

ICNRG
Internet-Draft
Intended status: Informational
Expires: January 2, 2015

D. Corujo
Instituto de Telecomunicacoes
K. Pentikousis
EICT
I. Vidal
J. Garcia-Reinoso
UC3M
S. Lederer
Alpen-Adria Universitat Klagenfurt
S. Spirou
Intracom Telecom
C. Westphal
Huawei
July 1, 2014

ICN Management Considerations
draft-corujo-icn-mgmt-05

Abstract

ICN has been proposing and evaluating novel ways for reaching on-line content in upcoming Future Internet environments, leveraging intrinsic capabilities such as naming, caching and built-in security. In order to fully realize the capabilities and vision provided by ICN, supportive management procedures need to be ensured, providing the architectures, and the elements that figure in them, with the means to facilitate the delivery of content and the operation of the network. In the current Internet, these management aspects have been being developed and enhanced in parallel to the existing data protocol and mechanisms, resulting in a plethora of different and hard-to-integrate approaches, but still fulfil indispensable roles and actions for the operation and well-being of the network. We consider that the availability of management mechanisms for ICN will foster deployment and, as such, should be tackled still in the design and experimentation phases. In this way, this document addresses and identifies ICN management considerations, under two different settings: a) achieving management operations using ICN-based mechanisms and, b) how to manage ICN procedures themselves. The ultimate goal is to provide the necessary breadth to establish management mechanisms deployment guidelines in a common way throughout the existing ICN ecosystem of architectures.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 2, 2015.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. ICN Management Approaches	4
2.1. ICN-assisted Management	4
2.1.1. Video Adaptation	4
2.1.1.1. Adaptive Delivery of Multimedia Content in ICN	4
2.1.2. Content Management	5
2.1.3. Network Policies	7
2.1.3.1. NetInf Management Considerations	7
2.1.4. Resource Management	8
2.2. Management for ICN Aspects	10
2.2.1. Caching	10
2.2.2. Information Freshness	11
2.3. Hybrid Approaches	11
2.3.1. Face Management	11
3. Acknowledgements	14
4. IANA Considerations	14
5. Security Considerations	14
6. Informative References	14
Authors' Addresses	17

1. Introduction

Information-centric networking (ICN) enables new ideas for naming and addressing, privacy, security, and trust, and should also lead us to think new ways for deploying, operating and managing networks in the future. By default, users, programs, information objects and hosts are in general untrustworthy and mobile in an information-centric network. This means that many of the assumptions in traditional network management, including all aspects of FCAPS (Fault, Configuration, Accounting, Performance, and Security) need to be rethought. However, despite the different instantiations of ICN architectures, and the plethora of novel research work built on top of them, little attention has been paid to management aspects so far. This includes both enabling "traditional" network management operations (which work well from small networks to large infrastructure networks), and supporting and optimizing intrinsic procedures of the ICN fabric.

This document aims to draw the attention of ICNMG to the importance of network management for real-world deployments. Today, network management is practically an add-on to host-centric deployments. We can do better as we move forward in ICN research considering the full range of deployments from home-office environments to challenged networks to tier-1 networks. To this end, we draft some first management considerations that, on the one hand, capitalize on ICN concepts for defining management procedures and, on the other, explore the possibilities for defining a common management framework irrespective of the ICN approach taken. We reckon that the latter is a much more formidable task and we are looking forward to tackling it together with other members of ICNMG.

We argue that addressing management at an early stage is not only important for real-world adoption and the successful future deployment of ICN, but also to deal with scenarios where management can simplify, enhance or optimize ICN network utilization and performance. The subject becomes particularly challenging, as disparate characteristics from different ICN approaches (e.g., in terms of namespace, granularity, routing, and so on) impact the definition and design of these management mechanisms. This document analyses ICN Management under three different perspectives. Firstly, in Section 2 it will provide considerations regarding the usage of ICN mechanisms for realizing management procedures. Secondly, in Section 2.2 will look into how the intrinsic procedures used for operating the ICN architecture can be managed. Finally, in Section 2.3 we will look in a combined way to the former two issues and identify the role of ICN when its own procedures will be used to manage ICN operations.

We plan to incrementally develop the draft, incorporating considerations on other ICN aspects as well as different approaches (e.g., [PURSUIT] and [NetInf]) as well as address other pertinent aspects as we receive feedback from the research group members.

2. ICN Management Approaches

In this part of the document, ICN management approaches will be addressed in respect to how ICN mechanisms can be used to realize management procedures, how to manage the specific ICN mechanisms themselves and hybrid approaches where the ICN mechanisms themselves are used to realize the management of ICN aspects.

2.1. ICN-assisted Management

This section addresses how the ICN operational mechanisms can be used to realize different kinds of network management procedures.

2.1.1. Video Adaptation

This section investigates ICN management considerations for the delivery of video data, and especially the adaptive delivery of video. From a content perspective, multimedia is omnipresent in the Internet, e.g., producing 62% of the total Internet traffic in North America's fixed access networks [GIPR2013].

Video, and multimedia content in general, has specific characteristics, which have to be considered and where network management consideration are necessary. The consumption of multimedia content comes along with timing requirements for the delivery of the content, for both, live and on-demand consumption. Long startup delays, buffering periods or poor quality, etc. should be avoided to achieve a good Quality of Experience of the consumer of the content. Of course, these requirements are heavily influenced by routing decision and caching, which are central parts of ICN, and which may be leveraged more efficiently by an intelligent network management.

2.1.1.1. Adaptive Delivery of Multimedia Content in ICN

Today's dominant streaming systems are based on the common approach of leveraging HTTP-based Internet infrastructures, which are consequently based on the Transmission Control Protocol (TCP) and the Internet Protocol (IP). Especially the adaptive multimedia streaming (AMS) via HTTP is gaining more and more momentum and resulted in the standardization of MPEG-DASH [MPEG-DASH], which stands for Dynamic Adaptive Streaming over HTTP. The basic idea of AHS is to split up the media file into segments of equal length, which can be encoded at

different resolutions, bitrates, etc. The segments are stored on conventional HTTP Web server and can be accessed through HTTP GET requests from the client. Due to this, the streaming system is pull based and the entire streaming logic is on the client side. This means that the client fully controls the bitrate of the streaming media on a per-segment basis, which has several advantages, e.g., the client knows its bandwidth requirements and capabilities best. As one can see, ICN and adaptive multimedia streaming have several elements in common, such as the client-initiated pull approach, the content being dealt with in pieces as well as the support of efficient replication and distribution of content pieces within the network. As ICN is a promising candidate for the Future Internet (FI) architecture, it is useful to investigate its suitability in combination with AMS systems and standards like MPEG-DASH as shown in [AdaptCCN][InterAdaptCCN], as well as the possibilities and benefits of intelligent network management to improve the performance of AMS in ICN as well as the resulting QoE at the client.

One of the most promising aspects in this context is the possibility of ICN to consume content from different origin nodes as well as over different network links in parallel, which can be seen as an intrinsic error resilience feature w.r.t. the network. This is a useful feature of ICN for adaptive multimedia streaming within mobile environments since most mobile devices are equipped with multiple network links. Here, a focus of ICN management could be in the load balancing of such traffic between the available links. This would increase the effective media throughput of the multimedia content, however, it could potentially lead to high variations of the resulting bandwidth which is available to the client. As DASH is designed for environments with dynamic bandwidth conditions, they can be compensated in general. However, more conservative adaptation algorithms may prevent too frequent switching between the content's bitrate representations as well as compensate short-term bandwidth drops caused by network link switches more smoothly.

2.1.2. Content Management

An ICN network aims to facilitate access to, and delivery of, information objects (content and services). Content (in particular, video) access and delivery seems to be the dominant use case in traditional, host-based networks, so ICN networking is forced to adopt content delivery as a minimum requirement. Indeed, virtually all ICN approaches so far target at least content delivery.

From the perspective of a content owner or provider, an ICN network functions essentially as a content delivery network. This creates a set of requirements for ICN. First of all, end-users and content providers alike should be able to Read (consume) a content object

available on the ICN network. In addition, content providers need the ability to Create (publish), Update, and Delete content. Finally, Accounting (logging) is necessary to support business models that typically require charging, analytics, and monitoring.

The Read operation has received the lion's share in ICN research. This is expected as content access and delivery is at the heart of ICN. Given a request for a named content object, the ICN network resolves that name to an object replica and proceeds with delivery to the end-user. Of course, different ICN approaches employ different mechanisms to achieve the Read operation. For example, name resolution can be done with a hierarchical system resembling DNS, with DHTs, or with flooding. Similarly, content delivery can be done over normal best-effort paths from the origin server, over dynamically computed provisioned paths, or from caches close to the end-user. Some approaches can even cater to mobile end-users and content hosts. ICN should be able to handle frequent Reads as well as Read spikes (flash crowds). In fact, it seems crucial for ICN's deployment chances to at least match the capabilities of incumbent content delivery systems.

ICN research has not addressed Create as much as Read, but some effort has been expended on mechanisms for publishing content. Much of this effort has focused on content naming schemes that enable global uniqueness of names and hence allow global addressing of the content objects. It has been difficult to balance human readability of names, efficiency in machine processing, and name aggregation (that can realistically enable request routing by name). Although a fully automated mechanism for (human-readable) name assignment would be desirable, so far it seems that a manual process, similar to that of domain name registration in DNS, is necessary to allocate at least namespaces. No other restrictions on naming have been seriously considered. The consensus seems to be that with ICN anyone should be able to publish anything. Content semantics are a higher layer issue. This might be a prudent approach when building a transport layer technology, but it could undermine the potential of ICN deployment. A content owner would not want copies of its content published on an ICN network under different names. In any case, once a name has been assigned, the Create operation is mainly about creating an entry in the name resolution system. This is obviously a security risk and furthermore, for highly distributed name resolution systems, it can suffer from considerable lag in availability. Fortunately, Create is a rare operation compared to Read.

Update is an operation that seeks to alter an already created object. A content provider would want to modify the data or the metadata of a published object either to rectify publication errors or to augment the object. It is debatable whether the provider should address the

later simply by creating a new object. Another use case for Update comes from the need to rebrand or alias an object when its rights have been sold to another party. Nevertheless, the Update operation has received minimal attention in ICN research. The main problem is one of consistency: once an update has been committed, an ICN network with highly distributed name resolution and content delivery (caching) would host both the old and new versions of the updated content object for some time. Security concerns for the Update and Create operations are similar. Update is normally rarer than Create, but this will not be the case for collaborative media.

Content providers may occasionally need to remove a published object. This is the goal of the Delete operation. An object might be deleted when it was published by mistake, because it's no longer useful or relevant, or because it's illegal. Consistency is a major challenge for the Delete operation as well. The high degree of distribution in ICN can sustain a network state where some data or metadata replicas of an object have been deleted, while others persist. On the other hand, this lag can be beneficial if deletion was initiated erroneously or maliciously. Like with the Update operation, Delete has not been properly investigated in ICN research. Deletes are typically less often than updates.

From the point of view of content providers and end users, an ICN network resembles a content directory and repository, with Create, Read, Update, and Delete as typical operations. As with any database system, the reliability of those operations (or transactions) depends on the properties of atomicity, consistency, isolation, and durability. The challenge for ICN research is to build systems at a massive scale that employ those properties.

2.1.3. Network Policies

Currently this section addresses Network Policies under the scope of NetInf. In future instantiations of this document, a more generic approach will be provided, besides highlighting specific ICN instantiations contributions.

2.1.3.1. NetInf Management Considerations

Early-phase work in NetInf management [NetInfSelfX] discussed a two-fold problem. The first question that arises is whether it is possible by adopting a new set of network primitives and in-network storage to usher a new type of network management. In other words, can network management become information-centric while handling often host-centric data? The second question is whether an information-centric network is more suitable for self-management mechanisms than IP-based networks are. In particular with respect to

the later, [NetInfSelfX] introduced some design considerations for adding self-management mechanisms in NetInf.

Of interest from this early work are two examples where network management can play a new role. First, network management can get involved in decisions about caching and (re)distribution of content, and not only whether an (inter)face is on or off, or what traffic limits should be enforced. Moreover, network policies can be distributed securely in the same way as other content in the network, removing the need for centralized management, and enabling improved recovery procedures. Second, network management can get involved in more intricate processes such as controlling multiaccess support, intermediating for content adaptation when deemed appropriate, and enabling richer tools for traffic engineering.

2.1.4. Resource Management

While caching has been the focus of much of the attention in ICN, one of the key advantages of the ICN architecture is that it allows a fine grained allocation of content to resources. This has been observed in [CB-TE] and [ICN-TE] for instance. Unlike IP, an ICN packet carries specific, explicit information about the content it carries. Further, this content is uniquely named, and different versions of the content will have different names.

ICN enables a shift in how to manage resources: instead of allocating open-ended flows to network resource, it allows to allocate well defined objects. This requires new network management tools beyond the current mechanisms which are specifically dedicated to ICN. NetFlow or the current TE mechanism do not take advantage of the ICN semantics, and of the benefits associated with these semantics.

In IP, a flow from a certain source address to a certain destination address can correspond to myriad potential applications: web traffic, video streaming, VoIP call all may use the same port 80 and be hosted by same servers. Therefore, providing appropriate resource to such a flow is a matter of guessing. The simple problem of identifying when a flow terminates is made unnecessarily complex in ICN: a timer is set-up, and when no packets match the flow filter, then the flow is over. Of course, multiple packets from different applications may match the same filter, and flows with different characteristics in terms of inter-arrival times could be broken down into multiple flows with an improper choice of time-out values.

In ICN, there is a unique mapping of the name to the content of the data stream going through the network. If a content object is requested, then it has well defined semantics, and the network management layer can identify exactly when the data stream starts and

ends based upon these semantics. Further, a content management layer can also learn the properties of the stream associated with a given identifier. [CB-TE] presented such a mechanism to learn the properties associated with a name, either by counting the bytes on the wire corresponding to this name, or by reading the footprint of the content with this name when stored in a cache. It is therefore possible for the management layer to gather meta-data pertaining to the content that goes through the network, and to use this meta-data to make proper resource allocations.

Of course, the resource manager should only acquire meta-data about content that is likely to be seen again (i.e., popular) or specific in any way (for instance, the name of an elephant flow). This considerably simplifies the task as the data of interest is concentrated on a few items. One potential usage of this meta-data is to keep track of what content is going through which link. In this scenario, each link keeps an aggregate tally of the amount of data that has been assigned to this link and subtracts the amount that has gone through. The resulting backlog can be then used to allocate new data streams to this or another link.

In [ICN-TE], it was shown that such a policy would significantly reduce the time spent in the network by content streams when considering a WAN topology and its corresponding end-to-end traffic matrix. The network load would stay the same when comparing with a min-MLU policy, but by splitting elephant flows across different paths, the completion time would be reduced. In the simulations of [ICN-TE], min-MLU is roughly 50% slower than a content-based policy.

This is an encouraging result and a step towards a management framework that assigns resource to content in a deterministic and fine-grained manner, unlike the probabilistic allocation of IP. The ICNMG should consider such a management framework, and evaluate the different proposals in light of this opportunity. For instance, an ICN architecture such as [PURSUIT] contains a natural mechanism to perform such allocation of content to paths as it assigns a source route to the content. On the other hand, an ICN architecture such as [NDN] needs to be expanded as the link allocation semantics are, in the current proposal, tied to the content resolution process: the interest homes into the content, and lays the reverse path for the content delivery at the same time. This semantics make the management for multiple link selection more difficult, as multiple interests would have to be sent over multiple links to provide path diversity. However, it could be an area of study for the ICNMG as solving such resource management problem would provide significant benefits to ICN architectures.

2.2. Management for ICN Aspects

This section will address management aspects for intrinsic ICN procedures.

2.2.1. Caching

Caching is a hot topic research nowadays in ICN. The challenges of caching in ICN are different than those of web caching, mainly because the former has to deal with high line rates and with a huge amount of content. Some ICN works propose to cache content in all ICN routers traversed by the data packet, in an LCE (Leave Copy Everywhere) fashion as in [NDN]. Some studies, like [L4M-ICN], have shown that other cache decision policies, focused to reduce the cache redundancy, may increase the overall caching performance. Some of these decision policies only use the local information available at the ICN routers, but others use the information available at other nodes to cache or not the incoming content. This is known as explicit cache coordination decision, and there are several proposals around this concept [ICN-CACHING]. The idea behind the explicit coordination is to exchange topological information, individual cache's state and content popularity view among a set of ICN routers, in order to coordinate caching decisions.

ICN may benefit of in-network caching, which consists of introducing content stores in ICN routers. The benefits are twofold: (1) improving the end-user experience by reducing the delay to retrieve content, and (2) reducing the overall aggregated bandwidth per request. On the other hand, caching in ICN presents several challenges like (1) centralized vs distributed management of the caches, (2) cache scalability, reducing the impact of the size of the total content catalog, (3) routing based on cache contents, (4) considering the different requirements imposed by different types of traffic (web/multimedia/IoT/etc.), etc. Most of these challenges can be solved, or at least minimized, by introducing management considerations in ICN proposals.

This way, a given ICN router may forward a request towards another router storing the requested content. In this context, the routing protocol is affected by the cache's state of surrounding neighbours. For example, in [CATT] the authors propose to distinguish between the source(s) and routers' caches that hold a copy of that content: the former paths are globally advertised, while the latter are only advertised within the router's neighbourhood. In all these cases, the use of a management framework may bring significant advantages, providing standard interfaces that allow the routers to dynamically manage their caches.

2.2.2. Information Freshness

One of the prime contributions of ICN-based designs, is the ability for the networking entities in charge of exercising the routing of the content, to actually store it, allowing it to serve content requests in a more readily fashion. However, there are scenarios where such facilities can raise issues, such as Internet of Things and Machine-to-Machine scenarios. Concretely, despite the caching capabilities of ICN contributing for, as an example, reducing the amount of networking stack fabric to be implemented in low-powered nodes and sensors, it can cause that the information consumed from caches is not up-to-date comparing to the information currently existing at the source. As an example, consider an accelerometer sensor which is providing the acceleration value of a car, and disseminates it via a ICN network towards different uses. When a consumer (e.g., a road traffic monitoring infrastructure) wishes to know the current speed of the car by requesting its name, it can be served by a stale content residing in a cache between the consumer and the information source.

These kinds of situations demand facilities and mechanisms to avoid the provision of stale content. For example, [ICN-FRESHNESS] considers the realization of an agreement mechanism, using ICN messaging exchanges, where both the source and the consumer agree on the minimal content freshness values for the information. Concretely, when the network entity determines that it has the content referring to a received name request, it will also evaluate the freshness value. If it is lower than the one previously agreed, it will then discard the content and rather forward the content request back to the source.

2.3. Hybrid Approaches

This section will analyse how ICN procedures can be used to manage ICN operations.

2.3.1. Face Management

The Named Data Networking [NDN] ICN architecture provides a new communication framework built on named data. Like other ICN counterparts, such as [NetInf], [PURSUIT] and [DONA], NDN intrinsically supports security, routing/forwarding, reliability, caching and even mobility, aiming at scalable and more efficient content-distribution than today's IP-based approaches. Fostered by an open-source implementation [CCNx], NDN has been at the heart of an active topic with several research contributions evaluating its deployment feasibility and performance in a number of scenarios [ICN-Scenarios].

NDN introduces the concept of a Strategy Layer, which can control Interest packet forwarding behavior. It basically determines which is the best interface (or set of interfaces) to send an Interest packet. The "strategy" component establishes a pre-configured algorithm for tackling Interest packet decisions, ranging from sending it sequentially on each interface until a Data packet is received, to evaluating which interfaces provide better performance (i.e., lower average RTT) in retrieving certain content (as discussed in [NDN]).

It is important to keep in mind that NDN replaces the commonly used term "interface" with the term "face", since packets can be forwarded over hardware network interfaces as well as between application interfaces, further acknowledging the information dissemination capabilities of ICN. This aspect is considered in [NDN] and [NDN-R], where programs can be associated to the NDN governing structures (like the FIB), defining configurations such as "sendToAll" and "sendToBest" with respect to managing the content reaching process. Corujo et al. [NDN-MGMT] exploit these concepts enabling management mechanisms to be deployed, and steer network operations and NDN operation, as described in the following section.

An important aspect supporting network management procedures is the interaction of network information residing at the network side with information about the network from the perspective of clients connected to it. The former includes, for instance, information stored in the network operator core about user profiles, associated policies, or data collected by the access network equipment, such as current and past traffic load levels, active flows, and maintenance information. Today, such information can be retrieved for management and operation support through dedicated signaling protocols (e.g., [RFC1157], [RFC6733]), or Operation Support Services (OSS) web services. The client point of view of the network includes information that, for example, a wireless terminal can provide, indicating wireless link quality, average return-trip times (RTT) or perceived Quality of Experience (QoE).

Both types of information can be capitalized upon allowing, for example, the network to coordinate network management procedures, considering as input information obtained from other network elements as well as from user nodes. One way to generate management information in network entities and at client nodes, as well as to consume and act upon it (i.e., using the management information exchange as a control channel) is to couple NDN nodes with Management Agent (MA) entities.

Fig. 1 (redrawn here from [NDN-MGMT] for convenience) illustrates how a MA can be deployed in both network and client entities, interfacing

with different operational aspects and protocol layers of an NDN node. By using NDN content reaching and disseminating mechanisms, management information can be consumed by the MA to steer not only the behavior of application processes and network interfaces, but also to interface with NDN supporting structures (i.e. Content Store (CS), Forward Information Base (FIB) and Pending Interest Table (PIT)). Effectively, different kinds of information can be conveyed to a network node responsible for managing the network (under different perspectives and processes), and resubmitted back towards client nodes, affecting the way applications interface with network interfaces and the NDN fabric.

