

Low Latency Communication Infrastructure for Synchronphasor Applications in Distribution Networks

Konstantinos V. Katsaros, Binxu Yang, Wei Koong Chai and George Pavlou
University College London
E-mail: {k.katsaros, binxu.yang,13, w.chai, g.pavlou}@ucl.ac.uk

Abstract—With the introduction of new power sources, such as distributed renewable energy resources, and loads, such as electric vehicles, electrical distribution networks must accommodate new energy flow patterns in a considerably dynamic environment. This leads to the need for increasing the *observability* of the grid to enable a series of mission-critical applications such as voltage/congestion control and fault detection/location. The deployment of Phasor Measurement Units appears to be a promising approach, offering high precision grid monitoring. However, while the low delay requirements of such applications raise a significant challenge to the communication infrastructure, there is currently no clear vision on the exact communication technologies and network topologies that could support these requirements. In this paper, we address this challenge by taking a systematic approach on the design of low latency communication infrastructures. Based on a large set of real medium voltage grid topologies from a European distribution network, we first perform a detailed analysis of the communication requirements. Guided by this analysis, we then propose two algorithms, PLeC and BW-PLeC algorithms, for the design of low latency communication infrastructures that enhance the currently available power-line communication technology with newer high-speed communication links at strategic points in the grid to satisfy the delay requirements while reducing deployment costs.

I. INTRODUCTION

The electrical distribution network undergoes a major transformation. The introduction of distributed renewable energy resources (DRERs), such as wind/solar farms, results in new energy flow patterns, departing from the traditional model of hierarchical unidirectional flows from centralized energy production sites. At the same time, the expected increase of electric vehicles (EVs) [14] will correspondingly increase variations in energy consumption, this time following different spatiotemporal patterns. The integration of these dynamic active components at the distribution level progressively introduces higher dynamics and poses new challenges to the stability of the system especially on power quality, voltage regulation, protection and reliability of the grid.

In such dynamic environment, the role of a Wide Area Monitoring, Protection and Control (WAMPAC) system gains a significant role for distribution network operators (DNOs), with the latest advances foreseeing the deployment of Synchronized Measurement Technology (SMT) across the grid. Such deployment, spearheaded by the state-of-the-art Phasor Measurement Units (PMUs), has a profound impact on the communication network. Applications such as real-time state

estimation (RTSE) [1] and protection [2] rely on the timely delivery of synchronized monitoring information to adjust the power grid operations to the current conditions and further prevent / respond to failures. Typically, these applications are based on the reporting of high rate (e.g., 50-60Hz [3]) synchronphasor measurements by PMUs, under stringent end-to-end communication delay requirements e.g., 20ms [1], [4].

Supporting such synchronphasor applications with stringent delay requirements highly depends on the capabilities of the underlying communication infrastructure. While studies have already analysed the communication requirements in the High Voltage (HV) domain [13], where high capacity optical fibre may be typically available, the communication network landscape in the Medium Voltage (MV) domain is far from clear. Recent works have investigated the use of wireless technologies, e.g., WiMAX and LTE, reporting concerns about the impact of control plane and medium access control (MAC) layer delays on the overall feasibility of delay sensitive applications [4]–[6]. At the same time, and as we show in this paper, the low data rates supported by the readily available power-line communication (PLC) [7] infrastructure appear as a bottleneck to the timely delivery of monitoring traffic. Hence, we identify a tradeoff between the expected performance gains from the deployment of high bandwidth technologies and the deployments costs (and/or MAC delay penalties) associated with wide scale deployments.

In this paper, we address this tradeoff by proposing the design of hybrid communication infrastructure utilizing both PLC and other higher bandwidth technologies (e.g., optical fiber or wireless). To this end, based on a large set of real MV grid topologies from a European DNO, we first perform a detailed analysis of the communication requirements and the various contributing factors to the perceived end-to-end latency, paying particular attention to the impact of the selected technology and the resulting communication topology characteristics (section III). Gaining insight from this analysis, we propose two algorithms for the design of hybrid communication infrastructures (section IV). The first algorithm aims at reducing the processing, transmission and propagation delays by constraining the lengths of the data delivery paths, while minimizing the scale of the deployment of higher bandwidth technologies. Taking a step further, we propose a second heuristic algorithm with the additional constraint

of minimizing queuing delays. Our packet-level simulations on top of the available topologies demonstrate the ability of the proposed algorithms to constrain the end-to-end delays while achieving a reduction of the high bandwidth links of the topology by 80%.

