Information-Centric Networking for the Internet of Things: Challenges and Opportunities

Marica Amadeo, Claudia Campolo, José Quevedo, Daniel Corujo, Antonella Molinaro, Antonio Iera, Rui L. Aguiar, and Athanasios V. Vasilakos

Abstract

In view of evolving the Internet infrastructure, ICN is promoting a communication model that is fundamentally different from the traditional IP address-centric model. The ICN approach consists of the retrieval of content *by (unique) names*, regardless of origin server location (i.e., IP address), application, and distribution channel, thus enabling in-network caching/replication and content-based security. The expected benefits in terms of improved data dissemination efficiency and robustness in challenging communication scenarios indicate the high potential of ICN as an innovative networking paradigm in the IoT domain. IoT is a challenging environment, mainly due to the high number of heterogeneous and potentially constrained networked devices, and unique and heavy traffic patterns. The application of ICN principles in such a context opens new opportunities, while requiring careful design choices. This article critically discusses potential ways toward this goal by surveying the current literature after presenting several possible motivations for the introduction of ICN in the context of IoT. Major challenges and opportunities are also highlighted, serving as guidelines for progress beyond the state of the art in this timely and increasingly relevant topic.

The definition of the Internet of Things (IoT) is still under debate, but there is a large consensus on attributing IoT a primary role in providing global access to services and information offered by billions of heterogeneous devices (or *things*), ranging from resource-constrained to powerful devices (and/or virtualized everyday life objects) in an interoperable way.

To this aim, *evolutionary* approaches that provide IP-based networking functionalities are typically pursued. In this arena, different Internet Engineering Task Force (IETF) working groups are very active (e.g., 6LoWPAN, ROLL, CoRE) [1], but despite great efforts and valuable achievements, the large-scale deployment of IP-based IoT solutions still provides challenges. The limited expressiveness of IP addressing simultaneously serving as locator and identifier, the need for a resolution system, complex mobility support, multicast, and massive access under the stringent performance requirements of IoT (e.g., scalability, energy efficiency) are just a few examples.

In parallel, the research community is currently exploring *cutting edge* approaches to transform the Internet, as we know it today, into a system more capable of and tailored for effective content distribution, according to today (and tomorrow's) needs. Information-centric networking (ICN) [2] has been recently proposed for this purpose and is inspiring the design of the future Internet architecture. Unlike the IP-address-centric networking of the current Internet, in ICN every piece of content has a *unique*, *persistent*, *location-independent name*, which is directly used by applications for accessing data. This *revolutionary* paradigm also provides content-based security regardless of the distribution channel and enables in-network data caching. Such features make ICN promising, not only for content distribution in the Internet, but also to support several IoT scenarios like the ones in Fig. 1, which involve different sensing and automation applications.

In fact, ICN matches a wide set of IoT applications that are *information-centric* in nature, since they target data regardless of the identity of the object that stores or originates them. For example, road traffic/environmental monitoring applications are oblivious to the specific car/sensor that provides the information. ICN *names* can directly address heterogeneous IoT contents and services, such as vehicular/home services and environmental data. Unlike IP addresses, such names are independent of the location of content/service producers, thus facilitating delivery operation in the presence of nodes' mobility.

By *caching data closer to consumers*, ICN can reduce data retrieval delay and network load, and limit massive access to resource-constrained devices. For instance, once home appliances have been triggered about their energy consumption, the retrieved information can be cached at intermediate nodes and be available for later requests.

Marica Amadeo, Claudia Campolo, Antonella Molinaro, and Antonio Iera are with University Mediterranea of Reggio Calabria.

José Quevedo, Daniel Corujo, and Rui Aguiar are with the Instituto de Telecomunicaçoes, Universidade de Aveiro.

Athanasios V. Vasilakos is with Luleå University of Technology.

0890-8044/16/$25.00 © 2016 IEEE

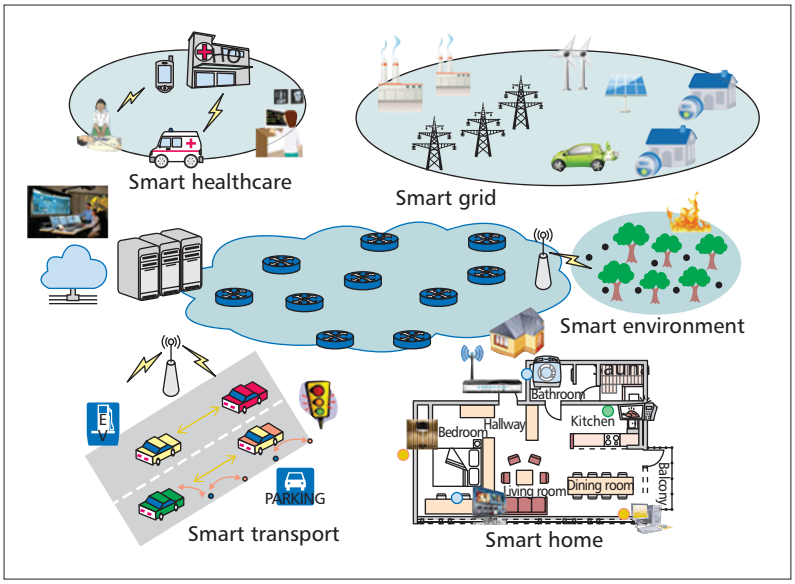

Figure 1. Main IoT scenarios: smart grid, smart environment, smart home, smart transport, and smart healthcare.

