Date of publication xxxx 00, 0000, date of current version xxxx 00, 0000.

Digital Object Identifier 10.1109/ACCESS.2019.DOI

Measurement and Analysis of PMU Reporting Latency for Smart Grid Protection and Control Applications

STEVEN M. BLAIR[1], (Member, IEEE), MAZHERUDDIN H. SYED[1], (MEMBER, IEEE), ANDREW J. ROSCOE[2], (Senior Member, IEEE), GRAEME M. BURT[1], (Member, IEEE), and JEAN-PIERRE BRAUN[3]

[1]Institute for Energy and Environment, Department of Electronic and Electrical Engineering, University of Strathclyde, Glasgow, G1 1RD, UK (e-mail: steven.m.blair@strath.ac.uk)
[2]Siemens Gamesa Renewable Energy, 389 Argyle Street, Glasgow, G2 8LR
[3]Federal Institute of Metrology (METAS), Lindenweg 50, CH-3003 Bern-Wabern, Switzerland

Corresponding author: Steven M. Blair (e-mail: steven.m.blair@strath.ac.uk).

This work received funding through the European Metrology Research Programme (EMRP), Project ENG52 "Measurement tools for Smart Grid stability and quality". The EMRP is jointly funded by the EMRP participating countries within EURAMET and the European Union. This project has also received funding from the EMPIR programme "FutureGrid II" (project number: 17IND06) co-financed by the Participating States and from the European Union's Horizon 2020 research and innovation programme. This paper has also been supported by the European Commission, under the FP7 project ELECTRA (grant no: 609687) and Horizon 2020 project ERIGrid (grant no: 654113). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the European Commission.

ABSTRACT Emerging power system protection and control applications require faster-responding measurements and more accurate knowledge of the actual latency of the measurement and communications systems. A new method for accurately determining the reporting latency of a phasor measurement unit (PMU) has been developed and demonstrated. This method operates in real-time, works passively for any existing PMU without requiring changes to the PMU hardware or software, and is very accurate—providing a measurement uncertainty of <500 ns in many cases, significantly surpassing the 0.002 s accuracy requirement in the most recent IEEE Synchrophasor standard. Only low-cost hardware and open source software are required. It is particularly important to understand end-to-end system latency, including the impact of local and wide-area communications, rather than just the latency of the PMU device; the proposed method also supports such practical measurements. It is therefore shown how this advance can be used to enable efficient, but realistic, cross-domain power system simulation studies which incorporate measurement and communications delays. These capabilities address complexity and uncertainty in the design and operation of future PMU-based protection and control functions for new smart grid services.

INDEX TERMS Communications, IEC 61850, IEEE 1588, IEEE C37.118, phasor measurement units (PMUs), Sampled Values, time synchronization.

I. INTRODUCTION

FAST-ACTING response to power system disturbances is becoming critical to ensuring system stability [1]. Wide-area phasor measurement unit (PMU) monitoring schemes are being utilized to enable new system functions, such as fast-acting frequency control [2], high-fidelity state estimation [3], wide-area protection [4], and decentralized control paradigms [5], [6]. In these applications, it is often critical that measurement latency is minimized [7]–[9], and it is therefore important that latency can be correctly characterized [10]. Furthermore, the North American SynchoPhasor Initiative (NASPI) has recommended avoiding the use of PMU data for system-critical operations unless timing accuracy and resiliency have been fully validated [11], and the lack of tools to perform this validation presents a significant barrier to exploiting PMU-based solutions. Latency must also be faithfully represented in simulation studies so that these novel systems can be comprehensively tested and derisked.

Although existing PMU calibrators have been designed to very accurately and automatically characterize the signal

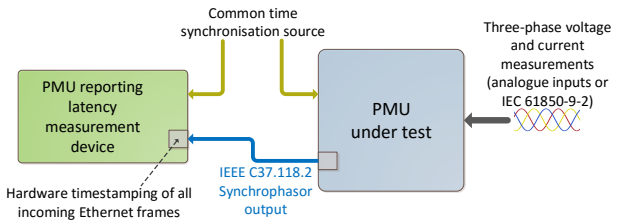

FIGURE 1. Conceptual overview of PMU reporting latency measurement method.

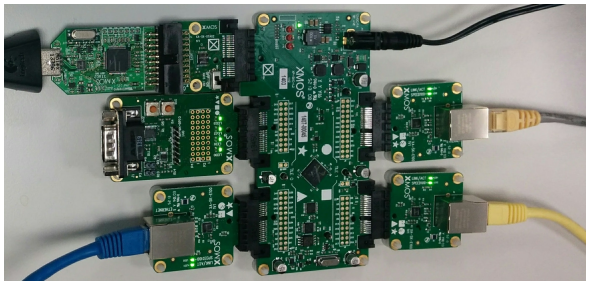

FIGURE 2. xCORE development board with three Ethernet interfaces

processing performance of a PMU under test [12], measuring PMU reporting latency automatically is complex due to its real-time nature and requirement for accurate time-stamping of data packets, and is typically not accommodated by PMU calibrators [13]. It is possible for a PMU to measure reporting latency internally, but this functionality may not be implemented or available to the user. Previously reported techniques [13] provide relatively low accuracy (of approximately 600-900 µs) and require expensive testing hardware. This paper describes a new method to accurately measure PMU reporting latency, without requiring specialized or expensive equipment. This method is convenient to apply retrospectively to any PMU, and the implementation provided by the authors [14] performs the latency measurement automatically.

Furthermore, smart grid applications inherently involve cross-domain challenges to integrate measurement technologies, communications, and real-time control systems. For time-sensitive applications, it is important to be able to 1) characterize the end-to-end latency of actual PMU installations, including the wide-area communications, and 2) validate complex PMU-based control and protection systems through simulation (including the use of real-time simulation). A major contribution of the work presented in this paper is to demonstrate the value of the novel PMU reporting latency measurement method in achieving these two additional objectives.

