1

An Information-centric Communication Infrastructure for Real-time State Estimation of Active Distribution Networks

W. K. Chai, N. Wang, K. V. Katsaros, G. Kamel, S. Melis, M. Hoefling, B. Vieira, P. Romano,
S. Sarri, T. Tesfay, B. Yang, F. Heimgaertner, M. Pignati, M. Paolone, M. Menth,
G. Pavlou, E. Poll, M. Mampaey, H. Bontius and C. Develder

Abstract—The evolution towards the emerging Active Distribution Networks (ADNs) can be realized via a Real-Time State Estimation (RTSE) application facilitated by the use of Phasor Measurement Units (PMUs). A critical challenge in PMU-based RTSE is the lack of a scalable and flexible communication infrastructure for the timely (i.e., sub-second) delivery of the high volume of synchronized and continuous synchrophasor measurements. We address this challenge by introducing a communication platform called C-DAX based on the information-centric networking (ICN) concept. With a topic-based publish-subscribe engine that decouples data producers and consumers in time and space, C-DAX enables efficient synchrophasor measurement delivery as well as flexible and scalable (re)configuration of PMU data communication for seamless full *observability* of power condition in complex and dynamic scenarios. Based on the derived set of requirements for supporting PMU-based RTSE in ADNs, we design the ICN-based C-DAX communication platform, together with a joint optimized physical network resource provisioning strategy in order to enable the agile PMU data communications in near real-time. In this paper, C-DAX is validated via a prototype implementation deployed over a sample feeder as a proof-of-concept and evaluated through simulation-based experiments using a large set of real medium voltage grid topologies operating in the Netherlands.

Index Terms—information-centric networking, communication infrastructure, RTSE, PMU, ADN

I. INTRODUCTION

WITH the advent of smart grid, the traditional view of distribution networks being mostly passive and static with marginal load variations will no longer hold. The introduction of multiple dynamic active components, e.g., distributed (renewable) energy resources (DERs) and electric vehicles, at distribution level poses new challenges to the system stability especially on power quality, voltage regulation, protection and reliability of the grid. This evolution of distribution networks towards the so-called Active Distribution Networks (ADNs)[1] require enhanced grid real-time monitoring systems. In this respect, the deployment of high-precision Phasor Measurement Units (PMUs) gains a significant role for distribution network operators (DNOs) yet such synchrophasor-based real-time state estimation (RTSE)

has stringent latency and stability requirements. The critical challenge here to the communication infrastructure at the distribution grid level is to be able to support reliable, flexible and scalable low-latency information access for the timely (i.e., tens of miliseconds) delivery of high volume of UTC-synchronized and continuous synchrophasor measurements, collected from geographically dispersed PMU locations. At transmission level, PMU-based monitoring applications have already seen much adoption [1], [2].

We address this communication challenge with an information-centric communication infrastructure which exploits the latest Information-Centric Networking (ICN) concept [3]. Instead of centralized point-to-point platforms such as SCADA, ICN communication is centric to information/data (i.e., *what*) rather than the host locations (i.e., *where*) as information consumers are mainly interested in the information itself rather than the explicit location of its source (e.g., the host IP address). Consequently, the ICN design principle puts information/data as the main focus where information is published, resolved, delivered and stored natively based on names instead of explicit host locations. This enables the spatiotemporal decoupling of communicating entities where data producers and consumers are agnostic to where and when the data will be published/consumed by their counterparts. ICN also offers intrinsic support for cyber-security as its operations are characterized by the location-independent feature that does not expose the network and physical locations of the communicating machines. Being an integral part of the underlying network infrastructure, securing the power grid communication becomes part and parcel of an ICN system and no *a posteriori* patches are required.

ICN benefits RTSE operations on several fronts. With its inherent publish-subscribe communication primitives, ICN provides greater flexibility in supporting PMU data delivery, as no explicit dedicated point-to-point communication sessions are needed between identified end devices. For instance, an asset-change (e.g., due to a line fault) may trigger a grid topology change. This will not affect ongoing data delivery operations in ICN, since data consumers (e.g., Phasor Data Concentrators, PDCs) only need to subscribe to the data of interest and the rest of the communication will be automatically set up. Such automation minimizes error-prone manual reconfigurations, while shortening the convergence time for the new data delivery patterns [4]. With information/data-

---

[1]ADNs are defined as distribution networks that have systems in place to control a combination of DERs like generators, loads and storage. Distribution system operators have the possibility of managing electricity flows via a flexible network topology. DERs take some responsibility for system support, depending on a suitable regulatory environment and connection agreement.

awareness in ICN, data caching [5] with local processing is also possible, which improves the data delivery efficiency. An example is the vision of on-demand islanding operations by DNOs. In this case, data can be cached and aggregated for later retrieval while the island is formed and all flows redirected to the new subscribers, e.g., a PDC at a different primary substation (P-SS), once the islanding manoeuvre is completed. This feature also contributes to the overall quality-of-service (QoS) support, which is especially critical to PMU data that has extremely stringent latency requirements. ICN network nodes may adapt measurement data rates targeting at different grid operations, rather than having such functions only at the end devices.

Recent works [4], [6]–[8] have advocated for the application of ICN in smart grid. Our work represents a further step to specifically realize PMU-based RTSE with the participation of a European DNO, providing real grid data and a live electric network for field trials of our prototype. Adhering to the derived RTSE requirements and design rationale in Section II, we present in Section III the specifications of our proposed C-DAX platform based on ICN principles. In Section IV, we validate C-DAX first via a proof-of-concept prototype showing the full functionalities of C-DAX along with the operation of a linear state estimator. This is complemented with a simulation campaign using real world grid data and topologies for large-scale evaluation. Our comprehensive evaluation shows how C-DAX effectively and robustly delivers the PMU measurements to the RTSE application and this further supports the adoption of ICN concepts in such a critical machine-to-machine communication scenario. We then delve deeper into the communication network revealing the networking performance against the derived RTSE requirements, highlighting the bandwidth savings and simplicity of C-DAX. We conclude and summarize our findings in Section V.

II. SYNCHROPHASOR-BASED REAL-TIME STATE ESTIMATION OF ACTIVE DISTRIBUTION NETWORKS

State estimation (SE) is a consolidated functionality extensively used in transmission networks (e.g., [9], [10]). For ADNs, a dedicated Energy Management System is required to support specific operations objectives [11], [12], [14]–[18]: (i) optimal voltage control / line congestion control, (ii) fault detection and location, (iii) post-fault management, (iv) local load balance, and (v) loss minimization. These operations need to know the system state in fine-granularity (e.g., [11]). Yet, existing SE processes have typical refresh rates in the order of a few minutes, whereas the time frames for the aforementioned functionalities range from a few hundred milliseconds for fault management, to potentially tens of seconds for voltage control and line congestions [19]. Therefore, the availability of 3-phase RTSE enables the possibility to properly track the higher ADN dynamics by assessing their state within the order of tens of milliseconds.

