

Internet Research Task Force (IRTF)
Request for Comments: 7945
Category: Informational
ISSN: 2070-1721

K. Pentikousis, Ed.
Traveling
B. Ohlman
Ericsson
E. Davies
Trinity College Dublin
S. Spirou
Intracom Telecom
G. Boggia
Politecnico di Bari
September 2016

Information-Centric Networking: Evaluation and Security Considerations

Abstract

This document presents a number of considerations regarding evaluating Information-Centric Networking (ICN) and sheds some light on the impact of ICN on network security. It also surveys the evaluation tools currently available to researchers in the ICN area and provides suggestions regarding methodology and metrics.

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 <insert_name> 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 7841.

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/rfc7945>.

Copyright Notice

Copyright (c) 2016 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
2.	Evaluation Considerations	4
2.1.	Topology Selection	5
2.2.	Traffic Load	6
2.3.	Choosing Relevant Metrics	10
2.3.1.	Traffic Metrics	13
2.3.2.	System Metrics	14
2.4.	Resource Equivalence and Trade-Offs	16
3.	ICN Security Aspects	16
3.1.	Authentication	17
3.2.	Authorization, Access Control, and Logging	18
3.3.	Privacy	19
3.4.	Changes to the Network Security Threat Model	20
4.	Evaluation Tools	21
4.1.	Open-Source Implementations	21
4.2.	Simulators and Emulators	22
4.2.1.	ndnSIM	22
4.2.2.	ccnSIM	23
4.2.3.	Icarus Simulator	23
4.3.	Experimental Facilities	24
4.3.1.	Open Network Lab (ONL)	24
4.3.2.	POINT Testbed	25
4.3.3.	CUTEi: Container-Based ICN Testbed	25
5.	Security Considerations	25
6.	Informative References	26
	Acknowledgments	37
	Authors' Addresses	38

1. Introduction

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

Different ICN approaches are evaluated in the peer-reviewed literature using a mixture of theoretical analysis, simulation and emulation techniques, and empirical (testbed) measurements. The specific methodology employed may depend on the experimentation goal, e.g., whether one wants to evaluate scalability, quantify resource utilization, or analyze economic incentives. In addition, though, we observe that ease and convenience of setting up and running experiments can sometimes be a factor in published evaluations. As discussed in [RFC7476], the development phase that ICN is going through and the plethora of approaches to tackle the hardest problems make this a very active and growing research area but, on the downside, it also makes it more difficult to compare different proposals on an equal footing.

Performance evaluation using actual network deployments has the advantage of realistic workloads and reflects the environment where the service or protocol is to be deployed. In the case of ICN, however, it is not currently clear what qualifies as a "realistic workload". Trace-based analysis of ICN is in its infancy, and more work is needed towards defining characteristic workloads for ICN evaluation studies. Accordingly, the experimental process and the evaluation methodology per se are actively being researched for different ICN architectures. Numerous factors affect the experimental results, including the topology selected; the background traffic that an application is being subjected to; network conditions such as available link capacities, link delays, and loss-rate characteristics throughout the selected topology; failure and disruption patterns; node mobility; and the diversity of devices used.

The goal of this document is to summarize evaluation guidelines and tools alongside suggested data sets and high-level approaches. We expect this to be of interest to the ICN community as a whole, as it can assist researchers and practitioners alike to compare and contrast different ICN designs, as well as with the state of the art in host-centric solutions, and identify the respective strengths and weaknesses. We note that, apart from the technical evaluation of the

functionality of an ICN architecture, the future success of ICN will be largely driven by its deployability and economic viability. Therefore, ICN evaluations should assess incremental deployability in the existing network environment together with a view of how the technical functions will incentivize deployers to invest in the capabilities that allow the architecture to spread across the network.

This document has been produced by the IRTF Information-Centric Networking Research Group (ICNRG). The main objective of the ICNRG is to couple ongoing ICN research in the above areas with solutions that are relevant for evolving the Internet at large. The ICNRG produces documents that provide guidelines for experimental activities in the area of ICN so that different, alternative solutions can be compared consistently, and information sharing can be accomplished for experimental deployments. This document incorporates input from ICNRG participants and their corresponding text contributions; it has been reviewed by several ICNRG active participants (see the Acknowledgments), and represents the consensus of the research group. That said, note that this document does not constitute an IETF standard; see also [RFC5743].

The remainder of this document is organized as follows. Section 2 presents various techniques and considerations for evaluating different ICN architectures. Section 3 discusses the impact of ICN on network security. Section 4 surveys the tools currently available to ICN researchers.

2. Evaluation Considerations

It is clear that the way we evaluate IP networks will not be directly applicable to evaluating ICN. In IP, the focus is on the performance and characteristics of end-to-end connections between a source and a destination. In ICN, the "source" responding to a request can be any ICN node in the network and may change from request to request. This makes it difficult to use concepts like delay and throughput in a traditional way. In addition, evaluating resource usage in ICN is a more complicated task, as memory used for caching affects delays and use of transmission resources; see the discussion on resource equivalents in Section 2.4.

There are two major types of evaluations of ICN that we see a need to make. One type is to compare ICN to traditional networking, and the other type is to compare different ICN implementations and approaches against each other.

In this section, we detail some of the functional components needed when evaluating different ICN implementations and approaches.

2.1. Topology Selection

There's a wealth of earlier work on topology selection for simulation and performance evaluation of host-centric networks. While the classic dumbbell topology is regarded as inappropriate for ICN, most ICN studies so far have been based on that earlier work for host-centric networks [RFC7476]. However, there is no single topology that can be used to easily evaluate all aspects of ICN. Therefore, one should choose from a range of topologies depending on the focus of the evaluation.

For scalability and resilience studies, there is a wide range of synthetic topologies, such as the Barabasi-Albert model [Barabasi99] and the Watts-Strogatz small-world topology [Watts98]. These allow experiments to be performed whilst controlling various key parameters (e.g., node degree). These synthetic topologies are appropriate in the general case, as there are no practical assurances that a future information-centric network will have the same topology as any of today's networks.

When studies look at cost (e.g., transit cost) or migration to ICN, realistic topologies should be used. These can be inferred from Internet traces, such as the CAIDA Macroscopic Internet Topology Data Kit (<http://www.caida.org/data/active/internet-topology-data-kit>) and Rocketfuel (<http://www.cs.washington.edu/research/networking/rocketfuel>). A problem is the large size of the topology (approximately 45K Autonomous Systems, close to 200K links), which may limit the scalability of the employed evaluation tool. Katsaros et al. [Katsaros15] address this problem by using scaled down topologies created following the methodology described in [Dimitropoulos09].

Studies that focus on node or content mobility can benefit from topologies and their dynamic aspects as used in the Delay-Tolerant Networking (DTN) community. As mentioned in [RFC7476], DTN traces are available to be used in such ICN evaluations.

As with host-centric topologies, defining just a node graph will not be enough for most ICN studies. The experimenter should also clearly define and list the respective matrices that correspond to the network, storage, and computation capacities available at each node as well as the delay characteristics of each link [Montage]. Real values for such parameters can be taken from existing platforms such as iPlane (<http://iplane.cs.washington.edu>). Synthetic values could be produced with specific tools [Kaune09].

