into smaller fragments for the transport over the network. One of the important factors determining the way that an information object is delivered to a receiver is the presence of chunks from several locations (e.g. caches) in the network and the reassembly of these chunks at the receiver. This leads to additional issues that need to be resolved in the ICN transport layer (e.g. receiver-driven synchronization and flow control between different chunk delivering sources, cache policies for chunks, security checks on chunks, should a chunk be identifiable, addressable and/or authenticatable? etc).

Overview of Research Activities on ICN

This section describes the most relevant results of different research activities on the topic in the literature.

CCN

CCN (Content-Centric Networking) [2] [3], designed at Palo Alto Research Center (PARC), is one of the pioneers to promote the ICN paradigm.

In CCN, object names are hierarchically organized in a lexicographic ordered tree. Leaves correspond to content of interest, and each internal node represents the common name prefix shared by a collection of content objects. While using this naming tree, CCN can support dynamically-managed contents by authorizing users to make requests based on a name prefix for one content that may not have not been created yet, allowing publishers to generate that content on demand. Objects can also be queried without knowing their full name. Suppose that a video, identified by the name prefix `/parc.com/videos/WidgetA.mpg` is split into small data chunks which are separately named by adding the suffixes `/$version number, chunk number$`.

These chunks, taken as a whole, form the video file. Users usually do not know the full names of different video segments and thereby rely on the name prefix to trigger the video download.

CCN communications are based on two packet types, Interest and Data, identified by the full or relative name of queried content. A consumer asks for one content by sending an Interest in that content. A CCN node hearing this Interest forwards it to its neighbours unless it owns the queried content and can immediately serve the consumer with a Data message. The latter case means that the name in the Interest is a prefix of the content name in the Data packet.

CCN forwarding is actually similar to the IP forwarding plane for fast lookup of content names

in the Interest packets. The Figure 3 describes the functional parts of a CCN node: the FIB to find the appropriate interface(s) to which arriving Interest packets should be forwarded to reach the providers of queried content, a Content Store that is the LRU buffer memory for content caching, and a Pending Interest Table (PIT) to keep track of the inbound interfaces of received Interest packets so that a Data packet sent back as a response to an Interest registered in the PIT table will be delivered to the right interface(s).

