

Internet Research Task Force (IRTF)
Request for Comments: 7476
Category: Informational
ISSN: 2070-1721

K. Pentikousis, Ed.
EICT
B. Ohlman
Ericsson
D. Corujo
Universidade de Aveiro
G. Boggia
Politecnico di Bari
G. Tyson
Queen Mary, University of London
E. Davies
Trinity College Dublin
A. Molinaro
UNIRC
S. Eum
NICT
March 2015

Information-Centric Networking: Baseline Scenarios

Abstract

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

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

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Research Task Force (IRTF). The IRTF publishes the results of Internet-related research and development activities. These results might not be suitable for deployment. This RFC represents the consensus of the Information-Centric Networking Research Group of the Internet Research Task Force (IRTF). Documents approved for publication by the IRSG are not a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7476>.

Copyright Notice

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

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

Table of Contents

1. Introduction	3
1.1. Baseline Scenario Selection	4
1.2. Document Goals and Outline	5
2. Scenarios	6
2.1. Social Networking	6
2.2. Real-Time Communication	7
2.3. Mobile Networking	9
2.4. Infrastructure Sharing	11
2.5. Content Dissemination	13
2.6. Vehicular Networking	14
2.7. Delay- and Disruption-Tolerance	17
2.7.1. Opportunistic Content Sharing	20
2.7.2. Emergency Support and Disaster Recovery	21
2.8. Internet of Things	22
2.9. Smart City	25
3. Cross-Scenario Considerations	27
3.1. Multiply Connected Nodes and Economics	27
3.2. Energy Efficiency	31
3.3. Operation across Multiple Network Paradigms	33
4. Summary	34
5. Security Considerations	35
6. Informative References	36
Acknowledgments	44
Authors' Addresses	44

1. Introduction

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

It is also worth noting that with the transition from a host-centric to an information-centric communication model the security paradigm changes as well. In a host-centric network, the basic idea is to create secure (remote-access) tunnels to trusted providers of data. In an information-centric network, on the other hand, any source (cache) should be equally usable. This requires some mechanism for

making each information item trustworthy by itself; this can be achieved, for example, by name-data integrity or by signing data objects.

Although interest in ICN is growing rapidly, ongoing work on different architectures, such as NetInf [NetInf], the original Content-Centric Networking [CCN], and its successors, Project CCNx [CCNx] and Named Data Networking (NDN) [NDNP], the Publish-Subscribe Internet (PSI) architecture [PSI], and the Data-Oriented Network Architecture [DONA] is far from being completed. One could think of ICN today as being at a stage of development similar to that of packet-switched networking in the late 1970s when different technologies, e.g., DECnet, Internetwork Packet Exchange (IPX), and IP, just to name a few, were being actively developed and put to the test. As such, ICN's current development phase 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 partially 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.

The first draft version of this document appeared in November 2012. It was adopted by ICNRG at IETF 87 (July 2013) as the document to address the work item on the definition of "reference baseline scenarios to enable performance comparisons between different approaches". Earlier draft versions of this document have been presented during the ICNRG meetings at IETF 85, IETF 86, IETF 87, IETF 88, IETF 89, and the ICNRG interim meeting in Stockholm in February 2013. This document has been reviewed, commented, and discussed extensively for a period of nearly two years by the vast majority of ICNRG members, which certainly exceeds 100 individuals. It is the consensus of ICNRG that the baseline scenarios described in this document should be published in the IRTF Stream of the RFC series. This document does not constitute a standard.

1.1. Baseline Scenario Selection

Earlier surveys [SoA1] [SoA2] note that describing ICN architectures is akin to shooting a moving target. We find that comparing these different approaches is often even more tricky. It is not uncommon that researchers devise different performance evaluation scenarios, typically with good reason, in order to highlight the advantages of their approach. This should be expected to some degree at this early stage of ICN development. Nevertheless, this document shows that

certain baseline scenarios seem to emerge in which ICN architectures could showcase their comparative advantages over current systems, in general, and against each other, in particular.

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

1.2. Document Goals and Outline

This document incorporates input from ICNRG participants and their corresponding text contributions, has been reviewed by several ICNRG active participants (see Section 7), and represents the consensus of the research group. However, this document does not constitute an IETF standard, but is an Informational document; see also [RFC5743]. As mentioned above, these scenarios are intended to provide a framework for evaluating different ICN approaches. The methodology for how to do these evaluations as well as definitions of metrics that should be used are described in a separate document [EVAL-METHOD]. In addition, interested readers should consider reviewing [CHALLENGES].

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

2. Scenarios

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

2.1. Social Networking

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

Along similar lines, Kim et al. [VPC] present an ICN-based social-networking platform in which a user shares content with her/his family and friends without the need for centralized content servers; see also Section 2.4, below, and [CBIS]. Based on the CCN naming scheme, [VPC] takes a user name to represent a set of devices that belong to the person. Other users in this in-network, serverless social sharing scenario can access the user's content not via a device name/address but with the user's name. In [VPC], signature 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 other words, the above-mentioned evaluation studies indicate that with ICN there may be no need for an end-to-end system design that intermediates between content providers and consumers in a hub-and-spoke fashion at all times.

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 Figure 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 Figure 1) obtain content put "on the network" (I1, ..., Im, and B1, B2) by a single user (e.g., Px) relying solely on network primitives.

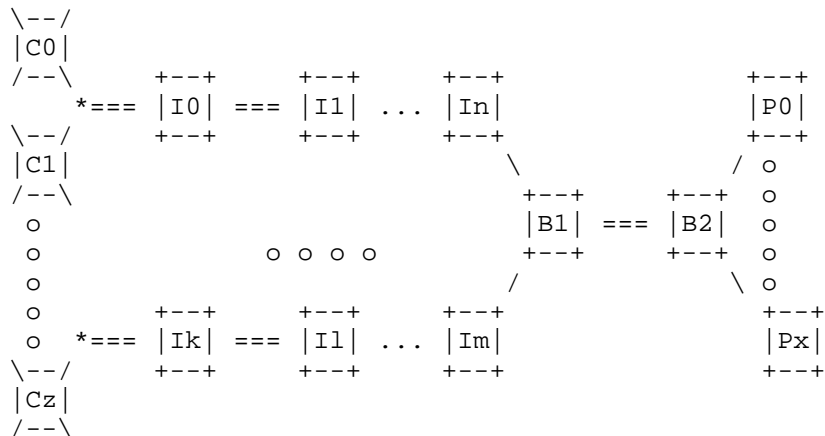


Figure 1. Dumbbell with Linear Daisy Chains

In summary, the social-networking scenario aims to exercise each ICN architecture in terms of network efficiency, multicast support, caching performance and its reliance on centralized mechanisms (or lack thereof).