2.2. Traffic Load

In this subsection, we provide a set of common guidelines, in the form of what we will refer to as a content catalog for different scenarios. This catalog, which is based on previously published work, can be used to evaluate different ICN proposals, for instance, on routing, congestion control, and performance, and can be considered as other kinds of ICN contributions emerge. As we are still lacking ICN-specific traffic workloads, we can currently only extrapolate from today's workloads. A significant challenge then relates to the identification of the applications contributing to the observed traffic (e.g., Web or peer-to-peer), as well as to the exact amount of traffic they contribute to the overall traffic mixture. Efforts in this direction can take heed of today's traffic mix comprising Web, peer-to-peer file sharing, and User-Generated Content (UGC) platforms (e.g., YouTube), as well as Video on Demand (VoD) services. Publicly available traces for these include those from web sites such as the MultiProbe Framework
<<http://multiprobe.ewi.tudelft.nl/multiprobe.html>>,
<<http://an.kaist.ac.kr/traces/IMC2007.html>> (see also [Cha07]), and the UMass Trace Repository
<<http://traces.cs.umass.edu/index.php/Network/Network>>.

Taking a more systematic approach, and with the purpose of modeling the traffic load, we can resort to measurement studies that investigate the composition of Internet traffic, such as [Labovitz10] and [Maier09]. In [Labovitz10], a large-scale measurement study was performed, with the purpose of studying the traffic crossing inter-domain links. The results indicate the dominance of Web traffic, amounting to 52% over all measured traffic. However, Deep Packet Inspection (DPI) techniques reveal that 25-40% of all HTTP traffic actually carries video traffic. Results from DPI techniques also reveal the difficulty in correctly identifying the application type in the case of P2P traffic: mapping observed port numbers to well-known applications shows P2P traffic constituting only 0.85% of overall traffic, while DPI raises this percentage to 18.32% [Labovitz10]. Relevant studies on a large ISP show that the percentage of P2P traffic ranges from 17% to 19% of overall traffic [Maier09]. Table 1 provides an overview of these figures. The "other" traffic type denotes traffic that cannot be classified in any of the first three application categories, and it consists of unclassified traffic and traffic heavily fragmented into several applications (e.g., 0.17% DNS traffic).

Traffic Type	Ratio
Web	31-39%
P2P	17-19%
Video	13-21%
Other	29-31%

Table 1: Traffic Type Ratios of Total Traffic [Labovitz10] [Maier09]

The content catalog for each type of traffic can be characterized by a specific set of parameters:

- a) the cardinality of the estimated content catalog
- b) the size of the exchanged contents (either chunks or entire named information objects)
- c) the popularity of objects (expressed in their request frequency)

In most application types, the popularity distribution follows some power law, indicating that a small number of information items trigger a large proportion of the entire set of requests. The exact shape of the power law popularity distribution directly impacts the performance of the underlying protocols. For instance, highly skewed popularity distributions (e.g., a Zipf-like distribution with a high slope value) favor the deployment of caching schemes, since caching a very small set of information items can dramatically increase the cache hit ratio.

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

$P(X=i) = (1 / i^{(\alpha)}) / C$, with $C = \text{SUM}(1 / j^{(\alpha)})$, $\alpha > 0$ where the sum is obtained considering all values of j , $1 \leq j \leq M$.

A recent analysis of HTTP traffic showed that content popularity is better reflected by a trimodal distribution model in which the head and tail of a Zipf distribution (with slope value 0.84) are replaced by two discrete Weibull distributions with shape parameter values 0.5 and 0.24, respectively [IMB2014].

A variation of the Zipf distribution, termed the Mandelbrot-Zipf distribution was suggested [Saleh06] to better model environments where nodes can locally store previously requested content. For example, it was observed that peer-to-peer file-sharing applications typically exhibited a 'fetch-at-most-once' style of behavior. This is because peers tend to persistently store the files they download, a behavior that may also be prevalent in ICN.

Popularity can also be characterized in terms of:

- a) The temporal dynamics of popularity, i.e., how requests are distributed in time. The popularity distribution expresses the number of requests submitted for each information item participating into a certain workload. However, they do not describe how these requests are distributed in time. This aspect is of primary importance when considering the performance of caching schemes since the ordering of the requests obviously affects the contents of a cache. For example, with a Least Frequently Used (LFU) cache replacement policy, if all requests for a certain item are submitted close in time, the item is unlikely to be evicted from the cache, even by a (globally) more popular item whose requests are more evenly distributed in time. The temporal ordering of requests gains even more importance when considering workloads consisting of various applications, all competing for the same cache space.
- b) The spatial locality of popularity i.e., how requests are distributed throughout a network. The importance of spatial locality relates to the ability to avoid redundant traffic in the network. If requests are highly localized in some area of the entire network, then similar requests can be more efficiently served with mechanisms such as caching and/or multicast, i.e., the concentration of similar requests in a limited area of the network allows increasing the perceived cache hit ratios at caches in the area and/or the traffic savings from the use of multicast. Table 2 provides an overview of distributions that can be used to model each of the identified traffic types i.e., Web, Video (based on YouTube measurements), and P2P (based on BitTorrent measurements). These distributions are the outcome of a series of modeling efforts based on measurements of real traffic workloads ([Breslau99] [Mahanti00] [Busari02] [Arlitt97] [Barford98] [Barford99] [Hefeeda08] [Guo07] [Bellissimo04] [Cheng08]

[Cheng13]). A tool for the creation of synthetic workloads following these models, and also allowing the generation of different traffic mixes, is described in [Katsaros12].

	Object Size	Temporal Locality	Popularity Distribution
Web	Concatenation of Lognormal (body) and Pareto (tail) [Barford98] [Barford99]	Ordering via the Least Recently Used (LRU) stack model [Busari02] Exact timing via exponential distribution [Arlitt97]	Zipf: $p(i)=K/i^a$ i: popularity rank N: total items K: $1/\text{Sum}(1/i^a)$ a: distribution slope values 0.64-0.84 [Breslau99] [Mahanti00]
VoD	Duration/size: Concatenated normal; most videos ~330 kbit/s [Cheng13]	No analytical models Random distribution across total duration	Weibull: $k=0.513$, $\lambda=6010$ Gamma: $k=0.372$, $\theta=23910$ [Cheng08]
P2P	Wide variation on torrent sizes [Heffeda08]. No analytical models exist: Sample a real BitTorrent distribution [Bellissimo04] or use fixed value	Mean arrival rate of 0.9454 torrents/hour Peers in a swarm arrive as $l(t)=l_0 * e^{-t/\tau}$ l0: initial arrival rate (87.74 average) tau: object popularity (1.16 average)* [Guo07]	Mandelbrot-Zipf [Heffeda08]: $p(i)=K/((i+q)/a)$ q: plateau factor, 5 to 100. Flatter head than in Zipf-like distribution (where $q=0$)

* Random ordering of swarm births (first request). For each swarm, calculate a different tau. Based on average tau and object popularity. Exponential decay rule for subsequent requests.

Table 2: Overview of Traffic Type Models

Table 3 summarizes the content catalog. With this shared point of reference, the use of the same set of parameters (depending on the scenario of interest) among researchers will be eased, and different proposals could be compared on a common base.

