ICNRG
Internet-Draft
Intended Status: Informational
Expires: January 16, 2014

K. Pentikousis, Ed.
Huawei Technologies
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
P. Mahadevan
PARC
S. Spirou
Intracom Telecom
A. Molinaro
UNIRC
D. Gellert
InterDigital
S. Eum
NICT

July 15, 2013

ICN Baseline Scenarios and Evaluation Methodology
draft-pentikousis-icn-scenarios-04

Abstract

This document aims at establishing a common understanding about 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 efficacy, 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. We then survey the evaluation tools currently available to researchers in this area and provide suggestions regarding methodology and metrics. Finally, this document sheds some light on the impact of ICN on network security.

Status of this Memo

Copyright and License Notice

Table of Contents

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. Anycast, multicast, and broadcast are natively supported, and energy efficiency is a design consideration from the very beginning.

Although interest in ICN is growing rapidly, ongoing work on different architectures, such as, for example, NetInf [NetInf], CCN and NDN [CCN], the publish-subscribe Internet (PSI) architecture [PSI], and the data-oriented architecture [DONA] is far from being completed. 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 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.

Ahlgren et al. [SoA] note that describing ICN architectures is akin to shooting a moving target. We find that comparing these different approaches is often even more tricky. in particular, we observe that a variety of performance evaluation scenarios has been devised, typically with good reason, in order to highlight the advantages of each ICN architecture. That is, there is no single scenario, use case, or reference topology which is employed as a benchmark consistently across the ICN literature. 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 superiority over current systems, in general, and against each other, in particular.

The remainder of this document is organized as follows. In Section 2 we review the peer-reviewed ICN literature and select 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. The list of scenarios has evolved since the first draft version of this document based on the input from the research group and the corresponding text contributions.

Section 3 presents currently available simulation tools and experimental testbeds that can be used in evaluating ICN, and outlines the key elements that should be considered in an ICN evaluation. Finally, Section 4 discusses the impact of ICN on network security.

2. Toward ICN Baseline Scenarios

This section presents a number of scenarios grouped into several categories. Note that certain evaluation scenarios span across these categories, so the boundaries between them should not be considered rigid and inflexible. There are two goals for this section. 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.

The overall aim is that each scenario is described at a sufficient level of detail so that it can serve as the base for comparative evaluations of different approaches. This will need to include reference configurations, topologies, specifications of traffic mixes and traffic loads. These specifications (or configurations) should preferably come as sets that describe extremes as well as "typical" usage scenarios.

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 and, therefore, we would expect that they are a "natural fit" for showcasing the superiority of ICN over 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 a centralized content server; see also subsection 2.4, below, and [JBDMM+12]. 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 verification does not require any centralized authentication server. Kim and Lee [VPC2] present a proof-of-concept evaluation in which users with ordinary smartphones can browse a list of members or content using a name, and download the content selected from the list.

In short, in both ICN-based social networking application scenarios there is no need for a classic client-server architecture (let alone a cloud-based infrastructure) to intermediate between content providers and consumers in a hub-and-spoke fashion.

Earlier work by Arianfar et al. [ANO10] considers a similar pull-based content retrieval scenario using a different architecture, pointing to significant performance advantages. Although the authors consider a network topology (redrawn in Fig. 1 for convenience) that has certain interesting characteristics, they do not explicitly address social networking in their evaluation scenario. Nonetheless, similarities are easy to spot: "followers" (such as C0, C1, ..., and Cz in Fig. 1) obtain content put "on the network" (I1, ..., Im, and B1, B2) by a single user (e.g. Px) relying solely on network primitives.

```
\--/
|C0|
/--\     +--+     +--+     +--+                  +--+
    *=== |I0| === |I1| ... |In|                  |P0|
\--/     +--+     +--+     +--+                  +--+
|C1|                           \             / o
/--\                            +--+     +--+  o
 o                              |B1| === |B2|  o
 o             o o o o          +--+     +--+  o
 o                             /             \ o
 o       +--+     +--+     +--+                  +--+
 o  *=== |Ik| === |Il| ... |Im|                  |Px|
\--/     +--+     +--+     +--+                  +--+
|Cz|
/--\
```

Figure 1. Dumbbell with linear daisy chains.

In summary, the social networking scenario aims to exercise each ICN

architecture in terms of network efficiency, multicast support, caching performance and its reliance on centralized mechanisms (or lack thereof).

2.2. Real-time Communication

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/UTP. 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, today 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 the 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 more work is recommended in this direction.

In short, 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. 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. For example, one could exploit the ability of nodes to better express their intended use of the network, i.e., the retrieval of a small subset of the global data corpus as discussed in [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 subsection 2.6 below.

