

On Information Exposure through Named Content

Kostantinos V. Katsaros, Lorenzo Saino, Ioannis Psaras, George Pavlou
Dept. of Electrical & Electronic Engineering
University College London
WC1E 7JE, Torrington Place, London, UK
Email: {k.katsaros, l.saino, i.psaras, g.pavlou}@ucl.ac.uk

Abstract—The proposed shift from host-centric to information-centric networking (ICN) has triggered extensive research in the area of content naming. Efforts have so far focused on the scalability and security properties that can make content objects routable and self-certifying. In this paper, we argue that the information that is exposed through explicitly naming content objects has been overlooked, although several operational and performance issues depend on the information that a name holds. We therefore revisit content naming design decisions taking into account information exposure and deployability of the ICN paradigm.

I. INTRODUCTION

Arguably, the main principle underpinning Information Centric Networking (ICN) is the shift of the main networking abstraction from *node identifiers to location-agnostic content identifiers*. Generally speaking, *content* in the ICN area is defined in terms of *chunks*, which need to be routable and self-certifying [1], [2]. The *content name* or *identifier* is then used to route contents, store them and look them up in in-network caches (i.e., not only in overlay, hierarchical caches or CDNs) and validate their origin.

A large corpus of literature focused on this specific aspect, that is, on how to name content objects, how to route them or how to resolve them to routable entities [1], [3], [4]. Research on naming so far has tackled the problem from several angles:

- **Security:** design names that allow for easy verification of content objects integrity, as well as for authentication of the content source [2], [5], [6].
- **Routability:** design names that allow (through either coupled or decoupled approaches) routing by name [1], [5], [4].
- **Scalability:** design names that can be aggregated or that do not require excessively large routing tables [7].
- **Extensibility:** design names with a flexible format supporting future adaptation (e.g., [8]).

In this paper we revisit ICN naming to investigate the issue of *information exposure* through directly and explicitly naming content objects. With the term information exposure we refer to the amount of information about a content object that is revealed to network entities by its name and/or through the name resolution process. We contend that information exposure aspects should be carefully considered in the design of naming formats, name resolution and routing mechanisms, as a number of desirable and undesirable features depend on this information. We argue that this perspective of directly named, self-certifying content has largely been ignored by the research community so far.

Looking closer at the properties of the information elements that can be accommodated in a content name, we find that these properties comprise both desirable and undesirable groups. Censorship for instance seems much easier if the source identity is revealed in the content name - this for example is the case with CCN/NDN permanent, hierarchical, human-readable names [5]. Information scoping, on the other hand (e.g., message delivery within or outside given boundaries) might prove very helpful in case of mobile, infrastructureless and opportunistic networking environments.

In this paper we take a closer look at different networking environments that can potentially benefit from ICN, with the purpose of identifying the functional requirements they impose regarding information exposure (Section II). Based on this analysis, we then identify the types of information that should or should not be exposed by content names and the name resolution process, elaborating on the benefits and downsides of exposing too much or too little information in content names (Section III). We then discuss several important aspects of information exposure, related to the role and actions of each actor in the considered networking environment, as well as the potential risks of malicious information exposure along with possible countermeasures (Section IV). Taking a step further, we elaborate on a series of technical aspects related to the realisation of information exposure in the context of ICN, highlighting potential pitfalls in the design of a semantically rich naming scheme and name resolution process (Section V). Finally, we present our conclusions in Section VI.

II. MOTIVATION, BACKGROUND AND REQUIREMENTS

Although information-centricity and location-independent content resolution and delivery have been primarily proposed for the fixed part of the network [5], [9] and especially for content distribution applications, they have been shown to also provide clear benefits in alternative environments. For instance, host centric, IP-based communication, has been repeatedly shown to be a poor fit for mobile environments, while a content-centric, *request-response* model seems to be meeting the requirements of client mobility [10] and network fragmentation [11]. Examples of other networking environments include machine-to-machine, smart grid applications [12], the *Internet of Things* (IoT) [13], vehicular [14] or home networking [15].

