

Network level Quality of Service (QoS) challenges for Smart Grid Measurement and Control systems

Kashinath Basu

Department of Computing and Communication
Technologies
Oxford Brookes University
Oxford, UK
kbasu@brookes.ac.uk

Ali Maqousi and Tatiana Balikhina

Faculty of Information Technology
University of Petra
Amman, Jordan
amaqousi@uop.edu.jo; tbalikhina@uop.edu.jo

Frank Ball

Frank Ball Consulting
Oxford, UK
frank_ball@ntlworld.com

Abstract— This paper focuses on the firm real-time requirements of Time-Critical Wide Area Measurement and Control systems, that are expected to play a major role in future Smart Grids. It analyses the operation of these systems and identifies their communication traffic characteristics. It shows that these characteristics are significantly different to those of the current near real-time Wide Area Measurement applications that provide visualization to support manual grid control. It then discusses the performance evaluation of these time critical systems and presents the second stage in an ongoing body of work aimed at developing models and techniques to carry out the performance evaluation process. It presents some preliminary results and outlines the direction for future work.

Keywords—smart grid; synchrophasor; phasor measurement units; phasor data concentrators; real-time; software defined network; wide area measurement systems; QoS

I. INTRODUCTION

Smart grid communication systems will be required to support a wide range of applications, some of which will have similar requirements to those of current Internet applications, and others that may fall into the remit of the Internet of Things (IoT). However, certain classes of smart grid applications, in particular, those intended to support automatic wide area control in the smart grids of the near future, have expectations that are significantly different from those of most existing wide area applications. This is due to their need for firm real-time delay bounds that have quite stringent delay targets. Furthermore, these delay bound apply to the combined delay resulting from both application level processing time and communication latency. The most time-critical of these applications require end-to end delay targets of 10ms or less. Although no separate targets has been set for the communication delay, it has been suggested that 1ms to 2ms

would be an appropriate goal for the delay component of the communication network [1] [2].

The drive to extend automatic control into the wide area is motivated by the two fundamental objectives of the Smart Grid: firstly, to provide greater efficiency in the use of current energy generation; and secondly, enable the inclusion of a wide range of renewable, but more variable, energy sources. However, extending automatic control into wide area presents the additional challenge of providing low latency in a larger scale network and over greater distances. Distances in the orders of 100km, 160km, 200kms, or even greater, are not unusual, and therefore, the effects of propagation delay will be significant. Furthermore, failures in the smart grid control system can lead to serious consequences [2], making it essential that delay targets for the system can be guaranteed prior to the system becoming operational. Also, given the potentially serious consequences of failing to meet these targets, it is quite possible that they may become subject to some form of mandate. Therefore, for these time-critical applications, performance evaluation must become an integral part of the system design process.

Due to the stringent nature of these latency requirements, it has suggested that point to point fibres between each monitoring device and the controller may be needed to minimize delay. Although not infeasible, this would result in fibre capacity being significantly underutilised. Furthermore, the data generated by these applications may also be required for historic purposes, such as post event analysis, and may need to be more widely distributed in non real-time communication. Therefore, using an integrated multiservice networking approach is desirable, provided that latency requirements can be guaranteed.

The aims of our investigation are: firstly, to derive generic and parameterized models to support the performance evaluation of Time-Critical Synchrophasor Measurement and

Control Systems; and secondly, to develop techniques and methods to evaluate the temporal performance of specific systems during their design phase. These models will be based on the generic concept of packet switching, so as to be applicable for both level 2 and level 3 switching. Therefore, throughout the discussion we will use the term forwarding device rather than router or switch. Our approach to the development of these models and their associated evaluation methods is to begin by studying the best case situation, and to follow this by a series of further studies in which there are relaxation of certain constraints that apply to this best case situation. We define the best case as being the situation where Phasor Measurement Units (PMUs), Phasor Data Concentrators (PDCs) and forwarding devices are high performance real-time devices that can process, transmit, receive and forward, as applicable, messages at line rate. Also in case there will be very strong isolation between real-time and non-real-time processing throughout. This is the case that will present the minimum of variation, and therefore, will be the least complex to evaluate.

This paper presents the first stage of our study and initial development of the evaluation process. The remainder of the paper is structured as follows: Section II describes the operation of a Synchrophasor Measurement and Control System and introduces its constituent devices, it briefly outlines the current operation of a current type of this system and then discusses how the more stringent delay requirements of proposed future systems result in significant changes to the traffic characteristics and QoS requirements; section III addresses the general conditions needed to meet the requirements of these time-critical systems; section IV provide an example of the performance evaluation of a best case Synchrophasor Measurement and Control System; section V outlines future work; and finally, section VI concludes.