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 which ever server (I2 or I3) can supply them without necessarily needing to access the servers in the core network "C" (P1 and P2). Lindgren [Lin11] 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.

```
+------------+      +-----------+
| Network N0 |      | Network C |
|            |      |           |
| +--+       | ==== |    +--+   |
| |I2|       |      |    |P1|   |
| +--+       |      |    +--+   |
|     \--/   |      |           |
+-----|C0|---+      |           |
|     /--\   |      |           |
| +--+       |      |           |
| |I3|       |      |      +--+ |
| +--+       | ==== |      |P2| |
|            |      |      +--+ |
| Network N1 |      |           |
+------------+      +-----------+
```

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 [AMR13]. 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 [MPZ10] and [OLG10] for ICN-based mobile ad-hoc networks (MANETs), and in [WAKVWZ12] and [ACM12] 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 subsection 2.9 and the economic consequences of nodes having multiple network interfaces is explored in subsection 2.7.

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 (perhaps) much faster than Metcalfe's law.

For example, Jacobson et al. [JBDMM+12] 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. [MCG11], 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. [CHPP12] 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 [L9] and BitTorrent-like traffic [TKMEMT12]. 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 (see also subsection 2.7).

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 to illustrate the superiority of ICN over other approaches. An obvious baseline to compare against in this regard is existing federations of IP-based Content Distribution Networks (CDNs).

2.5. Content Dissemination

Content dissemination has attracted more attention than other aspects of ICN, perhaps due to a misunderstanding of what the first "C" in CCN stands for. Scenarios in this category abound in the literature, including stored and streaming A/V distribution, file distribution, mirroring and bulk transfers, versioned content services (c.f., 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 subsection. 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 earlier, content dissemination scenarios in VANETs have a particular affinity to the mobility scenarios discussed earlier.

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] [JBDMM+12] [ANO10], 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. [FMP12] 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 [BK10] and [DMND], and additionally investigated in [WWKVZ12] [WAKVWZ12] [AKH11] [TL12] [ACM12] [CRoWN]. For example, Bai and Krishnamachari [BK10] 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.

```
+ - - _- - -_- - - -_- - _- - - +
|    /_\   /_\     /_\  /_\     |
|    o o   o o     o o  o o     |
|    +-------+     +-------+ _  |
|    |       |     |       |/_\ |
|  _ |       |     |       |o o |
| /_\|       |    |       |     |
| o o+--_----+\===/+--_----+     |
|       /_\    |RSU|  /_\        |
|       o o    /===\  o o        |
|    +-------+     +-------+ _  |
|    |       |     |       |/_\ |
| _  |       |     |       |o o |
|/_\ |       |     |       |     |
|o o +_-----_+     +_-----_+     |
|    /_\   /_\     /_\   /_\    |
+_ _ o_o_ _o_o_ _ _o_o_ _o_o_ _ +
```

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. [WWKVZ12], for instance, analyze the advantages of hierarchical naming for vehicular traffic information dissemination. Arnould et al. [AKH11] apply ICN principles to safety information dissemination between vehicles with multiple radio interfaces. In [TL12], TalebiFard and Leung use network coding techniques to improve content dissemination over multiple ICN paths. Amadeo et al. [ACM12][[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 [WAKVWZ12] or from RSUs [ACM12]. 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].

```
  \__/                    \__/
  |RU|                    |RU|
================================
     _     _     _     _
    /_\   /_\   /_\   /_\
_ _ o_o_ _o_o_ _o_o_ _o_o_ _ _ _
     _     _     _     _
    /_\   /_\   /_\   /_\
    o o   o o   o o   o o
================================
```

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. 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 networks. 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). 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. [OptSC] [FEMTO]. 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" demonstrator created by the SAIL project investigated this kind of scenario; more details are available in [SAIL-B3].

Jacobson et al. [CCN] include interactions between networks in their

overall system design, and mention both "an edge-driven, bottom-up incentive structure" and techniques based on evolutions of existing mechanisms both for ICN router discovery by the end-user and for interconnecting between autonomous systems (AS). For example, a BGP extension for domain-level content prefix advertisement can be used to enable efficient interconnection between AS's. Liu et al. [MLDHT] proposed to address the "suffix-hole" issue found in prefix-based name aggregation through the use of a combination of Bloom-filter based aggregation and multi-level DHT.

Name aggregation has been discussed for a flat naming design as well in [NCOA], which also notes that based on estimations in [DONA] flat naming may not require aggregation. This is a point that calls for further study. Scenarios evaluating name aggregation, or lack thereof, should take into account the amount of state (e.g. size of routing tables) maintained in edge routers as well as network efficiency (e.g. amount of traffic generated).

```
              +---------------+
  +---------->| Popular Video |
  |           +---------------+
  |             ^           ^
  |             |           |
  |           +-+-+ $0/MB +-+-+
  |           | A +-------+ B |
  |           ++--+       +-+-+
  |            | ^         ^ |
  |      $8/MB | |         | | $10/MB
  |            v |         | v
+-+-+  $0/MB  +--+---------+--+
| D +---------+       C       |
+---+         +---------------+
```

Figure 5. Relationships and transit costs between networks A to D.

