
Information-Centric Networking for Machine-to-Machine Data Delivery: A Case Study in Smart Grid Applications

Konstantinos V. Katsaros, Wei Koong Chai, Ning Wang, George Pavlou, Herman Bontius, and Mario Paolone

Abstract

Largely motivated by the proliferation of content-centric applications in the Internet, information-centric networking has attracted the attention of the research community. By tailoring network operations around named information objects instead of end hosts, ICN yields a series of desirable features such as the spatiotemporal decoupling of communicating entities and the support of in-network caching. In this article, we advocate the introduction of such ICN features in a new, rapidly transforming communication domain: the smart grid. With the rapid introduction of multiple new actors, such as distributed (renewable) energy resources and electric vehicles, smart grids present a new networking landscape where a diverse set of multi-party machine-to-machine applications are required to enhance the *observability* of the power grid, often in real time and on top of a diverse set of communication infrastructures. Presenting a generic architectural framework, we show how ICN can address the emerging smart grid communication challenges. Based on real power grid topologies from a power distribution network in the Netherlands, we further employ simulations to both demonstrate the feasibility of an ICN solution for the support of real-time smart grid applications and further quantify the performance benefits brought by ICN against the current host-centric paradigm. Specifically, we show how ICN can support real-time state estimation in the medium voltage power grid, where high volumes of synchrophasor measurement data from distributed vantage points must be delivered within a very stringent end-to-end delay constraint, while swiftly overcoming potential power grid component failures.

The rationale behind information-centric networking (ICN) [1] is that information consumers are mainly interested in the information itself rather than the explicit network location of the data/content source (e.g., the host IP address). As such, the primary concerns of the network will no longer be on the reachability between specific hosts but on efficient information dissemination and retrieval. Accordingly, the ICN design principle has put information/data at the center of the networking landscape where information is published, resolved, delivered, and stored natively based on names rather than explicit host locations. This in turn enables a series of desirable features such as the support of in-network

caching and multicast forwarding, as well as native security protection and mobility support, thanks to the spatiotemporal decoupling of the communicating entities where data producers and consumers are agnostic to where and when the data will be published/consumed by their counterparts.

Given that the ICN paradigm has mostly catered for supporting content distribution operations in the public Internet, the ICN concept has been arguably regarded as a key feature in the design of future Internet architectures [1]. While it is still debatable whether this will become a reality, proposals have also been made for applying ICN to alternative application domains such as machine-to-machine (M2M) smart grid communications [2–4]. Similar to the current Internet, today's power grid communications are based on the host-centric model for data exchange between specific machines in the centralized supervisory control and data acquisition (SCADA) environment. In this context, we highlight the following challenges in current smart grid communication infrastructures.

Decentralized large-scale data sharing: The SCADA system is expected to face distinct challenges with increased participation of new stakeholders (e.g., solar/wind farm owners) and active power *prosumers*¹ introduced into the grid operations. It will be very common that data originating from one device is of interest to multiple entities participating in differ-

Konstantinos V. Katsaros, Wei Koong Chai, and George Pavlou are with University College London.

Ning Wang is with the University of Surrey.

Herman Bontius is with Liandon B.V.

Mario Paolone is with École Polytechnique Fédérale de Lausanne.

¹ Entities acting both as producers and consumers of energy.

ent smart grid applications, including existing and future emerging services. Similarly, a single entity involved in one or multiple applications in the grid may also need to access data originated from a large number of devices. With the introduction of such one-to-many, many-to-one, or even many-to-many communications, the traditional host-centric model will suffer from increasing complexity stemming from the explicit (usually pairwise) communication between involved hosts.

Heterogeneous requirements in distributing smart grid data: Emerging smart grid applications present diverse requirements on quality of service (QoS) ranging from low data rate, delay/disruption tolerant (e.g., smart metering and energy pricing) to higher-data-rate ones with stringent delay/disruption requirements (e.g., synchrophasor measurements). Today's communication infrastructure is merely a "bit carrier" with awareness of data delivery requirements being expressed by (static) topological primitives (e.g., virtual private networks, VPNs). The lack of advanced features able to differentiate network behavior on an application/data level, such as in-network storage and processing, has led to higher complexity/overhead on the device side, as well as inefficient use of underlying network resources.

Network security requirements: Based on the host-centric model, communication parties need to know each other's network location (i.e., IP address) to transmit data. Such exposure of IP address in mission-critical power grid applications may introduce vulnerability to network intrusions and denial-of-service (DoS) attacks [4].

In this article, we advocate the introduction of ICN in smart grid applications, given that such applications are mainly interested in just *what* — in terms of the power grid data — rather than *where* (i.e., the specific network address of the data source). Taking grid measurement applications as a typical example, although the identity of measurement points is of importance, this is usually indicated by a standardized identifier code that uniquely identifies the data stream and the physical device but not its network IP address. This information forms part of the payload (i.e., the *what*) for the communication network, regardless of the applied networking paradigm. Moreover, in many popular smart grid applications (e.g., power consumption measurements for the support of demand response), data may be required in the form of aggregates (e.g., consumption over a particular power grid area), further decoupling the delivered information from the *exact* location of its origin in the network.