The scientific community is therefore debating ICN-IoT deployments within the ICN Research Group (ICNRG) of the Internet Radio Task Force (IRTF). Early documents such as [3, 4] are currently under discussion therein to define how to satisfy IoT requirements over existing ICN proposals. In the meantime, other research works were recently published [5–10] considering ICN as a promising networking solution for IoT, highlighting particular aspects of its feasibility.

To the best of the authors' knowledge, there is still a lack of properly addressing the topic of IoT integration with ICN and its inherent issues. This article contributes to fill this gap, while moving away from direct performance comparisons with IP-based IoT research, and paves the way for ICN usage in IoT by answering the following questions:

- *What are the motivations and main expected benefits of introducing the ICN paradigm in IoT?*
- *What are the solutions, addressed issues, and open challenges?*

Information-Centric Networking Basics

Several ICN architectures have been proposed [2], as summarized in Table 1, characterized by different protocol designs but sharing a common core of ICN principles that can be summarized as follows:

- Content-based *naming* and *security*
- In-network *caching*
- Name-based content *discovery* and *delivery*
- A connectionless receiver-driven communication model

As an example, a typical ICN data exchange is summarized in Fig. 2. Therein, ICN consumers (C_1 and C_2) specify *which* named content they seek and not *where* it is provided. Both *hierarchical* and *flat names* are possible in ICN, with the former appearing as uniform resource identifier (URI)-like identifiers with variable lengths, while the latter comprises fixed-length identifiers with no semantic structure. Moreover, the use of *unique names* makes each content packet a self-identifying unit and drives request forwarding toward content provider(s) (typically the closest one), thus enabling anycast retrieval.

Content-based security makes each data a self-authenticating unit, with protection and trust implemented at the packet level rather than at the communication channel level. The security mechanisms are closely related to the naming scheme. When hierarchical naming is used, security-related information (e.g., the publisher signature) is embedded into a separate field of the content unit, thus requiring a public key infrastructure (PKI) for integrity checks. Flat namespaces instead enable the use of self-certifying names, allowing integrity checks without the need for a PKI.

Since each data packet is self-consistent, *in-network caching* is enabled, with potentially every network element caching the processed data packets and making them available for future requests; for example, consumer C_2 in Fig. 2 is immediately served by router R_6. Distributed caching makes communication connectionless by not requiring consumers and producers to be simultaneously connected.

In contrast to the current Internet, where senders control data transmission, ICN data retrieval is *receiver-driven*, consisting of two phases: the *discovery* triggered by a consumer to find the content or its replication, and its *delivery* back to the interested consumer. Content discovery can be supported in two main ways: via *name-based routing* (NBR) or through a *look-up-based resolution system* (LRS).

With NBR, the consumer sends a content request packet (i.e., the so-called *Interest*), hop-by-hop relayed by the forwarding nodes by looking up a name match into their forwarding information base (FIB). Once the content is found, it follows the *soft-state* traces on the reverse path back. By recording the pending requests until the Data packets are received, each forwarder can measure delivery performance (e.g., round-trip time) and, in case of problems (e.g., when losses or delays are detected), promptly try alternative paths. Therefore, the forwarding plane can be considered *intelligent* and *adaptive*: it can deal with short-term churns, while the routing protocol only deals with long-term topology changes.

With LRS, the content request is delivered to a resolution system, the implementation of which varies depending on the ICN architecture. For example, SAIL defines a distributed name resolution system (NRS) based on hierarchical distributed hash tables (DHTs); and PURSUIT introduces the Rendezvous Network, implemented as a hierarchical DHT, which collects *publish* and *subscribe* messages and instructs the Topology Manager, which handles the network topology, to create *optimal* forwarding paths (Table 1). Therefore, the content is forwarded to the consumers by following the indications of the resolution system. For example, the NRS in SAIL identifies a set of host locators, and the Topology Manager in PURSUIT creates an in-packet Bloom filter that encodes the data delivery path in a compact manner.

None of the ICN architectures in Table 1 has been specifically designed with IoT in mind, but mainly to support general Internet services or a specific scenario (e.g., smart grid in C-DAX, emergency in GreenICN). Although the IoT applicability of some of these architectures (e.g., named delivery networking, NDN, and MobilityFirst) has been recently discussed [8], none of them can claim to fit all the IoT features in its native design, motivating the analysis in the next sections.

Why ICN for IoT?

ICN is still in the discussion phase, enabling it to consider IoT by design. Moreover, with the already occurring explosion of connected devices targeting smart environments (i.e., sensors and actuators), the information they produce can be regarded as content. In this sense, ICN provides a new opportunity, contributing to the shortening of the gap between the physical and digital worlds by addressing content by its name instead of the regular request-to-IP translation mechanisms used today. At this point, no clear ICN development path exists that could be