The diversity of these networking environments and applications results in a corresponding diversity in the requirements from the content names and the information these expose. This is because of the significantly different functional requirements of each environment. Below we elaborate on these require-

ments so as to subsequently identify the information that needs to be exposed (or not) for their support.

A. Content Distribution

Permanent content names have been adopted by all ICN proposals/architectures so far to enable and support name-based routing. Permanent names also allow for transparent (to the content provider (CP)) in-network caching at the ISP. Although serving content from caches located closer to users is a desirable feature for CDNs and CPs, doing so transparently (i.e., without the CDN or CP being aware of such actions) is not an attractive feature for them. The reason is that CDNs and CPs need to log information describing user access to their content so as to be able to assess the (spatiotemporal) popularity of their content, bill their customers, adapt their content to user needs and support personalised/targeted advertising. Therefore, *logging* at the CDN or CP side is an essential feature of any viable ICN proposal and is related to the information exposed through the content name and the name resolution process.

At the same time, exposing the identity of the CP can lead to violation of *content neutrality*, i.e., not discriminating (and assigning link resources) among different contents based on ownership (e.g., news from CP A vs. CP B) [4].

Furthermore, distributed and uncoordinated in-network content caching [16] can lead to stale content staying in caches, while updated versions have been published by the content owner. Cache purging through TTL-like techniques is the most straightforward way to deal with stale content, although its effectiveness is questionable [4], due to tradeoffs between freshness and cache hits. This information is again an attribute that should be exposed through explicit content naming.

Moreover, as network operators enter the content distribution market [17], there is an increasing need for intelligent forwarding and/or caching of delivered content through the differentiation of content/traffic types. This characterisation of content is also subject to the information exposed by content names.

B. Mobile, Opportunistic Networks

Infrastructureless environments have attracted wide attention by the community during the past decade [18], [19]. The importance of communication in absence of Internet connectivity might become of vital importance in emergency or disaster situations. In such cases, although some connectivity might still exist, networks often become fragmented. Local communication with neighbouring nodes (e.g., [20]) when networks get fragmented cannot be supported by the current host-centric communication model [11].

Information exposure through named content can assist in the replication or the *store, carry and forward* operations in opportunistic networks. As we showed in [11] for instance, time and space scoping can provide significant performance benefits. Name-based routing in case of infrastructureless environments takes the form of name-based replication, where messages are forwarded based on the information included in the message name.

C. IoT, Smart Cities/Grids and Vehicular Networks

In contrast to the scenarios discussed above, where communication is flat in the sense that information is transmitted

between any two devices, Smart City [21], Internet of Things [13] and Vehicular Networking [14] environments may pose further requirements in terms of end-devices. The information exchanged between two home appliances is radically different from the information exchanged between two vehicles on the highway. Furthermore, the information sent between two home appliances cannot be of interest to anyone (or anything) else outside that household, making requirements for caching vanish. Such messages will probably need to carry the identifier of the device within the message name in order to assist in the routing and delivery process.

Smart City applications will also require attention as to the information that messages will expose through names. For instance, application messages should carry spatial and temporal locality attributes to avoid them being forwarded to other geographic destinations e.g., the control (or frequency alteration) of traffic lights within a city based on ambulance sirens, or on real-time road traffic needs to propagate only locally and should not escape to other areas within the city. Similarly, energy requests of electric vehicles, in the context of smart grids, should be annotated so as to reach only nearby charging stations.

III. EXPOSED INFORMATION

A. From content identifiers to content characterisation

In the following we elaborate further on the concept of *information exposure* introduced in Section I. We expand the scope of information to be exposed through content naming beyond the mere identification of content items and their origin¹. This effort is driven by the need to differentiate treatment of delivered content in various application environments. We identify a (non-exhaustive) list of content attributes that leads us towards the *characterisation* of information delivered by information-centric networks.