Such RTSE assessment can be facilitated by the use of PMUs (e.g., [20], [21]). In order to correctly deploy such a paradigm into ADNs, PMU technology needs to satisfy additional requirements compared to those defined in the

TABLE I: PMUs number and computation time performance using the DKF algorithm for three different distribution test-feeders.

Network	PMUs number	Mean and standard deviation (in ms)
IEEE 7-bus	4	0.430 ± 0.026
IEEE 13-bus	7	0.850 ± 0.053
IEEE 34-bus	12	4.400 ± 0.520

IEEE Std. C37.118.1-2011 [23] and its latest amendment IEEE Std. C37.118.1a-2014 [24], that were originally intended for transmission network applications. In particular, dedicated PMUs for ADNs need to satisfy more stringent requirements in terms of phase accuracy and harmonic rejection, as well as to keep their performance stable during power system dynamics. Additionally, they also require reduced time-latencies and increased reporting rates, in the order of at least 50 to 60 fps.

In this respect, in what follows, we will make reference to a PMU prototype specifically developed by the authors to enable its use in ADNs. In particular, as demonstrated in [25], the proposed PMU prototype satisfies every requirement defined in [23] for class-P PMUs and the majority of class-M requirements, with the exception of the Out-Of-Band Interference tests. The prototype, together with the formulation of the so-called enhanced Interpolated DFT (e-IpDFT) synchrophasor estimation algorithm, presented in [25], have been further extended in [26] where an approach to reduce the measurement reporting latencies and consequently increase the reporting rates of the e-IpDFT algorithm is proposed. The PMU prototype is based on the National Instruments CompactRIO 9068, which comprises a reconfigurable Artix-7 FPGA and a dual-core ARM Cortex-A9 processor, and is loaded with a customized Linux-RT Operating System. As illustrated in [25], since the e-IpDFT algorithm has been deployed into the FPGA target only, the dual-core processor is let free to potentially host any third-party feature, such as the coupling with the C-DAX architecture, as shown later.

A. Latency and Delay Variation (Jitter) Requirements

The synchrophasor measurement data flow, based on the above-mentioned PMU technology, is characterized by the set of delays represented in Fig. 1.

- *PMU signal acquisition*: the delay associated with this operation typically corresponds to half of the window length used by the PMU to measure a GPS-timestamped synchrophasor, as the timestamp is usually placed in the middle of the time window. The window length adopted by the authors in both [25] and [26] contains three periods of a signal at the rated frequency of the power system (i.e., 60 ms at 50 Hz).
- *PMU synchrophasor estimation and data encapsulation*: the latencies introduced by the synchrophasor estimation processes typically depend on the number of input channels (voltages/currents) the PMU measures. In [26] a method that drastically reduces the latencies of any DFT-based algorithms has been presented and experimentally validated. In this respect, the PMU adopting the e-IpMSDFT algorithm described in [26] has demonstrated to be capable of estimating the synchrophasors of a

maximum number of 37 channels in less than 1 ms. The time for encapsulating data using the C37.118.2011 protocol is negligible.

- *Communication network delay*: this varies as a function of (i) the communication technologies used (e.g., optical fibers, WiFi, 3G, etc.), (ii) the communication protocol, (iii) the physical distances between the PMUs and the PDC, (iv) the processing delays incurred by functions related to medium adaptation, (de)coding, switching and routing. We derive the latency bounds for this delay which the communication infrastructure must achieve.
- *PDC data frame time alignment*: this is needed to compensate for network delay and jitter to have a consistent (i.e., synchronised timestamp) dataset feeding the state estimator. The time alignment adds an additional delay (in the range of 0–20 ms) to buffer the incoming PMU streamed data.
- *Bad data detection*: embedded in the SE process, it identifies and eventually replaces corrupted and/or missed measurements. Its latency is generally negligible.
- *State estimation*: the latency of this depends on a number of different parameters, i.e. the adopted SE algorithm, the size of the considered network and the total number of measurements. Table I shows the number of PMUs and the computation time, in terms of mean and standard deviation values (in milliseconds) for three distribution networks typically adopted in the technical literature. The IEEE 13-bus and the IEEE 34-bus distribution test feeders are two benchmark networks whose structure is given in [27], whereas the IEEE 7-bus network is the first section of the IEEE 34-bus distribution test feeder. The results shown in Table I are obtained by running a SE for 1500 time-steps. The adopted SE algorithm is the Discrete Kalman Filter (DKF), whose formulation is given in [28].

The number and location of the PMUs are, in general, selected by taking into account the following parameters:

- – Grid topology: radial or meshed,
- – availability of other conventional measurements, and
- – existence of the so-called "zero-injection buses"[2].

Once the above-mentioned factors are given, the PMU placement is determined by performing the observability analysis (e.g., [29], [30]).

As shown in Table I, the computation time increases when the size of the network becomes larger. However, even in the case of a quite large distribution level network, such as the IEEE 34-bus feeder, the timing performance is within the latency requirements of the RTSE process.

The time budget of the latency of the whole RTSE process is strongly dependent on the targeted application. In this respect, the most time-critical application we consider here is the use of RTSE for ADN relaying applications (e.g., for grid protection). In this case, we account for a maximum total latency of 100 ms. From Fig. 1, the time budget left for the
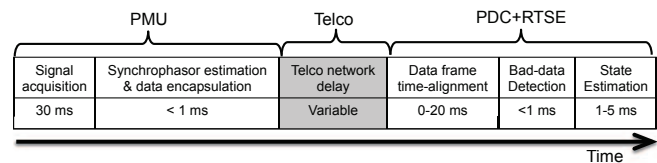

Fig. 1: Latency components of the entire RTSE process.

communication infrastructure is approximately 35–55 ms.

Furthermore, phasors measured at different grid locations at the same instance must be delivered within the range of 0–20 ms apart from each other (cf. Fig. 1) so as to be time-aligned. We consider the mid-point of 10 ms as the requirement for PMU data delay jitter.

B. Bandwidth Requirement

In terms of bandwidth requirements, and based on PMU data semantics [31], a realistic PMU message payload size is 86 bytes[3]. Further considering the overhead of UDP/IP headers, and a 32-byte SHA-256 message authentication code, the overall data rate for each PMU flow delivered to the link layer is 58.4 Kbps[4], for a fixed 50 Hz reporting rate. The bandwidth requirements on the communication links depend on the number and the location of the PMUs deployed. The selection of these locations constitutes a research area of its own (e.g., [32]). Without loss of generality, we consider a scenario involving PMU deployment at approximately every two secondary substations (S-SSes) along a feeder[5].

C. Resilience Requirement

Besides timely data access, seamless data delivery is also an important requirement for RTSE to keep track of the system state at all time. As such, the communication infrastructure must be able to ensure data continuity even when network congestions (e.g., due to improper dimensioning of the physical communication bandwidth resources) or anomalies (e.g., failure of communication equipment) occur. In this respect, detection of communication network congestion/failure and the corresponding recovery must be instantaneous so as to minimize the effect on the state estimation.