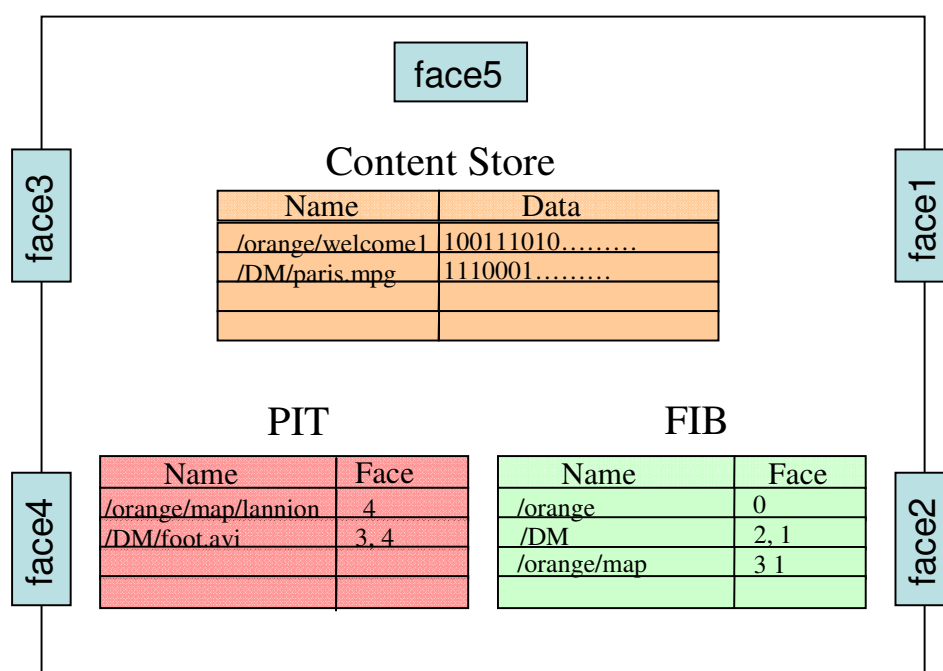


Figure 3: The CCN node

The FIB of a CCN node is populated with name prefix announcements encoded as type-length-value (TLV) elements within IP routing protocols. As messages of unrecognized types in the TLV scheme are currently ignored by these protocols, CCN nodes can be deployed with existing IGP or EGP routers.

PARC has started the development of a prototype, called CCNx [18], available as an open source

software. The PARC's idea is to develop a community around CCN to attract many people to work on this demonstrator, to make it evolve and being thus as a de facto standard for ICN solutions.

Named Data Networking (NDN) is a new project under NSF Future Internet Architecture (FIA) program based on CCN paradigm [15]. PARC also contributes to this project. The research activities for this project are related to routing and forwarding improvements for CCN. The first goal is to extend the existing routing protocols like BGP or OSPF with the compatibility of named content prefixes. Then, two approaches will be proposed to achieve routing scalability.

- Propose a provider assigned name, which looks like the provider assigned IP address, to achieve the aggregation.
- The user or application may choose a name which is easy to remember and NDN allows the user-selected name, so a mapping service is still necessary for mapping the user selected name to the provider-assigned name, like DNS domain name resolution.

Since the NDN uses CCN-like longest prefix matching for forwarding based on information stored in the tables FIB and PIT, this computationally intensive task, similar to current IP Lookup, is likely to be the performance bottleneck along the forwarding path due to a potentially huge number of objects in the network. So how to keep an effective forwarding with acceptable performance is an important challenge for the NDN network architecture. The goal consists in finding a trad-off between high speed longest name prefix match ability and an efficient content storing/deleting/replacing performance in FIB and PIT.

DONA

The DONA (Data-Oriented Network Architecture) [4] proposes a flat naming scheme and a name resolution using a distributed set of network entities, called Resolution Handlers, with caching capabilities to route requests towards the nearest copies of data.

Each information object in DONA belongs to a principal, uniquely identified by a public-private key pair. Names are then defined as $P:L$, where P is a hash value of the principal's public key and L is a label attributed by the principal to ensure that names are globally unique in the network.

The principal is responsible for organizing the structure of labels when naming the objects that it manages.

Content delivery in DONA relies on an overlay network of Resolution Handlers (RHs) using a route-by-name resolution through two primitives: $FIND(P:L)$ and $REGISTER(P:L)$. Each RH maintains a registration table containing entries of the form $\langle (P:L) \text{ or } (P:*) \text{, next-hop RH, distance to a copy} \rangle$, where $P:*$ means all data associated with the principal P . When a user asks for content $P:L$, it sends a $FIND(P:L)$ message to its local RH. If the registration table contains an entry for $(P:L)$, the message is forwarded to the corresponding next-hop RH; but if there exist several entries for $(P:L)$, the selection depends on local policies or the nearest copy of data, and if entries for both $(P:L)$ and $(P:*)$ exist, the longest matching label $(P:L)$ will be used. Now, if any entry exists for $(P:L)$ (or $(P:*)$), the $FIND$ message is sent to the local RH's parent. When a host is authorized by a principal P to serve a content object named $P:L$, it sends a $REGISTER(P:L)$ to its local RH (or $(P:*)$ if it is authorized to serve all principal's data). If any entry exists for $(P:L)$ in the registration table or if the new $REGISTER$ comes from a copy closer than an existing entry, the local RH creates (or updates) the entry for $(P:L)$ and forwards the

register message to its parents and only to its peers if local policies match. Otherwise, the local RH discards the REGISTER. Before forwarding the REGISTER message to the next-hop RH, the local RH adds to the message header: its signature to protect the authenticity of the message and the distance to the previous-hop RH to keep track of the total distance to the copy of data.