Traffic Load	Catalog Size	Mean Object Size	Popularity Distribution
	[Zhou11] [Fri12]	[Zhou11] [Fri12]	[Cha07] [Fri12] [Yu06]
	[Goog08]	[Marciniak08]	[Breslau99] [Mahanti00]
	[Zhang10a]	[Bellissimo04]	
	[Cha07]	[Psaras11]	
	[Fri12]	[Carofiglio11]	
=====			
Web	10 ¹²	Chunk: 1-10 KB	Zipf with 0.64 <= alpha <= 0.83

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

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

VoD (+HLS)	10 ⁴	Object: ~100 MB ~1 KB (*)	Zipf, 0.65 <= alpha <= 1
(+DASH)		~5.6 KB (*)	
=====			

UGC = User-Generated Content
 VoD = Video on Demand
 HLS = HTTP Live Streaming
 DASH = Dynamic Adaptive Streaming over HTTP

(*) Using adaptive video streaming (e.g., HLS and DASH), with an optimal segment length (10 s for HLS and 2 s for DASH) and a bitrate of 4500 kbit/s [RFC7933] [Led12]

Table 3: Content Catalog

2.3. Choosing Relevant Metrics

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

To arrive at a set of relevant metrics, it would be beneficial to look at the metrics used in existing ICN approaches, such as Content-Centric Networking (CCN) [Jacobson09] [VoCCN] [Zhang10b], NetInf [4WARD6.1] [4WARD6.3] [SAIL-B2] [SAIL-B3], PURSUIT [PRST4.5], COMET [CMT-D5.2] [CMT-D6.2], Connect [Muscariello11] [Perino11], and CONVERGENCE [Detti12] [Blefari-Melazzi12] [Salsano12]. The metrics used in these approaches fall into two categories: metrics for the approach as a whole, and metrics for individual components (name resolution, routing, and so on). Metrics for the entire approach are further subdivided into traffic and system metrics. It is important to note that the various approaches do not name or define metrics consistently. This is a major problem when trying to find metrics that allow comparison between approaches. For the purposes of exposition, we have tried to smooth over differences by classifying similarly defined metrics under the same name. Also, due to space constraints, we have chosen to report here only the most common metrics between approaches. For more details, the reader should consult the references for each approach.

Traffic metrics in existing ICN approaches are summarized in Table 4. These are metrics for evaluating an approach mainly from the perspective of the end user, i.e., the consumer, provider, or owner of the content or service. Depending on the level where these metrics are measured, we have made the distinction into user, application, and network-level traffic metrics. So, for example, network-level metrics are mostly focused on packet characteristics, whereas user-level metrics can cover elements of human perception. The approaches do not make this distinction explicitly, but we can see from the table that CCN and NetInf have used metrics from all levels, PURSUIT and COMET have focused on lower-level metrics, and Connect and CONVERGENCE opted for higher-level metrics. Throughput and download time seem to be the most popular metrics altogether.

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

Table 4: Traffic Metrics Used in ICN Evaluations

While traffic metrics are more important for the end user, the owner or operator of the networking infrastructure is normally more interested in system metrics, which can reveal the efficiency of an approach. The most common system metrics used are: protocol overhead, total traffic, transit traffic, cost savings, router cost, and router energy consumption.

Besides the traffic and systems metrics that aim to evaluate an approach as a whole, all surveyed approaches also evaluate the performance of individual components. Name resolution, request/data routing, and data caching are the most typical components, as summarized in Table 5. Forwarding Information Base (FIB) size and path length, i.e., the routing component metrics, are almost ubiquitous among approaches, perhaps due to the networking background of the involved researchers. That might be also the reason for the sometimes decreased focus on traffic and system metrics, in favor of component metrics. It can certainly be argued that traffic and system metrics are affected by component metrics; however, no approach has made the relationship clear. With this in mind and taking into account that traffic and system metrics are readily useful to end users and network operators, we will restrict ourselves to those in the following subsections.

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

Table 5: Component Metrics in Existing ICN Approaches

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

2.3.1. Traffic Metrics

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

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

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

2.3.2. System Metrics

Overall system metrics that need to be considered include reliability, scalability, energy efficiency, and delay/disconnection tolerance. In deployments where ICN is addressing specific scenarios, relevant system metrics could be derived from current experience. For example, in Internet of Things (IoT) scenarios, which are discussed in [RFC7476], it is reasonable to consider the current generation of sensor nodes, sources of information, and even measurement gateways (e.g., for smart metering at homes) or smartphones. In this case, ICN operation ought to be evaluated with respect not only to overall scalability and network efficiency, but

also the impact on the nodes themselves. Karnouskos et al. [SensReqs] provide a comprehensive set of sensor and IoT-related requirements, for example, which include aspects such as resource utilization, service life-cycle management, and device management.

Additionally, various specific metrics are also critical in constrained environments, such as processing requirements, signaling overhead, and memory allocation for caching procedures, in addition to power consumption and battery lifetime. For gateways (which typically act as a point of service to a large number of nodes and have to satisfy the information requests from remote entities), we need to consider scalability-related metrics, such as frequency and processing of successfully satisfied information requests.

Finally, given the in-network caching functionality of ICNs, efficiency and performance metrics of in-network caching have to be defined. Such metrics will need to guide researchers and operators regarding the performance of in-network caching algorithms. A first step on this direction has been made in [Psaras11]. The paper proposes a formula that approximates the proportion of time that a piece of content stays in a network cache. The model takes as input the rate of requests for a given piece of content (the Content of Interest (CoI)) and the rate of requests for all other contents that go through the given network element (router) and move the CoI down in the (LRU) cache. The formula takes also into account the size of the cache of this router.

The output of the model essentially reflects the probability that the CoI will be found in a given cache. An initial study [Psaras11] is applied to the CCN / Named Data Networking (NDN) framework, where contents get cached at every node they traverse. The formula according to which the probability or proportion is calculated is given by:

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

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

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

2.4. Resource Equivalence and Trade-Offs

As we have seen above, every ICN network is built from a set of resources, which include link capacities, and different types of memory structures and repositories used for storing named data objects and chunks temporarily (i.e., caching) or persistently, as well as name resolution and other lookup services. A range of engineering trade-offs arise from the complexity and processing requirements of forwarding decisions, management needs (e.g., manual configuration, explicit garbage collection), and routing needs (e.g., amount of state, manual configuration of routing tables, support for mobility).

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

3. ICN Security Aspects

The introduction of an information-centric networking architecture and the corresponding communication paradigm results in changes to many aspects of network security. These will affect all scenarios described in [RFC7476]. Additional evaluation will be required to ensure relevant security requirements are appropriately met by the implementation of the chosen architecture in the various scenarios.

The ICN security aspects described in this document reflect the ICN security challenges outlined in [RFC7927].

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

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

In practice, the work on security in the various ICN research projects has been heavily concentrated on authentication of content. Work on authorization, access control, and privacy and security threats due to the expanded role of in-network caches has been quite limited. For example, a roadmap for improving the security model in NetInf can be found in [Renault09]. As secure communications on the Internet are becoming the norm, major gaps in ICN security aspects are bound to undermine the adoption of ICN. A comprehensive overview of ICN security is also provided in [Tourani16].

In the following subsections, we briefly consider the issues and provide pointers to the work that has been done on the security aspects of the architectures proposed.

3.1. Authentication

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

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

All ICN architectures considered in this document primarily target the validity requirement using strong cryptographic means to tie the content request name to the content. Provenance and relevance are directly targeted to varying extents: There is a tussle or trade-off between simplicity and efficiency of access and level of assurance of all these traits. For example, maintaining provenance information can become extremely costly, particularly when considering (historic) relationships between multiple objects. Architectural decisions have therefore been made in each case as to whether the assessment is carried out by the information-centric network or left to the application.