NDN Fabric

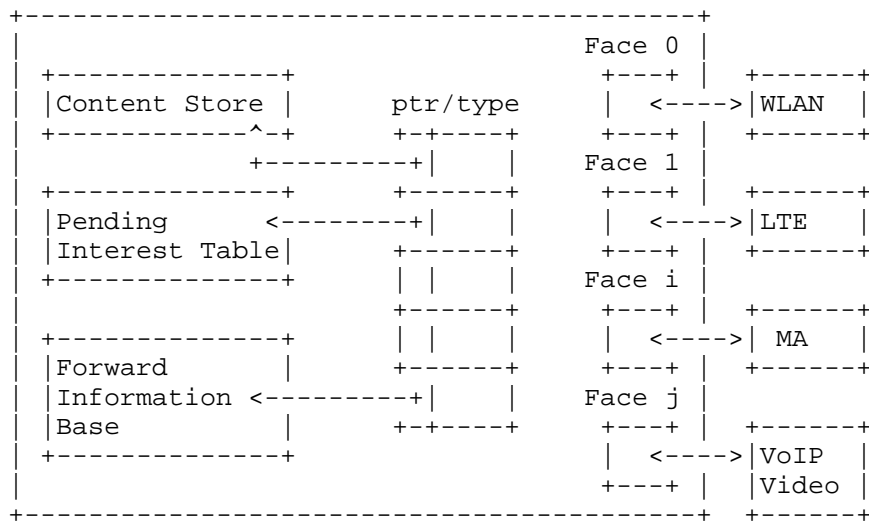


Figure 1. NDN Management Framework

MA can interface with the PIT and FIB structures, acting as a dynamic, application- and/or network-controlled interface to the strategy layer. This could also be used to direct how to forward NDN Interest and Data packets, in a configurable manner. Regarding network interfaces, the MA can interface with them not only to control (i.e., initiate wireless access scanning procedures), but also to collect information (i.e., an informational event regarding detected access points). Finally, the MA can also interface with application processes, drawing out information about the perceived QoS/QoE (e.g., lost packets or delay from a real-time video feed) and also to execute commands, such as selecting a better video codec when the network commands the video flow to be accessed from a different wireless access interface.

Conversely, MA entities residing in network equipment can provide informational events as well, but related to network-side link layer characteristics (such as number of attached nodes or load), as well as accepting commands from the network (i.e., activate maintenance procedures). Management processes residing in the network core can leverage information collected from applications, client terminals and network equipment, to drive optimization procedures. Such optimization procedures can also tap into other entities, containing complementary information such as policies and subscription information, and use it to produce an overall network decision, which can then be forwarded to multiple client nodes, in a policy enforcing way.

An important consideration from the NDN architecture, is the hierarchical namespace, allowing nodes to request and convey management data, by simply using an appropriate prefix (e.g., `ccn://domain/management/ME`).

By leveraging the NDN information-centric dissemination mechanisms to convey management information and commands as content, these management extensions inherit the intrinsic capabilities of the NDN architecture, including security and reliability, which are fundamental for management procedures.

3. Acknowledgements

This document has benefited from comments and/or text provided by the following members of ICNRG: TBD

4. IANA Considerations

This memo includes no request to IANA.

5. Security Considerations

TBD

6. Informative References

[AdaptCCN]

Lederer, S., Mueller, C., Rainer, B., Timmerer, C., and H. Hellwagner, "Adaptive Streaming over Content Centric Networks in Mobile Networks using Multiple Links", Proceedings of the IEEE International Conference on Communication (ICC), Budapest, Hungary , June 2013.

- [CATT] Eum, S., Nakauchi, K., Murata, M., Shoji, Y., and N. Nishinaga, "CATT: potential based routing with content caching for ICN", Workshop on Information-centric networking, pp 49-54 , 2012.
- [CB-TE] Chanda, A. et al., "Content Based Traffic Engineering in Software Defined Information Centric Networks", IEEE INFOCOM Workshop NOMEN , April 2013.
- [CCNx] PARC, "CCNx Project", 2013, <<http://www.ccnx.org>>.
- [DONA] Koponen, T. et al., "A Data-Oriented (and Beyond) Network Architecture", SIGCOMM, ACM , 2007.
- [GIPR2013] Sandvine, , "Global Internet Phenomena Report 1H 2013", Sandvine Intelligent Broadband Networks , 2013.
- [ICN-CACHING] Zhang, G., Li, Y., and T. Lin, "Caching in information centric networking: A survey", Computer Networks, vol. 57, no. 16, pp. 3128-3141, Nov , 2013.
- [ICN-FRESHNESS] Quevedo, J., Corujo, D., and R. Aguiar, "Consumer Driven Information Freshness Approach for Content Centric Networking", IEEE INFOCOM Workshop on Name-Oriented Mobility, Toronto, Canada, May , 2014.
- [ICN-Scenarios] Pentikousis, K., Ohlman, B., Corujo, D., and G. Boggia, "ICN Baseline Scenarios", draft-pentikousis-icn-scenarios (work in progress), February 2013.
- [ICN-TE] Su, K. et al., "On the Benefit of Information Centric Networks for Traffic Engineering", IEEE ICC , June 2014.
- [InterAdaptCCN] Grandl, R., Su, K., and C. Westphal, "On the Interaction of Adaptive Video Streaming with Content-Centric Networking", Proceedings of the 20th Packet Video Workshop 2013, San Jose, USA , December 2013.
- [L4M-ICN] Chai, W., He, D., Psaras, I., and G. Pavlou, "Cache "less for more" in information-centric networks", Lecture Notes in Computer Science Vol. 7289, Springer, pp. 27-40 , 2012.

- [MPEG-DASH] Sodagar, I., "The MPEG-DASH Standard for Multimedia Streaming Over the Internet", IEEE MultiMedia, IEEE, vol.18, no.4, pp.62-67 , 2011.
- [NDN] Jacobson, V., Smetters, D., Thornton, J., Plass, M., Briggss, N., and R. Braynard, "Networking Named Content", CoNEXT 2009, Rome , Dec 2009.
- [NDN-MGMT] Corujo, D., Vidal, I., Garcia-Reinoso, J., and R. Aguiar, "A named data networking flexible framework for management communications", Communications Magazine, IEEE , vol.50, no.12, pp.36-43 , Dec 2012.
- [NDN-R] Zhang, L. et al., "Named Data Networking (NDN) Project", NDN Report ndn-0001, Tech Report, PARC , 2010, <<http://www.named-data.net/techreport/TR001ndn-proj.pdf>>.
- [NDN-VOIP] Jacobson, V., Smetters, D., Briggss, N., Plass, M., Steward, P., and J. Thornton, "VoCCN: Voice Over Content-Centric Networks", ReARCH 2009, Rome , Dec 2009.
- [NDNFlexManager] UC3M and ITAV, "Framework for Flexible NDN Management", 2013, <<https://github.com/ndnflexmanager/framework>>.
- [NetInf] Ahlgren, B. et al., "Design considerations for a network of information", CoNEXT, Re-Arch Workshop, ACM , 2008.
- [NetInfSelfX] Pentikousis, K. et al., "Self-Management for a Network of Information", IEEE ICC Workshops 2009 , June 2009.
- [PURSUIT] Fotiou, N. et al., "Developing Information Networking Further: From PSIRP to PURSUIT", BROADNETS, ICST , 2010.
- [RFC1157] Case, J., Fedor, M., Schoffstall, M., and J. Davin, "Simple Network Management Protocol (SNMP)", STD 15, RFC 1157, May 1990.
- [RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, July 2003.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

[RFC6733] Fajardo, V., Arkko, J., Loughney, J., and G. Zorn, "Diameter Base Protocol", RFC 6733, October 2012.

Authors' Addresses

Daniel Corujo
Instituto de Telecomunicacoes
Campus Universitario de Santiago
Aveiro P-3810-193 Aveiro
Portugal

Phone: +351 234 377 900
Email: dcorujo@av.it.pt

Kostas Pentikousis
EICT GmbH
Torgauer Strabe 12-15
10829 Berlin
Germany

Email: k.pentikousis@eict.de

Ivan Vidal
UC3M
Av de la Universidad, 30
28911 Leganes, Madrid
Spain

Email: ivaldal@it.uc3m.es

Jaime Garcia-Reinoso
UC3M
Av de la Universidad, 30
28911 Leganes, Madrid
Spain

Email: jgr@it.uc3m.es

Stefan Lederer
Alpen-Adria Universitat Klagenfurt
Universitätsstrasse 65-67
Klagenfurt
Austria

Email: stefan.lederer@itec.aau.at

Spiros Spirou
Intracom Telecom
19.7 km Markopoulou Avenue
Peania 19002
Greece

Email: spis@intracom.com

Cedric Westphal
Huawei
2330 Central Expressway
Santa Clara, CA95050
USA

Email: cedric.westphal@huawei.com

ICNRG
Internet-Draft
Intended Status: Informational
Expires: April 24, 2014

K. Pentikousis, Ed.
EICT
B. Ohlman
Ericsson
E. Davies
Trinity College Dublin
S. Spirou
Intracom Telecom
G. Boggia
Politecnico di Bari
P. Mahadevan
PARC
October 21, 2013

Information-centric Networking: Evaluation Methodology
draft-irtf-icnrg-evaluation-methodology-00

Abstract

This document surveys the evaluation tools currently available to researchers in the information-centric networking (ICN) area and provides suggestions regarding methodology and metrics. Finally, this document sheds some light on the impact of ICN on network security.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

Copyright and License Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Table of Contents

1. Introduction	2
2. Evaluation Methodology	4
2.1. ICN Simulators and Testbeds	4
2.1.1. CCN and NDN	4
2.1.2. Publish/Subscribe Internet Architecture	6
2.1.3. NetInf	6
2.1.4. COMET	7
2.1.5. Large-scale Testing	7
2.2. Topology Selection	8
2.3. Traffic Load	9
2.4. Choosing Relevant Metrics	11
2.4.1. Traffic Metrics	13
2.4.2. System Metrics	15
2.5. Resource Equivalence and Tradeoffs	16
3. ICN Security Aspects	16
3.1. Authentication	17
3.2. Authorization, Access Control and Statistics	19
3.3. Privacy	19
3.4. Changes to the Network Security Threat Model	20
4. Security Considerations	21
5. IANA Considerations	21
6. Acknowledgments	21
7. Informative References	21
Authors' Addresses	26

1. Introduction

Information-centric networking (ICN) marks a fundamental shift in communications and networking. As discussed in [draft-irtf-icnrg-scenarios], the development phase that ICN is going through, and the plethora of approaches to tackle the hardest problems, make this a

very active and growing research area but, on the downside, it also makes it more difficult to compare different proposals on an equal footing. Different ICN approaches have been evaluated in the peer-reviewed literature using a mixture of theoretical analysis, simulation and emulation techniques, and empirical (testbed) measurements. These are all popular methods for evaluating network protocols, architectures, and services in the networking community. Typically, researchers follow a specific methodology based on the goal of their experiment, e.g., whether they want to evaluate scalability, quantify resource utilization, analyze economic incentives, and so on, as we have discussed earlier. In addition, though, we observe that ease and convenience of setting up and running experiments can sometimes be a factor in published evaluations.

It is worth pointing out that for well-established protocols, such as TCP, performance evaluation using actual network deployments has the benefit of realistic workloads and reflects the environment where the service or protocol will be deployed. However, results obtained in this environment are often difficult to replicate independently. Beyond this, the difficulty of deploying future Internet architectures and then engaging sufficient users to make such evaluation realistic is often prohibitive.

Moreover, for ICN in particular, it is not yet clear what qualifies as a "realistic workload". As such, trace-based analysis of ICN is in its infancy, and more work is needed towards defining characteristic workloads for ICN evaluation studies. Accordingly, the experimental process itself as well as the evaluation methodology are being actively researched for ICN architectures. Numerous factors affect the experimental results, including the topology selected, the background traffic that an application is being subjected to, network conditions such as available link capacities, link delays, and loss-rate characteristics throughout the selected topology; failure and disruption patterns; node mobility; as well as other aspects such as the diversity of devices used, and so on, as we explain in the remainder of this section.

Apart from the technical evaluation of the functionality of an ICN architecture, its future success will be largely driven by its deployability and economic viability. Thus any evaluation will also have to include an assessment of its incremental deployability in the existing network environment together with a view of how the technical functions will incentivize deployers to invest in the capabilities that allow the architecture to spread across the network.

This document incorporates input from ICNRG participants and their

corresponding text contributions, has been reviewed by several ICNMG active participants (see section 6), and represents the consensus of the research group. That said, note that this document does not constitute an IETF standard; see also [RFC5743].

The remainder of this document is organized as follows. Section 2 presents various techniques and considerations for evaluating different ICN architectures. Then, Section 3 discusses the impact of ICN on network security.

2. Evaluation Methodology

At this stage, we do not intend to develop a complete methodology or a benchmarking tool. Instead, this document proposes key guidelines alongside suggested data sets and high-level approaches that we expect to be of interest to the ICN community as a whole. Through this, researchers and practitioners alike would be able to compare and contrast different ICN designs against each other, as well as against the state of the art in host-centric solutions, and identify the respective strengths and weaknesses.

2.1. ICN Simulators and Testbeds

Since ICN is still an emerging area, the community is still in the process of developing effective evaluation environments, including simulators, emulators, and testbeds. To date, none of the available evaluation methodologies can be seen as the one and only community reference evaluation tool. Furthermore, no single environment supports all well-known ICN approaches. Simulators and emulators should be able to capture, faithfully, all features and operations of the respective ICN architecture(s). It is also essential that these tools and environments come with adequate logging facilities so that one can use them for in-depth analysis as well as debugging. Additional requirements include the ability to support mid- to large-scale experiments, the ability to quickly and correctly set various configurations and parameters, as well as to support the playback of traffic traces captured on a real testbed or network. Obviously, this does not even begin to touch upon the need for strong validation of any evaluated implementations.

The rest of this subsection summarizes the ICN simulators and testbeds currently available to the community.

2.1.1. CCN and NDN

The CCN project has open-sourced a software reference implementation of the architecture and protocol called CCNx (www.ccnx.org). CCNx is available for deployment on various operating systems and includes C and Java libraries that can be used to build CCN applications. CCN-lite (www.ccn-lite.net) is a lightweight implementation of the CCN protocol, supports most of the key features of CCNx, and is interoperable with CCNx. The core CCNx logic has been implemented in about 1000 lines of code and is ideal for classroom work and course projects as well as for quickly experimenting with CCNx extensions.

ndnSIM [ndnSIM] is a module that can be plugged into the ns-3 simulator and supports the core features of CCN. One can use ndnSIM to experiment with various CCN applications and services as well as components developed for CCN such as routing protocols, caching and forwarding strategies. The code for ns-3 and ndnSIM is openly available to the community and can be used as the basis for implementing ICN protocols or applications. For more details see <http://www.nsnam.org> and <http://www.ndnsim.net>.

ccnSim [ccnSim] is another CCN-specific simulator that was specially designed to handle forwarding of a large number of CCN-chunks. ccnSim is written in C++ for the OMNeT++ simulation framework (www.omnetpp.org). Interested readers could consider also the Content Centric Networking Packet Level Simulator [CCNPL]. Finally, CCN-Joker [CCNj] is an application-layer platform that can be used to build a CCN overlay. CCN-Joker emulates in user-space all basic aspects of a CCN node (e.g., handling of Interest and Data packets, cache sizing, replacement policies), including both flow and congestion control. The code is open source and is suitable for both emulation-based analyses and real experiments.

An example of a testbed that supports CCN is the Open Network Lab (see <https://onl.wustl.edu/>). The ONL testbed currently comprises 18 extensible gigabit routers and over a 100 computers representing clients and is freely available to the public for running CCN experiments. Nodes in ONL are preloaded with CCNx software. ONL provides a graphical user interface for easy configuration and testbed set up as per the experiment requirements, and also serves as a control mechanism, allowing access to various control variables and traffic counters. It is also possible to run and evaluate CCN over popular testbeds such as PlanetLab (www.planet-lab.org), Emulab (www.emulab.net), and Deter (www.isi.deterlab.net) by directly running the CCNx open-source code on PlanetLab and Deter nodes, respectively.

NEPI, the Network Experimentation Programming Interface, (<http://nepi.inria.fr>) is a tool developed for controlling and managing large-scale network experiments. NEPI provides an experiment

description language to design network experiments, describing topology, applications, and a controller to automatically deploy those experiments on target experimentation environments, such as PlanetLab. The controller is also capable of collecting result and log files during the experiment execution. NEPI also allows to specify node selection filters while designing the experiment, thereby supporting automatic discovery and provisioning of testbed nodes during experiment deployment, without the user having to hand-pick them. It is simple and efficient to use NEPI to evaluate CCNx on large-scale testbeds such as PlanetLab.

2.1.1.2. Publish/Subscribe Internet Architecture

The PSIRP project has open-sourced its Blackhawk publish-subscribe (Pub/Sub) implementation for FreeBSD; more details are available online at <http://www.psirp.org/downloads.html>. Despite being limited to one operating system, the code base also provides a virtual image to allow its deployment on other environments through virtualization.

The code distribution features a kernel module, a file system and scope daemon, as well as a set of tools, test applications and scripts. This work was extended as part of the PURSUIT project, resulting in the development of the Blackadder prototype for Linux and FreeBSD. It currently runs on a testbed across Europe, America (MIT) and Japan (NICT). All sites are connected via OpenVPN, which exports a virtual Ethernet device to all machines in the testbed. In total, 40 machines in a graph topology containing one Topology Manager and one Rendezvous node that handle all publish/subscribe and topology formation requests are interconnected [IEICE].

Moreover, the ICN simulation environment [ICN-Sim] allows the simulation of new techniques for topology management following the Publish-Subscribe paradigm and the PSIRP approach. The simulator is based on the OMNET++ simulator and the INET/MANET frameworks. It is currently publicly available at <http://sourceforge.net/projects/icnsim>. A design characteristic of this platform is the separation between the network and topology management policies. An interface is used to provide this functionality and policies can be imported and applied in the network as topology manager applications running on top of this interface.

2.1.1.3. NetInf

The EU FP7 4WARD and SAIL projects have made a set of open-source implementations available; see <http://www.netinf.org/open-source> for more details. Of note, two software packages are available. The first one is a set of tools for NetInf implementing different aspects

of the protocol (e.g., NetInf URI format, HTTP and UDP convergence layer) using different programming languages. The Java implementation provides a local caching proxy and client. The second one, is a OpenNetInf prototype from the 4WARD project. Besides a rich set of NetInf mechanisms implemented, it also provides a browser plug-in and video streaming software. The SAIL project developed a hybrid host-centric and information-centric network architecture called the Global Information Network (GIN). The prototype for this can be downloaded from <http://gin.ngnet.it>.

2.1.4. COMET

The EU FP7 COMET project developed a simulator, called Icarus, which implements ProbCache [PROBCACHE], centrality-based in-network caching [CL4M] and the hash-route-based algorithms detailed in [HASHROUTE]. The simulator is built in Python and makes use of the Fast Network Simulator Setup tool [FNSS] to configure the related parameters of the simulation. The simulator is available from:
<https://github.com/lorenzosaino/icarus/>

2.1.5. Large-scale Testing

An important consideration in the evaluation of any kind of future Internet mechanism, lies in the characteristics of that evaluation itself. Often, central to the assessment of the features provided by a novel mechanism, lies the consideration of how it improves over already existing technologies, and by "how much." With the disruptive nature of clean-slate approaches generating new and different technological requirements, it is complex to provide meaningful results for a network layer framework, in comparison with what is deployed in the current Internet. Thus, despite the availability of ICN implementations and simulators, the need for large-scale environments supporting experimental evaluation of novel research is of prime importance to the advancement of ICN deployment.

In this regard, initiatives such as the Future Internet Research and Experimentation Initiative (www.ict-fire.eu), enable researchers to test new protocols and architectures in real conditions over production networks (e.g., through virtualization and software-defined networking mechanisms), simplifying the validation of future evolutions and reducing the gap between research and deployment. Similarly, Future Internet Design (www.nets-find.net) is a long-term initiative along the same direction in the US. GENI (www.geni.net) also offers experimentation infrastructure as does PlanetLab (www.planet-lab.org), which likely offers the largest testbed available today. Those wishing to perform smaller, more controlled

experiments can also consider the Emulab testbed (www.emulab.net), which allows various topologies to be configured.

The Asia Future Internet Forum (www.asiafi.net) has also designed a testbed mainly used for ICN experiments. This testbed consists of multiple servers located in Asia and will be (presumably) expanded with servers in other locations. Each testbed server includes multiple Linux kernel-based LXCs. One container is called "bridge container" and the other container is called "user container". A bridge container has a global IP address used to connect to the physical network through bridge mode, and a local (private) IP address. A user container, which is assigned to each researcher, connects to the bridge container in the same server using their local IP addresses. A user container connects to the researcher's remote containers located in different servers via tunnels established between its local bridge container and the remote bridge containers.

Finally, the National Institute of Information and Communications Technology (NICT) builds and operates the high-performance testbed JGN-X (see <http://www.jgn.nict.go.jp/english/index.html>), which has cutting-edge network functions and technologies including those currently in development. JGN-X aims to establish new-generation network technology and accelerate the R&D in areas such as network virtualization and advanced operations of virtualized layers. JGN-X is used for collaboration among developers in order to foster the establishment and expansion of new-generation network technology.

2.2. Topology Selection

[draft-irtf-icnrg-scenarios] introduced several topologies that have been used in ICN studies so far but, to date and to the best of our understanding, there is no single topology that can be used to easily evaluate all aspects of the ICN paradigm. There is rough consensus that the classic dumbbell topology cannot serve well future evaluations of ICN approaches. Therefore, one should consider a range of topologies, each of which would stress different aspects, as outlined earlier in this document. Current Internet traces are also available to assist in this, e.g. see <http://www.caida.org/data/active/internet-topology-data-kit> and <http://www.cs.washington.edu/research/networking/rocketfuel>.

Depending on what is the focus of the evaluation, intra-domain topologies alone may be appropriate. However, those interested, for example, in quantifying transit costs will require inter-domain traces (note that the above CAIDA traces offer this). Scalability is an important consideration in this choice of this with CAIDA's ITDK traces recording millions of routers across thousands of domains.

Beyond these traces there is a wide range of synthetic topologies, such as the Barabasi-Albert model [BA] and the Watts-Strogatz small-world topology [WATTS]. These synthetic traces allow experiments to be performed whilst controlling various key parameters (e.g. degree).

Through this, different aspects can be investigated, such as inspecting resilience properties. For some research, this may be more appropriate as, practically speaking, there are no assurances that a future ICN will share the same topology with today's networks.

Besides defining the evaluation topology as a graph $G = (V, E)$, where V is the set of vertices (nodes) and E is the set of edges (links), one should also clearly define and list the respective matrices that correspond to the network, storage and computation capacities available at each node as well as the delay characteristics of each link, so that the results obtained can be easily replicated in other studies. Recent work by Hussain and Chen [Montage], although currently addressing host-centric networks, could also be leveraged and be extended by the ICN community. Measurement information can also be taken from existing platforms such as iPlane (<http://iplane.cs.washington.edu>), which can be used to provide configuration parameters such as access link capacity and delay. Alternatively, synthetic models such as [DELAY] can be used to configure such topologies.