Content forwarding in DONA is based on domain-level label switching. Within each domain, hosts are addressed with a label that is only unique to this domain. When a node sends a FIND message for some content, it appends its domain-specific label to the message header. When forwarding the FIND message from one RH to another RH, labels are pushed onto the stack within the header, so that reversing these labels allows the hosting entity to send queried data back to the user.

To implement caching, an RH has to replace the source address (or label) of an incoming FIND message with its address before forwarding it to the next-hop, so that the data object sent back as a response to the FIND request will be delivered to the RH that can then store data in its cache. Whenever cache is activated, upon reception of a FIND, if the RH has the queried data in its cache, it can serve the client directly. TTL field or metadata can also be associated with FIND messages to indicate how long a data object should be valid and to request updates when the timer expires.

PSIRP

The PSIRP (Publish-Subscribe Internet Routing Paradigm) [5] [6] [7] is a European FP7 research project, started in January 2008 and ended in June 2010, promoting an information-centric

publish-subscribe networking paradigm for the Future Internet.

The PSIRP architecture (figure 4) relies on a rendezvous system consisting of a set of interconnected physical devices, called Rendezvous Points (RPs), providing rendezvous functionalities for subscription and publication matching.

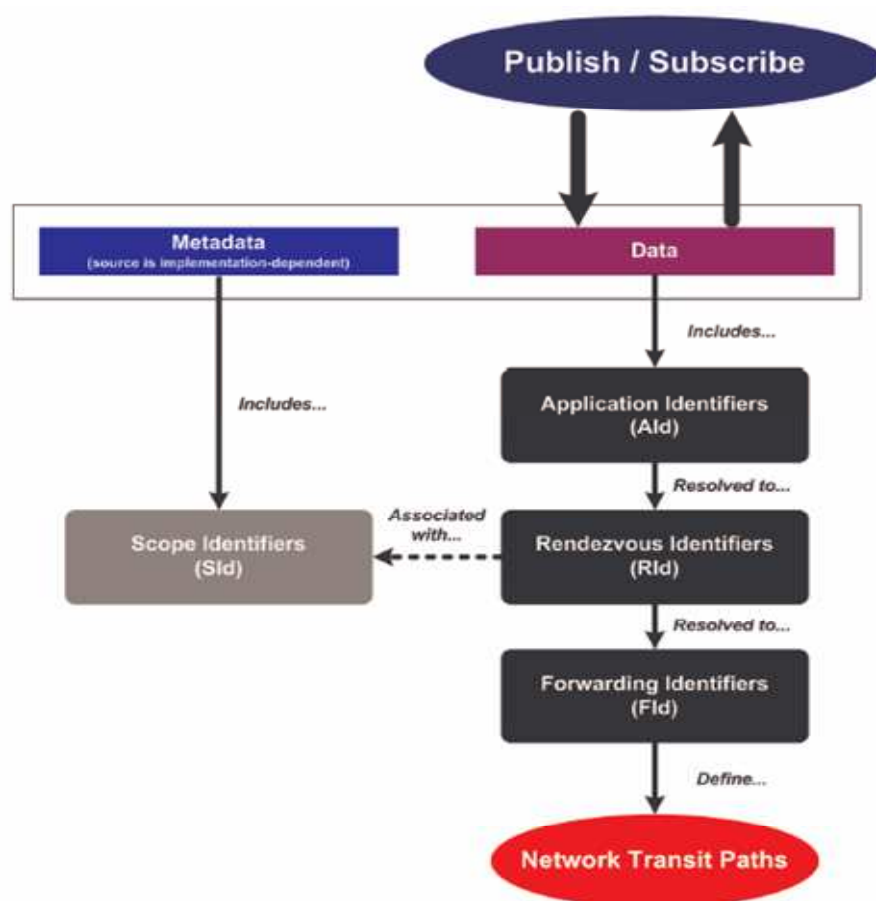


figure4: Components of the PSIRP architecture

At the highest level, each data object is related to an application and can be optionally named with an application identifier that is not required to be unique or universal in the network. But when communicating over the network, applications need to resolve application identifiers into a unique and persistent identification. PSIRP relies on a special class of identifiers, called