D. Flexibility Requirement

One unique feature of ADN is the possibility for DNOs — with possible participation of DERs — to manage electricity flows using a flexible grid topology where the power supply of a feeder may be moved from one P-SS to another. In this case, PMUs may stream synchrophasors concurrently to multiple PDCs. Moreover, maintenance operations requiring an asset change or islanding operations may also involve the switching of PMU data delivery from one PDC to another.

---

[2]A zero-injection bus is a bus that is not connected to a load or a DER and, as a consequence, has a nodal current (or its injected / absorbed power) of zero.

[3]Considering the case of PHNMR=6 and ANNMR=3, with 32-bit floating-point accuracy [31].

[4]In Section IV, we also consider an additional 16-byte C-DAX header overhead leading to an overall data rate of 64.8 Kbps per PMU.

[5]We use this only as a rule of thumb; C-DAX is not dependent on the exact number or placement of PMUs in the grid.
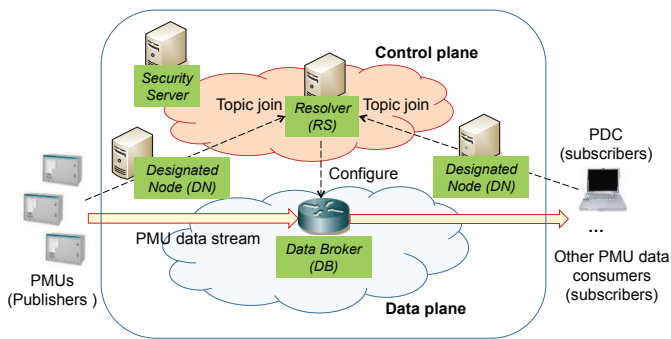
Fig. 2: C-DAX Architecture.

Per-flow configuration for each PMU-PDC communication pair is inefficient especially for large grids that involve hundreds of PMUs. Per-flow management requires cumbersome identification and communication of each and every affected PMU. Therefore, the communication infrastructure is also required to be flexible and simple for (re)configuring communicating parties.

III. C-DAX COMMUNICATION PLATFORM FOR RTSE

A. Design Rationale and Architecture

In this section, we present C-DAX[6], an ICN-based middleware platform with a topic-based publish-subscribe engine that decouples data producers and consumers in time and space. Utility equipment producing/consuming topic data such as PMUs as publishers, and PDCs as subscribers, are clients to the C-DAX platform. In C-DAX, a topic represents a group-based communication session for data distribution, e.g., a topic used by PMUs to publish their measurement data towards interested subscribers in RTSE operations. The platform enables scalable and flexible (re)configuration of PMU data communication in order to maintain seamless full observability of power conditions in complex and dynamic situations. Besides providing inherent security protection by obscuring target hosts, the decoupling of the communicating end hosts massively simplifies the configuration complexity of multipoint-to-multipoint communications in a decentralized manner. It eliminates the need to establish and maintain multiple and distinct point-to-point communication connections, thus achieving better system scalability in PMU-based RTSE communications. Furthermore, C-DAX enables the forming of topic groups, which facilitates PMU data delivery in which data can be replicated to all interested subscribers, saving the need for repeated transmission of the same data by emulating the function of multicast [33]. This is especially beneficial for the flexible topology feature of ADNs. Fig. 2 shows the C-DAX architecture for supporting a PMU-based RTSE application.

Fundamentally, C-DAX consists of both control and data planes, with the former being responsible for handling topic-based communication sessions, typically realised through a signalling protocol (cf. Section III-B), while the latter being dedicated to the forwarding of the PMU streaming data according to end-to-end QoS requirements (cf. Section II). In the data plane, PMUs, as legitimate data publishers, stream real-time measurement data under a given topic to an entity called Data Broker (DB) which acts as a rendezvous point connecting the publishers and the subscribers of this topic. All the data received at the DB are then forwarded to the legitimate subscribers of the topic such as a PDC. In such a publish-subscribe based approach, instead of explicitly handling pairs of communication endpoints, PMU data streams in C-DAX are manipulated based on a logical topic rooted at a common DB. Thus, the location of the DB in the data path may have a substantial influence on the end-to-end delay performance. In Section IV, we specifically study this issue through simulations based on real grid topologies. For scalability and resilience, especially in larger grids, multiple co-existing DBs can be deployed simultaneously for PMU streaming under a common topic. In this case, PMUs close to each other can publish to a nearby DB, and if this DB fails, an alternative nearby DB can take over its role for seamless failure recovery. Finally, the use of DBs enables the semantic grouping of PMUs in terms of management and configuration (e.g., PMUs in a certain feeder or part of a grid may need to be configured collectively with the same operating parameters). In this case, DBs may have extra functions in the data plane such as caching of PMU data for further usage and also rate adaptation of data streams towards subscribers with heterogeneous data rate requirements.

The C-DAX control plane facilitates clients to join and leave a topic. Before a PMU (which acts as a data publisher) starts streaming data pertaining to a given topic, it requires the location of the responsible DB for that topic. This is handled by the topic Resolver (RS) entity in the control plane which is responsible for mapping a topic name to the suitable DB. Specifically, a publisher joins a topic by sending a topic-join request towards the RS to obtain the location information of the corresponding DB. On the subscriber side, a topic-join request indicates to the RS the entity to which the published data under the given topic should be delivered. Such information is sent by the RS to the DB in the data plane by the RS (which knows the topic-DB mapping in the control plane). Since there can be potentially a large number of clients in the grid, we also introduce so-called Designated Node (DN) entities for the aggregation of topic-join requests sent towards the RS. In this case, all the topic-join request events originated from clients sharing a common DN are handled by that DN, such that the control signaling overhead can potentially be reduced. Taking a proxy role, each DN, for the purpose of security, is also responsible for performing local authentications of publishers/subscribers in the grid. Such operations require a dedicated security server (SecServ)[7] in the control plane to carry out the access control and client authentication operations upon the joining of new clients reported by a DN.

As can be seen, the management (including initial configuration and run-time reconfiguration) of the PMU-based RTSE communication is completely based on the logical topic rather

---

[6]http://www.cdax.eu

[7]For scalability and resilience, multiple security servers may be deployed (possibly organized in a hierarchical manner).

than the manipulation of explicit point-to-point communication points. Such a scheme offers increased scalability and flexibility in handling power grid topology dynamicity during normal operations.