II. SYNCHROPHASOR MEASUREMENT AND CONTROL SYSTEMS

A Synchrophasor is a measurement of the amplitude and angle of a sinusoidal waveform (in this case the waveform of power cycle) that is time-stamped using a UCT (Universally Coordinated Time) mechanism facilitated by GPS [3]. These synchronized measurements provide a comprehensive picture of the state of the power system. These measurements are taken by a PMU which is a specialized device that periodically samples the power cycle and calculates the synchrophasor measurement. Generally, six measurements are taken from the current and voltage for each of the three phases. These measurements are then encapsulated into a single fixed length message for transmission. Although that length may differ between different configurations of the device, generally, PMU devices are configured at the initialization stage of the system and remain unchanged once the system is operational. A message length (including protocol overheads) in the order of 1000bits is typical of many examples quoted in the literature. The frequency at which measurements are taken can vary depending on the requirements of the control application and the frequency of the power cycle, currently values of 10Hz, 30Hz, 50Hz, and 60Hz are employed with 120Hz being considered as a target for the future. In this paper, discussion

will be based on the case of a 60Hz power cycle and a 60Hz phasor sampling period.

PMUs are deployed throughout the grid, generally within substations, and are connected by direct communication links, or a substation LAN, to a local PDC. This device checks the validity of the messages before forwarding them as a batch, via a WAN, to a Super PDC (SPDC), which in turn has a direct connection to the Controller, as shown in Fig. 1. The end-to-end latency of the system is defined as the time between the timestamp value of the message and completion of the control decision process.




Fig. 1. Synchrophasor Data Flow Model Architecture.

PMU processing involves taking a number of evenly spaced samples over the duration of one power cycle with half the samples being taken before the UTC time stamp and the remainder taken after it. This means that there is delay of 8.35ms after the timestamp before further processing can take place. Following the sample phase a signal processing algorithm is used to calculate amplitude and angle of the synchrophasor. This information is then encapsulated into a message before being sent via the PDC, network and SPDC to the Controller. A PDC is a device whose primary role is to check for, and mark, any corrupted messages and to ensure that all messages from each set have the same timestamp. Once these checks are completed the PDC sends the messages via the network to the SPDC. This set of processes represent the time-critical functions of the PDC, however, these devices may also offer a number of other services, which for convenience of discussion we will refer to as auxiliary functions.

Currently, synchrophasor systems for wide area control applications focus mainly on providing visualization for wide area awareness. This is a near real-time process for which the latency requirements are in the order of 100ms. For these types of application the main constraint on the performance of the PMU, PDC and SPDC devices is that processing of the samples must be complete before the end of the next sampling period. In these less time-critical cases, the PDC devices can apply traffic shaping to their output. For example, for a local PDC serving 20 PMU devices that each produce one message of 1000bits, the PDC can smooth out the packet stream over the 6.7ms period resulting in a sustained rate of 1.2 Mb/s which is in the same order as figures widely reported in the literature [4]. In this case the traffic characteristics are very similar to that of other streaming applications. However, for more time-critical synchrophasor based applications, e.g. Automatic Wide Area Control, processing times will be subject to more stringent constraints and traffic smoothing will not be possible. It will result in a traffic profile that is significantly different from that of the near real-time case.

The ultimate requirement for synchrophasor based wide area automatic control applications is that they should be able

to carry out the measurement-to-decision process within one power cycle [1] [2], as shown in Fig. 2.




Fig. 2. PMU to Control Processing Cycle.

In this case the latency targets become an order of magnitude lower and therefore performance evaluation will need to be carried out at the micro second level. The evaluation process will need to consider the synchrophasor measurement and control system as a whole, although a degree of decomposition is possible, as will be discussed below. No specific delay target is assigned to each stage of the processing; however, there is some consensus that 2ms would be a reasonable network delay bound. Based on this figure the total activity period on the network due to this class of traffic would be in the order of 2ms for each cycle of 16.7ms.

There is also a consensus that this class of application requires a firm real-time guarantee that has a hard deadline but for which some missed deadlines and losses may be acceptable [1][2][4]. However, the term “some” is a rather vague parameter and as an alternative we propose the use of probabilistic hard real-time, that offers a more precise definition, i.e. for a latency bound T , a condition, such as the following, must hold:

$$P(t > T) \leq n10^{-x} \quad (1)$$

This allows the application to choose an allowable value for loss, which can then be related directly to high percentiles of delay distributions during the evaluation process. However, it should be noted, that this guarantee must apply to each individual message, and not to the aggregation of the batch in each cycle [1].