II. MODELLING OF THE COMMUNICATION ENVIRONMENT

A. Delay-sensitive Synchrophasor Monitoring Applications

Our work is motivated by the challenge to support 3-phase Real-Time State Estimation (RTSE) application in the MV grid where large volumes of raw synchrophasor measurement data are collected by geographically-distributed PMUs, strategically deployed in the power grid infrastructure (i.e., on a secondary substation (S-SS) level) to ensure full grid *observability*. The data is GPS-synchronized and continuously streamed to phasor data concentrators (PDCs), which are typically located at the primary substation (P-SS). PDCs collect and deliver synchrophasor data to applications ranging from simple visualization tools to more sophisticated functions such as protection, state estimation, voltage control, etc.

The time budget allowance for the communication infrastructure in delivering synchrophasor measurement data to PDCs is dependent on the application. Typical refresh rates of state estimation processes are in the order of a few minutes. However, the higher system dynamics due to the incorporation of DRERs and EVs require the fine-grained estimation of system state within a few tens/hundreds of ms [1]. PMU reporting frequencies of 50 or 60 frames-per-second [3] facilitate this fine-grained view of the power grid. However, timely delivery of these measurements is a challenge for the underlying communication infrastructure. Though the overall time budget depends on several aspects [1], such as the time required to perform the state estimation, an indicative rule foresees a delay such that each measurement is received at the PDC before the next one is available by the PMU [4]. For a 50Hz frequency, this results in a 20ms latency constraint (denoted as T_L).

In terms of bandwidth requirements, and based on PMU data semantics [8], a realistic PMU message payload size is of 102 bytes¹. Further considering the overhead of UDP and IP headers, and a 32-byte SHA-256 message authentication code, the overall data rate for each RTSE PMU flow delivered to the link layer is 64.8Kbps, for a fixed 50Hz 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 on its own (e.g., [9]). In this paper, without loss of generality, we consider a scenario involving the deployment of a PMU at approximately every two S-SSes along a feeder (see Fig. 1)².

¹Considering the case of PHNMR=6, ANNMR=6 and DGNMR=2, with 32-bit floating-point accuracy [8].

²We use this placement model only as a rule of thumb; our proposed solutions take no assumption on the exact placement of data sources.

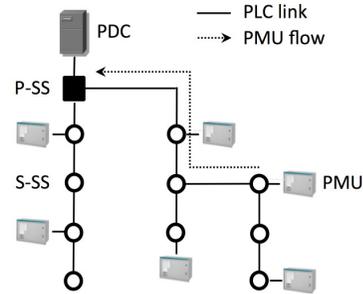


Fig. 1. Illustration of a MV grid with PMUs deployed across the topology. In the case of PLC, PMU flows follow the grid topology.

TABLE I
TOPOLOGY PROPERTIES OF SAMPLE MV GRID IN THE NETHERLANDS.

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

B. Communication Topology Based on Real MV Grids

We model the communication topology based on a set of MV power grid topologies operated by a DNO in the Netherlands. The topologies have tree-like structures rooted at primary substations. Each tree branch emanating from the P-SS corresponds to a distinct feeder. We represent the distribution grid as a tree graph, T , with N nodes $u_i, i \in [0, \dots, N-1]$ where node u_0 represents the root i.e., the P-SS. We define $d(u_i, u_j)$ as the communication cost between nodes u_i and u_j , measured as the length of the shortest path (in hops) between them. Table I summarizes the basic topological characteristics of the considered MV grids (aggregated).

Based on the real grid data, we derive a baseline communication network model enabled by PLC. In this model, DNOs make use of the existing power-line cables as the transmission medium. In the context of the considered set of applications, PMU flows reach the PDC by traversing their uphill links towards the root of the tree topology, as shown in Fig. 1. In other words, for the PLC case, the communication topology coincides with the power grid topology. The use of existing infrastructure i.e., power cables, offers a significant advantage to DNOs as it results in minimum investment for the deployment of a communication infrastructure. However, as we show in the following section, the low bandwidth capabilities of PLC (typically in the range of a few hundred Kbps [7]) poses significantly barriers to the support of PMU-based applications on MV networks.

III. END-TO-END DELAY ANALYSIS

Following the baseline PLC-based communication model, we first show that supporting the considered synchrophasor applications on top of existing (or readily available) communication infrastructures is not feasible. This investigation enables us to identify the exact bottlenecks in the baseline

model as input to the design of a low latency communication infrastructure suitable for the PMU applications.