In this context, the introduction of ICN, including its inherent publish/subscribe communication primitives, enables a higher degree of flexibility in supporting data sharing and smart grid control. In the long term, an ICN-based approach is expected to facilitate the support of complex and evolving data delivery structures (e.g., many-to-many communications for data sharing) due to the introduction of new applications/devices to the grid, since, in contrast to the current host-centric model, it does not focus on the establishment of explicit communication sessions. This is expected to significantly reduce system (re-)configuration complexity with the involvement of new services. It further applies in shorter time scales as well; that is, a change in the grid topology due to fault in a line or maintenance that requires an asset change will not affect the ongoing data delivery operations. For instance, a new data consumer is only required to subscribe to the data of interest and the rest of the communication will be automatically set up. Such automation minimizes manual reconfigurations and cuts down on possible human errors. Another example is the vision of on-demand islanding operations by distribution network operators (DNOs) that require dedicated monitoring and control infrastructures. With ICN,

the data can be cached for later retrieval while the island is formed and all flows redirected to the new subscribers once the islanding maneuver is completed. Data caching, possibly with local processing, as facilitated by a data/information-aware network, can also contribute to improving efficiency and facilitating QoS support. For instance, network nodes may adapt rates of measurement data targeting different grid operations rather than having such functions at the end devices. In addition, upon anomaly events (e.g., failures), faster response times can be achieved by allowing affected devices to locally fetch recovery instruction/data that is actively cached at nearby network entities, or quickly diverting affected data flows at intermediate nodes instead of reconfiguring the data producers themselves. As far as security is concerned, ICN offers intrinsic support for cyber-security in the power control system, as the identity along with the network and physical location of machines can be encrypted as part of the payload and are therefore not exposed [4]. In this case security becomes an integral part of the underlying network infrastructure rather than an *a posteriori* patch.

In this article we present an overlay ICN-based communication framework for supporting machine-to-machine (M2M)-oriented smart grid applications, based on publish/subscribe operations and the notion of location-independent *topics*. We shed light on the challenges faced in this emerging networking environment and elaborate on how these can be addressed by ICN communication primitives. Besides qualitative advantages, we also quantitatively illustrate the benefit of the proposed framework, focusing on the use case of phasor measurement units (PMUs)-based real-time state estimation (RTSE) in medium voltage (MV) power distribution networks. Accordingly, our simulations are based on two real European power grid topologies, and our results show that:

- With careful planning and provisioning of network resources, ICN can successfully support the requirements of such mission-critical communications.
- ICN communication primitives can substantially reduce the complexity of reconfiguration operations in cases of power grid component failures.

Machine-to-Machine Smart Grid Communication

The evolution of power distribution networks toward so-called active distribution networks (ADNs), shown in Fig. 1, requires the availability of suitable energy management systems (EMSs) to achieve specific operation objectives [5], such as:

- Optimal voltage/line-congestion controls
- Fault detection and location
- Post-fault management
- Local load balance
- Network losses minimization

These operations are significantly improved if the system state is known. In RTSE, large volumes of raw synchrophasor measurement data are collected by geographically distributed PMUs, strategically deployed in the power grid infrastructure to ensure full grid state *observability*. The UTC-synchronized data is continuously streamed to phasor data concentrators (PDCs), which collect synchrophasor data and other quantities (synchrophasor frequency, rate of change of frequency, powers, etc.) measured by PMUs and transmit them to other relevant applications.

With this specific real-time monitoring approach, PDCs located at different substations periodically report state estimation to the central SCADA entity at a lower rate than RTSE. In addition, power protection relays (that can be collo-

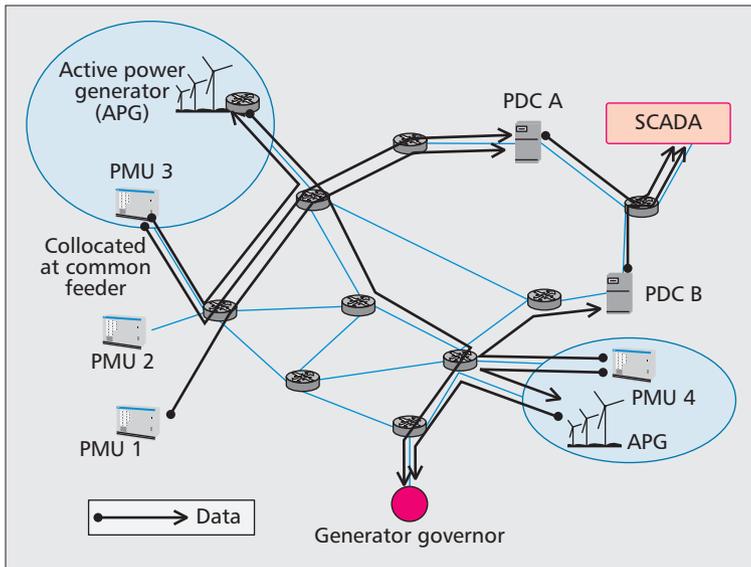


Figure 1. Current point-to-point data delivery in smart grid.

cated with PDC, and hence not shown in the figure) are also interested in such data to detect and react to potential anomalies ensuring the seamless operation of the power grid. Such anomalies (e.g., short circuits in the power grid) may result in changes of the power grid topology with the purpose of restoring operation and avoiding cascading effects (i.e., by opening/closing circuit breakers), resulting at the same time in changes for the monitoring data flows as well (e.g., a PMU device may need to direct its data to a different PDC). For these types of data receivers, there are stringent requirements on both data frequency (50 synchrophasor measurements/s) and end-to-end data latency (maximum 20 ms).