Service type. It provides a characterisation of the corresponding traffic e.g., in terms of elasticity. Example types are *voip*, or *video/audio* for telephony and *video* or *music streaming* respectively, *web* for short transfers and *fts* for long file transfer services (e.g., software update). Exposing this information enables link scheduling and multiplexing according to the traffic handled. For instance, different Active Queue Management (AQM) schemes may be applied depending on the type of traffic handled.

Ownership. Information can also be characterised by ownership. So far, work in this area has focused on the support of authenticity (or provenance), i.e., enabling users to verify that the received content came from the original source [9], [3]. However, recent research efforts have pointed out the need for supporting *content neutrality* (see Section II) [4].

Caching properties. Content names may also be characterised with respect to their prospective treatment by caches. First, transferred content can bear semantics regarding its *cacheability*, i.e., whether it should be cached by in-network caches or not. For instance, software updates is a typical example of content that is subject to caching, as opposed to live-streaming or telephony. Such information may also include

¹As discussed in Section V, whether the identified information should be included in the content names or not is still an open issue. In the following we follow the current convention, referring to information exposed through content names.

(geographic) scoping semantics as discussed later on. For example, caching the front page of a Greek news web site may or may not be preferred in a network in the UK exposing the trade-off between the efficiency of cache space utilisation and inter-domain traffic costs. Based on the exposed information, ISPs may apply any decision best fitting their goals. Moreover, caching-related information could further support the active control of cached copies of content in terms of *cache purging*, i.e., similarly to cacheability, content providers could indicate whether some version of their content is stale and can therefore be evicted from in-network caches.

Service class. Based on this information, routers can identify the content that is under a Service Level Agreement (SLA) and hence subject to preferential treatment. Given the concerns on ownership and content neutrality in the transport domain, ISPs can leverage their caching infrastructure for the provision of such services e.g., through preferential cache space allocation [22]. Note that this type of information goes beyond caching properties, which simply denote the desired caching behaviour of in-network caches that comes with no guarantees.

Scope. Limiting the reachability of content may be desirable in several different environments. Information can be scoped based on various criteria [8]. For example, it could be scoped in terms of Autonomous Systems (AS) to be included/excluded in the delivery of the content [23]. Time, as well as the geographic location of nodes could also represent scoping criteria in the context of smart cities/grids and opportunistic networks as well. For instance, in the latter case, information regarding the availability of emergency medical help, in the event of a physical disaster, should not be propagated beyond the impacted area, while short (ad-hoc) text messages should be characterised by a relatively short time-to-live.

Content format. The proliferation of mobile devices is characterised by substantial heterogeneity, i.e., a recent analysis of an Internet mobile streaming service reported “3,400 hardware models with 109 screen resolutions running 14 mobile OS, 3 audio codecs and 4 video codecs” [24]. Obviously the format of the content may be inadequate for certain devices (e.g., incompatible codecs, very low resolution) and/or it may waste network resources (e.g., streaming high resolution video to low resolution devices). Adapting content through in-network transcoding has been considered as a solution (e.g., [25]), however it goes against the location-independence property of ICN as it bounds content to transcoding host locations. It therefore becomes important to consider the exposure of the content format for named contents in order to support the adaptation of content to device characteristics.

B. Exposing information through name resolution

We identified above a set of information elements that characterise content and the corresponding service that can be explicitly exposed through the naming scheme. However, the semantics of a naming scheme cannot expose information related to content popularity and user access statistics in general. Even though this information is obviously important for content providers (see Section II), it is *hidden* from them in most ICN designs. The reason is that user requests for content may be served by any intermediate node either through caching or multicast (i.e., CCN/NDN [5]), or it can be served

by an ISP-operated name resolution system (i.e., DONA [9], PURSUIT [26], NetInf [27], COMET [23]). In both cases, the content provider is not involved in the name resolution procedure thus losing access to content access information. As a consequence, content providers are disincentivised to adopt ICN [4], despite the performance benefits brought by inherent ICN mechanisms such as caching and multicast [28].