2.2. Real-Time Communication

Real-time audio and video (A/V) communications include an array of services ranging from one-to-one voice calls to multiparty multimedia 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 (QoS) requirements that this type of communication imposes on network infrastructure are well known. Since one could argue that network primitives that are excellent for information dissemination are not well-suited for conversational services, ICN evaluation studies should consider real-time communication scenarios in detail.

Notably, Jacobson et al. [VoCCN] presented an early evaluation where the performance of a VoIP (Voice over IP) call using an information-centric approach was compared with that of an off-the-shelf VoIP implementation using RTP/UDP. The results indicated that despite the extra cost of adding security support in the ICN approach, performance was virtually identical in the two cases evaluated in

their testbed. However, the experimental setup presented is quite rudimentary, while the evaluation considered a single voice call only. Xuan and Yan [NDNpb] revisit the same scenario but are primarily interested in reducing the overhead that may arise in one-to-one communication employing an ICN architecture. Both studies illustrate that quality telephony services are feasible with at least one ICN approach. That said, future ICN evaluations should employ standardized call arrival patterns, for example, following well-established methodologies from the QoS and QoE (Quality of Experience) evaluation toolbox and would need to consider more comprehensive metrics.

Given the widespread deployment of real-time A/V communications, an evaluation of an ICN system should demonstrate capabilities beyond feasibility. For example, with respect to multimedia conferencing, Zhu et al. [ACT] describe the design of a distributed audio conference tool based on NDN. Their system includes ICN-based conference discovery, speaker discovery, 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 code and show that stringent temporal requirements can be met, while scalability is significantly improved when compared to a host-centric (IP-based) client-server system. This type of work points to benefits for both the data and control path of a modern network infrastructure.

Real-time communication also brings up the issue of named data granularity for dynamically generated content. In many cases, A/V data is generated in real-time and is distributed immediately. One possibility is to apply a single name to the entire content, but this could result in significant distribution delays. Alternatively, distributing A/V content in smaller "chunks" that 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 illustrating the benefits of adopting an information-centric approach as opposed to a host-centric one, and thus more work is recommended in this direction. Scenarios in this category should illustrate not only feasibility but reduced complexity, increased scalability, reliability, and capacity to meet stringent QoS/QoE requirements when compared to established host-centric solutions. Accordingly, the primary aim of this scenario is to exercise each ICN architecture in terms of its ability to satisfy real-time QoS requirements and provide improved user experience.

2.3. Mobile Networking

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

Dannewitz et al. [N-Scen] illustrate a scenario where a multiaccess end-host can retrieve email securely using a combination of cellular and Wireless Local Area Network (WLAN) connectivity. This scenario borrows elements from previous work, e.g., [DTI], and develops them further with respect to multiaccess. Unfortunately, Dannewitz et al. [N-Scen] do not present any results demonstrating that an ICN approach is, indeed, better. That said, the scenario is interesting as it considers content specific to a single user (i.e., her mailbox) and does point to reduced complexity. It is also compatible with recent work in the Distributed Mobility Management (DMM) Working Group within the IETF. Finally, Xylomenos et al. [PSIMob] as well as Pentikousis [EEMN] argue that an information-centric architecture can avoid the complexity of having to manage tunnels to maintain end-to-end connectivity as is the case with mobile anchor-based protocols such as Mobile IP (and its variants). Similar considerations hold for a vehicular (networking) environment, as we discuss in Section 2.6.

Overall, mobile networking scenarios have not been developed in detail, let alone evaluated at a large scale. Further, the majority of scenarios discussed so far have related to the mobility of the information consumer, rather than the source. 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 Figure 2) that can serve as the basis for other ICN evaluations. In this scenario, node "C0" has multiple network interfaces that can access local domains N0 and N1 simultaneously, allowing C0 to retrieve objects from whichever server (I2 or I3) can supply them without necessarily needing to access the servers in the core network "C" (P1 and P2). Lindgren [HybICN] explores this

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

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

To summarize, mobile networking scenarios should aim to provide service continuity for those applications that require it, decrease complexity and control signaling for the network infrastructure, as well as increase wireless capacity utilization by taking advantage of the broadcast nature of the medium. Beyond this, mobile networking scenarios should form a cross-scenario platform that can highlight how other scenarios can still maintain their respective performance metrics during periods of high mobility.

2.4. Infrastructure Sharing

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

network, although compensation schemes motivating users to contribute resources remain a research challenge primarily from a business perspective.

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

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

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

2.5. Content Dissemination

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

Decentralized content dissemination with on-the-fly aggregation of information sources was envisaged in [N-Scen], where information objects can be dynamically assembled based on hierarchically structured subcomponents. For example, a video stream could be associated with different audio streams and subtitle sets, which can all be obtained from different sources. Using the topology depicted in Figure 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 that may be available in different encoding formats. Effectively, this scenario has the information consumer issuing independent requests for content based on information identifiers, and stitching the pieces together irrespective of "where" or "how" they were obtained.

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

Content dissemination scenarios, in general, have a large overlap with those described in the previous sections and are explored in several papers, such as [DONA], [PSI], [PSIMob], [NetInf], [CCN], [CBIS], and [CCR], just to name a few. In addition, Chai et al. [CURLING] present a hop-by-hop hierarchical content resolution approach that employs receiver-driven multicast over multiple domains, advocating another content dissemination approach. Yet, largely, work in this area did not address the issue of access authorization in detail. Often, the distributed content is mostly assumed to be freely accessible by any consumer. Distribution of

paid-for or otherwise restricted content on a public ICN network requires more attention in the future. Fotiou et al. [ACDICN] consider a scheme to this effect, but it still requires access to an authorization server to verify the user's status after the (encrypted) content has been obtained. This may effectively negate the advantage of obtaining the content from any node, especially in a disruption-prone or mobile network.