DiBenedetto et al. [RP-NDN] study policy knobs made available by NDN to network operators. New policies, which are not feasible in the current Internet are described, including a "cache sharing peers" policy, where two peers have an incentive to share content cached in, but not originating from, their respective network. The simple example used in the investigation considers several networks and associated transit costs, as shown in Fig. 5. (based on Fig. 1 of [RP-NDN]). Agyapong and Sirbu [EconICN] further establish that ICN approaches should incorporate features that foster (new) business 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.

Ahlgren 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.

```
+----------+     +-----------------+     +------+
| Content  |     | Access Network/ |     | End  |
| Provider +---->|  ICN Provider   +---->| User |
+----------+     +-+-------------+-+     +------+
                   |             |
                   |             |
                   v             v
+-------------------+     +----------------+       +------+
| Interconnectivity |     | Access Network |       | End  |
|     Provider      +---->|     Provider   +------>| User |
+-------------------+     +----------------+       +------+
```

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 Internet, a traffic model taking into account business factors for each AS, and a routing model integrating the valley-free model and policy-compliance. A typical scenario topology is illustrated in Fig. 7, which is redrawn here based on Fig. 1 of [ICCP]. Note that it relates well with the topology illustrated in Fig. 1 of this

document.

```
                        o-----o
                  +-----+  J  +-----+
                  |     o--*--o     |
                  |        *        |
               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
                   |                 |         ^^  | route
             data  |           data  |    data ||  | to
                   |                 |         ||  | data
               o---v--o          o---v--o     o++--v-o
               | User |          | User |     | Data |
               o------o          o------o     o------o
```

```
Legend:
*****  Transit link
+---+  Peering link
+--->  Data delivery or route to data
```

Figure 7. Tier-based set of interconnections between AS A to J.

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.

2.8. Energy Efficiency

Advantage can be taken of some prominent ICN features to significantly reduce the energy footprint of future communication networks. A simple example of a potential energy-saving operation is caching. If a data object can be retrieved from within a network, rather than from a distant origin server, clearly, significant amounts of energy expenditure can be saved by avoiding several further hops. Alternatively, approaches that aim to simplify routers, such as [PURSUIT], could also reduce energy consumption by pushing routing decisions into more energy-efficient data centers.

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. 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 end-to-end always-on connections, introduces built-in energy inefficiencies argueing 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 ICN can be more energy efficient in delivering popular 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, as mentioned earlier, energy efficiency can be tackled by different ICN approaches in ways that it cannot in a host-centric paradigm. We already mentioned that in ICN, perpetual (always-on) connectivity is not necessary, therefore mechanisms that capitalize on powering down network interfaces are easier to accommodate. Since all 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. 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 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.

2.9. Delay- and Disruption-Tolerance

Delay- and Disruption-Tolerant Networking (DTN) [DTN] [TBB13] originated as a means to extend the Internet to interplanetary communications. 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 [DTN]. 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.

Third, 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 publications 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 (or 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 added to 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 recently developed a prototype implementation of NetInf running over the DTN Bundle Protocol.

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). These are discussed 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 the case of both 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 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. a period 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 ICN, the quality of service expectations are 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 that 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:

- o Satisfaction Probability: The probability by which an information request (e.g., Interest) will be satisfied (i.e., how often a Data response will be received).

- o Satisfaction Delay: 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).

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, this might occur on an underground train, in which people 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 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. Luckily, there exist a number of datasets that offer exactly this information required for microblogging.

In terms of the environment (i.e., mobility patterns), the Haggle 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. Sources include:

http://www.tweetarchivist.com/
http://twapperkeeper.com/
http://www.infochimps.com/collections/twitter-census
http://socialcomputing.asu.edu/datasets/Twitter

These datasets can therefore be used to extract information production, placement and consumption.

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 India flooding, as well as the 2011 Tohoku earthquake and tsunami. Power problems often exacerbate the issue, with communication problems 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. Here, we focus on the dissemination of standard broadcast information that should be received by all parties.

For the environmental setup, there are no commonly used mobility traces for disaster zones, unlike in the previous evaluative scenario. This is clearly due to the difficultly (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). [MODEL2] focuses on emergency responders, featuring the removal of nodes from the disaster zone, as well as things like obstacles (e.g. collapsed buildings). [MODEL3] also looks at emergency responders, but focuses 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 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). 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.

2.10. 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. [NDNl] 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 is 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.

2.11. 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 subsection, 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 [WAKVWZ12]. 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 test-bed 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).

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.

2.12. 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 subection 2.9 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 (see subsection 4.2) 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.

2.13. Summary

We conclude Section 2 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. subsections 2.1-2, 2.5-7, 2.10-11,) as does network, resource and energy efficiency (cf. subsections 2.1, 2.3-4, 2.7-8). Operational aspects such as network planing, manageability, reduced complexity and overhead (cf. subsection 2.2-4, 2.7, 2.10) 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. subsections 2.1, 2.3, 2.5-6). ICN must also be able to cope when the Internet expands to incorporate additional network paradigms (cf. subsection 2.12). 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. subsections 2.2, 2.9-10).