rendezvous identifiers, to name objects in a distributed way, using a flat naming structure. These identifiers could also include some security properties such as data integrity, self-certification, owner authentication or access control, using public-key cryptography whose the corresponding private key is only known to the publisher and the public key is widely distributed to receivers. In addition to have a globally unique rendezvous identifier, PSIRP data objects are also defined within a scope, so that the access to a particular set of data objects can be restricted to only a particular set of authorized users (publishers or subscribers). Scoping is also used for selecting the responsible RPs for Rendezvous identifiers.

PSIRP also makes use of metadata to provide additional semantic information on data objects; application-level metadata such as type of the document, size, author, access rights, caching, etc.; but also network-level metadata such as network access control, flow control, error or congestion notification, etc., are possible. For receivers, metadata can be used to describe what they want and to indicate their preferences in how to receive data that they request.

PSIRP forwarding uses a domain-level label switching based on forwarding identifiers, also called labels. A domain allocates a forwarding label for each rendezvous identifier that traverses the domain, and can aggregate several rendezvous identifiers under the same label for scalability. Typically, if some content is named with a rendezvous identifier RId, the publisher will send a publication related to RId, along with metadata and one or more scope identifiers Sid to limit the reachability of information. When receiving the publication, a rendezvous point forwards it to its neighbours until we reach the rendezvous node where the scope identifier Sid of the publication is registered, and forwarding instructions are updated in the routers on the way. Once receivers subscribe to the rendezvous identifier RId, subscriptions will be relayed from the local

rendezvous point to the rendezvous point where the publisher registers. As a subscription progresses to this rendezvous point, labels are appended to the subscription header, so that reversing these labels in the stack allows the publisher to send data back to the receiver. A forwarding tree is thus constructed for this rendezvous identifier: it represents the delivery paths over which content can be conveyed to receivers.

Whenever a RP receives multiples subscriptions with the same Rid, it only forwards a single subscription to the next-hop RP, so that there is also one single returned data flow from the publisher to this RP. This allows support of multicast forwarding trees.

Whenever a RP receives published data, it can store it in its cache to serve subsequent subscribers rapidly with low-latency. Caching makes it possible to guarantee high availability of content and services.

The PSIRP has ended, but the work will continue in PURSUIT, which is also a FP7 European project [15] [16]. This new project proposes to refine the PSIRP architecture, for both wireless and wireline networks. It is expected to handle further studies on important aspects such as caching mechanisms for better resource utilization and management, transport issues and enhancing mobility with network topologies based on rendez-vous points. Another main objective of PURSUIT is to provide various prototypes and development APIs for specifying a reference implementation of information-oriented protocols proposed in PSIRP.

NETINF

NetInf (Network of Information) [8] [9] is a network architecture, proposed by FP7 Project 4WARD.

NetInf makes a clear distinction between Information Objects (IO) which represent a piece of

content or information, identified by a globally unique identifier, and bit-level Objects (BO) which are the basic data object itself. The IO is composed of 3 fields: the identifier of the content, a set of metadata and the BO. The identifier contains the type of content and the hash of the owner's (or publisher) public key, hence providing authentication. The Metadata field provides semantic information about the IO, includes the security attributes, such public keys, content hashes and certificates and can be used by the search service.

Different versions of the same IO may have different authors. Authors are authorized to sign and modify IOs, or delegate modification and signing to other authors. Publishers are not authorized to modify or sign IOs but are authorized by the owners to distribute the IOs. In order to register/unregister an IO, an owner or publisher provides a signature of the IO and the registration time.