In summary, scenarios in this category aim to exercise primarily scalability and the 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, in both extreme as well as typical scenarios, can be specified by 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, and spatial and temporal scope, which pose great challenges to traditional networking solutions. VANETs, by their nature, are characterized by challenges such as fast-changing topology, intermittent connectivity, and high node mobility, but also by the opportunity to combine information from different sources as each vehicle does not care about "who" delivers the named data objects.

ICN is an attractive candidate solution for vehicular networking, as it has several advantages. First, ICN fits well to the nature of typical vehicular applications that are geography- and time-dependent (e.g., road traveler information, accident warning, point-of-interest advertisements) and usually target vehicles in a given area, regardless of their identity or IP address. These applications are likely to benefit from in-network and decentralized data caching and replication mechanisms. Second, content caching is particularly beneficial for intermittent on-the-road connectivity and can speed up data retrieval through content replication in several nodes. Caching can usually be implemented at relatively low cost in vehicles, as the energy demands of the ICN device are likely to be a negligible fraction of the total vehicle energy consumption, thus allowing for sophisticated processing, continuous communication, and adequate storage in the vehicle. Finally, ICN natively supports asynchronous

data exchange between end-nodes. By using (and redistributing) cached named information objects, a mobile node can serve as a link between disconnected areas. In short, ICN can enable communication even under intermittent network connectivity, which is typical of vehicular environments with sparse roadside infrastructure and fast-moving nodes.

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

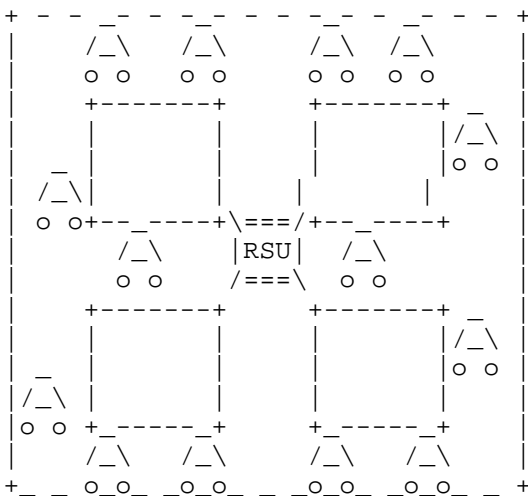


Figure 3. Urban Grid VANET Topology

As mentioned in the previous section, due to the short communication duration between a vehicle and the RSU, and the typically short time of sustained connectivity between vehicles, VANETs may be a good

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

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




Figure 4. Highway VANET Topology

To summarize, VANET scenarios aim to exercise ICN deployment from various perspectives, including scalability, caching, transport, and mobility issues. There is a need for further investigation in (i) challenging scenarios (e.g., disconnected segments); (ii) scenarios involving both consumer and provider mobility; (iii) smart caching techniques that take into consideration node mobility patterns, spatial and temporal relevance, content popularity, and social relationships between users/vehicles; (iv) identification of new

applications (beyond data dissemination and traffic monitoring) that could benefit from the adoption of an ICN paradigm in vehicular networks (e.g., mobile cloud, social networking).

2.7. Delay- and Disruption-Tolerance

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

The DTN architecture [RFC4838] is based on a "store, carry, and forward" paradigm that has been applied extensively to situations where data is carried between network nodes by a "data mule", which carries bundles of data stored in some convenient storage medium (e.g., a USB memory stick). With the advent of sensor and peer-to-peer (P2P) networks between mobile nodes, DTN is becoming a more commonplace type of networking than originally envisioned. Since ICN also does not rely on the familiar end-to-end communications paradigm, there are clear synergies [DTNICN]. It could therefore be argued that many of the key principles embodied within DTN also exist in ICN, as we explain next.

First, both approaches rely on in-network storage. In the case of DTN, bundles are stored temporarily on devices on a hop-by-hop basis. In the case of ICN, information objects are also cached on devices in a similar fashion. As such, both paradigms must provision storage within the network.

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

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

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

The above similarities suggest that integration of the two principles would be feasible. Beyond this, there are also a number of identifiable direct benefits. Through caching and replication, ICN offers strong information resilience, whilst, through store-and-forward, DTN offers strong connectivity resilience. As such, both architectures could benefit greatly from each other. Initial steps have already been taken in the DTN community to integrate ICN principles, e.g., the Bundle Protocol Query Block [BPQ] has been proposed for the DTN Bundle Protocol [RFC5050]. Similarly, initial steps have also been taken in the ICN community, such as [SLINKY]. In fact, the Scalable and Adaptive Internet Solutions (SAIL) project has developed a prototype implementation of NetInf running over the DTN Bundle Protocol.

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

For the purpose of evaluating the use of ICNs in a DTN setting, two key scenarios are identified in this document. (Note the rest of this section uses the term "ICDTN".) These are both prominent use cases that are currently active in both the ICN and DTN communities. The first is opportunistic content sharing, whilst the second is the use of ad hoc networks during disaster recovery (e.g., earthquakes). We discuss both types of scenarios in the context of a simulation-

based evaluation: due to the scale and mobility of DTN-like setups, this is the primary method of evaluation used. Within the DTN community, the majority of simulations are performed using the Opportunistic Network Environment (ONE) simulator [ONE], which is referred to in this document. Before exploring the two scenarios, the key shared components of their simulation are discussed. This is separated into the two primary inputs that are required: the environment and the workload.