An additional consideration for authentication is whether a name should be irrevocably and immutably tied to a static piece of preexisting content or whether the name can be used to refer to dynamically or subsequently generated content. Schemes that only target immutable content can be less resource-hungry as they can use digest functions rather than public key cryptography for generating and checking signatures. However, this can increase the load on

applications because they are required to manage many names, rather than use a single name for an item of evolving content that changes over time (e.g., a piece of data containing an age reference).

Data-Oriented Network Architecture (DONA) [DONA] and CCN [Jacobson09] [Smetters09] integrate most of the data needed to verify provenance into all content retrievals but need to be able to retrieve additional information (typically a security certificate) in order to complete the provenance authentication. Whether the application has any control of this extra retrieval will depend on the implementation. CCN is explicitly designed to handle dynamic content allowing names to be pre-allocated and attached to subsequently generated content. DONA offers variants for dynamic and immutable content.

Publish-Subscribe Internet Technology (PURSUIT) [Tagger12] appears to allow implementers to choose the authentication mechanism so that it can, in theory, emulate the authentication strategy of any of the other architectures. It is not clear whether different choices would lead to lack of interoperability.

NetInf uses the Named Information (ni) URI scheme [RFC6920] to identify content. This allows NetInf to assure validity without any additional information but gives no assurance on provenance or relevance. A "search" request allows an application to identify relevant content, and applications may choose to structure content to allow provenance assurance, but this will typically require additional network access. NetInf validity authentication is consequently efficient in a network environment with intermittent connectivity as it does not force additional network accesses and allows the application to decide on provenance validation if required. For dynamic content, NetInf can use, e.g., signed manifests. For more details on NetInf security, see [Dannewitz10].

3.2. Authorization, Access Control, and Logging

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

Access could be limited by encrypting the content, but the necessity of distributing keys out-of-band appears to negate the advantages of in-network caching. This also creates significant challenges when attempting to manage and restrict key access. An authorization delegation scheme has been proposed [Fotioul2]. This scheme allows semi-trusted entities (such as caches or CDN nodes) to delegate access control decisions to third-party access control providers that are trusted by the content publisher. The former entities have no access to subscriber-related information and should respect the decisions of the access control providers.

A recent proposal for an extra layer in the protocol stack [LIRA] gives control of the name resolution infrastructure to the publisher. This enables access logging as well some degree of active cache management, e.g., purging of stale content.

One possible technique that could allow for providing access control to heterogeneous groups and still allow for a single encrypted object representation that remains cacheable is Attribute-Based Encryption (ABE). A first proposal for this is presented in [Ion13]. To support heterogeneous groups and avoid having a single authority that has a master key multi-authority, ABE can be used [Lewkol1].

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

3.3. Privacy

Another area where the architectures have not been significantly analyzed is privacy. Caching implies a trade-off between network efficiency and privacy. The activity of users is significantly more exposed to the scrutiny of cache owners with whom they may not have any relationship. However, it should be noted that it is only the first-hop router/cache that can see who requests what, as requests are aggregated and only the previous-hop router is visible when a request is forwarded.

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

The impact of CCN on privacy has been investigated in [Lauinger10], and the analysis is applicable to all ICN architectures because it is mostly focused on the common caching aspect. The privacy risks of Named Data Networking are also highlighted in [Lauinger12]. Further work on privacy in ICNs can be found in [Chaabanel3]. Finally, Fotiou et al. define an ICN privacy evaluation framework in [Fotioul4].

3.4. Changes to the Network Security Threat Model

The architectural differences of the various ICN models versus TCP/IP have consequences for network security. There is limited consideration of the threat models and potential mitigation in the various documents describing the architectures. [Lauinger10] and [Chaabanel3] also consider the changed threat model. Some of the key aspects are:

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

- o Caches can be misused in the following ways:
 - + Attackers can use caches as storage to make their own content available.
 - + The efficiency of caches can be decreased by attackers with the goal of DoS attacks.
 - + Content can be extracted by any attacker connected to the cache, putting users' privacy at risk.

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

4. Evaluation Tools

Since ICN is an emerging area, the community is in the process of developing effective evaluation environments, including releasing open-source implementations, simulators, emulators, and testbeds. To date, none of the available evaluation tools can be seen as the one and only community reference evaluation tool. Furthermore, no single environment supports all well-known ICN approaches, as we describe below, hindering the direct comparison of the results obtained for different ICN approaches. The subsections that follow review the currently publicly available ICN implementations, simulators, and experimental facilities.

An updated list of the available evaluation tools will be maintained at the ICNRRG Wiki page: <<https://trac.tools.ietf.org/group/irtf/trac/wiki/IcnEvaluationAndTestbeds>>

4.1. Open-Source Implementations

The Named Data Networking (NDN) project has open-sourced a software reference implementation of the architecture and protocol called NDN (<http://named-data.net>). NDN is available for deployment on various operating systems and includes C and Java libraries that can be used to build applications.

CCN-lite (<http://www.ccn-lite.net>) is a lightweight implementation of the CCN protocol that supports most of the key features of CCNx and is interoperable with CCNx. CCN-lite implements the core CCN logic in about 1000 lines of code, so it is ideal for classroom work and course projects as well as for quickly experimenting with CCN extensions. For example, Baccelli et al. use CCN-lite on top of the RIOT operating system to conduct experiments over an IoT testbed [Baccelli14].

PARC is offering CCN source code under various licensing schemes, please see <<http://www.ccnx.org>> for details.

The PURSUIT project (<http://www.fp7-pursuit.eu>) has open-sourced its Blackhawk publish-subscribe (Pub/Sub) implementation for Linux and Android; more details are available at <<https://github.com/fp7-pursuit/blackadder>>. Blackadder uses the Click modular router for ease of development. The code distribution features a set of tools, test applications, and scripts. The POINT project (<http://www.point-h2020.eu>) is currently maintaining Blackadder.

The 4WARD and SAIL projects have open-sourced software that implements different aspects of NetInf, e.g., the NetInf URI format and HTTP and UDP convergence layer, using different programming languages. The Java implementation provides a local caching proxy and client. Further, an OpenNetInf prototype is available as well as a hybrid host-centric and information-centric network architecture called the Global Information Network (GIN), a browser plug-in and video-streaming software. See <<http://www.netinf.org/open-source>> for more details.

4.2. Simulators and Emulators

Simulators and emulators should be able to capture faithfully all features and operations of the respective ICN architecture(s) and any limitations should be openly documented. It is essential that these tools and environments come with adequate logging facilities so that one can use them for in-depth analysis as well as debugging. Additional requirements include the ability to support medium- to large-scale experiments, the ability to quickly and correctly set various configurations and parameters, as well as to support the playback of traffic traces captured on a real testbed or network. Obviously, this does not even begin to touch upon the need for strong validation of any evaluated implementations.

4.2.1. ndnSIM

The Named Data Networking (NDN) project (<http://named-data.net>) has developed ndnSIM [ndnSIM] [ndnSIM2]; this is a module that can be plugged into the ns-3 simulator (<https://www.nsnam.org>) and supports the core features of NDN. One can use ndnSIM to experiment with various NDN applications and services as well as components developed for NDN such as routing protocols and caching and forwarding strategies, among others. The code for ns-3 and ndnSIM is openly available to the community and can be used as the basis for implementing ICN protocols or applications. For more details, see <<http://ndnsim.net/2.0/>>.