In the case of renewable power generation units, the synchrophasor data may also be fed to local active power generators (APGs) for them to adjust power generation operations. The data frequency required for the APGs is much lower (e.g., 2 measurements/s) with no strict end-to-end latency requirement. Meanwhile, individual APGs also transmit information about their locally generated renewable power to the generation governor.

From Fig. 1, it is evident that the current power grid communication is suffering from several critical deficiencies. Since it is still based on the current host-to-host model, dedicated point-to-point communication sessions need to be maintained, and this requires complex per-machine configurations. At the same time, network resources may be wasted when different recipients request different rates of the same data flows in separate communication sessions. Complexity and the associated power grid control overheads further increase when anomalies call for rapid reconfigurations of data flows (i.e., diverting traffic of affected PMUs to alternative PDCs requires each PMU device to be reconfigured individually). Finally, the exposure of IP addresses of individual mission-critical entities also makes them vulnerable to DoS attacks.

ICN Framework for Smart Grids

Based on the aforementioned communication requirements in smart grids, we identify a favorable match with some major ICN design aspects. We exploit this match and apply ICN as an overlay for enabling resilient and seamless communication in smart grids. The decoupling of information from location and time fits the communication patterns of the considered applications, yielding opportunities for a simplified and efficient management of communication flows. Based on the

inherently supported pub/sub communication primitives, ICN introduces a degree of indirection between the communicating end hosts by enabling the network to actively mediate information delivery, similar to [9]. Apart from previously investigated security-related benefits [4], this practically translates to the ability to:

- Simplify both the establishment and reconfiguration of communication flows in the aforementioned multi-party communications (e.g., delivering synchrophasor measurements from PMUs at different feeders to PDCs and APGs), including the introduction of new devices interested in the data produced by legacy elements.
- Facilitate multi-criteria traffic management decisions, that is, selecting one or more indirection points based on the underlying transmission capabilities, application requirements, network conditions, topology characteristics, and so on (e.g., selecting only delay sensitive data to forward on high data rate links).

- Enable in-network management of smart grid data, including caching and processing such as rate adaptation, aggregation, filtering, and so on (e.g., enabling the rate adaptation of PMU measurements at an indirection point close to an APG).

- Enhance resilience of information delivery to protect the grid against anomalies/power failures and subsequently minimize power distribution disruption.
- Enhance security by avoiding the exposure of critical components' network locations through the means of indirection.

In these cases, the information-centrism of the network, expressed by the fundamental role of topics and their attributes, enables network operations to take place on an information level, bridging the gap between the application requirements, and the underlying technological and topological characteristics. Concurrently, enabling these ICN features in an overlay fashion facilitates adoption and deployment of the ICN principles, especially in a considerably heterogeneous environment where smart grid networks are often based on a set of diverse communication technologies.

In the following, we describe an instantiation of a smart grid communication platform based on the aforementioned ICN concepts that can support heterogeneous smart grid applications, including ones with stringent real-time requirements. We use common ICN building blocks and primitives, and describe the specific design requirements and challenges of the M2M and mission-critical applications described above.

Building an ICN-Enabled Smart Grid Communication Platform

Figure 2 presents a logical illustration of a topic-based ICN smart grid infrastructure to support data dissemination across heterogeneous entities. We follow the pub/sub paradigm that is inherent in ICN schemes for supporting communication between the smart grid entities. For instance, in one form or another, various ICN projects (e.g., PSIRP, PURSUIT, and COMET) employ similar pub/sub mechanisms to harness the benefits of ICN in terms of flexibility in communication and the added security features. In our proposal, communication is organized in location-independent topics that uniquely identify semantically related data. Each topic is associated with a set of attributes such as spatiotemporal information and reporting rate (where applicable). A topic resolution system handles data publication and subscription for interested receivers to access data published to a topic but without

directly contacting the publishers. In this system, which is conceptually similar to resolution systems in existing ICN schemes [1, 8–10], we follow the separation of control and data planes based on homogeneous entities. Specifically, for each topic, the resolution function can be located at a node other than the one that is responsible for the forwarding function, and each node can flexibly be responsible for resolution and data forwarding of different topics. In addition to plain forwarding on the data plane, we further build on the information-centric primitives of the architecture to enable in-network processing of data subject to the topic to which they belong, as well as their attributes. This may include aggregation of data being disseminated in multi-source-based topics, multi-criteria filtering (e.g., location, time), and rate adaptation according to heterogeneous receivers demanding different data reporting frequencies [11].

Each topic is dedicated to delivering data from their publishers to interested subscribers, possibly with different data requirements. PMU measurement data is one example topic. A subscriber can either subscribe to the entire topic (i.e., all data published under the topic) or a subset of the data by providing the attributes describing the desired data (e.g., measurements from specific feeders or the required rate of the measurements). Data publishers (e.g., PMUs) use the *publish* primitive to disseminate their measurement data to the topic. Information consumers (e.g., PDCs, APGs) employ the *subscribe* primitive to initiate reception of data from the topic.