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

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

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

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

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

2.7.1. Opportunistic Content Sharing

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

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

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

2.7.2. Emergency Support and Disaster Recovery

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

For the environmental setup, there are no commonly used mobility traces for disaster zones, unlike in the previous scenario. This is clearly due to the difficulty (near impossibility) of acquiring them in a real setting. That said, various synthetic models are available. The Post-Disaster Mobility Model [MODEL1] models civilians and emergency responders after a disaster has occurred,

with people attempting to reach evacuation points (this has also been implemented in the ONE simulator). Aschenbruck et al. [MODEL2] focus on emergency responders, featuring the removal of nodes from the disaster zone, as well as things like obstacles (e.g., collapsed buildings). Cabrero et al. [MODEL3] also look at emergency responders but focus on patterns associated with common procedures. For example, command and control centers are typically set up with emergency responders periodically returning. Clearly, the mobility of emergency responders is particularly important in this setting because they usually are the ones who will "carry" information into the disaster zone. It is recommended that one of these emergency-specific models be used during any evaluations, due to the inaccuracy of alternate models used for "normal" behavior.

The workload input in this evaluative scenario is far simpler than for the previous scenario. In emergency cases, the dissemination of individual pieces of information to all parties is the norm. This is often embodied using things like the Common Alert Protocol (CAP), which is an XML standard for describing warning message. It is currently used by various systems, including the Integrated Public Alert & Warning System and Google Crisis Response. As such, small objects (e.g., 512 KB to 2 MB) are usually generated containing text and images; note that the ONE simulator offers utilities to easily generate these. These messages are also always generated by central authorities, therefore making the placement problem easier (they would be centrally generated and given to emergency responders to disseminate as they pass through the disaster zone). The key variable is therefore the generation rate, which is synonymous with the rate that microblogs are written in the previous scenario. This will largely be based on the type of disaster occurring; however, hourly updates would be an appropriate configuration. Higher rates can also be tested, based on the rate at which situations change (landslides, for example, can exhibit highly dynamic properties).

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

2.8. Internet of Things

Advances in electronics miniaturization combined with low-power wireless access technologies (e.g., ZigBee, Near Field Communication (NFC), Bluetooth, and others) have enabled the coupling of interconnected digital services with everyday objects. As devices with sensors and actuators connect into the network, they become "smart objects" and form the foundation for the so-called Internet of

Things (IoT). IoT is expected to increase significantly the amount of content carried by the network due to machine-to-machine (M2M) communication as well as novel user-interaction possibilities.

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

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

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

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

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

IoT exposes ICN concepts to a stringent set of requirements that are exacerbated by the quantity 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) that is compared with the Data-Oriented Network Architecture described in [DONA]. Their results show how topological routing information has a positive impact on resolution, at the expense of memory and processing overhead.

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

However, one important concern in IoT scenarios, also motivated by this strongly heterogeneous environment, is how content dissemination will be affected by the different semantics of the disparate

information and content being shared. In fact, this is already a difficult problem that goes beyond the scope of ICN [SEMANT]. With the ability of the network nodes to cache forwarded information to improve future requests, a challenge arises regarding whether the ICN fabric should be involved in any kind of procedure (e.g., tagging) that facilitates the relationship or the interpretation of the different sources of information.

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

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

2.9. Smart City

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

In a smart city, ICT plays a leading role and acts as the glue bringing together all actors, services, resources (and their interrelationships) that the urban environment is willing to host and provide [MVM]. ICN appears particularly suitable for these scenarios. Domains of interest include intelligent transportation systems, energy networks, health care, A/V communications, peer-to-

peer and collaborative platforms for citizens, social inclusion, active participation in public life, e-government, safety and security, and sensor networks. Clearly, this scenario has close ties to the vision of IoT, discussed in the previous section, as well as to vehicular networking.

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

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

A different scenario is considered in [ACC], where DHTs are employed for distributed, scalable, and geographically aware service lookup in a smart city. Also in this case, the ICN application is validated by considering a small-scale testbed: a small number of nodes are emulated with simple embedded PCs or specific hardware boards (e.g., for some sensor nodes); other nodes (which connect 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 that is publicly available. Some software applications are designed on real systems (e.g., PCs and smartphones).

With respect to evaluating ICN approaches in smart-city scenarios, it is necessary to consider generic metrics useful to track and monitor progress on services results and also for comparing localities between themselves and learn from the best [ISODIS]. In particular, it is possible to select a specific set of Key Performance Indicators (KPIs) for a given project in order to evaluate its success. These

KPIs may reflect the city's environmental and social goals, as well as its economic objectives, and they can be calculated at the global, regional, national, and local levels. Therefore, it is not possible to define a unique set of interesting metrics, but in the context of smart cities, the KPIs should be characterized with respect to the developed set of services offered by using the ICN paradigm.

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

3. Cross-Scenario Considerations

This section discusses considerations that span multiple scenarios.

3.1. Multiply Connected Nodes and Economics

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

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

Infrastructure networks are generally under the control of separate economic entities that may have different policies about the information of an ICN deployed within their network caches. As ICN shifts the focus from nodes to information objects, the interaction

between networks will likely evolve to capitalize on data location independence, efficient and scalable in-network named object availability, and access via multiple paths. These interactions become critical in evaluating the technical and economic impact of ICN architectural choices, as noted in [ArgICN]. Beyond simply adding diversity in deployment options, these networks have the potential to alter the incentives among existing (and future, we may add) network players, as noted in [EconICN].

Moreover, such networks enable more numerous internetwork relationships where exchange of information may be conditioned on a set of multilateral policies. For example, shared SCNs are emerging as a cost-effective way to address coverage of complex environments such as sports stadiums, large office buildings, malls, etc. Such networks are likely to be a complex mix of different cellular and WLAN access technologies (such as HSPA, LTE, and Wi-Fi) as well as ownership models. It is reasonable to assume that access to content generated in such networks may depend on contextual information such as the subscription type, timing, and location of both the owner and requester of the content. The availability of such contextual information across diverse networks can lead to network inefficiencies unless data management can benefit from an information-centric approach. The "Event with Large Crowds" demonstrator created by the SAIL project investigated this kind of scenario; more details are available in [SAIL-B3].

Jacobson et al. [CCN] include interactions between networks in their overall system design and mention both "an edge-driven, bottom-up incentive structure" and techniques based on evolutions of existing mechanisms both for ICN router discovery by the end-user and for interconnecting between Autonomous Systems (ASes). For example, a BGP extension for domain-level content prefix advertisement can be used to enable efficient interconnection between ASes. 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, for example, in [NCOA], in which the authors note 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).




