

Cost-Efficient Low Latency Communication Infrastructure for Synchrophasor Applications in Smart Grids

Binxu Yang, Konstantinos Vasileios Katsaros, Wei Koong Chai, and George Pavlou

Abstract—With the introduction of distributed renewable energy resources and new loads, such as electric vehicles, the power grid is evolving to become a highly dynamic system that necessitates continuous and fine-grained *observability* of its operating conditions. In the context of the medium voltage (MV) grid, this has motivated the deployment of phasor measurement units (PMUs), that offer high-precision synchronized grid monitoring, enabling mission-critical applications such as fault detection/location. However, PMU-based applications present stringent delay requirements, raising a significant challenge to the communication infrastructure. In contrast to the high voltage domain, there is no clear vision for the communication and network topologies for the MV grid; a full-fledged optical fiber-based communication infrastructure is a costly approach due to the density of PMUs required. In this study, we focus on the support of low-latency PMU-based applications in the MV domain, identifying and addressing the tradeoff between communication infrastructure deployment costs and the corresponding performance. We study a large set of real MV grid topologies to get an in-depth understanding of the various key latency factors. Building on the gained insights, we propose three algorithms for the careful placement of high capacity links, targeting a balance between deployment costs and achieved latencies. Extensive simulations demonstrate that the proposed algorithms result in low-latency network topologies while reducing deployment costs by up to 80% in comparison to a ubiquitous deployment of costly high capacity links.

Index Terms—Delay, medium voltage power grid, phasor measurement units (PMUs), real topologies, synchronization.

I. INTRODUCTION

THE energy sector has been undergoing major transformative changes in recent years in order to address pressing concerns in improving energy efficiency of the grid and to reduce overall carbon emissions. The increasing penetration of distributed renewable energy sources (DRER) (e.g., solar/wind farms), the rising deployment of electric vehicles (EV) [1], [2], and active consumer participation into power grid operations (e.g., interactive consumer applications) are pushing today's power grid infrastructure to the limit. The progressive integration of these active components introduces significantly higher system volatility, posing new challenges to system stability, with respect to power quality, voltage regulation, protection [3], and

fault location. In fact, this constitutes a major shift from passive to active distribution networks (ADNs).¹

To cope with this increasing volatility, distribution network operators (DNOs) aim at the design and development of enhanced cyber-physical systems enabling both the fine-grained *monitoring* and *control* of their power grid infrastructure. In the envisioned systems, a communication infrastructure supports the near-real time *observability* of the power grid conditions, enabling in turn the control of the power grid infrastructure in terms of the aforementioned control operations. In this context, the deployment of high-precision phasor measurement units (PMUs) [4] gains a significant role for DNOs. By supporting high rate, synchronized monitoring of key system parameters, PMUs enable the synchrophasor-based real-time state estimation (RTSE) [5] of the grid, opening the way for fine grained and timely control of the overall system [6]. For example, fault localization enables the instant identification and the subsequent opening/closing of the appropriate breakers, isolating the fault. It has become apparent that the close synergy of communications and the power grid will enable its fine-grained management, supporting the timely adaptation to increasingly dynamic operating conditions.

However, such applications come with stringent end-to-end communication delay requirements, i.e., in the order of a few tens of milliseconds [6]–[8]. In turn, the expected benefits from the envisioned cyber-physical system depend on the ability of the communication infrastructure to actually support these requirements. While high-capacity optical fiber may be typically available on the transmission level (i.e., high voltage (HV) domain), adopting a similar approach on the distribution level (i.e., in the medium voltage (MV) domain) raises significant concerns with respect to the associated costs. Our analysis of a large set of real topologies (cf., Section II-B) shows that the mostly urban environment of the distribution grid calls for a dense deployment of high-capacity communication links, as opposed to the HV domain [9]. As a result, the full-fledge fiber optic communication deployment in urban environment for MV distribution grid is currently not practical and plagued with various difficulties and prohibitive costs. Recent works have alternatively investigated the use of wireless technologies such as WiMAX and LTE, reporting however concerns about the impact of control plane and medium access control (MAC) layer delays, which is directly

Manuscript received November 20, 2015; revised March 27, 2016; accepted April 10, 2016. This work supported by the European Commission's Seventh Framework Program FP7-ICT-2011-8 under Grant 318708 (C-DAX) and the CHIST-ERA/EPSRC UK project CONCERT under the Grant EP/L01835/1.

The authors are with the University College London, London WC1E 6EA, U.K. (e-mail: binxu.yang@ucl.ac.uk; k.katsaros@ucl.ac.uk; w.chai@ucl.ac.uk; g.pavlou@ucl.ac.uk).

Digital Object Identifier 10.1109/JSYST.2016.2556420

¹ADNs are defined as distribution networks that have systems in place to control a combination of distributed renewable energy resources like generators, loads and storage. DNOs have the possibility of managing electricity flows via a flexible network topology.

affected by the number of devices accessing the high-capacity wireless channel(s) [8], [10], [11]. On the other hand, the readily available power-line communication (PLC) [12] infrastructure has relatively low costs, but the typically low PLC bandwidth appears as a bottleneck to the timely delivery of delay sensitive monitoring traffic.² Based on the above observations, we identify the tradeoff between the performance gains from the deployment of high bandwidth technologies and the deployment costs (and/or MAC/signaling delay penalties in the case of wireless technologies) associated with wide scale PMU deployment in the MV domain. We highlight that this is the first work in MV domain investigating low latency communication infrastructure for PMU-based applications.

We address this tradeoff by considering the design of a hybrid communication infrastructure, where the existing PLC infrastructure is utilized to reduce the number of high capacity links required to satisfy the low latency requirements along with the associated costs. Our problem resembles a facility location problem, where we seek the minimum number and location of high-capacity links in the MV grid to satisfy our application-level latency constraints. As the problem is known to be NP-hard [13], [14], we turn our attention to heuristic-based solutions. To this end, and in order to guide the design of our solution, we engage in an in-depth analysis of the end-to-end delay (T_{e2e}) components. Based on a large set of 14 real MV grid topologies operated by a large DNO in the Netherlands, we perform an analysis of important topological characteristics of the MV domain [15], while also paying attention to PMU communication specificities such as the impact of precise PMU data synchronization. Our analysis yields valuable and pragmatic insights for the design of both low-cost and low-latency communication infrastructures for the MV grid, which we embody in the design of three different heuristic-based optimization algorithms. An extensive set of detailed packet level simulations demonstrate the effectiveness of our algorithms.

Our study here is based on our preliminary study in [16]. Summarizing, the contributions of this study are as follows.

- 1) We identify the tradeoff between communication infrastructure deployment costs and application-level latencies for low-latency PMU-based applications in the MV domain (see Section II). To the best of our knowledge, this is the first study focusing on this tradeoff for the MV domain of the grid.
- 2) We conduct an in-depth analysis of real MV grid topology characteristics, based on real MV grids comprising 14 primary-substations (P-SS) and 1323 secondary-substations (S-SS). Our analysis yields specific guidelines for the design of low latency communication infrastructures in the MV domain (see Section III).
- 3) We identify and analyze the impact of PMU synchronization on the communication latencies (see Section III).
- 4) We propose three heuristic algorithms for the design of a low-cost and low-latency communication infrastructure in the MV grid (see Section IV). The proposed algorithms are general in nature and applicable to both optical fiber

²As also demonstrated by our detailed packet-level simulations in Section II.