Finally, the dynamic aspects of a topology, such as node and content mobility, disruption patterns, packet loss rates as well as link and node failure rates, to name a few, should also be carefully considered. As mentioned in [draft-irtf-icnrg-scenarios], for example, contact traces from the DTN community could also be used in ICN evaluations.

2.3. Traffic Load

As we are still lacking ICN-specific traffic workloads we can currently only extrapolate from today's workloads. In this subsection we provide a first draft of a set of common guidelines, in the form of what we will refer to as a content catalog for different scenarios. This catalog, which is based on previously published work, could be used to evaluate different ICN proposals, for example, on routing, congestion control, and performance, and can be considered as other kinds of ICN contributions emerge.

We take scenarios from today's Web, file sharing (BitTorrent-like) and User Generated Content (UGC) platforms (e.g., YouTube), as well as Video on Demand (VoD) services. Publicly available traces for these include those available from web sites such as http://mikel.tlm.unavarra.es/~mikel/bt_pam2004,

<http://multiprobe.ewi.tudelft.nl/multiprobe.html>,
<http://an.kaist.ac.kr/traces/IMC2007.html>, and
<http://traces.cs.umass.edu/index.php/Network/Network>.

The content catalog for each type of traffic can be characterized by a specific set of parameters: the cardinality of the estimated content catalog, the average size of the exchanged contents (either chunks or entire named information objects), and the statistical distribution that best reflect the popularity of objects and their request frequency. Table I summarizes the content catalog. With this shared point of reference, the use of the same set of parameters (depending on the scenario of interest) among researchers will be eased, and different proposals could be compared on a common base.

Table I. Content Catalog

Traffic Load	Catalog Size [L1][L2] [L3][L5]	Mean Object Size [L4][L5][L7][L8] [L9][L10]	Popularity Distribution [L3][L5][L6][L11][L12]
=====			
Web	10 ¹²	Chunk: 1-10 kB	Zipf with 0.64 <= alpha <= 0.83

File sharing	5x10 ⁶	Chunk: 250-4096 kB Object: ~800 MB	Zipf with 0.75 <= alpha <= 0.82

UGC	10 ⁸	Object: ~10 MB	Zipf, alpha >= 2

VoD	10 ⁴	Object: ~100 MB	Zipf, 0.65 <= alpha <= 1
=====			

* UGC = User Generated Content ** VoD = Video on Demand

Several studies in the past years have stated that Zipf's law is the discrete distribution that best represents the request frequency in a number of application scenarios, ranging from the Web to VoD services. The key aspect of this distribution is that the frequency of a content request is inversely proportional to the rank of the content itself, i.e., the smaller the rank, the higher the request frequency. If we denote with M the content catalog cardinality and with $1 \leq i \leq M$ the rank of the i-th most popular content, we can express the probability of requesting the content with rank "i" as:

$$P(X=i) = (1/i^{\alpha}) / C, \text{ with } C = \sum (1 / j^{\alpha}), \alpha > 0$$

where the sum is obtained considering all values of j, $1 \leq j \leq M$.

Further, a variation of the Zipf distribution, termed the Mandelbrot-Zipf distribution, has been suggested by [P2PMod] to better model environments where nodes can locally store previously requested content. For example, it was observed that peer-to-peer file sharing applications typically exhibited a 'fetch-at-most-once' style of behavior. This is because peers tend to persistently store the files they download, a behavior that may also be prevalent in ICN.

2.4. Choosing Relevant Metrics

ICN is a networking concept that spun out of the desire to align the operation model of a network with the model of its typical use. For TCP/IP networks, this means to change the mechanisms of data access and transport from a host-to-host model to a user-to-information model. The premise is that the effort invested in changing models will be offset, or even surpassed, by the potential of a "better" network. However, such a claim can be validated only if it is quantified.

Quantification of network performance requires a set of standard metrics. These metrics should be broad enough so they can be applied equally to host-centric and information-centric (or other) networks. This will allow reasoning about a certain ICN approach in relation to an earlier version of the same approach, to another ICN approach or to the incumbent host-centric approach. It will therefore be less difficult to gauge optimization and research direction. On the other hand, the metrics should be targeted to network performance only and should avoid unnecessary expansion into the physical and application layers. Similarly, at this point, it is more important to capture as metrics only the main figures of merit and to leave more esoteric and less frequent cases for the future.

To arrive at a set of relevant metrics, it would be beneficial to look at the metrics used in existing ICN approaches, such as CCN [CCN] [VoCCN] [NDNP], NetInf [4WARD6.1] [4WARD6.3] [SAIL-B2] [SAIL-B3], PURSUIT [PRST4.5], COMET [CMT-D5.2] [CMT-D6.2], Connect [SHARE] [RealCCN], and CONVERGENCE [ICN-Web] [ICN-Scal] [ICN-Tran]. The metrics used in these approaches fall into two categories: metrics for the approach as a whole, and metrics for individual components (resolution, routing, etc.). Metrics for the entire approach are further subdivided into traffic and system metrics. It is important to note that the various approaches do not name or define metrics consistently. This is a major problem when trying to find metrics that allow comparison between approaches. For the purposes of exposition, in what follows we have tried to smooth differences by pitting similarly defined metrics under the same name. Also, due to space constraints, we have chosen to report here only the most common

metrics between approaches. For more details the reader should consult the references for each approach.

Traffic metrics in existing ICN approaches are summarized in Table II. These are metrics for evaluating an approach mainly from the perspective of the end user, i.e., the consumer, provider, or owner of the content or service. Depending on the level where these metrics are measured, we have made the distinction into user, application and network-level traffic metrics. So for example, network-level metrics are mostly focused on packet characteristics, whereas user-level metrics can cover elements of human perception. The approaches don't make this distinction explicitly, but we can see from the table that CCN and NetInf have used metrics from all levels, PURSUIT and COMET have focused on lower-level metrics, and Connect and CONVERGENCE prefer higher-level metrics. Throughput and download time seem to be the most popular metrics altogether.

Table II. Traffic metrics used in ICN evaluations

	User	Application		Network	
	Download time	Goodput	Startup latency	Throughput	Packet delay
CCN	x	x		x	x
NetInf	x		x	x	x
PURSUIT			x	x	x
COMET			x	x	
Connect	x				
CONVERGENCE	x	x			

While traffic metrics are more important for the end user, the owner or operator of the networking infrastructure is normally more interested in system metrics, which can reveal the efficiency of an approach. The various ICN approaches have used system metrics, but unfortunately the situation is not as coherent as with the traffic metrics. The most common system metrics used are: protocol overhead, total traffic, transit traffic, cost savings, router cost, and router energy consumption.

Besides the traffic and systems metrics that aim to evaluate an

approach as a whole, all of the surveyed approaches also evaluate the performance of individual components. The name resolution, request/data routing, and data caching are the most typical components, so Table III presents the popular metrics for each of those components. FIB size and path length, i.e., the routing component metrics, are almost ubiquitous among approaches, perhaps due to the networking background of the involved researchers. That might be also the reason for the sometimes decreased focus on traffic and system metrics, in favor of component metrics. It can certainly be argued that traffic and system metrics are affected by component metrics, however no approach has made the relationship clear. With this in mind, and also taking into account that traffic and system metrics are readily useful to end users and network operators, we will restrict ourselves to those in the following sections.

Table III. Component metrics in existing ICN approaches

	Resolution		Routing		Cache	
	Resolution time	Request rate	FIB size	Path length	Size	Hit ratio
CCN	x		x	x	x	x
NetInf	x	x		x		x
PURSUIT			x	x		
COMET	x	x	x	x		x
CONVERGENCE		x	x		x	

Before proceeding, we should note that we'd like our metrics to be applicable to host-centric networks as well. Standard metrics already exist for IP networks and it would certainly be beneficial to take them into account. It is encouraging that many of the metrics used by existing ICN approaches can also be used on IP networks and that all of the approaches have tried on occasion to draw the parallels.

2.4.1. Traffic Metrics

At their core, host-centric and information-centric networking function as data transport services. Information of interest to a user resides in one or more storage points connected to the network

and, on the user's request, the network transports this information to the user for consumption. We could therefore do worse than to quantify the data transport performance of the network in terms of Quality of Service (QoS) metrics.

The IETF has been working for more than a decade on devising metrics and methods for measuring the performance of IP networks. The work has been carried out largely within the IPPM WG, guided by a relevant framework [RFC2330]. IPPM metrics include delay, delay variation, loss, reordering, and duplication. While the IPPM work is certainly based on packet-switched IP networks, it is conceivable that it can be modified and extended to cover ICN networks as well. However, more study is necessary to turn this claim into a certainty. Many experts have toiled for a long time on devising and refining the IPPM metrics and methods, so it would be an advantage to use IPPM on measuring ICN performance. In addition, IPPM works already for host-centric networks, so comparison with information-centric networks would entail only the ICN extension of the IPPM framework. Finally, an important benefit of measuring the transport performance of a network at its output, using QoS metrics such as IPPM, is that it can be done mostly without any dependence to applications.

Another option for measuring transport performance would be to use Quality of Service metrics, not at the output of the network like with IPPM, but at the input to the application. So for an application like live video streaming the relevant metrics would be startup latency, playout lag and playout continuity. The benefit of this approach is that it abstracts away all details of the underlying transport network, so it can be readily applied to compare between networks of different concepts (host-centric, information-centric, or other). As implied earlier, the drawback of the approach is its dependence on the application, so it is likely that different (types of) applications will require different metrics. It might be possible to identify standard metrics for each type of application, but the situation is not as clear as with IPPM metrics and further investigation is necessary.

At a higher level of abstraction, we could measure the network's transport performance at the application output. This entails measuring the quality of the transported and reconstructed information as perceived by the user during consumption. In such an instance we would use Quality of Experience (QoE) metrics, which are by definition dependent on the application. For example, the standardized methods for obtaining a Mean Opinion Score (MOS) for VoIP (e.g., ITU-T P.800) is quite different from those for IPTV (e.g., PEVQ). These methods are notoriously hard to implement, as they involve real users in a controlled environment. Such constraints can be relaxed or dropped by using methods that model

human perception under certain environments, but these methods are typically intrusive. The most important drawback of measuring network performance at the output of the application is that only one part of each measurement is related to network performance. The rest is related to application performance, e.g., video coding, or even device capabilities, both of which are irrelevant to our purposes here and are generally hard to separate. We therefore see the use of QoE metrics in measuring ICN performance as a poor choice.

2.4.2. System Metrics

Overall system metrics that need to be considered include reliability, scalability, energy efficiency, and delay/disconnection tolerance. In deployments where ICN is addressing specific scenarios, relevant system metrics could be derived from current experience. For example, in IoT scenarios, which were discussed earlier in [draft-irtf-icnrg-scenarios], it is reasonable to consider the current generation of sensor nodes, sources of information, and even measurement gateways (e.g., for smart metering at homes) or smartphones. In this case, ICN operation ought to be evaluated with respect not only to overall scalability and network efficiency, but also the impact on the nodes themselves. Karnouskos et al. [SensReqs] provide a comprehensive set of sensor and IoT-related requirements, for example, which include aspects such as resource utilization, service life-cycle management and device management.

Additionally, various specific metrics are also critical in constrained environments, such as CPU processing requirements, signaling overhead, and memory allocation for caching procedures in addition to power consumption and battery lifetime. Also, in nodes acting as gateways, which typically not only act as a point of service to a large number of nodes, but also have to satisfy the information requests from remote entities; they need to consider scalability-related metrics, such as frequency and processing of successfully satisfied information requests.

Finally, given the in-network caching functionality of Information-Centric Networks, metrics for the efficiency and performance of in-network caching have to be defined. Such metrics will need to guide researchers and operators regarding the performance of in-network caching algorithms. A first step on this direction has been made in [L9]. The paper proposes a formula that approximates the proportion of time that a content stays in a network cache. The model takes as input the rate of requests for a given content (the Content of Interest) and the rate of requests for all other contents that go through the given network element (router) and move the CoI down in the (LRU) cache. The formula takes also into account the size of the

cache of this router.

The output of the model essentially reflects the probability that the CoI will be found in a given cache. The initial study [L9] is applied to the CCN/NDN framework, where contents get cached at every node they traverse, while efforts are underway to assess the accuracy of the model for other caching strategies. The formula according to which the probability or proportion is calculated is given by:

$$pi = [\mu/(\mu+\lambda)]^N,$$

where λ is the request rate for CoI, μ is the request rate for contents that move CoI down the cache and N is the size of the cache (in slots).

The formula can be used to assess the caching performance of the system and can also potentially be used to identify the gain of the system due to caching. This can then be used to compare against gains by other factors, e.g., addition of extra bandwidth in the network.

2.5. Resource Equivalence and Tradeoffs

As we have seen above, every ICN network is built from a set of resources, which include link capacities, different types of memory structures and repositories used for storing named information objects and chunks temporarily (i.e. caching) or persistently, as well as name resolution and other lookup services. Complexity and processing needs in terms of forwarding decisions, management (e.g. need for manual configuration, explicit garbage collection, and so on), and routing (i.e. amount of state needed, need for manual configuration of routing tables, support for mobility, etc.) set the stage for a range of engineering tradeoffs.

In order to be able to compare different ICN approaches it would be beneficial to be able to define equivalence in terms of different resources which today are considered incomparable. For example, would provisioning an additional 5 Mb/s link capacity lead to better performance than adding 100 GB of in-network storage? Within this context one would consider resource equivalence (and the associated tradeoffs) for example for cache hit ratios per GB of cache, forwarding decision times, CPU cycles per forwarding decision, and so on.

3. ICN Security Aspects

The introduction of an information-centric networking architecture

and the corresponding communication paradigm changes many aspects of network security. These will affect all the scenarios described in [draft-irtf-icnrg-scenarios]. Additional evaluation will be required to ensure relevant security requirements are appropriately met by the implementation of the chosen architecture in the various scenarios.

The various ICN architectures that are currently proposed have concentrated on authentication of delivered content to ensure the integrity of the content. However the approaches are primarily applicable to freely accessible content that does not require access authorization, although they will generally support delivery of encrypted content.

The introduction of widespread caching mechanisms may also provide additional attack surfaces. The caching architecture to be used also needs to be evaluated to ensure that it meets the requirements of the usage scenarios.

In practice, the work on security in the various ICN research projects has been heavily concentrated on authentication of content. Work on authorization, access control, privacy and security threats due to the expanded role of in-network caches has been quite limited. A roadmap for improving the security model in NetInf can be found in [NETINFSC]. In the rest of this section we briefly consider the issues and provide pointers to the work that has been done on the security aspects of the architectures proposed.

3.1. Authentication

For fully secure content distribution, content access requires that the receiver needs to be able to reliably assess:

- validity: is it a complete, uncorrupted copy of what was originally published;
- provenance: can the receiver identify the publisher, and, if so, whether it and the source of any cached version of the document can be adequately trusted; and
- relevance: is the content an answer to the question that the receiver asked.

All the ICN architectures considered in this document primarily target the validity requirement using strong cryptographic means to tie the content request name to the content. Provenance and relevance are directly targeted to varying extents: There is a tussle or trade-off between simplicity and efficiency of access and

level of assurance of all these traits. For example, maintaining provenance information can become extremely costly, particularly when considering (historic) relationships between multiple objects. Architectural decisions have therefore been taken in each case as to whether the assessment is carried out by the ICN or left to the application.

An additional consideration for authentication is whether a name should be irrevocably and immutably tied to a static piece of preexisting content or whether the name can be used to refer to dynamically or subsequently generated content. Schemes that only target immutable content can be less resource hungry as they can use digest functions rather than public key cryptography for generating and checking signatures. However, this can increase the load on applications. This is because they are required to manage many names, rather than using a single name for an item of evolving content that changes over time (e.g. a piece of data containing an age reference).

NetInf uses the Named Information (ni) URI scheme [RFC6920] to identify content. This allows NetInf to assure validity without any additional information but gives no assurance on provenance or relevance. A "search" request allows an application to identify relevant content and applications may choose to structure content to allow provenance assurance but this will typically require additional network access. NetInf validity authentication is consequently efficient in a network environment with intermittent connectivity as it does not force additional network accesses and allows the application to decide on provenance validation if required. NetInf primarily targets static content, but an extension would allow dynamic content to be handled. The immutable case only uses digest functions.

DONA [DONA] and CCN [CCN], [SECCONT] integrate most of the data needed to verify provenance into all content retrievals but need to be able to retrieve additional information (typically a security certificate) in order to complete the provenance authentication. Whether the application has any control of this extra retrieval will depend on the implementation. CCN is explicitly designed to handle dynamic content allowing names to be pre-allocated and attached to subsequently generated content. DONA offers variants for dynamic and immutable content.

PURSUIT [PSTSEC] appears to allow implementers to choose the authentication mechanism so that it can, in theory, emulate the authentication strategy of any of the other architectures. It is not clear whether different choices would lead to lack of interoperability.

3.2. Authorization, Access Control and Statistics

A potentially major concern for all the ICN architectures considered here is that they do not provide any inbuilt support for an authorization framework or for statistics monitoring. Once content has been published and cached in servers, routers or end points not controlled by the publisher, the publisher has no way to enforce access control, determine which users have accessed the content or revoke its publication. In fact, in some cases, it is even difficult for the publishers themselves to perform access control, where requests do not necessarily contain host/user identifier information.

Access could be limited by encrypting the content but the necessity of distributing keys out-of-band appears to negate the advantages of in-network caching. This also creates significant challenges when attempting to manage and restrict key access. An authorization delegation scheme has been proposed [ACDICN] but this requires access to a server controlled by the publisher to obtain an access token making it essentially just an out-of-band key distribution system.

Evaluating the impact of the absence of these features will be essential for any scenario where an ICN architecture might be deployed. It may have a seriously negative impact on the applicability of ICN in commercial environments unless a solution can be found.

3.3. Privacy

Another area where the architectures have not been significantly analyzed is privacy. Caching implies a trade-off between network efficiency and privacy. The activity of users is significantly more exposed to the scrutiny of cache owners with whom they may not have any relationship.

Although in many ICN architectures, the source of a request is not explicitly identified, an attacker may be able to obtain considerable information if s/he can monitor transactions on the cache and obtain details of the objects accessed, the topological direction of requests and information about the timing of transactions. The persistence of data in the cache can make life easier for an attacker by giving a longer timescale for analysis.

The impact of CCN on privacy has been investigated in a useful master's thesis [CCNSEC]. The analysis in this thesis is mostly applicable to all of the ICN architectures because it is mostly focused on the common caching aspect. The privacy risks of named data networking are also highlighted in [CCNPRIV]. Further work on

privacy in ICNs can be found in [CONPRV].

3.4. Changes to the Network Security Threat Model

The architectural differences of the various ICN models as compared to TCP/IP have consequences for network security. There is limited consideration of the threat models and potential mitigation in the various documents describing the architectures. The references [CCNSEC] and [CONPRV] also consider the changed threat model. Some of the key aspects are:

- o Caching implies a tradeoff between network efficiency and user privacy as discussed in Section 3.3.
- o More powerful routers upgraded to handle persistent caching increase the network's attack surface. This is particularly the case in systems (e.g., CCN) that may need to perform cryptographic checks on content that is being cached. For example, not doing this could lead routers to disseminate invalid content.
- o ICNs makes it difficult to identify the origin of a request as mentioned in Section 4.3 slowing down the process of blocking requests and requiring alternative mechanisms to differentiate legitimate requests from inappropriate ones as access control lists (ACLs) will probably be of little value for ICN requests.
- o Denial-of-service (DoS) attacks may require more effort on ICN than on TCP/IP but they are still feasible. One reason for this is that it is difficult for the attacker to force repeated requests for the same content onto a single node; ICNs naturally spread content so that after the initial few requests, subsequent requests will generally be satisfied by alternative sources, blunting the impact of a DoS attack. That said, there are many ways around this, e.g., generating random suffix identifiers that always result in cache misses.
- o Per-request state in routers can be abused for DoS attacks.
- o Caches can be misused in the following ways:
 - + Attackers can use caches as storage to make their own content available.
 - + The efficiency of caches can be decreased by attackers with the goal of DoS attacks.
 - + Content can be extracted by any attacker connected to the cache, putting users' privacy at risk.

Appropriate mitigation of these threats will need to be considered in each scenario.

4. Security Considerations

This document does not impact the security of the Internet.

5. IANA Considerations

This document presents no IANA considerations.

6. Acknowledgments

Daniel Corujo and Gareth Tyson contributed to an earlier version of this document.

This document has benefited from reviews, pointers to the growing ICN literature, suggestions, comments and proposed text provided by the following members of the IRTF Information-Centric Networking Research Group (ICNRG), listed in alphabetical order: Marica Amadeo, Hitoshi Asaeda, Claudia Campolo, Suyong Eum, Dorothy Gellert, Luigi Alfredo Grieco, Myeong-Wuk Jang, Ren Jing, Will Liu, Antonella Molinaro, Ioannis Psaras, Dirk Trossen, Jianping Wang, Yuanzhe Xuan, and Xinwen Zhang.

7. Informative References

- [RFC5743] Falk, A., "Definition of an Internet Research Task Force (IRTF) Document Stream", RFC 5743, December 2009.
- [RFC2330] Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics", RFC 2330, May 1998.
- [RFC6920] Farrell, S., Kutscher, D., Dannewitz, C., Ohlman, B., Keranen, A., and P. Hallam-Baker, "Naming Things with Hashes", RFC 6920, April 2013.
- [ndnSIM] Afanasyev, A. et al., ndnSIM: NDN simulator for NS-3 NDN Technical Report NDN-0005, Revision 2, October 2012.
- [ccnSim] Rossini, G. and D. Rossi, "Large scale simulation of CCN networks", Proc. Algotel 2012, La Grande Motte, France, May 2012.

