Information Centric Networking: A Model for an Improved Internet (and Mobile/Ad Hoc Networks)

Nicola Blefari Melazzi
Department of Electronic Engineering, University of Rome Tor Vergata / CNIT
Via del Politecnico, 1
Italy

blefari@uniroma2.it

We describe the main features of Information Centric Networking (ICN), a networking paradigm devised to overcome some intrinsic limitations of the current Internet. ICN improves network efficiency, naturally supports mobility and multicast communications, eases the operation of fragmented networks (e.g. ad hoc networks) and provides an information-oriented security model. Then, we show some applications of the ICN concept, including mobile, ad-hoc operational scenarios.

ABSTRACT

The Internet provides users with communication channels between hosts identified by an IP address. Information Centric Networking (ICN) is an alternative paradigm, emerged to overcome some intrinsic limitations of the current Internet, in which all information (e.g. a picture) is given a name, which does not include references to its location; then, user's requests for a specific information are routed toward the "closest" copy of such information, which could be stored in a server, in a cache contained in a network node or even in another user's device; then, the information is delivered to the requesting user. With ICN, the communication network becomes aware of the name of the information that it provides and the routing decisions are made on the basis of the information name. In addition, ICN secures the information package itself, instead of securing the communication channels. As a result, ICN: i) improves network efficiency, thanks to in-network caching and information-based routing; ii) naturally supports mobility and multicast communications; iii) eases the operation of fragmented networks, or sets of devices temporarily or permanently disconnected from the rest of the network (e.g. ad hoc networks); iv) offers simpler application programming interfaces; v) provides an information-oriented security model which is rapidly becoming essential, in a world where all traffic in being encrypted, wreaking havoc with established network mechanisms. In this paper, we survey our work in the field of ICN. First we briefly recall how ICN works and present its expected advantages; then we show some applications of the ICN concept including mobile, ad-hoc operational scenarios; we conclude by listing open research issues, including application of ICN to Intelligent Transportation Systems.

1.0 INTRODUCTION

The highly heterogeneous, distributed and mobile nature of modern networks calls for a new networking model. The current Internet model is based on the Internet Protocol (IP) and provides users with communication channels (sockets) between hosts (e.g., a client and a server) that are identified by an IP address. IP network nodes, or routers, forward data among users' hosts on the basis of their IP addresses, which statically determine where they are topologically located in the network. Mobility and multicast solutions have been designed and standardized, but are not currently supported by the network, forcing application developers to implement them at layers higher than IP, as an overlay, introducing inefficiency and complexity. IP routers are blind as to what they are forwarding. Security is provided by securing the communication channels.

Information Centric Networking (ICN) is an alternative paradigm emerged to overcome some intrinsic limitations of the current Internet [1][2][3][4]. In ICN, the network provides users with access to information by name, instead of providing communication channels between hosts. The idea is to provide "access to named data" as the fundamental network service. This means that all information (e.g. a document, a picture) is given a name, which does not include references to its location; then, user's requests for a specific information are routed toward the "closest" copy of such information, which could be stored in a server, in a cache contained in a network node or even in another user's device; finally the content is delivered to the requesting user by the network. With ICN, the communication network becomes aware of the name of the information that it provides and the routing decisions are made on the basis of the information name. This enables nodes to carry out advanced delivery services, like caching and multicasting, thus reducing the resources needed on servers, and improve responsiveness and reliability of applications. In addition, ICN secures the information package itself, instead of securing the communication channels, thus information can be trustily delivered also by untrusted servers or nodes and remain protected also when emerges from a communication channel (e.g. a picture is protected not only while it travels into the network but also after arriving at destination). As a result, ICN: i) improves network efficiency, thanks to in-network caching and information-based routing; ii) naturally supports mobility and multicast communications; iii) eases the operation of fragmented networks, or sets of devices temporarily or permanently disconnected from the rest of the network[1]; iv) offers simpler application programming interfaces; v) provides an information-oriented security and access control model which is rapidly becoming essential, in a world where all traffic in being encrypted, wreaking havoc with established network mechanisms. ICN secures the content itself, instead of securing the communication channels, allowing for a more flexible and customizable protection of content and user privacy.

