ICNRG
Internet-Draft
Intended Status: Informational
Expires: December 20, 2013

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
D. Gellert
InterDigital
P. Mahadevan
PARC
S. Spyrou
Intracom Telecom
A. Molinaro
UNIRC

June 18, 2013

ICN Baseline Scenarios and Evaluation Methodology
draft-pentikousis-icn-scenarios-03

Abstract

This document aims at establishing a common understanding about potential experimental setups where different information-centric networking (ICN) approaches can be tested and compared against each other while showcasing their advantages. Towards this end, we review the ICN literature and document scenarios which have been considered in previous performance evaluation studies. The scenarios presented aim to exercise a variety of aspects that an ICN solution can address. On the one hand, we consider 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. On the other hand, we focus on ICN-specific aspects, such as information security and trust, persistence, availability, provenance, and location independence.

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. Any-, multi-, and broadcasting are supported by default, and energy efficiency is a design consideration from the 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. note [SoA] that describing ICN architectures is akin to shooting a moving target. We find that comparing these different approaches is often even more tricky. It is not uncommon that researchers devise different performance evaluation scenarios, typically with good reason, in order to highlight the advantages of their approach. This should be expected to some degree at this early stage of development. Nevertheless, we argue that certain scenarios seem to emerge where ICN architectures could showcase their superiority over current systems, in general, and against each other, in particular. In Section 2 we review several scenarios from the published ICN literature and use them 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 can obviously change, as input from the research group is received. For example, this revision adds scenarios stemming from recent work exploring "Vehicular Networking" in ICN.

Section 3 of this document is a first outline of the key elements that should be considered in an ICN evaluation.

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. As such, 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 centralized content server; see also section 2.4, below, and [CBIS]. Based on the CCN naming scheme, [VPC] takes a user name to represent a set of devices that belong to the person. Other users in this in-network, serverless social sharing scenario can access the user's content not via a device name/address but with the user's name. In [VPC], signature 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 evaluations 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. [CCR] 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.

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).

```
\--/
|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

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 is well-known. Some would argue that network primitives which are excellent for information dissemination are not well-suited for conversational services.

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 former case, performance was virtually identical in the two cases evaluated in a testbed. However, the experimental setup presented is quite rudimentary, while the evaluation considered a single voice call only. This scenario does, nonetheless, illustrate that quality telephony services are feasible with at least one ICN approach, but it would need to be further enhanced to include more comprehensive metrics, as well as standardized call arrival patterns, for example, following well-established methodologies from the quality of service/experience (QoS/QoE) evaluation toolbox.

Given the wide-spread deployment of real-time A/V communications, an

ICN approach 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. The design 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 an IP client-server system. This type of work points to benefits both in the data path and the 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 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. Arguably, more work is needed 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. Primarily, this scenario aims to therefore 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 environment, as we discuss in subsection 2.6.

Overall, mobile networking scenarios have not been developed in detail, let alone evaluated at a large scale. Further, the majority of scenarios discussed so far have related to information consumer, rather than source, mobility. We expect that in the coming period more papers will address this topic. Earlier work [mNetInf] argues that for mobile and multiaccess networking scenarios we need to go beyond the current mobility management mechanisms in order to capitalize on the core ICN features. They present a testbed setup (redrawn in Fig. 2) which can serve as the basis for other ICN evaluations. Lindgren [HybICN] explores this scenario further using simulation for an urban setting and reports sizable gains in terms of reduction of object retrieval times and core network capacity use.

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 the possible gains in terms of energy efficiency [COMCOM]. Obvioulsy, ICN architectures must avoid broadcast storms. Early work in this area considers distributed packet suppression techniques which exploit delayed transmissions and overhearing; examples can be found in [MobiA] and [WDays] for ICN-based mobile ad-hoc networks (MANETs), and in [WAK] and [ACMV] 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 also considered in subsection 2.9.