- [CCNPL] Muscariello, L., "Content centric networking packet level simulator", available online at <http://perso.rd.francetelecom.fr/muscariello/sim.html>
- [CCNj] Cianci, I. et al. "CCN - Java Opensource Kit EmulatoR for Wireless Ad Hoc Networks", Proc. 7th ACM Int. Conf. on Future Internet Technologies, Seoul, Korea, Sept., 2012.
- [IEICE] G. Parisi, D. Trossen, and H. Asaeda, "A Node Design and a Framework for Development and Experimentation for an Information-Centric Network", IEICE Trans. Commun., vol. E96-B, no. 7, pp.1650-1660, July 2013.
- [ICN-Sim] N. Vastardis et al., "Simulation Tools Enabling Research on Information-centric Networks", Proc. ICC FutureNet Workshop. IEEE, 2012.
- [PROBCACHE] I. Psaras, W. Chai, G. Pavlou, "Probabilistic In-Network Caching for Information-Centric Networks", Proc. SIGCOMM ICN Workshop. ACM, 2012.
- [CL4M] Chai, W. K. et al., "Cache 'Less for More' in Information-centric Networks", Proc. Networking. IFIP, 2012.
- [HASHROUTE] L. Saino, I. Psaras, G. Pavlou, "Hash-routing Schemes for Information-Centric Networking", Proc. SIGCOMM ICN Workshop. ACM, 2013.
- [FNSS] L. Saino, C. Cocora and G. Pavlou, "A Toolchain for Simplifying Network Simulation Setup", Proc. SIMUTOOLS. ACM, 2013.
- [BA] Barabasi, A. and R. Albert, "Emergence of scaling in random networks", Science, vol. 286, no. 5439, pp. 509-512, 1999.
- [WATTS] Watts, D. J. and S. H. Strogatz, "Collective dynamics of small-world networks", Nature, vol. 393, no. 6684, pp. 40-41, 1998.
- [Montage] Hussain, A. and J. Chen, "Montage Topology Manager: Tools for Constructing and Sharing Representative Internet Topologies", DETER Technical Report, ISI-TR-684, Aug. 2012.
- [DELAY] Kaune, S. et al., "Modelling the Internet Delay Space Based on Geographical Locations", Proc. Euromicro, Weimar, Germany, 2009.

- [L1] <http://googleblog.blogspot.it/2008/07/we-knew-web-was-big.html>
- [L2] Zhang, C., Dhungel, P., and K. Di Wu., "Unraveling the BitTorrent ecosystem", IEEE Transactions on Parallel and Distributed Systems, pp. 1164-1177, 2010.
- [L3] Cha, M., Kwak, H., Rodriguez, P., Ahn, Y.-Y., and S. Moon, "I tube, you tube, everybody tubes: analyzing the world's largest user generated content video system", Proc. ACM SIGCOMM conference on Internet measurement (IMC), San Diego (CA), USA, Oct. 2007.
- [L4] Zhou, J., Li, Y., Adhikari, K., and Z.-L. Zhang, "Counting YouTube videos via random prefix sampling", In Proc. of IMC'11, Berlin, Germany, Nov. 2011.
- [L5] Fricker, C., Robert, P., Roberts, J. and N. Sbihim, "Impact of traffic mix on caching performance in a content-centric network", In Proc. of IEEE NOMEN 2012, Workshop on Emerging Design Choices in Name-Oriented Networking, Orlando, USA, Mar. 2012.
- [L6] Yu, H., Zheng, D., Zhao, B. Y. and W. Zheng, "Understanding user behavior in large-scale video-on-demand systems", In SIGOPS Oper. Syst. Rev., Vol. 40, pp. 333-344, April 2006.
- [L7] Marciniak, P., Liogkas, N., Legout, A. and E. Kohler, "Small is not always beautiful", In Proc. of IPTPS, International Workshop of Peer-to-Peer Systems, Tampa Bay, Florida (FL), USA, Feb. 2008.
- [L8] Bellissimo, A., Levine, B. and P. Shenoy, "Exploring the use of BitTorrent as the basis for a large trace repository", University of Massachusetts, Tech. Rep., 2004.
- [L9] Psaras, I. et al., "Modelling and Evaluation of CCN-Caching Trees", In Proc. of the 10th international IFIP conference on Networking, Valencia, Spain, May 2011.
- [L10] Carofiglio, G., Gallo, M., Muscariello, L., and D. Perino, "Modeling Data Transfer in Content-Centric Networking", In Proc. of ITC, San Francisco, USA, Sep. 2011.
- [L11] Breslau, L., Cao, P., Fan, L., Phillips, G. and S. Shenker, "Web caching and zipf-like distributions:

- evidence and implications", In Proc. of INFOCOM '99, New York (NY), USA, Mar. 1999.
- [L12] Mahanti, A., Williamson, C., and D. Eager., "Traffic analysis of a web proxy caching hierarchy", IEEE Network, Vol.14, No.3, pp.16-23, May/June 2000.
- [P2PMod] Saleh, O., and M. Hefeeda, "Modeling and caching of peer-to-peer traffic", Proc. ICNP. IEEE, 2006.
- [CCN] Jacobson, V. et al., "Networking Named Content", Proc. CoNEXT. ACM, 2009.
- [VoCCN] Jacobson, V. et al., "VoCCN: Voice-over Content-Centric Networks", Proc. CoNEXT Re-Arch Workshop. ACM, 2009.
- [NDNP] Zhang, L. et al., "Named Data Networking (NDN) Project", NDN Technical Report NDN-0001, Oct. 2010. Available: <http://named-data.net/publications/techreports/>
- [4WARD6.1] Ohlman, B. et al., "First NetInf Architecture Description", 4WARD Project Deliverable D-6.1, Apr. 2009.
- [4WARD6.3] Ahlgren, B. et al., "NetInf Evaluation", 4WARD Project Deliverable D-6.3, June 2010.
- [SAIL-B2] SAIL, "NetInf Content Delivery and Operations", SAIL Project Deliverable D-B.2, May. 2012.
- [SAIL-B3] Kutscher, D. (ed.) et al., "Final NetInf Architecture", SAIL Project Deliverable D-B.3, Jan. 2013. Available: <http://www.sail-project.eu/deliverables/>
- [PRST4.5] Riihijarvi, J. et al., "Final Architecture Validation and Performance Evaluation Report", PURSUIT Project Deliverable D4.5, Jan. 2013.
- [CMT-D5.2] Ben, A. et al., "Scalability of COMET System", COMET Project Deliverable D5.2, Feb. 2013.
- [CMT-D6.2] Georgiades, M. et al., "Prototype Experimentation and Demonstration", COMET Project Deliverable D6.2, Feb. 2013.
- [SHARE] Muscariello, L. et al., "Bandwidth and storage sharing performance in information centric networking", Proc. SIGCOMM ICN Workshop. ACM, 2011.
- [RealCCN] Perino, D. et al., "A Reality Check for Content Centric

Networking", Proc. SIGCOMM ICN Workshop. ACM, 2011.

- [ICN-Web] Detti, A. et al., "Supporting the Web with an Information Centric Network that Routes by Name", Elsevier Computer Networks, vol. 56, no. 17, Nov. 2012.
- [ICN-Scal] Blefari Melazzi, N. et al., "Scalability Measurements in an Information-Centric Network", Springer Lecture Notes in Computer Science (LNCS), vol. 7586, 2012.
- [ICN-Tran] Salsano, S. et al., "Transport-Layer Issues in Information Centric Networks ", Proc. SIGCOMM ICN Workshop. ACM, 2012.
- [SensReqs] Karnouskos, S. et al., "Requirement considerations for ubiquitous integration of cooperating objects", Proc. NTMS. IFIP, 2011.
- [NETINFSC] Renault, E, Ahmad, A., and M. Abid, "Towards a Security Model for the Future Network of Information", Proc. Conf. Ubiquitous Information Technologies and Applications, IEEE, 2009.
- [DONA] Koponen, T. et al., "A Data-Oriented (and Beyond) Network Architecture", Proc. SIGCOMM. ACM, 2007.
- [SECCONT] Smetters, D., and V. Jacobson, "Securing network content", Technical Report TR-2009-01, PARC, 2009.
- [PSTSEC] Tagger, B., et al, "Update on the Architecture and Report on Security Analysis", Deliverable 2.4, PURSUIT EU FP7 project, April 2012.
- [ACDICN] Fotiou, N. et al., "Access control enforcement delegation for information-centric networking architectures", Proc. SIGCOMM ICN Workshop. ACM, 2012.
- [CCNSEC] Lauinger, T., "Security and Scalability of Content-Centric Networking", Masters Thesis, Technische Universitaet Darmstadt and Eurecom, Sep. 2010.
- [CCNPRIV] Lauinger, Y., et al, "Privacy Risks in Named Data Networking: What is the Cost of Performance?", ACM SIGCOMM Computer Communication Review Editorial Note, vol. 42, iss. 5, 2012
- [CONPRV] Chaabane, A et al, "Privacy in Content-Oriented Networking: Threats and Countermeasures", arXiv:1211.5183, 2012.

Authors' Addresses

Kostas Pentikousis (editor)
EICT GmbH
Torgauer Strasse 12-15
10829 Berlin
Germany

Email: k.pentikousis@eict.de

Borje Ohlman
Ericsson Research
S-16480 Stockholm
Sweden

Email: Borje.Ohlman@ericsson.com

Elwyn Davies
Trinity College Dublin/Folly Consulting Ltd
Dublin, 2
Ireland

Email: davieseb@scss.tcd.ie

Spiros Spirou
Intracom Telecom
19.7 km Markopoulou Avenue
19002 Peania, Athens
Greece

Email: spis@intracom.com

Gennaro Boggia
Dep. of Electrical and Information Engineering
Politecnico di Bari
Via Orabona 4
70125 Bari
Italy

Email: g.boggia@poliba.it

Priya Mahadevan

Palo Alto Research Center
3333 Coyote Hill Rd
Palo Alto, CA 94304
USA

Email: Priya.Mahadevan@parc.com

ICNRG
Internet-Draft
Intended Status: Informational
Expires: February 9, 2015

K. Pentikousis, Ed.
EICT
B. Ohlman
Ericsson
D. Corujo
Universidade de Aveiro
G. Boggia
Politecnico di Bari
G. Tyson
Queen Mary, University of London
E. Davies
Trinity College Dublin
A. Molinaro
UNIRC
S. Eum
NICT
August 8, 2014

Information-centric Networking: Baseline Scenarios
draft-irtf-icnrg-scenarios-03

Abstract

This document aims at establishing a common understanding about a set of scenarios that can be used as a base for the evaluation of different information-centric networking (ICN) approaches so that they can be tested and compared against each other while showcasing their own advantages. Towards this end, we review the ICN literature and document scenarios which have been considered in previous performance evaluation studies. We discuss a variety of aspects that an ICN solution can address. This includes general aspects, such as, network efficiency, reduced complexity, increased scalability and reliability, mobility support, multicast and caching performance, real-time communication efficiency, energy consumption frugality, and disruption and delay tolerance. We detail ICN-specific aspects as well, such as information security and trust, persistence, availability, provenance, and location independence.

This document is a product of the IRTF Information-Centric Networking Research Group (ICNRG).

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

Copyright and License Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Table of Contents

1. Introduction	4
1.1. Baseline Scenario Selection	5
1.2. Document Goals and Outline	5
2. Scenarios	6
2.1. Social Networking	6
2.2. Real-time Communication	7
2.3. Mobile Networking	9
2.4. Infrastructure Sharing	12
2.5. Content Dissemination	13
2.6. Vehicular Networking	14
2.7. Delay- and Disruption-Tolerance	17
2.7.1. Opportunistic Content Sharing	21
2.7.2. Emergency Support and Disaster Recovery	21
2.8. Internet of Things	23
2.9. Smart City	26
3. Cross-scenario Considerations	27
3.1. Multiply-connected Nodes and Economics	27
3.2. Energy Efficiency	32
3.3. Operation across Multiple Network Paradigms	33
4. Summary	34
5. Security Considerations	36
6. IANA Considerations	36
7. Acknowledgments	36
8. Informative References	36
Authors' Addresses	44

1. Introduction

Information-centric networking (ICN) marks a fundamental shift in communications and networking. In contrast with the omnipresent and very successful host-centric paradigm, which is based on perpetual connectivity and the end-to-end principle, ICN changes the focal point of the network architecture from the end host to "named information" (or content, or data). In this paradigm, connectivity may well be intermittent. End-host and in-network storage can be capitalized upon transparently, as bits in the network and on storage devices have exactly the same value. Mobility and multiaccess are the norm and anycast, multicast, and broadcast are natively supported.

It is also worth noting that with the transition from a host-centric to an information-centric communication model the security paradigm changes as well. In a host-centric network, the basic idea is to create secure (remote-access) tunnels to trusted providers of data. In an information-centric network, on the other hand, any source (cache) should be equally usable. This requires some mechanism for making each information item trustworthy by itself, which can be achieved, for example, by name-data-integrity or by signing data objects.

Although interest in ICN is growing rapidly, ongoing work on different architectures, such as, for example, NetInf [NetInf], CCN [CCN] and NDN [NDNP], the publish-subscribe Internet (PSI) architecture [PSI], and the data-oriented architecture [DONA] is far from being completed. One could think of ICN today as being at an equivalent stage of development similar to the one that packet-switched networking was in the late 70's when different technologies, e.g. DECnet, IPX, and IP, just to name a few, were actively developed and put to the test. As such, the development phase that ICN is going through, and the plethora of approaches to tackle the hardest problems, make this a very active and growing research area but, on the downside, it also makes it more difficult to compare different proposals on an equal footing. This document aims to address this partially by establishing a common understanding about potential experimental setups where different ICN approaches can be tested and compared against each other while showcasing their advantages.

The first version of this document appeared in November 2012. It was adopted by ICNRG at IETF 87 (July 2013) as the document to address the work item on the definition of "reference baseline scenarios to enable performance comparisons between different approaches". Earlier versions of this document have been presented during the ICNRG meetings at IETF 85, IETF 86, IETF 87, IETF 88, IETF 89 and at the ICNRG interim meeting in Stockholm in February 2013. This document

has been reviewed, commented, and discussed extensively for a period of nearly two years by the vast majority of ICNRG members, which certainly exceeds 100 individuals. It is the consensus of ICNRG that the baseline scenarios described in this document should be published in the IRTF Stream RFC Series. This document does not constitute a standard.

1.1. Baseline Scenario Selection

Ahlgren et al. [SoA1][SoA2] note that describing ICN architectures is akin to shooting a moving target. We find that comparing these different approaches is often even more tricky. It is not uncommon that researchers devise different performance evaluation scenarios, typically with good reason, in order to highlight the advantages of their approach. This should be expected to some degree at this early stage of ICN development. Nevertheless, this document shows that certain baseline scenarios seem to emerge in which ICN architectures could showcase their comparative advantage over current systems, in general, and against each other, in particular.

This document surveys the peer-reviewed ICN literature and presents prominent evaluation study cases as a foundation for the baseline scenarios to be considered by the IRTF Information-Centric Networking Research Group (ICNRG) in its future work. There are two goals for this document. First, to provide a set of use cases and applications that highlight opportunities for testing different ICN proposals. Second, to identify key attributes of a common set of techniques that can be instrumental in evaluating ICN. Further, these scenarios are intended to equip researchers with sufficient configuration data to effectively evaluate their ICN proposals in a variety of settings, particularly extending beyond scenarios focusing simply on traditional content delivery. The overall aim is that each scenario is described at a sufficient level of detail, and with adequate references to already published work, so that it can serve as the base for comparative evaluations of different approaches. Example code which implements some of the scenarios and topologies included in this document is available from <http://telematics.poliba.it/icn-baseline-scenarios>.

1.2. Document Goals and Outline

This document incorporates input from ICNRG participants and their corresponding text contributions, has been reviewed by several ICNRG active participants (see section 7), and represents the consensus of the research group. However, this document does not constitute an IETF standard, but is indented as an informational document; see also

[RFC5743]. As mentioned above, these scenarios are intended to provide a framework for evaluating different ICN approaches. The methodology for how to do these evaluations as well as definitions of metrics that should be used will be described in a separate document [draft-irtf-icnrg-evaluation-methodology]. In addition, interested readers should consider reviewing [draft-kutscher-icnrg-challenges].

The remainder of this document presents a number of scenarios grouped into several categories in section 2, followed by a number of cross-scenario considerations in section 3. Overall, note that certain evaluation scenarios span across these categories, so the boundaries between them should not be considered rigid and inflexible. Section 4 summarizes in a concise manner the main evaluation aspects across the range of scenarios discussed in this document.

2. Scenarios

This section presents nine scenario categories based on use cases and evaluations which have appeared in the peer-reviewed literature.

2.1. Social Networking

Social networking applications have proliferated over the past decade based on overlay content dissemination systems that require large infrastructure investments to rollout and maintain. Content dissemination is at the heart of the ICN paradigm. Therefore, we would expect that social networking scenarios are a "natural fit" for comparing ICN performance with traditional client-server TCP/IP-based systems. Mathieu et al. [ICN-SN], for instance, illustrate how an Internet Service Provider (ISP) can capitalize on CCN to deploy a short-message service akin to Twitter at a fraction of the complexity of today's systems. Their key observation is that such a service can be seen as a combination of multicast delivery and caching. That is, a single user addresses a large number of recipients, some of which receive the new message immediately as they are online at that instant, while others receive the message whenever they connect to the network.

Along similar lines, Kim et al. [VPC] present an ICN-based social networking platform in which a user shares content with her/his family and friends without the need for centralized content servers; see also section 2.4, below, and [CBIS]. Based on the CCN naming scheme, [VPC] takes a user name to represent a set of devices that belong to the person. Other users in this in-network, serverless social sharing scenario can access the user's content not via a device name/address but with the user's name. In [VPC], signature

Real-time audio and video (A/V) communications include an array of services ranging from one-to-one voice calls to multi-party multi-media conferences with support ranging from whiteboards to augmented reality. Real-time communications have been studied and deployed in the context of packet- and circuit-switched networks for decades. The stringent quality of service requirements that this type of communication imposes on network infrastructure are well known. Since one could argue that network primitives which are excellent for information dissemination are not well-suited for conversational services, ICN evaluation studies should consider real-time communication scenarios in detail.

Notably, Jacobson et al. [VoCCN] presented an early evaluation where the performance of a VoIP (voice over IP) call using an information-centric approach was compared with that of an off-the-shelf VoIP implementation using RTP/UDP. The results indicated that despite the extra cost of adding security support in the ICN approach, performance was virtually identical in the two cases evaluated in their testbed. However, the experimental setup presented is quite rudimentary, while the evaluation considered a single voice call only. Xuan and Yan [NDNpb] revisit the same scenario but are primarily interested in reducing the overhead that may arise in one-to-one communication employing an ICN architecture. Both studies illustrate that quality telephony services are feasible with at least one ICN approach. That said, future ICN evaluations should employ standardized call arrival patterns, for example, following well-established methodologies from the quality of service/experience (QoS/QoE) evaluation toolbox and would need to consider more comprehensive metrics.

Given the wide-spread deployment of real-time A/V communications, an evaluation of an ICN system should demonstrate capabilities beyond feasibility. For example, with respect to multimedia conferencing, Zhu et al. [ACT] describe the design of a distributed audio conference tool based on NDN. Their system includes ICN-based conference discovery, discovery of speakers and voice data distribution. The reported evaluation results point to gains in scalability and security. Moreover, Chen et al. [G-COPSS] explore the feasibility of implementing a Massively Multiplayer Online Role Playing Game (MMORPG) based on CCNx and show that stringent temporal requirements can be met, while scalability is significantly improved when compared to a host-centric (IP-based) client-server system. This type of work points to benefits for both the data and control path of a modern network infrastructure.

Real-time communication also brings up the issue of named data granularity for dynamically generated content. For instance, in many cases A/V data is generated in real-time and is distributed

immediately. One possibility is to apply a single name to the entire content, but this could result in significant distribution delays. Alternatively, distributing A/V content in smaller "chunks" which are named individually may be a better option with respect to real-time distribution but raises naming scalability concerns.

We observe that, all in all, the ICN research community has hitherto only scratched the surface of this area with respect to illustrating the benefits of adopting an information-centric approach as opposed to a host-centric one, and thus more work is recommended in this direction. Scenarios in this category should illustrate not only feasibility but reduced complexity, increased scalability, reliability, and capacity to meet stringent QoS/QoE requirements when compared to established host-centric solutions. Accordingly, the primary aim of this scenario is to exercise each ICN architecture in terms of its ability to satisfy real-time QoS requirements and improved user experience.

2.3. Mobile Networking

IP mobility management relies on anchors to provide ubiquitous connectivity to end-hosts as well as moving networks [MMIN]. This is a natural choice for a host-centric paradigm that requires end-to-end connectivity and a continuous network presence for hosts [SCES]. An implicit assumption in host-centric mobility management is therefore that the mobile node aims to connect to a particular peer, as well as to maintain global reachability and service continuity [EEMN]. However, with ICN new ideas about mobility management should come to the fore capitalizing on the different nature of the paradigm, such as native support for multihoming, abstraction of network addresses from applications, less dependence on connection-oriented sessions, and so on [MOBSURV].

Dannewitz et al. [N-Scen] illustrate a scenario where a multiaccess end-host can retrieve email securely using a combination of cellular and wireless local area network (WLAN) connectivity. This scenario borrows elements from previous work, e.g., [DTI], and develops them further with respect to multiaccess. Unfortunately, Dannewitz et al. [N-Scen] do not present any results demonstrating that an ICN approach is, indeed, better. That said, the scenario is interesting as it considers content specific to a single user (i.e., her mailbox) and does point to reduced complexity. It is also compatible with recent work in the Distributed Mobility Management (DMM) Working Group within the IETF. Finally, Xylomenos et al. [PSIMob] as well as [EEMN] argue that an information-centric architecture can avoid the complexity of having to manage tunnels to maintain end-to-end connectivity as is the case with mobile anchor-based protocols such

as Mobile IP (and its variants). Similar considerations hold for a vehicular (networking) environment, as we discuss in section 2.6.

Overall, mobile networking scenarios have not been developed in detail, let alone evaluated at a large scale. Further, the majority of scenarios discussed so far have related to information consumer, rather than source, mobility. We expect that in the coming period more papers will address this topic. Earlier work [mNetInf] argues that for mobile and multiaccess networking scenarios we need to go beyond the current mobility management mechanisms in order to capitalize on the core ICN features. They present a testbed setup (redrawn in Fig. 2) which can serve as the basis for other ICN evaluations. In this scenario, node "C0" has multiple network interfaces that can access local domains N0 and N1 simultaneously allowing C0 to retrieve objects from whichever server (I2 or I3) can supply them without necessarily needing to access the servers in the core network "C" (P1 and P2). Lindgren [HybICN] explores this scenario further for an urban setting. He uses simulation and reports sizable gains in terms of reduction of object retrieval times and core network capacity use.




Figure 2. Overlapping wireless multiaccess.