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. subsections 2.2-3, 2.6, 2.10-11).

3. Evaluation Methodology

As we have seen in the previous section, 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 evaluations should also

include an assessment of 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.

In this section, we present various techniques and considerations for evaluating different ICN architectures. 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 ICNRG and 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, and identify the respective strengths and weaknesses.

3.1. 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.

3.1.1. 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 [CGB12] 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 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.

3.1.2. PSI

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 [PTA13].

Moreover, the ICN simulation environment [VBYR12] 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.

3.1.3. 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 code for this can be downloaded from http://gin.ngnet.it.

3.1.4. COMET

The EU FP7 COMET project developed a simulator, called Icarus, which implements ProbCache [PCP12], centrality-based in-network caching [CHPP12] and the hash-route-based algorithms detailed in [SPP13]. The simulator is built in Python and makes use of the Fast Network Simulator Setup tool [SCP13] to configure the related parameters of the simulation. The simulator is available from: https://github.com/lorenzosaino/icarus/

3.1.5. 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.

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.

3.2. Topology Selection

Section 2 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 aspect in such an evaluation. For instance, CAIDA's ITDK traces record millions of routers across thousands of domains.

Beyond these "real-world" traces there is a wide range of synthetic topologies, such as the Barabasi-Albert model [BA99] and the Watts-Strogatz small-world topology [WS98]. 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 lines of ICN 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 [KPLKTS09] 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 subsection 2.9, for example, contact traces from the DTN community could also be used in ICN evaluations.

3.3. 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 such as 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.

Several previous studies have stated that Zipf's law is the discrete distribution that best represents the request frequency in a number of application scenarios, ranging from typical Web access 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 <= i <= 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, with C = SUM(1 / j^(alpha)), alpha > 0

where the sum is obtained considering all values of j, 1 <= j <= M.

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

Further, a variation of the Zipf distribution, termed the Mandelbrot-Zipf distribution, has been suggested by [SH06] 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.

3.4. 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 a fundamental change in data access and transport mechanisms 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 that 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 a different ICN approach, and to the incumbent host-centric approach. It will therefore be less difficult to gauge optimization and research direction. On the other hand, metrics should be targeted to network performance primarily 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 considered in previously published evaluations for several ICN approaches, such as CCN [CCN] [VoCCN] [NDNProj]; NetInf [4WARD6.1] [4WARD6.3] [SAIL-B2] [SAIL-B3]; PURSUIT [PRST4.5], COMET [CMT-D5.2] [CMT-D6.2]; Connect [MCG11] [RealCCN]; and CONVERGENCE [ICN-Web] [ICN-Scal] [ICN-Tran]. The metrics used in these studies 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 sometimes ICN 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 provided in this document for each approach.

Traffic metrics in existing ICN approaches are summarized in Table II. These metrics capture mainly 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 do not 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 evaluation studies

	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. Although different ICN approaches have used system metrics in their respective evaluation studies, 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 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. 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 between different studies, 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 current ICN evaluations

	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 would 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.

3.4.1. 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 quantify the data transport performance of the network in terms of well-established 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 on particular applications.

Another option for measuring transport performance would be to use QoS 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 networks of different architecture (host-centric, information-centric, or other). As implied earlier, the drawback of this 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 at this stage.

3.4.2. 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 subsection 2.10, 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. [KVHHM11] 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 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. For gateway nodes, which are typically a point of service to a large number of nodes and have to satisfy the information requests from remote entities, we need to consider scalability-related metrics, such as frequency of requests received and processing of successfully satisfied information requests.

Finally, given the in-network caching functionality in ICN, metrics for the efficiency and performance of in-network caching have to be defined, which can quantify 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 data object is stored in a network cache. The model takes as input the rate of requests for a given content, named the Content of Interest (CoI), and the rate of requests for all other data objects 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 lambda is the request rate for CoI, mu 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.

3.5. Resource Equivalence and Tradeoffs

As we have seen above, every information-centric 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 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.

3.6. Technology Evolution Assumptions

TBD

4. Security Considerations

The introduction of an information-centric networking architecture and the corresponding communication paradigm changes many aspects of network security. Additional evaluation will be required to ensure relevant security requirements are appropriately met by the implementation of the chosen architecture in the scenarios presented in Section 2.

The various ICN architectures that are currently proposed have concentrated on authentication of delivered content to ensure content integrity. 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. In general, 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 [RAA09]. In the rest of this section we briefly consider the issues at hand and provide pointers to the work that has been done on the security aspects of the architectures proposed.