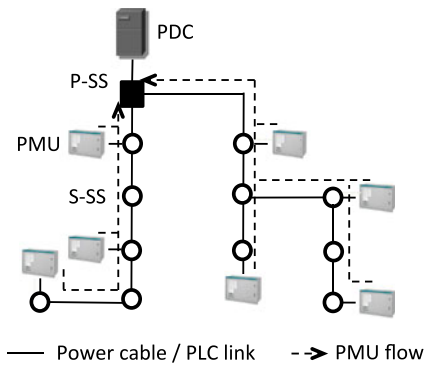


Fig. 1. MV power grid.

and wireless technologies-based communication infrastructures.

- 5) We demonstrate the effectiveness of our algorithms, i.e., the resulting network topologies satisfy the delay requirements while requiring up to 80% less high capacity links compared to the HV model of ubiquitous high-capacity link deployment (see Section V).

II. BACKGROUND AND PROBLEM STATEMENT

Designing a communication infrastructure for the support of a purpose specific cyber-physical system, such as the smart power grid, necessitates a good understanding of the operational context, in our case of the MV power grid. Fig. 1 provides a high level illustration of a typical MV grid, i.e., the (power) distribution network. A typical MV grid topology has a tree-like structure rooted at a P-SS, which is responsible for stepping down the transmission voltage from HV to MV. Each tree branch emanating from a P-SS corresponds to a distinct feeder (cable) further distributing the MV power to the desired areas through a series of S-SSes, responsible for further stepping down the voltage. The power distribution network consists of multiple such trees rooted at different P-SSes.

A. Delay-Sensitive Synchrophasor Monitoring Applications

Our study is motivated by the challenge to support three-phase RTSE application. RTSE is considered as an important tool for DNOs as it supports particularly important energy management and protection operations, such as fault detection/localization, postfault management, and voltage control [17], [18]. PMUs enable the support of such applications by monitoring power system parameters (e.g., phase angle, voltage, rate of change of frequency (ROCOF), etc.) at strategically selected S-SSes in the MV grid³(see Fig 1). All PMUs are GPS-synchronized and stream their measurements to phasor data concentrators (PDCs), which are typically located at the P-SS. PDCs collect, time align, and deliver synchrophasor data to applications such as the RTSE.

³The selection of PMU locations constitutes a research area on its own (e.g., [19]). Without loss of generality, we consider a scenario with a PMU deployed at approximately every two S-SSes along a feeder (see Fig. 1).

Although typical refresh rates of state estimation processes are of the order of a few minutes, the high system dynamics of ADNs, due to DRERs and EVs, necessitate the fine-grained estimation of system state within a few tens/hundreds of ms [5]. PMU reporting frequencies (F) of 50 or 60 frames/s facilitate this detailed view of the power grid [20]. Based on PMU data semantics [21], a realistic PMU message payload size is 102 B.⁴ Further considering UDP and IP headers, and a 32-B SHA-256 message authentication code, the overall data rate for each RTSE PMU flow delivered to the link layer is 64.8 Kb/s, for $F = 50$ Hz.

The timely delivery of these measurements is a challenge for the underlying communication infrastructure. In this study, we account for RTSE applications a maximum total latency of 100 ms [5], [7], [8], including latencies for PMU signal acquisition, PMU synchrophasor estimation and data encapsulation, communication network delay, PDC data frame time alignment, bad data detection, and state estimation [7]. The time budget left for telecommunication network delay (T_{e2e}) depends on these latency components and has typically a constraint (denoted as T_{max}) of 20 ms, at a PMU reporting rate of 50 Hz [5], [8]. It was recently shown that the telecommunication network delay constraint could be further relaxed to 35–55 ms due to new advancements in state estimation algorithms [7]. Nevertheless, in this study, we focus on $T_{max} = 20$ ms, as a more stringent requirement.⁵ At this point, it is important to stress that at the application level, RTSE necessitates the availability of *all* synchronized PMU measurements within the defined T_{max} . Otherwise, state estimation will suffer in terms of accuracy; hence, T_{max} stands for the worst case T_{e2e} acceptable.

B. Problem Statement

The support of the identified latency requirements depends heavily on the underlying communication infrastructure, which in turn is largely determined by the locations of the communicating entities and the selected transmission technology. We first consider a baseline communication network model enabled by PLC technologies [12], which, by allowing DNOs to make use of their existing power-line cables as the transmission medium, constitute the most straightforward and low-cost option for the support of communications in the power grid. In this baseline scenario, the communication network topology coincides with the MV power grid topology. We investigate the topological properties of the resulting communication network model based on a set of 14 MV power grid topologies operated by a DNO in the Netherlands. Table I summarizes the basic aggregated topological characteristics of the considered MV grids. Furthermore, in Table II, we present the topological properties per area. Our dataset shows close agreement with literature (e.g., as surveyed in [22]) and thus, representative to general MV grids.

We represent the distribution grid, and the corresponding baseline communication network model, as a set of tree graphs,

⁴Considering PHNMR=6, ANNMR=6 and DGNMR=2, with 32-bit floating-point accuracy [21].

⁵We note though that this is only an input parameter to the proposed algorithms (see Section IV), not affecting their general applicability.

TABLE I
SUMMARY OF REAL MV GRID TOPOLOGICAL PROPERTIES OF A LARGE EUROPEAN DNO

P-SS	14
S-SS	1323
Number of edges (cables)	1426
Average cable length	498 m
Average node degree	2.02

TABLE II
REAL MV GRID TOPOLOGICAL PROPERTIES PER AREA

Grid	Number of nodes	Number of edges	Mean node degree	Link ^a density	Mean ^b path length	Mean ^c betweenness
Area 1	187	223	2.0744	0.0128	7.3105	6.2774
Area 2	112	134	2.7077	0.0216	7.745	6.6869
Area 3	36	43	1.9545	0.0683	6.2778	5.1351
Area 4	155	177	2.1718	0.0148	7.1290	6.0897
Area 5	89	102	2.125	0.02604	7.1290	6.2247
Area 6	82	82	1.9759	0.02469	3.7195	2.6867
Area 7	22	22	1.9130	0.0952	3.4545	2.3478
Area 8	177	177	1.9887	0.01136	5.3728	4.3483
Area 9	28	28	1.9887	0.07407	5.7857	4.6207
Area 10	50	51	2	0.04163	4.5	3.4313
Area 11	101	101	1.9803	0.02	5.5049	4.4608
Area 12	98	98	1.9798	0.02061	4.55102	3.5152
Area 13	41	41	1.9524	0.05	2.5854	1.5476
Area 14	145	147	2.0119	0.0141	5.2897	4.2603

^aPath length represents the number of hops from an S-SS to the P-SS.

^bLink density = $\frac{|E|}{(|V|-1) \cdot |V|/2}$.

^cBetweenness represents the number of shortest paths between an S-SS and the P-SS that involve the measured node.

$G(V, E)$, with $v \in V$ as substations where node v_0 represents the root (i.e., the P-SS).⁶ The edges, $e \in E$, represent physical cables connecting different S-SSes. Then, we denote the distance in hop count between v_i and $v_{i'}$ as $d(v_i, v_{i'})$ with $i \neq i'$. Further, let U be the set of nodes (S-SSes) equipped with PMUs, comprising PMU-enabled nodes v_i^j , where $j \in [0..|U| - 1]$ is the PMU index and $i \in [0..|V| - 1]$ is the node index. We define $P(v_i^j)$ as the shortest path comprising the consecutive edges connecting PMU-enabled node, v_i^j , to v_0 (see dashed lines in Fig. 1). The length of $P(v_i^j)$ is $|P(v_i^j)| = d(v_i^j, v_0) = d(v_i, v_0)$.⁷

In the PLC-enabled baseline model, PMU flows (dashed arrows in Fig. 1) reach the PDC by traversing their uphill PLC links toward the root of the tree topology. Following the PMU deployment scheme described in Section II-A, for the available MV grid topologies, we simulate the operation of 795 PMUs in a detailed packet-level simulation environment (see Section V). Fig. 2 shows the cumulative fraction of the T_{e2e} observed at the PDC for a duration of 10 min with PLC bandwidth values: 100 and 500 Kb/s [12].⁸ The vast majority of PMU messages

⁶On the communication level, v represent routing/switching devices located at the corresponding S-SSes, forwarding data packets.