Project	Naming and security	Service model	Discovery	Delivery	Main application scenarios
Named Data Networking (NDN): www.named-data.net	Hierarchical names; publisher signature with PKI	Pull-based	Name-based routing of requests, which are maintained as pending in traversed nodes	Soft-state forwarding (content units follow the pending requests back to the consumers)	• General data dissemination • Traditional applications (e.g., conferencing) • VANETs • Building management systems
MobilityFirst: mobilityfirst.winlab.rutgers.edu/	Self-certifying flat names	Pull-based	A distributed global name resolution service maps names to locators (i.e., a set of network addresses)	Forwarding based on the discovered locators	• Data dissemination under mobility conditions • IoT services
Pursuing a Pub/Sub Internet (PURSUIT) (previously PSIRP): www.fp7-pursuit.eu/	Self-certifying flat names consisting of scope and rendezvous parts (scopes organized hierarchically)	Publish-subscribe	Content requests are subscriptions managed by rendezvous nodes, which instruct the network Topology Manager to create forwarding paths	Source-routing (path information is stored in a Bloom filter included in the packet)	General data dissemination
Scalable and Adaptive Internet Solutions (SAIL) (previously 4WARD): www.sail-project.eu/	Self-certifying flat names with possible explicit aggregation	Pull-based	A distributed name resolution system based on hierarchical DHT maps names to locators, but direct name-based routing (as in NDN) is also supported	Forwarding based on the discovered locators or soft-state forwarding if name-based routing is used for discovery	General data dissemination
Convergence: http://www.ict-convergence.eu/	• Self-certifying flat names • Hierarchical names • Publisher signature with PKI	Pull-based and publish-subscribe	Direct name-based routing or subscription to the publish/subscribe system	Soft-state forwarding; or forwarding managed by the publish/subscribe system	• General data dissemination • Video streaming
Content Mediator Architecture for Content-Aware Networks (COMET): http://www.comet-project.org/	Name consists of two human-readable parts: • The naming authority • The content name under its naming authority Public key cryptography	Pull-based	The content mediator entity (CME) is able to locate all the content copies. Best server and delivery path are selected by the CME and related instructions included in the COMET packet header	Source-routing (content-aware forwarding entities forward the packet by following the instructions in the COMET packet header)	General data dissemination
Architecture and Applications of Green ICN (GreenICN): http://greenicn.org	• self-certifying flat names; • hierarchical names plus attributes for user-defined priority, space/temporal-validity; publisher signature with PKI	Pull based and topic-based publish/subscribe	Direct name-based routing with prioritization or subscription to the publish/subscribe system	Soft-state forwarding or forwarding managed by the publish/subscribe system	• Emergency (i.e., the aftermath of a disaster) • Video streaming
Cyber-Secure Data and Control Cloud for Power Grids (C-DAX): http://cdax.eu/	Information organized in topics, each one with a flat identifier, a set of attributes, and publishing keys; access control and key management handled by a C-DAX security Server	Pull-based and topic-based publish/subscribe	Direct queries or subscriptions to topics resolved by the C-DAX resolver discovery system	FIBs for topic identifiers and subscribers (entries updated upon join/leaving events) maintained by C-DAX broker nodes	Smart grid

Table 1. An overview of ICN research projects.

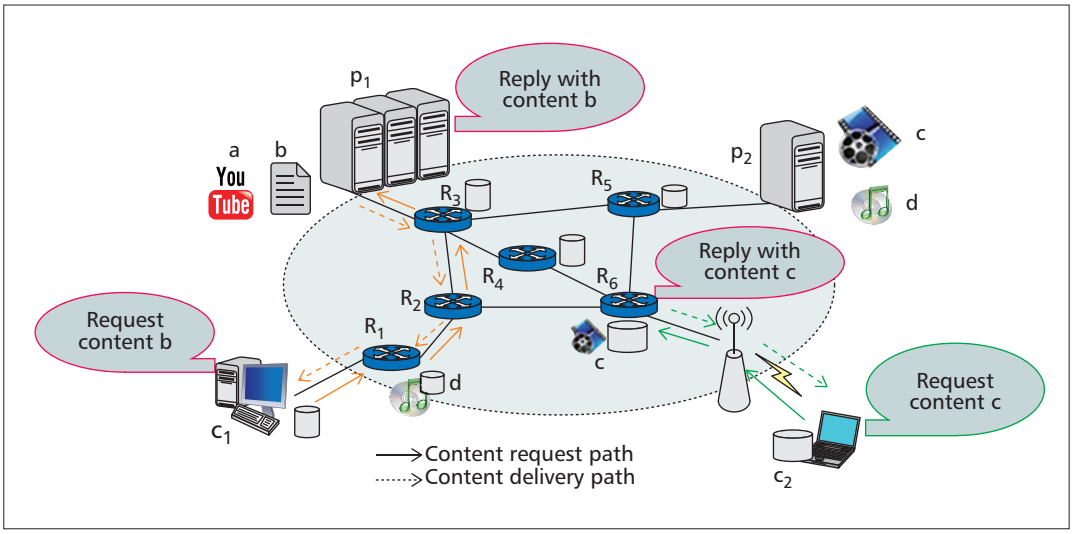

Figure 2. Content exchange in ICN: consumers request contents that can come from any source holding a copy of them (either the original provider or a caching node).

used for direct comparison against IP in order to assess performance increase. This is because on one hand, ICN is still being progressed, and on the other hand, IP has unfolded into multiple variants of IoT solutions. In this sense, this article does not intend to provide a comparison between IP IoT and ICN IoT, but rather to identify IoT as an important deployment scenario for the utilization of ICN mechanisms, and their key benefits in such environments. What is important to highlight, however, is that ICN can benefit from its exposure to different scenarios of what can be considered content retrieval, allowing new concepts to be developed based on named requests, new applications to be created, and the enrichment of the base ICN mechanisms to cater to a broader range of scenarios, as a true future Internet network layer.

The previously discussed ICN core principles have the potential to fulfill the main IoT requirements summarized in Table 2 and discussed in the following.

Scalability

IoT by itself provides stringent scalability challenges, still being addressed by the research community, in the presence of an upcoming and increasing explosion of data/signaling packets generated by billions of connected devices. The forefront of typical IP-based content retrieval mechanisms (e.g., peer-to-peer, P2P, and content delivery networking, CDN) poses complex issues, such as suboptimal peer selection or their incapability to leverage in-network storage, in such scenarios. The inherent operating mechanisms of ICN, despite not being specifically targeted with IoT in mind, offer promising scalability aspects for its deployment capability in such environments.