B. Basic C-DAX Operations

Clients need to join topics before publishing or consuming data pertaining to those topics. The C-DAX publish-subscribe communication functions rely on three core operations; namely (i) joining, (ii) leaving, and (iii) configuring a topic. Each operation utilizes two main primitives in a request and respond manner. A C-DAX client wanting to join a topic first sends a `TopicJoinReq` message to its DN stating its purpose (i.e., either data publication or data subscription), the interested topic, and its own Client ID. We further differentiate two types of DN: Publisher DN (PubDN) and Subscriber DN (SubDN), defined according to whether a publisher client or subscriber client is attached to it. For security, a client signature is attached to the message. The `TopicJoinReq` message is forwarded to the security server for access control. In response, the security server will reply to the DN with a `TopicJoinResp` message that corresponds to the initial join request. This message will indicate whether the request from the client to join the specific topic is allowed or denied. The protocol for *leaving* a topic is similar but with the request and reply primitives instead being `TopicLeaveReq` and `TopicLeaveResp`, respectively.

The signaling for configuring a topic in C-DAX depends on the role of the client. Since data flows outwards from a publisher towards a subscriber, the correct forwarding states must be installed at the right locations to properly direct the data flows. Once a DN receives a response message from the security server, it will trigger the topic configuration process, beginning with a `TopicConfigReq` message. This is routed to the RS to resolve the topic to the DB according to predefined rules (e.g., using a distributed hash table approach [8]). Upon reception of the `TopicConfigReq` message, the RS will generate a corresponding `TopicConfig` message. For a publisher topic join, the RS will forward this `TopicConfig` message to the original PubDN while for a subscriber topic join, the message will be forwarded to the DB responsible for the interested topic instead. In both cases, next-hop information will be installed. Table II summarizes the primitives for the core C-DAX operations while Fig. 3 shows the sequence diagrams for publisher topic join, subscriber topic join, publisher topic leave and subscriber topic leave operations, respectively. The entities involved in the data plane all maintain their own forwarding table which essentially maps topics to the correct next hop.

C. C-DAX Features

The simplicity of C-DAX in accommodating changes in the communication flow pattern constitutes the C-DAX flexibility feature. Dynamicity of ADNs is not restricted to the volatility of the system state but also the possibilities of changes in the energy flow directions (e.g., due to the flexible topology of ADN, asset change for maintenance, islanding operations

TABLE II: C-DAX operations.

| Operation	Primitive	Usage Parameters
Topic Join	`TopicJoinReq`	Mode, Topic, ClientID, ClientSig Mode = Publish, Subscribe
	`TopicJoinResp`	Mode, Topic, Response, ServerSig Mode = Publish, Subscribe Response = Yes, No
Topic Leave	`TopicLeaveReq`	Mode, Topic, ClientID, ClientSig Mode = Publish, Subscribe
	`TopicLeaveResp`	Mode, Topic, ClientID, ClientSig Mode = Publish, Subscribe
Configuration	`TopicConfigReq`	Mode, Topic, DN-ID, AuthToken Mode = Publish, Subscribe, Unsub-scribe
	`TopicConfig`	Mode, Topic, TargetID, ResolverSig Mode = Fwd, Stop_Fwd TargetID =DB-ID, DN-ID, ...
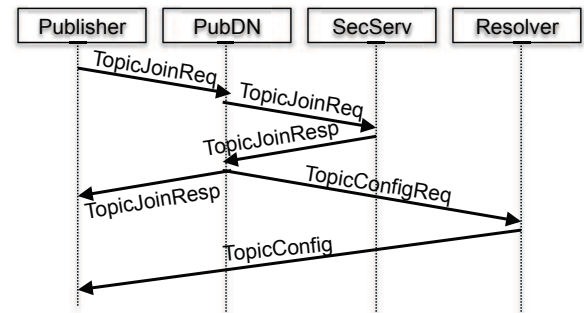
(a) Publisher join.
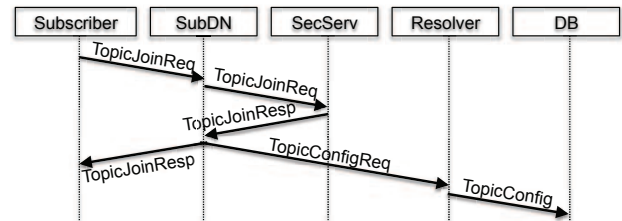
(b) Subscriber join.
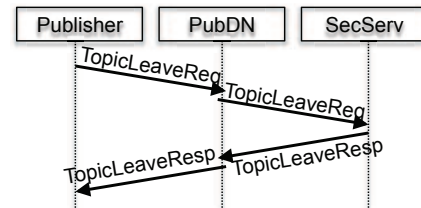
(c) Publisher leave.
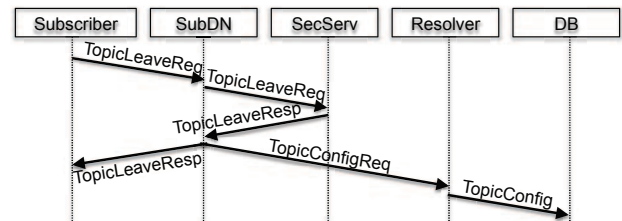
(d) Subscriber leave.

Fig. 3: Topic join/leave signalling sequences for publishers (e.g., PMUs) and subscribers (e.g., PDCs) who are decoupled spatiotemporally.

etc.). A direct consequence of this change is that it also affects the destination of synchrophasor measurements.

Consider a distribution network with M P-SSes and N S-SSes. As described in Section II-B, we consider the scenario involving the deployment of a PMU at every other S-SS along a feeder. Assuming each P-SS is equipped with a PDC, then we have M potential subscribers and $\frac{N}{2}$ publishers. For point-to-point (PtP) communication, to receive PMU measurements, each PDC must explicitly communicate with $\frac{N}{2}$ PMUs, resulting in $\frac{N}{2} \times M$ independent communication pairs. However, with C-DAX, each PDC (and similarly PMU) needs to communicate with only the responsible DB ($= \frac{N}{2} + M$ communication pairs). In other words, for every additional subscriber, the extra communication overhead required for a point-to-point solution increases by a factor of M while with C-DAX, the increase is only linear with M. As demonstrated in Section IV-B, this feature reduces the requirements on the physical communication infrastructure, enabling the corresponding reduction of the associated communication infrastructure deployment capital expenditure (CAPEX).

Besides enhancing system scalability, the elimination of pair-wise publisher-subscriber communication reduces the re-configuration overhead. In fact, C-DAX clients are never required to contact their counterparts, e.g., between a PMU as the publisher client, and a PDC as a subscriber client. All new clients joining a specific topic always follow the same procedure (see Fig. 3) regardless of other clients in the topic group. In addition, none of the clients already joined to the topic group need to know of changes in the group membership. For instance, a PMU does not need to react to any newly joined/left PDC, but can continue to send its data to the same DB.