4.2.2. ccnSIM

ccnSim [ccnSim] is a CCN-specific simulator that was specially designed to handle forwarding of a large number of CCN-chunks (<http://www.infres.enst.fr/~drossi/index.php?n=Software.ccnSim>). ccnSim is written in C++ for the OMNeT++ simulation framework (<https://omnetpp.org>). Other CCN-specific simulators include the CCN Packet-Level Simulator [CCNPL] and CCN-Joker [Ciancil2]. CCN-Joker emulates in user space all basic aspects of a CCN node (e.g., handling of Interest and Data packets, cache sizing, replacement policies), including both flow and congestion control. The code is open source and is suitable for both emulation-based analyses and real experiments. Finally, Cabral et al. [MiniCCNx] use container-based emulation and resource isolation techniques to develop a prototyping and emulation tool.

4.2.3. Icarus Simulator

The Icarus simulator [ICARUS] focuses on caching in ICN and is agnostic with respect to any particular ICN implementation. The simulator is implemented in Python, uses the Fast Network Simulator Setup tool [Saino13], and is available at <http://icarus-sim.github.io>. Icarus has several caching strategies implemented, including among others ProbCache [Psaras12], node-centrality-based caching [Chail2], and hash-route-based caching [HASHROUT].

ProbCache [Psaras12] is taking a resource management view on caching decisions and approximates the available cache capacity along the path from source to destination. Based on this approximation and in order to reduce caching redundancy across the path, it caches content probabilistically. According to [Chail2], the node with the highest "betweenness centrality" along the path from source to destination is responsible for caching incoming content. Finally, [HASHROUT] calculates the hash function of a content's name and assigns contents to caches of a domain according to that. The hash space is split according to the number of caches of the network. Then, upon subsequent requests, and based again on the hash of the name included in the request, edge routers redirect requests to the cache assigned with the corresponding hash space. [HASHROUT] is an off-path caching strategy; in contrast to [Psaras12] and [Chail2], it requires minimum coordination and redirection overhead. In its latest update, Icarus also includes implementation of the "Satisfied Interest Table" (SIT) [Sourlas15]. The SIT points in the direction where content has been sent recently. Among other benefits, this enables information resilience in case of network fragmentation (i.e., content can still

be found in neighbor caches or in users' devices) and inherently supports user-assisted caching (i.e., P2P-like content distribution).

Tortelli et al. [ICNSIMS] provide a comparison of ndnSIM, ccnSim, and Icarus.

4.3. Experimental Facilities

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

Different experimental facilities have different characteristics and capabilities, e.g., having low cost of use, reproducible configuration, easy-to-use tools, and available background traffic, and being sharable.

4.3.1. Open Network Lab (ONL)

An example of an experimental facility that supports CCN is the Open Network Lab [ONL] that currently comprises 18 extensible gigabit routers and over a 100 computers representing clients and is freely available to the public for running CCN experiments. Nodes in ONL are preloaded with CCNx software. ONL provides a graphical user interface for easy configuration and testbed setup as per the experiment requirements, and also serves as a control mechanism, allowing access to various control variables and traffic counters.

Further, it is also possible to run and evaluate CCN over popular testbeds [PLANETLAB] [EMULAB] [DETERLAB] [OFELIA] by directly running, for example, the CCNx open-source code [Salsano13] [Carofiglio13] [Awiphan13] [Bernardin14]. Also, the Network Experimentation Programming Interface (NEPI) [NEPI] is a tool developed for controlling and managing large-scale network experiments. NEPI can be used to control and manage large-scale CCNx experiments, e.g., on PlanetLab [Quereilha14].

4.3.2. POINT Testbed

The POINT project is maintaining a testbed with 40 machines across Europe, North America (Massachusetts Institute of Technology (MIT)), and Japan (National Institute of Information and Communications Technology (NICT)) interconnected in a topology containing one Topology Manager and one rendezvous node that handle all publish/subscribe and topology formation requests [Paris13]. All machines run Blackadder (see Section 4.1). New nodes can join, and experiments can be run on request.

4.3.3. CUTEi: Container-Based ICN Testbed

NICT has also developed a testbed used for ICN experiments [Asaeda14] comprising multiple servers located in Asia and other locations. Each testbed server (or virtual machine) utilizes a Linux kernel-based container (LXC) for node virtualization. This testbed enables users to run applications and protocols for ICN in two experimentation modes using two different container designs:

1. application-level experimentation using a "common container" and
2. network-level experimentation using a "user container."

A common container is shared by all testbed users, and a user container is assigned to one testbed user. A common container has a global IP address to connect with other containers or external networks, whereas each user container uses a private IP address and a user space providing a closed networking environment. A user can login to his/her user containers using SSH with his/her certificate, or access them from PCs connected to the Internet using SSH tunneling.

This testbed also implements an "on-filesystem cache" to allocate caching data on a UNIX filesystem. The on-filesystem cache system accommodates two kinds of caches: "individual cache" and "shared cache." Individual cache is accessible for one dedicated router for the individual user, while shared cache is accessible for a set of routers in the same group to avoid duplicated caching in the neighborhood for cooperative caching.

5. Security Considerations

This document does not impact the security of the Internet, but Section 3 outlines security and privacy concerns that might affect a deployment of a future ICN approach.

6. Informative References

- [4WARD6.1] Ohlman, B., et al., "First NetInf Architecture Description", 4WARD Project Deliverable D-6.1, April 2009.
- [4WARD6.3] Ahlgren, B., et al., "NetInf Evaluation", 4WARD Project Deliverable D-6.3, June 2010.
- [Arlitt97] Arlitt, M. and C. Williamson, "Internet web servers: workload characterization and performance implications", IEEE/ACM Transactions on Networking, vol. 5, pp. 631-645, DOI 10.1109/90.649565, 1997.
- [Asaeda14] Asaeda, H., Li, R., and N. Choi, "Container-Based Unified Testbed for Information-Centric Networking", IEEE Network, vol. 28, no. 6, pp. 60-66, DOI 10.1109/MNET.2014.6963806, 2014.
- [Awiphan13] Awiphan, S., et al., "Video streaming over content centric networking: Experimental studies on PlanetLab", Proc. Computing, Communications and IT Applications Conference (ComComAp), IEEE, DOI 10.1109/ComComAp.2013.6533602, 2013.
- [Baccelli14] Baccelli, E., et al., "Information Centric Networking in the IoT: Experiments with NDN in the Wild", Proceedings of the 1st international conference on Information-centric networking (ICN '14), ACM, DOI 10.1145/2660129.2660144, 2014.
- [Barabasi99] Barabasi, A. and R. Albert, "Emergence of Scaling in Random Networks", Science, vol. 286, no. 5439, pp. 509-512, DOI 10.1126/science.286.5439.509, 1999.
- [Barford98] Barford, P. and M. Crovella, "Generating representative web workloads for network and server performance evaluation", in ACM SIGMETRICS '98 / PERFORMANCE '98, pp. 151-160, DOI 10.1145/277851.277897, 1998.
- [Barford99] Barford, P., Bestavros, A., Bradley, A., and M. Crovella, "Changes in web client access patterns: Characteristics and caching implications", World Wide Web, vol. 2, pp. 15-28, DOI 10.1023/A:1019236319752, 1999.