Concretely, recent standardization efforts[1] highlight the potential of ICN-based IoT solutions to draw away from the current typical centralized service discovery of devices and services by mapping named information to an object or the information generated by it (e.g., sensor measurements). In fact, associating IoT content to names enables information to be structured into scopes and allows users to specifically request the content that they really want (instead of locating it in a specific node amid all the other content available therein). This naming flexibility exploits the higher addressing potential of ICN, allowing a name in the IoT context to identify not only a content, but also a service or a device function.

[1] IRTF RFC 7476 — Information-Centric Networking: Baseline Scenarios, https://tools.ietf.org/html/rfc7476.

By offering name resolution at the network layer and forwarding content by its name, ICN also has the potential to reduce the signaling footprint in IoT deployments. Concretely, ICN nodes have the ability to identify requests for the same named information, avoiding the need to forward them differently on the same path. In addition, content becomes cached in traversing nodes, allowing requests to be satisfied by the first available copy, preventing source over-querying and supporting connectionless scenarios. Finally, ICN allows data to be transmitted to multiple consumers by using native anycasting and multicasting.

However, the utilization of these mechanisms in IoT environments also has the potential to raise scalability issues of their own, under debate by the ICN community [2]. Concretely, ICN name-based mechanisms are made available regardless of the content location, which can limit different scenarios. Considerations on extending information naming to also identify devices can actually draw solutions that reduce the applicability scope of ICN, as it would be trying to mimic the host-based behavior of TCP/IP. Moreover, the amount of content names is orders of magnitude larger than the number of hosts connected to the current Internet, meaning that the routing and naming capabilities of ICN face a much more difficult task when compared to the current global routing and Domain Name Service (DNS) resolution services.

Notwithstanding, ICN research has already been putting forth solutions, such as the utilization of DHTs, late-binding mechanisms, and routing information aggregation, that have a substantial impact in naming resolution procedures, albeit inducing greater memory and processing costs. However, further practical deployment analysis is needed to thoroughly assess both the scalability benefits and hindrances of the current installments of ICN, as well as of upcoming IoT-supportive enhancements to its ongoing design.

Quality of Service

Due to the high heterogeneity of IoT use cases, quality of service (QoS) requirements can be very different. For instance, sensing requires the exchange of typically small data, either in an event triggered (e.g., an alarm) or periodical (e.g., traffic monitoring) manner. Some sensing data require timely reception (e.g., in case of an alarm), while others may tolerate longer delivery delays (e.g., home temperature monitoring). Some IoT applications also account for data freshness needs, for example, when consumers are interested in the latest instance

	Named data	Anycasting	Multicasting	In-network caching	Content-based security	Connectionless mode
Scalability: avoiding the explosion of data/signaling packets in the presence of billions of devices	✓	✓	✓	✓		✓
Quality of service: support of different application requirements (e.g., reduced access latency)		✓	✓	✓		
Security: integrity, privacy, authentication, authorization, trustworthiness	✓				✓	
Energy efficiency: managing communication toward sleeping devices, including mechanisms for reducing energy consumption		✓	✓	✓		✓
Mobility: consumer and producer mobility support, multi-homing supportt	✓	✓		✓		✓
Heterogeneity: managing heterogeneous devices/technologies/services	✓					✓

Table 2. Main IoT requirements and basic ICN support.

of a constantly upgraded content (e.g., a hospital needs updated vital signs of a remotely monitored patient) vs. an available older copy in a nearby cache point.

IP networks apply QoS through the execution of different extensions done over the base protocol, such as multiprotocol label switching (MPLS) and Resource Reservation Protocol (RSVP) under the integrated/differential services (IntServ/Diffserv) paradigms. In these cases, resources are reserved at each hop between the source and the content requester, requiring extensive signaling, flow identification, and queue processing at the forwarding entities. Ultimately, this leads to complexity in routers, with the potential to increase to a larger extension with the explosion of IoT traffic due to the unprecedented amount of connected nodes, different device characteristics, and traffic requirements.

ICN has the potential to improve the quality of content retrieval and manage different QoS demands. The native support of in-network caching, anycasting, and multicasting all together contributes to speed up data retrieval and reduce traffic congestion. Moreover, every ICN design is able to perform advanced and efficient forwarding mechanisms. For instance, architectures with LRS may leverage knowledge of the network topology to compute optimal delivery paths (e.g., PURSUIT). Vice versa, architectures with NBR may leverage the adaptive forwarding capability to react to early signs of network problems (e.g., NDN).

Security

Enabling security services in IoT is fundamental, since most IoT applications have the potential to affect our personal daily lives, and are not deployed in isolation but are exposed to external controls on the Internet.

In IP, security was not conceived by design, with its support being introduced later on to allow for authentication and data integrity. In this way, aspects such as wireless communications or low-powered nodes can hinder the performance of existing protocols. IP-based security protocols (e.g., IKEv2/IPsec, TLS/SSL, DTLS, HIP, PANA, or EAP), even though discussed as solutions in IPv6 over Low-Power Wireless Personal Network (6LoWPAN) and Connection Oriented Route Establishment (CoRE), are dependent on the location identification of nodes. In reality, it is the communication channel between a specific pair of communication nodes that is being secured rather than the content. Moreover, when security has to be combined with other requirements, such as mobility, their joint operation generates even more complex scenarios.

By offering *security support at the network layer*, ICN facilitates content sharing between nodes since data authentication and integrity can be verified locally, removing the need for trusting in intermediary nodes. In addition, by securing the content itself, ICN can restrict data access to a specific user or a group of users.