The NetInf Name Resolution System (NRS) takes as input either an identifier or a set of attributes describing some properties of the searched object and returns a set of binding records for IOs that matches the input. The IOs include a reference that directly or indirectly can be used to retrieve the BO. This means that a two-step resolution is possible where the application or user may choose in the list of returned IOs, which to select for asking corresponding BOs (contents), based different costs or criteria for retrieval (download speed, definition, quality...).

In NetInf, name persistency is ensured as the ID is independent of the location, only the Name Resolution System's entry will be updated. NetInf also allows persistency regarding evolution of content through versioning which is useful for dynamic objects such as streams.

Netinf defines name resolution zones to separate level of trust for different operators and/or customers. Each zone is responsible for persistently storing a BO with corresponding identifying IO. It is also responsible for caching strategy (pre-population...).

The main routing mechanism used in Netinf is Multiple DHT (MDHT) which is basically a way to implement recursive lookups. When a client asks for an object, a first DHT lookup is made at the first level (e.g., its access networks zone). If it is not found, another DHT lookup is issued at a upper level (e.g., POP zone). If it is still not found, another DHT lookup is made at a upper level (e.g., domain level), etc. When the DHT lookup is successful at a given level, the result is returned to the client. It is to be noted, that despite the hop-by-hop routing and local resolution this provides, the top DHT level has to contain bindings for all data registered in a domain, with possible scalability and possibly performance issues.

The SAIL project [17] is a follow-up project for 4WARD and continues work in the ICN activities. Within the SAIL project, the Netinf solution is under evolution to take into consideration some issues, such as the scalability of the solution, the way to resolve names, etc. In SAIL, an objective is also to develop a prototype for the netinf solution, called OpenNetInf [18], a published and open source software. Like CCNx, the idea is to attract people working for this prototype.

Plexus

The Plexus [10] solution could not be really considered as an ICN solution, because it is not designed with this objective, but it is worth mentioning it since it can help in designing alternative solutions to current ICN proposals, mainly for content naming and retrieving in the network.

Plexus, is based on previous work called DPMS [11], and is designed with the primary intention to look for contents in a P2P system without exact matching semantic but rather searching a pattern in the content names. Indeed they argue their work telling that current P2P networks need

to know exactly the names of the searched contents (e.g. searching "Lord of the Rings" or "Lord Rings" will lead to different identifiers after the hashing function and thus no mean to get it if we have the wrong name as the input parameter). In their approach, they propose to split the names into trigrams and use a Bloom Filter to set bits to indicate whether this trigram is present or not in the name of the content. Thus the user can find the content even if the searched name is not exactly the one mentioned when inserting the content in the network. The following picture 5 presents this concept using the Bloom Filter.