⁷For clarity, for the rest of the paper, we simply refer $P(v_i^j)$ as P_j since there is only one unique path from a PMU to the PDC.

⁸PLC encompasses a diverse set of technical realizations with different bandwidth values, broadband 500 Kb/s being one of them. Our methodology can be applied for different bandwidth values. For extremely low values, lower datarate PMU configurations should obviously be considered.




Fig. 2. T_{e2e} of PMU flows with PLC and optical fiber.

delivered exceeds T_{max} . Clearly, the considered set of applications cannot be supported by PLC technology alone. However, we will show in Section III that limited bandwidth is not the only key delay factor.

We further consider and simulate an optical-fiber based communication network model, following the current practice in HV deployments [9]. In particular, we consider 10 Gb/s optical fiber links directly connecting PMU-enabled S-SSes to the PDC at P-SS. As shown in Fig. 2, this communication infrastructure fully conforms to the T_{max} constraint. However, it necessitates the deployment of 795 optical fiber links in total, representing a significant capital expenditure (CAPEX).

Recent studies have also shown that the adoption of wireless technologies may lead to an increase of medium access delays due to the contention for access to the shared wireless medium, even in cases where no other background traffic is served [8], [10], [11]. This contention and the corresponding delays increase with the number of wireless transmitting devices, i.e., subject to the selected wireless technology, an increased volume of attempts to transmit increases the collision probability, leading to back-off/scheduling delays. Given that existing wireless networks (e.g., cellular (A-)LTE, WiMAX) have been dimensioned for a particular access load, the introduction of additional devices (i.e., PMUs) raises concerns about the aforementioned performance penalties. Of course, increasing frequency reuse with the deployment of smaller cells would reduce contention, for a certain access demand. However, this would obviously come at a significant deployment cost for communication network operators.⁹ The synchronization of PMUs only further exacerbates the contention issue, since it increases collision probabilities and/or limits scheduling flexibility. For all these reasons, it follows that the number of wireless transmitting devices should also be kept to a minimum.

In short, PLC, though readily available, appears unable to support the considered low latency applications, urging for alternative solutions such as the use of modern wireless or high-speed wired technologies. However, the deployment of such technologies incurs a non-negligible capital expenditure and/or performance penalties. In this respect, it becomes apparent that the scale of deployment of high capacity links needs to be carefully considered. Considering this tradeoff between deployment

⁹Dedicated, private wireless networks constitute another option for DNOs. However, they are associated with other types of deployment costs, e.g., spectrum license costs. We consider this particular aspect out of the scope of this paper.




Fig. 3. Hybrid communication infrastructure.

costs and performance, we propose the design of hybrid communication infrastructures that exploit the existing low cost PLC capabilities, while also employing higher bandwidth technologies. The rationale is to take advantage of the availability of PLC to *partially* accomplish the task of delivering the PMU data flows to the PDC, thus reducing the number of high capacity links in the overall network. Starting from our baseline network model, the objective then becomes to select the minimum sub-set of S-SSes to be equipped with high capacity direct links to the P-SS/PDC (e.g., optical fiber) and act as *sink* nodes, i.e., aggregating PMU traffic through PLC links. The envisioned scenario is illustrated in Fig. 3.

Let X_i with $i \in [0..|V| - 1]$ be a binary decision variable, set to 1 if node v_i is equipped with a high-capacity communication link; we denote such a node with v_i^k with $k \in [0..|S| - 1]$, where S be the set of sink nodes. Let also Y_{jk} with $j \in [0..|U| - 1]$ and $k \in [0..|S| - 1]$ be a binary variable set to 1 if a PMU flow from v_i^j is delivered to a sink node v_i^k .¹⁰ Then, denoting the end-to-end delay of each PMU flow over path P_j (with $j \in [0..|U| - 1]$), as T_{e2e}^j , our objective can be loosely expressed as follows:

$$\text{minimize } \sum_i X_i,$$

$$\text{subject to } T_{e2e}^{P_j} \leq T_{max} \quad \forall j \in [0..|U| - 1] \quad (1)$$

$$\sum_k Y_{jk} = 1 \quad \forall j \in [0..|U| - 1] \quad (2)$$

$$k \in [0..|S| - 1]$$

$$X_i \in \{0, 1\} \quad \forall i \in [0..|V| - 1] \quad (3)$$

$$Y_{jk} \in \{0, 1\} \quad \forall j, k. \quad (4)$$

The exact nature of the problem and the corresponding solution obviously depend on the first constraint which only roughly expresses the low latency requirement. The second constraint ensures that each PMU-enabled node sends its flow to a single sink node. To assess the hardness of our problem, we can merely express the first constraint by setting an upper limit (i.e., d_{max})

¹⁰Note that $i = i'$ is allowed i.e., an S-SS can be equipped with both a PMU and a high capacity link.

for the distance between a PMU-enabled node v_i^j and the corresponding sink location $v_{i'}^k$. This results in constraint (1) to be re-written as follows (see also Section III-B1):

$$d(v_i^j, v_{i'}^k) \leq d_{\max}.$$

Even in this simple case, the resulting problem is a typical NP-hard *facility location* optimization problem [13], [14] thus turning our attention to heuristic-based solutions. To further explore the problem space and guide the design of our heuristics, we first decompose T_{e2e} into its constituents and investigate the most important factors impacting them (see Section III). In this effort, we get valuable input from the detailed investigation of our large set of MV topologies. Our analysis yields important insights for the subsequent design of the proposed heuristic algorithms (see Section IV).

III. ANALYSIS OF IMPACT FACTORS

Our analysis of the various latency impact factors is enabled by the identification of the various components of T_{e2e} .

- 1) Processing delay (proc): The time used for operations such as medium adaptation, (de)coding, switching, routing, message authentications codes generation/verification.
- 2) Propagation delay (prop): Depends on the transmission medium and the distance travelled by the signal. For copper cable, this is typically 5 ns/m.
- 3) Transmission delay (trans): The time required to transmit the data and is subject to the bandwidth of the underlying transmission technology.
- 4) Queuing delay (queue): The time spent by data waiting for transmission at the transmitting devices.

We consider for this analysis a discrete time domain divided into slots with each slot capable of containing exactly one PMU packet. For each delay component $x \in \{\text{proc}, \text{prop}, \text{trans}, \text{queue}\}$, we consider the corresponding per hop delay t_x . Additionally, we define the aggregate T_x of each delay component x over a path P_j as $T_x = \sum^{|P_j|} t_x$.

In the following sections, we investigate the impact of the key factors affecting the aforementioned delay components in order to get insights on where to place high capacity links to achieve the low latency requirement in a cost-efficient manner.