Figure 5. Relationships and Transit Costs between Networks A to D

DiBenedetto et al. [RP-NDN] study policy knobs made available by NDN to network operators. New policies that 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 Figure 5 (based on Figure 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.

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

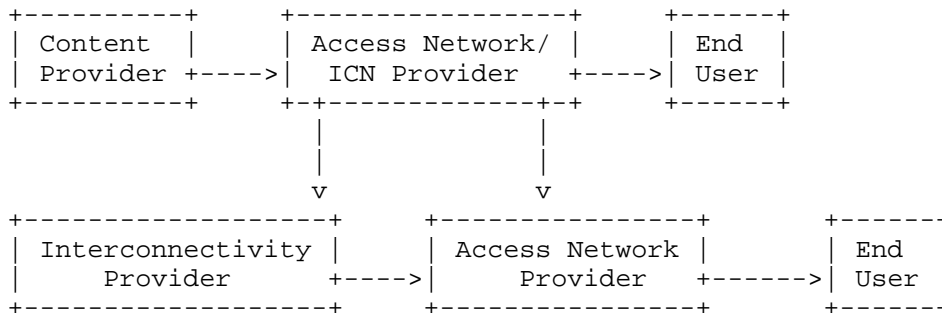
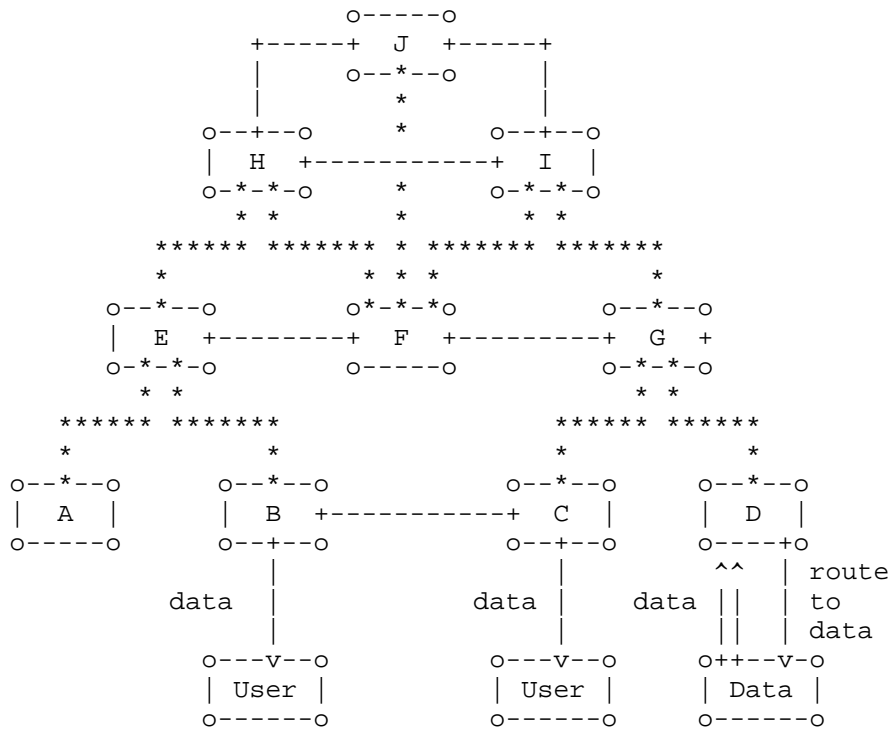


Figure 6. Setup and Operating Costs of Network Entities

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

Rajahalme et al. [ICCP] point out that ICN architectural changes may conflict with the current tier-based peering model. For example, changes leading to shorter paths between ISPs are likely to meet resistance from Tier-1 ISPs. Rajahalme [IDMcast] shows how incentives can help shape the design of specific ICN aspects, and in [IDArch] he presents a modeling approach to exploit these incentives. This includes a network model that describes the relationship between Autonomous Systems based on data inferred from the current Internet, a traffic model taking into account business factors for each AS, and a routing model integrating the valley-free model and policy compliance. A typical scenario topology is illustrated in Figure 7, which is redrawn here based on Figure 1 of [ICCP]. Note that it relates well with the topology illustrated in Figure 1 of this document.



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

Figure 7. Tier-Based Set of Interconnections between AS A to J

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

3.2. Energy Efficiency

ICN has prominent features that can be taken advantage of in order to significantly reduce the energy footprint of future communication networks. Of course, one can argue that specific ICN network elements may consume more energy than today's conventional network

equipment due to the potentially higher energy demands for named-data processing en route. On balance, however, ICN introduces an architectural approach that may compensate on the whole and can even achieve higher energy efficiency rates when compared to the host-centric paradigm.

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

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

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

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

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

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

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

3.3. Operation across Multiple Network Paradigms

Today the overwhelming majority of networks are integrated with the well-connected Internet with IP at the "waist" of the technology hourglass. However, there is a large amount of ongoing research into alternative paradigms that can cope with conditions other than the standard set assumed by the Internet. Perhaps the most advanced of these is Delay- and Disruption-Tolerant Networking (DTN). DTN is

considered as one of the scenarios for the deployment in Section 2.7, but here we consider how ICN can operate in an integrated network that has essentially disjoint "domains" (a highly overloaded term!) or regions that use different network paradigms and technologies, but with gateways that allow interoperation.

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

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

The NetInf architecture, in particular, targets the inter-domain scenario by the use of a convergence-layer architecture [SAIL-B3], and Publish-Subscribe Internet Routing Paradigm (PSIRP) and/or Publish-Subscribe Internet Technology (PURSUIT) is envisaged as a candidate for an IP replacement.