Energy Efficiency

Resource-constrained IoT devices have severe limitations on power and computing capabilities, as well as on networking functionalities. Most embedded devices spend a great part of their lifetime in sleep mode and only awake when they need to exchange data. Therefore, energy-efficient operation design is crucial for any IoT networking solution.

Current energy efficiency approaches are not handled at the network layer, being targeted at the medium access control (MAC) layer or above the transport layer. For example, the Constrained Application Protocol (CoAP) from CoRE provides a web framework realized through a subset of Representational State Transfer (REST) primitives. By running over UDP, it provides a lightweight transport solution with no connection establishment phase and small overhead. However, CoAP targets a limited class of applications (i.e., manipulation of simple resources) and requires devices to support a full web stack implementation, which might be prohibitive for different devices. Moreover, strategies such as header compression done in 6LoWPAN can imprint processing requirements over low-powered devices.

The receiver-driven communication model of ICN, coupled with anycasting and in-network caching, can help retrieve contents even in constrained networks with low duty-cycle providers. In fact, a request can be satisfied by another node, holding a copy of the data, when the producer is in doze/sleeping mode. Furthermore, distributed caching may avoid massive data access to constrained devices, thus saving energy resources. Native multicasting also matches the goal of reducing the amount of traffic and interactions with energy-constrained nodes.

Mobility

Mobility support is a key requirement, for example, when IoT devices move aboard vehicles or are carried by humans.

IP mobility management solutions (e.g., Mobile IP) have been under continuous research, especially due to the explosion of mobile terminals. However, they have been commonly associated with scalability problems, leading to more efficient solutions (e.g., distributed mobility management), which have yet to reach adoption by mobile operators. In any case, the validity of such approaches in IoT scenarios has yet to be proved.

Thanks to its receiver-driven nature, ICN supports *consumer mobility*: when a consumer relocates, it can simply re-issue any unsatisfied request/subscription and be served by a different node. Moreover, ICN natively supports *host multi-homing*, so content requests or data delivery can use any of the interfaces (or even all simultaneously) available at the device.

In general, *producer mobility* entails additional signaling in ICN: it requires updates in intermediate forwarders (in the NBR case) or in entities managing name resolution. Such procedures may generate delays and disruption periods; however, anycasting, in-network caching, and multi-homing may greatly help in coping with the issue.

Heterogeneity

IoT is expected to be a highly heterogeneous environment, with a rich variety of devices, technologies, and services involving different stakeholders and manufacturers. The Internet will be traversed by huge amounts of IoT data generated by networked devices with widely different traffic characteristics. This implies added challenges to network providers regarding infrastructure planning, considering that the full extent of upcoming global IoT traffic is still unknown. Despite the flexibility of the narrow-waist design of IP and its ability to maximize interoperability, it becomes complex to apply common network functionality to the explosive number of technologies involved in the upcoming IoT.

Standardized ICN naming schemes for IoT would allow abstracting services and contents in order to hide the heterogeneity in underlying networks and devices, and facilitate interoperability among different players. For example, ICN naming has the potential to allow entities to request content by its name, independent of the type of service that provides and transports it from the source — message queuing telemetry transport (MQTT), CoAP, and Advanced Message Queuing Prptocol (AMQP). Furthermore, by decoupling consumers and producers and delivering self-consistent data packets, ICN can interconnect information, devices, and services under heterogeneous network scenarios.

ICN toward IoT

In spite of the ICN prospects for IoT, the uniqueness and complexity of IoT requirements raise challenges that require adaptations to the design of ICN protocols.

In the following we present such challenges and we discuss how they are addressed in the literature, by identifying remaining open issues.

Naming

An ICN naming scheme for IoT should be highly expressive and customizable, and it should expose service (e.g., sensing and action) and data features.

Hierarchical names have been mainly considered in the literature to support such properties [7, 11–13]. The basic idea is to define a hierarchy of name components that identify the IoT application (e.g., building management system, energy control) and the attributes that describe the related contents/services. In [11], for instance, the name of a sensor data */ucla.edu/bms/building/melnitz/studio/1/data/panel/J/voltage/<timestamp>* indicates the application (a building management system deployed at the University of California, Los Angeles), the physical location of the sensor (Panel J inside Studio 1, Melnitz Hall), the type of data (voltage), and the time of acquisition.

Similarly, in the case of actuation applications, commands/management parameters can be provided as named components [12]. As an example, the name carried in a request packet commanding an actuator to turn off the light in the Laboratory of Telecommunications could be */campus.edu/lighting/building1/floor2/tlclab/OFF*, where "campus.edu/lighting" indicates the application name prefix, "/building1/floor2/tlclab/" is the location of the actuator, and the component "OFF" means the action to perform.

Flat names are typically obtained through hash algorithms applied to (already existing) contents and can hardly be assigned to dynamic IoT contents that are not yet published. Differently, hierarchical names facilitate the request of dynamic contents that are generated *on demand* (e.g., a parameter measured by a sensor), provided that naming conventions have been specified during the system configuration/setup.

Through the hierarchy of name components, a simple *versioning* system can be deployed to manage those cases where a producer constantly updates the content value, like the temperature in a room. This would help to manage the freshness requirement of some applications [5].

However, hierarchical names are subject to length constraints, for instance, to fit the maximum payload size of some protocols such as ZigBee [6]. In parallel, variable-length names make line-speed name lookup extremely challenging. Especially under large-scale scenarios, naming schemes should be designed together with processing techniques (e.g., name component encoding) that accelerate name lookup [4]. This is fundamental to reduce content access latencies, crucial in safety-critical applications (e.g., smart transport and healthcare).