Due to the nature of PMU-PDC interaction, some degree of decomposition in the performance evaluation process is made possible. The primary traffic sources are the PMUs that produce a single message for each cycle, in synchrony with each other. PDCs wait until all the messages they expect to receive in a given cycle have arrived before starting to process them. To allow for message losses the PDC sets a waiting-time. Once this time has expired, it starts to process the messages that have arrived in time. Messages from that cycle that arrive later are discarded, and therefore, missing the deadline set by the waiting-time, is equivalent to loss. The same process is employed by the SPDC, however, in this case the messages will have been subjected to network delays. Therefore, the SPDC waiting-time setting will be influenced by a prediction of combined PMU, PDC and network delay. Once the PDC has finished performing its internal functions it

will start to transmit the messages over the network. The output from the PDC will be in form of a short burst of packets, the duration of which will depend on the number of messages and the rate at which the PDC can operate, which ideally should be designed to operate at the line rate of the communication link. Once the burst has been sent, there will be no further transmissions until the next cycle.

For example, in the case of a local PDC serving 20 PMU devices that each produce one message of 1000bits, and a link rate of 500Mb/s the burst duration would be 40 μ s. Alternatively, for a link rate of 100Mb/s (or in the case of a PDC that can only operate at that speed) the burst duration would be 200 μ s. In both cases the burst duration is very short in comparison to the cycle time of 16.7ms. In cases such as this, the concept of a stream with an average, or sustained, rate is not relevant. Also the transmission process from the PDCs to the SPDC is a case of many-to-one communication, and since all the messages are created at generally the same time, the bursts may interact with each other as they pass through the forwarding devices along paths that fan-in to the SPDC (Fig. 3). Queuing behaviour due to these interactions will generally be transitory due to the low ratio of burst duration to cycle period. Also, due to the waiting-time operation of the SPDC the most significant QoS metric will be the latency of the last message to arrive at the SPDC.

III. MEETING THE REQUIREMENTS OF TIME-CRITICAL APPLICATIONS

Bakken et al [1] present a thorough and extensive survey of wide area control in a smart grid and a detailed analysis of the overall requirements of Wide Area Measurement Systems for Data Delivery (WAMS-DD), of which Synchrophasor Measurement and Control Systems are a component. From this analysis they produce a comprehensive set of both requirements and guidelines for the implementation of a WAMS-DD. One requirement that is particularly relevant to time-critical control applications is that firm end-to-end guarantees must be provided over the entire grid, and these guarantees must be given to each individual message, and not based on a weaker aggregation over long periods of time. The guidelines that follow on from this requirements effectively propose: strict priority queuing on an end-to-end basis; forwarding decisions based on packet header only; the avoidance of retransmission mechanisms for reliable delivery; and the use of static routing for this class of traffic. It should be noted that the constraint of static routing need only apply to time-critical classes of traffic, and other classes could be served by dynamic routing in cases where the forwarding devices have the capability of providing both simultaneously. A Software Defined Networking (SDN) approach could help to facilitate such a capability due to its separation of the control and forwarding processes. To meet the requirement of reliability the guidelines propose sending each message over multiple disjoint paths. This redundant approach fits in with an existing practice of deploying PMU devices in redundant pairs [3].

Although attempting to meet the stringent requirements for future Synchrophasor Measurement and Control Systems may appear to be a difficult task, there are a few factors that help to mitigate the problem. Firstly, the system will be based on static infrastructure, and mobility will not be an issue; secondly, device configuration can occur prior to the system becoming operational; thirdly, full information regarding the number of PMUs, PDCs, forwarding devices and their interconnections, together with link distances, can be made available prior to evaluation; finally, apart from distance, these networks will be of relatively small scale. Also, in some cases relative high rate data fibre optic carriers may be available as part of the existing infrastructure, as for example within an OPGW (Optical Ground Wire Systems) cable [5]. These cables provide both protection and communications and can be installed on high voltage pylons. They can contain up to 24 optical fibre links that each have a link rate of 500Mb/s. One further point is that the cost of using more expensive high performance equipment throughout, and redundant equipment for robustness and reliability, may not be a major issue. It has been reported that equipment costs only account for about 5% of the total cost of installing a synchrophasor measurement system [6]. Therefore, if the findings of this report represent a general case, then as an example, trebling the current cost of equipment should only add about 10% to the total bill.