II. METHOD AND OPEN SOURCE IMPLEMENTATION

A. BACKGROUND AND METHOD OVERVIEW

The IEEE 1588 standard [15], also known as the Precision Time Protocol (PTP), enables high-quality time synchronization over Ethernet networks. An important aspect of the standard is the use of hardware timestamping, which involves recording the exact time when the start of a PTP Ethernet frame (i.e. the first bit following the Ethernet preamble) enters or leaves a node in the network. In some devices, the timestamping occurs in the physical layer (PHY) of the Ethernet interface for the highest accuracy. This convention is also conveniently aligned with the requirements for PMU reporting latency defined in the IEEE C37.118.1a Synchrophasor standard [16]: the reporting latency is the time difference between the first bit of a PMU report message and the timestamp contained in the report. This means that a device with an Ethernet interface which supports PTP hardware timestamping, and the accompanying software stack, can be used to very precisely measure reporting latency, according to this definition [17]. Fig. 1 illustrates this method. The measurement device and the PMU under test are both synchronized to a common time source, and the measurement device receives Synchrophasor data (encoded in IEEE C37.118.2 format) from the PMU.

It should be noted that this method measures the time for the first bit to be received, rather than the time of transmission as specified in the standard; however, as noted in [13], this difference—comprising the propagation delay of the physical layer medium—is negligible. Therefore, the measurement device should use a dedicated Ethernet network interface to directly connect to the PMU under test (rather than via an Ethernet switch) for the highest accuracy, according the Synchrophasor standard requirements. However, for convenience an ordinary Ethernet switch could be used while still remaining well within the measurement accuracy requirement of 0.002 s defined in [16] (as is demonstrated in Section III-B).

B. REAL-TIME IMPLEMENTATION

The XMOS xCORE platform has been used to implement the PMU reporting latency measurement method, according to Fig. 1. This hardware platform is well-suited to real-time, deterministic applications involving Ethernet [18], has been previously demonstrated for use as real-time Ethernet delay emulation for time-critical protection applications [19], [20] and IEC 61850-9-2 Sampled Value encoding performance analysis [7]. As shown in Fig. 2, the xCORE supports multiple Ethernet network interfaces, with IEEE 1588 hardware timestamping in the PHY, and a development board is available at a relatively low cost (~$150). There are existing open source software libraries for the xCORE, including a PTP communications stack [21]. These libraries have been extended by the authors to perform the PMU latency measurement calculation; this additional code is also open source and available at [14].

Fig. 3a illustrates the software configuration of the xCORE device, and Fig. 3b provides more detail for the PMU reporting latency calculation method. The open source implementation supports the IEEE C37.118.2 protocol over User Datagram Protocol (UDP), which is better suited for real-time applications than using Transmission Control Protocol

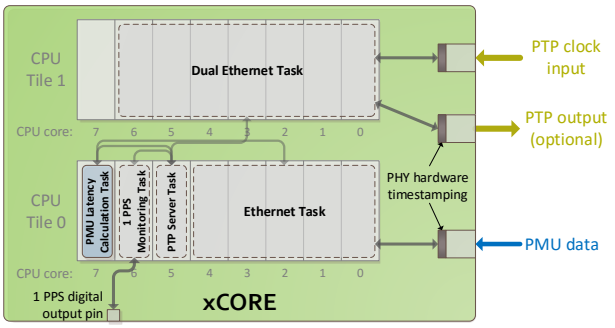

(a) Software task CPU core assignment and hardware mapping

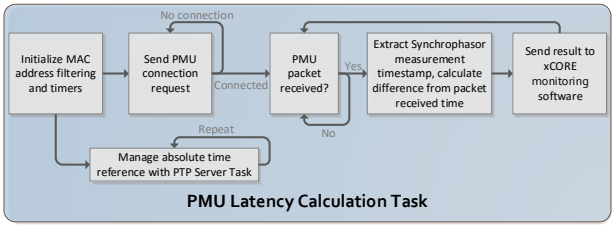

(b) Detailed PMU Latency Calculation Task process

FIGURE 3. Overview of xCORE software configuration

(TCP). The process for testing PMU reporting latency is fully automated. The xCORE initiates the process by sending the appropriate command to the PMU under test to start transmission of data (unless the multicast protocol is used, where this handshake is not required). Once the PMU starts transmitting reports, the xCORE is able to receive these Ethernet frames (with hardware timestamping of the exact arrival time) and extract the PMU report measurement timestamp value from the data. Through a connection to the PTP Server task (see Fig. 3a) to provide a mapping to absolute time, this timestamp can be compared with the frame arrival timestamp from the Ethernet interface to calculate the PMU reporting latency for each packet.

C. CONFIGURATION OPTIONS

Fig. 4 illustrates different options for using the developed platform for measuring PMU reporting latency. All options are supported by the open source implementation. The configuration chosen will depend on the capabilities of the PMU under test (i.e. its time synchronization interface) and other available hardware. In summary, Options 1 and 2 are very similar, except for the method used to synchronize the PMU. Option 3 is useful if an absolute time reference source is not available, and Option 4 can be used to improve timing accuracy—compared to Option 1—if an Ethernet switch supporting transparent clock functionality is not available. The timing accuracy of each configuration option is analyzed in Section III-B.