4.1. Authentication

For fully secure content distribution, content access requires that the receiver needs to be able to reliably assess validity, provenance, and relevance. Validity relates to whether the received data object is a complete, uncorrupted copy of what was originally published. Provenance means that the receiver can identify the publisher, and, if so, whether it and the source of any cached version of the document can be adequately trusted. Relevance ascertains that the content obtained answers the question that the receiver asked.

All 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 as 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, for instance, 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. 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] [SJ09] 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.

PSI allows the authentication mechanism to be chosen so that it can, in theory, adopt the authentication strategy of any other other architectures [PRST2.4]. In the light of such choices, however, interoperability (if required by the chosen deployment) needs to be taken care of through dedicated solutions.

4.2. Authorization, Access Control and Statistics

A potentially major concern for all 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 in [FMP12] 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.

4.3. 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 [Lau10]. The analysis in this master's thesis is to a large degree applicable to all of ICN architectures because it is mostly focused on the common caching aspect. The privacy risks of named data networking are also highlighted in [LLRSBK12]. Further work on privacy in ICNs can be found in [CDKU12].

4.4. 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 changed threat model is also discussed in [Lau10] and [CDKU12]. Some of the key aspects are:

- Caching implies a tradeoff between network efficiency and user privacy as discussed in subsection 4.3.

- 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 ICN makes it difficult to identify the origin of a request as mentioned in subsection 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 in 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. Information-centric networks 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.

5. IANA Considerations

This document presents no IANA considerations.

6. Acknowledgments

This document has benefited from 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, Will Liu, Ioannis Psaras, Dirk Trossen, Jianping Wang, Yuanzhe Xuan, and Xinwen Zhang.

7. Informative References

[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.

[RFC6920] Farrell, S., Kutscher, D., Dannewitz, C., Ohlman, B., Keranen, A., and P. Hallam-Baker, "Naming Things with Hashes", RFC 6920, April 2013.

[RFC2330] Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics", RFC 2330, May 1998.

[RFC5050] Scott, K. and S. Burleigh, "Bundle Protocol Specification", RFC 5050, November 2007.

[ndnSIM] Afanasyev, A., Moiseenko, I., and L. Zhang, "ndnSIM: NDN simulator for NS-3", NDN Technical Report NDN-0005, Revision 2, October 2012.

[NetInf] Ahlgren, B., D'Ambrosio, M., Marchisio, M., Marsh, I., Dannewitz, D., Ohlman, B., Pentikousis, K., Strandberg, O., Rembarz, R., and V. Vercellone, "Design considerations for a network of information", Proc. CoNEXT Re-Arch Workshop, p. 66. ACM, 2008.

[ACM12] Amadeo, M., Campolo, C. and A. Molinaro, "Content-centric networking: is that a solution for upcoming vehicular networks?." In Int. workshop on Vehicular inter-networking, systems, and applications, pp. 99-102. ACM, 2012.

[CRoWN] Amadeo, M., Campolo, C. and A. Molinaro, "CRoWN: Content-Centric Networking in Vehicular Ad Hoc Networks", Communications Letters, IEEE, vol. 16, no. 9 (2012): 1380-1383.

[AMR13] Amadeo, M., Molinaro, A., and G. Ruggeri, "E-CHANET: Routing, forwarding and transport in Information-Centric

multihop wireless networks", Computer Communications (2013).

[ANO10] Arianfar, S., Nikander, P., and J. Ott, "On content-centric router design and implications", Proc. Re-Architecting the Internet Workshop, p. 5. ACM, 2010.

[AKH11] Arnould, G., Khadraoui D., and Z. Habbas, "A self-organizing content centric network model for hybrid vehicular ad-hoc networks." Proc. ACM Int. Symp. Design and analysis of intelligent vehicular networks and applications, pp. 15-22. ACM, 2011.

[BK10] Bai, F. and B. Krishnamachari, "Exploiting the wisdom of the crowd: localized, distributed information-centric VANETs", Communications Magazine, IEEE, vol. 48, no. 5 (2010): 138-146.

[BA99] Barabasi, A. and R. Albert, "Emergence of scaling in random networks", Science, vol. 286, no. 5439, pp. 509-512, 1999.

[CDKU12] Chaabane, A., De Cristofaro, E., Kaafar, M. A., and E. Uzun, "Privacy in Content-Oriented Networking: Threats and Countermeasures", arXiv:1211.5183, 2012.

[CURLING] Chai, W. K., Wang, N., Psaras, I., and G. Pavlou, "CURLING: Content-Ubiquitous Resolution and Delivery Infrastructure for Next-Generation Services", Communications Magazine, IEEE, vol. 49, no. 3 (2011): 112-120.

[CHPP12] Chai, W. K., He, D., Psaras, I., and G. Pavlou, "Cache 'less for more' in information-centric networks," Proc. NETWORKING 2012, pp. 27-40. Springer, 2012.

[CGB12] Cianci, I. Grieco, L. A., and G. Boggia, "CCN - Java Opensource Kit EmulatoR for Wireless Ad Hoc Networks", Proc. 7th Int. Conf. on Future Internet Technologies, pp. 7-12. ACM, 2012.