Consider Fig. 2 again, which illustrates both the publication and subscription procedures. All publishers and subscribers connect to the network via a designated node (DN), which is ICN-aware for handling topics. One or more ICN-aware nodes act as rendezvous points (RPs) for each topic. Publishers publish their data to an RP, which is responsible for forwarding data to subscribers (red solid arrows), as well as potentially storing, caching, and/or processing it.

- A publisher (e.g., in Fig. 2, PMUs in topic 1, APGs in topic 2 and PDCs in topic 3) wishing to disseminate data first sends out a publish request (blue dotted arrows) to its DN with the topic identifier. The topic resolution from the DN to the resolver can be based on specific techniques such as geographical hash table (GHT)-based approaches [4]. Then the resolver replies to the DN with information about the corresponding RP that hosts the corresponding topic. The publisher then publishes its data to the RP from where the data can be forwarded to subscribers.

- A receiver (e.g., in Fig. 2, PDCs in topic 1, generator governor in topic 2, and SCADA in topic 3) sends its subscribe request (purple dashed arrows) to its local DN, specifying the identifier of the topic in which it is interested. Based on the same topic resolution mechanism as the publisher, the request is routed from the DN to the resolver. The resolver then finds and informs the corresponding RP for forwarding the published data to the subscriber's DN.

Complex topic subscription is also supported. For instance, in Fig. 2, three PMUs publish their data stream to *topic 1*, for which PDC A and APG A are subscribers. However, APG A, being collocated in the same feeder with PMU 3, is only inter-

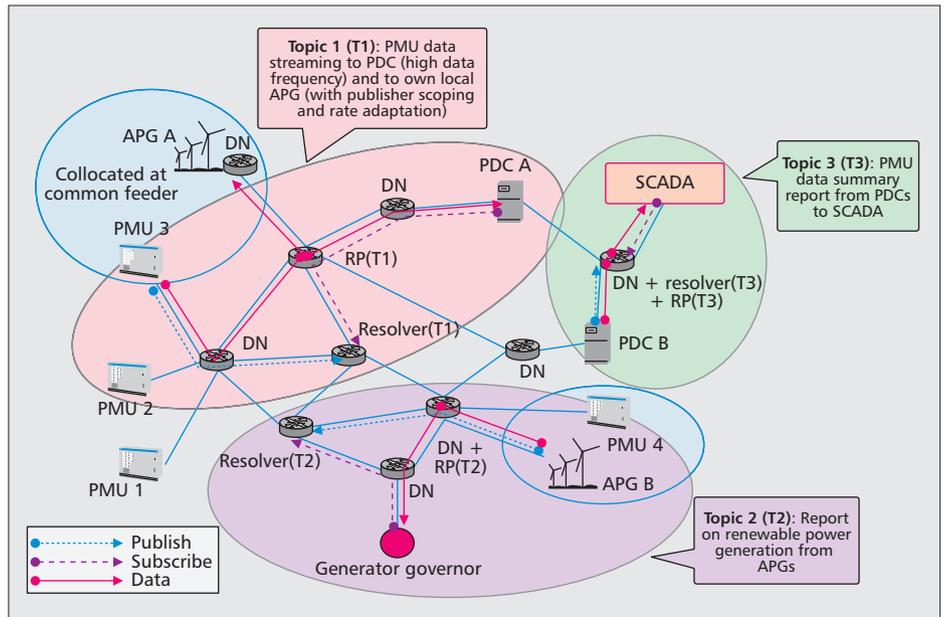


Figure 2. ICN-enabled M2M communication in smart grid for wide area monitoring.

ested in the data published from that specific PMU. As such, the RP for this topic needs to perform in-network publisher scoping to ensure the subscriber receives only the interested data. This scoping function can be flexibly configured using topic attributes denoted in the subscribe requests. This scoping is based on the semantically structured attribute information instead of network addresses. In the case of the current host-centric model, PMU 3 would instead be configured to establish one separate communication session per recipient (i.e., PDC A and APG A), thus increasing configuration complexity and bandwidth consumption. Furthermore, subscribers may require data in different granularity/rate. For instance, PDC A needs comprehensive fine-grained data at high frequency (e.g., 50 messages/sec), while APG A requires data at much lower frequency (e.g., 2 messages/sec). To deal with these heterogeneous requirements, the RP performs in-network rate adaptation according to the specific subscriber requirements. Such in-network data processing functions offer extra advantages for both bandwidth resource conservation compared to the current dedicated host-to-host communications, and the distribution of processing load in the network. The other two topics in the figure are used respectively for the PDCs to report their own summarized PMU data to the central SCADA, and for the APGs to report the generated power information to the generator governor. The corresponding RPs can also cache passing-through topic data for future localized data access on new subscriber requests. Moreover, cases of power grid anomalies, can be handled efficiently by simply reconfiguring subscriptions (e.g., in Fig. 2), a power grid component failure may result in a change of the power grid topology such that PMUs 1 and 2 need to deliver their traffic to PDC B instead; in this case, it suffices for a single subscription request issued by PDC B to the corresponding RP, whereas separate reconfiguration message per PMU would be required in the case of host-centric communications, increasing the complexity and the associated delays, as shown later.