The configuration options presented in Fig. 4 focus on measuring PMU reporting latency as defined in the Syn-

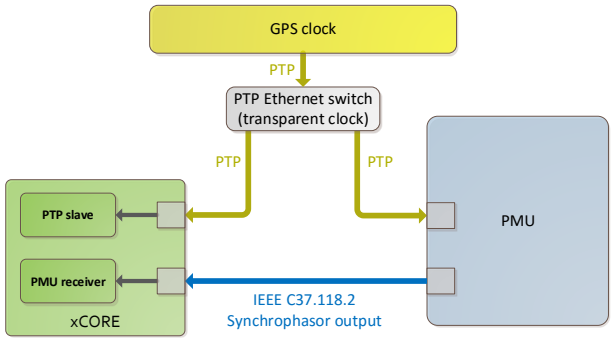

(a) Option 1: xCORE and PMU synchronized on same PTP network

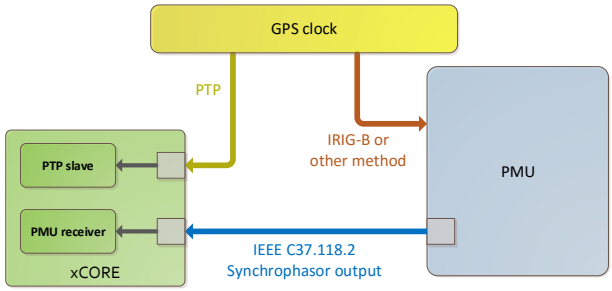

(b) Option 2: xCORE and PMU synchronized by different methods

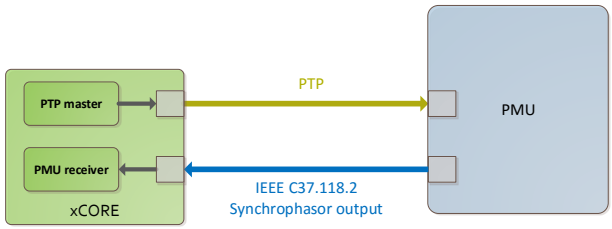

(c) Option 3: xCORE synchronizes PMU to xCORE local time

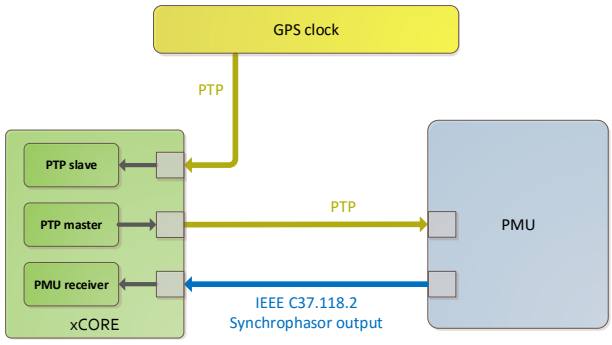

(d) Option 4: xCORE distributes time to PMU using PTP (i.e. xCORE acts like a PTP boundary clock)

FIGURE 4. Supported configuration options for measuring PMU reporting latency

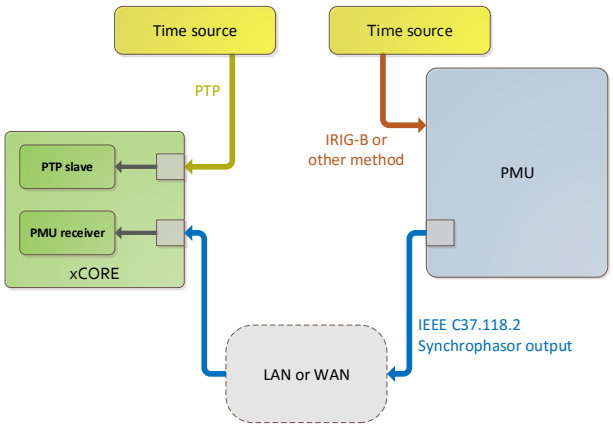

FIGURE 5. Configuration Option 2 for measurement of PMU reporting latency through a LAN or WAN

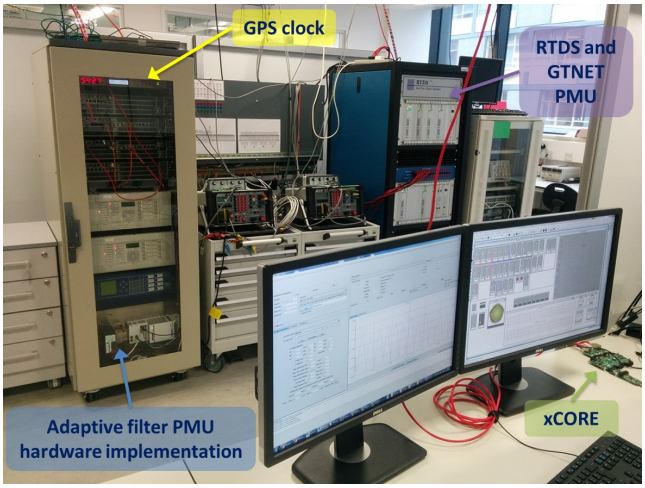

FIGURE 6. Laboratory configuration for real-time validation

chrophasor standards. However, it is of critical importance for many smart grid applications to be able to characterize the end-to-end latency of the full system, i.e. including the communications due to the local area network (LAN) and wide-area network (WAN), if applicable. The proposed platform also supports this, as illustrated in Fig. 5 for configuration Option 2, with the only requirement being that both sites (i.e. the PMU location and the reporting latency measurement location) have access to a common time source.

III. REAL-TIME VALIDATION OF METHOD

A. OVERVIEW

This section proves that the proposed measurement platform meets the required accuracy, and is flexible to be applied in different practical situations. The platform has been validated and demonstrated using an Arbiter 1201C GPS clock (with 100 ns rated accuracy), two PMU implementations, and a Real Time Digital Simulator (RTDS) to supply controllable signals to the PMUs under test. The laboratory configuration is shown in Fig. 6. Within the RTDS, time synchronization is managed by a "GTSYNC" card, which supports IRIG-B and PTP inputs, and can be used to distribute time to other devices (e.g. using IRIG-B or 1 PPS signals). The RTDS supplies analogue waveforms (representing signals from voltage and current transformers) to the PMU inputs, with new values being calculated every simulation time-step (50 µs). The RTDS also has the ability to emulate a PMU in real-time, including the IEEE C37.118.2 data output, using the "GTNET" hardware card. A GTNET card can also be configured to digitally output the voltage and current waveform signals using the IEC 61850-9-2 Sampled Value (SV) protocol.