In effect, we argue that content access information should be exposed to content providers. To this end, we foresee two potential approaches. The first approach is tailored for architectures that decouple name resolution from routing and forwarding ([9], [26], [27], [23]). This decoupling adheres to the “design for tussle” principle [29], thus enabling content providers to potentially establish an interface with the name-resolution system for logging purposes. The second approach foresees the direct involvement of the CPs, i.e., as detailed in [4], CPs can be responsible for directly providing users with the name/identifier of the content they are interested in, which is subsequently used by intermediate routers to reach the content (i.e., name-based routing). Though introducing additional complexity at the content provider, this approach comes with a series of benefits including the support of content provider-based cache purging (see Section V), the support for content neutrality and the re-use of existing network infrastructure and signalling protocols [4].

C. Ephemeral names

Another aspect of the information exposed through content names and name resolution in ICN, relates to its lifetime. When CPs are directly involved in the name resolution process, as discussed above, a mechanism is required to force each content request to reach the CP so as to be logged. This can be enabled by the use of *ephemeral* names, i.e., content names are transient and their value is controlled (i.e., changed) by the CP [4]. In effect, users need to reach the CP to get the current identifier for the requested content. Note that permanent names cannot satisfy the logging requirement as their value can be cached and subsequently reused without the involvement of the CP [4].

Though as discussed above logging can also be supported in other ways, the adoption of ephemeral names enables additional desirable features. First, since a routable identifier can be known only after a resolution, cached contents can be withdrawn/updated by content providers (i.e., cache purging) at any time simply by resolving subsequent requests to a new ephemeral name. Additionally, frequent name changes and the decoupling from the content provider’s identity make it prohibitive for operators to discriminate content based on the content provider without their consent, hence enabling *content neutrality*.

Moreover, it must be noted that the transient character of names plays an important role in the exposure of service class information (see Section III). As further elaborated in Section IV-B, ephemeral names can be used to mitigate the malicious use of service class information.

IV. HANDLING OF EXPOSED INFORMATION

In this section, we focus on analysing how network entities involved in processing named content can handle the information exposed by names. To this end, we next elaborate on practical aspects such as the possible value ranges and lifetime

of each type of exposed information element, the actors that manage the exposed information, issues related to inter-domain networking and the risks stemming from potential malicious usage of this information.

A. Handling service type information

As mentioned in Section III, we foresee a limited range of fixed service type values, such as `voip` and `fts`. These values would be subject to standardisation (e.g., IANA) as is the case now with Internet Media Types² [30], so as to ensure cross-domain compatibility of applications. Obviously, service type assignment would be made by content providers and have a permanent character.

The risk of content providers setting an inaccurate service type value to maliciously gain in performance is limited. For example, if a content provider labelled bulk download traffic as real-time traffic, it would achieve lower latency but also likely lower throughput and hence worse user experience. On the contrary, service providers are incentivised to set correct service types so that all involved ISPs in a session handle the content optimally.

B. Handling service class information

The service class information is also set by the content provider, however in this case a closer interaction with ISPs is required. To support premium services, ISPs are expected to configure accordingly their caches so as to identify premium content and treat it in a differentiated way (e.g., cache space allocation for a specific amount of time). In this respect, ISPs should communicate the service class information to be used to their customer content providers. The selected value range would need first to reflect the different classes of premium services provided by the ISP. However, service class information should be conveyed with ephemeral, non human readable names so as to mitigate free-riding, i.e., permanent service class values would enable malicious users to label their content accordingly in order to get premium treatment. At the same time however, apart from the frequent communication of the updated service class values to the content provider, the transient character of service class values would possibly also necessitate the frequent reconfiguration of caches so as to enable the identification of premium content. A potential approach to address these issues relies on the use of algorithmic identifiers [8], [2]. Instead of ISPs communicating each new ephemeral value of the service class identifier, they can communicate an algorithm through which content providers may produce subsequent ephemeral values themselves. In-network caches could be then configured in a similar manner. Moreover, since service class values are ISP specific, they obviously get invalidated when crossing domain borders. In this case, inter-domain quality of service (QoS) support would necessitate a mechanism for the negotiation of service class values, upon the establishment of the corresponding inter-domain Service Level Agreement.