Advanced ICN Data Forwarding with Multiple RPs

We also propose the option of having multiple RPs for handling data dissemination within a topic. In this case, each publisher in a topic always publishes its data to one of the RPs determined by the resolver, while a subscriber may be connected to either all or a subset of RPs for that topic. For

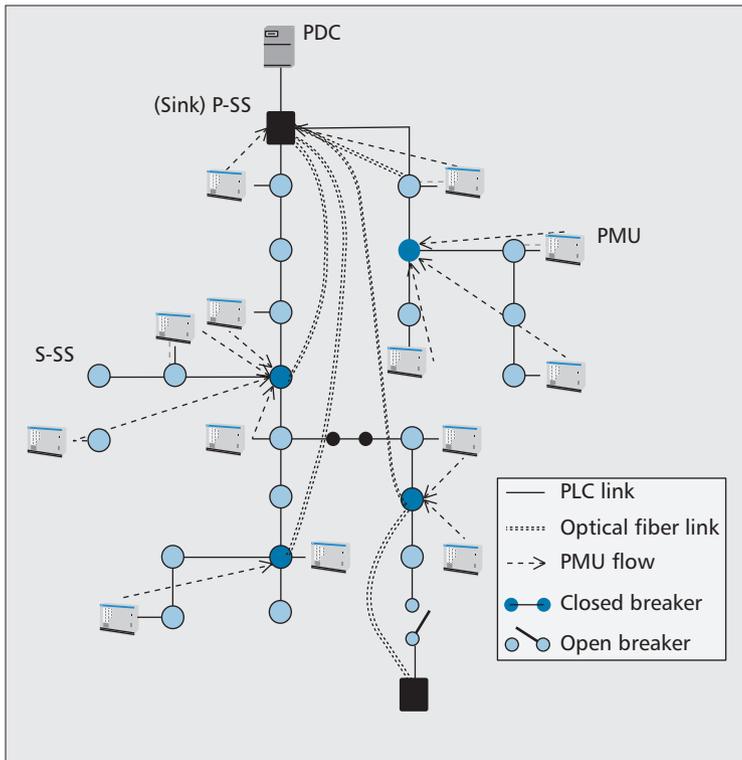


Figure 3. Example MV distribution network.

instance, an APG may receive data from one RP if it is only interested in the measurements from the PMU(s) publishing to that specific RP. From this perspective, having multiple RPs naturally enables the publisher scoping function.

We exploit the multiplicity of RPs for both resiliency and resource management in smart grids. By distributing the in-network data handling function to different RPs, we add redundancy and availability and thus eliminate a single point of failure for transmitting and processing data. At the same time, by having multiple RPs for each topic, both data transmission and processing loads can be appropriately shared.

The determination of the optimal number and locations of the RPs must take into account various factors. A key factor relates to the locations of publishers and subscribers. For instance, placing a topic RP strategically at a point that maximizes the opportunity for common delivery paths (e.g., near many clustered subscribers) can save network bandwidth. This provisioning function is computed by the resolver without client awareness. The overall traffic and processing load per topic (e.g., topic data rate), as well as across topics (i.e., current network/node conditions), also constitutes an important input factor, enabling a series of load balancing decisions. The selection of the RPs is also constrained by the underlying *physical* network topologies and technologies. For instance, both the number and locations of RPs for the topics of delay-sensitive applications (e.g., PMU data) need to be selected taking into account the data transmission capabilities (specifically, delay and bandwidth support) of the underlying communication technology. This is considered as a particularly beneficial capability in the context of the technological and topological heterogeneity of smart grid communication networks i.e., a diverse set of technologies coexists in current power distribution networks, including power line communication (PLC) and GSM/GPRS. At the same time, newer technologies, such as optical fiber, WiMAX, and Long Term Evolution (LTE), are being considered [13, 14].

It is important to note that the selection of RPs is taking place on a topic level. Directly derived from the ICN nature

of the proposed framework, this feature inherently allows the formation of an information-centric management plane where application requirements (expressed through different topics) meet the varying underlying network capabilities/conditions. This is considered a valuable tool in smart grids, enabling the handling of diverse applications under diverse network conditions/capabilities.

Feasibility Study and Performance Benefits Brought by ICN

To gain insights on the benefits of the aforementioned multi-RP selection mechanism and further derive data plane resource provisioning strategies for supporting ICN-based communications, we have conducted a feasibility study on a real-time smart grid application with an extremely low latency requirement. Focusing on a network infrastructure with heterogeneous communication technologies, we demonstrate the ability of our ICN-based approach to adapt to the available network substrate capabilities and select the communication technologies that best serve the application requirements.

For our feasibility study, we focus on the case of the RTSE application for MV power distribution networks [6]. State estimation is based on data reported by PMUs deployed at a secondary substation (S-SS) level

(Fig. 3). A typical deployment scheme to guarantee the observability of the power grid would involve a PMU (as a topic publisher) at approximately every two S-SSs along a cable line (i.e., feeder). The PDC (as the topic subscriber) is typically connected to the primary substation (P-SS). The P-SS and the S-SSs act as DNs for the directly attached PMU/PDC devices.

To achieve the stringent delay requirement (i.e., 20 ms), we first assess the capabilities of the underlying communication infrastructure. To this end, our study is based on real MV power grid topologies from two European distribution networks. The topologies have a tree structure rooted at the P-SS, which performs the high-to-medium voltage transformation. S-SSs are responsible for delivering the power to the low voltage network. Circuit breakers deployed at strategically selected grid locations are used in cases of failures to change the power grid topology, effectively connecting the affected S-SSs to an alternative P-SS (Fig. 3). Table 1 summarizes the topological characteristics of the distribution networks we use.