The benefits from capitalizing on the broadcast nature of wireless access technologies has yet to be explored to its full potential in the ICN literature, including quantifying possible gains in terms of energy efficiency [E-CHANET]. Obviously, ICN architectures must avoid broadcast storms. Early work in this area considers distributed packet suppression techniques which exploit delayed transmissions and overhearing; examples can be found in [MobiA] and [CCNMANET] for ICN-based mobile ad-hoc networks (MANETs), and in [RTIND] and [CCNVANET] for vehicular scenarios.

One would expect that mobile networking scenarios will be naturally coupled with those discussed in the previous sections, as more users access social networking and multimedia applications through mobile devices. Further, the constraints of real-time A/V applications create interesting challenges in handling mobility, particularly in terms of maintaining service continuity. This scenario therefore spans across most of the others considered in this document with the likely need for some level of integration, particularly considering the well-documented increases in mobile traffic. Mobility is further considered in section 2.7 and the economic consequences of nodes having multiple network interfaces is explored in section 3.1.

Host-centric mobility management has traditionally used a range of metrics for evaluating performance on a per-node and network-wide level. The first metric that comes to mind is handover latency, defined in [RFC5568] as the "period during which the mobile node is unable to send or receive packets". This metric should be considered in ICN performance evaluation studies dealing with mobility. Note that in IP-based networks handover latency has been addressed by the introduction of mobility management protocols, which aim to hide node mobility from the correspondent node, and often follow a make-before-break approach in order to ensure seamless connectivity, and minimize or eliminate altogether handover latency. The "always-on" and "always best connected" [ABC] paradigms have guided mobility management research and standardization for a good decade or so. One can argue that such mechanisms are not particularly suited for ICN. That said, there has been a lot of interest recently in distributed mobility management schemes (see [MMIN] for a summary), where mobility management support is not "always on" by default. Such schemes may be more suitable for ICN. As a general recommendation ICN designs should aim to minimize handover latency so that the end-user and service Quality of Experience (QoE) is not affected adversely.

Network overhead, such as, for instance, the amount of signaling necessary to minimize handover latency, is also a metric that should be considered when studying ICN mobility management. In the past, network overhead has been seen as one of the main factors hindering the deployment of various mobility solutions. In IP-based networks, network overhead includes, but is not limited to, tunneling overhead, in-band control protocol overhead, mobile terminal and network equipment state maintenance and update. ICN designs and evaluation studies should clearly identify the network overhead associated with handling mobility. Alongside network overhead, deployment complexity should also be studied.

To summarize, mobile networking scenarios should aim to provide service continuity for those applications that require it, decrease complexity and control signaling for the network infrastructure, as

well as increase wireless capacity utilization by taking advantage of the broadcast nature of the medium. Beyond this, mobile networking scenarios should form a cross-scenario platform that can highlight how other scenarios can still maintain their respective performance metrics during periods of high mobility.

2.4. Infrastructure Sharing

A key idea in ICN is that the network should secure information objects per se, not the communications channel that they are delivered over. This means that hosts attached to an information-centric network can share resources on an unprecedented scale, especially when compared to what is possible in an IP network. All devices with network access and storage capacity can contribute their resources increasing the value of an information-centric network, although compensation schemes motivating users to contribute resources remain a research challenge primarily from a business perspective.

For example, Jacobson et al. [CBIS] argue that in ICN the "where and how" of obtaining information are new degrees of freedom. They illustrate this with a scenario involving a photo sharing application which takes advantage of whichever access network connectivity is available at the moment (WLAN, Bluetooth, and even SMS) without requiring a centralized infrastructure to synchronize between numerous devices. It is important to highlight that since the focus of communication changes, keep-alives in this scenario are simply unnecessary, as devices participating in the testbed network contribute resources in order to maintain user content consistency, not link state information as is the case in the host-centric paradigm. This means that the notion of "infrastructure" may be completely different in the future.

Muscariello et al. [SHARE], for instance, presented early work on an analytical framework that attempts to capture the storage/bandwidth tradeoffs that ICN enables and can be used as foundation for a network planning tool. In addition, Chai et al. [CL4M] explore the benefits of ubiquitous caching throughout an information-centric network and argue that "caching less can actually achieve more." These papers also sit alongside a variety of other studies that look at various scenarios such as caching HTTP-like traffic [CCNCT] and BitTorrent-like traffic [BTCACHE]. We observe that much more work is needed in order to understand how to make optimal use of all resources available in an information-centric network. In real-world deployments, policy and commercial considerations are also likely to affect the use of particular resources and more work is expected in this direction as well.

In conclusion, scenarios in this category, would cover the communication-computation-storage tradeoffs that an ICN deployment must consider. This would exercise features relating to network planning, perhaps capitalizing on user-provided resources, as well as operational and economical aspects of ICN and contrast them with other approaches. An obvious baseline to compare against in this regard is existing federations of IP-based Content Distribution Networks (CDNs), such as the ones discussed in the IETF CDNI WG.

2.5. Content Dissemination

Content dissemination has attracted more attention than other aspects of ICN. Scenarios in this category abound in the literature, including stored and streaming A/V distribution, file distribution, mirroring and bulk transfers, versioned content services (cf. Subversion-type revision control), as well as traffic aggregation.

Decentralized content dissemination with on-the-fly aggregation of information sources was envisaged in [N-Scen], where information objects can be dynamically assembled based on hierarchically structured subcomponents. For example, a video stream could be associated with different audio streams and subtitle sets, which can all be obtained from different sources. Using the topology depicted in Fig. 1 as an example, an application at C1 may end up obtaining, say, the video content from I1, but the user-selected subtitles from Px. Semantics and content negotiation, on behalf of the user, were also considered, e.g., for the case of popular tunes which may be available in different encoding formats. Effectively this scenario has the information consumer issuing independent requests for content based on information identifiers, and stitching the pieces together irrespective of "where" or "how" they were obtained.

A case in point for content dissemination are vehicular ad-hoc networks (VANETs), as an ICN approach may address their needs for information dissemination between vehicles better than today's solutions, as discussed in the following section. The critical part of information dissemination in a VANET scenario revolves around "where" and "when". For instance, one may be interested in traffic conditions 2 km ahead while having no interest in similar information about the area around the path origin. VANET scenarios may provide fertile ground for showcasing the ICN advantage with respect to content dissemination especially when compared with current host-centric approaches. That said, information integrity and filtering are challenges that must be addressed. As mentioned above, content dissemination scenarios in VANETs have a particular affinity to the mobility scenarios discussed in section 2.3.

Content dissemination scenarios, in general, have a large overlap with those described in the previous sections and are explored in several papers, such as [DONA] [PSI] [PSIMob] [NetInf] [CCN] [CBIS] [CCR], just to name a few. In addition, Chai et al. [CURLING] present a hop-by-hop hierarchical content resolution approach, which employs receiver-driven multicast over multiple domains, advocating another content dissemination approach. Yet, largely, work in this area did not address the issue of access authorization in detail. Often, the distributed content is mostly assumed to be freely accessible by any consumer. Distribution of paid-for or otherwise restricted content on a public ICN network requires more attention in the future. Fotiou et al. [ACDICN] consider a scheme to this effect but it still requires access to an authorization server to verify the user's status after the (encrypted) content has been obtained. This may effectively negate the advantage of obtaining the content from any node, especially in a disruption-prone or mobile network.

In summary, scenarios in this category aim to exercise primarily scalability, cost and performance attributes of content dissemination. Particularly, they should highlight the ability of an ICN to scale to billions of objects, while not exceeding the cost of existing content dissemination solutions (i.e., CDNs) and, ideally, increasing performance. These should be shown in a holistic manner, improving content dissemination for both information consumers and publishers of all sizes. We expect that in particular for content dissemination both extreme as well as typical scenarios can be specified drawing data from current CDN deployments.

2.6. Vehicular Networking

Users "on wheels" are interested in road safety, traffic efficiency, and infotainment applications that can be supported through vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) wireless communications. These applications exhibit unique features in terms of traffic generation patterns, delivery requirements, spatial and temporal scope, which pose great challenges to traditional networking solutions. VANETs, by their nature, are characterized by challenges such as fast-changing topology, intermittent connectivity, high node mobility, but also by the opportunity to combine information from different sources as each vehicle does not care about "who" delivers the named data objects.

ICN is an attractive candidate solution for vehicular networking, as it has several advantages. First, ICN fits well to the nature of typical vehicular applications that are geography- and time-dependent (e.g., road traveler information, accident warning, point-of-interest advertisements) and usually target vehicles in a given area,

regardless of their identity or IP address. These applications are likely to benefit from in-network and decentralized data caching and replication mechanisms. Second, content caching is particularly beneficial for intermittent on-the-road connectivity and can speed up data retrieval through content replication in several nodes. Caching can usually be implemented at relatively low cost in vehicles as the energy demands of the ICN device are likely to be a negligible fraction of the total vehicle energy consumption, thus allowing for sophisticated processing, continuous communication and adequate storage in the vehicle. Finally, ICN natively supports asynchronous data exchange between end-nodes. By using (and redistributing) cached named information objects, a mobile node can serve as a link between disconnected areas. In short, ICN can enable communication even under intermittent network connectivity, which is typical of vehicular environments with sparse roadside infrastructure and fast moving nodes.

The advantages of ICN in vehicular networks were preliminarily discussed in [EWC] and [DMND], and additionally investigated in [DNV2V] [RTIND] [CCNHV] [CCDIVN] [CCNVANET] [CRoWN]. For example, Bai and Krishnamachari [EWC] take advantage of the localized and dynamic nature of a VANET to explore how a road congestion notification application can be implemented. Wang et al. [DMND] consider data collection where Road-Side Units (RSUs) collect information from vehicles by broadcasting NDN-like INTEREST packets. The proposed architecture is evaluated using simulation in a grid topology and is compared against a host-centric alternative based on Mobile IP. See Fig. 3 for an indicative example of an urban VANET topology. Their results indicate high efficiency for ICN even at high speeds. That said, the authors point out that as this work is a preliminary exploration of ICN in vehicular environments, many issues remain to be evaluated, such as system scalability to large numbers of vehicles and the impact of vehicles forwarding Interests and relaying data for other vehicles.




Figure 3. Urban grid VANET topology.

As mentioned in the previous section, due to the short communication duration between a vehicle and the RSU, and the typically short time of sustained connectivity between vehicles, VANETs may be a good showcase for the ICN advantages with respect to content dissemination. Wang et al. [DNV2V], for instance, analyze the advantages of hierarchical naming for vehicular traffic information dissemination. Arnould et al. [CCNHV] apply ICN principles to safety information dissemination between vehicles with multiple radio interfaces. In [CCDIVN], TalebiFard and Leung use network coding techniques to improve content dissemination over multiple ICN paths. Amadeo et al. [CCNVANET][CRoWN] propose an application-independent ICN framework for content retrieval and distribution where the role of provider can be played equivalently by both vehicles and RSUs. ICN forwarding is extended through path-state information carried in Interest and Data packets, stored in a new data structure kept by vehicular nodes, and exploited also to cope with node mobility.

Typical scenarios for testing content distribution in VANETs may be highways with vehicles moving in straight lines, with or without RSUs along the road, as shown in Fig. 4. With a NDN approach in mind, for example, RSUs may send Interests to collect data from vehicles [DMND], or vehicles may send Interests to collect data from other peers [RTIND] or from RSUs [CCNVANET]. Fig. 2 applies to content dissemination in VANET scenarios as well, where C0 represents a vehicle which can obtain named information objects via multiple wireless peers and/or RSUs (I2 and I3 in the figure). Grid topologies such as the one illustrated in Fig. 3 should be considered in urban scenarios with RSUs at the crossroads or co-located with

traffic lights as in [CRoWN].




Figure 4. Highway VANET topology.

To summarize, VANET scenarios aim to exercise ICN deployment from various perspectives, including scalability, caching, transport, and mobility issues. There is a need for further investigation in (i) challenging scenarios (e.g., disconnected segments); (ii) scenarios involving both consumer and provider mobility; (iii) smart caching techniques which take into consideration node mobility patterns, spatial and temporal relevance, content popularity, and social relationships between users/vehicles; (iv) identification of new applications (beyond data dissemination and traffic monitoring) that could benefit from the adoption of an ICN paradigm in vehicular networks (e.g., mobile cloud, social networking).

2.7. Delay- and Disruption-Tolerance

Delay- and Disruption-Tolerant Networking (DTN) originated as a means to extend the Internet to interplanetary communications [DTN]. However, it was subsequently found to be an appropriate architecture for many terrestrial situations as well. Typically, this was where delays were greater than protocols such as TCP could handle, and where disruptions to communications were the norm rather than occasional annoyances, e.g. where an end-to-end path does not necessarily exist when communication is initiated. DTN has now been applied to many situations, including opportunistic content sharing, handling infrastructural issues during emergency situations (e.g., earthquakes) and providing connectivity to remote rural areas without existing Internet provision and little or no communications or power infrastructure.

The DTN architecture [RFC4838] is based on a "store, carry and forward" paradigm that has been applied extensively to situations where data is carried between network nodes by a "data mule", which carries bundles of data stored in some convenient storage medium

(e.g., a USB memory stick). With the advent of sensor and peer-to-peer (P2P) networks between mobile nodes, DTN is becoming a more commonplace type of networking than originally envisioned. Since ICN also does not rely on the familiar end-to-end communications paradigm, there are, thus, clear synergies [DTNICN]. It could therefore be argued that many of the key principles embodied within DTN also exist in ICN, as we explain next.

First, both approaches rely on in-network storage. In the case of DTN, bundles are stored temporarily on devices on a hop-by-hop basis.

In the case of ICN, information objects are also cached on devices in a similar fashion. As such, both paradigms must provision storage within the network.

Second, both approaches espouse late binding of names to locations due to the potentially large interval between request and response generation. In the case of DTN, it is often impossible to predict the exact location (in a disconnected topology) where a node will be found. Similarly, in the case of ICN, it is also often impossible to predict where an information object might be found. As such, the binding of a request/bundle to a destination (or routing locator) must be performed as late as possible.

Finally, both approaches treat data as a long-lived component that can exist in the network for extended periods of time. In the case of DTN, bundles are carried by nodes until appropriate next hops are discovered. In the case of ICN, information objects are typically cached until storage is exhausted. As such, both paradigms require a direct shift in the way applications interact with the network.

Through these similarities, it becomes possible to identify many DTN principles that are already in existence within ICN architectures. For example, ICN nodes will often retain information objects locally, making them accessible later on, much as DTN bundles are handled. Consequently, these synergies suggest strong potential for marrying the two technologies. This, for instance, could include building new integrated Information-Centric Delay Tolerant Network (ICDTN) protocols or, alternatively, building ICN schemes over existing DTN protocols (and vice versa).

The above similarities suggest that integration of the two principles would be certainly feasible. Beyond this, there are also a number of direct benefits identifiable. Through caching and replication, ICN offers strong information resilience, whilst, through store-and-forward, DTN offers strong connectivity resilience. As such, both architectures could benefit greatly from each other. Initial steps have already been taken in the DTN community to integrate ICN principles, e.g., the Bundle Protocol Query Block [BPQ] has been

proposed for the DTN Bundle Protocol [RFC5050]. Whilst, similarly, initial steps have also been taken in the ICN community, such as [SLINKY]. In fact, the SAIL project has developed a prototype implementation of NetInf running over the DTN Bundle Protocol.

Of course, in many circumstances, information-centricity is not appropriate for use in delay- and disruption-tolerant environments. This is particularly the case when information is not the key communications atom transmitted. Further, situations where a single sink is always used for receiving information may not warrant the identification and routing of independent information objects. However, there are a number of key scenarios where clear benefits could be gained by introducing information-centric principles into DTNs, two of which we describe later in this section.

For the purpose of evaluating the use of ICNs in a DTN setting, two key scenarios are identified in this document (note the rest of this section uses the term ICDTN). These are both prominent use cases that are currently active in both the ICN and DTN communities. The first is opportunistic content sharing, whilst the second is the use of ad hoc networks during disaster recovery (e.g., earthquakes). We discuss both types of scenarios in the context of a simulation-based evaluation: due to the scale and mobility of DTN-like setups, this is the primary method of evaluation used. Within the DTN community, the majority of simulations are performed using the Opportunistic Network Environment (ONE) simulator [ONE], which is referred to in this document. Before exploring the two scenarios, the key shared components of their simulation are discussed. This is separated into the two primary inputs that are required: the environment and the workload.

In both types of scenarios the environment can be abstractly modeled by a time series of active connections between device pairs. Unlike other scenarios in this document, an ICDTN scenario therefore does not depend on (relatively) static topologies but, rather, a set of time-varying disconnected topologies. In opportunistic networks, these topologies are actually products of the mobility of users. For example, if two users walk past each other, an opportunistic link can be created. There are two methods used to generate these mobility patterns and, in turn, the time series of topologies. The first is synthetic, whereby a (mathematical) model of user behavior is created in an agent-based fashion, e.g., random waypoint, Gauss-Markov. The second is trace-driven, whereby the mobility of real users is recorded and used. In both cases, the output is a sequence of time-stamped "contacts", i.e. periods of time in which two devices can communicate. An important factor missing from typical mobility traces, however, is the capacity of these contacts: how much data can be transferred? In both approaches to modeling mobility, links are

usually configured as Bluetooth or WiFi (ONE easily allows this, although lower layer considerations are ignored, e.g., interference).

This is motivated by the predominance of these technologies on mobile phones.

The workload in an ICDTN is modeled much like the workload within the other scenarios. It involves object creation/placement and object retrieval. Object creation/placement can either be done statically at the beginning of the simulations or, alternatively, dynamically based on a model of user behavior. In both cases, the latter is focused on as it models far better the characteristics of the scenarios.

Once the environment and workload has been configured, the next step is to decide the key metrics for the study. Unlike traditional networking, the quality of service expectation is typically far lower in an ICDTN, thereby moving away from metrics such as throughput. At a high-level, it is of clear interest to evaluate different ICN approaches with respect to both their delay- and disruption-tolerance, i.e., how effective is the approach when used in an environment subject to significant delay and/or disruption; and to their active support for operations in a DTN environment.

The two most prominent metrics considered in a host-centric DTN are delivery probability and delivery delay. The former relates to the probability by which a sent message will be received within a certain delay bound, whilst the latter captures the average length of time it takes for nodes to receive the message. These metrics are similarly important in an ICDTN, although they are slightly different due to the request-response nature of ICN. Therefore, the two most prominent evaluative metrics are satisfaction probability and satisfaction delay. The former refers to the probability by which an information request (e.g., Interest) will be satisfied (i.e., how often a Data response will be received). Satisfaction delay refers to the length of time it takes an information request to be satisfied.

Note that the key difference between the host-centric and information-centric metrics is the need for a round-trip rather than a one-way communication. Beyond this, depending on the focus of the work, other elements that may be investigated include name resolution, routing and forwarding in disconnected parts of the network; support for unidirectional links; number of round trips needed to complete a data transfer; long-term content availability (or resilience); efficiency in the face of disruption, and so on. It is also important to weigh these performance metrics against the necessary overheads. In the case of an ICDTN, this is generally measured by the number of message replicas required to access content. Note that routing in a DTN is often replication-based,

which leads to many copies of the same message.

2.7.1. Opportunistic Content Sharing

The first key baseline scenario in this context is opportunistic content sharing. This occurs when mobile nodes create opportunistic links between each other to share content of interest. For example, people riding on an underground train can pass news items between their mobile phones. Equally, content generated on the phones (e.g., tweets [TWIMIGHT]) could be stored for later forwarding (or even forwarded amongst interested passengers on the train). Such scenarios, clearly, must be based around either the altruistic or incentivized interaction amongst users. The latter is a particularly active area of research. These networks are often termed pocket-switched networks, as they are independently formed between the user devices. Here, the evaluative scenario of ICDTN microblogging is proposed. As previously discussed, the construction of such an evaluative scenario requires a formalization of its environment and workload. Fortunately, there exist a number of datasets that offer exactly this information required for microblogging.

In terms of the environment (i.e., mobility patterns), the Haggie project produced contact traces based on conference attendees using Bluetooth. These traces are best targeted at application scenarios in which a small group of (50-100) people are in a relatively confined space. In contrast, larger scale traces are also available, most notably MIT's Reality Mining project. These are better suited for cases where longer-term movement patterns are of interest.

The second input, workload, relates to the creation and consumption of microblogs (e.g. tweets). This can be effectively captured because subscriptions conveniently formalize who consumes what. For bespoke purposes, specific data can be directly collected from Twitter for trace-driven simulations. Several Twitter datasets are already available to the community containing a variety of data, ranging from Tweets to follower graphs. See <http://www.tweetarchivist.com>, <http://twapperkeeper.com>, <http://www.infochimps.com/collections/twitter-census>, and <http://socialcomputing.asu.edu/datasets/Twitter>. These datasets can therefore be used to extract information production, placement and consumption.

2.7.2. Emergency Support and Disaster Recovery

The second key baseline scenario in this context relates to the use of ICDTNs in emergency scenarios. In these situations it is typical

for infrastructure to be damaged or destroyed, leading to the collapse of traditional forms of communications (e.g., cellular telephone networks). This has been seen in the recent North Indian flooding, as well as the 2011 Tohoku earthquake and tsunami. Power problems often exacerbate the issue, with communication failures lasting for days. Therefore, in order to address this, DTNs have been used due to their high levels of resilience and independence from fixed infrastructure. The most prominent use of DTNs in disaster areas would be the dissemination of information, e.g., warnings and evacuation maps. Unlike the previous scenario, it can be assumed that certain users (e.g., emergency responders) are highly altruistic. However, it is likely many other users (e.g., endangered civilians) might become far more conservative in how they use their devices for battery conserving purposes. Here, we focus on the dissemination of standard broadcast information that should be received by all parties; this is something generally led by emergency responders.

For the environmental setup, there are no commonly used mobility traces for disaster zones, unlike in the previous scenario. This is clearly due to the difficulty (near impossibility) of acquiring them in a real setting. That said, various synthetic models are available. The Post Disaster Mobility Model [MODEL1] models civilians and emergency responders after a disaster has occurred, with people attempting to reach evacuation points (this has also been implemented in ONE). Aschenbruck et al. [MODEL2] focus on emergency responders, featuring the removal of nodes from the disaster zone, as well as things like obstacles (e.g., collapsed buildings). Cabrero et al. [MODEL3] also look at emergency responders, but focus on patterns associated with common procedures. For example, command and control centers are typically set up with emergency responders periodically returning. Clearly, the mobility of emergency responders is particularly important in this setting because they usually are the ones who will "carry" information into the disaster zone. It is recommended that one of these emergency-specific models are used during any evaluations, due to the inaccuracy of alternate models used for "normal" behavior.