In terms of resilience, C-DAX detects communication failures using `Hello` messages and timers, such that a C-DAX node periodically probes other C-DAX nodes with which it is communicating to check if they are still alive. In case a C-DAX node fails to receive `Hello` messages from a neighbor for a period of time, then the associated timer expires. At this point, it assumes that this neighbor has become unavailable and hence, necessary recovery operations need to be activated accordingly. The C-DAX operator can configure the frequency of such probing messages according to application requirements. In the context of RTSE, this frequency is also dependent on whether there is a bad/missing data filtering algorithm (e.g., [34]) present for which the impact of data loss can be mitigated. It is beyond the scope of our work to explore the performance of different mitigation algorithms. As a conservative baseline, assuming the absence of such an algorithm, the timers can be set to expire within 50 ms to ensure the failure is detected and recovered sufficiently early (see Section IV-A2).

The critical component type in the C-DAX architecture is the DB as they are the intermediaries connecting the publishers and subscribers, and without it, data cannot reach the subscribers. To minimize the effect of a DB failure and thus the recovery time, a PMU may pre-establish a connection with an alternative DB indicated by the RS. With this, once the original DB is detected to have failed, the PMUs attached can immediately switch the transmission of the data to the second DB. To combat the effect of a possible cascading
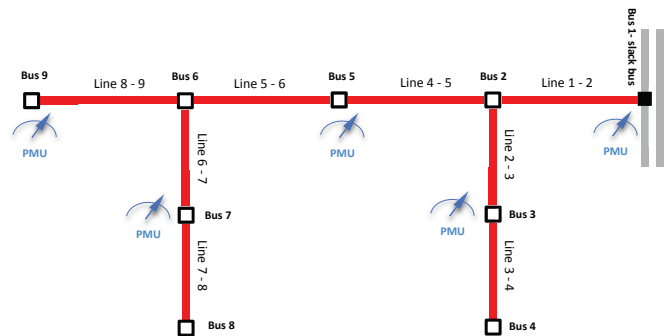

Fig. 4: Feeder used for proof-of-concept prototype experimentation.

failure, the RS at the control plane must elect a new backup DB whenever the failure of the current primary DB (PrimDB) is detected. When this is done, the current secondary DB (SecDB) becomes the primary one[8].

IV. VALIDATION AND PERFORMANCE EVALUATION

In this section, we proceed to evaluate C-DAX by means of a prototype as well as a series of packet-level simulations. First, in Section IV-A we validate our C-DAX prototype by deploying it to a sample feeder operating in the Netherlands. Then, in order to assess the performance of the proposed architecture in large-scale scenarios involving hundreds of substations and PMUs, we resort to extensive packet-level simulations over a large set of real medium-voltage (MV) grids (cf. Section IV-B).

A. C-DAX Prototype for Supporting RTSE

To validate C-DAX, we have developed a prototype encompassing the complete set of C-DAX functionalities. Each C-DAX node is implemented modularly such that it can perform the functions of different entities depending on the C-DAX operator's need (i.e., a physical C-DAX node can assume the role of multiple entities, such as a DN and DB, simultaneously). Such a feature offers high configuration flexibility of C-DAX nodes according to specific grid topologies. Phasor measurements are obtained using the PMU prototype based on the National Instruments CompactRIO 9068 that was introduced in Section II and first presented in [25]. On the other hand, the RTSE has been implemented using the DKF algorithm presented in Section II and described in detail in [28].

We deploy our prototype on a virtualized network testbed providing support for the configuration of custom network topologies. For our purpose, we emulate the operation of C-DAX on top of the feeder shown in Fig. 4, consisting of 1 P-SS and 8 S-SSes. For full feeder *observability*, we deploy a total of five PMUs, one at each of the buses 1, 3, 5, 7 and 9. The corresponding C-DAX overlay deployment is shown in Fig. 5 with four physical machines assuming the full set of C-DAX entities. PMU measurements are taken and streamed to the PDC at the P-SS over C-DAX. The communication infrastructure follows the power grid topology, with high speed 1 Gbps links resembling a future optical fiber communication

---

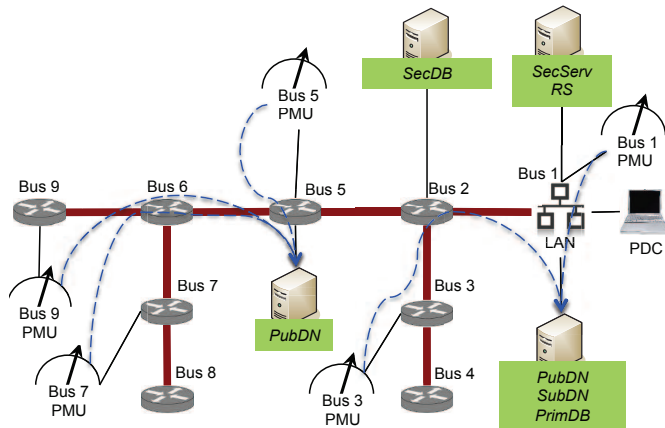[8]A similar principle can apply to DNs, as and when required.

7

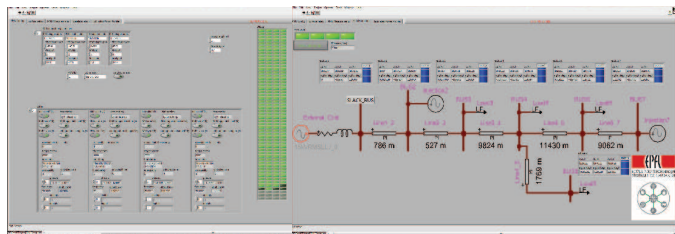Fig. 5: C-DAX deployment on the feeder.

Fig. 6: Labview-based RTSE monitoring a typical feeder.

infrastructure[9]. After time alignment, the PDC forwards the measurements to our Labview-based RTSE algorithm for the computation of the states. Fig. 6 shows the RTSE operating on a typical feeder.

1) Delay measurements: Figures 7(a) and 7(b) present the cumulative distribution function (CDF) of the end-to-end delay perceived at the PDC. The observed delay values are well below the 35–55 ms time budget discussed in Section II-A, with an average value of 1.43 ms. This firstly owes to the high capacity of the communication links which is several orders of magnitude larger than the considered data rates (see Section II-B). However, it is important to observe that the introduction of intermediate overlay forwarding devices such as the DNs and the DB, in the context of C-DAX, does not impose a prohibitive processing delay, despite the transition of packets from the kernel space to the user space of the operating system of C-DAX nodes. It is noted as well, that PMU data can also be stored/cached in databases by intermediate nodes (i.e., at the DB or even at the DN) for future reference (typically for other non-real-time applications outside the consideration of this paper), however only once they have first been directly streamed in real time towards the state estimator in PDC without involving database. This means that any input/output (I/O) operation delays[10] at the database level do not impact the perceived end-to-end latency at the state estimator side. Fig. 7(b) reveals slightly higher delays for the flows of PMUs located further away from the

Fig. 8: Packet loss as a function of `Hello` message frequency. (Color in online version.)

TABLE III: Topology properties of a sample MV grid in the Netherlands.

