

Performance Evaluation of Time-Critical Smart Grid Applications

Frank Ball
Frank Ball Consulting
Oxford UK
frank_ball@ntlworld.com

Kashinath Basu
School of Computing
Oxford Brookes University
Oxford UK
kbasu@brookes.ac.uk

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 first stage in a 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, time-critical applications, performance evaluation

I. INTRODUCTION

A communication network is an essential component of a smart grid system. Its role is to support a wide range of applications, many of which have very similar requirements to those of current Internet applications. In particular, they have the general requirements for security, resilience, reliability and wide area interconnectivity. However, a number of classes of smart grid applications, particularly those intended for controlling smart grids in the near future, have requirements that are significantly different from those of any existing Internet application. This is due to their Firm-real-time end-to-end delay bound requirements. Furthermore the delay bound specified for each individual type of time-critical application refers to the combined delay resulting from both application level processing time and communication latency. For the most time-critical of these applications, delay targets in the orders of 3ms, 10ms and 16ms have been stipulated. Although no separate targets has been set for communication delay, it has been suggested by some that 1ms to 2ms would be an appropriate goal for the delay component of the communication network [1] [2].

Time-critical smart grid applications are responsible for state estimation, control, protection, and ensuring the stability of power generation and distribution. Their domains of operation include both the local area for internal sub-station control, and the wide area for protection, control, and maintaining wide area awareness. Currently, when operating

in the wide area, the role of these applications is generally limited to providing visualization for wide area awareness. Applications that provide visualization have near-real-time requirement and can tolerate latencies in the order of 100ms. Currently, automatic control, for which the more stringent delay requirements apply, is mainly limited to the local area. However, the goal for the future is to extend automatic control into the wide area. This goal is motivated by the fundamental objectives of the Smart Grid. That are: 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. 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, making it essential that delay targets for the system can be guaranteed prior to the system becoming operational.

Due to the stringent nature of the latency requirements, it has suggested that point to point fibers between each device and the controller may be needed to minimize delay. However, this would result fiber capacity being significantly underutilized. Furthermore, the data generated by these application is also used for historic purposes, such as post event analysis. Therefore, using an integrated multiservice networking approach would be 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 Synchronphasor Measurement and Control Systems; and secondly, to develop techniques and methods to evaluate the temporal performance of models based on specific systems. These models will be based on the generic concept of packet switching, so as to be applicable both level 2 and level 3 switching. Therefore, throughout the discussion we will use the term forwarding device rather than router or switch. This paper present the first stage of this investigation in which we analyze the data flow patterns of Synchronphasor Measurements Systems and identify the characteristics of the network traffic they will produce. We

show these characteristics to be significantly different to those of the Near-Real-Time control applications in current use, and that a different interpretation of the QoS parameters to that generally applied to streaming applications, will be needed. The remainder of the paper is structured as follows: section 2 provides a brief overview of requirements of Wide Area Measurement System and highlights the points that relate to time-critical applications; Section 3 describes the operation of a Synchrophasor Measurement System, introduces their constituent devices, and outlines the current operation of this type of system. It then discusses how the more stringent delay requirements of proposed future systems present a significant challenge to the performance of the system; section 4 presents an analysis of the traffic characteristics and performance parameters that relate to these systems and briefly presents some results from a preliminary performance evaluation. Finally; section 5 concludes and outlines future work.

II. WIDE AREA MEASUREMENT SYSTEMS

A Wide Area Measurement System (WAMS) [3] is an advanced sensing and measurement system used to continuously monitor the power grid. System state and power quality are monitored using information obtained from Phasor Measurement Units (PMUs), which are devices deployed throughout the grid. To support robustness and reliability, it is current practice to deploy PMUs as redundant pairs. PMUs provide accurate system state measurements in real-time. The information generated within a WAMS is used not only by time-critical applications, but may also be required as historical information to be used, for example, in post event analysis.