C. Handling ownership information

As discussed in [3], the exposure of ownership information depends on the selected naming scheme. With self-certifying

names, the binding between the content name and the owner identity must be established by an external authority, while with human-readable hierarchical names the binding is intrinsic in the name. In this respect, self-certifying names are preferred, as they do not expose the identity of the owner at the network layer, thus supporting content neutrality. However, users should still be able to verify the provenance of the received content. To this end, content-based name resolution as described in [4] can constitute a viable approach. Namely, in a two-level name resolution process, users first resolve human readable identifiers to (ICN) network layer self-certifying ones. This resolution step is carried out involving the CP and enables the verification of provenance due to the human-readable hierarchical nature of the name³. Subsequently, self-certifying identifiers can be used to guide forwarding decisions inside the network.

D. Handling caching properties

Exposing this information enables content providers to define the desired behaviour they would like their content to receive from the in-network caches operated by ISPs. Content cacheability, would be simply expressed by a boolean value of a permanent character as this depends on whether it is meaningful to cache an item or not. However, as discussed in Section III, scoping information could also be included. Spatial scopes can be represented in several different ways such as a radius around a set of coordinates `<type=circle; pos=x,y;radius=r>`, or global hierarchical map format, e.g., `<country/state/city/postal-code>` [11]. Temporal scoping is usually expressed through some time-to-live (TTL) value. When no explicit cache purging mechanism is available then it can be an absolute time value e.g., `Thu, 01 Aug 2014 16:00:00 GMT`. Coarse grained synchronisation of caches is required in this case.

Cache purging can be supported in the simplest case by just using ephemeral names. The life-time of an ephemeral name should be selected by the content provider, subject to the lifecycle of content versions. However, as ephemeral names also allow content providers to mitigate censorship, their values should be changed with a frequency that strikes a balance between the time needed by ISPs to classify their content and potentially throttle it and the time needed to make in-network caching and multicast efficient. Taking a step further, an advanced form of cache purging could also include an explicit reference to the identifier/name of the previous version of the content that is to be purged from the cache reducing the time stale content resides in caches [4].

On another front, the exposure of caching properties raises some important security concerns. First, malicious users/content providers can frequently change the name of the content so as to fill the caches with stale data, i.e., each new name obsoletes the previous one, resulting in all previous versions of the content to unnecessarily occupy the cache space until they get evicted. This can be avoided with the explicit purging of the old version. However, in this approach it is important to make sure that malicious users cannot achieve the purging of cached content they do not own, i.e., by first retrieving the content identifier of the targeted content. In this case, the cache purging mechanism should allow for the

²We note though that Internet Media Types do not present semantics related to the behaviour and requirements of the corresponding traffic so it is not considered adequate candidates for expressing service types.

³Such names can become available through search engines.

verification of the common ownership of both the new and the old version of the content. It should be noted though, that such functionality results in additional processing overheads that would obstruct line-speed operation.

E. Handling scoping information

As discussed previously, scoping information can have different forms. When information dissemination is scoped in terms of network areas, we can envision the naming information to refer to the included/excluded areas in the form Autonomous System (AS) numbers, along with an indicator of the filtering or scoping function [23]. Geographic and temporal scopes could be expressed as discussed in Section IV-D.

At this point it is important to note the distinction between receiver-based and sender-based scoping in the context of push-based or pull-based communication models. While ICN has in principle shifted towards receiver-driven, pull-based communications, scoping of information dissemination may in some scenarios be placed in the hands of senders, in push-based applications (e.g., emergency signals for medical help). This obviously raises concerns for denial of service (DoS) attacks based on the flooding of the network.