Primary Substations (P-SS)	5
Secondary Substations (S-SS)	579
Number of edges (cables)	679
Average cable length	698 m
Average node degree	2.08

PDC, a result attributed to the transmission, propagation and processing delays accumulated from the longer delivery paths. The difference observed is in the order of 1 ms, implying that the buffering requirements at the PDC (for time-alignment) are very low. Fig. 7(c) and Table IV take a closer look at the perceived delay variation across the PMU flows, which has an average value of 0.1 ms, with negligible differences between flows. As such, it is well within the delay variation requirement derived in Section II-A.

2) Resilience: To investigate the efficiency of the C-DAX resilience mechanism, we consider the scenario where the PrimDB for the PMU topic fails at some random point in time. Fig. 8 shows the number of packets lost (with 95% confidence interval) during a failure event as a function of the frequency of the `Hello` messages. At the lowest frequency (i.e., 800 ms interval between consecutive `Hello` messages), less than 50 messages are lost per PMU, reaching an aggregate of approximately 160 lost measurements. Considering the 50 Hz rate of PMUs, this loss corresponds to less than 1 s of measurements. Increasing the `Hello` message rate reduces the number of lost messages, up to an average of 37 messages across all PMUs. However, this is at the cost of 20 `Hello` messages per second for each PubDN-DB pair, which translates to a data rate of 5.44 Kbps delivered to the link layer[11].

B. Simulation-based Evaluation of C-DAX with Real MV Grids

While the performance evaluation of the C-DAX prototype provides useful insights with respect to the expected performance of the architecture in the context of RTSE, it is also important to investigate performance aspects of a larger scale deployment, considering deployment issues and the impact of the overall communication infrastructure. To this end, we focus our investigations on a set of MV grid topologies which are currently in live operation in the Netherlands. Table III

---

[9]Though optical fiber technology supports higher data rates, we are limited by the capabilities of our testbed.

[10]Considering that todays storage equipment allows for I/O rates in the order of several tens to several hundreds of megabytes per second, it becomes apparent that database I/O operations can easily support several hundreds of PMU flows — which as described in Section II-B, are in the order of several tens of Kbps — without introducing any performance bottleneck.
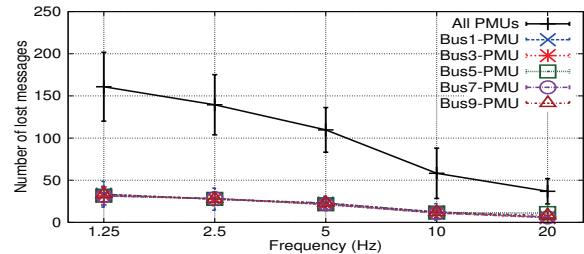
[11]Considered overhead: message type (8 bits), source and destination node type (4 bits each), source node ID (32 bits), UDP/IP header (28 Bytes).

(a) End-to-end delay across all PMU flows.

(b) End-to-end delay per PMU flow.
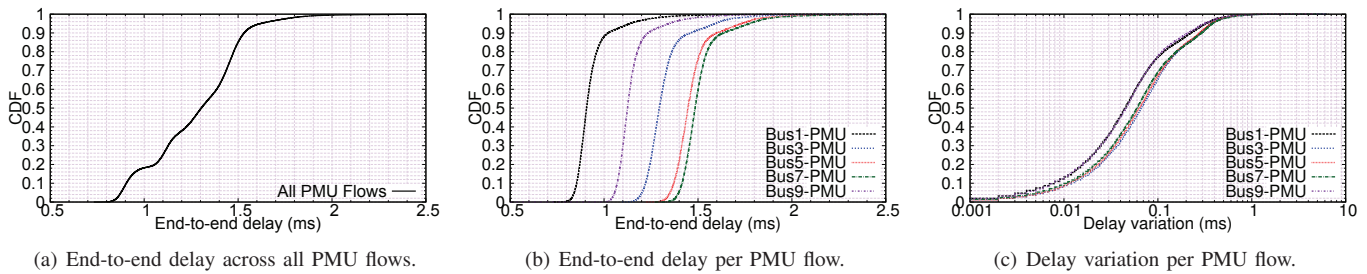
(c) Delay variation per PMU flow.

Fig. 7: End-to-end delay and delay variation of PMU flows in the prototype. (Color in online version.)

TABLE IV: Delay variations recorded (in ms)

	Bus 1	Bus 5	Bus 8	Bus 12	Bus 15	All
Mean	0.087	0.113	0.109	0.110	0.085	0.297
Max	6.554	6.336	4.678	5.484	5.850	5.000
Min	0.000	0.000	0.000	0.000	0.000	0.000
Median	0.045	0.067	0.063	0.060	0.046	0.263

shows the basic properties of the topologies. Our study is based on a detailed packet-level simulation environment we have developed based on OMNET++ [35].

We highlight here that to evaluate C-DAX, we need to first describe its dependency on the actual physical communication infrastructure available in the grid as the underlying technologies is a strong determinant on the performance of the communication architecture. In Section IV-A, we considered a high bandwidth communication infrastructure, however, this raises concerns regarding the costs associated with a full-scale deployment. At the same time, we have shown in [4] that the readily available power-line communications (PLC) alone cannot support RTSE in MV grids, regardless of the overlay communication architecture deployed. DNOs are currently investigating the best way to modernise their communication infrastructure for supporting future smart grid applications which have more stringent requirements. The aim is to minimize the investment cost on deploying new communication infrastructure. As such, we consider the concurrent use of PLC and optical fiber technologies yet keeping in mind the need to minimize infrastructure cost in the sense that optical fiber is used only when necessary. A limited number of *sink* nodes are selected as traffic aggregation points. These nodes are directly connected to one or more P-SSes with optical fiber link(s). PMUs use the low-cost PLC capabilities to send data to the sink nodes, reducing the corresponding optical fiber deployment costs while the sink nodes exploit the high bandwidth of optical fiber for last-leg speedy delivery of the measurement data to PDCs.

We formulate the sink node selection process as a distance-constraint version of the *p*-center facility location problem [36], i.e., defining the minimum (hop) distance of each PMU to its closest sink node as the communication cost. The solution is then the minimum number of sink nodes such that the maximum communication cost (d_{max}) is minimized. We solve this via a sequential location procedure [36] and obtain 214, 132, 95, 75 and 64 sink nodes/DBs for $d_{max} = 1, 2, 3, 4, 5$ respectively. On top of the resulting topologies, we consider the deployment of one PMU every other S-SS (which results in a total of 363 PMUs), one PubDN per S-SS and one

DB per sink node. In the following, we consider 500 Kbps PLC links [37] and 10 Gbps optical fiber links for the sink-to-PDC connections.