The capabilities of ICN are particularly valuable as we move towards an increasingly mobile connected world, where information, end-points and people are continually connecting to different points, requiring in-built mobility support from the network. The Internet's coupling of the IP address for both identifying a device (and related content) and for determining where it is topologically located in the network, resulted in conflicting goals: for routing to be efficient, the address must be assigned topologically; for collections of devices to be easily and effectively managed, without the need for renumbering in response to topological change or mobility events, the address must not be explicitly tied to the topology [5]. While this weakness has been well recognized in many cases in the past, solutions have been addressed higher up in the protocol stack and/or implemented as an overlay, for fear of touching IP. But we have had to pay the price, in added complexity and inefficiency. ICN offers a clean solution by logically separating network locators from identifiers, not only of devices but also of content, and potentially of users and functions.

Another motivation for ICN is that the cloud concept is transforming the Internet, which is becoming a network of data centers, with a communication model consisting of computer-to-cloud-to-computer interactions; in this scenario, abstractions and interfaces become fundamental, and ICN is just the perfect solution to transfer data from users to the cloud and viceversa, and within the cloud, easing data analysis and manipulation. The same concept can be applied to services and functions, by giving names also to network functions, so easing the task of allocating and moving them and their components where and when they are needed.

A final motivation for ICN (actually the first in the mind of its proponents) is that the Internet is more and more a multimedia information delivery platform. Information access on the Internet is exploding, with usage shifting to social networking and Internet of Things. Video in particular is increasingly demanding of bandwidth, and multi-party interactive multimedia communications, with stringent quality of experience guarantees from mission-critical applications, would certainly benefit from ICN's caching capability and general better efficiency. Content Delivery Networks are expected to handle 75% of the total traffic and ICN can be seen as nothing else that a *democratic* CDN, built in in the network.

[1] For instance. sensors networks, vehicular networks, social gatherings, mobile networks on board trains, planes, or networks stricken by disaster; note also that ICN could be applied to the whole Internet but also to a subset of it, this is especially useful in ad hoc scenarios.

Last but not least, from a feasibility point of view, we are at an opportune time in the evolution of the Internet to introduce ICN: i) network softwarization (i.e., Software Defined Networking (SDN) + Network Function Virtualization (NFV)) will make it easier to change/add functionality; ii) the research community is now busy designing 5G, so we are on the eve of a new generation of the network; in particular, the ability to partition the network infrastructure in slices, enables support of innovation while production and mission critical functions continue unimpeded.

In this paper, we survey our work in the field of ICN. First we briefly recall how ICN works; then we present its expected main advantages; after that we show some exploitations of the ICN concept in specific scenarios, surveying our work in this area and including an ad hoc use case; we conclude by listing open research issues.

2.0 HOW IT WORKS

The basic functions of an ICN infrastructure are to: i) address information contents, adopting an addressing scheme based on names (identifiers), which do not include references to their location; ii) route a user request, which includes a "destination" content-name, toward the "closest" copy of the content with such a name; iii) deliver the content back to the requesting host (see Figure 1).

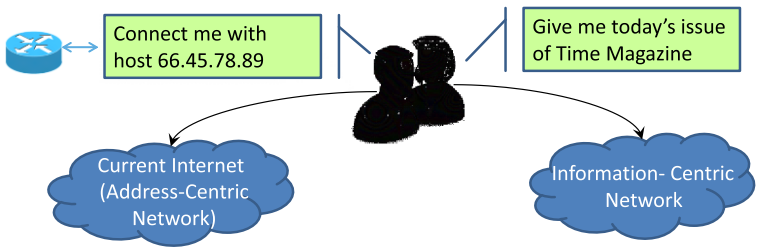

Figure 1: ICN concept

So far several ICN architectures have been proposed [4], which differ on how they implement key functionality, such as:

- Naming: different architectures adopt different naming schemes; some of them use hierarchical names, whereas other ones use flat names.
- Routing: to forward a content request towards a content source, some architectures use on-path name-based routing tables, whereas other ones use an off-path resolution system e.g. based on DHT.
- Security: some architectures decouple security information from the names and this allows using human readable names, while security information is inserted in the header of network data units and the system requires a PKI. Other architectures use self-certified names that include an hash of the public key of the publisher in the name; these architectures do not requires a PKI, but names are not human readable.
- Communication primitives: most architectures use a request-response communication primitive; possibly, publish-subscribe primitives are provided by added functionality located in the middleware or application layer; others inherently supports publish-subscribe communication [4].