[FMP12] Fotiou, N., Marias, G. F., and G. C. Polyzos, "Access control enforcement delegation for information-centric networking architectures", Proc. ACM SIGCOMM Workshop on Information-centric networking, pp. 85-90. ACM, 2012.

[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.

[CCN] Jacobson, V., Smetters, D. K., Thornton, J. D., Plass, M. F., Briggs, N. H., and R. L. Braynard, "Networking Named Content", Proc. CoNEXT, pp. 1-12. ACM, 2009.

[VoCCN] Jacobson, V., Smetters, D. K., Briggs, N. H., Plass, M. F., Stewart, P., Thornton, J. D., and R. L. Braynard, "VoCCN: Voice-over Content-Centric Networks", Proc. CoNEXT Re-Arch Workshop, pp. 1-6. ACM, 2009.

[JBDMM+12] Jacobson, V., Braynard, R. L., Diebert, T., Mahadevan, P., Mosko, M., Briggs, N. H., Barber, S., et al., "Custodian-based information sharing", Communications Magazine, IEEE, vol. 50, no. 7 (2012): 38-43.

[KVHHM11] Karnouskos, S., Villasenor-Herrera, V., Haroon, M., Handte, M. and P. J. Marron, "Requirement considerations for ubiquitous integration of cooperating objects", Proc. IFIP NTMS, pp. 1-5. IEEE, 2011.

[KPLKTS09] Kaune, S., Pussep, K., Leng, C., Kovacevic, A., Tyson, G., and R. Steinmetz, "Modelling the internet delay space based on geographical locations", Proc. Euromicro, pp. 301-310. IEEE, 2009.

[SLINKY] Kawadia, V., Riga, N., Opper, J., and D. Sampath, "Slinky: An adaptive protocol for content access in disruption-tolerant ad hoc networks", Proc. MobiHoc Workshop on Tactical Mobile Ad Hoc Networking, 2011.

[DONA] Koponen, T., Chawla, M., Chun, B. G., Ermolinskiy, A., Kim, K. H., Shenker, S., and I. Stoica, "A Data-Oriented (and Beyond) Network Architecture", ACM SIGCOMM Computer Communication Review, vol. 37, no. 4, pp. 181-192. ACM, 2007.

[Lau10] Lauinger, T., "Security and Scalability of Content-Centric Networking", Masters Thesis, Technische Universitaet Darmstadt and Eurecom, Sep. 2010.

[LLRSBK12] Lauinger, Y., Laoutaris, N., Rodriguez, P., Strufe, T., Biersack, E., and E. Kirda, "Privacy Risks in Named Data Networking: What is the Cost of Performance?", ACM SIGCOMM Computer Communication Review 42, no. 5 (2012): 54-57.

[Lin11] Lindgren, A., "Efficient content distribution in an

information-centric hybrid mobile network", Proc. Consumer Communications and Networking Conference (CCNC), 2011 IEEE, pp. 1134-1138. IEEE, 2011.

[MPZ10] Meisel, M., Pappas, V., and L. Zhang, "Ad hoc networking via named data", Proc. Int. Workshop on Mobility in the Evolving Internet Architecture, pp. 3-8. ACM, 2010.

[MCG11] Muscariello, L., Carofiglio, G., and M. Gallo, "Bandwidth and storage sharing performance in information centric networking", Proc. ACM SIGCOMM Workshop on Information-centric networking, pp. 26-31. ACM, 2011.

[CCNPL] Muscariello, L. and M. Gallo, "Content centric networking packet level simulator", available online at http://code.google.com/p/ccnpl-sim/

[OLG10] Oh, S.-Y., Lau, D., and M. Gerla. "Content centric networking in tactical and emergency manets", Proc. Wireless Days (WD), 2010 IFIP, pp. 1-5. IEEE, 2010.

[PTA13] Parisis, G., Trossen, D., 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.

[PCP12] Psaras, I., Chai, W., and G. Pavlou, "Probabilistic In-Network Caching for Information-Centric Networks", Proc. ACM SIGCOMM Workshop on Information-centric networking, pp. 55-60. ACM, 2012.

[RAA09] Renault, E., Ahmad, A., and M. Abid, "Towards a Security Model for the Future Network of Information", Proc. Int. Conf. Ubiquitous Information Technologies and Applications, pp. 1-6. IEEE, 2009.

[ccnSim] Rossini, G. and D. Rossi, "Large scale simulation of CCN networks", Proc. Algotel 2012 , La Grande Motte, France, May 2012.

[SCP13] Saino, L., Cocora, C., and G. Pavlou, "A Toolchain for Simplifying Network Simulation Setup", Proc. SIMUTOOLS. ACM, 2013.

[SPP13] Saino, L., Psaras, I. , and G. Pavlou, "Hash-routing Schemes for Information-Centric Networking", Proc. ACM SIGCOMM Workshop on Information-centric networking. ACM, 2013.