It is crucial to note that the exact form of the communication network depends on the employed communication technology. In the simplest case, the use of PLC results in a tree topology for the communications network, coinciding that of the power grid. However, due to the bandwidth constraints of PLC technology (i.e., at the order of a few hundred kilobits per second [12]), as well as the lengths of the S-SS to P-SS delivery paths, which accumulate the corresponding packet processing and propagation delays, we resort to alternative network infrastructures for satisfying stringent delay requirements.

To this end, we consider a hybrid network infrastructure employing both PLC and optical fiber technologies. A limited number of S-SSs is selected as traffic aggregation points, or sink nodes, connected to one or more P-SSs with direct optical fiber links.² This is to take advantage of the existing low-cost PLC capabilities so as to aggregate traffic to a limited number of sink nodes, thus reducing the corresponding optical fiber deployment costs. We formulate the S-SS selection process as a distance constraint version of the p-center facility location problem [15]; that is, defining the minimum (hop)

distance of each S-SS to its closest sink node as the communication cost, our target is to find the minimum number of sink nodes such that the maximum communication cost (d_{\max}) in the network is minimized. We solve this optimization problem following a sequential location procedure tailored for tree topologies [15]. An example solution is presented in Fig. 3 for $d_{\max} = 3$.

Based on the resulting set of sink nodes/RPs, we simulate the delivery of RTSE traffic within the proposed ICN framework. Following the PMU deployment scenario described above, we simulate the operation of 118 and 69 PMUs on topologies 1 and 2, respectively. Considering a PMU message payload of 86 bytes, as well as all protocol stack header overheads, we calculate the overall data rate for each RTSE PMU flow at the link layer to be 64.8 kb/s, for a fixed 50 Hz reporting rate.

To support the low latency requirements of this topic, the resolver selects the RP of each PMU to coincide with the closest sink node. Note that even though this RP selection takes topological (i.e., hop distance to PMU) and technological (i.e., employing optical fiber instead of PLC links) characteristics into account, it is still performed on a topic basis (i.e., regarding delay-sensitive information).

Figure 4 shows the cumulative distribution function (CDF) of the end-to-end delay observed in the case of 500 kb/s PLC for a duration of 20 min. For each optical fiber link, we assume 10 Gb/s of available bandwidth. We show the end-to-end delay for the following scenarios:

1. *Plain* PLC scenario (i.e., no optical fiber)
2. *Hybrid*, $d_{\max} = 1$, which results in the use of 67 and 40 sink nodes in topologies 1 and 2, respectively.
3. *Hybrid*, $d_{\max} = 2$, which results in the use of 41 and 27 sink nodes in topologies 1 and 2, respectively.
4. *Hybrid*, $d_{\max} = 3$, which results in the use of 30 and 2. The methodology followed in this study is also applicable in the case of wireless technologies. However, we focus here on optical fiber technologies as recent studies have shown that wireless technologies such as WiMAX and LTE introduce significant control plane and MAC layer delays [13, 14]. The investigation of such aspects is beyond the scope of this article. 18 sink nodes in the topologies 1 and 2, respectively.
5. *Optical fiber*, where all S-SSs act as sink nodes; effectively, this scenario results in the selection of one RP per publisher PMU.

We observe that *hybrid* scenarios with $d_{\max} = 2$ can satisfy the delay requirement of the selected application, resulting in the deployment of a total of 68 optical fiber links. The *optical fiber* architecture satisfies the requirement but at the cost of the deployment of 306 optical fiber connections for both areas.³ However, it is important to note that PLC data rates raise barriers in the support of RTSE delay requirements, with the *plain* PLC evidently failing to support the desired latency. For even lower PLC data rates, such as 100 kb/s (not shown here due to length limitations), we observed that even $d_{\max} = 1$ cannot keep the total delay below 20 ms, although it keeps it always below 21.53 ms.

Having examined the feasibility of supporting the end-to-end delay requirements of the considered RTSE application

	Topology 1	Topology 2
Primary substations (P-SS)	1	1
Secondary substations (S-SS)	190	114
Number of edges (cables)	190	114
Average cable length	684 m	763 m
Average S-SS to P-SS path length (hops)	Avg.: 7.29; Max: 19	Avg.: 7.69; Max: 20
Average node degree	1.989	1.983

Table 1. Topology characteristics of two MV grid areas in the Netherlands.

in the context of the above described ICN framework/architecture, we further look into the benefits introduced by its adoption. For this, we consider the scenario of power grid component failures requiring the redirection of data flows of the affected PMUs to an alternative PDC, in the context again of the RTSE application. As argued earlier, the current host-centric networking paradigm would necessitate one reconfiguration message per affected PMU. In contrast, the proposed multi-RP ICN framework requires only a subscription message per involved RP.

Utilizing the same MV grid topologies as in the previous experiments, we measure the number of reconfiguration messages in a series of random failures. We consider the exact same physical infrastructure for both cases of plain IP and ICN scenarios, with each sink node being connected to both the default and backup P-SSs. For $d_{\max} = 1$, we measure an average number of 6.63 messages for plain IP vs. 2.73 messages in the case of ICN. This translates into 58.75 percent less messages for ICN (the reduction rises to 72.55 percent and 76.56 percent for $d_{\max} = 2$ and $d_{\max} = 3$, respectively) obviously significantly simplifying reconfiguration operations.