A. Bandwidth

Bandwidth availability impacts both $t_{\text{trans}}/T_{\text{trans}}$ and $t_{\text{queue}}/T_{\text{queue}}$. Obviously, $t_{\text{trans}}/T_{\text{trans}}$ increase with lower bandwidth values. Moreover, queuing delays perceived at a node increase when the available output bandwidth is lower than the incoming data rate at the node.¹¹ Fig. 4(a) shows the cumulative fraction of the T_{trans} across all PMU-to-PDC paths, for the cases of PLC and optical fiber-based communication infrastructures. For the PLC-based case, T_{trans} exceeds T_{\max} for 92.91% and 26.59% of the transmitted packets for the cases of 100 and 500 Kb/s, respectively. In the case of optical fiber, we

see a considerable reduction of accumulated T_{trans} compared to the PLC case, leaving abundant delay budgets for other delay components. This is a direct consequence of the reduction of t_{trans} values from 13.52 or 2.70 ms for 100 and 500 Kb/s PLC datarates, respectively, to only $t_{\text{trans}} = 13.52 \mu\text{s}$ for the case of optical fiber (for the considered payload size and header overheads; see Section II-A).

Our simulations for the baseline PLC scenario (see Section II) also indicate that on average, T_{queue} accounts for 96.88% of T_{e2e} , with T_{trans} and T_{proc} accounting only for 2.18% and 0.94%, respectively.¹² This domination of T_{queue} on T_{e2e} implies the lack of sufficient bandwidth to support the PMU traffic. Although the perceived T_{queue} and T_{trans} evidently demonstrate the role of the adopted technology's bandwidth, they are still dependent on a series of other factors including the communication network topology and the tight synchronization of PMUs. We further investigate these aspects next.

B. Topology

1) *Path Length*: The path length, $|P_j|$, has an important impact on perceived aggregate T_x latencies, since lengthy paths accumulate delays on multiple hops. We further use our set of MV topologies to realistically quantify this impact. Fig. 4(b) shows the cumulative fraction of the processing delays accumulated by data packets across all PMU-to-PDC PLC paths (i.e., $P_j : \forall j \in [0, \dots, |U| - 1]$), for a range of per node processing delay values, t_{proc} . These values depend on the computational resources of the forwarding devices and can vary significantly, ranging from a few micro-seconds to even milliseconds per packet [23]. If we consider recent overlay approaches [7], [24], these delays may further increase due to the transition of packets from the kernel to the user space. We notice that, subject to t_{proc} , the overall delay penalty T_{proc} may get close or even exceed T_{\max} , even though T_{proc} constitutes only 0.94% of T_{e2e} (for $t_{\text{proc}} = 1$ ms). Similarly, as previously discussed, Fig. 4(a) shows T_{trans} values close to T_{\max} , though T_{trans} constitutes only 2.18% of T_{e2e} . This is a direct effect of path lengths, which in our topologies have an average and maximum value of 5.84 and 20 hops, respectively.

In essence, these measurements yield an important guideline for the design of low latency communication networks: in the presence of high t_{proc} values (i.e., in the order of 1 ms), bandwidth availability alone may not suffice in keeping T_{e2e} values low, when paths are considerably long, e.g., interconnecting P/S-SS with optical fiber, following the power grid topology. Moreover, the provisioning of computational resources at each forwarding node should be carefully considered.

Building on these observations, we reformulate the first constraint of our optimization problem (1). Namely, to limit the effect of path lengths on T_{e2e} , we constrain the maximum number of PLC hops by limiting the distance between a PMU and

¹¹As queuing delays are also related to both topological aspects of the communication network and the synchronization of PMUs, we discuss them in detail in Section III-B

¹²Due to the short distances between S-SSes (see Table I), we omit $t_{\text{prop}}/T_{\text{prop}}$ in the following as it is only in the order of microseconds.




Fig. 4. Impact of topology on PMU application performance (a) Accumulated T_{trans} when using various technologies and t_{trans} , (b) Accumulated T_{proc} when using purely PLC technology with various t_{proc} , (c) PMU bandwidth requirements; traffic volume exceeds PLC capabilities.

its sink node (d_{max})¹³:

$$d(v_i^j, v_i^k) \leq d_{\text{max}} = \left\lfloor \frac{T_{\text{max}} - t_{\text{proc}}}{t_{\text{trans}} + t_{\text{proc}}} \right\rfloor$$

$$\forall i \in [0..|V| - 1] \forall j \in [0..|U| - 1] \quad \forall k \in [0..|S| - 1] \quad (5)$$

For the cases of 100 and 500 Kb/s PLC, we get $d_{\text{max}} = 1$ and 5 hops, respectively, as d_{max} 's limit values.

2) *Application-Level Betweenness*: As previously mentioned, T_{queue} constitutes 96.88% of T_{e2e} . This delay component depends on the relation between the available and the required bandwidth at each forwarding device. While the former depends on the selected transmission technology, the latter depends on topological aspects of the communication network. Fig. 4(c) shows the cumulative fraction of the total PMU traffic volume aggregated at each PLC link toward the PDC in our MV topologies. Again, we see that a PLC-based infrastructure fails to accommodate the resource requirements as for more than half of the communication nodes, the bandwidth requirements exceed a typical bandwidth value of 100 Kb/s ($\approx 10\%$ for 500 Kb/s links).

To better understand this aspect, we introduce the concept of *application-level betweenness*, $b(v_i)$, as the number of shortest paths P_j crossing node, v_i . Note that $b(v_i)$ is determined both by the topology structure and the placement of the PMUs. In the considered set of MV grid topologies, we observe an average and maximum $b(v_i)$ value of 3.24 and 32, respectively. Considering a 64.8 Kb/s data rate per PMU flow, it is easy to understand the domination of T_{queue} in T_{e2e} .

Building on this observation, we formulate the next constraint for the design of our hybrid communication network topologies, i.e., we impose an upper bound on application-level betweenness (b_{max}) throughout the topology:

$$b(v_i) \leq b_{\text{max}} = \left\lfloor \frac{\text{BW}}{\text{DR}} \right\rfloor \quad \forall i \in [0..|V| - 1] \quad (6)$$

where BW is the available PLC bandwidth and data rate, DR = 64.8 Kb/s. For BW = 100 and 500 Kb/s, this yields $b_{\text{max}} = 1$ and 7, respectively, significantly lower than the observed $b(v_i)$ values in the baseline network model.

¹³ t_{prop} (average $\leq 3 \mu\text{s}$ in the considered topologies) and t_{trans} on the sink-to-PDC link ($\leq 2 \mu\text{s}$ for a 10Gb/s optical fiber link) are considered negligible. However, we account the t_{proc} for the sink-to-PDC hop.




Fig. 5. Example of path $P_{j'}$ joining path P_j at node u_i .

C. PMU Synchronization

Another factor with significant impact on the T_{e2e} is the synchronized nature of PMU flows.¹⁴ As briefly mentioned in Section II, such synchronization may significantly impact the delays for access to the wireless medium. However, PMU synchronization also has an important impact on the baseline PLC network model. Packets originating at different PMUs reach the same forwarding device at (almost) the same time. Consequently, packets wait in the transmission queue for a time linear to t_{trans} , i.e., waiting until all interfering packets from other PMU(s) get transmitted. Our simulation results show that approximately 20% of PMU flows experience such synchronization problem across 12.45% of forwarding nodes, inflating the overall observed T_{queue} .

To assess the impact of synchronization, we follow the approach proposed in [25]. Specifically, we focus on the worst-case scenario, i.e., a packet has to wait for all other packets (almost) simultaneously arriving the same node, to get transmitted first. We consider this worst-case scenario as our target is to limit the maximum T_{e2e} perceived.