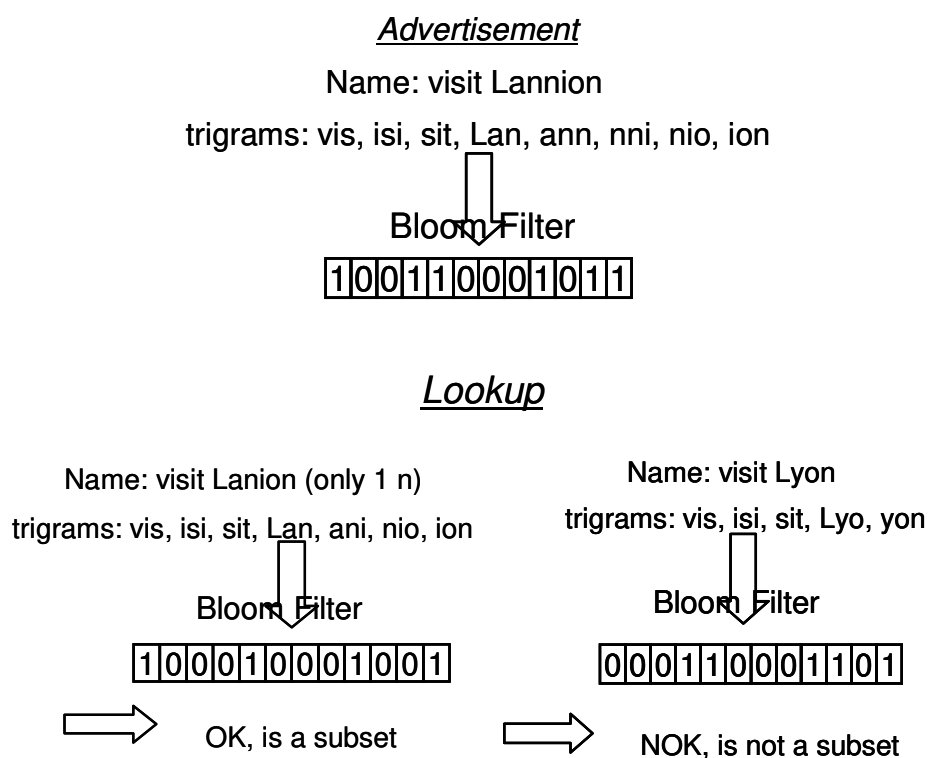


Figure 5: Bloom Filter in Plexus

It is well-known that with Bloom filters there may be false-positive (well-known formula), but it

is lower to expect to have good results in most of the times.

This Plexus system could be seen as a flat system, such as others DHT systems and could lead to inefficient search. However, in the design of DPMS, the authors define a kind of trees: propagation tree, where patterns generated by one peer are propagated in the network and the aggregation tree, where nodes aggregate patterns coming from different nodes below. This second tree avoids a large volume of advertisement data but due to aggregation, information content reduces as we go up in the tree. It is then a trade-off.

We can see that this system could be adapted to fit with ICN requirements where content could be advertised via the trigrams and bloom filters to intermediate nodes that could store advertised patterns and thus be able to route to the appropriate sources when an end-user is looking for a given content, using the same mechanism

Challenges

The aforementioned ICN solutions are still in their early stage and do not address efficiently all the features the future Internet should provide. In particular, some important aspects such as security, QoS considerations, scalability and reliability, and network management need to be addressed in better details by an information-centric internetworking architecture. We are conscious that there are many other issues, not discussed here, that also need attention, but we think the following challenges are essential for a first approach of the ICN model;

Security: Security is an essential part of the ICN paradigm. Instead of securing forwarding paths, the ICN model enforces security directly on content objects by placing provision for protection of the authenticity of content objects or against unauthorized access and privacy issues. The use of self-certifying naming based on public-key cryptography implies some issues

such as key compromise, revocation and management that need to be resolved, sometimes by external public key infrastructures and in other cases by the naming scheme itself. For example, CCN advocates for a distribution of keys as a specific type of content objects using the SDSI/SPKI model, and in DONA, key revocations for a principal P can be handled by publications of the form P:L for some reserved name L. While ICN provides content-level security based on cryptographic binding between object naming and corresponding content, there remain many other security concerns to be addressed, and most of them are related to denial of service and fuzzy pub/sub message flooding or disturbance. As a possible solution, it may be interesting to explore how far probabilistic data streaming algorithms can be used in ICN to enforce rate limiting or signature analysis on content objects or messages exchanged between users. Policy-based routing with access control carried by network metadata or directly as labels within naming may also contribute to guarantee better security in ICN. Van Jacobson et al. mention for instance the notion of content firewall to filter objects based on specific range of the namespace in addition to digital signatures of Interest requests.