Time synchronization accuracy, for various configurations, is established in Section III-B, and the reporting latency measurements for two PMU implementations are presented in Section III-C.

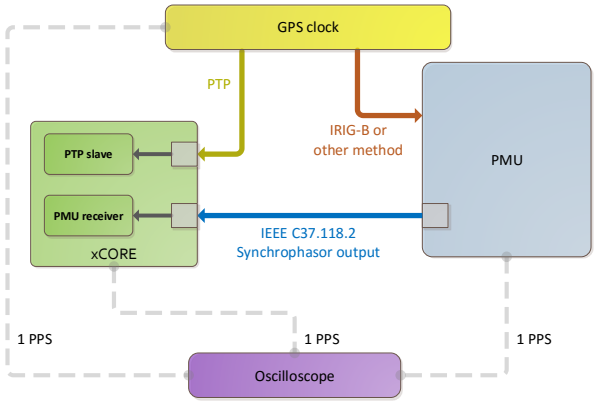

FIGURE 7. Configuration Option 2 (see Fig. 4b) with timing accuracy validation using 1 PPS signals

B. TIME SYNCHRONIZATION ACCURACY

Timing accuracy has been measured by comparing the 1 Pulse Per Second (PPS) signal from the master GPS clock with the 1 PPS signals recreated by the xCORE and the GTSYNC timing card within the RTDS. This process is illustrated for configuration Option 2 in Fig. 7, with the full results given in Table 1.

The results in Table 1 compare the four configuration options and, where appropriate, sub-options (e.g. 2a, 2b, and 2c) with different Ethernet switch hardware and calibration adjustments. These results show that it is possible achieve time synchronization accuracy within <500 ns between the xCORE and the PMU under test in several configurations. Furthermore, by monitoring the 1 PPS signals, the xCORE can be calibrated (i.e. an additional time offset can be manually added) to further reduce the timing uncertainty to <100 ns (Option 2c and Option 4b). The impact of using non-PTP Ethernet switches within these configurations is shown through Option 1, Option 2a, and Option 3a; this results in an error of approximately 10 µs between the xCORE and absolute time. This error is well within the 0.002 s accuracy

TABLE 1. Comparison of time synchronization accuracy under different configurations

Configuration option (see Fig. 4)	xCORE synchro-nization source	GTSYNC synchro-nization source	xCORE synchro-nized using non-PTP switch?	GTSYNC synchronized using non-PTP switch?	1 PPS error between xCORE and GTSYNC (ns)	1 PPS jitter between xCORE and GTSYNC (ns)	Calibration of xCORE clock offset to match GTSYNC 1 PPS signal	Notes
1	GPS clock using PTP	GPS clock using PTP	Yes	Yes	1300	500	None	Both xCORE and GTSYNC are ~10 µs off absolute time due to non-PTP switch.
2a	GPS clock using PTP	GPS clock using IRIG-B	Yes	n/a	9000	100	None	GTSYNC within ~50 ns of absolute time.
2b	GPS clock using PTP	GPS clock using IRIG-B	No	n/a	800	100	None	
2c	GPS clock using PTP	GPS clock using IRIG-B	No	n/a	25	100	-810 ns	
3a	GPS clock using PTP	xCORE using PTP	Yes	No	400	200	None	Both xCORE and GTSYNC ~10 µs off absolute time.
3b	GPS clock using PTP	xCORE using PTP	No	No	380	100	None	Both xCORE and GTSYNC ~1 µs off absolute time.
4a	Local time	xCORE using PTP	n/a	No	380	50	None	Not using absolute time
4b	Local time	xCORE using PTP	n/a	No	0	50	-380 ns	Not using absolute time

requirement of the Synchrophasor standard, but clearly PTP transparent clocks should be used for the best accuracy.

C. PMU TESTING RESULTS

Two PMU implementations have been used to demonstrate the use of the measurement method proposed in this paper: the RTDS GTNET simulated PMU (which is based on the IEEE C37.118.1 reference PMU) and an adaptive filter-based PMU implementation described in [22]–[24]. The results are summarized in Table 2, and are discussed in the following subsections. In all cases, 7000 samples have been taken, based on the amount of memory available on the xCORE device for storing results. For convenience, a non-PTP Ethernet switch has been used for some of the tests involving the adaptive filter PMU (i.e. using configuration Option 1 without a transparent clock). This means that these latency measurements have an error of approximately 10 µs, but this is well within the Synchrophasor standard requirements and only comprises a small proportion (i.e. 10 µs / 20.234 ms = 0.05%) of the actual measurement values. In all cases, a 50 Hz nominal power system frequency is used (except where the actual frequency is deliberately modified for some tests), and the measured reporting latency is well within the standard requirements of $2/F_s$ (for P class) or $7/F_s$ (for M class), where F_s is the PMU reporting rate. The results are analyzed in detail in the following subsections.

1) RTDS GTNET Simulated PMU Results

The RTDS GTNET PMU provides an implementation of the basic reference PMU provided in the Synchrophasor standard. The GTNET results in Table 2 use a Hamming window, and an emulated sampling rate of 16 samples per nominal cycle (i.e. $16 \times 50 = 800$ Hz). Note that for practical constraints, the GTNET PMU implementation requires relatively high additional latency of 1.5-3 ms, as is reflected in the measured results.

2) Adaptive Filter PMU Implementation Results

The adaptive filter PMU algorithm is based on the Discrete Fourier Transform with adaptive filtering and other enhancements as described in [22]–[24]. The algorithm has been implemented on a Beckhoff hardware platform which provides time synchronization using PTP [25]. The analogue sampling operates at 10 kHz (with signals provided by the RTDS analogue outputs) and is tightly regulated in hardware, with the samples aligned with the PTP synchronization clock reference. The measurement modules can also be physically distributed using EtherCAT. The PMU algorithm processing occurs in a "soft real-time" manner; this is the cause of the relatively high standard deviation of latency, compared to the RTDS GTNET PMU, given in Table 2.