We focus on a node of interest, v , with L inbound links $e_l, l \in [0..|L| - 1]$ and e_f outbound link. Further, a path, $P_{j'}$, is said to join path P_j when they share the same outgoing edge e_f but not an incoming edge e_l at the node of interest (see Fig. 5). Let $R_{P_j}(e_f)$ be the number of paths, $P_{j'}$, that join P_j : $\forall j' \neq j; j', j \in [0..|U| - 1]$, at edge e_f . Then, the *route interference number* (RIN) of path P_j is defined as follows:

$$R(P_j) = \sum_{e_f \in P_j} R_{P_j}(e_f).$$

¹⁴It is worth noting that the synchronization issue did not draw much attention in the HV domain because of the low PMU deployment density and the high bandwidth of the adopted transmission technologies [9].

By counting the number of interfering paths at each forwarding node toward the PDC, RIN allows us to derive the maximum number of times a PMU packet can be delayed due to synchronization in the case where all PMUs send a single packet. In this case, [25] showed that the overall end-to-end T_{queue} of the packet sent on P_j is bounded by $R(P_j)$.

When the b_{max} constraint is met, the aforementioned single packet case can be generalized into a multipacket case where PMUs send one packet at each measurement interval, $1/F$. This generalization is possible because measurements taken at one interval will only arrive after all measurement packets from preceding intervals have been transmitted. Furthermore, when the number of flows of each link e_f is lower than the maximum $b(v_i) : \forall i \in [0..|V| - 1]$, the worst-case queuing delay of path P_j , $T_{\text{queue}}^{P_j}$, is bounded by a tighter upper bound compared to RIN [26], [27]. To state this delay bound, let $\beta(e_f)$ denote the number of interfering packets at edge, e_f and Q_l denote the number of paths from inbound edge e_l . Then, we express $\beta(e_f)$ and the corresponding worst-case queuing delay, $T_{\text{queue}}^{P_j}$, as follows:

$$\beta(e_f) = \sum_l Q_l - \max_l \{Q_l\} \quad (7)$$

$$T_{\text{queue}}^{P_j} = \sum_{e_f \in P_j} \beta(e_f) t_{\text{trans}} \quad (8)$$

where function $\max\{Q_l\}$ selects at the outbound edge, e_f , the maximum number of Q_l from all inbound, e_l .

We can then extend the notion of worst-case queuing delay bound to include t_{proc} and t_{trans} along the path to the PDC. Then, $T_{\text{e2e}}^{P_j}$ is calculated as follows:

$$T_{\text{e2e}}^{P_j} = \sum_{e_f \in P_j} \{\beta(e_f) t_{\text{trans}} + t_{\text{trans}} + t_{\text{proc}}\}. \quad (9)$$

Based on this formulation, we take into account synchronization when satisfying the constraint

$$T_{\text{e2e}}^{P_j} < T_{\text{max}}. \quad (10)$$

IV. ALGORITHM DESIGN

Building on the identified constraints, we next describe three heuristic-based algorithms for the design of low latency and low cost hybrid communication infrastructures. Each algorithm is tailored for specific application environments.

- 1) The *path length constraint (PLeC) algorithm* selects sink node locations by constraining the length of data delivery paths [with d_{max} , (5)], so that the accumulated T_{proc} , T_{prop} , and T_{trans} are also capped (see Section III). Since it does not cater for bandwidth availability, this algorithm is most suitable for low DR applications (e.g., low DR PMU reporting) and can be employed for environments where multiple (low DR) applications share the same communication infrastructure.
- 2) The *application-level betweenness and PLeC (AB-PLeC) algorithm* selects the sink locations by constraining both path lengths and the number of PMU flows on each PLC

link [with b_{max} , (6)]; therefore, explicitly targeting the reduction of T_{queue} . By adjusting the b_{max} constraint according to the residual bandwidth of each link, AB-PLeC can be easily adapted to cater for background traffic, i.e., from applications expected to share the same communication infrastructure (e.g., intelligent electronic device-based monitoring).

- 3) The *flow interference and bandwidth constraint (FIB) algorithm* selects the sink locations by explicitly seeking the nodes at which a PMU packet exceeds T_{max} in the worst-case scenario [see (10)], limiting both b_{max} and $\beta(e_f)$ values. In contrast to the first two algorithms, FIB takes synchronization into account; however, it is tailored for cases of dedicated communication infrastructure, i.e., no background traffic.

A. PLeC Algorithm

For the PLeC algorithm, we follow the distance constraint formulation of the p -center facility location problem [14]. We define $S = s_1, \dots, s_m$ as the set of sink nodes, with $1 \leq m \leq |V|$. Further, let $D(S, v_i) = \min\{d(s, v_i) : s \in S\}$, the distance between each node v_i and its nearest sink node. Our objective is to find the minimum set S such that for all $D(S, v_i) \leq d_{\text{max}}$. We solve this problem via the *sequential location procedure* proposed in [14]. Our algorithm (see Algorithm 1) takes as input the tree topology, G and the distance constraint, d_{max} , and outputs the set of selected sink nodes, S , along with set M (see next). For all nodes v_i , we define a distance value $a_i, i \in [0..|V| - 1]$ and a set M_i , which contains the nodes that can use node v_i as their sink node, under the d_{max} constraint. We further set $M = \bigcup_{i \in [0..|V|-1]} M_i$.

The algorithm starts by randomly selecting a leaf node, v_l from G , along with its parent node v_p . Traversing the tree hierarchy toward its root, the algorithm updates the distance value a_p of nodes v_p as in line 9, until it reaches 0. Note that the hierarchy is traversed by removing the visited leaf nodes from the topology. When $a_p = 0$, node v_p is added to the sink node set (function `ADD_SINK(G, S, v_m)`, line 20). In this step, all nodes v_i whose minimum hop distance to the new sink v_m is below their a_i value are added to the M_m set. All nodes assigned to the new sink are also removed from the tree.¹⁵

The outcome of the algorithm consists of the sets M_i for each selected sink node v_i . These sets may overlap with each other in cases where more than one sink nodes reside within the d_{max} range of some node. At the same time, subject to the exact topological characteristics of tree G , sets M_i may not all have the same size. This means that a careless assignment of nodes to sinks may result in the overloading of some sink nodes both with respect to their processing and bandwidth capabilities. We address this through a simple node assignment procedure which balances the load between sink nodes. Based on the available M sets, the procedure first produces sets L_i which hold the set of all sink nodes within d_{max} range of each node v_i . The members

¹⁵This process may result in a forest. Structure G is used for all trees, and `getRandomLeafNode()` (line 6) returns a leaf node randomly selected from any of the trees.

Algorithm 1: PLeC algorithm.

Input: G, d_{\max}
Output: S

- 1: $S \leftarrow \emptyset$
- 2: **for all** i **do**
- 3: $a_i \leftarrow d_{\max}$
- 4: **end for**
- 5: **while** $G \neq \emptyset$ **do**
- 6: $v_l \leftarrow G.getRandomLeafNode()$
- 7: **if** $v_l \neq v_0$ **then**
- 8: $v_p \leftarrow v_l.getParentNode()$
- 9: $a_p \leftarrow \min(a_p, a_l - 1)$
- 10: $G.removeNode(v_l)$
- 11: **if** $a_p = 0$ **then**
- 12: $ADD_SINK(G, S, v_p)$
- 13: **end if**
- 14: **else**
- 15: $ADD_SINK(G, S, v_l)$
- 16: **end if**
- 17: **end while**
- 18: **return** S, M
- 19:
- 20: **function** $AddSinkG, S, v_s$
- 21: $S \leftarrow S \cup v_s$
- 22: $R \leftarrow \emptyset$
- 23: **for all** $v_i \in G$ **do**
- 24: **if** $d(u_i, u_s) \leq a_i$ **then**
- 25: $M_s \leftarrow M_s \cup v_i$
- 26: $R \leftarrow R \cup v_i$
- 27: **end if**
- 28: **end for**
- 29: $T \leftarrow G \setminus R$
- 30: **end function**

of each L_i set are ordered in increasing hop distance to v_i . The sink node at the smallest distance is selected. When multiple sink nodes are located at the same distance, the algorithm selects the preferred sink node v_m with the minimum M_m size so as to not overload other sinks which can possibly serve more nodes.