```
+-----------+      +-----------+
| Network 0 |      | Network C |
|           |      |           |
| +--+      | ==== |    +--+   |
| |I2|      |      |    |P1|   |
| +--+      |      |    +--+   |
|     \--/  |      |           |
+-----|C0|--+      |           |
|     /--\  |      |           |
| +--+      |      |           |
| |I3|      |      |      +--+ |
| +--+      | ==== |      |P2| |
|           |      |      +--+ |
| Network 1 |      |           |
+-----------+      +-----------+
```

Figure 2. Overlapping wireless multiaccess

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. [CBIS] argue that in ICN the "where and how" of obtaining information are new degrees of freedom. They

illustrate this with a scenario involving a photo sharing application which takes advantage of whichever access network connectivity is available at the moment (WLAN, Bluetooth, and even SMS) without requiring a centralized infrastructure to synchronize between numerous devices. It is important to highlight that since the focus of the 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.

Carofiglio et al. [SHARE], for instance, also present early work on an analytical framework that attempts to capture the storage/bandwidth tradeoffs that ICN enables and can be used as foundation for a network planning tool. In addition, Chai et al. [CL4M] explore the benefits of ubiquitous caching throughout an information-centric network and argue that "caching less can actually achieve more." These papers also sit alongside a variety of other studies that look at various scenarios such as caching HTTP-like traffic [L9] and BitTorrent-like traffic [BTCACHE]. We observe that much more work is needed in order to understand better how to use optimally all resources available in an information-centric network. In real-world deployments, policy and commercial considerations are also likely to affect the use of particular resources and more work is expected in this direction as well.

In conclusion, scenarios in this category, would cover the communication/computation/storage tradeoffs that an ICN deployment must consider. This would exercise features relating to network planning, perhaps capitalizing on user-provided resources, as well as operational and economical aspects 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, SVN-type of services, 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] [CBIS] [CCR], just to name a few. In addition, Chai et al. [CURLING] present a hop-by-hop hierarchical content resolution approach, which employs receiver-driven multicast over multiple domains, advocating another content dissemination approach. Yet, largely, work in this area did not address the issue of access authorization in detail. Often the distributed content is mostly assumed to be freely accessible by any consumer. Distribution of paid-for or otherwise restricted content on a public ICN network requires more attention in the future. Fotiou et al. [ACDICN] consider a scheme to this effect but it still requires access to an authorization server to verify the user's status after they have obtained the (encrypted) content. 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 fast-changing topology, intermittent connectivity, high node mobility, but also the possibility 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 location- and time-dependent (e.g., road traveler information, collision warning, point-of-interest advertisements) and usually target vehicles in a given area, regardless of their identity or IP address. These applications are likely to benefit from in-network and decentralized data caching and replication mechanisms. Second, content caching is particularly beneficial for intermittent on-the-road connectivity and can speed up data retrieval through content replication in several nodes. Caching can usually be implemented at relatively low cost in vehicles as the energy demands of the ICN device are likely to be a negligible fraction of the total vehicle energy consumption, thus allowing for sophisticated processing, continuous communication and adequate storage in the vehicle. Finally, ICN natively supports asynchronous data exchange between end-nodes. By using (and redistributing) cached named information objects, a mobile node can serve as a link between disconnected areas. In short, ICN can enable communication even under intermittent network connectivity, which is typical of vehicular environments with sparse roadside infrastructure and fast moving nodes.

The advantages of ICN in vehicular networks were preliminarily discussed in [EWC] and [DMND], and additionally investigated in [NOMEN][WAK][DIVA][DIVA2][ACMV][CRoWN]. For example, Bai and

Krishnamachari [EWC] take advantage of the localized and dynamic nature of a VANET to explore how a road congestion notification application can be implemented. Wang et al. [DMND] consider data collection where Road-Site 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 indicating high efficiency 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 the scalability to large numbers of vehicles and the impact of vehicles forwarding Interests and relaying data for other vehicles.

As mentioned in the previous section, due to the short sojourn time between a vehicle and the RSU and the short time of sustained connectivity between vehicles, VANET may be a good showcase for the ICN advantages with respect to content dissemination. In [NOMEN] Wang et al. analyze the advantages of hierarchical naming for vehicular traffic information dissemination. Arnould et al. [DIVA] apply ICN principles to safety information dissemination between vehicles with multiple radio interfaces. In [DIVA2], TalebiFard and Leung use network coding techniques to improve content dissemination over multiple ICN paths. Amadeo et al. [ACMV][[CRoWN] propose an application-independent ICN framework for content retrieval and distribution where the role of providers can be indifferently played by 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 and with or without RSUs along the road. With a ICN/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 [WAK] or from the RSUs [ACMV]. Fig. 2 could apply to content dissemination in VANET scenarios where C0 represents a vehicle which can obtain named information objects via multiple wireless peers and/or RSUs (I2 and I3 in the figure). Also grid topologies can be considered in a urban scenario with RSUs at the crossroads or co-located with traffic lights as in [CRoWN].

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 (i) in more challenging scenarios (e.g., disconnected segments); (ii) when considering both consumer and provider mobility; (iii) designing

smart caching techniques accounting for node mobility patterns, spatial- and time- relevance and popularity of content, and also social relationships between users/cars; (iv) identification of new applications (beyond data dissemination and traffic monitoring) that could benefit of the ICN paradigm in vehicular networks (e.g., mobile cloud, social networking).

2.7. Network Interaction

As ICN shifts the focus from nodes to information objects, the interaction between networks evolves to capitalize on data location independence, efficient and scalable in-network named object availability and multi-access functionality. These interactions become critical in evaluating the technical and economic impact of ICN architectural choices, as noted in [ArgICN]. Additional challenges are presented by the emergence of new types of networks, such as Small Cell Networks (SCN), Heterogeneous Networks (HetNet) and virtual/overlay networks. 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 requestor of the content. The availability of such contextual information across diverse networks can lead to network inefficiencies and data management issues that can benefit from an information-centric approach.

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).

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. 3. (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.

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

Figure 3. Relationships and transit costs between networks A to D

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. 4 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 4. 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; notably, it 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, which includes a network model describing the relationship between AS 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. 5, redrawn here based on Fig. 1 of [ICCP]. Note that it relates well with the topology illustrated in Fig. 1 of this document.

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.