By sharing a common name prefix for multiple contents/services, hierarchical names scale better than flat names, since they facilitate the definition of name aggregation rules in the FIB, which is critical for *big data*. This implies that IoT applications operating in the same domain and handling information/services with *global* scopes should be designed by developers with common (shared) name-prefixes. In ICN deployments dealing with Internet contents, name prefixes are usually related to the top-level and second-level domain names that identify websites and their contents; for example, the prefix *youtube.com* is associated with every Youtube video. In IoT, instead, the name prefix can identify application types, physical locations, or other macro-categories that broadly identify groups of data and services. We are still far from a leading naming solution tackling the mentioned issues, and stakeholders are required to agree on some basic naming conventions.

Security

ICN security mechanisms that consider the unique features of IoT applications and device limitations must be defined.

First, some IoT applications require *queries from consumers to be authenticated*; for example, an actuator will execute an action, such as turning-on/off appliances, only if this is required by a trusted authorized entity. Currently, ICN security mechanisms are only applied over data packets and do not support request authentication. Second, IoT devices with low processing and memory capabilities hardly use resource-intensive public key cryptography.

Preliminary solutions to the raised issues can be found in the literature. According to [12, 13], security information can be embedded in request packets as the last name component

ICN component	Main references	Guidelines
Naming	[6, 7, 11–13]	• Identifying both services and contents associated with objects • Interoperability in the presence of multiple crossed domains and interested consumers • Short length names in the presence of resource-constrained networks • Support of on-demand generated and dynamic contents • Efficient name aggregation rules • Support for data morphing
Caching	[5, 7, 14]	• Caching/replacement policies accounting for IoT traffic peculiarities, for example, freshness requirements from both producers and consumers • Support for storage/battery-constrained nodes (e.g., off-path caching via cloud, IoT-specific decision policies)
Discovery and delivery	[6–8, 15]	• Hybrid NBR-LRS approaches • Cloud-assisted name resolution
Security	[12, 13]	• Authenticated data and request • Flexibility in cryptography algorithm selection
Morphing	[9, 10]	• Efficient data aggregation and filtering support without information loss

Table 3. Research solutions and directions for ICN components in IoT.

by leveraging the virtually unrestricted hierarchical name composition. However, authenticated requests increase the complexity of the security framework and at the same time increase the overall name length, which could not suit the payload size of low-power access technologies.

Specific lightweight solutions for encryption and authentication become fundamental for resource-constrained devices. In this context, *symmetric cryptography* can be useful [13]. The disadvantage lies in the inflexibility with respect to key management, as it requires pre-distribution of keys. A good trade-off between complexity and resource saving can be obtained by elliptic curve cryptography, the prevalent public key scheme currently considered for small devices.

Generally, ICN security functions for IoT must be flexible in the selection of cryptographic techniques, since there is no one-size-fits-all solution; the most appropriate one shall be chosen depending on device capability and application domain.

Caching

In-network caching acquires special significance in IoT domains. On one hand, caching is generally beneficial because it speeds up data retrieval and increases its availability. On the other hand, caching and related replacement operations can be quite expensive in terms of both processing and energy consumption. Therefore, a first question is whether caching should be enabled in any IoT device or only in powerful nodes. A simple design choice would forbid constrained devices to cache contents [3], but in [7] caching proved to be highly beneficial even when enabled in IoT nodes with small storage capacity. In fact, it reduces the number of (lossy) hops toward the producer by limiting the network load and the overall energy consumption. In addition, caching is viable since data generated by IoT devices typically have a small size and a short lifetime.

Overall, IoT data can be cached in *network routers* and *resource-constrained devices* by implementing caching decision and replacement policies that account for the peculiarities of IoT traffic, for example, compatibly with freshness requirements [5, 14] and device capabilities (e.g., residual battery level and storage).

Specifically, the indiscriminate "cache everything everywhere" approach has been proven inefficient due to the high level of content redundancy and poor utilization of the available cache resources. Alternative caching decision policies (e.g., probabilistic, as in [14, references therein]) have been proposed for more efficient usage of the available caching space, alleviating the load of nodes and bandwidth consumption. Many of these policies address the space usage issue by also considering the content *popularity* for topology-related centrality metrics. These could also be viable approaches in IoT, where the same content is requested by different applications. However, they work well under the assumption of static content, which is not entirely valid for IoT, where contents are usually *transient*. Therefore, caching decision policies for IoT should focus on improving dissemination speed rather than long-lasting caching, while data replacement policies that behave according to the content freshness requirements of the consumer applications and to the content generation pattern are suggested, being mindful that this additional elaboration may affect the correct line-speed packet processing [5].

In addition, ICN may resort to off-path caching (according to which caching points are along alternative paths) to alleviate the load on constrained IoT devices and proactively distribute contents in specific locations (e.g., the cloud) by preventing data redundancy at the cost of additional overhead for cache management.

Discovery and Delivery

Name-based routing and lookup-based resolution systems offered by ICN for content discovery may suit specific IoT scenarios, mainly depending on the content characteristics (e.g., popularity, dynamic generation) and network features (e.g., infrastructureless vs. infrastructured).

NBR, coupled with data delivery performed by maintaining some soft-state at each Interest forwarder, is suitable to access popular contents in infrastructured scenarios, while it is the only viable solution in isolated networks with an intermittently available or nonexistent infrastructure. Its inherent benefits are:

- Robust and resilient retrieval through adaptive forwarding coupled with in-network caching [6]
- Easy resource discovery in infrastructureless networks through direct inter-device communication by means of Interest packet broadcasting [7]

The downside of deploying NBR is mainly related to the growth of the FIB size and routing updates in the case of a huge number of names, and the overhead of maintaining the soft-state. However, name-prefix aggregation can successfully cope with such challenges, together with adaptive forwarding.