To this end, we have developed a detailed packet-level simulation environment based on OMNeT++ [10]. We translate, as input to the simulation, the available power grid data to their equivalent PLC-based communications topology. We consider two PLC bandwidth values: 100Kbps and 500Kbps [7]. Following to the PMU deployment scenario described in section II-B, we simulate the operation of 795 PMUs deployed across the MV grid. Fig. 2(a) shows the cumulative distribution function (CDF) of the end-to-end delay observed at the PDC for a simulation period of 10 minutes. The vast majority of PMU messages is delivered with a considerably high delay that exceeds the threshold of 20ms. Obviously, the considered set of applications cannot be supported by purely using the existing PLC technology.

To get a better understanding of this result and more importantly gain input on how to design a suitable communication infrastructure, we next engage in the decomposition of the observed delays, investigating the impact of the various delay contributing factors, namely:

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

The properties of the communication topology can have significant impact on the end-to-end delays. This is because the topology largely determines the actual traffic volume traversing each link, as well as the total length of the delivery paths, with an obvious impact on the accumulated t_{queue} , t_{trans} , t_{proc} and t_{prop} . At the same time, the selected technology affects the perceived delays in terms of t_{trans} and t_{queue} .

We illustrate the impact of these characteristics on the feasibility of RTSE in Figures 2(b) and 2(c). Fig. 2(b) shows the CDF of the processing delays accumulated by data packets across all PMU-to-PDC PLC paths in the available MV topology, 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, from a few micro-seconds to even milliseconds per packet [11]. If we further consider recent overlay approaches [12], these delays may further increase due to the transition of packets from the kernel space to the user space.

We notice that, subject to t_{proc} value, the overall delay penalty may get close or even exceed T_L . In essence, these measurements directly reflect the impact of the delivery path lengths on the perceived performance, which can be considerably high when the t_{proc} value reaches / exceeds the value of approximately 1ms. This is an important guideline on the

design of a low-latency communication infrastructure with respect to the provisioning of computational resources at each forwarding node. At the same time, it is important to note that lengthy delivery paths also increase the total time spent by each packet in transmission (t_{trans}); each packet needs to be received in its entirety before it gets transmitted again on the next hop.

Besides the effect of t_{proc} , we further shed light on the impact of t_{queue} . Fig. 2(c) shows the CDF of the total PMU traffic volume aggregated at each PLC link towards the PDC. This corresponds to the CDF of the bandwidth requirements on the uplink of each node. 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 100Kbps ($\approx 10\%$ for 500Kbps links). This limitation results in high t_{queue} ³ largely contributing to the end-to-end latencies observed in Fig. III. It must be stressed that this result does not only owe to the low bandwidth capabilities of PLC, but is further attributed to the tree-like structure of the communication topology i.e., an increasing number of PMU flows is aggregated as we get closer to the P-SS.

IV. DESIGNING A LOW LATENCY COMMUNICATION INFRASTRUCTURE

The aforementioned implications contribute to a set of tradeoffs with respect to the design of a low latency smart grid communication infrastructure. PLC technology, though readily available, appears not able to support the considered low latency applications. This urges for alternative solutions such as the use of modern wireless or high-speed wired technologies such as optical fiber. However, the deployment of such technologies incurs a non-negligible capital expenditure (CAPEX), necessitating the careful dimensioning of the network, i.e., the overall deployment costs should be kept low while ensuring the ability of the network to support the application layer requirements. Recent studies have also shown that the adoption of wireless technologies may lead to an increase of medium access delays due to the shared nature of the wireless medium [4]–[6]. The contention for access to the wireless medium increases with the number of wireless transmitting devices, leading to increased waiting times for a permission to transmit data. In this respect, the number of wireless transmitting devices should also be kept to a minimum for performance reasons as well.

Based on these observations, 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, while ensuring conformance to application performance requirements with enhanced transmission technologies. To this end, a limited

³Due to length limitations, we only show the bandwidth requirements as it provides an assessment of the communication requirements, which are orthogonal to the selected technology and topology.