1) Bandwidth savings and impact on communication infrastructure: As discussed in Section III-C, the design of C-DAX enables multicast, lowering the bandwidth requirements of the underlying communication infrastructure. To demonstrate this bandwidth saving, we consider an example scenario where each PMU flow is delivered to two or more recipients concurrently. Fig. 9 shows[12] the CDF of the required per PLC link bandwidth for the cases of both C-DAX and plain point-to-point IP communications, for $d_{max} = 4$. C-DAX consistently requires the least per PLC bandwidth with less than 1% of the PLC link found to be overloaded (i.e., requiring more than 500Kbps to transfer the PMU data). On the contrary, the best and worst cases for the point-to-point solution for $d_{max} = 4$ involves approximately 10% and 20% of congested PLC links, respectively. The impact of this overloading in the case of point-to-point IP communications can then be seen on the application domain. Figures 10 and 11 show the CDF of the end-to-end delay recorded both for C-DAX and point-to-point solutions, respectively, with $d_{max} = 1, 2, 3, 4, 5$, where the upper and lower delay thresholds of the delay budget (cf. Section II-A) are indicated with vertical lines. From Fig. 10, we see that when $d_{max} \leq 4$, all measurement packets are successfully delivered within the delay requirements. However, for $d_{max} > 4$, the latency exceeds even the lower threshold of 35 ms, and therefore cannot support RTSE. For the case of the plain point-to-point solution, the acceptable d_{max} value is below two.

The clear implication of these results for C-DAX is that only 75 sink nodes are required, while for the point-to-point solution 132 sink nodes are needed. Considering that each sink node deployment corresponds to the deployment of an optical fiber connection to the P-SS, this represents approximately a 43% cost reduction for the deployment of the communication infrastructure in the case of C-DAX. This means that even though the typically low data rates of PMU flows could easily be accommodated by high capacity links (e.g., optical fibers), still, the adoption of C-DAX and the careful selection of the sink locations would result in significant CAPEX savings for DSOs, without compromising performance.

---

[12]We get similar results for different d_{max} values, but this is not shown due to paper length limitations.
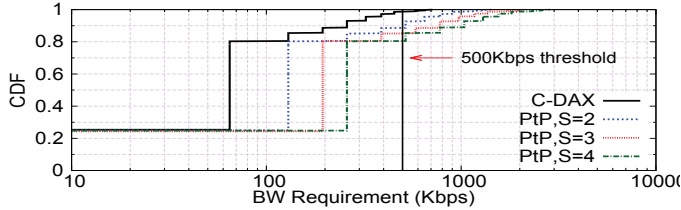
9

Fig. 9: Bandwidth required per PLC link for $d_{max} = 4$ and different number of concurrent subscribers, S. (Color in online version.)
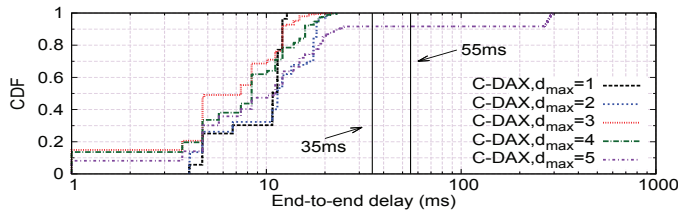

Fig. 10: End-to-end delay for C-DAX. (Color in online version.)

2) Simplification of re-configuration operations: We next look into the (re)configuration message overhead. In the case of the point-to-point approach, a re-configuration event (e.g., a new PDC is deployed), each and every involved PMU must be explicitly configured (cf. Section III-C) so as to update the list of destination hosts, and each of these configuration messages must traverse the entire PDC-PMU path. In the context of the previously described hybrid topologies, this path will also include a sink node. On the contrary, in the case of C-DAX, only one single message is required to add/remove a subscriber to PMU data. This message need only reach the DB corresponding to the selected PMU flows. It is important then to note that the number of re-configuration messages is determined by the number of DBs instead of the number of PMUs involved in the re-configuration. To gain an understanding of the simplifications achieved in this respect, a total number of 75 DBs would be available in the considered hybrid topologies, serving a total of 363 PMUs; implying a reduction in the number of re-configuration message in the order of 80%. At the same time, in the case of C-DAX, these messages need only traverse one hop, i.e., from a P-SS to a sink node. However, this is not the case with point-to-point communications where the corresponding messages would need to reach the PMUs themselves, traversing longer
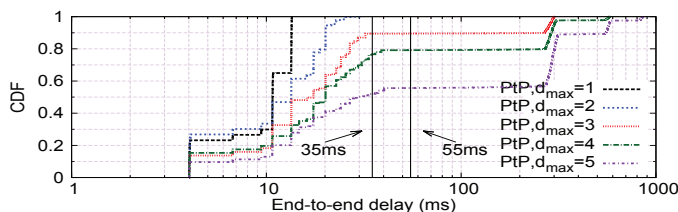

Fig. 11: End-to-end delay for the point-to-point (PtP) approach. (Color in online version.)
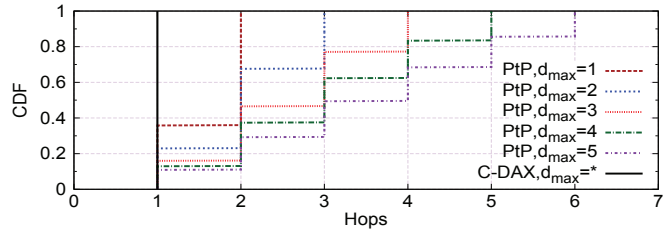

Fig. 12: PMU reconfiguration: distance travelled by control messages. (Color in online version.)

network paths. Fig. 12 shows the CDF of the PMU to sink node distance for different d_{max} values. We see that for $d_{max} = 2$, for which point-to-point communications can satisfy the delay requirements, 76.91% of re-configuration messages would traverse more than one hop. In this respect, re-configuration of PMU flows in C-DAX requires less signalling, and at the same time takes place faster.

V. Conclusion

In this paper, we propose an ICN-based communication platform, C-DAX, for the scalable and flexible support of PMU-based RTSE in ADNs. Employing a topic-based publish-subscribe communication paradigm, C-DAX spatiotemporally decouples data producers (PMUs) and consumers (PDCs, historians) through the use of DBs and effectively eliminates the need for pair-wise communication. This simplifies communication (re)configuration especially in dynamic scenarios (e.g., in flexible grid topology or asset change). It also enables bandwidth savings through emulation of multicast, thus significantly reducing the infrastructure investment costs to upgrade current aged facilities for future smart grid operations. Our experiments, based on a realistic prototype deployment, show that C-DAX indeed is capable of meeting all the RTSE requirements. The platform enables correct state estimations at the RTSE while also being capable of seamlessly handling component failures. Through our further analysis based on large-scale simulation scenarios with real grid data, we also found that C-DAX provides significant bandwidth savings while operating well within the RTSE requirements. Hence, from our study, C-DAX offers both qualitative (in terms of flexibility, scalability and resilience) and quantitative (in terms of latency and bandwidth savings) benefits against the existing solutions for supporting PMU-based RTSE at distribution level networks.