In this paper, we follow the model proposed in [2], namely Content-Centric Networking. CCN is supported by an implementation, named CCNx [6] available for Linux, Mac OS and Android platforms. Another important ICN architecture is Named Data Networking (NDN) [7] which used CCNx as its codebase, but as of 2013 it has forked a different version, and by now also developed new functionality and protocols. However, the basic concepts of CCN and NDN are the same and their differences are not of

importance in the context of this paper, thus we will use the terms CCN and NDN interchangeably.

A CCN/NDN network provides users with information items exposed by unique names. Information items (or contents) may be whole documents but also chunks of files. A name that uniquely identifies an information item is formed by a hierarchy of strings, aka *components*, separated by a "/" character. The generic naming scheme is ccnx:/component#1/…/component#n. To fetch an information item, a client sends an *Interest* message that includes the name of the information item. Then, the CCN network finds a source and sends back the information item to the client within a *Data* message (aka Content Object).

Figure 2 shows the model of a CCN node. A name-based Forwarding Information Base (FIB) is used to *route-by-name* Interest messages using a prefix match strategy. A FIB entry contains a prefix and a list of *upstream* faces on which the Interest message may be forwarded to reach a source of information. A face is a generalization of the concept of interface and may be a connection to a next-hop CCN node or directly to an application party. For instance, in Figure 2 Interest messages for information items whose names contain the prefix "foo.news/P1" are forwarded on face 2. The FIB entries are configured through the CCN API, either manually or by a name-based routing protocol.

During the forwarding process of Interest messages, a CCN node leaves reverse path information <name – *downstream* faces> in a Pending Interest Table (PIT), where a downstream face is the face from which an Interest comes from. For instance, in Figure 2 the node has received from the faces 0 and 1 the Interest messages for "foo.news/P1/2" and "foo.news/P1/3". When an Interest reaches a node having the requested information item (either in the cache or in a local application), the node sends back the information item within a Data message. The Data message is routed on the downstream path by consuming the reverse path information previously left in the PITs. If not consumed by a Data message, PIT entries expire after a time out period.

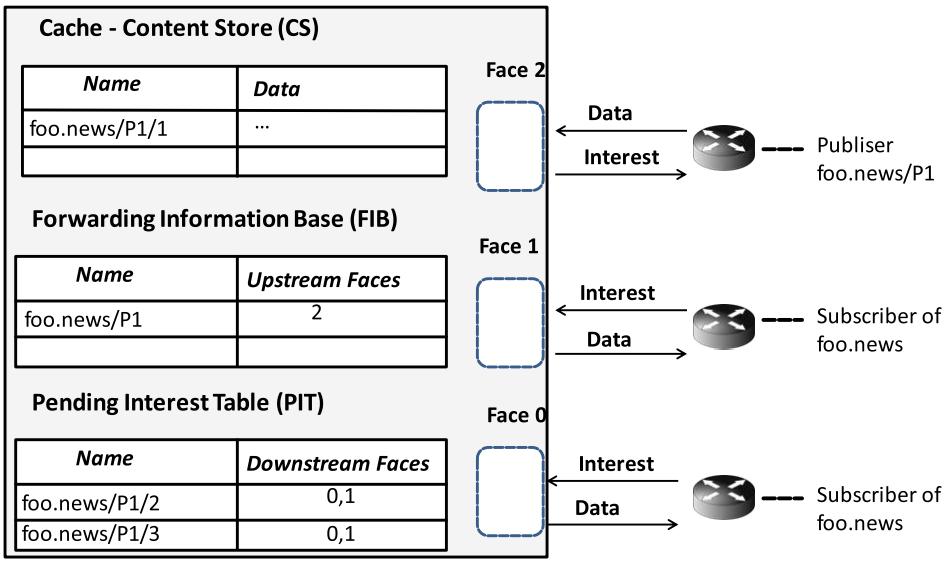

Figure 2: Model of a CCN node

CCN nodes temporarily store the forwarded Data messages in a cache memory (e.g. foo.new/P1/1 in Figure 2), named Content Store. The content store has a limited storing capacity and its use is controlled by a cache replacement policy. If a node receives an Interest message related to a cached information item, the item is sent back immediately, without further forwarding the Interest message. Since caching occurs on all network nodes, this caching approach is usually referred to as *in-network caching*.