At the same time, we anticipate important benefits brought by ICN with respect to the response time to failures. Figure 5 shows the CDF of the hop distance traveled by the considered reconfiguration messages. In all cases, ICN results in a single-hop transmission of the reconfiguration message to each affected RP. In contrast, the corresponding messages in the host-centric case need to reach the PMUs, resulting in longer paths (by 165, 214, and 257 percent for $d_{\max} = 1, 2, \text{ and } 3$, respectively), which obviously result in longer delays for the reconfiguration of the network. This is considered of particular importance in view of the real-time characteristic of the RTSE application, and comes as a direct gain of an ICN design that focuses on denoting the required information rather than reconfiguring the targeted devices (i.e., PMUs).

Summary and Conclusions

In this article, we revisit ICN in the context of smart grid communications and machine-to-machine applications in the domain of active electrical distribution networks. As opposed to the broadly studied content distribution applications in the Internet, such applications present, in several cases, real-time, often complex, communication patterns with stringent delay requirements critical for the operation and protection of the power grid. At the same time, the underlying communication infrastructure is heterogeneous, with multiple technologies

² The methodology followed in this study is also applicable in the case of wireless technologies. However, we focus here on optical fiber technologies as recent studies have shown that wireless technologies such as WiMAX and LTE introduce significant control plane and MAC layer delays [13, 14]. The investigation of such aspects is beyond the scope of this article.

³ These figures do not include links to backup P-SSs.

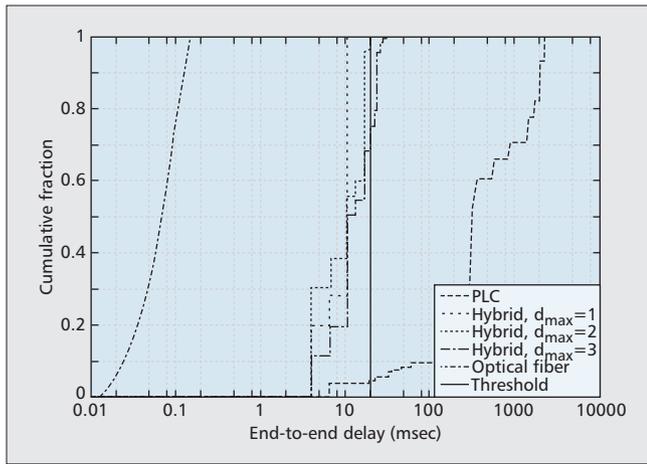


Figure 4. Cumulative distribution function of end-to-end delay for 500kb/s PLC links.

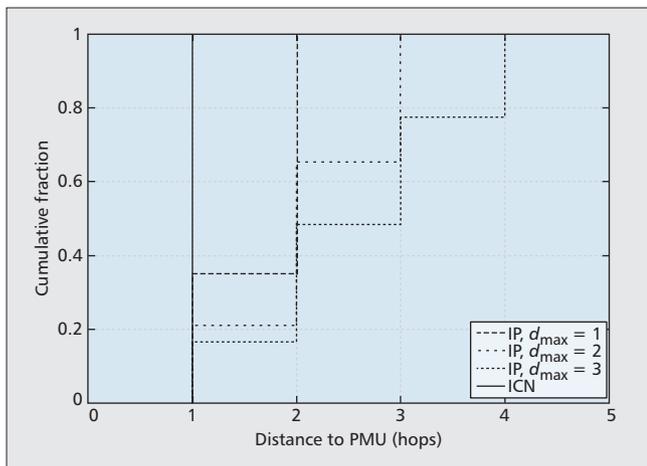


Figure 5. Cumulative distribution function of number of hops required by data flow reconfiguration messages.

being currently deployed or considered for the support of a diverse set of applications. In this context, we have shown how ICN can address the emerging challenges based on a series of design features such as multi-RP selection and in-network processing. Focusing on the topological characteristics of a real power distribution network in the Netherlands, we demonstrate the ability of an ICN approach to address the aforementioned challenges through a series of simulations. A proof-of-concept prototype of the presented functionality has been implemented in order to further experimentally investigate the proposed approach in field trials.

Acknowledgments

The research leading to these results has received funding from the European Community's Seventh Framework Programme FP7-ICT-2011-8 under grant agreement no. 318708 (C-DAX). The authors alone are responsible for the content of this article.

References

- [1] G. Xylomenos *et al.*, "A Survey of Information-Centric Networking Research," *IEEE Commun. Surveys and Tutorials*; available online since 19 July 2013.
- [2] Y. J. Kim *et al.*, "A Data-Centric Information Infrastructure for Smart Grid," *IEEE Commun. Mag.*, vol. 48, no. 11, 2010.
- [3] J. Zhang *et al.*, "iHEMs: An Information-Centric Approach to Secure Home Energy Management," *Proc. IEEE SmartGridComm*, 2012.
- [4] Y. J. Kim *et al.*, "SeDAX: A Scalable, Resilient, and Secure Platform for Smart Grid Communications," *IEEE JSAC*, vol. 30, no. 6, 2012.