B. AB-PLeC Algorithm

The AB-PLeC algorithm finds the set of sink locations that constrains the number of PMU flows being forwarded by each PLC link while maintaining the d_{\max} constraint. It takes as input the tree graph topology G, d_{\max}, b_{\max} and the maximum number of PMU flows that can be accommodated by a high bandwidth link connecting a sink node to the PDC, b'_{\max} . b'_{\max} is set in a similar way to b_{\max} , considering the available high-capacity link bandwidth value, and it is, therefore, normally expected to be considerably higher than b_{\max} . In addition to a_i , for each node v_i , we define b_i as the current $b(v_i)$. All b_i values are initialized to 0, unless a PMU is attached to the corresponding node (line 4). The tree topology is traversed from the leafs toward the root node, allowing the forwarding of PMU flows over PLC links up to the point where the uplink capacity of a

visited node is exceeded (line 22). This node is then selected to act as a sink location (line 23). PMU flows from additional descendants in the tree may be added, subject to the b'_{\max} value (line 15). Visited nodes and sinks are removed from G and the algorithm terminates when all nodes have been removed. Then, each node in the tree can forward its traffic to its closest ancestor sink node.

C. FIB Algorithm

Based on the delay bound formulation (9), we propose a heuristic algorithm that constrains b_{\max} and the number of interfering packets of each flow (via T_{\max}), precisely identifying the required sink locations.

The algorithm takes as input the tree topology $G, b_{\max}, T_{\max}, t_{\text{trans}}$, and t_{proc} and outputs the set S of sink node locations. The algorithm first creates a set U of the nodes equipped with a PMU and computes b_i for all $i \in [0..|V| - 1]$ (line 8). In the second stage, for each PMU-enabled node v_i^j , FIB parses G toward the PDC accumulating the worst-case delay at each node (line 15). When either the calculated delay at a node v_m in P_j reaches T_{\max} or b_{\max} is violated, the algorithm marks v_m as a hotspot.

Algorithm 2: AB-PLeC algorithm.

Input: $G, d_{\max}, b_{\max}, b'_{\max}$
Output: S

- 1: $S \leftarrow \emptyset$
- 2: **for all** i **do**
- 3: $a_i \leftarrow d_{\max}$
- 4: $b_i \leftarrow (v_i.hasPMU())?1 : 0$
- 5: **end for**
- 6: **while** $G \neq \emptyset$ **do**
- 7: $v_l \leftarrow G.getRandomLeafNode()$
- 8: $v_p \leftarrow v_l.getParentNode()$
- 9: $x \leftarrow \min(a_p, a_l - 1)$
- 10: $y \leftarrow b_p + b_l$
- 11: **if** $v_l.markedAsSink()$ **then**
- 12: $S \leftarrow S \cup v_l$
- 13: $G.removeNode(v_l)$
- 14: **else**
- 15: **if** $y > b'_{\max}$ **and** $v_p \neq v_0$ **then**
- 16: $S \leftarrow S \cup v_p$
- 17: $G.removeNode(v_p)$
- 18: **else**
- 19: $a_p \leftarrow x$
- 20: $b_p \leftarrow y$
- 21: $G.removeNode(v_l)$
- 22: **if** $b_p > b_{\max}$ **or** $a_p \leq 0$ **then**
- 23: $v_l.markAsSink()$
- 24: **end if**
- 25: **end if**
- 26: **end if**
- 27: **end while**
- 28: **return** S

After parsing all nodes in U , the FIB algorithm finds the hotspot of each PMU flow. In the third stage, the algorithm selects a

Algorithm 3: FIB algorithm.

Input: $G, T_{\text{threshold}}, b_{\text{max}}, t_{\text{trans}}, t_{\text{proc}}$
Output: S

- 1: $S \leftarrow \emptyset$
- 2: **for all** v_i in G **do**
- 3: $U \leftarrow \cup(v_i.\text{hasPMU}())?v_i : 0$
- 4: **end for**
- 5: **while** $U \neq \emptyset$ **do**
- 6: $\text{Hotspot} \leftarrow \emptyset$
- 7: **for all** $i \in [0..|V| - 1]$ **do**
- 8: $b_i \leftarrow \text{calculateApplicationBetweenness}()$
- 9: **end for**
- 10: $U.\text{disableHotspotLabel}()$
- 11: **for all** v_i^j in U **do**
- 12: $T_j \leftarrow 0$
- 13: $P_j \leftarrow v_i^j.\text{getShortestPathToPDC}()$
- 14: **for all** v_m in P_j **do**
- 15: $T_j \leftarrow T_j + \text{WorstCaseDelayAt}(v_m)$
- 16: **if** $T_j \geq T_{\text{max}}$ or $\text{check}(b_{\text{max}})$ **then**
- 17: $v_m.\text{markAsHotspot}()$
- 18: $\text{Hotspots} \leftarrow \text{Hotspots} \cup v_m$
- 19: Break
- 20: **end if**
- 21: **if** $v_m.\text{isHotspot}()$ **then**
- 22: Break
- 23: **end if**
- 24: **end for**
- 25: **end for**
- 26: $\text{leafhotspot} \leftarrow \text{Hotspots}.\text{getLeafHotspot}()$
- 27: $S \leftarrow S \cup \text{leafhotspot}$
- 28: $U.\text{removeChildrenPMUs}(\text{leafhotspot})$
- 29: **end while**
- 30: **return** S



leafhotspot (i.e., a hotspot with no hotspot descendants) that is farthest to v_0 and adds it into the sink set, S . In the fourth stage, FIB removes the sub-tree rooted at the selected sink location from G . The above four stages are repeated until all PMU-enabled nodes have been removed from U .

V. PERFORMANCE EVALUATION

We apply the proposed algorithms on the available tree-like MV power grid topologies and derive a series of alternative communication network topologies under specific constraints. Based on the derived topologies, we perform an extensive set of detailed packet level simulations. We focus on the case of 500 Kb/s but similar conclusions apply for the case of 100 Kb/s. We consider each sink node to be connected to the P-SS with a 10 Gb/s optical fiber link and set $t_{\text{proc}} = 1$ ms. Based on the above, we then get $d_{\text{max}} = 5$, $b_{\text{max}} = 7$, and $b'_{\text{max}} = 147$ as the topological metrics that would conform to the desired T_{e2e} requirement. Table III summarizes the results for the various derived topologies. We denote the constraints considered by each algorithm as $\text{PLeC}(d_{\text{max}})$ and $\text{AB-PLeC}(d_{\text{max}}, b_{\text{max}})$. For each topology, we show the percentage of packets measured

TABLE III
SUMMARY OF RESULTING TOPOLOGIES

Sink deployment	% packets $> T_{\text{max}}$	max T_{e2e} (ms)	# sink nodes	% gain
PLeC(2)	0%	14.8	309	61.13%
PLeC(3)	0.25%	22	236	70.31%
PLeC(4)	2.16%	25.6	188	76.35%
PLeC(5)	15.94%	290	160	79.87%
AB-PLeC(3,7)	0%	17.5	256	67.79%
AB-PLeC(4,7)	0.67%	22.9	194	75.59%
AB-PLeC(5,7)	20.51%	39.1	147	81.5%
FIB	0%	19.56	163	79.49%
Full optical fiber	0%	1.14	795	0%

Fig. 6. CDF of T_{e2e} for 500 Kb/s PLC links: PLeC algorithm.Fig. 7. CDF of T_{e2e} for 500 Kb/s PLC links: AB-PLeC and FIB algorithm.

to exceed T_{max} , the maximum T_{e2e} , the total number of sink node locations, i.e., the number of high capacity links required, and the gain in terms of the reduction percentage of fiber links compared to the full optical fiber scheme. Figs. 6 and 7 further show the cumulative fraction of T_{e2e} of all packets, for the various topologies.