In case of concurrent Interest messages regarding the same information item, a CCN node forwards only the first received Interest message and stores in the PIT the set of downstream faces from which it received the next Interests. When the node receives the related Data message, it relays a copy of it towards

all such downstream faces. In doing so a CCN node provides a network-level *multicast* distribution. Of course this is just the basic way of operation, the full specification of CCN/NDN is way beyond the scope of this paper.

3.0 EXPECTED ADVANTAGES

In our opinion, expected advantages of ICN include [10]:

1) Efficient content-routing: ICN would enable Internet Service Providers (ISPs) to perform native content routing. This would be a built-in facility of the network, which would transform the Internet to a native content distribution network, as opposed to Content Delivery Networks (CDN) overlays, which are the current practice.

2) In-network caching: caching enabled today by HTTP proxies requires performing complex stateful operations, which are increasingly difficult, owing to the fact that more and more traffic is nowadays encrypted at the source. ICN would significantly improve efficiency and scalability of caching by judiciously using caches at critical points in the network, exploiting the knowledge of the name of content.

3) Simplified handling of mobile and multicast communication: with ICN, unlike current mobility architectures, when a user changes point of attachment to the network, she will simply ask for the next chunk of the content she is interested in, without the need for maintaining tunnels or re-routing from an anchor point and maintaining state in the network; the next chunk may be provided by a different node than the one that it would have been used before the handover. Furthermore, multicast is an inherent capability in ICN, with content requested from a node being delivered to all interested receivers, without using overlays.

4) Simplified support for time/space-decoupled communications: allowing fragmented networks, or sets of devices to operate even when disconnected from the rest of the 'network' (e.g. sensors networks, ad-hoc networks, vehicular networks, social gatherings, mobile networks on board trains, planes, or networks stricken by disaster).

5) Simplified support for peer-to-peer communications: ICN inherently supports communications between peers, without the need for application-layer overlays.

6) Content-oriented security model: securing the content itself, instead of securing the communication channels, allows for a more flexible and customizable protection of content and user privacy and protects in-network caches from fake content.

7) Content-oriented access control: ICN can provide access to content as a function of time, place (e.g., country), or profile of user requesting the item. This functionality would allow implementing: i) access revocation (also known as digital forgetting), so that content may be removed from the system by its creator, ii) garbage collection, deleting from the network 'expired'/obsolete contents.

8) Content-oriented quality of service differentiation (and possibly pricing): ICN would enable ISPs to differentiate the quality perceived by users from different services.

9) Create, deliver and consume contents in a modular and personalized way: ICN provides opportunities for better customization of the interests of users and the content that is published by providers. This would enable more efficient consumption of content because of better "granularity" in how content is described and identified.

10) Network awareness of transferred content, allowing network operators to better control information and related revenue flows, favoring competition between operators in the inter-domain market and better balancing the equilibrium of power across the eco-system, including

over-the-top players.

A final overall advantage of ICN, which in a way comprehends the specific advantages listed above, is a simplification of network design, operation and management. Currently, content and service providers have to "patch" shortcomings and deficiencies of IP data delivery by using several "extra-IP" functionalities, such as HTTP proxies, CDNs, multi-homing and intra-domain multicast delivery, to name a few. This implies the involvement of several parties, the use of several specific protocols, the deployment of ad hoc devices and the interplay of different functionalities, often offered and managed by different companies and businesses. Apart from technical complexity, such operations also add management and administrative complexity. In an ICN environment, such diverse functions can be integrated in the network, e.g. by inherently supporting data replication, caching, multi-homing and multicast delivery.