F. Handling content format information

Content format information should be expressive enough to reflect the current device heterogeneity (see Section III). Similarly to the case of service types (see Section IV-A), the allowed values would be subject to standardisation (e.g., by IANA) as is the case now with Internet Media Types. Content providers should choose the correct value to assign a permanent semantic annotation to each version of their content, describing attributes such as screen resolution, audio/video codec and audio/video bitrate.

It is noted that improper use of this information would result in the wrong format to be served to the users. Serving a lower quality version of the content would obviously increase the disutility of users, thus providing no benefit for a misbehaving content provider. Serving higher quality content, on the other hand, would be expected to translate to the overconsumption of network resources with a potential impact on the charges of the (wireless) network operator to the content providers and/or the users. In this respect, content providers would be disincentivised to deliver content of higher quality format.

V. IMPLEMENTATION IMPLICATIONS

Realising the exposure of information as discussed in previous sections has important implications on the implementation level. Here we make some observations regarding the challenges posed by the exposure of information on this front.

1) *Information exposed in content names results in excessive header sizes:* It becomes obvious from the previous discussion that several types of information can be exposed through content names so as to enable a wide set of network services. As this information is eventually to be handled by an information-centric network, its footprint can become a concern, i.e., ensuring the uniqueness of the content identifiers via an appropriately selected length and at the same time conveying rich meta-data is expected to result in large-size headers which would consume a significant part of the available bandwidth. For instance, the default header overhead in the

CCNx implementation is 650 bytes [5]. Further considering, for instance, 3 bits for the representation of 5 different service types, 8 bits to represent different service classes⁴, 3 bytes to represent the content format⁵, at least 8 bytes for temporal scoping⁶ and at least 12 bytes for spatial scoping⁷, we get to a conservative estimation of an approximate additional 25 byte overhead, leading to a total header size of 675 bytes. Though large part of the overall overhead is due to security information (i.e., signature to authenticate the content), an increase of 25 bytes may be significant, especially in environments where packets may have considerably low size, i.e., sensor readings in the IoT.

2) *Variable lengths of exposed information do not facilitate line speed operations:* For instance, service class information is of no use in the case of opportunistic, infrastructure-less networks, or certain content providers may not have engaged in a premium service with their ISPs. In this respect, the overall information to be exposed can vary both across and within different networking environments. Dynamically adapting the length of the corresponding content identifiers/names would complicate line speed operations usually performed by hardware on fixed length header fields. In this respect, a trade-off emerges between the expressiveness of the naming scheme and the resulting header overheads.

3) *Utilising information exposed in content names increases the processing load at in-network devices:* The purpose of exposing this additional information in content names is obviously the support of enhanced in-network services such as link scheduling, premium cache services, scoped forwarding, media adaptation, etc. However, this comes at the cost of additional processing of the extended ICN packet header, which may challenge the quality and reliability of the network. For instance, a light-weight IoT device may not have the necessary resources (i.e., computational, energy) to reason and apply geographical scoping on a per packet level.

A way to overcome this limitation is to expose the desired information through variable sets of content *attributes*, that are not included in content names. In this approach baseline name-based forwarding can be performed at line speed, as the size of the name can remain fixed, regardless the networking environment. At the same time, devices of limited capacity can simply ignore the attributes of the data, letting more intelligent forwarding and/or caching decisions to be made by more powerful devices, or simply devices that have the required resources (e.g., memory) at the time of forwarding (i.e., this may vary in time) [11]. However, due to the variable size of the attributes list, advanced forwarding and caching decisions cannot be supported at line speed.

VI. CONCLUSIONS

In this paper we have revisited design tradeoffs for content naming in ICNs. We have focused on the information exposed by content names and have argued that this information can

⁴An 8-bit Differentiated Services (DS) field is currently used in DiffServ [31]

⁵Note that results in [24] indicate the potential existence of more than 1000 combinations of screen resolution and audio/video codec.

⁶Here only considering the representation of a single deadline date for the forwarding of content.