We see that FIB, PLeC(2), and AB-PLeC(3,7) fully satisfy the delay constraint while requiring only 163, 309, and 256 sink nodes, respectively. This constitutes a reduction in the order of up to 80% compared to the case of ubiquitous optical fiber deployment, requiring 795 such links. PLeC(2) achieves an overall better performance with median and maximum delay values of 4.7 and 14.8 ms, against 7.4 and 20 ms of AB-PLeC(3,7) respectively. AB-PLeC(4,7) and PLeC(3) closely follow, only slightly exceeding T_{max} for $< 1\%$ of the measured packets, i.e., by 2.9 and 2 ms, respectively. Also, we see that PLeC(4) and AB-PLeC(5,7) achieve a maximum delay value of 25.6 and 39.1 ms. As discussed in Section II-A, these latencies could be acceptable in cases of improved delay budgets [7], lowering the number of




Fig. 8. Each point corresponds to one flow and denotes the length of the path traversed toward the PDC and the number of times the flow may encounter synchronization delays.

sink nodes to 194 and 236, respectively, i.e., an improvement in the other delay components could reduce the high capacity links by approximately 24% and 23%, respectively.

PLeC(5), PLeC(4), AB-PLeC(4,7), and AB-PLeC(5,7) exceed T_{\max} , even though we enforce the d_{\max} and b_{\max} constraint values derived from the considered MV topologies. In the case of PLeC(5), T_{e2e} reaches a maximum of 290 ms. This is because the PLeC algorithm does not take into account the b_{\max} constraint. Indeed, $b(u_i)$ values (for nonsink nodes) in PLeC(5) topologies reach a maximum value of 15, resulting in overloaded uplinks. However, this does not hold for AB-PLeC.

For AB-PLeC(4,7) and AB-PLeC(5,7), the nonconformance is attributed to PMU synchronization. Fig. 8 shows for each flow the relation between the length of the corresponding path to the PDC and the number of times the flow may suffer synchronization events, i.e., its packets arrive at a node (almost) at the same time with packets of other flows.¹⁶ Topology AB-PLeC(4,7) allows a maximum of four hops to a sink node for all PMU flows (hence five to the PDC), which leads to a delay of 15.82 ms including $d_{\max}(t_{\text{trans}} + t_{\text{proc}})$ from PMUs to sinks and a $t_{\text{proc}} + t_{\text{trans}_{10\text{-Gb/s}}}$ from sinks to the PDC. This leaves 4.184 ms as the remaining budget for T_{\max} . Given this time budget, the maximum number of t_{trans} a packet could afford to wait in the queue due to synchronization in AB-PLeC(4,7) is therefore 1 (i.e., 2.704 ms). However, we observe that for AB-PLeC(4,7), some flows may experience synchronization delays twice, thus exceeding T_{\max} .

In contrast, the FIB algorithm presents the advantage of explicitly and precisely identifying the locations where T_{\max} is reached. Compared to AB-PLeC(4,7), we see that FIB may yield even longer paths than AB-PLeC, however, only for cases of limited synchronization events. For instance, Fig. 8 shows a six-hop path with only one synchronization event. In essence, FIB postpones the selection of a sink location as much as possible, leading to sink nodes closer to the PDC, i.e., utilizing PLC as much as possible. In contrast, AB-PLeC(3,7) constrains the number of hops to sinks to 3 (four hops to PDC), forcing packets that could still use PLC, to use the high capacity links of sink nodes. As a result, a higher number of sink nodes must be unnecessarily deployed, i.e., a 56% increase of sink nodes against the FIB algorithm.

¹⁶Our analysis in Section III-C focuses on the worst-case scenario, which is experienced by only one of the flows.

¹⁷Obviously, multiple data points coincide in each case.

VI. CONCLUSION AND FUTURE WORK

The paradigm shift toward ADNs introduces new challenges for the support of smart grid monitoring and control, including low latency communications for mission-critical applications. Though similar challenges have been addressed in the HV domain through the use of high-speed optical fiber links [9], the communication infrastructure landscape in the MV domain remains far from clear. In this paper, based on an extensive set of real MV grid dataset from a large European DNO, we show that the increased density of the MV domain renders the support of low latency applications particularly costly, as opposed to the much sparser HV domain. Proposing the limited use of the readily available PLC technologies, we identify and examine the tradeoff between performance and deployment costs. Drawing on the insights gained from a detailed study of the available topologies, we derive a set of practical guidelines for the design of hybrid, low latency communication infrastructures. Our investigation explicitly identifies, quantifies, and addresses the effect of PMU-device synchronization on queuing delays. Building on our empirical observations, we propose and evaluate three heuristic algorithms that identify the locations in a given grid that should be equipped with high capacity links, striking a balance between low latencies and deployment costs. Enforcing our algorithms on the available MV topologies and additionally engaging in extensive packet-level simulations, we show that the proposed algorithms can indeed satisfy the targeted low latencies while reducing the extend of high capacity link deployment by up to 80% in comparison to ubiquitous deployment of direct interconnection between PMUs and PDC. We believe our provisioning framework and associated algorithms can help power grid operators in a cost-effective transition toward a smart grid infrastructure.

To the best of our knowledge, this is the first study that identifies, quantifies, and further investigates the tradeoff between performance and deployment costs in the MV domain. Previous work in the area has been centered around the investigation of the topological properties of MV grids [15], [28], or the design of overlay communication networks [7], [24] focused on reliability and reconfiguration issues.

Our future work includes a detailed investigation of additional key performance indicators such as packet loss. To this end, the interplay between transport layer error control mechanisms (such as TCP) and their impact on the achieved latencies and throughput calls for a detailed investigation, in the context of synchrophasor applications. Moreover, the impact of link layer, forward error correction and error control mechanisms (e.g., automatic repeat request) needs to be taken into account in the design of future smart grid communication infrastructures in the MV domain.

ACKNOWLEDGMENT

The authors alone are responsible for the content of this paper.

REFERENCES

- [1] K. V. Katsaros, W. K. Chai, B. Vieira, and G. Pavlou, "Supporting smart electric vehicle charging with information-centric networking," in *Proc. 10th Int. Conf. Heterogeneous Netw. Quality, Rel., Security Robustness*, 2014, pp. 174–179.