QoS considerations: Quality of Service refers to the ability to provide the optimal use of shared resources by scheduling these resources among different classes of traffic carried by the network. The primary goal is to apply different priority to different flows to increase the utility of the network and to guarantee the required level of performance. From the ICN viewpoint, the notion of end-to-end flow does not exist and any communication is an exchange of named content between producers and consumers. The fundamental concern in ICN is then how to prioritize information objects and how to define criteria or mechanisms that allow the network to provide QoS differentiation between objects. In other words, we need to formalize the unifying thread

between information objects to define what a flow means in the ICN paradigm, or maybe we should define a novel notion of Quality of Information for a clean slate thinking. The challenge is especially difficult since contents are pervasive, replicated, cached, distributed and may be accessed or originated from many sources (possibly from different access networks, having different capabilities) via many different paths, traversing ICN routers and non ICN routers as well. None of the ICN solutions in the literature clearly addresses the QoS issue or proposes a study about resource management. Some just argue that the quality will be better because the content will be retrieved from closer sources; others say that ICN will improve QoS but without detailing why and how. A satisfactory solution to ensure quality of service in ICN, depending on the context that binds content objects, might be based in defining a new identification scheme that could allow to embed QoS (and possibly context) information in content identifiers and allow its routing in the ICN. It could also be integrated with routing/queuing/caching mechanisms that enable the delivery with respect to the required QoS. Dynamic routing solutions based on the naming of the content as well as using network topology or other concepts such as network metadata may also be used to help in the distribution of information with the required Quality of Service.

Scalability and reliability: The main challenge ICN will have to face is to distribute billions of objects to billions of interconnected devices. A related problem is how to synchronize between content to avoid naming conflict or how to guarantee that names are globally unique. ICN must propose an efficient procedure for name allocation that should be distributed and self-manageable to cope with the huge number of objects in the network and to favour dynamically-generated content. This reliability feature for naming has however a price regarding scalability.

Due to the large spatial distribution of objects in the network, scalability is actually an important challenge for content storage and cache policy [13] as well as for naming, whether it is a flat naming as in DONA, a recursive naming as in CCN or a Multiple DHT as in NetInf where at the upper level every IO must be referenced and a total number of objects to be referenced reaching a factor of 10 over current addressing space. Scalability of routing (or name resolution capacity) in the ICN paradigm is also a concern. For example in Van Jacobson's CCN, unbounded namespace could result in order of magnitude more prefixes in the FIB table. A CCN name is actually designed as a URL-like chain of characters, so how to name an object is quite an open issue, and how to efficiently build the routing and forwarding tables (PIT & FIB) for fast and efficient lookups needs to be analyzed to ensure system scalability. In particular, we need to find out if existing IP lookup techniques (such as hash-based design, using TCAM, trie-based schemes, etc.) could be brought to the CCN network. All these aspects should be carried out in a further research work. Proposed routing solutions based on DHT like NetInf or PSIRP may also suffer from the burden of resource discovery overheads. In actual fact, an important open question is the potential issues of backward compatibility with the current Internet architecture: how the ICN forwarding layer would behave and scale over the IP layer since it seems hard to do without IP?

Network Management and better manageability: Network management is lacking in most of the existing ICN solutions and is only mentioned in PSIRP but without further description. As the term content object is generic in ICN, it does not necessary refer to application-level data, but can also name network entities (host, link, domain, etc) or more generally anything material or mental that may be perceived by the senses. Based on this extended definition, PSIRP argues that

we can reuse the ICN model for network monitoring and management. It seems promising to explore deeply into this direction as ICN can be used to capture the knowledge about context information and may thus be more efficient to design networks with ease of configuration for management, leading to self-configured and self-optimized networking. Translating network protocols for collecting IP traffic information such as NetFlow or SNMP into the specific ICN context can also be of interest, for example for content-based network billing. This appears particularly problematic since identity is separated from locator of content providers or consumers.