```
                     +-----+
               ------+  J  +------
               |     +--+--+     |
               |        *        |
            +--+--+     *     +--+--+
            |  H  +-----------+  I  |
          **+-----+ **  *  ** +-----+***
         *            * * *             *
         *            * * *             *
      +--+--+        ++-+-++         +--+--+
      |  E  +--------+  F  +---------+  G  +
    **+-----+***     +-----+       **+-----+**
   *            *                 *           *
   *            *                 *           *
+--+--+      +--+--+           +--+--+     +--+--+
|  A  |      |  B  +-----------+  C  |     |  D  |
+-----+      +--+--+           +--+--+     +----++
                |                 |         ^^  | route
          data  |            data |    data ||  | to
                v                 v         ||  v data
             +----+            +----+      +++--+
             |User|            |User|      |Data|
             +----+            +----+      +----+

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

Figure 5. Tier-based set of interconnections between AS A to J

2.8. **Energy Efficiency**

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. For example, the work by Guan et al. [EECCN] indicates that CCN may be much more energy-efficient than traditional CDNs for delivering popular content given the current networking equipment energy consumption levels.

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, large amounts of energy expenditure can be saved by avoiding several further hops. Alternatively, approaches that aim to simplify routers [PURSUIT]

could also reduce energy consumption by pushing routing decisions into more energy-efficient data centers.

Evaluating energy efficiency does not require the definition of new scenarios, 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] [DTNICN] 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]. First, both approaches rely on in-network storage. Second, both approaches espouse late binding of names to locations and, third, both approaches treat data as a long-term component that can exist in the network for extended periods of time.

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.

A 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). Another key example of what is essentially the same scenario is use in emergency and disaster situations where the local infrastructure has either been destroyed or is otherwise inaccessible to first responders. Being able to exchange and cache information without the need for any installed infrastructure could greatly improve the effectiveness of emergency responders. These kind of scenarios bode well with those introduced earlier in Section 2.4 about (re)defining what "infrastructure" may mean in practice in an information-centric network.

Especially in the context of the scenarios discussed above, it is of clear interest to evaluate different ICN approaches with respect both to their delay- and disruption-tolerance, i.e., how effective is the approach when used in a delay tolerant network situation; and to their active support for operations in a DTN environment. Important aspects to be evaluated in support of this application include, but are not limited to, 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.

To assist in this evaluation, within the DTN community, a number of important contact traces have emerged as de-facto evaluative tools. They include Haggle's INFOCOM traces and MIT's Reality Mining. Typically, these are used with the Opportunistic Network Environment (ONE) simulator [ONE] to evaluate the above types of metrics. Based

on this, and with proper extensions, a strong platform for evaluating the delay and disruption tolerance properties of different ICN approaches could be developed.

In summary, the key evaluative metric that DTN scenarios aim to exercise is resilience. This, on the one hand, includes connectivity resilience as offered currently in the DTN community (via store-and-forward) as well as information resilience that can be offered through ICN's use of caching and replication.