LRS is useful in infrastructured scenarios with unpopular and popular contents. LRS can also be beneficial with off-

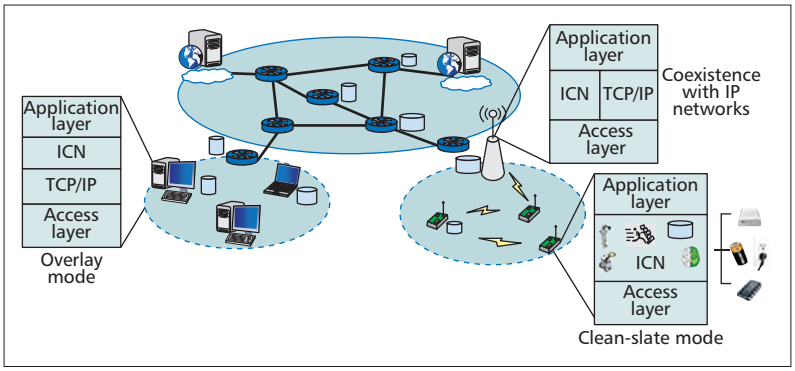

Figure 3. Main deployment options.

path caching; in fact, an intermittently connected IoT source can push data in a predefined always-on location (e.g., the cloud) accessible by consumers. In these cases, deploying a global name resolution service based on a hierarchically organized DHT is desirable, such as in the SAIL and PURSUIT architectures, while additional scalability properties for name lookup can be obtained by using data center capacities in the cloud, as suggested in [15].

In summary, NBR and LRS solutions may complement each other. Hence, by leveraging cloud computing, multi-level DHT, name-prefix aggregation and adaptive forwarding, an effective discovery and delivery platform can be provided with the potential to scale even for a huge number of IoT resources.

Morphing

According to [3], an ICN node does not (and should not) provide any data transformation (aggregation, filtering, etc.), because it should be kept outside the ICN domain to reduce inside complexity and function overloading. Notwithstanding, even if not explicitly mentioned among the ICN core principles, ICN could enable lightweight in-network data manipulation (we call it *morphing*) at intermediate nodes, by embedding semantics awareness at the networking layer [4, 9, 10].

The motivations for data morphing are manifold and particularly strong for IoT. First, there is a high abundance of raw data in the network, but consumers are likely to want to receive manipulated data. Second, morphing greatly simplifies data post-processing for those applications that either operate on aggregated data or need a complete knowledge base without information loss. Third, filtering/aggregating data helps to improve the scalability in content retrieval and reduce the network and device resource usage.

Hierarchical names could facilitate data aggregation and enable intermediate nodes to track and process packets, while meeting latency and accuracy demands. The requester could explicitly ask to retrieve aggregated data, allowing the network to select the best nodes providing them.

Otherwise, more powerful nodes perform such tasks in a transparent way for the consumers (e.g., a concentrator replies to the utility company with aggregated energy consumption data from multiple users). Morphing would be expected only at carefully selected locations in order to trade off between effectiveness and computational resource demands.

Conclusions

The surveyed literature provides preliminary ICN design solutions to face IoT requirements and opens several research opportunities. Although this area is still taking shape, we strive to summarize the main research directions for ICN core components in Table 3.

We have learned that the features and requirements of IoT, along with the forming nature of the IoT concept and recent Internet evolutions, pave the way to deeper investigation.

Finally, an important question still remains concerning the practical deployment of ICN. Broadly speaking, ICN solutions can be deployed as overlay over the existing IP infrastructure, or as clean-slate implementation directly over access layer technologies to replace IP. Overlay solutions are discouraged due to their complexity and the overhead for overlay management and encapsulation inside IP protocols. This is also inadvisable for resource-constrained devices, representing a high percentage of IoT objects. A clean-slate solution can easily be deployed where there is no need to communicate with IP-based nodes (e.g., in isolated vehicular environments) or to maintain backward compatibility, but it raises concerns when global access and connectivity are required.

A most likely short-term design would allow coexistence with IP-based technologies. Similarities between ICN hierarchical names and URIs of web resources could facilitate such coexistence. The translation between them may be implemented easily in the node (e.g., a smart home gateway) interfacing ICN islands and the rest of the Internet (Fig. 3).

In conclusion, we can state that ICN holds promise as a candidate networking solution for IoT. Outstanding issues must be addressed in order to meet expectations, among which is the classic scepticism in pursuing a revolutionary instead of an evolutionary approach. We are confident that an open-minded view of ICN's value and role would bring great benefits to the challenging IoT research field.

Acknowledgments

This work was partially funded under grants: PON03PE_00050_2 DOMUS, MIUR, and UID/EEA/50008/2013, FCT.