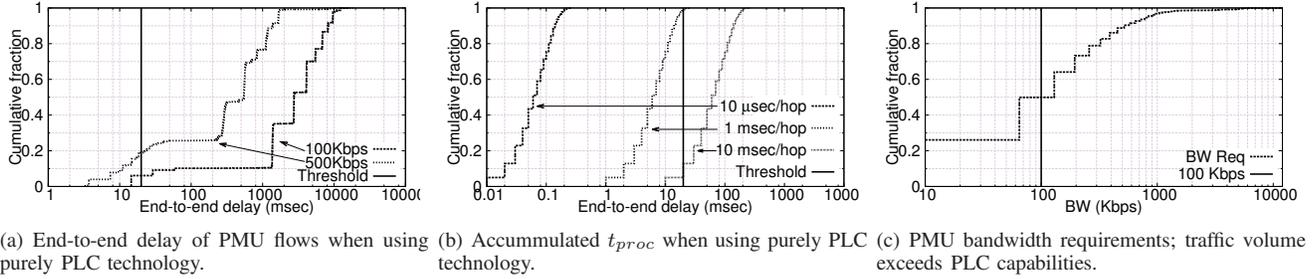


Fig. 2. Impact of topology on PMU application performance.

set of *sink* communication nodes at S-SSes is selected to be connected to the PDC via a high capacity direct link (e.g., optical fibre). These nodes act as traffic aggregation points for the remainder of the communication nodes, which employ PLC to deliver their traffic there. The objective then is two-fold. First, we target the selection of the appropriate sink locations, such that the resulting topology can adhere to the design guidelines derived from the end-to-end delay analysis (section III). Second, we try to minimize the number of aggregation points, so as to conform to the aforementioned cost and performance requirements. To this end we present the following two algorithms:

- The *path length constraint (PLeC) algorithm* minimizes the number sink nodes while constraining the length of data delivery paths, so that the accumulated t_{proc} , t_{prop} and t_{trans} are also capped (see section III).
- The *bandwidth and path length constraint (BW-PLeC) algorithm* is heuristic and constrains both path lengths and the number of PMU flows on each PLC link; therefore, explicitly targeting the reduction of t_{queue} .

A. Path Length Constraint (PLeC) Algorithm

For the PLeC algorithm, we follow the distance constraint formulation of the p -center facility location problem [15]. In this context, we define $S = s_1, s_2, \dots, s_p$ as the set of sink nodes, with $1 \leq p \leq N$. Further, let $D(S, u_i) = \min\{d(s, u_i) : s \in S\}$, the distance between each node u_i and its nearest sink node. Our objective is to find the minimum set S such that $D(S, u_i) \leq d_{max}, \forall i \in [0, \dots, N-1]$. We solve this problem via the *sequential location procedure* proposed in [15]. The d_{max} value can be derived as: $d_{max} = \left\lceil \frac{T_L - t_{proc}}{t_{trans} + t_{proc}} \right\rceil^4$. Our algorithm (see Algorithm 1) takes as input the tree topology, T and the distance constraint, d_{max} , and outputs the set of selected sink nodes, S , along with set M (see next). For all nodes $u_i, i \in [0, \dots, N-1]$, we define a distance value $a_i, i \in [0, \dots, N-1]$ and a set M_i , which contains the nodes that can use node u_i as their sink node, under the d_{max} constraint. Sets M_i is denoted as M .

The algorithm starts by randomly selecting a leaf node, u_l from T , along with its parent node u_p . Traversing the

Algorithm 1 PLeC algorithm

Input: T, d_{max}
Output: S

- 1: $S \leftarrow \emptyset$
- 2: **for all** i **do**
- 3: $a_i \leftarrow d_{max}$
- 4: **end for**
- 5: **while** $T \neq \emptyset$ **do**
- 6: $u_l \leftarrow T.getRandomLeafNode()$
- 7: **if** $u_l \neq u_0$ **then**
- 8: $u_p \leftarrow u_l.getParentNode()$
- 9: $a_p \leftarrow \min(a_p, a_l - 1)$
- 10: $T.removeNode(u_l)$
- 11: **if** $a_p = 0$ **then**
- 12: $ADD_SINK(T, S, u_p)$
- 13: **end if**
- 14: **else**
- 15: $ADD_SINK(T, S, u_l)$
- 16: **end if**
- 17: **end while**
- 18: **return** S, M
- 19:
- 20: **function** $ADD_SINK(T, S, u_s)$
- 21: $S \leftarrow S \cup u_s$
- 22: $R \leftarrow \emptyset$
- 23: **for all** $u_i \in T$ **do**
- 24: **if** $d(u_i, u_s) \leq a_i$ **then**
- 25: $M_s \leftarrow M_s \cup u_i$
- 26: $R \leftarrow R \cup u_i$
- 27: **end if**
- 28: **end for**
- 29: $T \leftarrow T \setminus R$
- 30: **end function**

tree hierarchy, the algorithm updates the distance value a_p of nodes u_p as shown 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 u_p is added to the sink node set (function $ADD_SINK(T, S, u_s)$, line 20). In this step, all nodes u_i whose minimum hop distance to the new sink node u_s is below their a_i value are added to the M_s set. All nodes assigned to the new sink are also removed from the tree⁵.