[SH06] Saleh, O., and M. Hefeeda, "Modeling and caching of peer-to-peer traffic", Proc. ICNP, pp. 249-258. IEEE, 2006.

[SJ09] Smetters, D., and V. Jacobson, "Securing network content", Technical Report TR-2009-01, PARC, 2009.

[TL12] TalebiFard, P. and V. Leung, "A content centric approach to dissemination of information in vehicular networks", Proc. ACM Int. Symp. Design and analysis of intelligent vehicular networks and applications, pp. 17-24. ACM, 2012.

[PSI] Trossen, D. and G. Parisis, "Designing and realizing an information-centric internet", Communications Magazine, IEEE, vol. 50, no. 7 (2012): 60-67.

[TKMEMT12] Tyson, G., Kaune, S., Miles, S., El-khatib, Y., Mauthe, A., and A. Taweel, "A trace-driven analysis of caching in content-centric networks", Proc. Computer Communications and Networks (ICCCN), 2012 21st International Conference on, pp. 1-7. IEEE, 2012.

[TBB13] Tyson, G., Bigham J., and E. Bodanese, "Towards an Information-Centric Delay-Tolerant Network", Proc. IEEE INFOCOM Workshop on Emerging Design Choices in Name-Oriented Networking (NOMEN). 2013.

[VBYR12] Vastardis, N., Bontozoglou, A., Yang, K., and M. Reed, "Simulation tools enabling research on Information-centric Networks", Proc. IEEE ICC, pp. 5833-5838. IEEE, 2012.

[DMND] Wang, J., Wakikawa, R., and L. Zhang, "DMND: Collecting data from mobiles using named data", Proc. Vehicular Networking Conference (VNC), 2010 IEEE, pp. 49-56. IEEE, 2010.

[WAKVWZ12] Wang, L., Afanasyev, A., Kuntz, R., Vuyyuru, R., Wakikawa, R., and L. Zhang, "Rapid traffic information dissemination using named data", In Workshop on Emerging Name-Oriented Mobile Networking Design-Architecture, Algorithms, and Applications, pp. 7-12. ACM, 2012.

[WWKVZ12] Wang, L., Wakikawa, R., Kuntz, R., Vuyyuru, R., and L. Zhang, "Data naming in vehicle-to-vehicle communications", Proc. Computer Communications Workshops (INFOCOM WKSHPS), 2012 IEEE Conference on, pp. 328-333. IEEE, 2012.

[WS98] Watts, D. J. and S. H. Strogatz, "Collective dynamics of small-world networks", Nature, vol. 393, no. 6684, pp. 40-

"10, 1998.

[NDNProj] Zhang, L., Estrin, D., Burke, J., Jacobson, V., Thornton, J. D., Smetters, D. K., Zhang, B., et al. , "Named Data Networking (NDN) Project", NDN Technical Report NDN-0001, Oct. 2010.

[SoA] Ahlgren, B. et al., "A survey of information-centric networking", IEEE Commun. Mag., vol. 50, no. 7, July 2012.

[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.

[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.

[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.

[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.

[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.

[OptSC] Paolini, M. Optimizing the Small-Cell Business ROI, SmallCell Americas 2012, available online at www.smallcellsamericas.com/files/ monica_paolini_senza_fili_consulting.pdf

[FEMTO] Rioridan, R. Using FemtoCloud technology to deliver femtocell-as-a-service, SmallCell Americas 2012, available online at www.smallcellsamericas.com/files/ 1110rob_riordan_femto_america_2012.pdf

[MLDHT] Liu H. et al., "A multi-level DHT routing framework with aggregation", Proc. SIGCOMM ICN Workshop. ACM, 2012.

[RP-NDN] DiBenedetto S. et al., "Routing policies in named data networking", Proc. SIGCOMM ICN Workshop. ACM, 2011.

[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.

[SAIL-B2] SAIL, "NetInf Content Delivery and Operations", SAIL Project Deliverable D-B.2, May. 2012.

[SAIL-B3] SAIL, "Final NetInf Architecture", SAIL Project Deliverable D-B.3 , Jan. 2013.

[SAIL-A7] SAIL, "New business models and business dynamics of the future networks", SAIL Project Deliverable D-A.7, Aug. 2011.

[SAIL-A8] SAIL, "Evaluation of business models", SAIL Project Deliverable D-A.8, Jan. 2013.

[EECCN] Guan, K. et al., "On the Energy Efficiency of Content Delivery Architectures", Proc. ICC Workshops. IEEE, 2011.

[DTN] Fall, K., "A delay-tolerant network architecture for challenged internets", Proc. SIGCOMM. ACM, 2003.

[BPQ] S. Farrell, A. Lynch, D. Kutscher, and A. Lindgren, "Bundle protocol query extension block", draft-irtf-dtnrg-bpq-00 (work in progress), May 2012.