Due to adaptive filtering which is unique to this implementation, the PMU reporting latency depends on the measured

TABLE 2. Measured PMU reporting latency for various configurations: clocking configurations, PMU class, reporting rate, and signal input frequency

PMU device	PMU input type	Signal input	Configuration option (see Fig. 4 and Table 1)	Reporting rate, F_s (Hz)	PMU class	Mean latency (ms)	Max latency (ms)	Std. dev. of latency (μs)	Theoretical latency, based on window length (ms)	Difference between measured mean latency and theoretical latency (ms)
RTDS GTNET	Digital	50 Hz	2a	50	P	21.595	21.626	8.7	20.0	1.595
RTDS GTNET	Digital	50 Hz	2a	50	M	91.846	91.871	8.0	88.75	3.096
RTDS GTNET	Digital	50 Hz	2a	100	P	21.594	21.619	6.4	20.0	1.594
RTDS GTNET	Digital	50 Hz	2a	100	M	44.344	44.373	6.4	41.25	3.094
Adaptive Filter	Analogue	50 Hz	1a	50	P	20.234	20.285	28.9	20.0	0.234
Adaptive Filter	Analogue	50 Hz	1a	50	M	100.231	100.286	29.3	100.0	0.231
Adaptive Filter	Analogue	50 Hz	1a	100	P	20.240	20.286	27.6	20.0	0.240
Adaptive Filter	Analogue	50 Hz	1a	100	M	60.230	60.284	32.4	60.0	0.230
Adaptive Filter	IEC 61850 SV	50 Hz	2b (PMU does not require synchronization)	50	M	101.001	101.055	29.6	100.0	1.001
Adaptive Filter	Analogue	55 Hz	1a	100	M	54.780	54.830	31.8	54.545	0.234
Adaptive Filter	Analogue	45 Hz	1a	100	M	66.898	66.950	24.9	66.667	0.232

frequency value. The impact of this is shown through the tests at off-nominal frequency for the M class implementation in Table 2, where the reporting latency decreases as the system frequency increases (due to the reduced window length).

3) Impact of Processing Time

The proposed measurement method can also be used to estimate the impact of the processing time of the PMU under test. For example, the adaptive filter M class PMU algorithm uses a ten-cycle window length (i.e. the total filter group delay) for a 50 Hz reporting rate, which equates to 200 ms at nominal frequency. The Synchrophasor report timestamp is defined as corresponding to the middle of the window; therefore the theoretical PMU reporting latency, at nominal frequency, is 200 ms / 2 = 100 ms. From the measured reporting latency results in Table 2, it can be calculated that the additional latency due to measurement acquisition, algorithm processing, and generating valid PMU report Ethernet frames is approximately 100.231 ms−100 ms = 0.231 ms. The results for each test are given in the final column in Table 2; the range in values demonstrates how the choice of the implementation platform and protocol can influence the overall latency.

4) Impact of IEC 61850 SV Input

The RTDS GTNET card can be configured to represent a Merging Unit which supplies voltage and current waveform data using the IEC 61850-9-2 SV protocol. It can be observed from Table 2 that the use of SV as the input to the PMU adds approximately 800 µs to the overall reporting latency due the additional stage involving a Merging Unit digitizing and packetizing the waveform data; this is dependent on the Merging Unit implementation, performance, and the number of samples per packet [7]. The SV latency results are also significantly higher than the encoding performance given in [26] (even considering the difference in dataset size) due to practical restrictions of the RTDS GTNET implementation.

The platform presented in this paper has been augmented to measure SV latency directly, in addition to measuring PMU reporting latency, using a method similar to [27]. For the GTNET Merging Unit, a mean SV latency of 825 µs has been measured which is consistent total reporting latency with the SV PMU result in Table 2. Note that the Second of Century (SOC) value—which is not normally included in SV messages—is encoded within each sample contained within each SV frame. The PMU therefore does not need to be synchronized to absolute time in this configuration because the timestamp is recorded by the Merging Unit.

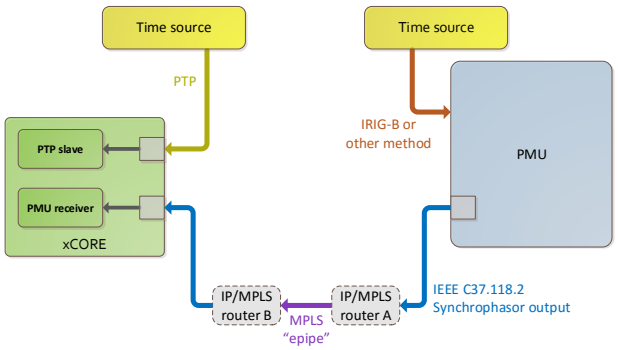

FIGURE 8. IP/MPLS WAN configuration

IV. WIDE-AREA NETWORK LATENCY MEASUREMENT AND SIMULATION

A. LABORATORY-BASED WAN DEMONSTRATION

A major advantage of the measurement method proposed in this paper is that it can be applied to measure the full end-to-end latency of PMU measurements in real utility communications networks (i.e. the time for PMU reports, relative to the report timestamp, to reach the end user or application). A similar approach for testing distributed power system protection performance is described in [28], albeit requiring relatively expensive hardware and software, and not tailored to the requirements of PMUs. This capability has been proven for representative wide-area networks (WANs) in two ways:

1) Using a modern packet-based WAN, implemented with Internet Protocol/Multiprotocol Label Switching (IP/MPLS) [19], [29]. This configuration, using two commercially-available IP/MPLS routers, is illustrated in Fig. 8. An "epipe" service has been used to transport the IEEE C37.118.2 PMU data over the IP/MPLS network.
2) Using an additional xCORE device to emulate a large network, by delaying Ethernet traffic by configurable amounts in real-time, using the method demonstrated in [20]. To mimic the potential for jitter in large WANs, a delay characteristic has been used with a fixed delay of 1 ms, plus a variable delay (with a mean of 10 ms and std. dev. of 2 ms).