The outcome of the algorithm consists of the sets M_i for each selected sink node u_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 T , sets M_i may not

⁴ t_{prop} (average $\leq 3\mu$ s in the considered topologies) and t_{trans} on the sink-to-PDC link ($\leq 2\mu$ s for a 10Gbps optical fiber link) are considered negligible. However, we account the t_{proc} for the sink-to-PDC hop.

⁵This process may result in a forest. Structure T is used for all trees, and $getRandomLeafNode()$ (line 6) returns a leaf node randomly selected from any of the trees.

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 u_i . The members of each L_i set are ordered in increasing hop distance to u_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 u_s with the minimum M_i size so as to not overload other sinks which can possibly serve more nodes.

B. Bandwidth and Path Length Constraint (BW-PLeC) Algorithm

The BW-PLeC algorithm focuses on finding 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 topology T , the maximum number of PMU flows that can be accommodated by a PLC link, bw_{max} , the maximum number of PMU flows that can be accommodated by a high bandwidth link connecting a sink node to the PDC, bw'_{max} , and d_{max} . The value of bw_{max} is determined by the ratio of the considered bandwidth on each link (BW) and the data rate of the PMU flows (DR): $bw_{max} = \lfloor \frac{BW}{DR} \rfloor$. bw'_{max} is set in a similar way, and it is normally expected to be considerably higher than bw_{max} . In addition to a_i , for each node $u_i, i \in [0, \dots, N-1]$, we define b_i as the current number of flows to be forwarded by the node. All b_i are initialized to 0, unless a PMU is attached to the corresponding node (line 4). The tree topology is traversed from the leafs towards 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 bw'_{max} value (line 15). Visited nodes and sinks are removed from T and the algorithm terminates when all nodes have been removed. Then, each node in the tree can forward its traffic to each closest ancestor sink node.

C. Simulation results

By applying the proposed algorithms on the available MV power grid topologies, we derive a series of alternative communication network topologies under specific constraints. We assess the performance of the derived topologies with packet-level simulations (see section III). Due to length limitations, we focus in the following, on the case of PLC bandwidth=500Kbps, but similar conclusions applies for the case of 100Kbps. We consider each sink node to be connected to the P-SS with a 10Gbps optical fibre link and $t_{proc} = 1$ ms. Based on the above, we then get $d_{max} = 5$, $bw_{max} = 7$ and $bw'_{max} = 147$. We then derive the following topologies: (1) PLeC with d_{max} values 5, but also values 4, 3 and 2 (denoted by PLeC(d_{max})), which yields topologies with a total of 160,

Algorithm 2 BW-PLeC algorithm

Input: $T, d_{max}, bw_{max}, bw'_{max}$
Output: S

- 1: $S \leftarrow \emptyset$
- 2: **for all** i **do**
- 3: $a_i \leftarrow d_{max}$
- 4: $b_i \leftarrow (u_i.hasPMU())?1 : 0$
- 5: **end for**
- 6: **while** $T \neq \emptyset$ **do**
- 7: $u_l \leftarrow T.getRandomLeafNode()$
- 8: $u_p \leftarrow u_l.getParentNode()$
- 9: $x \leftarrow \min(a_p, a_l - 1)$
- 10: $y \leftarrow b_p + b_l$
- 11: **if** $u_l.markedAsSink()$ **then**
- 12: $S \leftarrow S \cup u_l$
- 13: $T.removeNode(u_l)$
- 14: **else**
- 15: **if** $y > bw'_{max}$ **and** $u_p \neq u_0$ **then**
- 16: $S \leftarrow S \cup u_p$
- 17: $T.removeNode(u_p)$
- 18: **else**
- 19: $a_p \leftarrow x$
- 20: $b_p \leftarrow y$
- 21: $T.removeNode(u_l)$
- 22: **if** $b_p > bw_{max}$ **or** $a_p \leq 0$ **then**
- 23: $u_l.markAsSink()$
- 24: **end if**
- 25: **end if**
- 26: **end if**
- 27: **end while**
- 28: **return** S

188, 236 and 309 sink nodes respectively, (2) BW-PLeC with d_{max} values 5, 4 and 3, $bw_{max} = 7$ and $bw'_{max} = 147$ (denoted as BW-PLeC(d_{max}, bw_{max})), which yields topologies with a total of 147, 194 and 256 sink nodes respectively.