The workload input in this evaluative scenario is far simpler than for the previous scenario. In emergency cases, the dissemination of individual pieces of information to all parties is the norm. This is often embodied using things like the Common Alert Protocol (CAP), which is an XML standard for describing warning message. It is currently used by various systems, including the Integrated Public Alert & Warning System and Google Crisis Response. As such, small objects (e.g., 512KB to 2MB) are usually generated containing text and images; note that the ONE simulator offers utilities to easily generate these. These messages are also always generated by central

authorities, therefore making the placement problem easier (they would be centrally generated and given to emergency responders to disseminate as they pass through the disaster zone). The key variable is therefore the generation rate, which is synonymous with the rate that microblogs are written in the previous scenario. This will largely be based on the type of disaster occurring, however, hourly updates would be an appropriate configuration. Higher rates can also be tested, based on the rate at which situations change (lands slides, for example, can exhibit highly dynamic properties).

To summarize, this section has highlighted the applicability of ICN principles to existing DTN scenarios. Two evaluative setups have been described in detail, namely, mobile opportunistic content sharing (microblogging) and emergency information dissemination.

2.8. Internet of Things

Advances in electronics miniaturization combined with low-power wireless access technologies (e.g., ZigBee, NFC, Bluetooth and others) have enabled the coupling of interconnected digital services with everyday objects. As devices with sensors and actuators connect into the network, they become "smart objects" and form the foundation for the so-called Internet of Things (IoT). IoT is expected to increase significantly the amount of content carried by the network due to machine-to-machine (M2M) communication as well as novel user interaction possibilities.

Yet, the full potential of IoT does not lie in simple remote access to smart object data. Instead, it is the intersection of Internet services with the physical world that will bring about the most dramatic changes. Burke [IoTeX], for instance, makes a very good case for creating everyday experiences using interconnected things through participatory sensing applications. In this case, inherent ICN capabilities for data discovery, caching, and trusted communication are leveraged to obtain sensor information and enable content exchange between mobile users, repositories, and applications.

Kutscher and Farrell [IWMT] discuss the benefits that ICN can provide in these environments in terms of naming, caching, and optimized transport. The Named Information URI scheme (ni) [RFC6920], for instance, could be used for globally unique smart object identification, although an actual implementation report is not currently available. Access to information generated by smart objects can be of varied nature and often vital for the correct operation of large systems. As such, supporting timestamping, security, scalability, and flexibility need to be taken into account.

Ghodsi et al. [NCOA] examine hierarchical and self-certifying naming schemes and point out that ensuring reliable and secure content naming and retrieval may pose stringent requirements (e.g., the necessity for employing PKI), which can be too demanding for low-powered nodes, such as sensors. That said, earlier work by Heidemann et al. [nWSN] shows that, for dense sensor network deployments, disassociating sensor naming from network topology and using named content at the lowest level of communication in combination with in-network processing of sensor data is feasible in practice and can be more efficient than employing a host-centric binding between node locator and the content existing therein.

Burke et al. [NDN1] describe the implementation of a lighting control building automation system where the security, naming and device discovery NDN mechanisms are leveraged to provide configuration, installation and management of residential and industrial lighting control systems. The goal is an inherently resilient system, where even smartphones can be used for control. Naming reflects fixtures with evolved identification and node reaching capabilities thus simplifying bootstrapping, discovery, and user interaction with nodes. The authors report that this ICN-based system requires less maintenance and troubleshooting than typical IP-based alternatives.

Biswas et al. [CIBUS] visualize ICN as a contextualized information-centric bus (CIBUS) over which diverse sets of service producers and consumers co-exist with different requirements. ICN is leveraged to unify different platforms to serve consumer-producer interaction in both infrastructure and ad hoc settings. Ravindran et al. [Homenet], show the application of this idea in the context of a home network, where consumers (residents) require policy-driven interactions with diverse services such as climate control, surveillance systems, and entertainment systems. Name-based protocols are developed to enable zero-configuration node and service discovery, contextual service publishing and subscription, policy-based routing and forwarding with name-based firewall, and hoc device-to-device communication.

IoT exposes ICN concepts to a stringent set of requirements which are exacerbated by the amount of nodes, as well as by the type and volume of information that must be handled. A way to address this is proposed in [IoTScope], which tackles the problem of mapping named information to an object, diverting from the currently typical centralized discovery of services and leveraging the intrinsic ICN scalability capabilities for naming. It extends the base [PURSUIT] design with hierarchically-based scopes, facilitating lookup, access, and modifications of only the part of the object information that the user is interested in. Another important aspect is how to efficiently address resolution and location of the information objects, particularly when large numbers of nodes are connected, as

in IoT deployments. In [ICN-DHT], Katsaros et al. propose a Distributed Hash Table (DHT) which is compared with DONA [DONA]. Their results show how topological routing information has a positive impact on resolution, at the expense of memory and processing overhead.

The use of ICN mechanisms in IoT scenarios faces the most dynamic and heterogeneous type of challenges, when taking into consideration the requirements and objectives of such integration. The disparity in technologies (not only in access technologies, but also in terms of end-node diversity such as sensors, actuators and their characteristics) as well as in the information that is generated and consumed in such scenarios, will undoubtedly bring about many of the considerations presented in the previous sections. For instance, IoT shares similarities with the constraints and requirements applicable to vehicular networking. Here, a central problem is the deployment of mechanisms that can use opportunistic connectivity in unreliable networking environments (similarly to the vehicular networking and DTN scenarios).

However, one important concern in IoT scenarios, also motivated by this strongly heterogeneous environment, is how content dissemination will be affected by the different semantics of the disparate information and content being shared. In fact, this is already a difficult problem that goes beyond the scope of ICN [SEMANT]. With the ability of the network nodes to cache forwarded information to improve future requests, a challenge arises regarding whether the ICN fabric should be involved in any kind of procedure (e.g., tagging) that facilitates the relationship or the interpretation of the different sources of information.

Another issue lies with the need for having energy-efficiency mechanisms related to the networking capabilities of IoT infrastructures. Often, the devices in IoT deployments have limited battery capabilities, and thus need low power consumption schemes working at multiple levels. In principle, energy efficiency gains should be observed from the inherent in-network caching capability. However, this might not be the most usual case in IoT scenarios, where the information (particularly from sensors, or controlling actuators) is more akin to real-time traffic, thus reducing the scale of potential savings due to ubiquitous in-network caching.

ICN approaches, therefore, should be evaluated with respect to their capacity to handle the content produced and consumed by extremely large numbers of diverse devices. IoT scenarios aim to exercise ICN deployment from different aspects, including ICN node design requirements, efficient naming, transport, and caching of time-restricted data. Scalability is particularly important in this

regard as the successful deployment of IoT principles could expand both device and content numbers dramatically beyond all current expectations.

2.9. Smart City

The rapid increase in urbanization sets the stage for the most compelling and challenging environments for networking. By 2050 the global population will reach nine billion people, 75% of which will dwell in urban areas. In order to cope with this influx, many cities around the world have started their transformation toward the Smart City vision. Smart cities will be based on the following innovation axes: smart mobility, smart environment, smart people, smart living, and smart governance. In development terms, the core goal of a smart city is to become a business-competitive and attractive environment, while serving citizen well being [CPG].

In a smart city, ICT plays a leading role and acts as the glue bringing together all actors, services, resources (and their interrelationships), that the urban environment is willing to host and provide [MVM]. ICN appears particularly suitable for these scenarios. Domains of interest include intelligent transportation systems, energy networks, health care, A/V communications, peer-to-peer and collaborative platforms for citizens, social inclusion, active participation in public life, e-government, safety and security, sensor networks. Clearly, this scenario has close ties to the vision of IoT, discussed in the previous section, as well as to vehicular networking.

Nevertheless, the road to build a real information-centric digital ecosystem will be long and more coordinated effort is required to drive innovation in this domain. We argue that smart city needs and ICN technologies can trigger a virtuous innovation cycle toward future ICT platforms. Recent concrete ICN-based contributions have been formulated for home energy management [iHEMS], geo-localized services [ACC], smart city services [IB], and traffic information dissemination in vehicular scenarios [RTIND]. Some of the proposed ICN-based solutions are implemented in real testbeds while others are evaluated through simulation.

Zhang et al. [iHEMS] propose a secure publish-subscribe architecture for handling the communication requirements of Home Energy Management Systems (HEMS). The objective is to safely and effectively collect measurement and status information from household elements, aggregate and analyze the data, and ultimately enable intelligent control decisions for actuation. They consider a simple experimental testbed for their proof-of-concept evaluation, exploiting open source code

for the ICN implementation, and emulating some node functionality in order to facilitate system operation.

A different scenario is considered in [ACC], where DHTs are employed for distributed, scalable, and geographically-aware service lookup in a smart city. Also in this case, the ICN application is validated by considering a small-scale testbed: a small number of nodes are realized with simple embedded PCs or specific hardware boards (e.g., for some sensor nodes); other nodes realizing the network connecting the principal actors of the tests are emulated with workstations. The proposal in [IB] draws from a smart city scenario (mainly oriented towards waste collection management) comprising sensors and moving vehicles, as well as a cloud computing system that supports data retrieval and storage operations. The main aspects of this proposal are analyzed via simulation using open source code which is publicly available. Some software applications are designed on real systems (e.g., PCs and smartphones).

With respect to evaluating ICN approaches in smart city scenarios, it is necessary to consider generic metrics useful to track and monitor progress on services results and also for comparing localities between themselves and learn from the best [ISODIS]. In particular, it is possible to select a specific set of Key Performance Indicators (KPIs) for a given project in order to evaluate its success. These KPIs may reflect the city's environmental and social goals, as well as its economic objectives, and they can be calculated at the global, regional, national, and local levels. Therefore, it is not possible to define a unique set of interesting metrics, but in the context of smart cities the KPIs should be characterized with respect to the developed set of services offered by using the ICN paradigm.

To sum up, smart city scenarios aim to exercise several ICN aspects in an urban environment. In particular, they can be useful to (i) analyze the capacity of using ICN for managing extremely large data sets; (ii) study ICN performance in terms of scalability in distributed services; (iii) verify the feasibility of ICN in a very complex application like vehicular communication systems; and (iv) examine the possible drawbacks related to privacy and security issues in complex networked environments.

3. Cross-scenario Considerations

This section discusses considerations that span multiple scenarios.

3.1. Multiply-connected Nodes and Economics

The evolution of, in particular, wireless networking technologies has

resulted in a convergence of the bandwidth and capabilities of various different types of network. Today a leading edge mobile telephone or tablet computer will typically be able to access a Wi-Fi access point, a 4G cellular network and the latest generation of Bluetooth local networking. Until recently a node would usually have a clear favorite network technology appropriate to any given environment. The choice would, for example, be primarily determined by the available bandwidth with cost as a secondary determinant. Furthermore, it is normally the case that a device only uses one of the technologies at a time for any particular application.

It seems likely that this situation will change so that nodes are able to use all of the available technologies in parallel. This will be further encouraged by the development of new capabilities in cellular networks including Small Cell Networks (SCN) and Heterogeneous Networks (HetNet) [SCN] [HetNet]. Consequently, mobile devices will have similar choices to wired nodes attached to multiple service providers allowing "multi-homing" via the various different infrastructure networks as well as potential direct access to other mobile nodes via Bluetooth or a more capable form of ad hoc Wi-Fi.

Infrastructure networks are generally under the control of separate economic entities that may have different policies about the information of an ICN deployed within their network caches. As ICN shifts the focus from nodes to information objects, the interaction between networks will likely evolve to capitalize on data location independence, efficient and scalable in-network named object availability and access via multiple paths. These interactions become critical in evaluating the technical and economic impact of ICN architectural choices, as noted in [ArgICN]. Beyond simply adding diversity in deployment options, these networks have the potential to alter the incentives among existing, and future, we may add, network players, as noted in [EconICN].

Moreover, such networks enable more numerous inter-network relationships where exchange of information may be conditioned on a set of multilateral policies. For example, shared SCNs are emerging as a cost-effective way to address coverage of complex environments such as sports stadiums, large office buildings, malls, etc. Such networks are likely to be a complex mix of different cellular and WLAN access technologies (such as HSPA, LTE, and Wi-Fi) as well as ownership models. It is reasonable to assume that access to content generated in such networks may depend on contextual information such as the subscription type, timing, and location of both the owner and requester of the content. The availability of such contextual information across diverse networks can lead to network inefficiencies unless data management can benefit from an information-centric approach. The "Event with Large Crowds"

relationships. For example, publishers should be able to indicate their willingness to partake in the caching market, proper reporting should be enabled to avoid fraud, and content should be made cacheable as much as possible to increase cache hit ratios.

Kutscher et al. [SAIL-B3] enable network interactions in the NetInf architecture using a name resolution service at domain edge routers, and a BGP-like routing system in the NetInf Default Free Zone. Business models and incentives are studied in [SAIL-A7] and [SAIL-A8], including scenarios where the access network provider (or a virtual CDN) guarantees QoS to end users using ICN. Fig. 6 illustrates a typical scenario topology from this work which involves an interconnectivity provider.

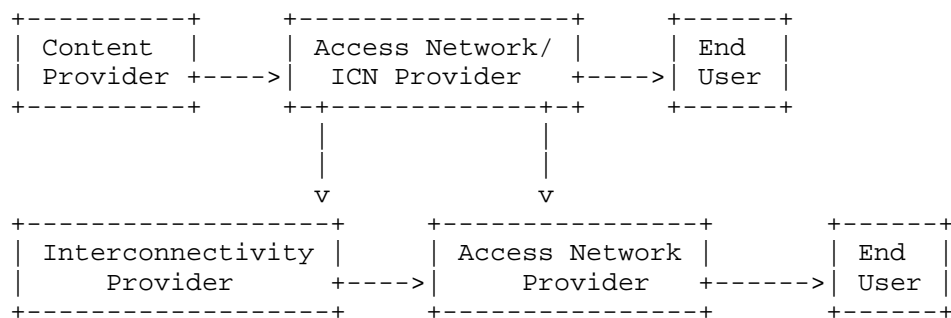


Figure 6. Setup and operating costs of network entities.

Jokela et al. [LIPSIN] propose a two-layer approach where additional rendezvous systems and topology formation functions are placed logically above multiple networks and enable advertising and routing content between them. Visala et al. [LANES] further describe an ICN architecture based on similar principles, which, notably, relies on a hierarchical DHT-based rendezvous interconnect. Rajahalme et al. [PSIRP1] describe a rendezvous system using both a BGP-like routing protocol at the edge and a DHT-based overlay at the core. Their evaluation model is centered around policy-compliant path stretch, latency introduced by overlay routing, caching efficacy, and overlay routing node load distribution.

Rajahalme et al. [ICCP] point out that ICN architectural changes may conflict with the current tier-based peering model. For example, changes leading to shorter paths between ISPs are likely to meet resistance from Tier-1 ISPs. Rajahalme [IDMcast] shows how incentives can help shape the design of specific ICN aspects, and in [IDArch] he presents a modeling approach to exploit these incentives. This includes a network model which describes the relationship between Autonomous Systems based on data inferred from the current

```

      O-----O
      +-----+ J +-----+
      |         |         |
      |         |         |
      O---+---O   *   O---+---O
      |   H   +-----+   I   |
      O-*-*--O   *   O-*-*--O
      * *         *         *
*****
*               * * *               *
O---*---O       O*-*-*O       O---*---O
|   E   +-----+   F   +-----+   G   +
O-*-*--O       O-----O       O-*-*--O
* *               * *
*****
*               *               *               *
O---*---O       O---*---O       O---*---O       O---*---O
|   A   |       |   B   +-----+   C   |       |   D   |
O-----O       O---+---O       O---+---O       O---+---O
data           data           data
|             |             |
O---V---O       O---V---O       O+++V---O
|   User   |       |   User   |       |   Data   |
O-----O       O-----O       O-----O

```

To sum up, the evaluation of ICN architectures across multiple network types should include a combination of technical and economic aspects, capturing their various interactions. These scenarios aim to illustrate scalability, efficiency and manageability, as well as traditional and novel network policies. Moreover, scenarios in this category should specifically address how different actors have proper incentives, not only in a pure ICN realm, but also during the migration phase towards this final state.

3.2. Energy Efficiency

ICN has prominent features which can be taken advantage of in order to significantly reduce the energy footprint of future communication networks. Of course, one can argue that specific ICN network elements may consume more energy than today's conventional network equipment due to the potentially higher energy demands for named-data processing en route. On balance, however, ICN introduces an architectural approach which may compensate on the whole and can even achieve higher energy efficiency rates when compared to the host-centric paradigm.

We elaborate on the energy efficiency potential of ICN based on three categories of ICN characteristics. Namely, we point out that a) ICN does not rely solely on end-to-end communication, b) ICN enables ubiquitous caching, and c) ICN brings awareness of user requests (as well as their corresponding responses) at the network layer thus permitting network elements to better schedule their transmission patterns.

First, ICN does not mandate perpetual end-to-end communication, which introduces a whole range of energy consumption inefficiencies due to the extensive signaling, especially in the case of mobile and wirelessly connected devices. This opens up new opportunities for accommodating sporadically connected nodes and could be one of the keys to an order of magnitude decrease in energy consumption over and above what other technological advances can contribute. For example, web applications often need to maintain state at both ends of a connection in order to verify that the authenticated peer is up and running. This introduces keep-alive timers and polling behavior with a high toll on energy consumption. Pentikousis [EEMN] discusses several related scenarios and explains why the current host-centric paradigm, which employs perpetual end-to-end connections, introduces built-in energy inefficiencies arguing that patches to make currently deployed protocols energy-aware cannot provide for an order of magnitude increase in energy efficiency.

Second, ICN network elements come with built-in caching capabilities, which is often referred to as ubiquitous caching. Pushing data objects to caches closer to end user devices, for example, could significantly reduce the amount of transit traffic in the core network, thereby reducing the energy used for data transport. Guan et al. [EECCN] study the energy efficiency of CCNx (based on their proposed energy model) and compare it with conventional content dissemination systems such as CDNs and P2P. Their model is based on the analysis of the topological structure and the average hop-length from all consumers to the nearest cache location. Their results show that an information-centric approach can be more energy efficient in

delivering popular and small size content. In particular, they also note that different network element design choices (e.g. the optical bypass approach) can be more energy-efficient in delivering infrequently accessed content.

Lee et al. [EECD] investigate the energy efficiency of various network devices deployed in access, metro, and core networks for both CDNs and ICN. They use trace-based simulations to show that an ICN approach can substantially improve the network energy efficiency for content dissemination mainly due to the reduction in the number of hops required to obtain a data object, which can be served by intermediate nodes in ICN. They also emphasize that the impact of cache placement (in incremental deployment scenarios) and local/cooperative content replacement strategies need to be carefully investigated in order to better quantify the energy efficiencies arising from adopting an ICN paradigm.

Third, ICN elements are aware of the user request and its corresponding data response, due to the nature of name-based routing, they can employ power consumption optimization processes for determining their transmission schedule or powering down inactive network interfaces. For example, network coding [NCICN] or adaptive video streaming [COAST] can be used in individual ICN elements so that redundant transmissions, possibly passing through intermediary networks, could be significantly reduced, thereby saving energy by avoiding to carry redundant traffic.

Alternatively, approaches that aim to simplify routers, such as [PURSUIT], could also reduce energy consumption by pushing routing decisions to a more energy-efficient entity. Along these lines, Ko et al. [ICNDC] design a data center network architecture based on ICN principles and decouple the router control-plane and data-plane functionalities. Thus, data forwarding is performed by simplified network entities while the complicated routing computation is carried out in more energy-efficient data centers.

To summarize, energy efficiency has been discussed in ICN evaluation studies but most published work is preliminary in nature. Thus, we suggest that more work is needed in this front. Evaluating energy efficiency does not require the definition of new scenarios or baseline topologies, but does require the establishment of clear guidelines so that different ICN approaches can be compared not only in terms of scalability, for example, but also in terms to power consumption.

3.3. Operation across Multiple Network Paradigms

Today the overwhelming majority of networks are integrated with the well-connected Internet with IP at the "waist" of the technology hourglass. However there is a large amount of ongoing research into alternative paradigms that can cope with conditions other than the standard set assumed by the Internet. Perhaps the most advanced of these is Delay- and Disruption-Tolerant Networking (DTN). DTN is considered as one of the scenarios for the deployment in section 2.7 but here we consider how ICN can operate in an integrated network that has essentially disjoint "domains" (a highly-overloaded term!) or regions that use different network paradigms and technologies, but with gateways that allow interoperation.

ICN operates in terms of named data objects so that requests and deliveries of information objects can be independent of the networking paradigm. Some researchers have contemplated some form of ICN becoming the new waist of the hourglass as the basis of a future reincarnation of the Internet, e.g., [ArgICN], but there are a large number of problems to resolve, including authorization and access control and transactional operation for applications such as banking, before some form of ICN can be considered as ready to take over from IP as the dominant networking technology. In the meantime, ICN architectures will operate in conjunction with existing network technologies as an overlay or in cooperation with the lower layers of the "native" technology.

It seems likely that as the reach of the "Internet" is extended, other technologies such as DTN will be needed to handle scenarios such as space communications where inherent delays are too large for TCP/IP to cope with effectively. Thus, demonstrating that ICN architectures can work effectively in and across the boundaries of different networking technologies will be important.

The NetInf architecture in particular targets the inter-domain scenario by the use of a convergence layer architecture [SAIL-B3] and PSIRP/PURSUIT is envisaged as a candidate for an IP replacement.

The key items for evaluation over and above the satisfactory operation of the architecture in each constituent domain will be to ensure that requests and responses can be carried across the network boundaries with adequate performance and do not cause malfunctions in applications or infrastructure because of the differing characteristics of the gatewayed domains.

4. Summary

This document presents a wide range of different application areas in which the use of information-centric network designs have been

evaluated in the peer-reviewed literature. Evidently, this broad range of scenarios illustrates the capability of ICN to potentially address today's problems in an alternative and better way than host-centric approaches as well as to point to future scenarios where ICN may be applicable. We believe that by putting different ICN systems to the test in diverse application areas, the community will be better equipped to judge the potential of a given ICN proposal and therefore subsequently invest more effort in developing it further. It is worth noting that this document collected different kinds of considerations, as a result of our ongoing survey of the literature and the discussion within ICNRP, which we believe would have otherwise remained unnoticed in the wider community. As a result, we expect that this document can assist in fostering the applicability and future deployment of ICN over a broader set of operations, as well as possibly influencing and enhancing the currently-available base ICN proposals and possibly assist in defining new scenarios where ICN would be applicable.