References

[1] North American Synchrophasor Initiative, www.naspi.org.
[2] M. Chenine, K. Zhu and L. Nordström, "Survey on priorities and communication requirements for PMU-based applications in the Nordic region," IEEE Power Tech Conference, Bucharest, Romania, June 2009.

[3]   G. Xylomenos *et. al.*, "A Survey of Information-Centric Networking Research," IEEE Communications Surveys and Tutorials (DOI 10.1109/SURV.2013.070813.00063; available online since 19 Jul. 2013)

[4]   K. V. Katsaros, *et. al.*, "Information-centric networking for machine-to-machine data delivery – A case study in smart grid applications," IEEE Networks Magazine, 2014 (to appear).

[5]   W. K. Chai, *et. al.*, "Cache 'less for more' in information-centric networks (extended version)," Elsevier Computer Communications, vol. 36, no. 7, pp. 758-770, 2013.

[6]   Y. J. Kim *et. al.*, "A Data-centric Information Infrastructure for Smart Grid," IEEE Communication Magazine, Vol. 48, Issue 11, 2010

[7]   J. Zhang *et. al.*, "iHEMs: An Information-centric Approach to Secure Home Energy Management," Proc. IEEE SmartGridComm, 2012

[8]   Y. J. Kim *et. al.*, "SeDAX: A Scalable, Resilient, and Secure Platform for Smart Grid Communications," IEEE Journal on Selected Areas in Communications, Vol. 30, No. 6, 2012

[9]   A. Abur and A. Gómez-Expósito, "Power System State Estimation: Theory and Implementation," New York: Marcel Dekker, 2004.

[10] A. Monticelli, "State estimation in electric power systems," New York, Springer 2000.

[11] K. Christakou, *et. al.*, "Efficient computation of sensitivity coefficients of node voltages and line currents in unbalanced radial electrical distribution networks," IEEE Trans. on Smart Grid, vol. 4, no. 2, pp. 741-750, 2013.

[12] G. T. Heydt, "The next generation of power distribution systems," IEEE Transactions on Smart Grid, vol. 1, no. 3, Dec. 2010.

[13] IEEE Guide for Design, Operation, and Integration of Distributed Resource Island Systems with Electric Power Systems, IEEE Standard 1547.4, 2011.

[14] IEEE Recommended Practice for Interconnecting Distributed Resources with Electric Power Systems Distribution Secondary Networks, IEEE Standard 1547.6, 2011.

[15] IEEE Guide for Smart Grid Interoperability of Energy Technology and Information Technology Operation with the Electric Power System (EPS), End-Use Applications, and Loads, IEEE Standard 2030, 2011.

[16] Development and Operation of Active Distribution Networks, Cigré Working Group C6.11, Apr. 2011.

[17] R. Singh, B. C. Pal, and R. A. Jabr, "Choice of estimator for distribution system state estimation," Generation, Transmission & Distribution, IET, vol. 3, no. 7, pp. 666-678, Jul. 2009.

[18] L. Junqi, *et. al.*, "Trade-Offs in PMU Deployment for State Estimation in Active Distribution Grids," IEEE Trans. on Smart Grid, vol.3, no.2, pp.915-924, June 2012.

[19] 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.

[20] A. G. Expósito, *et. al.*, "A multilevel state estimation paradigm for smart grids," in Proc. of the IEEE, vol.99, no.6, pp.952-976, June 2011.

[21] S. Sarri, *et. al.*, "State Estimation of Active Distribution Networks: Comparison Between WLS and Kalman-Filter Algorithms Integrating PMUs," 3rd IEEE PES Innovative Smart Grid Technologies (ISGT) Europe Conference, Berlin, Germany, October 14-17, 2012.

[22] K. Jones, *et. al.*, "Three-Phase Linear State Estimation with Phasor Measurements," IEEE PES Gen. Meet., Vancouver, Canada, Jul., 2013.

[23] "IEEE Standard for Synchrophasor Measurements for Power Systems," IEEE Std C37.118.1-2011 (Revision of IEEE Std C37.118-2005) , pp.1,61, Dec. 28 2011

[24] "IEEE Standard for Synchrophasor Measurements for Power Systems – Amendment 1: Modification of Selected Performance Requirements," IEEE Std C37.118.1a-2014 (Amendment to IEEE Std C37.118.1-2011) , pp.1,25, April 30 2014

[25] P. Romano and M. Paolone, "Enhanced Interpolated-DFT for Synchrophasor Estimation in FPGAs: Theory, Implementation, and Validation of a PMU Prototype", Instrumentation and Measurement, IEEE Transactions on, 2014.

[26] P. Romano and M. Paolone, "An Enhanced Interpolated-Modulated Sliding DFT for High Reporting Rate PMUs', Applied Measurements for Power Systems (AMPS), 2014 IEEE International Workshop on, 2014

[27] IEEE Distribution Planning Working Group, "Radial distribution test feeders," *IEEE Trans. on Power Syst.*, vol. 6, pp. 975-985, 1991.

[28] L. Zanni, *et. al.*, "Probabilistic assessment of the process-noise covariance matrix of discrete Kalman filter state estimation of active distribution networks," in Proc. Intern. Conf. of Prob. Methods Applied to Power Syst., Durham, UK, July 7-10, 2014.

[29] N. M. Manousakis, *et. al.*, "Taxonomy of PMU placement methodologies," *IEEE Trans. on Power Syst.*, vol. 27, no. 2, pp. 1070-1077, May 2012.

[30] B. Xu, and A. Abur, "Observability analysis and measurement placement for systems with PMUs," in Proc. of the IEEE PES Power Syst. Conf. and Exp. 2004, New York, USA, vol. 2, pp. 943-946, Oct. 10-13, 2004.

[31] "IEEE Standard for Synchrophasor Data Transfer for Power Systems," IEEE Std C37.118.2-2011 (Revision of IEEE Std C37.118-2005) , pp.1,53, Dec. 28 2011

[32] Q. Li, *et. al.*, "An information-theoretic approach to PMU placement in electric power systems," IEEE Trans. on Smart Grid, vol. 4, no. 1, 2013.

[33] M. Seewald, "Building an architecture based on IP-Multicast for large phasor measurement unit (PMU) networks," Innovative Smart Grid Technologies (ISGT), 2013 IEEE PES, Feb. 2013.

[34] M.Pignati, *et. al.*, "A Pre-Estimation Filtering Process of Bad Data for Linear Power Systems State Estimators using PMUs", 18th Power Systems Computation Conference (PSCC), Wroclaw, Poland, Aug. 2014.

[35] A. Varga, *OMNeT++ Simulator Home Page*, http://www.omnetpp.org.

[36] R. Francis *et al.*, "Distance Constraints for Tree Network Multifacility Location Problems," Operations Research, vol. 26, no. 4, 1978.

[37] S. Galli *et al.*, "Power Line Communications and the Smart Grid," in *Proc. of IEEE SmartGridComm*, 2010, pp. 303–308.