The research community is working to achieve the advantages promised by ICN, while addressing its two biggest open issues: i) scalability of the naming and routing functionality (e.g. by caching within network nodes a subset of all possible routes and querying a remote database for less usual requests [11]) and ii) the need to devise a credible migration path from the communication paradigm of the current network infrastructure or, more practically, to allow the coexistence of different network operating modes, so that different sections of the network can operate according to different paradigms (e.g. by resorting to recent software networks solutions such as SDN and NFV [12]).

4.0 INTERNAMES: AN EVOLUTION OF ICN

In [8] we proposed *Internames*, which evolves from ICNs host(s)-to-name model to a name-to-name principle in which names identify both source and destination entities, and the name of all communication entities is not statically bound to their physical location. In addition, names are used to identify all entities involved in communication: content, users, devices, logical points; also services are bound to a service-identifier (i.e. a name) rather than to an IP address, to easily enable re-location or duplication or anycast search of service (components), easing the support of emerging service-centric networking architecture. Internames is envisioned to be an overarching name-to-name communication primitive that is fully compatible with ICN principles, accommodates the co-existence (or gradual migration) of different network realms (e.g., IP, ICN, VANET) and is suitable for application scenarios where ICN is somehow limited by its reliance on a "host-to-name" approach.

4.1 Disaster management scenario

Internames further increases the utility and efficiency of ICN.

For instance, let us consider a representative use case developed in the framework of the GreenICN joint EU-Japan project [9]: that of Disaster Management, to limit the effects of disasters on ICT infrastructures. The project considered in particular the requirements derived by analysis of the consequences of the enormous earthquake that hit Northeastern Japan on March 11, 2011. The lack of information and means of communication caused the isolation of several Japanese cities, impacting the safety and well-being of residents, and affected rescue work, evacuation activities, and the supply chain for food and other essential items. The earthquake showed the vulnerability of current networks and that mobile phones have become the lifelines for communication, including safety confirmation [14].

ICN can be very helpful in disaster situations since the features that we described in the previous Section particularize as explained below [14]:

- Routing-by-name: ICN protocols natively route and identify content by names, moving the process of name resolution from the application layer to the network layer. This functionality is very handy in a fragmented network where reference to location-based, fixed addresses may not work as a consequence of disruptions. In addition, ICN does not necessarily rely on the

reachability of application-layer servers (e.g. DNS resolvers).

- Sessionless: ICN does not require full end-to-end connectivity. This feature facilitates a seamless aggregation between a normal network and a fragmented network, which needs Delay Tolerant Networks-like message forwarding.
- Caching: Caching helps in handling huge amounts of traffic, and can help to avoid congestion events typical of the aftermath of a disaster.
- Authentication of named data objects: ICN is built around the concept of named data objects. Several proposals exist for integrating the concept of 'self-certifying data' into a naming scheme. With such approaches, the origin of data retrieved from the network can be authenticated without relying on a trusted third party or PKI.
- Support of group communications; ICN's natural support of multicasting is very useful to send data to groups, which arise very often after a disaster.
- Content-based access control: ICN can regulate access to data objects (e.g. only to a specific user or class of users) by means of techniques that integrate confidentiality, integrity, authenticity and access-control *within the content itself*. In this field, we exploited a technique named Ciphertext-Policy Attribute Based Encryption (CP-ABE) [19]. The idea of CP-ABE is that a content (e.g. instructions for rescuers) can be encrypted along with an access control policy, so that only users satisfying such policy can decrypt the received information, and no server-based access control infrastructure is needed (policies travel with the data, and decryption is bound to the user possession of a set of attributes). Fully decentralized CP-ABE schemes permit decryption policies over attributes released by multiple independent authorities, thus getting rid of the need to deploy hardly viable central administrative entities. CP-ABE can be used to encrypt content to be sent to different categories of users. For instance, we can encrypt a ciphertext such that a message can only be decrypted by someone with attributes "Public Official" and "Rank> Executive" or "Emergency Team" and "Any Rank". It is clear that this functionality could facilitate trusted communications among peer users in isolated areas of the network.