The results are summarized in Table 3. In all cases, the RTDS GTNET PMU implementation with a reporting rate of $F_s = 50$ Hz has been tested, and 7000 PMU reports have been sampled. It can be observed that the IP/MPLS network adds relatively small latency, much of which is due to additional Ethernet link transmission times (adding ~7.4 µs per link, for 92 byte Ethernet frames at 100 Mbps). As would be expected, the impact of real-time network emulation on PMU latency is much more significant than for the two-node IP/MPLS network.

Fig. 9 illustrates the impact of the emulated communications network delay on total PMU latency (i.e. including measurement and communications). It can be observed that the distributions are slightly skewed, compared to a Gaussian distribution. This is because the order of packets is maintained regardless of the random delay applied to a given packet i.e. packets are queued in the xCORE device until all prior packets have been transmitted. Similarly, the deviation of the mean latency increase given in Table 3 from the theoretical values (mean of 11 ms, std. dev. of 2 ms) is due to the skewed distribution, not due to accuracy of the PMU latency measurement platform.

TABLE 3. Summary of impact of WANs on PMU latency

Communications network type	PMU class	Mean latency (ms)	Std. dev. of latency (μs)	Mean latency increase (compared to Table 2) due to communications network (μs)
IP/MPLS (two nodes)	M	91.893	8.01	47
IP/MPLS (two nodes)	P	21.643	7.51	48
Real-time network emulation	M	102.925	1955	11079
Real-time network emulation	P	32.670	1954	11075

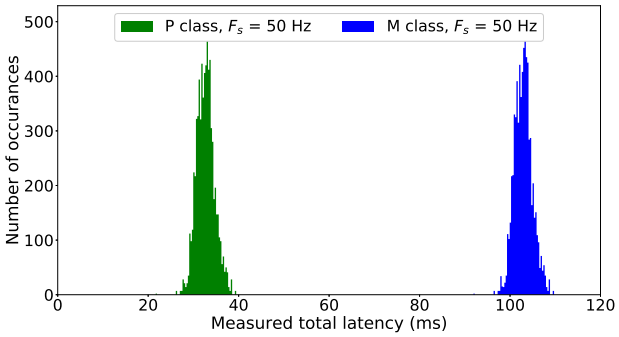

FIGURE 9. Distributions of measured RTDS GTNET PMU latency in emulated WAN

B. APPLICATIONS IN POWER SYSTEM SIMULATION

This section demonstrates the practical use of the latency characterization data acquired, as given in Section IV-A, to significantly improve the realism of power system simulation studies, whilst also enabling simpler models to be used—thereby enabling complex smart grid solutions to be conveniently and comprehensively designed and validated. Fig. 10 illustrates a hypothetical PMU-based differential protection scheme, where PMU data is transferred over a WAN. Each protection Intelligent Electronic Device (IED) receives local and remote current phasors which are compared according to a typical line differential protection algorithm for a 400 kV transmission system [19]. The objective is to realistically

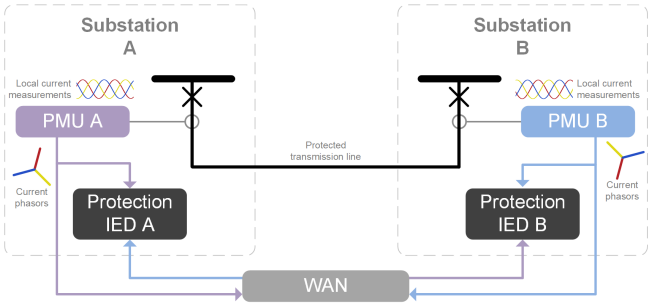

FIGURE 10. Simulated PMU-based differential protection scheme

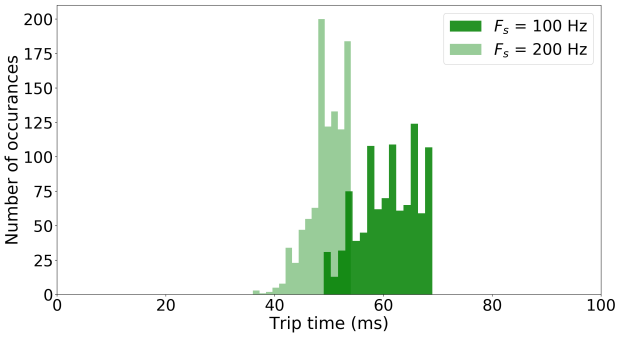

FIGURE 11. Trip time distributions using P class PMU-based protection scheme

characterize the protection operation time (i.e. the trip time) of this system following a simulated fault, catering for as many practical considerations as possible.

The power system model and protection IED logic have been implemented in MATLAB Simulink. The use of a phasor-based power system simulation, rather than a detailed transient simulation [30], greatly simplifies the model, reduces execution time, and avoids the need to explicitly implement a PMU algorithm because the simulation intrinsically generates current phasors at each time-step. However, using the simulation phasor data directly for protection or other real-time applications is very unrealistic because it does not incorporate the latency associated with the measurement window of the PMU (i.e. the reporting latency), the reporting rate of the PMU (which dictates strict intervals for sending Synchrophasor data), or communications delays. Therefore, an additional communications emulation logic block has been added to the simulation to do the following:

1) Down-sample the simulation time-step (1 ms) to map to appropriate PMU reporting rates (e.g. 100 or 200 reports per second). This emulates the periodic, packetized nature of the PMU data stream.
2) For the PMU data transferred over the WAN, the data is queued with a random delay to represent the measurement and communications latency. The latency data acquired for the "real-time network emulation" P class PMU from Table 3, with a mean total latency of 32.7 ms, has been used. To maintain the original order of data within the queue, the delay applied to a given set of phasors is forced to be greater than that for data already in the queue; this also represents the level of service which can be achieved using modern packet-based WANs [14]. It should be noted that a mean of 11 ms to represent the WAN is significantly larger than would be expected in practice for such a protection scheme, but has been chosen to be consistent with Section IV-A. Alternatively, the maximum latency value from the measured results could be used to further simplify the simulation study.
3) Pass the PMU data to the remote Protection IED after the simulation time reaches the computed delay time.