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 on 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] 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.

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 (as in the vehicular 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 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 [WAK]. 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. 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). 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 any evaluation will also have to include an assessment of its incremental deployability in the existing network environment together with a view of how the technical functions will incentivize deployers to invest in the capabilities that allow the architecture to spread across the network.

In this section, we present various techniques and considerations for evaluating different ICN architectures. At this stage, we do not intend to develop a complete methodology or a benchmarking tool. Instead, this document proposes key guidelines alongside suggested data sets and high-level approaches that we expect to be of interest to the ICN community as a whole. Through this, researchers and practitioners alike would be able to compare and contrast different ICN designs against each other, 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 [CCNj] is an application-layer platform that can be used to build a CCN overlay. CCN-Joker emulates in user-space all basic aspects of a CCN node (e.g., handling of Interest and Data packets, cache sizing, replacement policies), including both flow and congestion control. The code is open source and is suitable for both emulation-based analyses and real experiments.

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

NEPI, the Network Experimentation Programming Interface, (http://nepi.inria.fr) is a tool developed for controlling and managing large-scale network experiments. NEPI provides an experiment description language to design network experiments, describing topology, applications, and a controller to automatically deploy those experiments on target experimentation environments, such as PlanetLab. The controller is also capable of collecting result and log files during the experiment execution. NEPI also allows to specify node selection filters while designing the experiment, thereby supporting automatic discovery and provisioning of testbed nodes during experiment deployment, without the user having to hand-pick them. It is simple and efficient to use NEPI to evaluate CCNx on large-scale testbeds such as PlanetLab.

3.1.2. Publish/Subscribe Internet Architecture

The PSIRP project has open-sourced its Blackhawk publish-subscribe (Pub/Sub) implementation for FreeBSD; more details are available online at http://www.psirp.org/downloads.html. Despite being limited to one operating system, the code base also provides a virtual image to allow its deployment on other environments through virtualization. The code distribution features a kernel module, a file system and scope daemon, as well as a set of tools, test applications and scripts. This work was extended as part of the PURSUIT project, resulting in the development of the Blackadder prototype for Linux

and FreeBSD. It currently runs on a testbed across Europe and America (MIT) comprising over 25 nodes. Moreover, the ICN simulation environment [ICN-Sim] allows the simulation of new techniques for topology management following the Publish-Subscribe paradigm and the PSIRP approach. The simulator is based on the OMNET++ simulator and the INET/MANET frameworks. It is currently publicly available at http://sourceforge.net/projects/icnsim. A design characteristic of this platform is the separation between the network and topology management policies. An interface is used to provide this functionality and policies can be imported and applied in the network as topology manager applications running on top of this interface.

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

3.1.4. 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 AKARI program (see http://akari-project.nict.go.jp/) is an Architecture Design Project from the National Institute of Information and Communications Technology (NICT) of Japan. AKARI fosters the development of a new network architecture and design to support future technologies. As with the other initiatives, it addresses a number of research questions, considering novel approaches on optical and wireless networks, transport, identifier/locator split, security, routing with quality of service, virtualization, among others.

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 consideration in this choice of this with CAIDA's ITDK traces recording millions of routers across thousands of domains. Beyond these traces there is a wide range of synthetic topologies, such as the Barabasi-Albert model [BA] and the Watts-Strogatz small-world topology [WATTS]. These synthetic traces allow experiments to be performed whilst controlling various key parameters (e.g. degree). Through this, different aspects can be investigated, such as inspecting resilience properties. For some research, this may be more

appropriate as, practically speaking, there are no assurances that a future ICN will share the same topology with today's networks.

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

Finally, the dynamic aspects of a topology, such as node and content mobility, disruption patterns, packet loss rates as well as link and node failure rates, to name a few, should also be carefully considered. As mentioned in 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 the content catalog. With this shared point of reference, the use of the same set of parameters (depending on the scenario of interest) among researchers will be eased, and different proposals could be compared on a common base.