Now, all the above features are provided by ICN. Internames can provide added utility. When there is a single, static, sender of information to one or multiple recipients identified by a name, current ICN host-to-name communications are sufficient. However, in our disaster use case, each sender may have different roles, persona and responsibilities (as an individual, as an authority). When that person wishes to send some information, e.g., related to a disaster, the initial communication could be viewed as coming from an authority (identified by a name) to a designated set of recipients (identified by another name). Unlike ICN, the return message could be addressed to the authority (i.e., a name) and could follow a different path for communication than the original path in the reverse direction. The responses would be delivered to the original name that transmitted the initial message, even if the named entity moves from its original location, or the original named entity is mapped to a different entity with a related location (e.g. because the original one is not reachable anymore), going beyond simple source mobility. Moreover, the response could be delivered to more than one entity associated with that name (in a sense, it is the reverse of a traditional multicast, having the information flow from receivers to a group of/all senders).

We conclude the section by briefly describing two use cases in the framework of the considered scenario.

Use case a) . The government of a country wants to be prepared in case a disaster strikes a large region, affecting thousands of citizens. For this reason, it prepares a package of information containing three different sets of documents to be distributed to government representatives, rescuers and citizens. The structure and some parts of the documents are defined before-hand but other critical parts are completed with real data collected right after the disaster. The government, playing the role of content provider creates the documents by exploiting information coming from different sources (e.g. firemen, policemen, but also network operators), setting rights on who, where and when can access such documents. A telecom operator distributes the documents, suitably exploiting routing, caching, multicast, quality of service

differentiation and energy saving functionalities provided by ICN. Government representatives, rescuers and citizens receive only the documents meant for each of such classes, with differentiated performance and priorities. Rescuers can inject copies of same documents after reaching a disconnected area. Citizens can exchange documents among themselves, facilitating the spreading of information to disconnected portions of the network. Developers can design applications leveraging all the above and even design simple ad hoc applications right after a disaster to face unexpected/new needs.
Use case b) A telecom operator owning and managing both a fixed and a cellular network is interested in monitoring the connectivity status of the network, the use of network resources, the status of user terminals, the energy consumption and battery levels of network and of users devices. It also wants to use all this information to manage the network in normal conditions and to re-establish network operation in an isolated area, following a disaster/breakdown and manage it. With ICN, network devices and even network entities can be identified and addressed with a name, facilitating: i) entities discovery; ii) anycast routing; iii) end-to-end management, including user terminal; and iv) operation in fragmented networks.

5.0 APPLICATION SCENARIOS

In this section we present two applications of ICN, to better show its potential.

5.1 A cooperative peer-to-peer application for live streaming of videos

The first example of exploitation of ICN concerns a peer-to-peer application for live streaming of videos encoded at multiple bit rates (adaptive live video streaming, we use the MPEG-DASH streaming standard) [13]. Peers are assumed to be a small set of neighbouring mobile cellular devices that cooperatively download a live video stream from the cellular interface and share downloaded video segments through a proximity channel (e.g. Wi-Fi Direct). The cooperation logic is designed to improve the playback quality perceived by a peer, with respect to the quality that the same peer could achieve by downloading the stream only by itself. For instance, we can imagine a situation in which passengers of a train are interested in watching news with their mobile phones, or alternatively a pay-per-view scenario in which all the mobile devices available in an apartment are concurrently used to improve the video quality of a stream offered by a content provider like Netflix. We show how to exploit an ICN API, namely the CCN API, to simplify the application development. Indeed, if we had used the plain TCP/IP API, we would have had the burden of implementing and orchestrating on top of it routing-by-name, caching and multicast functionalities, which instead are built-in in CCN. In [13] we show both the feasibility and the performance of the application.