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

4. Summary

This document presents a wide range of different application areas in which the use of information-centric network designs have been evaluated in the peer-reviewed literature. Evidently, this broad range of scenarios illustrates the capability of ICN to potentially address today's problems in an alternative and better way than host-centric approaches as well as to point to future scenarios where ICN may be applicable. We believe that by putting different ICN systems

to the test in diverse application areas, the community will be better equipped to judge the potential of a given ICN proposal and therefore subsequently invest more effort in developing it further. It is worth noting that this document collected different kinds of considerations, as a result of our ongoing survey of the literature and the discussion within ICNRG, which we believe would have otherwise remained unnoticed in the wider community. As a result, we expect that this document can assist in fostering the applicability and future deployment of ICN over a broader set of operations, as well as possibly influencing and enhancing the base ICN proposals that are currently available and possibly assist in defining new scenarios where ICN would be applicable.

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

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

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

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

5. Security Considerations

This document does not impact the security of the Internet.

6. Informative References

- [RFC5743] Falk, A., "Definition of an Internet Research Task Force (IRTF) Document Stream", RFC 5743, December 2009, <<http://www.rfc-editor.org/info/rfc5743>>.
- [RFC5050] Scott, K. and S. Burleigh, "Bundle Protocol Specification", RFC 5050, November 2007, <<http://www.rfc-editor.org/info/rfc5050>>.
- [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, <<http://www.rfc-editor.org/info/rfc4838>>.
- [RFC5568] Koodli, R., Ed., "Mobile IPv6 Fast Handovers", RFC 5568, July 2009, <<http://www.rfc-editor.org/info/rfc5568>>.
- [RFC6920] Farrell, S., Kutscher, D., Dannewitz, C., Ohlman, B., Keranen, A., and P. Hallam-Baker, "Naming Things with Hashes", RFC 6920, April 2013, <<http://www.rfc-editor.org/info/rfc6920>>.
- [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.
- [CCNx] Mosko, M., Solis, I., and E. Uzun, "CCN 1.0 Protocol Architecture", Project CCNx documentation, Xerox Palo Alto Research Center, 2013, <http://ccnx.org/pubs/ICN_CCN_Protocols.pdf>.
- [NDNP] Zhang, L., et al., "Named Data Networking (NDN) Project", NDN Technical Report NDN-0001, Oct. 2010, <<http://named-data.net/publications/techreports/>>.
- [PSI] Trossen, D. and G. Parisis, "Designing and realizing an information-centric internet", IEEE Commun. Mag., vol. 50, no. 7, July 2012.
- [DONA] Koponen, T., et al., "A Data-Oriented (and Beyond) Network Architecture", Proc. SIGCOMM, ACM, 2007.
- [SoA1] Ahlgren, B., et al., "A survey of information-centric networking", IEEE Commun. Mag., vol. 50, no. 7, July 2012.

- [SoA2] Xylomenos, G., et al. "A survey of information-centric networking research", *IEEE Communications Surveys and Tutorials* (2013): 1-26.
- [ICN-SN] Mathieu, B., et al., "Information-centric networking: a natural design for social network applications", *IEEE Commun. Mag.*, vol. 50, no. 7, July 2012.
- [VPC] Kim, J., et al., "Content Centric Network-based Virtual Private Community", *Proc. ICCE, IEEE*, 2011.
- [CBIS] Jacobson, V., et al., "Custodian-Based Information Sharing", *IEEE Commun. Mag.*, vol. 50, no. 7, July 2012.
- [VPC2] Kim, D. and J. Lee, "CCN-based virtual private community for extended home media service", *IEEE Trans. Consumer Electronics*, vol. 57, no. 2, May 2011.
- [CCR] Arianfar, S., et al., "On content-centric router design and implications", *Proc. CoNEXT Re-Arch Workshop, ACM*, 2010.
- [VoCCN] Jacobson, V., et al., "VoCCN: Voice-over Content-Centric Networks", *Proc. CoNEXT Re-Arch Workshop, ACM*, 2009.
- [NDNpb] Xuan, Y. and Z. Yan, "Enhancing Routing Efficiency in Named Data Network with Piggybacked Interest", *Proc. CFI, ACM*, 2013.
- [ACT] Zhu, Z., et al., "ACT: Audio Conference Tool Over Named Data Networking", *Proc. SIGCOMM ICN Workshop, ACM*, 2011.
- [G-COPSS] Chen, J., et al., "G-COPSS: A Content Centric Communication Infrastructure for Gaming Applications", *Proc. ICDCS, IEEE*, 2012.
- [MMIN] Pentikousis, K., and P. Bertin., "Mobility Management in Infrastructure Networks", *IEEE Internet Computing* 17, no. 5 (2013): 74-79.
- [SCES] Allman, M., et al., "Enabling an Energy-Efficient Future Internet through Selectively Connected End Systems", *Proc. HotNets-VI, ACM*, 2007.
- [EEMN] Pentikousis, K., "In Search of Energy-Efficient Mobile Networking", *IEEE Commun. Mag.*, vol. 48, no. 1, Jan. 2010.

- [MOBSURV] Tyson, G., et al., "A Survey of Mobility in Information-Centric Networks: Challenges and Research Directions", Proc. MobiHoc Workshop on Emerging Name-Oriented Mobile Networking Design, ACM, 2012.
- [N-Scen] Dannewitz, C., et al., "Scenarios and research issues for a Network of Information", Proc. MobiMedia, ICST, 2012.
- [DTI] Ott, J. and D. Kutscher, "Drive-thru Internet: IEEE 802.11b for 'automobile' users", Proc. INFOCOM, IEEE, 2004.
- [PSIMob] Xylomenos, G., et al., "Caching and Mobility Support in a Publish-Subscribe Internet Architecture", IEEE Commun. Mag., vol. 50, no. 7, July 2012.
- [mNetInf] Pentikousis, K. and T. Rautio, "A Multiaccess Network of Information", Proc. WoWMoM, IEEE, 2010.
- [HybICN] Lindgren, A., "Efficient content distribution in an information-centric hybrid mobile networks", Proc. CCNC, IEEE, 2011.
- [E-CHANET] M. Amadeo, et al., "E-CHANET: Routing, Forwarding and Transport in Information-Centric Multihop Wireless Networks", Computer Communications, Elsevier, Jan. 2013 online.
- [MobiA] Meisel, M., et al., "Ad Hoc Networking via Named Data", Proc. MobiArch, ACM 2010.
- [CCNMANET] Oh, S. Y., et al., "Content Centric Networking in Tactical and Emergency MANETs", Proc. Wireless Days, IFIP, 2010.
- [RTIND] Wang, L., et al., "Rapid Traffic Information Dissemination Using Named Data", Proc. MobiHoc NoM workshop, ACM, 2012.
- [CCNVANET] Amadeo, M., et al., "Content-Centric Networking: is that a Solution for Upcoming Vehicular Networks?", Proc. VANET, ACM, 2012.
- [ABC] Gustafsson, E., and A. Jonsson. "Always best connected", Wireless Communications, IEEE 10.1 (2003): 49-55.
- [SHARE] Muscariello, L., et al., "Bandwidth and storage sharing performance in information centric networking", Proc. SIGCOMM ICN Workshop, ACM, 2011.