- [2] Y. Cao, N. Wang, G. Kamel, and Y.-J. Kim, "An electric vehicle charging management scheme based on publish/subscribe communication framework," *IEEE Syst. J.*, DOI: 10.1109/JSYST.2015.2449893.
- [3] D. Laverty, D. Morrow, R. Best, and P. Crossley, "Differential ROCOF relay for loss-of-mains protection of renewable generation using phasor measurement over Internet Protocol," in *Proc. CIGRE/IEEE PES Joint Symp. Integr. Wide-Scale Renewable Resources Into Power Del. Syst.*, 2009, pp. 1–7.
- [4] P. Myrda and K. Koellner, "NASPInet—The internet for synchrophasors," in *Proc. 43rd Hawaii Int. Conf. Syst. Sci.*, 2010, pp. 1–6.
- [5] 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," in *Proc. Cigré SC6 Colloquium*, 2014, pp. 1–6.
- [6] A. Phadke and J. Thorp, "Communication needs for wide area measurement applications," in *Proc. 5th Int. Conf. Critical Infrastructure*, 2010, pp. 1–7.
- [7] W. K. Chai, "An information-centric communication infrastructure for real-time state estimation of active distribution networks," *IEEE Trans. Smart Grid*, vol. 6, no. 4, pp. 2134–2146, Jul. 2015.
- [8] R. Khan and J. Khan, "Wide area PMU communication over a WiMAX network in the smart grid," in *Proc. 3rd Int. Conf. IEEE SmartGridComm*, 2012, pp. 187–192.
- [9] M. Chenine, K. Zhu, and L. Nordstrom, "Survey on priorities and communication requirements for PMU-based applications in the Nordic Region," in *Proc. IEEE Bucharest PowerTech*, 2009, pp. 1–8.
- [10] P. Cheng, L. Wang, B. Zhen, and S. Wang, "Feasibility study of applying LTE to smart grid," in *Proc. IEEE Smart Grid Modeling Simul.*, 2011, pp. 108–113.
- [11] J. Brown and J. Khan, "Performance comparison of LTE FDD and TDD based Smart Grid communications networks for uplink biased traffic," in *Proc. IEEE SmartGridComm*, 2012, pp. 276–281.
- [12] S. Galli, A. Scaglione, and Z. Wang, "Power line communications and the smart grid," in *Proc. IEEE SmartGridComm*, 2010, pp. 303–308.
- [13] M. S. Daskin, "What you should know about location modeling," *Naval Res. Logistics*, vol. 55, no. 4, pp. 283–294, 2008.
- [14] R. L. Francis, T. J. Lowe, and H. D. Ratliff, "Distance constraints for tree network multifacility location problems," *Operations Res.*, vol. 26, no. 4, pp. 570–596, 1978.
- [15] G. A. Pagani and M. Aiello, "Towards decentralization: A topological investigation of the medium and low voltage grids," *IEEE Trans. Smart Grid*, vol. 2, no. 3, pp. 538–547, Sep. 2011.
- [16] K. Katsaros, B. Yang, W. K. Chai, and G. Pavlou, "Low latency communication infrastructure for synchrophasor applications in distribution networks," in *Proc. IEEE SmartGridComm*, 2014, pp. 392–397.
- [17] K. Christakou, J. LeBoudec, M. Paolone, and D.-C. Tomozei, "Efficient computation of sensitivity coefficients of node voltages and line currents in unbalanced radial electrical distribution networks," *IEEE Trans. Smart Grid*, vol. 4, no. 2, pp. 741–750, Jun. 2013.
- [18] *IEEE Guide for Design, Operation, and Integration of Distributed Resource Island Systems with Electric Power Systems*, IEEE Std 1547.4-2011, pp. 1–54, 2011.
- [19] Q. Li *et al.*, "An information-theoretic approach to PMU placement in electric power systems," *IEEE Trans. Smart Grid*, vol. 4, no. 1, pp. 446–456, Mar. 2013.
- [20] IEEE, *IEEE Standard for Synchrophasor Measurements for Power Systems*, IEEE Standard C37.118.1, 2011.
- [21] IEEE, *IEEE Standard for Synchrophasor Data Transfer for Power Systems*, IEEE Standard C37.118.2, 2011.
- [22] G. A. Pagani and M. Aiello, "The power grid as a complex network: a survey," *Physica A: Statist. Mech. Appl.*, vol. 392, no. 11, pp. 2688–2700, 2013.
- [23] R. Ramaswamy, N. Weng, and T. Wolf, "Characterizing network processing delay," in *Proc. Global Telecommunications Conf.*, 2004, vol. 3, pp. 1629–1634.
- [24] K. V. Katsaros *et al.*, "Information-centric networking for machine-to-machine data delivery—A case study in smart grid applications," *IEEE Netw. Mag.*, vol. 28, no. 3, pp. 58–64, May/June 2014.
- [25] I. Chlamtac, A. Farago, H. Zhang, and A. Fumagalli, "A deterministic approach to the end-to-end analysis of packet flows in connection-oriented networks," *IEEE/ACM Trans. Netw.*, vol. 6, no. 4, pp. 422–431, Aug. 1998.
- [26] J.-Y. Le Boudec and G. Hebuterne, "Comments on 'a deterministic approach to the end-to-end analysis of packet flows in connection oriented networks,'" *IEEE/ACM Trans. Netw.*, vol. 8, no. 1, pp. 121–124, Feb. 2000.
- [27] H. Zhang, "A note on deterministic end-to-end delay analysis in connection oriented networks," in *Proc. IEEE Int. Conf. Commun.*, 1999, vol. 2, pp. 1223–1227.
- [28] W. K. Chai, V. Kyritsis, K. Katsaros, and G. Pavlou, "Resilience of inter-dependent communication and power distribution networks against cascading failures," presented at the 15th IFIP Networking, Vienna, Austria, 2016.



Binxu Yang received the B.S. degree in telecommunications from Xidian University, Xian, China, the French Engineering degree from Telecom Bretagne, Plouzane, France, and the M.Res. degree (Distinction) from University College London (UCL), London, U.K., in 2010, 2013, and 2014, respectively. He is currently working toward the Ph.D. degree from the Department of Electronic and Electrical Engineering, UCL.

His current research interests include smart grid communications, mobile edge computing, and resource allocation in network function virtualization.



Konstantinos Vasileios Katsaros received the B.Sc. degree in informatics in 2003, the M.Sc. (Hons.) degree in 2005, and the Ph.D. degree in computer science from the Athens University of Economics and Business, Athens, Greece, in 2010.

He has worked in the areas of cloud networking, mobile grid computing, and multicast/broadcast service provision over cellular networks. His research interests include information-centric networking and smart grid communications. He is a currently a Research Associate at the Department of Electronic and Electrical Engineering, University College London, London, U.K.



Wei Koong Chai received the B.Eng. degree in electrical engineering from the Universiti Teknologi Malaysia, Johor Bahru, Malaysia, in 2000, and both the M.Sc. (Distinction) and the Ph.D. degrees from the University of Surrey, Surrey, U.K., in 2002 and 2008, respectively.

He is currently a Senior Research Associate at the Department of Electronic and Electrical Engineering, University College London, London, U.K. 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, quality of service, resource management (e.g., for satellite networks and wireless mesh networks), cross-layer design (specifically on interaction of protocols at different layers), traffic engineering, and network optimization.



George Pavlou received the Diploma degree in engineering from the National Technical University of Athens, Athens, Greece, and the M.Sc. and Ph.D. degrees in computer science from University College London, London, U.K.

He is a Professor of communication networks in the Department of Electronic and Electrical Engineering, University College London, where he coordinates research activities in networking and network management. His research interests include networking and network management, including aspects such as traffic engineering, quality of service management, autonomic networking, information-centric networking, grid networking, and software-defined networks. He has been instrumental in a number of European and U.K. research projects that produced significant results with real-world uptake and has contributed to standardization activities in ISO, ITU-T, and IETF. He has been on the editorial board of a number of key journals in these areas, he is the Chief Editor of the bi-annual IEEE Communications network and service management series and in 2011 he received the Daniel Stokesbury Award for distinguished technical contribution to the growth of the network management field.