Table I. Content catalog

Traffic Load	Catalog Size [L1][L2][L3][L5]	Mean Object Size [L4][L5][L7][L8][L9][L10]	Popularity Distribution [L3][L5][L6][L11][L12]
Web	10^12	Chunk: 1-10 kB	Zipf, 0.64 <= alpha <= 0.83
File sharing	5x10^6	Chunk: 250-4096 kB Object: ~800 MB	Zipf, 0.75 <= alpha<= 0.82
UGC	10^8	Object: ~10 MB	Zipf, alpha >= 2
VoD	10^4	Object: ~100 MB	Zipf, 0.65 <= alpha <= 1

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

Several studies in the past years have stated that Zipf's law is the discrete distribution that best represents the request frequency in a number of application scenarios, ranging from the Web to VoD services. The key aspect of this distribution is that the frequency of a content request is inversely proportional to the rank of the content itself, i.e., the smaller the rank, the higher the request frequency. If we denote with M the content catalog cardinality and with 1 <= 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.

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

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 to change the mechanisms of data access and transport from a host-to-host model to a user-to-information model. The premise is that the effort invested in changing models will be offset, or even surpassed, by the potential of a "better" network. However, such a claim can be validated only if it is quantified.

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

To arrive at a set of relevant metrics we could survey the various ICN approaches and their design requirements (as metrics should normally correspond to requirements). Furthermore, as we want our metrics to be applicable to host-centric networks as well, we should also look at the capabilities and design requirements of IP networks. Standard metrics already exist for IP networks and it would certainly be beneficial to take them into account.

Depending on the type of evaluation and the focal area of interest, e.g. name resolution vs. routing efficiency vs. congestion control and fair sharing of resources vs. QoS for A/V communications, the metrics that are of prime importance may vary. That said, we should in general consider two broad categories: traffic-related metrics and system metrics.

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 do worse than to quantify the data transport performance of the network in terms of

Quality of Service (QoS) metrics.

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

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

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

each measurement is related to network performance. The rest is related to application performance, e.g., video coding, or even device capabilities, both of which are irrelevant to our purposes here and are generally hard to separate. We therefore see the use of QoE metrics in measuring ICN performance as a poor choice.

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. [SensReqs] provide a comprehensive set of sensor and IoT-related requirements, for example, which include aspects such as resource utilization, service life-cycle management and device management.

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

3.5. Resource Equivalence and Tradeoffs

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

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

3.6. Technology Evolution Assumptions

TBD

4. Security Considerations

TBD

5. IANA Considerations

This document presents no IANA considerations.

6. Acknowledgments

This document has benefited from comments and proposed text provided by the following members of the IRTF Information-Centric Networking Research Group (ICNRG): Marica Amadeo, Claudia Campolo, Luigi Alfredo Grieco, Myeong-Wuk Jang, Ren Jing, Will Liu, and Jianping Wang.

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.

[NetInf] Ahlgren, B. et al., "Design considerations for a network of information", Proc. CoNEXT Re-Arch Workshop. ACM, 2008.

[CCN] Jacobson, V. et al., "Networking Named Content", Proc. CoNEXT. ACM, 2009.

[PSI] Trossen, D. and G. Parisis, "Designing and realizing an information-centric internet", IEEE Commun. Mag., vol. 50, no. 7, July 2012.

[DONA] Koponen, T. et al., "A Data-Oriented (and Beyond) Network Architecture", Proc. SIGCOMM. ACM, 2007.

[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.

[CCR] Arianfar, S. et al., "On content-centric router design and implications", Proc. CoNEXT Re-Arch Workshop. ACM, 2010.

[VoCCN] Jacobson, V. et al., "VoCCN: Voice-over Content-Centric Networks", Proc. CoNEXT Re-Arch Workshop. ACM, 2009.

[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.

[HybICN] Lindgren, A., "Efficient content distribution in an information-centric hybrid mobile networks", Proc. CCNC. IEEE, 2011.

[MobiA] Meisel, M. et al., "Ad Hoc Networking via Named Data", Proc. MobiArch. ACM 2010.

[WDays] Oh, S. Y. et al., "Content Centric Networking in Tactical and Emergency MANETs", Proc. Wireless Days. IFIP, 2010.

[CBIS] Jacobson, V. et al., "Custodian-Based Information Sharing", IEEE Commun. Mag., vol. 50, no. 7, July 2012.

[SHARE] Carofiglio, G. et al., "Bandwidth and storage sharing performance in information centric networking", Proc. SIGCOMM ICN Workshop. ACM, 2011.

[CL4M] Chai, W. K. et al., "Cache 'Less for More' in Information-centric Networks", Proc. Networking. IFIP, 2012.