- [CL4M] Chai, W. K., et al., "Cache 'Less for More' in Information-centric Networks", Proc. Networking, IFIP, 2012.
- [CCNCT] Psaras, I., et al., "Modelling and Evaluation of CCN-Caching Trees", In Proc. of the 10th international IFIP conference on Networking, Valencia, Spain, May 2011.
- [BTCACHE] Tyson, G., et al., "A Trace-Driven Analysis of Caching in Content-Centric Networks", Proc. ICCCN, IEEE, 2012.
- [CURLING] Chai, W. K., et al., "CURLING: Content-Ubiquitous Resolution and Delivery Infrastructure for Next-Generation Services", IEEE Commun. Mag., vol. 49, no. 3, Mar. 2011.
- [ACDICN] Fotiou, N., et al., "Access control enforcement delegation for information-centric networking architectures", Proc. SIGCOMM ICN Workshop, ACM, 2012.
- [EWC] Bai, F. and B. Krishnamachari, "Exploiting the wisdom of the crowd: localized, distributed information-centric VANETs", IEEE Commun. Mag., vol. 48, no. 5, May 2010.
- [DMND] Wang, J., R. Wakikawa, and L. Zhang, "DMND: Collecting data from mobiles using Named Data", Proc. Vehicular Networking Conference (VNC), IEEE, 2010.
- [DNV2V] Wang, L., et al., "Data Naming in Vehicle-to-Vehicle Communications", Proc. INFOCOM NOMEN workshop, IEEE, 2012.
- [CCNHV] Arnould, G., et al., "A Self-Organizing Content Centric Network Model for Hybrid Vehicular Ad-Hoc Networks". Proc. DIVANet, ACM, 2011.
- [CCDIVN] TalebiFard, P. and V.C.M. Leung, "A Content Centric Approach to Dissemination of Information in Vehicular Networks", Proc. DIVANet, ACM, 2012.
- [CROWN] Amadeo, M., et al., "CROWN: Content-Centric Networking in Vehicular Ad Hoc Networks", IEEE Communications Letters, vol. 16, no. 9, Sept. 2012.
- [SCN] Hoydis, J., et al., "Green small-cell networks", IEEE Vehicular Technology Magazine, vol. 6, no.1, pp. 37-43, March 2011.

- [HetNet] Li, H., et al. "Efficient HetNet implementation using broadband wireless access with fiber-connected massively distributed antennas architecture", IEEE Wireless Communications, vol. 18, no.3, pp. 72-78, June 2011.
- [ArgICN] Trossen, D., et al., "Arguments for an information centric internetworking architecture", ACM SIGCOMM CCR, 40:26-33, Apr. 2010.
- [EconICN] Agyapong, P. and M. Sirbu, "Economic Incentives in Information Centric Networking: Implications for Protocol Design and Public Policy", IEEE Commun. Mag., vol. 50, no. 12, Dec. 2012.
- [SAIL-B3] Kutscher, D., Ed., et al., "Final NetInf Architecture", SAIL Project Deliverable D-B.3, Jan. 2013, <<http://www.sail-project.eu/deliverables/>>.
- [MLDHT] Liu, H., et al., "A multi-level DHT routing framework with aggregation", Proc. SIGCOMM ICN Workshop, ACM, 2012.
- [NCOA] Ghodsi, A., et al., "Naming in Content-oriented Architectures", Proc. SIGCOMM ICN Workshop, ACM, 2011.
- [RP-NDN] DiBenedetto, S., et al., "Routing policies in named data networking", Proc. SIGCOMM ICN Workshop, ACM, 2011.
- [SAIL-A7] Salo, J., Ed., et al., "New business models and business dynamics of the future networks", SAIL Project Deliverable D-A.7, Aug. 2011, <<http://www.sail-project.eu/deliverables/>>.
- [SAIL-A8] Zhang, N., Ed., et al., "Evaluation of business models", SAIL Project Deliverable D-A.8, Jan. 2013, <<http://www.sail-project.eu/deliverables/>>.
- [LIPSIN] Jokela, P., et al., "LIPSIN: line speed publish/subscribe inter-networking", Proc. of ACM SIGCOMM 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 Data-Oriented Networks", Proc. CoNEXT Re-Arch Workshop, ACM, 2008.
- [IDMcast] Rajahalme, J., "Incentive-informed Inter-domain Multicast", Proc. Global Internet Symposium 2010.
- [IDArch] Rajahalme, J., "Inter-domain incentives and Internet architecture", PhD. Dissertation, Aalto University, Aug. 2012.
- [PURSUIT] Fotiou, N., et al., "Developing Information Networking Further: From PSIRP to PURSUIT", Proc. BROADNETS, ICST, 2010.
- [EECCN] Guan, K., et al., "On the Energy Efficiency of Content Delivery Architectures", Proc. ICC Workshops, IEEE, 2011.
- [EECD] Lee, U., Rimac, I., Kilper, D., and V. Hilt, "Toward energy-efficient content dissemination", IEEE Network, vol. 25, no. 2, March-April 2011.
- [NCICN] Montpetit, M. J., Westphal, C., and D. Trossen, "Network coding meets information-centric networking: an architectural case for information dispersion through native network coding", Proc. MOBIHOC NoM Workshop, ACM, 2012.
- [COAST] Gruneberg, K., et al., "File format specification and 3D video codec", COAST Project Deliverable D5.1, July 2011, <<http://www.coast-fp7.eu/deliverables.html>>.
- [ICNDC] Ko, B. J., Pappas, V., Raghavendra, R., Song, Y., Dilmaghani, R. B., Lee, K.-w., and D. Verma, "An information-centric architecture for data center networks", Proc. SIGCOMM ICN Workshop, ACM, 2012.
- [DTN] Fall, K., "A delay-tolerant network architecture for challenged internets", Proc. SIGCOMM, ACM, 2003.
- [DTNICN] Tyson, G., Bigham, J., and E. Bodanese, "Towards an Information-Centric Delay-Tolerant Network", Proc. IEEE INFOCOM NOMEN 2013, Turin, Italy.
- [BPQ] Farrell, S., Lynch, A., Kutscher, D., and A. Lindgren, "Bundle Protocol Query Extension Block", Work in Progress, draft-irtf-dtnrg-bpq-00, May 2012.