We conclude this document with a brief summary of the evaluation aspects we have seen across a range of scenarios.

The scalability of different mechanisms in an ICN architecture stands out as an important concern (cf. sections 2.1, 2.2, 2.5, 2.6, 2.8, 2.9, 3.1) as does network, resource and energy efficiency (cf. sections 2.1, 2.3, 2.4, 3.1, 3.2). Operational aspects such as network planing, manageability, reduced complexity and overhead (cf. sections 2.2, 2.3, 2.4, 2.8, 3.1) should not be neglected especially as ICN architectures are evaluated with respect to their potential for deployment in the real world. Accordingly, further research in economic aspects as well as in the communication, computation, and storage tradeoffs entailed in each ICN architecture is needed.

With respect to purely technical requirements, support for multicast, mobility, and caching lie at the core of many scenarios (cf. sections 2.1, 2.3, 2.5, 2.6). ICN must also be able to cope when the Internet expands to incorporate additional network paradigms (cf. section 3.3). We have also seen that being able to address stringent QoS requirements and increase reliability and resilience should also be evaluated following well-established methods (cf. sections 2.2, 2.8, 2.9).

Finally, we note that new applications that significantly improve the end user experience and forge a migration path from today's host-centric paradigm could be the key to a sustained and increasing deployment of the ICN paradigm in the real world (cf. sections 2.2, 2.3, 2.6, 2.8, 2.9).

5. Security Considerations

This document does not impact the security of the Internet.

6. IANA Considerations

This document presents no IANA considerations.

7. Acknowledgments

Dorothy Gellert contributed to an earlier version of this document.

This document has benefited from reviews, pointers to the growing ICN literature, suggestions, comments and proposed text provided by the following members of the IRTF Information-Centric Networking Research Group (ICNRG), listed in alphabetical order: Marica Amadeo, Hitoshi Asaeda, Claudia Campolo, Luigi Alfredo Grieco, Myeong-Wuk Jang, Ren Jing, Hongbin Luo, Priya Mahadevan, Will Liu, Ioannis Psaras, Spiros Spirou, Dirk Trossen, Jianping Wang, Yuanzhe Xuan, and Xinwen Zhang.

The authors would like to thank Mark Stapp, Juan Carlos Zuniga, and G.Q. Wang for their comments and suggestions as part of their open and independent review of this document within ICNRG.

8. Informative References

- [RFC5743] Falk, A., "Definition of an Internet Research Task Force (IRTF) Document Stream", RFC 5743, December 2009.
- [RFC5050] Scott, K. and S. Burleigh, "Bundle Protocol Specification", RFC 5050, November 2007.
- [RFC4838] Cerf, V., Burleigh, S., Hooke, A., Torgerson, L., Durst, R., Scott, K., Fall, K., and H. Weiss, "Delay-Tolerant Networking Architecture", RFC 4838, April 2007.
- [RFC5568] Koodli, R., Ed., "Mobile IPv6 Fast Handovers", RFC 5568, July 2009.
- [RFC6920] Farrell, S., Kutscher, D., Dannewitz, C., Ohlman, B., Keranen, A., and P. Hallam-Baker, "Naming Things with Hashes", RFC 6920, April 2013.
- [NetInf] Ahlgren, B. et al., "Design considerations for a network of information", Proc. CoNEXT Re-Arch Workshop. ACM,

2008.

- [CCN] Jacobson, V. et al., "Networking Named Content", Proc. CoNEXT. ACM, 2009.
- [NDNP] Zhang, L. et al., "Named Data Networking (NDN) Project", NDN Technical Report NDN-0001, Oct. 2010. Available: <http://named-data.net/publications/techreports/>
- [PSI] Trossen, D. and G. Parisis, "Designing and realizing an information-centric internet", IEEE Commun. Mag., vol. 50, no. 7, July 2012.
- [DONA] Koponen, T. et al., "A Data-Oriented (and Beyond) Network Architecture", Proc. SIGCOMM. ACM, 2007.
- [SoA1] Ahlgren, B. et al., "A survey of information-centric networking", IEEE Commun. Mag., vol. 50, no. 7, July 2012.
- [SoA2] Xylomenos, G., et al. "A survey of information-centric networking research", IEEE Communications Surveys and Tutorials (2013): 1-26.
- [ICN-SN] Mathieu, B. et al., "Information-centric networking: a natural design for social network applications", IEEE Commun. Mag., vol. 50, no. 7, July 2012.
- [VPC] Kim, J. et al., "Content Centric Network-based Virtual Private Community", Proc. ICCE. IEEE, 2011.
- [CBIS] Jacobson, V. et al., "Custodian-Based Information Sharing", IEEE Commun. Mag., vol. 50, no. 7, July 2012.
- [VPC2] Kim, D. and J. Lee, "CCN-based virtual private community for extended home media service", IEEE Trans. Consumer Electronics, vol. 57, no. 2, May 2011.
- [CCR] Arianfar, S. et al., "On content-centric router design and implications", Proc. CoNEXT Re-Arch Workshop. ACM, 2010.
- [VoCCN] Jacobson, V. et al., "VoCCN: Voice-over Content-Centric Networks", Proc. CoNEXT Re-Arch Workshop. ACM, 2009.
- [NDNpb] Xuan, Y. and Z. Yan, "Enhancing Routing Efficiency in Named Data Network with Piggybacked Interest", Proc. CFI. ACM, 2013.
- [ACT] Zhu, Z. et al., "ACT: Audio Conference Tool Over Named

- Data Networking", Proc. SIGCOMM ICN Workshop. ACM, 2011.
- [G-COPSS] Chen, J. et al., "G-COPSS: A Content Centric Communication Infrastructure for Gaming Applications", Proc. ICDCS. IEEE, 2012.
- [MMIN] Pentikousis, K., and P. Bertin., "Mobility Management in Infrastructure Networks." IEEE Internet Computing 17, no. 5 (2013): 74-79.
- [SCES] Allman, M. et al., "Enabling an Energy-Efficient Future Internet through Selectively Connected End Systems", Proc. HotNets-VI. ACM, 2007.
- [EEMN] Pentikousis, K., "In Search of Energy-Efficient Mobile Networking", IEEE Commun. Mag., vol. 48, no. 1, Jan. 2010.
- [MOBSURV] Tyson, G. et al., "A Survey of Mobility in Information-Centric Networks: Challenges and Research Directions", Proc. MobiHoc Workshop on Emerging Name-Oriented Mobile Networking Design. ACM, 2012.
- [N-Scen] Dannewitz, C. et al., "Scenarios and research issues for a Network of Information", Proc. MobiMedia. ICST, 2012.
- [DTI] Ott, J. and D. Kutscher, "Drive-thru Internet: IEEE 802.11b for 'automobile' users", Proc. INFOCOM. IEEE, 2004.
- [PSIMob] Xylomenos, G. et al., "Caching and Mobility Support in a Publish-Subscribe Internet Architecture", IEEE Commun. Mag., vol. 50, no. 7, July 2012.
- [mNetInf] Pentikousis, K. and T. Rautio, "A Multiaccess Network of Information", Proc. WoWMoM. IEEE, 2010.
- [HybICN] Lindgren, A., "Efficient content distribution in an information-centric hybrid mobile networks", Proc. CCNC. IEEE, 2011.
- [E-CHANET] M. Amadeo, et al., "E-CHANET: Routing, Forwarding and Transport in Information-Centric Multihop Wireless Networks", Computer Communications. Elsevier, Jan. 2013 online.
- [MobiA] Meisel, M. et al., "Ad Hoc Networking via Named Data", Proc. MobiArch. ACM 2010.

- [CCNMANET] Oh, S. Y. et al., "Content Centric Networking in Tactical and Emergency MANETs", Proc. Wireless Days. IFIP, 2010.
- [RTIND] Wang, L. et al., "Rapid Traffic Information Dissemination Using Named Data", Proc. MobiHoc NoM workshop. ACM, 2012.
- [CCNVANET] Amadeo, M. et al., "Content-Centric Networking: is that a Solution for Upcoming Vehicular Networks?", Proc. VANET. ACM, 2012.
- [ABC] Gustafsson, E., and A. Jonsson. "Always best connected." Wireless Communications, IEEE 10.1 (2003): 49-55.
- [SHARE] Muscariello, L. et al., "Bandwidth and storage sharing performance in information centric networking", Proc. SIGCOMM ICN Workshop. ACM, 2011.
- [CL4M] Chai, W. K. et al., "Cache 'Less for More' in Information-centric Networks", Proc. Networking. IFIP, 2012.
- [CCNCT] Psaras, I. et al., "Modelling and Evaluation of CCN-Caching Trees", In Proc. of the 10th international IFIP conference on Networking, Valencia, Spain, May 2011.
- [BTCACHE] Tyson, G. et al., "A Trace-Driven Analysis of Caching in Content-Centric Networks", Proc. ICCCN. IEEE, 2012.
- [CURLING] Chai, W. K. et al., "CURLING: Content-Ubiquitous Resolution and Delivery Infrastructure for Next-Generation Services", IEEE Commun. Mag., vol. 49, no. 3, Mar. 2011.
- [ACDICN] Fotiou, N. et al., "Access control enforcement delegation for information-centric networking architectures", Proc. SIGCOMM ICN Workshop. ACM, 2012.
- [EWC] Bai, F. and B. Krishnamachari, "Exploiting the wisdom of the crowd: localized, distributed information-centric VANETs", IEEE Commun. Mag., vol. 48, no. 5, May 2010.
- [DMND] Wang, J., R. Wakikawa, and L. Zhang, "DMND: Collecting data from mobiles using Named Data", Proc. Vehicular Networking Conference (VNC). IEEE, 2010.
- [DNV2V] Wang, L. et al., "Data Naming in Vehicle-to-Vehicle Communications", Proc. INFOCOM NOMEN workshop. IEEE, 2012.
- [CCNHV] Arnould, G. et al., "A Self-Organizing Content Centric

- Network Model for Hybrid Vehicular Ad-Hoc Networks". Proc. DIVANet. ACM, 2011.
- [CCDIVN] TalebiFard, P. and V.C.M. Leung, "A Content Centric Approach to Dissemination of Information in Vehicular Networks". Proc. DIVANet. ACM, 2012.
- [CROWN] Amadeo, M. et al., "CROWN: Content-Centric Networking in Vehicular Ad Hoc Networks", IEEE Communications Letters, vol. 16, no. 9, Sept. 2012.
- [SCN] Hoydis, J., et al., "Green small-cell networks", IEEE Vehicular Technology Magazine, vol. 6, no.1, pp. 37-43, March 2011.
- [HetNet] Li, H., et al. "Efficient HetNet implementation using broadband wireless access with fiber-connected massively distributed antennas architecture." IEEE Wireless Communications, vol. 18, no.3, pp. 72-78, June 2011.
- [ArgICN] Trossen, D. et al., "Arguments for an information centric internetworking architecture", ACM SIGCOMM CCR, 40:26-33, Apr. 2010.
- [EconICN] Agyapong, P. and M. Sirbu, "Economic Incentives in Information Centric Networking: Implications for Protocol Design and Public Policy", IEEE Commun. Mag., vol. 50, no. 12, Dec. 2012.
- [SAIL-B3] Kutscher, D. (ed.) et al., "Final NetInf Architecture", SAIL Project Deliverable D-B.3 , Jan. 2013. Available: <http://www.sail-project.eu/deliverables/>
- [MLDHT] Liu, H. et al., "A multi-level DHT routing framework with aggregation", Proc. SIGCOMM ICN Workshop. ACM, 2012.
- [NCOA] Ghodsi, A. et al., "Naming in Content-oriented Architectures", Proc. SIGCOMM ICN Workshop. ACM, 2011.
- [RP-NDN] DiBenedetto, S. et al., "Routing policies in named data networking", Proc. SIGCOMM ICN Workshop. ACM, 2011.
- [SAIL-A7] Salo, J. (ed.) et al., "New business models and business dynamics of the future networks", SAIL Project Deliverable D-A.7, Aug. 2011. Available: <http://www.sail-project.eu/deliverables/>
- [SAIL-A8] Zhang, N. (ed.) et al., "Evaluation of business models",

- SAIL Project Deliverable D-A.8, Jan. 2013. Available:
<http://www.sail-project.eu/deliverables/>
- [LIPSIN] Jokela, P. et al., "LIPSIN: line speed publish/subscribe inter-networking", Proc. of ACM SIG-COMM 2009.
- [LANES] Visala, K. et al., "LANES: An Inter-Domain Data-Oriented Routing Architecture", Proc. CoNEXT Re-Arch Workshop. ACM, 2009.
- [PSIRP1] Rajahalme, J. et al., "Inter-Domain Rendezvous Service Architecture", PSIRP Technical Report TR09-003, Dec. 2009.
- [ICCP] Rajahalme J. et al., "Incentive-Compatible Caching and Peering in DataOriented Networks", Proc. CoNEXT Re-Arch Workshop. ACM, 2008.
- [IDMcast] Rajahalme, J., "Incentive-informed Inter-domain Multicast", Proc. Global Internet Symposium 2010.
- [IDArch] Rajahalme, J., "Inter-domain incentives and Internet architecture", PhD. Dissertation, Aalto University, Aug. 2012.
- [PURSUIT] Fotiou, N. et al., "Developing Information Networking Further: From PSIRP to PURSUIT", Proc. BROADNETS. ICST, 2010.
- [EECCN] Guan, K. et al., "On the Energy Efficiency of Content Delivery Architectures", Proc. ICC Workshops. IEEE, 2011.
- [EECD] Lee, U., Rimal, I., Kilper, D., and V. Hilt, "Toward energy-efficient content dissemination", IEEE Network, vol. 25, no. 2, March-April 2011.
- [NCICN] Montpetit, M. J., Westphal, C., and D. Trossen, "Network coding meets information-centric networking: an architectural case for information dispersion through native network coding", Proc. MOBIHOC NoM Workshop. ACM, 2012.
- [COAST] Gruneberg, K. et al., "File format specification and 3D video codec", COAST Project Deliverable D5.1, July 2011. Available: <http://www.coast-fp7.eu/deliverables.html>
- [ICNDC] Ko, B. J., Pappas, V., Raghavendra, R., Song, Y., Dilmaghani, R. B., Lee, K.-w., and D. Verma, "An information-centric architecture for data center

- networks", Proc. SIGCOMM ICN Workshop. ACM, 2012.
- [DTN] Fall, K., "A delay-tolerant network architecture for challenged internets", Proc. SIGCOMM. ACM, 2003.
- [DTNICN] Tyson, G., Bigham, J. and E. Bodanese, "Towards an Information-Centric Delay-Tolerant Network", Proc. IEEE INFOCOM NOMEN 2013, Turin, Italy.
- [BPQ] Farrell, S., Lynch, A., Kutscher, D., and A. Lindgren, "Bundle protocol query extension block", draft-irtf-dtnrg-bpq-00 (work in progress), May 2012.
- [SLINKY] Kawadia, V., Riga, N., Opper, J., and D. Sampath, "Slinky: An adaptive protocol for content access in disruption-tolerant ad hoc networks", in Proc. MobiHoc Workshop on Tactical Mobile Ad Hoc Networking, 2011.
- [ONE] The Opportunistic Network Environment simulator.
Available: <http://www.netlab.tkk.fi/tutkimus/dtn/theone>
- [TWIMIGHT] Hossmann, T., et al. "Twitter in disaster mode: smart probing for opportunistic peers", Proc. 3rd ACM International Workshop on Mobile Opportunistic Networks. ACM, 2012.
- [MODEL1] Uddin, M.Y.S., Nicol, D.M., Abdelzaher, T.F., and R.H. Kravets, "A post-disaster mobility model for Delay Tolerant Networking", Simulation Conference (WSC), Proceedings of the 2009 Winter , vol., no., pp.2785,2796, 13-16 Dec. 2009
- [MODEL2] Aschenbruck, N., Gerhards-Padilla, E., and P. Martini, "Modeling mobility in disaster area scenarios.", Performance Evaluation 66, no. 12 (2009): 773-790.
- [MODEL3] Cabrero, S., Paneda, X.G., Plagemann, T., Melendi, D., and R. Garcia, "An Overlay Routing Protocol for Video over sparse MANETs in Emergencies", Cadernos de Informatica 6, no. 1 (2011): 195-202.
- [IoTEx] Burke, J., "Authoring Place-based Experiences with an Internet of Things: Tussles of Expressive, Operational, and Participatory Goals", Proc. Interconnecting Smart Objects with the Internet Workshop. IAB, 2011.
- [IWMT] Kutscher, D. and S. Farrell, "Towards an Information-Centric Internet with more Things", Proc. Interconnecting

- Smart Objects with the Internet Workshop. IAB, 2011.
- [nWSN] Heidemann, J. et al., "Building efficient wireless sensor networks with low-level naming", Proc. SOSP. ACM, 2001.
- [NDN1] Burke, J. et al., "Authenticated Lighting Control Using Named Data Networking", NDN Technical Report NDN-0011, Oct. 2012.
- [CIBUS] Biswas, T. et al., "Contextualized Information-Centric Home Network", Proc. ACM SIGCOMM. ACM, 2013.
- [Homenet] Ravindran, R. et al., "Information-centric Networking based Homenet", Proc. International Workshop on Management of the Future Internet (ManFI). IFIP/IEEE, 2013.
- [IoTScope] Marias, G.F. et al., "Efficient information lookup for the Internet of Things", Proc. WoWMoM. IEEE, 2012.
- [ICN-DHT] Katsaros, K. et al., "On inter-domain name resolution for information-centric networks", Proc. Networking. Springer, 2012.
- [SEMANT] Sheth, A. et al., "Semantic Sensor Web," Internet Computing, IEEE , vol.12, no.4, pp.78,83, July-Aug. 2008
- [CPG] Cianci, I. et al., "Content Centric Services in Smart Cities", Proc. NGMAST. IEEE, 2012.
- [MVM] Hernandez-Munoz, J.M. et al., "Smart cities at the forefront of the future Internet", The Future Internet. Springer, 2011.
- [iHEMS] Zhang, J. et al., "iHEMS: An Information-Centric Approach to Secure Home Energy Management", Proc. SmartGridComm. IEEE, 2012.
- [ACC] Andreini, F. et al., "A scalable architecture for geo-localized service access in smart cities", Proc. Future Network and Mobile Summit. IEEE, 2011.
- [IB] Idowu, S. and N. Bari, "A Development Framework for Smart City Services, Integrating Smart City Service Components", Master Thesis. Lulea University of Technology, 2012.
- [ISODIS] ISO/DIS 37120, Sustainable development and resilience of communities --Indicators for city services and quality of life, 2013.

Authors' Addresses

Kostas Pentikousis (editor)
EICT GmbH
Torgauer Strasse 12-15
10829 Berlin
Germany

Email: k.pentikousis@eict.de

Borje Ohlman
Ericsson Research
S-16480 Stockholm
Sweden

Email: Borje.Ohlman@ericsson.com

Daniel Corujo
Instituto de Telecomunicacoes
Campus Universitario de Santiago
P-3810-193 Aveiro
Portugal

Email: dcorujo@av.it.pt

Gennaro Boggia
Dep. of Electrical and Information Engineering
Politecnico di Bari
Via Orabona 4
70125 Bari
Italy

Email: g.boggia@poliba.it

Gareth Tyson
School and Electronic Engineering and Computer Science
Queen Mary, University of London
United Kingdom

Email: gareth.tyson@eeecs.qmul.ac.uk

Elwyn Davies

Trinity College Dublin/Folly Consulting Ltd
Dublin, 2
Ireland

Email: davieseb@scss.tcd.ie

Antonella Molinaro
Dep. of Information, Infrastructures, and Sustainable
Energy Engineering
Universita' Mediterranea di Reggio Calabria
Via Graziella 1
89100 Reggio Calabria
Italy

Email: antonella.molinaro@unirc.it

Suyong Eum
National Institute of Information and Communications Technology
4-2-1, Nukui Kitamachi, Koganei
Tokyo 184-8795
Japan

Phone: +81-42-327-6582
Email: suyong@nict.go.jp

ICNRG
Internet-Draft
Intended status: Informational
Expires: July 1, 2016

J. Seedorf
NEC
M. Arumaithurai
University of Goettingen
A. Tagami
KDDI R&D Labs
K. Ramakrishnan
University of California
N. Blefari Melazzi
University Tor Vergata
December 29, 2015

Using ICN in disaster scenarios
draft-seedorf-icn-disaster-05

Abstract

Information Centric Networking (ICN) is a new paradigm where the network provides users with named content, instead of communication channels between hosts. This document outlines some research directions for Information Centric Networking with respect to applying ICN approaches for coping with natural or human-generated, large-scale disasters.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on July 1, 2016.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Disaster Scenarios	3
3. Research Challenges and Benefits of ICN	4
3.1. High-Level Research Challenges	4
3.2. How ICN can be Beneficial	5
3.3. ICN as Starting Point vs. Existing DTN Solutions	7
4. Use Cases and Requirements	8
5. Solution Design	9
6. The GreenICN Project	11
7. Conclusion	12
8. References	12
8.1. Normative References	12
8.2. Informative References	12
Appendix A. Acknowledgment	14
Authors' Addresses	14

1. Introduction

This document summarizes some research challenges for coping with natural or human-generated, large-scale disasters. In particular, the document discusses potential directions for applying Information Centric Networking (ICN) to address these challenges.

There are existing research approaches (for instance, see further the discussions in the IETF DTN Research Group [dtnrg]) and an IETF specification [RFC5050] for disruptant tolerant networking, which is a key necessity for communicating in the disaster scenarios we are considering in this document (see further Section 3.1).

'Disconnection tolerance' can thus be achieved with these existing DTN approaches. However, while these approaches can provide independence from an existing communication infrastructure (which indeed may not work anymore after a disaster has happened), ICN offers as key concepts suitable naming schemes and multicast communication which together enable many key (publish/subscribe-based) use cases for communication after a disaster (e.g. message prioritisation, one-to-many delivery of important messages, or group

communication among rescue teams, see further Section 4). One could add such features to existing DTN protocols and solutions; however, in this document we explore the use of ICN as starting point for building a communication architecture that works well before and after a disaster. We discuss the relationship between the ICN approaches (for enabling communication after a disaster) discussed in this document with existing work from the DTN community in more depth in Section 3.3.