Fig. 3 and 4 show the CDF of the end-to-end delay for all considered topologies. We see that BW-PLeC(3,7) clearly satisfies the delay constraint with 256 high bandwidth links / sink nodes. PLeC(2) achieves the same goal with 309 sink nodes. Despite PLeC's optimality, the d_{max} constraint in this case is more stringent, also resulting in an overall better performance i.e., median and maximum delay values of 4.7ms and 14.8ms against 7.4ms and 20ms of BW-PLeC(3,7) respectively. BW-PLeC(4,7) and PLeC(3) closely follow, only slightly exceeding T_L for $< 1\%$ of the measured packets, i.e., by 2.9ms and 2ms respectively. Interestingly, we see that PLeC(5), PLeC(4) and BW-PLeC(5,7) fail to satisfy the delay requirement, even though the distance and bandwidth constraints are enforced. In the case of the PLeC algorithm, this is due to the fact that the algorithm does not take into account the bandwidth constraint; hence, it is possible that certain links are overloaded. However, this does not hold for BW-PLeC. A closer inspection has shown that, for both algorithms, delay values exceeding T_L are a consequence of the synchronisation of the PMU flows which results in forwarding nodes receiving packets from equidistant PMUs at approximately the same time. This results in a delay penalty linearly related to t_{trans} , i.e., since only one packet can be transmitted at a time, packets arriving relatively later must wait in the queue until the preceding ones get transmitted.

From the above, it becomes apparent that for the considered

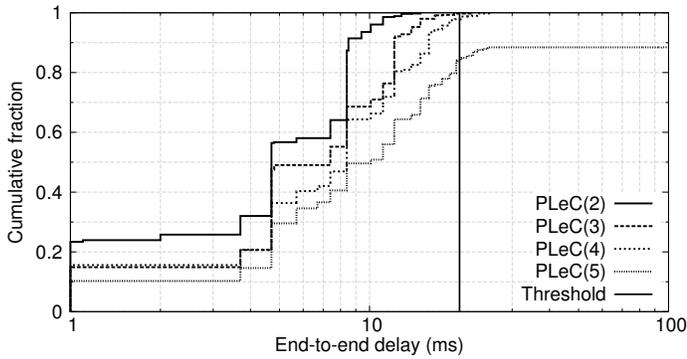


Fig. 3. CDF of end-to-end delay for 500Kbps PLC links: PLeC algorithm.

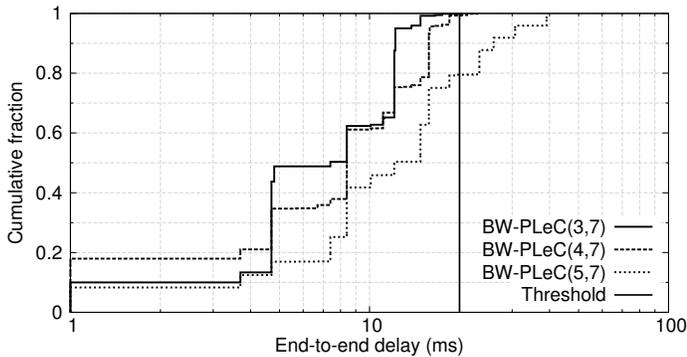


Fig. 4. CDF of end-to-end delay for 500Kbps PLC links: BW-PLeC algorithm.

real MV topologies, the BW-PLeC algorithm provides a low latency communication network topology with a limited number of high bandwidth links / sink locations (i.e., 256). This presents a far more attractive approach compared to scenarios of ubiquitous optical fibre deployment where either all S-SSes are directly connected to the P-SS, resulting in 1323 optical fibre links, or the optical fibre infrastructure follows the grid topology resulting in 1426 links (see Table I). In both cases, the reduction of optical fibre links would be in the order of 80%. However, still, the provided solution is not the optimal one. This can be realized by considering that sink node's PLC uplinks (connecting them to their ancestors in T) are not utilized by downstream PMU flows, as in the PLeC algorithm. Finally, we must note that the optimality of the PLeC algorithm presents a significant advantage for cases where the bandwidth constraint can be relaxed. Such cases include the use of lower PMU reporting rates (e.g., 10 or 25 Hz [3]), smaller PMU data frame [8] and/or other low bandwidth but delay stringent applications, such as system protection functions including emergency reporting and responding (e.g., circuit breaker control operations) [16].