- [Bellissimo04] Bellissimo, A., Levine, B., and P. Shenoy, "Exploring the Use of BitTorrent as the Basis for a Large Trace Repository", University of Massachusetts Amherst, Tech. Rep. 04-41, 2004.
- [Bernardini14] Bernardini, C., et al., "Socially-aware caching strategy for content centric networking", Proc. IFIP Networking Conference, DOI 10.1109/IFIPNetworking.2014.6857093, 2014.
- [Blefari-Melazzi12] Blefari Melazzi, N., et al., "Scalability Measurements in an Information-Centric Network", Springer Lecture Notes in Computer Science (LNCS), vol. 7586, DOI 10.1007/978-3-642-41296-7_6, 2012.
- [Breslau99] Breslau, L., Cao, P., Fan, L., Phillips, G., and S. Shenker, "Web caching and zipf-like distributions: evidence and implications", Proc. of INFOCOM '99, New York (NY), USA, DOI 10.1109/INFOCOM.1999.749260, March 1999.
- [Busari02] Busari, M. and C. Williamson, "ProWGen: a synthetic workload generation tool for simulation evaluation of web proxy caches", Computer Networks, vol. 38, no. 6, pp. 779-794, DOI 10.1016/S1389-1286(01)00285-7, 2002.
- [Carofiglio11] Carofiglio, G., Gallo, M., Muscariello, L., and D. Perino, "Modeling Data Transfer in Content-Centric Networking", Proceedings of the 23rd International Teletraffic Congress (ITC '11), San Francisco, USA, September 2011.
- [Carofiglio13] Carofiglio, G., et al., "Optimal multipath congestion control and request forwarding in Information-Centric Networks", Proc. 2013 21st IEEE International Conference on Network Protocols (ICNP), DOI 10.1109/ICNP.2013.6733576, 2013.
- [CCNPL] Institut de Recherche Technologique (IRT) SystemX, "CCNPL-SIM", <<http://systemx.enst.fr/ccnpl-sim>>.
- [ccnSim] Rossini, G. and D. Rossi, "Large scale simulation of CCN networks", Proc. AlgoTel 2012, La Grande Motte, France, May 2012.

- [Cha07] Cha, M., Kwak, H., Rodriguez, P., Ahn, Y.-Y., and S. Moon, "I tube, you tube, everybody tubes: analyzing the world's largest user generated content video system", Proceedings of the 7th ACM SIGCOMM conference on Internet measurement (IMC '07), San Diego (CA), USA, DOI 10.1145/1298306.1298309, October 2007.
- [Chaabane13] Chaabane, A., De Cristofaro, E., Kaafar, M., and E. Uzun, "Privacy in Content-Oriented Networking: Threats and Countermeasures", ACM SIGCOMM Computer Communication Review, Vol. 43, Issue 3, DOI 10.1145/2500098.2500102, July 2013.
- [Cheng08] Cheng, X., Dale, C., and J. Liu, "Statistics and social network of YouTube videos", 16th International Workshop on Quality of Service (IWQoS 2008), IEEE, pp. 229-238, DOI 10.1109/IWQOS.2008.32, 2008.
- [Cheng13] Cheng, X., Liu, J., and C. Dale, "Understanding the Characteristics of Internet Short Video Sharing: YouTube as a Case Study", IEEE Transactions on Multimedia, vol. 15, issue 5, DOI 10.1109/TMM.2013.2265531, 2013.
- [Chai12] Chai, W., He, D., Psaras, I., and G. Pavlou, "Cache 'Less for More' in Information-centric Networks", Proceedings of the 11th international IFIP TC 6 conference on Networking (IFIP '12), DOI 10.1007/978-3-642-30045-5_3, 2012.
- [Cianci12] Cianci, I. et al. "CCN - Java Opensource Kit EmulatoR for Wireless Ad Hoc Networks", Proc. of the 7th International Conference on Future Internet Technologies (CFI '12), Seoul, Korea, DOI 10.1145/2377310.2377313, September 2012.
- [CMT-D5.2] Beben, A., et al., "Scalability of COMET System", COMET Project Deliverable D5.2, February 2013.
- [CMT-D6.2] Georgiades, M., et al., "Prototype Experimentation and Demonstration", COMET Project Deliverable D6.2, February 2013.
- [Dannewitz10] Dannewitz, C., Golic, J., Ohlman, B., B. Ahlgren, "Secure Naming for A Network of Information", IEEE Conference on Computer Communications Workshops (INFOCOM), San Diego, CA, DOI 10.1109/INFCOMW.2010.5466661, March 2010.

- [DETERLAB] Benzel, T., "The Science of Cyber-Security Experimentation: The DETER Project", Proceedings of the 27th Annual Computer Security Applications Conference (ACSAC '11), DOI 10.1145/2076732.2076752, December 2011.
- [Dimitropoulos09] Dimitropoulos, X., et al., "Graph annotations in modeling complex network topologies", ACM Transactions on Modeling and Computer Simulation (TOMACS), vol. 19, no. 4, DOI 10.1145/1596519.1596522, November 2009.
- [DONA] Koponen, T., et al., "A Data-Oriented (and Beyond) Network Architecture", Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications (SIGCOMM '07), ACM, DOI 10.1145/1282380.1282402, 2007.
- [EMULAB] Eide, E., et al., "An Experimentation Workbench for Replayable Networking Research", Proceedings of the 4th USENIX conference on Networked systems design & implementation (NSDI '07), 2007.
- [Fotiou12] Fotiou, N., et al., "Access control enforcement delegation for information-centric networking architectures", Proceedings of the second edition of the ICN workshop on Information-centric networking (ICN '12), ACM, DOI 10.1145/2342488.2342507, 2012.
- [Fotiou14] Fotiou, N., et al., "A framework for privacy analysis of ICN architectures", Proc. Second Annual Privacy Forum (APF), Springer, DOI 10.1007/978-3-319-06749-0_8, 2014.
- [Fri12] Fricker, C., Robert, P., Roberts, J. and N. Sbihi, "Impact of traffic mix on caching performance in a content-centric network", 2012 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPs), Orlando, USA, DOI 10.1109/INFOCOMW.2012.6193511, March 2012.
- [Goog08] Google, "Official Google Blog: We knew the web was big...", July 2008, <<http://googleblog.blogspot.it/2008/07/we-knew-web-was-big.html>>.
- [Guo07] Guo, L., Chen, S., Xiao, Z., Tan, E., Ding, X., and X. Zhang, "A performance study of BitTorrent-like peer-to-peer systems", IEEE Journal on Selected Areas in Communication, vol. 25, no. 1, pp. 155-169, DOI 10.1109/JSAC.2007.070116, 2007.

- [HASHROUT] Saino, L., Psaras, I., and G. Pavlou, "Hash-routing Schemes for Information-Centric Networking", Proceedings of the 3rd ACM SIGCOMM workshop on Information-centric networking (ICN '13), DOI 10.1145/2491224.2491232, 2013.
- [Hefeeda08] Hefeeda, M. and O. Saleh, "Traffic Modeling and Proportional Partial Caching for Peer-to-Peer Systems", IEEE/ACM Transactions on Networking, vol. 16, no. 6, pp. 1447-1460, DOI 10.1109/TNET.2008.918081, 2008.
- [ICARUS] Saino, L., Psaras, I., and G. Pavlou, "Icarus: a Caching Simulator for Information Centric Networking (ICN)", Proceedings of the 7th International ICST Conference on Simulation Tools and Techniques (SimuTools '14), DOI 10.4108/icst.simutools.2014.254630, 2014.
- [Detti12] Detti, A., et al., "Supporting the Web with an Information Centric Network that Routes by Name", Elsevier Computer Networks, vol. 56, no. 17, DOI 10.1016/j.comnet.2012.08.006, November 2012.
- [ICNSIMS] Tortelli, M., et al., "CCN Simulators: Analysis and Cross-Comparison", Proceedings of the 1st international conference on Information-centric networking (ICN '14), ACM, DOI 10.1145/2660129.2660133, 2014.
- [IMB2014] Imbrenda, C., Muscariello, L., and D. Rossi, "Analyzing Cacheable Traffic in ISP Access Networks for Micro CDN Applications via Content-Centric Networking", Proceedings of the 1st international conference on Information-centric networking (ICN '14), DOI 10.1145/2660129.2660146, 2014.
- [Ion13] Ion, M., Zhang, J., and E. Schooler, "Toward content-centric privacy in ICN: attribute-based encryption and routing", Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM (SIGCOMM '13), ACM, DOI 10.1145/2486001.2491717, 2013.
- [Jacobson09] Jacobson, V., et al., "Networking Named Content", Proceedings of the 5th international conference on Emerging networking experiments and technologies (CoNEXT '09), DOI 10.1145/1658939.1658941, 2009.