References

[1] Z. Sheng *et al.*, "A Survey on the IETF Protocol Suite for the Internet of Things: Standards, Challenges, and Opportunities," *IEEE Wireless Commun.*, vol. 20, no. 6, 2013, pp. 91–98.
[2] B. Ahlgren *et al.*, "A Survey of Information-Centric Networking," *IEEE Commun. Mag.*, vol. 50, no. 7, 2012, pp. 26–36.
[3] A. Lindgren *et al.*, "Applicability and Trade-Offs of Information-Centric Networking for Efficient IoT," IETF Internet Draft, Jan. 2015.
[4] Y. Zhang *et al.*, "ICN Based Architecture for IoT — Requirements and Challenges," IETF Internet Draft, Nov. 2014.
[5] J. Quevedo, D. Corujo, and R. Aguiar, "Consumer Driven Information Freshness Approach for Content Centric Networking," *Proc. IEEE NOM*, 2014.
[6] M. Amadeo *et al.*, "Named Data Networking for IoT: An Architectural Perspective," *Proc. European Conf. Networks and Commun.*, Bologna, Italy, 2014.
[7] E. Baccelli *et al.*, "Information Centric Networking in the IoT: Experiments with NDN in the Wild," *ACM Conf. Information-Centric Networking*, 2014.
[8] S. Li *et al.*, "A Comparative Study of MobilityFirst and NDN Based ICN-IoT Architectures," *Proc. 10th IEEE Int'l. Conf. Heterogeneous Networking for Quality, Reliability, Security and Robustness*, 2014, pp. 158–63.
[9] K. V. Katsaros *et al.*, "Information-Centric Networking for Machine-to-Machine Data Delivery: A Case Study in Smart Grid Applications," *IEEE Network*, vol. 28, no. 3, 2014, pp. 58–64.
[10] N. Fotiou and G. C. Polyzos, "Realizing the Internet of Things Using Information-Centric Networking," *Proc. 10th IEEE Int'l. Conf. Heterogeneous Networking for Quality, Reliability, Security and Robustness*, 2014, pp. 193–94.
[11] W. Shang *et al.*, "Securing Building Management Systems Using Named Data Networking," *IEEE Network*, vol. 3, no. 28, 2014, pp. 50–56.
[12] J. Burke *et al.*, "Securing Instrumented Environments over Content-Centric Networking: The Case of Lighting Control and NDN," *Proc. IEEE NOMEN Wksp.*, 2013.
[13] J. Burke *et al.*, "Secure Sensing over Named Data Networking," *Proc. IEEE Network Computing and Applications*, 2014, pp. 175–80.
[14] S. Vural *et al.*, "In-Network Caching of Internet-of-Things Data," *Proc. IEEE ICC*, 2014.
[15] X. Vasilakos, K. Katsaros, and G. Xylomenos, "Cloud Computing for Global Name-Resolution in Information-Centric Networks," *Proc. 2nd Symp. Network Cloud Computing and Applications*, 2012, IEEE, 2012, pp. 88–94.

Biographies

MARICA AMADEO is a postdoctoral researcher at University Mediterranea of Reggio Calabria, Italy. She received a B.S. degree (2005) and an M.S. degree (2008) in telecommunications engineering from the University Mediterranea of Reggio Calabria, and a Ph.D. degree in 2013 from the same university. Her major research interests are in information-centric networking and wireless ad hoc networks.

CLAUDIA CAMPOLO is an assistant professor of telecommunications at University Mediterranea of Reggio Calabria, Italy. She received an M.S. degree in telecommunications engineering (2007) and a Ph.D. degree (2011) from the same university. She was a visiting Ph.D. student at Politecnico di Torino (2008) and a DAAD fellow at University of Paderborn, Germany (2015). Her main research interests are in vehicular networking and future Internet architectures.

JOSÉ QUEVEDO graduated in 2009 from the Polytechnic University of Havana (CUJAE), Cuba, where he studied telecommunications and electronics engineering. He is currently pursuing his Ph.D. and working as a researcher in the Advanced Telecommunications and Networks Group at the Instituto de Telecomunicacõoes, Universidade de Aveiro, Portugal. His current research areas include information-centric networking and the Internet of Things.

DANIEL CORUJO is a research fellow in the University of Aveiro, where he concluded his Ph.D. in communication models for the future mobile Internet in 2013. He previously worked at Nokia Siemens Networks and as an IMS deployment executive for the research branch of Portugal Telecom. He is pursuing research areas in mobility mechanisms for heterogeneous networks and the future Internet. He contributes to the IEEE 802.21 and IRTFG's ICNRG standardization work groups.

ANTONELLA MOLINARO has been an associate professor of telecommunications with the University Mediterranea of Reggio Calabria since 2005. Previously, she was an assistant professor with the University of Messina (1998–2001) and the University of Calabria (2001–2004), and a research fellow at the Polytechnic of Milan (1997–1998). She was with Telesoft, Rome (1992–1993) and with Siemens, Munich (1994–1995) as a CEC Fellow in the RACE-II program. Her current research focuses on vehicular networking, information-centric networking, and future Internet.

ANTONIO IERA [SM'07] graduated in computer engineering from the University of Calabria, Italy (1991), and received his Master's diploma in information technology from CEFRIEL/Politecnico di Milano (1992) and Ph.D. degree from the University of Calabria (1996). Since 1997, he has been with the University of Reggio Calabria, and is currently a full professor of telecommunications and director of the Laboratory for Advanced Research into Telecommunication Systems. His recent research focuses on RFID systems and the Internet of Things.

RUI L. AGUIAR received a Ph.D. degree in electrical engineering in 2001 from the University of Aveiro. He is currently a full professor at the same university, responsible for networking aspects. His current research interests are centered on the implementation of advanced networks/systems with special emphasis on future Internet and mobile architectures. He has more than 400 published papers in those areas. He is a member of the 5GPP Infrastructure Association, and is/has been extensively involved in several conferences.

ATHANASIOS V. VASILAKOS is a professor with the Luleå University of Technology, Sweden. He has served or is serving as an Editor for many technical journals, including *IEEE Transactions on Network and Service Management*, *IEEE Transactions on Cloud Computing*, *IEEE Transactions on Information Forensics and Security*, *IEEE Transactions on Cybernetics*, *IEEE Transactions on Nanobioscience*, *IEEE Transactions on Information Technology in Biomedicine*, *ACM Transactions on Autonomous and Adaptive Systems*, and *IEEE Journal on Selected Areas In Communications*. He is General Chair of the European Alliances for Innovation.