⁷This corresponds to the representation of a 32-bit unsigned integer radius around a set of 32-bit unsigned integer GPS co-ordinates.

be crucial for the deployability and sustainability of a future information-centric Internet. Our discussion covered a range of networking environments as well as how the design properties of a name can be handled in order to enable or disable certain operations. We showed that name design decisions that do not take into account information exposure aspects can have the side effect of enabling undesirable features and disabling desirable ones. Further work is needed in order to crystallise the factors that need to be taken into account and realise them through a naming scheme for information-centric networks. We hope that this study serves as a stepping stone on this direction.

ACKNOWLEDGMENT

This work was supported by EPSRC UK (COMIT project) grant no. EP/K019589/1 and EU FP7/NICT (GreenICN project) grant no. (EU) 608518/(NICT)167.

REFERENCES

- [1] A. Detti, M. Pomposini, N. Blefari-Melazzi, and S. Salsano, "Supporting the web with an information centric network that routes by name," *Computer Networks*, vol. 56, no. 17, 2012.
- [2] W. Wong and P. Nikander, "Secure naming in information-centric networks," in *ACM ReArch 2010*, pp. 12:1–12:6.
- [3] A. Ghodsi, T. Koponen, J. Rajahalme, P. Sarolahti, and S. Shenker, "Naming in content-oriented architectures," in *ACM SIGCOMM ICN workshop*, New York, NY, USA, 2011, pp. 1–6.
- [4] I. Psaras, K. Katsaros, L. Saino, and G. Pavlou, "Lira: A location independent routing layer based on source-provided ephemeral names," 2014. [Online]. Available: <http://www.ee.ucl.ac.uk/~lsaino/publications/lira-techrep.pdf>
- [5] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard, "Networking named content," in *ACM CoNEXT 2009*, New York, NY, USA, pp. 1–12.
- [6] C. Dannewitz, T. Rautio, O. Strandberg, and B. Ohlman, "Secure naming structure and p2p application interaction," 2011.
- [7] C. Tsilopoulos, G. Xylomenos, and Y. Thomas, "Reducing forwarding state in content-centric networks with semi-stateless forwarding," in *Proceedings of IEEE INFOCOM'14*.
- [8] D. Trossen, B. Tagger, G. Parisi, M. Skjogstad, N. Fotiou, X. Vasilakos, C. Stais, Y. Thomas, M. Al-Naday, M. Reed, V. Vassilakis, A. Bontozoglou, B. Gajic, P. Sarolahti, P. Flegkas, and V. Sourlas, "D2.5 Final Updated Architecture," PSIRP Project, Tech. Rep., July 2010.
- [9] T. Koponen, M. Chawla, B.-G. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker, and I. Stoica, "A data-oriented (and beyond) network architecture," in *ACM SIGCOMM 2007*, pp. 181–192.
- [10] G. Tyson, N. Sastry, R. Cuevas, I. Rimac, and A. Mauthe, "A survey of mobility in information-centric networks," *Commun. ACM*, vol. 56, no. 12, pp. 90–98, Dec. 2013.
- [11] I. Psaras, L. Saino, M. Arumathurai, K. Ramakrishnan, and G. Pavlou, "Name-based replication priorities in disaster cases," in *IEEE NOM workshop 2014*.
- [12] K. V. Katsaros, W. K. Chai, N. Wang, G. Pavlou, H. Bontius, and M. Paolone, "Information-centric Networking for Machine-to-Machine Data Delivery - A Case Study in Smart Grid Applications," *IEEE Network Magazine*, to appear 2014.
- [13] A. Rayes, M. Morrow, and D. Lake, "Internet of things implications on icn," in *IEEE CTS*, 2012, pp. 27–33.
- [14] G. Grassi, D. Pesavento, G. Pau, R. Vuyyuru, R. Wakikawa, and L. Zhang, "Vanet via named data networking," in *IEEE INFOCOMM 2013 NOMEN Workshop*, 2014.