[BTCACHE] Tyson, G. et al., "A Trace-Driven Analysis of Caching in Content-Centric Networks", Proc. ICCCN. IEEE, 2012.

[CURLING] Chai, W. K. et al., "CURLING: Content-Ubiquitous Resolution and Delivery Infrastructure for Next-Generation Services", IEEE Commun. Mag., vol. 49, no. 3, Mar. 2011.

[DMND] J. Wang, et al., "DMND: collecting data from mobiles using named data", Proc. VNV. IEEE, 2010.

[NOMEN] L. Wang, et al., "Data Naming in Vehicle-to-Vehicle Communications", Proc. INFOCOM NOMEN workshop. IEEE, 2012.

[WAK] L. Wang, et al., "Rapid Traffic Information Dissemination Using Named Data", Proc. MobiHoc NoM workshop. ACM, 2012.

[DIVA] G. Arnould, et al., "A Self-Organizing Content Centric Network Model for Hybrid Vehicular Ad-Hoc Networks". Proc. DIVANet. ACM, 2011.

[DIVA2] P. TalebiFard and V.C.M. Leung, "A Content Centric Approach to Dissemination of Information in Vehicular Networks". Proc. DIVANet. ACM, 2012.

[ACMV] M. Amadeo, et al., "Content-Centric Networking: is that a Solution for Upcoming Vehicular Networks?", Proc. VANET. ACM, 2012.

[CRoWN] M. Amadeo, et al., "CRoWN: Content-Centric Networking in Vehicular Ad Hoc Networks", IEEE Communications Letters, vol. 16, no. 9, Sept. 2012.

[COMCOM] M. Amadeo, et al., "E-CHANET: Routing, Forwarding and Transport in Information-Centric Multihop Wireless Networks", Computer Communications. Elsevier, Jan. 2013 on line.

[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-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

[EWC] Bai, F. and B. Krishnamachari, "Exploiting the wisdom of the crowd: localized, distributed information-centric VANETs", IEEE Commun. Mag., vol. 48, no. 5, May 2010.

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

[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.

[DTNICN] G. Tyson, J. Bigham and E. Bodanese, "Towards an Information-Centric Delay-Tolerant Network", Proc. IEEE INFOCOM NOMEN 2013, Turin, Italy.

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

[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.

[ndnSIM] Afanasyev, A. et al., ndnSIM: NDN simulator for NS-3 NDN Technical Report NDN-0005, Revision 2, October 2012.

[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.

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

[CCNj] Cianci, I. et al. "CCN - Java Opensource Kit EmulatoR for Wireless Ad Hoc Networks", Proc. 7th ACM Int. Conf. on Future Internet Technologies, Seoul, Korea, Sept., 2012.

[CCNPL] Muscariello, L., "Content centric networking packet level simulator", available online at http://perso.rd.francetelecom.fr/muscariello/sim.html

[ICN-Sim] N. Vastardis et al., "Simulation Tools Enabling Research on Information-centric Networks", Proc. ICC FutureNet Workshop. IEEE, 2012.

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

[WATTS] Watts, D. J. and S. H. Strogatz, "Collective dynamics of small-world networks", Nature, vol. 393, no. 6684, pp. 40-"10, 1998.

[Montage] Hussain, A. and J. Chen, "Montage Topology Manager: Tools for Constructing and Sharing Representative Internet Topologies", DETER Technical Report, ISI-TR-684, Aug. 2012.

[DELAY] Kaune, S. et al., "Modelling the Internet Delay Space Based on Geographical Locations", Proc. Euromicro, Weimar, Germany, 2009.

[P2PMod] Saleh, O., and M. Hefeeda, "Modeling and caching of peer-to-peer traffic", Proc. ICNP. IEEE, 2006.

[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", In Proc. of IEEE NOMEN 2012, Workshop on Emerging Design Choices in Name-Oriented Networking, Orlando, USA, Mar. 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.

[SensReqs] Karnouskos, S. et al., "Requirement considerations for ubiquitous integration of cooperating objects", Proc. NTMS. IFIP, 2011.

[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

[ACDICN] Fotiou, N. et al., "Access control enforcement delegation for information-centric networking architectures", Proc. SIGCOMM ICN 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

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

Email: dorothy.gellert@interdigital.com

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

Email: Priya.Mahadevan@parc.com

Spiros Spyrou
Intracom Telecom
19.7 km Markopoulou Ave.
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