- [Katsaros12] Katsaros, K., Xylomenos, G., and G. Polyzos, "GlobeTraff: a traffic workload generator for the performance evaluation of future Internet architectures", 2012 5th International Conference on New Technologies, Mobility and Security (NTMS), DOI 10.1109/NTMS.2012.6208742, 2012.
- [Katsaros15] Katsaros, K., et al., "On the Inter-domain Scalability of Route-by-Name Information-Centric Network Architectures", Proc. IFIP Networking Conference, DOI 10.1109/IFIPNetworking.2015.7145308, 2015.
- [Kaune09] Kaune, S. et al., "Modelling the Internet Delay Space Based on Geographical Locations", 17th Euromicro International Conference on Parallel, Distributed and Network-based Processing, Weimar, Germany, DOI 10.1109/PDP.2009.44, 2009.
- [Labovitz10] Labovitz, C., Iekel-Johnson, S., McPherson, D., Oberheide, J., and F. Jahanian, "Internet inter-domain traffic", In Proceedings of the ACM SIGCOMM 2010 conference (SIGCOMM DOI 10.1145/1851182.1851194, 2010.
- [Lauinger10] Lauinger, T., "Security and Scalability of Content-Centric Networking", Masters Thesis, Technische Universitaet Darmstadt and Eurecom, September 2010.
- [Lauinger12] Lauinger, Y., et al, "Privacy Risks in Named Data Networking: What is the Cost of Performance?", ACM SIGCOMM Computer Communication Review, Vol. 42, Issue 5, DOI 10.1145/2378956.2378966, 2012.
- [Led12] Lederer, S., Muller, C., and C. Timmerer, "Dynamic adaptive streaming over HTTP dataset", Proceedings of the ACM Multimedia Systems Conference (MMSys '12), pp. 89-94, DOI 10.1145/2155555.2155570, 2012.
- [Lewko11] Lewko, A. and B. Waters, "Decentralizing attribute-based encryption", Proc. of EUROCRYPT 2011, Lecture Notes in Computer Science (LNCS), vol. 6632, pp. 568-588, DOI 10.1007/978-3-642-20465-4_31, 2011.

- [LIRA] Psaras, I., Katsaros, K., Saino, L., and G. Pavlou, "Lira: A location independent routing layer based on source-provided ephemeral names", Electronic and Electrical Eng. Dept., UCL, London, UK, Tech. Rep. 2014, <<http://www.ee.ucl.ac.uk/comit-project/publications.html>>.
- [Mahanti00] Mahanti, A., Williamson, C., and D. Eager., "Traffic analysis of a web proxy caching hierarchy", IEEE Network, Vol. 14, No. 3, pp. 16-23, DOI 10.1109/65.844496, May/June 2000.
- [Maier09] Maier, G., Feldmann, A., Paxson, V., and M. Allman, "On dominant characteristics of residential broadband internet traffic", In Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference (IMC '09), New York, NY, USA, 90-102. DOI 10.1145/1644893.1644904, 2009.
- [Marciniak08] Marciniak, P., Liogkas, N., Legout, A., and E. Kohler, "Small is not always beautiful", In Proc. of IPTPS, International Workshop of Peer-to-Peer Systems, Tampa Bay, Florida (FL), USA, February 2008.
- [MiniCCNx] Cabral, C., et al., "High fidelity content-centric experiments with Mini-CCNx", 2014 IEEE Symposium on Computers and Communications (ISCC), DOI 10.1109/ISCC.2014.6912537, 2014.
- [Montage] Hussain, A. and J. Chen, "Montage Topology Manager: Tools for Constructing and Sharing Representative Internet Topologies", DETER Technical Report, ISI-TR-684, August 2012.
- [Muscariello11] Muscariello, L., Carofiglio, G., and M. Gallo, "Bandwidth and storage sharing performance in information centric networking", Proceedings of the ACM SIGCOMM workshop on Information-centric networking (ICN '11), DOI 10.1145/2018584.2018593, 2011.
- [ndnSIM] Afanasyev, A., et al., "ndnSIM: NDN simulator for NS-3", NDN Technical Report NDN-0005, Revision 2, October 2012.
- [ndnSIM2] Mastorakis, S., et al., "ndnSIM 2.0: A new version of the NDN simulator for NS-3", NDN Technical Report NDN-0028, Revision 1, January 2015.

- [NEPI] Quereilhac, A., et al., "NEPI: An integration framework for Network Experimentation", 2011 19th International Conference on Software, Telecommunications and Computer Networks (SoftCOM), IEEE, 2011.
- [OFELIA] Sune, M., et al., "Design and implementation of the OFELIA FP7 facility: The European OpenFlow testbed", Computer Networks, vol. 61, pp. 132-150, DOI 10.1016/j.bjp.2013.10.015, March 2014.
- [ONL] DeHart, J., et al., "The open network laboratory: a resource for networking research and education", ACM SIGCOMM Computer Communication Review (CCR), vol. 35, no. 5, pp. 75-78, DOI 10.1145/1096536.1096547, 2005.
- [Parisis13] Parisis, G., Trossen, D., and H. Asaeda, "A Node Design and a Framework for Development and Experimentation for an Information-Centric Network", IEICE Transactions on Communications, vol. E96-B, no. 7, pp. 1650-1660, July 2013.
- [Perino11] Perino, D. and M. Varvello, "A Reality Check for Content Centric Networking", Proceedings of the ACM SIGCOMM workshop on Information-centric networking (ICN '11), DOI 10.1145/2018584.2018596, 2011.
- [PLANETLAB] Chun, B., et al., "Planetlab: an overlay testbed for broad-coverage services", ACM SIGCOMM Computer Communication Review (CCR), vol. 33, no. 3, pp. 3-12, DOI 10.1145/956993.956995, 2003.
- [PRST4.5] Riihijarvi, J., et al., "Final Architecture Validation and Performance Evaluation Report", PURSUIT Project Deliverable D4.5, January 2013.
- [Psaras11] Psaras, I., Clegg, R., Landa, R., Chai, W., Pavlou, G., "Modelling and Evaluation of CCN-Caching Trees", Proceedings of the 10th international IFIP TC 6 conference on Networking, Valencia, Spain, May 2011.
- [Psaras12] Psaras, I., Chai, W., and G. Pavlou, "Probabilistic In-Network Caching for Information-Centric Networks", Proceedings of the second edition of the ICN workshop on Information-centric networking (ICN '12), DOI 10.1145/2342488.2342501, 2012.