Section 2 gives some examples of what can be considered a large-scale disaster and what the effects of such disasters on communication networks are. Section 3 outlines why ICN can be beneficial in such scenarios and provides a high-level overview on corresponding research challenges. Section 4 describes some concrete use cases and requirements for disaster scenarios. In Section 5, some concrete ICN-based solutions approaches are outlined. Related research activities are ongoing in the GreenICN research project; Section 6 provides an overview of this project.

2. Disaster Scenarios

An enormous earthquake hit Northeastern Japan (Tohoku areas) on March 11, 2011, and caused extensive damages including blackouts, fires, tsunamis and a nuclear crisis. The lack of information and means of communication caused the isolation of several Japanese cities. This impacted the safety and well-being of residents, and affected rescue work, evacuation activities, and the supply chain for food and other essential items. Even in the Tokyo area that is 300km away from the Tohoku area, more than 100,000 people became 'returner' refugees, who could not reach their homes because they had no means of public transportation (the Japanese government has estimated that more than 6.5 million people would become returner refugees if such a catastrophic disaster were to hit the Tokyo area).

That earthquake in Japan also showed that the current network is vulnerable against disasters and that mobile phones have become the lifelines for communication including safety confirmation. The aftermath of a disaster puts a high strain on available resources due to the need for communication by everyone. Authorities such as the President/Prime-Minister, local authorities, Police, fire brigades, and rescue and medical personnel would like to inform the citizens of possible shelters, food, or even of impending danger. Relatives would like to communicate with each other and be informed about their wellbeing. Affected citizens would like to make enquiries of food distribution centres, shelters or report trapped, missing people to the authorities. Moreover, damage to communication equipment, in addition to the already existing heavy demand for communication highlights the issue of fault-tolerance and energy efficiency.

Additionally, disasters caused by humans such as a terrorist attack may need to be considered, i.e. disasters that are caused deliberately and willfully and have the element of human intent. In such cases, the perpetrators could be actively harming the network by launching a Denial-of-Service attack or by monitoring the network passively to obtain information exchanged, even after the main disaster itself has taken place. Unlike some natural disasters that are predictable using weather forecasting technologies and have a slower onset and occur in known geographical regions and seasons, terrorist attacks may occur suddenly without any advance warning. Nevertheless, there exist many commonalities between natural and human-induced disasters, particularly relating to response and recovery, communication, search and rescue, and coordination of volunteers.

The timely dissemination of information generated and requested by all the affected parties during and the immediate aftermath of a disaster is difficult to provide within the current context of global information aggregators (such as Google, Yahoo, Bing etc.) that need to index the vast amounts of specialized information related to the disaster. Specialized coverage of the situation and timely dissemination are key to successfully managing disaster situations. We believe that network infrastructure capability provided by Information Centric Networks can be suitable, in conjunction with application and middleware assistance.

3. Research Challenges and Benefits of ICN

3.1. High-Level Research Challenges

Given a disaster scenario as described in Section 2, on a high-level one can derive the following (incomplete) list of corresponding technical challenges:

- o Enabling usage of functional parts of the infrastructure, even when these are disconnected from the rest of the network: Assuming that parts of the network infrastructure (i.e. cables/links, routers, mobile bases stations, ...) are functional after a disaster has taken place, it is desirable to be able to continue using such components for communication as much as possible. This is challenging when these components are disconnected from the backhaul, thus forming fragmented networks. This is especially true for today's mobile networks which are comprised of a centralised architecture, mandating connectivity to central entities (which are located in the core of the mobile network) for communication. But also in fixed networks, access to a name resolution service is often necessary to access some given content.

- o Decentralised authentication: In mobile networks, users are authenticated via central entities. In order to communicate in fragmented or disconnected parts of a mobile network, the challenge of decentralising such user authentication arises. Independently of the network being fixed or mobile, data origin authentication of content retrieved from the network is challenging when being 'offline' (e.g. disconnected from servers of a security infrastructure such as a PKI).
- o Delivering/obtaining information and traffic prioritization in congested networks: Due to broken cables, failed routers, etc., it is likely that in a disaster scenario the communication network has much less overall capacity for handling traffic. Thus, significant congestion can be expected in parts of the infrastructure. It is therefore a challenge to guarantee message delivery in such a scenario. This is even more important as in the case of a disaster aftermath, it may be crucial to deliver certain information to recipients (e.g. warnings to citizens) with higher priority than other content.
- o Delay/Disruption Tolerant Approach: Fragmented networks makes it difficult to support end-to-end communication. However, communication in general and especially during disaster can tolerate some form of delay. E.g. in order to know if his/her relatives are safe or a 'SOS' call need not be supported in an end-to-end manner. It is sufficient to improve communication resilience in order to deliver such important messages.
- o Energy Efficiency: Long-lasting power outages may lead to batteries of communication devices running out, so designing energy-efficient solutions is very important in order to maintain a usable communication infrastructure.
- o Contextuality: Like any communication in general, disaster scenarios are inherently contextual. Aspects of geography, the people affected, the rescue communities involved, the languages being used and many other contextual aspects are highly relevant for an efficient realization of any rescue effort and, with it, therealization of the required communication.

The list above is most likely incomplete; future revisions of this document intend to add additional challenges to the list.

3.2. How ICN can be Beneficial

Several aspects of ICN make related approaches attractive candidates for addressing the challenges described in Section 3.1. Below is an

(incomplete) list of considerations why ICN approaches can be beneficial to address these challenges:

- o Routing-by-name: ICN protocols natively route by named data objects and can identify objects by names, effectively moving the process of name resolution from the application layer to the network layer. This functionality is very handy in a fragmented network where reference to location-based, fixed addresses may not work as a consequence of disruptions. For instance, name resolution with ICN does not necessarily rely on the reachability of application-layer servers (e.g. DNS resolvers). In highly decentralised scenarios (e.g. in infrastructureless, opportunistic environments) the ICN routing-by-name paradigm effectively may lead to a 'replication-by-name' approach, where content is replicated depending on its name.
- o Authentication of named data objects: ICN is built around the concept of named data objects. Several proposals exist for integrating the concept of 'self-certifying data' into a naming scheme (see e.g. [RFC6920]). With such approaches, the origin of data retrieved from the network can be authenticated without relying on a trusted third party or PKI.
- o Content-based access control: ICN can regulate access to data objects (e.g. only to a specific user or class of users) by means of content-based security; this functionality could facilitate trusted communications among peer users in isolated areas of the network.
- o Caching: Caching content along a delivery path is an inherent concept in ICN. Caching helps in handling huge amounts of traffic, and can help to avoid congestion in the network (e.g. congestion in backhaul links can be avoided by delivering content from caches at access nodes).
- o Sessionless: ICN does not require full end-to-end connectivity. This feature facilitates a seamless aggregation between a normal network and a fragmented network, which needs DTN-like message forwarding.
- o Potential to run traditional IP-based services (IP-over-ICN): While ICN and DTN promote the development of novel applications that fully utilize the new capabilities of the ICN/DTN network, work in [Trossen2015] has shown that an ICN-enabled network can transport IP-based services, either directly at IP or even at HTTP level. With this, IP- and ICN/DTN-based services can coexist, providing the necessary support of legacy applications to affected

users, while reaping any benefits from the native support for ICN in future applications.

- o Opportunities for traffic engineering and traffic prioritization: ICN provides the possibility to perform traffic engineering based on the name of desired content. This enables priority based replication depending on the scope of a given message [Psaras2014]. In addition, as [Trossen2015], among others, have pointed out, the realization ICN services and particularly of IP-based services on top of ICN provide further traffic engineering opportunities. The latter not only relate to the utilization of cached content, as outlined before, but to the ability to flexibly adapt to route changes (important in unreliable infrastructure such as in disaster scenarios), mobility support without anchor points (again, important when parts of the infrastructure are likely to fail) and the inherent support for multicast and multihoming delivery.

The list above is most likely incomplete; future revisions of this document intend to add more considerations to the list and to argue in more detail why ICN is suitable for addressing the aforementioned research challenges.

3.3. ICN as Starting Point vs. Existing DTN Solutions

There has been quite some work in the DTN (Delay Tolerant Networking) community on disaster communication (for instance, see further the discussions in the IETF DTN Research Group [dtnrg]). However, most DTN work lacks important features such as publish/subscribe (pub/sub) capabilities, caching, multicast delivery, and message prioritisation based on content types, which are needed in the disaster scenarios we consider. One could add such features to existing DTN protocols and solutions, and indeed individual proposals for adding such features to DTN protocols have been made (e.g. [Greifenberg2008] [Yoneki2007] propose the use of a pub/sub-based multicast distribution infrastructure for DTN-based opportunistic networking environments).

However, arguably ICN---having these intrinsic properties (as also outlined above)---makes a better starting point for building a communication architecture that works well before and after a disaster. For a disaster-enhanced ICN system this would imply the following advantages: a) ICN data mules would have built-in caches and can thus return content for interests straight on, b) requests do not necessarily be routed to a source (as with existing DTN protocols), instead any data mule or end-user can in principle respond to an interest, c) Built-in multi-cast delivery implies energy-efficient large-scale spreading of important information which is crucial in disaster scenarios, and d) pub/sub extension for

popular ICN implementations exist [COPSS2011] which are very suitable for efficient group communication in disasters and provide better reliability, timeliness and scalability as compared to existing pub/sub approaches in DTN [Greifenberg2008] [Yoneki2007].

Finally, most DTN routing algorithms have been solely designed for particular DTN scenarios. By extending ICN approaches for DTN-like scenarios, one ensures that a solution works in regular (i.e. well-connected) settings just as well (which can be important in reality, where a routing algorithm should work before and after a disaster). It is thus reasonable to start with existing ICN approaches and extend them with the necessary features needed in disaster scenarios.

4. Use Cases and Requirements

This Section describes some use cases for the aforementioned disaster scenario (as outlined in Section 2) and discusses the corresponding technical requirements for enabling these use cases.

- o Delivering Messages to Relatives/Friends: After a disaster strikes, citizens want to confirm to each other that they are safe. For instance, shortly after a large disaster (e.g., Earthquake, Tornado), people have moved to different refugee shelters. The mobile network is not fully recovered and is fragmented, but some base stations are functional. This use case imposes the following high-level requirements: a) People must be able to communicate with others in the same network fragment, b) people must be able to communicate with others that are located in different fragmented parts of the overall network. More concretely, the following requirements are needed to enable the use case: a) a mechanism for scalable message forwarding scheme that dynamically adapts to changing conditions in disconnected networks, b) DTN-like mechanisms for getting information from disconnected island to another disconnected island, c) data origin authentication so that users can confirm that the messages they receive are indeed from their relatives or friends, and d) the support for contextual caching in order to provide the right information to the right set of affected people in the most efficient manner.
- o Spreading Crucial Information to Citizens: State authorities want to be able to convey important information (e.g. warnings, or information on where to go or how to behave) to citizens. These kinds of information shall reach as many citizens as possible. i.e. Crucial content from legal authorities shall potentially reach all users in time. The technical requirements that can be derived from this use case are: a) Data origin authentication, such that citizens can confirm the authenticity of messages sent

by authorities, b) mechanisms that guarantee the timeliness and loss-free delivery of such information, which may include techniques for prioritizing certain messages in the network depending on who sent them, and c) DTN-like mechanisms for getting information from disconnected island to another disconnected island.

It can be observed that different key use cases for disaster scenarios imply overlapping and similar technical requirements for fulfilling them. As discussed in Section 3.2, ICN approaches are envisioned to be very suitable for addressing these requirements with actual technical solutions. In [Robitzsch2015], a more elaborate set of requirements is provided that addresses, among disaster scenarios, a communication infrastructure for communities facing several geographic, economic and political challenges.

5. Solution Design

This Section outlines some ICN-based approaches that aim at fulfilling the previously mentioned use cases and requirements.

- o ICN 'data mules': To facilitate the exchange of messages between different network fragments, mobile entities can act as ICN 'data mules' which are equipped with storage space and move around the disaster-stricken area gathering information to be disseminated. As the mules move around, they deliver messages to other individuals or points of attachment to different fragments of the network. These 'data mules' could have a pre-determined path (an ambulance going to and from a hospital), a fixed path (drone/robot assigned specifically to do so) or a completely random path (doctors moving from one camp to another).
- o Priority dependent Name-based replication: By allowing spatial and temporal scoping of named messages, priority based replication depending on the scope of a given message is possible. Clearly, spreading information in disaster cases involves space and time factors that have to be taken into account as messages spread. A concrete approach for such scope-based prioritisation of ICN messages in disasters, called 'NREP', has been proposed [Psaras2014], where ICN messages have attributes such as user-defined priority, space, and temporal-validity. These attributes are then taken into account when prioritizing messages. In [Psaras2014], evaluations show how this approach can be applied to the use case 'Delivering Messages to Relatives/Friends' described in Section 4.
- o Information Resilience through Decentralised Forwarding: In a dynamic or disruptive environment, such as the aftermath of a

disaster, both users and content servers may dynamically join and leave the network (due to mobility or network fragmentation). Thus, users might attach to the network and request content when the network is fragmented and the corresponding content origin is not reachable. In order to increase information resilience, content cached both in in-network caches and in end-user devices should be exploited. A concrete approach for the exploitation of content cached in user devices is presented in [Sourlas2015]. The proposal in [Sourlas2015] includes enhancements to the NDN router design, as well as an alternative Interest forwarding scheme which enables users to retrieve cached content when the network is fragmented and the content origin is not reachable. Evaluations show that this approach is a valid tool for the retrieval of cached content in disruptive cases and can be applied to tackle the challenges presented in Section 3.1.

- o Energy Efficiency: A large-scale disaster causes a large-scale blackout and thus a number of base stations (BSs) will be operated by their batteries. Capacities of such batteries are not large enough to provide cellular communication for several days after the disaster. In order to prolong the batteries' life from one day to several days, different techniques need to be explored: Priority control, cell-zooming, and collaborative upload. Cell zooming switches-off some of the BSs because switching-off is the only way to reduce power consumed at the idle time. In cell zooming, areas covered by such inactive BSs are covered by the active BSs. Collaborative communication is complementary to cell zooming and reduces power proportional to a load of a BS. The load represents cellular frequency resources. In collaborative communication, end-devices delegate sending and receiving messages to and from a base station to a representative end-device of which radio propagation quality is better. The design of an ICN-based publish/subscribe protocol that incorporates collaborative upload is ongoing work. In particular, the integration of collaborative upload techniques into the COPSS (Content Oriented Publish/Subscribe System) framework is envisioned [COPSS2011].
- o Data-centric confidentiality and access control: In ICN, the requested content is not anymore associated to a trusted server or an endpoint location, but it can be retrieved from any network cache or a replica server. This call for 'data-centric' security, where security relies on information exclusively contained in the message itself, or, if extra information provided by trusted entities is needed, this should be gathered through offline, asynchronous, and non interactive communication, rather than from an explicit online interactive handshake with trusted servers. The ability to guarantee security without any online entities is particularly important in disaster scenarios with fragmented

networks. One concrete cryptographic technique is 'Ciphertext-Policy Attribute Based Encryption' (CP-ABE), allowing a party to encrypt a content specifying a policy, which consists in a Boolean expression over attributes, that must be satisfied by those who want to decrypt such content. Such encryption schemes tie confidentiality and access-control to the transferred data, which can be transmitted also in an unsecured channel, enabling the source to specify the set of nodes allowed to decrypt.

- o Decentralised authentication of messages: Self-certifying names provide the property that any entity in a distributed system can verify the binding between a corresponding public key and the self-certifying name without relying on a trusted third party. Self-certifying names thus provide a decentralized form of data origin authentication. However, self-certifying names lack a binding with a corresponding real-world identity. Given the decentralised nature of a disaster scenario, a PKI-based approach for binding self-certifying names with real-world identities is not feasible. Instead, a Web-of-Trust can be used to provide this binding. Not only are the cryptographic signatures used within a Web-of-Trust independent of any central authority; there are also technical means for making the inherent trust relationships of a Web-of-Trust available to network entities in a decentralised, 'offline' fashion, such that information received can be assessed based on these trust relationships. A concrete scheme for such an approach has been published in [Seedorf2014], where also concrete examples for fulfilling the use case 'Delivering Messages to Relatives/Friends' with this approach are given.

6. The GreenICN Project

This section provides a brief overview of the GreenICN project. You can find more information at the project web site <http://www.greenicn.org/>

The recently formed GreenICN project, funded by the EU and Japan, aims to accelerate the practical deployment of ICN, addressing how ICN networks and devices can operate in a highly scalable and energy-efficient way. The project will exploit the designed infrastructure to support multiple applications including the following two broad exemplary scenarios: 1) The aftermath of a disaster, e.g. hurricane, earthquake, tsunami, or a human-generated network breakdown when energy and communication resources are at a premium and it is critical to efficiently distribute disaster notification and critical rescue information. Key to this is the ability to exploit fragmented networks with only intermittent connectivity, the potential exploitation of multiple modalities of communication and use of query/response and pub/sub approaches; 2) Scalable, efficient pub/sub

video delivery, a key requirement in both normal and disaster situations.

GreenICN will expose a functionality-rich API to spur the creation of new applications and services expected to drive industry and consumers, with special focus on the EU and Japanese environments, into ICN adoption. Our team, comprising researchers with diverse expertise, system and network equipment manufacturers, device vendors, a startup, and mobile telecommunications operators, is very well positioned to design, prototype and deploy GreenICN technology, and validate usability and performance of real-world GreenICN applications, contributing to create a new, low-energy, Information-Centric global communications infrastructure. We also plan to make contributions to standards bodies to further the adoption of ICN technologies.

7. Conclusion

This document outlines some research directions for Information Centric Networking (ICN) with respect to applying ICN approaches for coping with natural or human-generated, large-scale disasters. The document describes high-level research challenges as well as a general rationale why ICN approaches could be beneficial to address these challenges. One main objective of this document is to gather feedback from the ICN community within the IETF and IRTF regarding how ICN approaches can be suitable to solve the presented research challenges. Future revisions of this draft intend to include additional research challenges and to discuss what implications this research area has regarding related, future IETF standardisation.

8. References

8.1. Normative References

- [RFC5050] Scott, K. and S. Burleigh, "Bundle Protocol Specification", RFC 5050, DOI 10.17487/RFC5050, November 2007, <<http://www.rfc-editor.org/info/rfc5050>>.
- [RFC6920] Farrell, S., Kutscher, D., Dannewitz, C., Ohlman, B., Keranen, A., and P. Hallam-Baker, "Naming Things with Hashes", RFC 6920, DOI 10.17487/RFC6920, April 2013, <<http://www.rfc-editor.org/info/rfc6920>>.

8.2. Informative References

- [COPSS2011]
Chen, J., Arumaithurai, M., Jiao, L., Fu, X., and K. Ramakrishnan, "COPSS: An Efficient Content Oriented Publish/Subscribe System", Seventh ACM/IEEE Symposium on Architectures for Networking and Communications Systems (ANCS), 2011.
- [dtnrg] Fall, K. and J. Ott, "Delay-Tolerant Networking Research Group - DTNRG", <https://irtf.org/dtnrg>.
- [Greifenberg2008]
Greifenberg, J. and D. Kutscher, "Efficient publish/subscribe-based multicast for opportunistic networking with self-organized resource utilization", Advanced Information Networking and Applications-Workshops, 2008.
- [Psaras2014]
Psaras, I., Saino, L., Arumaithurai, M., Ramakrishnan, K., and G. Pavlou, "Name-Based Replication Priorities in Disaster Cases", 2nd Workshop on Name Oriented Mobility (NOM), 2014.
- [Robitzsch2015]
Robitzsch, S., Trossen, D., Theodorou, C., Barker, T., and A. Sathiaseel, "D2.1: Usage Scenarios and Requirements", H2020 project RIFE, public deliverable, 2015.
- [Seedorf2014]
Seedorf, J., Kutscher, D., and F. Schneider, "Decentralised Binding of Self-Certifying Names to Real-World Identities for Assessment of Third-Party Messages in Fragmented Mobile Networks", 2nd Workshop on Name Oriented Mobility (NOM), 2014.
- [Sourlas2015]
Sourlas, V., Tassiulas, L., Psaras, I., and G. Pavlou, "Information Resilience through User-Assisted Caching in Disruptive Content-Centric Networks", 14th IFIP NETWORKING, May 2015.
- [Trossen2015]
Trossen, D., "IP-over-ICN", To appear in EUCNC 2015.

[Yoneki2007]

Yoneki, E., Hui, P., Chan, S., and J. Crowcroft, "A socio-aware overlay for publish/subscribe communication in delay tolerant networks", Proceedings of the 10th ACM Symposium on Modeling, Analysis, and Simulation of Wireless and Mobile Systems, 2007.

Appendix A. Acknowledgment

The authors would like to thank Ioannis Psaras for useful comments. Further, the authors would like to thank Joerg Ott and Dirk Trossen for valuable comments and input, in particular regarding existing work from the DTN community which is highly related to the ICN approaches suggested in this document.

This document has been supported by the GreenICN project (GreenICN: Architecture and Applications of Green Information Centric Networking), a research project supported jointly by the European Commission under its 7th Framework Program (contract no. 608518) and the National Institute of Information and Communications Technology (NICT) in Japan (contract no. 167). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the GreenICN project, the European Commission, or NICT.

Authors' Addresses

Jan Seedorf
NEC
Kurfuerstenanlage 36
Heidelberg 69115
Germany

Phone: +49 6221 4342 221
Fax: +49 6221 4342 155
Email: seedorf@neclab.eu

Mayutan Arumaithurai
University of Goettingen
Goldschmidt Str. 7
Goettingen 37077
Germany

Phone: +49 551 39 172046
Fax: +49 551 39 14416
Email: arumaithurai@informatik.uni-goettingen.de

Atsushi Tagami
KDDI R&D Labs
2-1-15 Ohara
Fujimino, Saitama 356-85025
Japan

Phone: +81 49 278 73651
Fax: +81 49 278 7510
Email: tagami@kddilabs.jp

K. K. Ramakrishnan
University of California
Riverside CA
USA

Email: kkramakrishnan@yahoo.com

Nicola Blefari Melazzi
University Tor Vergata
Via del Politecnico, 1
Roma 00133
Italy

Phone: +39 06 7259 7501
Fax: +39 06 7259 7435
Email: blefari@uniroma2.it