Currently, there is a wide range in the performance capabilities of PMU and PDC devices. [7]. Although PMUs are subject to compliance testing for correctness of measurement and quality of data, as yet there are no compliance requirements for their temporal performance. PDCs are not subject to compliance testing, and in some cases they have been implemented on general purpose computational engine including windows PCs. Clearly, if these more stringent delay requirements are to be met, then PMUs, PDCs, the SPDC and the controller will need to become, not only faster, but true real-time devices with performance requirements being built into their design. Any auxiliary function that they provide, including reconfiguration and device updating should not be allowed to interfere with the time critical functions. Similar requirements will apply to networking equipment. In particular forwarding devices will need to provide strict priority queuing to time-critical class of traffic. Ideally, forwarding should operate at line rate, with queuing only taking place at the output links. If this is not the case, full details of internal operation and performance at the microsecond level may be required for accurate evaluation. Also, in all cases it will be essential that packet classification operates at line rate for strict priority queuing to be maintained.

IV. EVALUATION OF TIME-CRITICAL APPLICATIONS

This section presents the process used to evaluate the performance of the PMU to SPDC stages of a Synchrophasor Measurement and Control System, the message Fan-in for which is shown in Fig. 3. This example does not represent any particular real system, nor do we claim it to be typical. Therefore, the results presented here are only relevant with the context of this example.

The evaluation is carried out under the best case assumption (as outlined in section 1), and the settings are as follows: All PMUs are connected by a direct link to their PDC; each PDC serves 20 PMUs, each of which produces 1 message/cycle; the PDC output message length is 1000bits; the link rate based on OPGW fibre of 500Mb/s for all links; all forwarding devices operating at line rate, with queuing at the output link only, using a pre-emptive strict priority discipline; and a fixed forwarding latency of 10 μ s for all forwarding devices. The propagation delays, based on the widely quoted value of 5 μ s/km are shown on the links in Fig. 3.

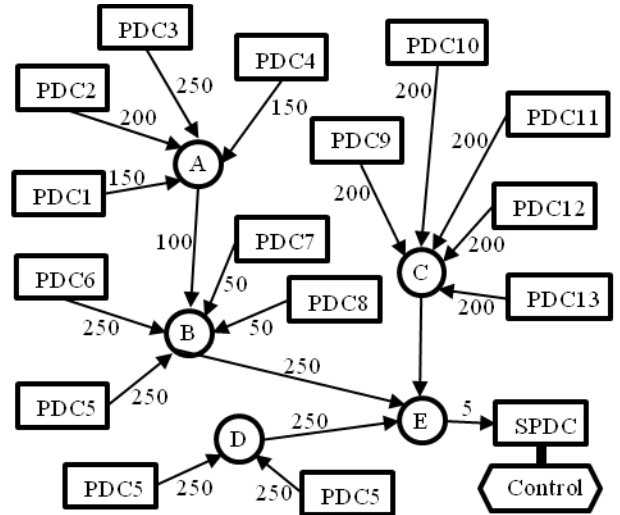


Fig. 3. PMU Message Fan-in to SPDC.

Although the main objective is to evaluate the end-to-end latency of the entire system, because the PDC process the messages as a batch before beginning to transmit them, it is possible to evaluate the communication latency between the PDC and the SDPC in isolation. From the perspective of the network the PDC is just a source that periodically creates short burst of packets. Firstly, however, it will be necessary to determine the time from the time-stamp that each PDC starts its output. To simplify discussion we will assume that this first stage has been carried out and the burst start time is the same for all PDCs.

Evaluating network latency involves working down through the levels of the fan-in, shown in Fig. 3, starting at the PDC level and evaluating the effects that each set of PDCs has on the forwarding node to which it is connected. To do this we use an output function for each PDC which is based on its transmission start time and burst length duration (in this particular example the start time will be zero for all PDCs). Using the PDC's output function in combinations with the propagation delay to the forwarding node and the node's forwarding latency value, we can derive the input function to the node's output queue. Using the set of input functions from all the attached PDCs we need to evaluate the queuing behaviour and the output function for the node. In the case of deterministic or relatively low variability, input from the PDCs, evaluation of the queuing behavior and deriving the node's output function can be readily achieved using a combination of arithmetic and basic network calculus [8].

Once the processing has been completed for all forwarding nodes at that level, the set of derived output processes can then be used to continue the process at the next level down, and so on until the SPDC is reached. The evaluation results for this example are show bellow in table 1.

TABLE I. LATENCY VALUES IN MICROSECONDS

Latency Type (μsec)	PCD Set		
	{1..8}	{9..13}	{14..15}
Total Max	1005	1015	515
Propagation Max	650	605	380
Forwarding Max	30	20	20

The maximum latency for this example is $1015\mu\text{s}$ that, together with the PMU-PDC latencies, would advise the setting of the SPDC waiting time. We estimate that the accuracy level for these results is plus or minus $2\mu\text{s}/\text{queue}$ due to the discrete nature of the evaluation, although further work is required to confirm this. Although, as would be expected, propagation delay is generally the most significant factor, it is not always the case that the PDC with the highest propagation delay has the greatest overall latency, as is shown in this example. The maximum queue occupancy of 80 messages occurred at forwarding device C, and relates to delay of $160\mu\text{s}$ which is about 16% of the total delay. Although this is not very high, it is not insignificant.