- [SLINKY] Kawadia, V., Riga, N., Opper, J., and D. Sampath, "Slinky: An adaptive protocol for content access in disruption-tolerant ad hoc networks", in Proc. MobiHoc Workshop on Tactical Mobile Ad Hoc Networking, 2011.
- [ONE] "The Opportunistic Network Environment simulator", <<http://www.netlab.tkk.fi/tutkimus/dtn/theone>>.
- [TWIMIGHT] Hossmann, T., et al. "Twitter in disaster mode: smart probing for opportunistic peers", Proc. 3rd ACM International Workshop on Mobile Opportunistic Networks, ACM, 2012.
- [MODEL1] Uddin, M.Y.S., Nicol, D.M., Abdelzaher, T.F., and R.H. Kravets, "A post-disaster mobility model for Delay Tolerant Networking", Simulation Conference (WSC), Proceedings of the 2009 Winter, pp. 2785-2796, 13-16 Dec. 2009.
- [MODEL2] Aschenbruck, N., Gerhards-Padilla, E., and P. Martini, "Modeling mobility in disaster area scenarios", Performance Evaluation 66, no. 12 (2009): 773-790.
- [MODEL3] Cabrero, S., Paneda, X.G., Plagemann, T., Melendi, D., and R. Garcia, "An Overlay Routing Protocol for Video over sparse MANETs in Emergencies", Cadernos de Informatica 6, no. 1 (2011): 195-202.
- [IoTEx] Burke, J., "Authoring Place-based Experiences with an Internet of Things: Tussles of Expressive, Operational, and Participatory Goals", Position Paper, Interconnecting Smart Objects with the Internet Workshop, IAB, 2011.
- [IWMT] Kutscher, D. and S. Farrell, "Towards an Information-Centric Internet with more Things", Position Paper, Interconnecting Smart Objects with the Internet Workshop, IAB, 2011.
- [nWSN] Heidemann, J., et al., "Building efficient wireless sensor networks with low-level naming", Proc. SOSP, ACM, 2001.
- [NDN1] Burke, J., et al., "Authenticated Lighting Control Using Named Data Networking", NDN Technical Report NDN-0011, Oct. 2012.
- [CIBUS] Biswas, T., et al., "Contextualized Information-Centric Home Network", Proc. ACM SIGCOMM, ACM, 2013.

- [Homenet] Ravindran, R., et al., "Information-centric Networking based Homenet", Proc. International Workshop on Management of the Future Internet (ManFI), IFIP/IEEE, 2013.
- [IoTScope] Marias, G.F., et al., "Efficient information lookup for the Internet of Things", Proc. WoWMoM, IEEE, 2012.
- [ICN-DHT] Katsaros, K., et al., "On inter-domain name resolution for information-centric networks", Proc. Networking, Springer, 2012.
- [SEMANT] Sheth, A., et al., "Semantic Sensor Web," Internet Computing, IEEE , vol. 12, no. 4, pp.78-83, July-Aug. 2008
- [CPG] Cianci, I., et al., "Content Centric Services in Smart Cities", Proc. NGMAST, IEEE, 2012.
- [MVM] Hernandez-Munoz, J.M., et al., "Smart cities at the forefront of the future Internet", The Future Internet, Springer, 2011.
- [iHEMS] Zhang, J., et al., "iHEMS: An Information-Centric Approach to Secure Home Energy Management", Proc. SmartGridComm, IEEE, 2012.
- [ACC] Andreini, F., et al., "A scalable architecture for geo-localized service access in smart cities", Proc. Future Network and Mobile Summit, IEEE, 2011.
- [IB] Idowu, S. and N. Bari, "A Development Framework for Smart City Services, Integrating Smart City Service Components", Master's Thesis, Lulea University of Technology, 2012.
- [ISODIS] ISO/DIS, "Sustainable development and resilience of communities --Indicators for city services and quality of life", ISO/DIS 37120, 2013.
- [EVAL-METHOD] Pentikousis, K., Ed., Ohlman, B., Davies, E., Spirou, S., Boggia, G., and P. Mahadevan, "Information-centric Networking: Evaluation Methodology", Work in Progress, draft-irtf-icnrg-evaluation-methodology-00, July 2014.
- [CHALLENGES] Kutscher, D., Ed., Eum, S., Pentikousis, K., Psaras, I., Corujo, D., Saucez, D., Schmidt, T., and M. Waehlich, "ICN Research Challenges", Work in Progress, draft-irtf-icnrg-challenges-01, February 2015.

Acknowledgments

Dorothy Gellert contributed to an earlier draft version of this document.

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

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

Authors' Addresses

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

EEmail: k.pentikousis@eict.de

Borje Ohlman
Ericsson Research
S-16480 Stockholm
Sweden

EEmail: Borje.Ohlman@ericsson.com

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

EEmail: dcorujo@av.it.pt

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

EEmail: g.boggia@poliba.it

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

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

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

EEmail: davieseb@scss.tcd.ie

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

EEmail: antonella.molinaro@unirc.it

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

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