V. CONCLUSIONS

In this paper, we focus on the design problem of low latency communication infrastructures for smart grids. In particular, we study synchrophasor applications utilizing PMUs in MV grids which, due to the strict latency requirements, represent the set of the most demanding applications for the communication infrastructure. We first show the inadequacy of the

readily available PLC technologies for such applications and identify the critical communication bottlenecks. Based on the insights from this feasibility study, we propose and evaluate two algorithms (i.e., PLeC and BW-PLeC algorithms) for the design of hybrid topologies that deploys limited number of aggregation/sink nodes at S-SSes with high bandwidth links directly to P-SS in conjunction with the exploitation of the low bandwidth PLC technologies. Our investigation draws conclusions from a large set of real MV power grid topologies and shows that the proposed algorithms can substantially reduce the number of high bandwidth links required ($\approx 80\%$) to support the stringent delay requirements of PMU-based applications. From our investigation, we also discovered a unique contributing factor to the end-to-end delay that is specifically due to the precise synchronisation of PMU flows. We will account this aspect in our subsequent work, further considering its relation to the PMU placement problem.

ACKNOWLEDGMENTS

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

REFERENCES

- [1] 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. of the 2013 Cigré SC6 Colloquium*, 2014.
- [2] D. Laverty *et al.*, "Differential ROCOF relay for Loss-of-Mains protection of Renewable Generation using phasor measurement over Internet Protocol," in *Integration of Wide-Scale Renewable Resources Into the Power Delivery System, 2009 CIGRE/IEEE PES Joint Symposium*, 2009.
- [3] IEEE, "IEEE Standard for Synchrophasor Measurements for Power Systems," *IEEE Standard C37.118.1*, 2011.
- [4] R. Khan *et al.*, "Wide area PMU communication over a WiMAX network in the smart grid," in *Proc. of IEEE SmartGridComm*, 2012.
- [5] P. Cheng *et al.*, "Feasibility study of applying LTE to Smart Grid," in *Proc. of IEEE SGMS*, 2011.
- [6] J. Brown *et al.*, "Performance comparison of LTE FDD and TDD based Smart Grid communications networks for uplink biased traffic," in *Proc. of IEEE SmartGridComm*, 2012.
- [7] S. Galli *et al.*, "Power Line Communications and the Smart Grid," in *Proc. of IEEE SmartGridComm*, 2010.
- [8] IEEE, "IEEE Standard for Synchrophasor Data Transfer for Power Systems," *IEEE Standard C37.118.2-2011*, 2011.
- [9] Q. Li *et al.*, "An Information-Theoretic Approach to PMU Placement in Electric Power Systems," *IEEE Transactions on Smart Grid*, vol. 4, no. 1, pp. 446–456, 2013.
- [10] A. Varga, *OMNeT++ Simulator Home Page*, <http://www.omnetpp.org>.
- [11] R. Ramaswamy *et al.*, "Characterizing network processing delay," in *Proc. of IEEE GLOBECOM*, 2004.
- [12] K. V. Katsaros *et al.*, "Information-centric Networking for Machine-to-Machine Data Delivery – A Case Study in Smart Grid Applications," *IEEE Network Magazine*, vol. 28, no. 3, pp.58–64, 2014.
- [13] R. Hasan *et al.*, "Analyzing NASPInet data flows," in *Proc. of IEEE/PES PSCE*, 2009.
- [14] K. V. Katsaros *et al.*, "Supporting Smart Electric Vehicle Charging with Information-Centric Networking," in *Proc. of Q-ICN*, 2014.
- [15] R. L. Francis *et al.*, "Distance Constraints for Tree Network Multifacility Location Problems," *Operations Research*, vol. 26, no. 4, pp. 570–596, 1978.
- [16] W. Wang *et al.*, "A survey on the communication architectures in smart grid," *Computer Networks*, vol. 55, no. 15, pp. 3604–3629, 2011.