5.2 Topic-based, publish-subscribe MANET systems

The second example of exploitation of ICN regards publish-subscribe Mobile Ad-hoc NETworks (MANETs), i.e. systems that connect mobile wireless devices without an underlying communication infrastructure. Communications occur in a multi-hop fashion, using mobile devices as routers. Many MANET applications require to exchange data (GPS position, messages, pictures, etc.) by using a publish-subscribe interaction scheme, which is information-centric in nature. Publishers characterize their information items with a set of attributes, and subscribers register their interests in receiving only those information items whose attribute match a given criterion, regardless of who is the publisher. Publishers and subscribers are decoupled in space (they do not need to know of each other), and in time (they do not need to be active at the same time). Such decoupling simplifies mobility and disconnected operations, which are typical of MANETs. Moreover, one-to-many delivery can exploit intrinsic broadcast properties of the wireless channel.
The simplest publish-subscribe scheme is the so called “topic-based” one, in which the only attribute of an information item is its being part of a given topic. Participants of a topic-based publish-subscribe system can publish information items on a topic and can subscribe to a topic to receive the related published information items. Topics are identified by a name (e.g. “foo.news”), and publishing an information on a topic *T* implies its distribution to the users subscribed to *T*.

Several MANET publish-subscribe systems have been proposed so far in the literature assuming an underlying TCP/IP network. In the paper [15], we discuss the benefits of building a MANET publish-subscribe system exploiting ICN, rather than TCP/IP, and show how ICN functionality can be exploited to easily set up an efficient and resilient topic-based publish-subscribe system for MANETs. The in-network caching and multicast functionality offered by ICN results in low delay and delivery efficiency. Moreover, through a dynamic configuration of ICN forwarding tables, we carry out a data muling functionality. Moving devices can use this functionality to transport information items among disconnected parts of the network, thus realizing the delay tolerant delivery typical of so-called Delay Tolerant Networks (DTNs).
Our findings are based on the practical development experience that we gained by using ICN as the underlying layer of the BEE DDS middleware [17], which is a specific implementation of the Object Management Group (OMG) specification of the Data Distribution Service for Real Time System [18]. In [15] we discuss design approaches, present details of our system and report the results of a performance evaluation study carried out with real software in an emulated environment, based on Linux virtual machines.

6.0 CONCLUSIONS

In this paper, we argued that Information Centric Networking (ICN) is a promising paradigm to augment the functionality provided by the Internet and to reduce its complexity. The main open issue of ICN is the design of a credible migration path from the current network infrastructure or of realistic solutions that would allow the coexistence of different network operating modes. The issue is somehow made more approachable thanks to recent advances in network softwarization techniques that ease the evolution of network protocols. In addition, it would be interesting and worthwhile to adopt ICN in specific and isolated sections of the network, without necessarily targeting a worldwide pervasive deployment, like for instance in ad hoc networks.

As for other applications, we are currently working on exploiting ICN for realizing distributed spatial databases. Our ICN-based database, named OpenGeoBase, allows anyone to publish data relevant to a specific geographic area, ranging from transport schedules to sensor-generated or user generated real-time information, but also, point of interests, etc. Then, interested users/travel operators can search and retrieve all data available in such geographic area, which are needed e.g. to plan an optimal multimodal trip. Publishers are not forced to upload their data in a central repository but they can keep them in local, distributed repositories, under their control. OpenGeoBase logically puts together all individual repositories and make it easy for users to search for and retrieve the data they are interested in. OpenGeoBase can quickly and efficiently provide massive information to database users; easily operate in a distributed way, deploying and using many database engines in parallel; secure every piece of content in a customizable way; naturally slice resources, so that several tenants and users can use the database in parallel and independently. We are also exploiting OpenGeoBase to support an Intelligent Transport System application, by enabling fast responses to continent-wide queries about stops, routes, trips, schedules, real-time updates, fares, etc. The application is meant as a support to multimodal real time journey planning, but the ability of OpenGeoBase to gather/share distributed data (e.g. from sensors) very efficiently and quickly is of great interest also for automated driving and for other applications.

7.0 ACKNOWLEDGMENT

This work was performed in the context of the project BonVoyage, which received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 635867. The Author thanks Andrea Detti as principal co-author of several referenced papers recalled in this work.

8.0 REFERENCES

[1] T. Koponen, M. Chawla, B.G. Chun, et al.: "A data-oriented (and beyond) network architecture", ACM

SIGCOMM 2007

[2] V. Jacobson, D. K. Smetters, J. D. Thornton et al., ”Networking named content”, ACM CoNEXT 2009

[3] A. Detti, N. Blefari Melazzi, S. Salsano, M. Pomposini, “CONET: A Content Centric Inter-Networking Architecture”, ACM SIGCOMM Workshop on Information-Centric Networking, ICN 2011.