V. WORK IN PROGRESS AND FUTURE DIRECTIONS

We do not consider the process that we have used in the evaluation of the best case to be a general solution to the problem of evaluating the performance of Synchrophasor Measurement and Control Systems. However, we do see it as a useful starting point toward the development of a more comprehensive solution. Furthermore, as the process is not very time consuming it could be used for a first estimate in cases where the propagation delay from a source is approaching the desired latency target. If the delay targets cannot be met in the best case, then it is unlikely they can be met when there is greater variability. Elimination at this stage could avoid the more time costly process of evaluating a scenario with greater variability.

Although there is no principled reason why PMUs and PDCs could not be designed in such a way as to produce deterministic output, implementation convenience and other pragmatic factors will inevitably result in some degree of variability. Therefore the next stage of our investigation will be to modify and extend the evaluation process to accommodate variation. However, for accuracy, this will require that the values for the parameters of variability are made available. Ideally, any such information should include probability distributions. Furthermore, variability within the devices will need to be stable. Clearly the property of stability is something that should be expected from real-time devices, i.e. it should not be possible for any auxiliary operations to interfere with real-time processing. In the case of PMUs, PDCs and forwarding devices, the required information could

be obtained as part of compliance testing, and the viability of extending compliance tests to include these requirements will be part of our ongoing investigation.

We have recently developed the model of PDC output functions with greater variability and produced a number of potential representative distribution functions. Applying convolution to these distributions we have produced for burst length distributions which in turn we have used to construct distributions for the resulting combined burst lengths from the merger of the simultaneous output from a number of PDCs. This work is still in progress and is not yet ready for publication. However, this stage of the investigation is pointing towards a combination of discrete event simulation and Monte Carlo simulation for the evaluation of cases with greater variability.

VI. CONCLUSION

This paper has addressed the performance requirements of time-Critical Synchrophasor Measurement and Control Systems intended for the smart grids of the future. It has outlined the operation of these systems and analysed the characteristics of the network traffic they will produce. This analysis has shown that these characteristic are significantly different to that of the current Synchrophasor Measurement and Control Systems. It has reviewed the conditions, mechanisms and approaches needed to meet the requirements of these system and introduced the initial stages of a process (based on generic and parameterized models) that are being developed in order evaluate their performance.

References

- [1] D.E. Bakken, A. Bose, C.H. Hauser, D.E. Whitehead and G.C. Zweigle, "Smart Generation and Transmission With Coherent, Real-Time Data," Proceedings of the IEEE, vol. 99, no. 6, June, 2011.
- [2] K. C. Budka., J. G. Deshpande, T. L. Doumi., M. Maddan and T. Mew, "Communication Network Architecture and Design Principles for Smart Grids," Bell Labs Technical Journal, vol. 15 no. 2, pp 205-228, 2010.
- [3] M. Popovic, P. Gao, D. Tomozei, J. Le Boudec, "On the Necessity of Traffic Shaping for PMU Measurement Data Streams," Power and Energy Automation Conference, Spokane, Washington, USA, March 26-28, 2013.
- [4] F.Ball. K Basu, " Performance Evaluation of Time-Critical Smart Grid Applications," Proc. of 11th International Network Conference (INC 2016), Frankfurt, Germany, July 2016.
- [5] M. H. Yaghmaee, Z. Yousefi, M. Zabihi and S. Alishahi, "Quality of Service Guarantee in Smart Grid Infrastructure Communication using Traffic Classification," C I R E D 22nd International Conference on Electricity Distribution. Stockholm, Sweden, June, 2013.
- [6] "Factors affecting PMU installation Cost," US Department of Energy, Smartgrid.Gov, Oct 2014.
- [7] D. M. Laverty, R. J. Best, P. Brogan, I. A. Khatib, L. Vanfretti, and D. J. Morrow, "The Open PMU Platform for Open-Source Phasor Measurements," IEEE Transactions on Instrumentation and Measurement, vol. 62, no. 4, April 2013.
- [8] W.M. Sofack, and M. Boyer, "Non Preemptive Static Priority with Network Calculus," IEEE 16th Conference on Emerging Technologies & Factory Automation (ETFA 2011), pp. 1-9, Sept. 2011.