The delay for the local PMU measurements to reach the Protection IED is much smaller than the delay of the remote measurements, and is therefore ignored. The protection algorithm compares local and remote current phasors with the same timestamp, and must "trigger" (i.e. detect fault conditions) for three consecutive measurements before a trip is issued. Fig. 11 illustrates the results for the trip time of the protection scheme, for PMU reporting rates of 100 and 200 reports per second, following the initiation of a simulated three-phase short-circuit fault within the protected zone. Due to the stochastic nature of the latency characterization, the entire simulation is executed for 1000 iterations to provide a distribution of the trip time. The distributions illustrate the combined effects of the PMU reporting period, the random delay (according to the Table 3 data) to represent the measurement and communications latency, and the requirement for three trip confirmations. As noted in Section IV-A, the distributions are skewed due to packet order being maintained.

Without the emulation of measurement and communications (but still catering for the PMU reporting rate), the simulation would yield a constant trip time of 59 ms (F_s =100 Hz) or 44 ms (F_s =200 Hz)—both of which are incorrect estimations of the maximum time. Using static parameters may be acceptable for some applications, but for protection schemes it is important to understand the worst-case behavior. Furthermore, real utility communications networks may not fit the relatively simple assumption of Gaussian latency characteristics [31], and therefore direct measurement is required to accurately determine the characteristics. This example therefore demonstrates how the PMU latency measurement method introduced in this paper can be used to create more realistic simulations—informed by actual data—whilst also reducing simulation complexity because phasor representation models can be used instead of transient models. Furthermore, this enables larger simulations—such as investigations of the scalability of wide-area control, protection, and automation systems—to be implemented more conveniently and accurately. This does not fully replace the need for laboratory validation with real PMUs (to cater for measurement phenomena such as perceived frequency deviations during phase step changes [32]), but enables the

rapid prototyping and validation of novel control and protection schemes which require realistic representation of power system, measurement, and communications domains.

V. CONCLUSIONS

Timely measurements are critical for addressing many challenges associated with power system operation, such as the increasing requirement for fast-acting reserves to stabilize frequency following a significant system disturbance. This paper has presented a new method to very accurately and conveniently characterize the actual latency performance of PMU measurements. The open source software is readily available at [14] for use in further research and development activities. The timing accuracy achieved is typically <500 ns—significantly more accurate and more cost-effective than the method presented in [13] (with accuracy of approximately 600-900 µs). It has also been shown how the proposed method can be used to measure end-to-end latency of PMU applications in representative wide-area communications networks, and how this information is valuable to improve the convenience and realism of cross-domain simulation studies; this advance therefore enables demanding, time-critical PMU-based systems to be designed and validated.

By significantly improving the ease and accuracy of measuring PMU reporting latency, this work may attract future changes to the IEEE C37.118 Synchrophasor standard to require stricter PMU reporting latency measurement accuracy for real-time PMU applications. The use of the platform for PMU and IEC 61850 Sampled Value latency measurement has been demonstrated in this paper, and measuring the latency of other time-critical applications—such as IEC 61850 GOOSE messaging performance for power system protection applications—could also be supported in future work.

REFERENCES

[1] A. J. Roscoe, M. Yu, R. Ierna, J. Zhu, A. Dyśko, H. Urdal, and C. Booth, "A vsm (virtual synchronous machine) convertor control model suitable for rms studies for resolving system operator/owner challenges," in 15th Wind Integr. Work., 2016.

[2] P. Wall, N. Shams, V. Terzija, V. Hamidi, C. Grant, D. Wilson, S. Norris, K. Maleka, C. Booth, Q. Hong, and A. Roscoe, "Smart frequency control for the future gb power system," in IEEE PES ISGT Eur., 2016.

[3] G. Rietveld, J.-P. Braun, R. Martin, P. Wright, W. Heins, N. Ell, P. Clarkson, and N. Zisky, "Measurement infrastructure to support the reliable operation of smart electrical grids," IEEE Trans. Instrum. Meas., vol. 64, no. 6, pp. 1355 – 1363, 2015.

[4] S. Blair, G. Burt, N. Gordon, and P. Orr, "Wide area protection and fault location: review and evaluation of pmu-based methods," in 14th Int. Conf. Dev. Power Syst. Prot. (DPSP 2018), Mar. 2018.

[5] K. Kok and S. Widergren, "A society of devices: Integrating intelligent distributed resources with transactive energy," IEEE Power Energy Mag., vol. 14, no. 3, pp. 34–45, May 2016.

[6] E. G. Sansano, M. H. Syed, A. Roscoe, G. Burt, M. Stanovich, and K. Schoder, "Controller hil testing of real-time distributed frequency control for future power systems," in IEEE PES Innov. Smart Grid Technol. Eur., 2016.

[7] S. M. Blair, A. J. Roscoe, and J. Irvine, "Real-time compression of iec 61869-9 sampled value data," in 2016 IEEE Int. Work. Appl. Meas. Power Syst. IEEE, 2016, pp. 1–6.

[8] A. Monti, C. Muscas, and F. Ponci, Phasor Measurement Units and Wide Area Monitoring Systems. Elsevier, 2016.