[4] G. Xylomenos, C. Ververidis, V. Siris, N. Fotiou; C. Tsilopoulos, X. Vasilakos, K. Katsaros, G. Polyzos George, "A Survey of Information-Centric Networking Research," Communications Surveys & Tutorials, IEEE , vol.16, no.2, pp.1024-1049, Second Quarter 2014

[5] David Meyer and Darrel Lewis. The locator/id separation protocol (lisp). In RFC 6830, 2013.

[6] CCNx project web site: http://www.ccnx.org

[7] Named Data Networking (NDN) project. http://named-data.net.

[8] N. Blefari Melazzi, A. Detti, M. Arumaithurai, K.K.Ramakrishnan: ”Internames: a name-to-name principle for the future Internet”, International Workshop on Quality, Reliability, and Security in Information-Centric Networking (Q-ICN), 20 August 2014 Rhodes island, Greece.

[9] GreenICN: Architecture and Applications of Green Information Centric Networking, a joint EU-Japan research project. www.greenicn.org

[10] N. Blefari Melazzi, L. Chiariglione: “The potential of Information Centric Networking in two illustrative use scenarios: mobile video delivery and network management in disaster situations”, IEEE Multimedia Communications Technical Committee E-letter, special issue on "Multimedia Services in Information Centric Networks", Vol. 8, N. 4, July 2013, pp. 25-28, http://committees.comsoc.org/mmc/e-news/E-Letter-July13.pdf

[11] A. Detti, M. Pomposini, N. Blefari-Melazzi, S. Salsano: “Supporting the Web with an Information Centric Network that Routes by Name”, Computer Networks: The International Journal of Computer and Telecommunications Networking, Elsevier North-Holland, Inc. New York, NY, USA, Vol. 56, 2012, pp. 3705-3722

[12] N. Blefari-Melazzi, A. Detti, G. Morabito, S. Salsano, L. Veltri: “Information Centric Networking over SDN and OpenFlow: Architectural Aspects and Experiments on the OFELIA Testbed”, Computer Networks: The International Journal of Computer and Telecommunications Networking, Elsevier North-Holland, Inc. New York, NY, USA, Volume 57, Issue 16, 13 November 2013, Pages 3207–3221

[13] Detti; B. Ricci; N. Blefari Melazzi: “Mobile Peer-To-Peer Video Streaming over Information-Centric Networks”, Computer Networks: The International Journal of Computer and Telecommunications Networking, Elsevier North-Holland, Inc. New York, NY, USA. Volume 81 Issue C, April 2015

[14] J. Seedorf, D. Kutscher, A. Tagami, K. Sugiyama, M. Arumaithurai, Y. Koizumi, T. Hasegawa, N. Blefari-Melazzi, T. Asami, K. K. Ramakrishnan, T. Yagyu, I. Psaras: “The Benefit of Information Centric Networking for Enabling Communications in Disaster Scenarios“, IEEE GC 2015 Workshop on Information Centric Network Solutions for Real-World Applications, Globecom 2015, December 6-10, 2015, San Diego, CA, USA

[15] A. Detti, D. Tassetto, N. Blefari Melazzi, F. Fedi: “Exploiting Content Centric Networking to develop topic-based, publish-subscribe MANET systems”, Elsevier, Ad Hoc Networks, Volume 24, Part B, January 2015

[16] G. Bianchi, A. Caponi, A. Detti, N. Blefari Melazzi: “Check before storing: what is the performance price of content integrity verification in LRU caching?”, ACM SIGCOMM Computer Communication Review, Volume 43 Issue 3, July 2013

[17] BEE DDS website http://www.beedds.com

[18] Object Management Group, “Data Distribution Services for Real Time Systems”, Version 1.2, OMG, January 2007

[19] GreenICN project, Deliverable D2.3.3: “Final solution for Access control and management in fragmented networks”, http://www.greenicn.org/deliverables/deliverables/