- [5] IEEE Std 2030, "IEEE Guide for Smart Grid Interoperability of Energy Technology and Information Technology Operation with the Electric Power System (EPS), End-Use Applications, and Loads," 2011.
- [6] M. Paolone *et al.*, "A Hardware-in-the-Loop Test Platform for the Real-Time State Estimation of Active Distribution Networks Using Phasor Measurement Units," *Proc. Cigré SC6 Colloquium*, Oct. 2013.
- [7] IEEE Std C37.118.1, "Synchrophasor Measurements for Power Systems," 2011.
- [8] T. Koponen *et al.*, "A Data-Oriented (and Beyond) Network Architecture" *Proc. ACM SIGCOMM '07*, Aug. 2007.
- [9] W.K. Chai *et al.*, "Curling: Content-Ubiquitous Resolution and Delivery Infrastructure for Next-Generation Services," *IEEE Commun. Mag.*, vol. 49, no. 3, Mar. 2011, pp. 112–20.
- [10] K.V. Katsaros *et al.*, "On Inter-Domain Name Resolution for Information-Centric Networks," *Proc. IFIP TC6 Networking*, 2012.
- [11] V. Arya *et al.*, "CPS-Net: In-Network Aggregation for Synchrophasor Applications," *Proc. Commun. Systems and Networks*, Jan. 2011.
- [12] S. Galli *et al.*, "Power Line Communications and the Smart Grid," *Proc. IEEE SmartGridComm*, 2010.
- [13] R.H. Khan *et al.*, "Wide Area PMU Communication over a WiMAX Network in the Smart Grid," *Proc. IEEE SmartGridComm*, 2012.
- [14] P. Cheng *et al.*, "Feasibility Study of Applying LTE to Smart Grid," *Proc. IEEE Smart Grid Modeling and Simulation*, Oct. 2011.
- [15] R. L. Francis *et al.*, "Distance Constraints for Tree Network Multifacility Location Problems," *Operations Research*, vol. 26, 1978, pp. 570–96.

Biographies

KONSTANTINOS V. KATSAROS (k.katsaros@ucl.ac.uk) received his B.Sc. in informatics (2003), and his M.Sc. (Honours, 2005) and Ph.D. (2010) degrees in computer science from A.U.E.B. Greece. He has worked in the areas of cloud networking, mobile grid computing, and multicast/broadcast service provision over cellular networks. He has participated in many EU projects including PSIRP and PURSUIT. His research interests focus on information-centric networking and smart grid communications. Currently he is a research associate at the Department of Electronic and Electrical Engineering, University College London, United Kingdom.

WEI KOONG CHAI (w.chai@ucl.ac.uk) was awarded a B.Eng. degree in electrical engineering from the Universiti Teknologi Malaysia in 2000, and both M.Sc. (Distinction) and Ph.D. degrees from the University of Surrey, United Kingdom, in 2002 and 2008, respectively. He is currently a senior research associate at the Department of Electronic and Electrical Engineering, University College London. His research spans across heterogeneous networks including wired/wireless networks and cyber physical systems such as smart grid. His current research interests include information-centric networking, smart grid communication, complex network analysis, QoS, and resource management.

NING WANG (n.wang@surrey.ac.uk) is a senior lecturer at the Centre for Communication Systems Research (CCSR), University of Surrey. He received his B.Eng. (Honours) degree from Changchun University of Science and Technology, P.R. China, in 1996, his M.Eng. degree from Nanyang University, Singapore, in 2000, and his Ph.D. degree from the University of Surrey in 2004. His research interests mainly include information-centric networking, network resource management and optimization, and smart grid communications. Since 2009 he has been the principal investigator of a number of EU and U.K. research grants in the areas of future Internet design, and network management and control.

GEORGE PAVLOU (g.pavlou@ucl.ac.uk) is a professor of communication networks in the Department of Electronic and Electrical Engineering, University College London. He received a Diploma in engineering from the National Technical University of Athens, Greece, and M.Sc. and Ph.D. degrees in computer science from University College London. His research interests focus on networking and network management, including aspects such as traffic engineering, QoS management, autonomic networking, information-centric networking, and software-defined networks. He has been instrumental in a number of research projects that produced significant results with real-world uptake and has contributed to standardization activities in ISO, ITU-T, and IETF.

HERMAN H.I. BONTIUS (herman.bontius@alliander.com) has a B.Sc. in electrotechnical engineering and telecommunication. He has 20 years of engineering and consultancy experience in requirements analysis, innovation, and development of data communication solutions, mainly for power utilities. Currently he is a consultant working on requirements analysis, architecture, and specification of grid solutions for monitoring, protection, and control.

MARIO PAOLONE (mario.paolone@epfl.ch) received his M.Sc. degree (with honours) in electrical engineering and his Ph.D. degree from the University of Bologna, Italy, in 1998 and 2002, respectively. In 2005, he was appointed a researcher in electric power systems at the University of Bologna, where he was with the Power Systems Laboratory until 2011. In 2010, he received associate professor eligibility from the Politecnico di Milano, Italy. Currently, he is an associate professor at the Swiss Federal Institute of Technology, Lausanne, Switzerland, where he has accepted the EOS Holding Chair of the Distributed Electrical Systems Laboratory.