Conclusion

In this paper, we have introduced the new promising Information Centric Networking paradigm that aims at overcoming limitations of the current Internet. We then presented and compared the main ICN solutions that currently exist, some having demonstrators proving the feasibility of the solution. Even if, for pragmatic issues and more short-term deployment, those prototypes are currently running over Internet, the defined ICN solutions have been designed to be run without Internet, but by replacing it (as the Internet did with the telephony networks). We then may imagine that future Internet will be based on ICN concept in some years; that's the reason why network operators as well as network equipment providers carefully investigate this paradigm, directly and via collaborative projects with universities. But, even if existing solutions are promising, several issues and challenges we have identified at the end of this paper, still remain to be addressed accurately before a real and successful deployment can happen.

REFERENCES

- [1] P. Eugster, P. Felber, R. Guerraoui and A.-M. Kermarrec, The many faces of publish/subscribe, *ACM Computing Surveys*, 35(2):114-131, 2003.
- [2] V. Jacobson, D.K. Smetters, J.D Thornton, M. Plass, N. Briggs and R. L. Braynard, Networking Named Content. In *Proceedings of ACM CoNEXT 2009*, Dec. 2009.
- [3] V. Jacobson, D. K. Smetters, N. H. Briggs, M. F. Plass, P. Stewart, J. D. Thornton and R. L. Braynard, VoCCN: Voice-over Content-Centric Networks. In *Proceedings of ACM ReArch'09*, Dec. 2009.
- [4] 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 *Proceedings of SIGCOMM'07*, Aug. 2007.
- [5] PSIRP project, Publish-Subscribe Internet Routing Paradigm, <http://www.psirp.org/>
- [6] Conceptual Architecture of PSIRP Including Subcomponent Descriptions, Public Deliverable (D2.2) of the PSIRP project, <http://www.psirp.org/files/Deliverables/FP7-INFISO-ICT-216173-PSIRP-D2.2 ConceptualArchitecture v1.1.pdf>
- [7] P. Jokela, A. Zahemszky, C. Esteve, S. Arianfar, and P. Nikander, LIPSIN: Line Speed Publish/Subscribe Inter-Networking. In *Proceedings of SIGCOMM'09*, Aug. 2009.
- [8] Netinf Website, Network of Information, <http://www.netinf.org>
- [9] Deliverable D6.2 of the 4WARD project, M.D. Ambrosio, M. Marchisio, V. Vercellone et al., Second NetInf Architecture Description, Jan.2010, [http://www.4ward-project.eu/index.php?s=file download&id=70](http://www.4ward-project.eu/index.php?s=file%20download%26id%3D70)
- [10] R. Ahmed and R. Boutaba, Plexus: a scalable peer-to-peer protocol enabling efficient subset search. In *IEEE/ACM Transactions on Networking*, Vol. 17, Issue 1, PP. 130-143, February

2009

[11] R. Ahmed and R. Boutaba, Distributed Pattern Matching: A Key to Flexible and Efficient P2P Search. In IEEE Journal on Selected Areas in Communications (JSAC) issue on Peer-to-Peer Communications and Applications, Vol. 25 (1), pp. 73-83, January 2007

[12] S. Arianfar, J. Ott, L. Eggert, P. Nikander and W. Wong, A Transport Protocol for Content-Centric Networks. In IEEE International on Network Protocols (ICNP) Poster Session, Kyoto, Japan, October 5-8, 2010.

[13] S. Arianfar, P. Nikander and J. Ott, On content-centric router design and implications. In Proceedings of ReArch 2010, Philadelphia, USA.

[14] PARC Technical Report NDN-0001, "Named Data Networking (NDN) Project", October 2010.

[15] N. Fotiou, P. Nikander, D. Trossen, G.C. Polyzos, Developing Information Networking Further: From PSIRP to PURSUIT, International ICST Conference on Broadband Communications, Networks, and Systems (BROADNETS), October 2010

[16] www.fp7-pursuit.eu/

[17] <http://www.sail-project.eu/>

[18] <http://www.ccnx.org/>

[19] <http://code.google.com/p/opennetinf/>