[9] B. J. Pierre, F. Wilches-Bernal, D. A. Schoenwald, R. T. Elliott, J. C. Neely, R. H. Byrne, and D. J. Trudnowski, "Open-loop testing results for the pacific dc intertie wide area damping controller," in 2017 IEEE Manchester PowerTech. IEEE, Jun. 2017, pp. 1–6.

[10] NERC, "Reliability guideline - pmu placement and installation," Tech. Rep., 2016.

[11] NASPI Time Synchronization Task Force, "Time synchronization in the electric power system," NASPI, Tech. Rep., 2017.

[12] J. Braun and S. Siegenthaler, "The calibration of static and dynamic performances of pmus," in 17th Int. Congr. Metrol., B. Larquier, Ed. Les Ulis, France: EDP Sciences, Sep. 2015, p. 12002.

[13] P. Castello, C. Muscas, P. A. Pegoraro, and S. Sulis, "Automated test system to assess reporting latency in pmus," in 2016 IEEE Int. Instrum. Meas. Technol. Conf. Proc. IEEE, May 2016, pp. 1–6.

[14] S. M. Blair, "Real-time measurement of pmu reporting latency," 2017. [Online]. Available: https://doi.org/10.5281/zenodo.400934

[15] IEEE, "1588-2008 ieee standard for a precision clock synchronization protocol for networked measurement and control systems," 2008.

[16] Synchrophasor Measurements for Power Systems Working Group, "C37.118.1a-2014 - ieee standard for synchrophasor measurements for power systems – amendment 1: Modification of selected performance requirements," Tech. Rep., 2014.

[17] K. E. Martin, "Synchrophasor measurements under the ieee standard c37.118.1-2011 with amendment c37.118.1a," IEEE Trans. Power Deliv., vol. 30, no. 3, pp. 1514–1522, Jun. 2015.

[18] G. Martins, D. Lacey, A. Moses, M. J. Rutherford, and K. P. Valavanis, "A case for i/o response benchmarking of microprocessors," in IECON 2012 - 38th Annu. Conf. IEEE Ind. Electron. Soc. IEEE, Oct. 2012, pp. 3018–3023.

[19] S. M. Blair, C. D. Booth, B. De Valck, D. Verhulst, C. Kirasack, K. Y. Wong, and S. Lakshminarayanan, "Validating secure and reliable ip/mpls communications for current differential protection," in Dev. Power Syst. Prot., 2016.

[20] S. M. Blair, C. D. Booth, B. De Valck, D. Verhulst, and K.-Y. Wong, "Modeling and analysis of asymmetrical latency in packet-based networks for current differential protection application," IEEE Trans. Power Deliv., vol. 33, no. 3, pp. 1185–1193, Jun. 2018.

[21] XMOS, "Time sensitive networking library," 2016. [Online]. Available: https://github.com/xmos/lib{_}tsn

[22] A. J. Roscoe, "Exploring the relative performance of frequency-tracking and fixed-filter phasor measurement unit algorithms under c37.118 test procedures, the effects of interharmonics, and initial attempts at merging p-class response with m-class filtering," IEEE Trans. Instrum. Meas., vol. 62, no. 8, pp. 2140–2153, Aug. 2013.

[23] A. J. Roscoe, I. F. Abdulhadi, and G. M. Burt, "P and m class phasor measurement unit algorithms using adaptive cascaded filters," IEEE Trans. Power Deliv., vol. 28, no. 3, pp. 1447–1459, Jul. 2013.

[24] A. J. Roscoe and S. M. Blair, "Choice and properties of adaptive and tunable digital boxcar (moving average) filters for power systems and other signal processing applications," in 2016 IEEE Int. Work. Appl. Meas. Power Syst. IEEE, Sep. 2016, pp. 1–6.

[25] M. J. Kiekebusch, N. Di Lieto, S. Sandrock, D. Popovic, and G. Chiozzi, "Mathworks simulink and c++ integration with the new vlt plc-based standard development platform for instrument control systems," in SPIE Astron. Telesc. + Instrum., G. Chiozzi and N. M. Radziwill, Eds. International Society for Optics and Photonics, Jul. 2014, p. 91522B.

[26] S. M. Blair, F. Coffele, C. D. Booth, and G. M. Burt, "An open platform for rapid-prototyping protection and control schemes with iec 61850," IEEE Trans. Power Deliv., vol. 28, no. 2, pp. 1103–1110, 2013.

[27] D. M. E. Ingram, F. Steinhauser, C. Marinescu, R. R. Taylor, P. Schaub, and D. A. Campbell, "Direct evaluation of iec 61850-9-2 process bus network performance," IEEE Trans. Smart Grid, vol. 3, no. 4, pp. 1853–1854, Dec. 2012.

[28] F. Steinhauser, "Assessing communication networks for distributed protection and automation systems with time synchronized and distributed measurement systems," in 13th Int. Conf. Dev. Power Syst. Prot. Institution of Engineering and Technology, 2016.

[29] S. M. Blair and C. D. Booth, "Real-time teleprotection testing using ip/mpls over xdsl," Glasgow, 2013. [Online]. Available: https://pure.strath.ac.uk/portal/files/26184600/001{_}DSL{_}Testing.pdf

[30] S. Abourida, J. Bélanger, and V. Jalili-Marandi, "Real-time power system simulation: Emt vs. phasor," OPAL-RT Technologies Inc., Tech. Rep., 2016.

[31] A. Derviskadic, P. Romano, M. Pignati, and M. Paolone, "Architecture and experimental validation of a low-latency phasor data concentrator," IEEE Trans. Smart Grid, vol. 9, no. 4, pp. 2885–2893, Jul. 2018.
[32] A. J. Roscoe, A. Dysko, B. Marshall, M. Lee, H. Kirkham, and G. Rietveld, "The case for redefinition of frequency and rocof to account for ac power system phase steps," in 2017 IEEE Int. Work. Appl. Meas. Power Syst. IEEE, Sep. 2017.