- [15] T. Biswas, A. Chakraborti, R. Ravindran, X. Zhang, and G. Wang, "Contextualized information-centric home network," in *ACM SIGCOMM 2013*, pp. 461–462.
- [16] I. Psaras, W. K. Chai, and G. Pavlou, "Probabilistic in-network caching for information-centric networks," in *ACM ICN Sigcomm workshop 2012*, pp. 55–60.
- [17] J. Wulf, R. Zarnekow, T. Hau, and W. Brenner, "Carrier activities in the cdn market - an exploratory analysis and strategic implications," in *Proc. of the 2010 International Conference on Intelligence in Next Generation Networks (ICIN)*, oct. 2010, pp. 1–6.
- [18] I. Psaras, L. Wood, and R. Tafazolli, "Delay-/disruption-tolerant networking: State of the art and future challenges," in *Technical Report, University of Surrey, 2009*. [Online]. Available: <http://www.ee.ucl.ac.uk/~uceedips/dtn-srv-ipsaras.pdf>
- [19] M. Khabbaz, C. Assi, and W. Fawaz, "Disruption-tolerant networking: A comprehensive survey on recent developments and persisting challenges," 2012.
- [20] J. Ott, E. Hyytia, P. Lassila, T. Vaegs, and J. Kangasharju, "Floating content: Information sharing in urban areas," in *IEEE PerCom*, 2011.
- [21] G. Piro, I. Cianci, L. A. Grieco, G. Boggia, and P. Camarda, "Information centric services in smart cities," *J. Syst. Softw.*, vol. 88, pp. 169–188, Feb. 2014.
- [22] P. Agyapong and M. Sirbu, "Economic incentives in information-centric networking: implications for protocol design and public policy," *Communications Magazine, IEEE*, vol. 50, no. 12, pp. 18–26, 2012.
- [23] W. K. Chai, N. Wang, I. Psaras, G. Pavlou, C. Wang, G. de Blas, F. Ramon-Salguero, L. Liang, S. Spirou, A. Beben, and E. Hadjioannou, "Curling: Content-ubiquitous resolution and delivery infrastructure for next-generation services," *IEEE Communications Magazine*, vol. 49, no. 3, pp. 112–120, march 2011.
- [24] Y. Liu, F. Li, L. Guo, B. Shen, S. Chen, and Y. Lan, "Measurement and analysis of an internet streaming service to mobile devices," *IEEE TPDS*, vol. 24, no. 11, pp. 2240–2250, Nov 2013.
- [25] R. Han, P. Bhagwat, R. LaMaire, T. Mummert, V. Perret, and J. Rubas, "Dynamic adaptation in an image transcoding proxy for mobile web browsing," *Personal Communications, IEEE*, 1998.
- [26] K. V. Katsaros, N. Fotiou, X. Vasilakos, C. N. Ververidis, C. Tsilopoulos, G. Xylomenos, and G. C. Polyzos, "On inter-domain name resolution for information-centric networks," in *IFIP Networking 2012*, Berlin, Heidelberg, pp. 13–26.
- [27] C. Dannewitz, M. D'Ambrosio, and V. Vercellone, "Hierarchical DHT-based name resolution for information-centric networks," *Elsevier Computer Communications*, vol. 36, no. 7, pp. 736–749, Apr. 2013.
- [28] G. Xylomenos, C. N. Ververidis, V. A. Siris, N. Fotiou, C. Tsilopoulos, X. Vasilakos, K. V. Katsaros, and G. C. Polyzos, "A survey of information-centric networking research," *Communications Surveys Tutorials, IEEE*, vol. 16, no. 2, pp. 1024–1049, Second 2014.
- [29] A. Kostopoulos, I. Papafili, C. Kalogiros, T. Levä, N. Zhang, and D. Trossen, "A Tussle Analysis for Information-Centric Networking Architectures," in *The Future Internet*, 2012, pp. 6–17.
- [30] N. Freed and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part Two: Media Types," Nov. 1996.
- [31] K. Nichols, S. Blake, F. Baker, and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers," RFC 2474 (Proposed Standard), Dec. 1998.