- [Quereilhac14] Quereilhac, A., et al., "Demonstrating a unified ICN development and evaluation framework", Proceedings of the 1st international conference on Information-centric networking (ICN '14), ACM, DOI 10.1145/2660129.2660132, 2014.
- [Renault09] Renault, E., Ahmad, A., and M. Abid, "Toward a Security Model for the Future Network of Information", Proceedings of the 4th International Conference on Ubiquitous Information Technologies & Applications (ICUT '09), IEEE, DOI 10.1109/ICUT.2009.5405676, 2009.
- [RFC2330] Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics", RFC 2330, DOI 10.17487/RFC2330, May 1998, <<http://www.rfc-editor.org/info/rfc2330>>.
- [RFC5743] Falk, A., "Definition of an Internet Research Task Force (IRTF) Document Stream", RFC 5743, DOI 10.17487/RFC5743, December 2009, <<http://www.rfc-editor.org/info/rfc5743>>.
- [RFC6920] Farrell, S., Kutscher, D., Dannewitz, C., Ohlman, B., Keranen, A., and P. Hallam-Baker, "Naming Things with Hashes", RFC 6920, DOI 10.17487/RFC6920, April 2013, <<http://www.rfc-editor.org/info/rfc6920>>.
- [RFC7476] Pentikousis, K., Ed., Ohlman, B., Corujo, D., Boggia, G., Tyson, G., Davies, E., Molinaro, A., and S. Eum, "Information-Centric Networking: Baseline Scenarios", RFC 7476, DOI 10.17487/RFC7476, March 2015, <<http://www.rfc-editor.org/info/rfc7476>>.
- [RFC7927] Kutscher, D., Ed., Eum, S., Pentikousis, K., Psaras, I., Corujo, D., Saucez, D., Schmidt, T., and M. Waehlis, "Information-Centric Networking (ICN) Research Challenges", RFC 7927, DOI 10.17487/RFC7927, July 2016, <<http://www.rfc-editor.org/info/rfc7927>>.
- [RFC7933] Westphal, C., Ed., Lederer, S., Posch, D., Timmerer, C., Azgin, A., Liu, W., Mueller, C., Detti, A., Corujo, D., Wang, J., Montpetit, M., and N. Murray, "Adaptive Video Streaming over Information-Centric Networking (ICN)", RFC 7933, DOI 10.17487/RFC7933, August 2016, <<http://www.rfc-editor.org/info/rfc7933>>.

- [SAIL-B2] SAIL, "NetInf Content Delivery and Operations", SAIL Project Deliverable D-B.2, May 2012.
- [SAIL-B3] Kutscher, D., Ed., et al., "Final NetInf Architecture", SAIL Project Deliverable D-B.3, January 2013, <<http://www.sail-project.eu/deliverables/>>.
- [Saino13] Saino, L., Cocora, C., and G. Pavlou, "A Toolchain for Simplifying Network Simulation Setup", Proceedings of the 6th International ICST Conference on Simulation Tools and Techniques (SimuTools '13), 2013.
- [Saleh06] Saleh, O., and M. Hefeeda, "Modeling and caching of peer-to-peer traffic", Proceedings of the 2006 IEEE International Conference on Network Protocols (ICNP), DOI 10.1109/ICNP.2006.320218, 2006.
- [Salsano12] Salsano, S., et al., "Transport-Layer Issues in Information Centric Networks", Proceedings of the second edition of the ICN workshop on Information-centric networking (ICN '12), ACM, DOI 10.1145/2342488.2342493, 2012.
- [Salsano13] Salsano, S., et al., "Information Centric Networking over SDN and OpenFlow: Architectural Aspects and Experiments on the OFELIA Testbed", Computer Networks, vol. 57, no. 16, pp. 3207-3221, DOI 10.1016/j.comnet.2013.07.031, November 2013.
- [SensReqs] Karnouskos, S., et al., "Requirement considerations for ubiquitous integration of cooperating objects", 2011 4th IFIP International Conference on New Technologies, Mobility and Security (NTMS), DOI 10.1109/NTMS.2011.5720605, 2011.
- [Smetters09] Smetters, D., and V. Jacobson, "Securing network content", Technical Report TR-2009-01, PARC, 2009.
- [Sourlas15] Sourlas, V., Tassiulas, L., Psaras, I., and G. Pavlou, "Information Resilience through User-Assisted Caching in Disruptive Content-Centric Networks", 14th IFIP Networking Conference, Toulouse, France, DOI 10.1109/IFIPNetworking.2015.7145301, May 2015.

- [Tagger12] Tagger, B., et al., "Update on the Architecture and Report on Security Analysis", Deliverable 2.4, PURSUIT EU FP7 project, April 2012.
- [Tourani16] Tourani, R., Mick, T., Misra, S., and G. Panwar, "Security, Privacy, and Access Control in Information-Centric Networking: A Survey", arXiv:1603.03409, March 2016.
- [VoCCN] Jacobson, V., et al., "VoCCN: Voice-over Content-Centric Networks", Proceedings of the 2009 workshop on Re-architecting the internet (ReArch '09), DOI 10.1145/1658978.1658980, 2009.
- [Watts98] Watts, D. J. and S. H. Strogatz, "Collective dynamics of 'small-world' networks", Nature, vol. 393, no. 6684, pp. 440-442, DOI 10.1038/30918, April 1998.
- [Yu06] Yu, H., Zheng, D., Zhao, B., and W. Zheng, "Understanding user behavior in large-scale video-on-demand systems", ACM SIGOPS Operating Systems Review - Proceedings of the 2006 EuroSys conference, Vol. 40, Issue 4, pp. 333-344, DOI 10.1145/1218063.1217968, April 2006.
- [Zhang10a] Zhang, C., Dhungel, P., Wu, D., and K. Ross, "Unraveling the BitTorrent Ecosystem", IEEE Transactions on Parallel and Distributed Systems, vol. 22, issue 7, pp. 1164-1177, DOI 10.1109/TPDS.2010.123, 2010.
- [Zhang10b] Zhang, L., et al., "Named Data Networking (NDN) Project", NDN Technical Report NDN-0001, October 2010, <<http://named-data.net/publications/techreports/>>.
- [Zhou11] Zhou, J., Li, Y., Adhikari, K., and Z.-L. Zhang, "Counting YouTube videos via random prefix sampling", Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference (IMC '11), Berlin, Germany, DOI 10.1145/2068816.2068851, November 2011.

Acknowledgments

Konstantinos Katsaros contributed the updated text of Section 2.2 along with an extensive set of references.

Priya Mahadevan, Daniel Corujo, and Gareth Tyson contributed to a 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, E. Baccelli, Claudia Campolo, Christian Esteve Rothenberg, Suyong Eum, Nikos Fotiou, Dorothy Gellert, Luigi Alfredo Grieco, Myeong-Wuk Jang, Ren Jing, Will Liu, Antonella Molinaro, Luca Muscariello, Ioannis Psaras, Dario Rossi, Stefano Salsano, Damien Saucez, Dirk Trossen, Jianping Wang, Yuanzhe Xuan, and Xinwen Zhang.

The IRSG review was provided by Aaron Falk.

Authors' Addresses

Kostas Pentikousis (editor)
Travelping
Koernerstr. 7-10
10785 Berlin
Germany

Email: k.pentikousis@travelping.com

Borje Ohlman
Ericsson Research
S-16480 Stockholm
Sweden

Email: Borje.Ohlman@ericsson.com

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

Email: daviese@scss.tcd.ie

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

Email: spis@intracom-telecom.com

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

Email: g.boggia@poliba.it