[TWIMIGHT] Hossmann, Theus, et al. "Twitter in disaster mode: smart probing for opportunistic peers", Proc. of the third ACM international workshop on Mobile Opportunistic Networks. ACM, 2012.

[ONE] The Opportunistic Network Environment simulator. Available at http://www.netlab.tkk.fi/tutkimus/dtn/theone

[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.

[RFC6920] Farrell, S. et al., "Naming Things with Hashes", RFC 6920, April 2013.

[NCOA] Ghodsi, A. et al., "Naming in Content-oriented

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

[nWSN] Heidemann, J. et al., "Building efficient wireless sensor networks with low-level naming", Proc. SOSP. ACM, 2001.

[NDNl] Burke, J. et al., "Authenticated Lighting Control Using Named Data Networking", NDN Technical Report NDN-0011, Oct. 2012.

[IoTScope] Marias, G.F. et al., "Efficient information lookup for the Internet of Things", Proc. WoWMoM. IEEE, 2012.

[PURSUIT] Fotiou, N. et al., "Developing Information Networking Further: From PSIRP to PURSUIT", Proc. BROADNETS. ICST, 2010.

[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] Hernndez-Muoz, 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.

[L1] http://googleblog.blogspot.it/2008/07/we-knew-web-was-big.html

[L2] C. Zhang, P. Dhungel, and K. Di Wu., "Unraveling the BitTorrent ecosystem", IEEE Transactions on Parallel and Distributed Systems, pp. 1164-1177, 2010.

[L3] M. Cha, H. Kwak, P. Rodriguez, Y.-Y. Ahn, 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] J. Zhou, Y. Li, K. Adhikari, and Z.-L. Zhang, "Counting YouTube videos via random prefix sampling", In Proc. of IMC'11, Berlin, Germany, Nov. 2011.

[L5] C. Fricker, P. Robert, J. Roberts, and N. Sbihim, "Impact of traffic mix on caching performance in a content-centric network", Proc. Computer Communications Workshops (INFOCOM WKSHPS), 2012 IEEE Conference on, pp. 310-315. IEEE, 2012.

[L6] H. Yu, D. Zheng, B. Y. Zhao, 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] P. Marciniak, N. Liogkas, A. Legout, 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] A. Bellissimo, B. Levine, and P. Shenoy, "Exploring the use of BitTorrent as the basis for a large trace repository", University of Massachusetts, Tech. Rep., 2004.

[L9] I. Psaras, R. G. Clegg, R. Landa, W. K. Chai, and G. Pavlou, "Modelling and Evaluation of CCN-Caching Trees", In Proc. of the 10th international IFIP conference on Networking, Valencia, Spain, May 2011.

[L10] G. Carofiglio, M. Gallo, L. Muscariello, and D. Perino, "Modeling Data Transfer in Content-Centric Networking", In Proc. of ITC, San Francisco, USA, Sep. 2011.

[L11] L. Breslau, P. Cao, L. Fan, G. Phillips, 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., C. Williamson, and D. Eager., "Traffic analysis of a web proxy caching hierarchy", IEEE Network, Vol.14, No.3, pp.16-23, May/June 2000.

[802.11p] "Information technology - Telecommunications and

information exchange between systems - Local and metropolitan area networks - Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications Amendment 6: Wireless Access in Vehicular Environments", IEEE Standard 802.11p, 2010.

[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.

[PRST2.4] Tagger, B., et al, "Update on the Architecture and Report on Security Analysis", Deliverable 2.4, PURSUIT EU FP7 project, April 2012

[PRST4.5] Riihijarvi, J. et al., "Final Architecture Validation and Performance Evaluation Report", PURSUIT Project Deliverable D4.5, Jan. 2013.

[CMT-D5.2] B 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.

[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.

[NDNpb] Xuan, Y. and Z. Yan, "Enhancing Routing Efficiency in Named Data Network with Piggybacked Interest", Proc. CFI. ACM, 2013.

[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.

[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 Inform??tica 6, no. 1 (2011): 195-202.

[EECD] Lee, U., Rimac, I., Kilper, D., and V. Hilt, "Toward energy-efficient content dissemination", IEEE Network, vol. 25, no. 2, March-April 2011.

[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.

[COAST] COAST, D5.1- File format specification and 3D video codec. Available: http://www.coast-fp7.eu/deliverables.html

[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.

Authors' Addresses

Kostas Pentikousis (editor)
Huawei Technologies
Carnotstrasse 4
10587 Berlin
Germany

Email: k.pentikousis@huawei.com

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@eecs.qmul.ac.uk

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

Email: davieseb@scss.tcd.ie

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

Email: Priya.Mahadevan@parc.com

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

Email: spis@intracom.com

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

Dorothy Gellert
InterDigital Communications, LLC
781 Third Avenue
King Of Prussia, PA 19406-1409
USA

Email: dorothy.gellert@interdigital.com

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