To support the range of communication needed by a WAMS, Wide Area Measurement Systems for Data Delivery (WAMS-DD) have been proposed. 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 a WAMS-DD. From this analysis they produce a comprehensive set of both requirements and guidelines for the implementation of a WAMS-DD. The requirements that are particularly relevant to the time-critical applications can be summarized as follows: the smart grid communication system must provide a wide range of QoS and a "one sizes fits all" approach is not possible; the time-critical control applications require firm end-to-end deterministic guarantees that 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 these requirements are as follow: don't depend on priority guarantees that are based on preferential treatment in times of heavy traffic, as they cannot provide firm end-to-end delay bounds, i.e. use mechanisms that provide strong class isolation; avoid post-error recovery mechanisms since these can add considerable latency in the case of dropped packets, this guideline proposes that a better alternative would be to send each message over multiple disjoint paths; use static routing not dynamic, much stronger latency guarantees can be given using complete knowledge of the network topology; forwarding decisions should be based on packet header only; exploit a priori

knowledge of predicable traffic; exploit the much smaller scale of a WAMS-DD system in comparison to the Internet. It should be noted that these particular guidelines are intended mainly for supporting the time-critical traffic classes. In particular the constraint of static routing need not apply to other classes of traffic, provided that the forwarding devices can support a combination of both static and dynamic routing.

III. SYNCHROPHASOR MEASUREMENT 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 timestamped using a UCT (Universally Coordinated Time) mechanism facilitated by GPS [4]. These synchronized measurements provide a comprehensive picture of 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 Phasor Data Concentrator (PDC). This device checks the validity of the messages before forwarding them as a batch, via a WAN, to a Super Phasor Data Concentrator (SPDC), which in turn has a direct connection to the Controller, as shown in figure 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.

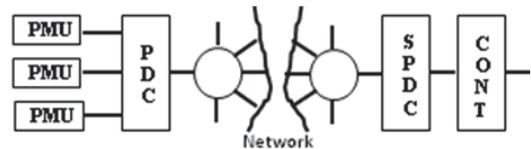


Fig. 1. Synchrophasor Data Flow Architecture

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 transmitted. In the case of less time-critical synchrophasor based applications, e.g. visualization, that is a Near Real-Time process, the latency requirements are in the order of 100ms. For this types of application the main constraint on the performance of the PMU is that processing of the samples must be complete before the end of the next sampling period. See below in Figure 2.

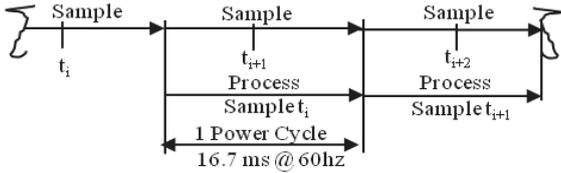


Fig. 2. PMU Processing Cycle

Also in these less time-critical cases the PCD device can apply traffic shaping to its 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 the time-critical synchrophasor based application, e.g. Wide Area Automatic Control, processing times will be subject to more stringent constraints and traffic smoothing will not generally, be possible. This 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 control applications is to carry out the measurement-to-decision process within one power cycle [1] [2]. Figure 3 shows how meeting this requirement significantly changes the processing cycle from that shown in figure 2.

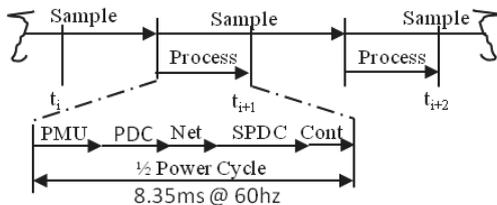


Fig. 3. PMU to Control Processing Cycle

The latency targets become an order on 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 shown in section 4. There is 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 acceptable [1][2][3]. 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, the following condition must hold,

$$P(t > T) \leq 10^{-x} \quad (1)$$

This allows the application to choose a desired value for x, and during evaluation the requirement can be related directly to high percentiles of a delay distribution. However, in order to provide this guarantee to each individual message, as advised in the guidelines presented in section 2, this condition must be applied to each source individually, and not to the aggregation of the batch in each cycle.

Currently, there is a wide range in the performance capabilities of PMU and PDC devices. [5]. 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 be implemented on general purposes computational engine including windows PCs. Clearly, if these more stringent requirements are to be met 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 be provide 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 line rate for strict priority queuing to be maintained.

Due to the nature of PMU and PDC operations there is some interdependency between latency in the network and the latency of the devices. In particular, 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. Clearly, setting an appropriate value for the waiting time will require information relating to network latency, and in turn the waiting time settings will affect the overall latency of the system. Although attempting to meet these stringent requirements may appear to be a difficult task, there are a few points that help to mitigate the problem. Also, most of these points aid the application of the guidelines outlined in section 2.

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. 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.

IV. TRAFFIC CHARACTERISTICS AND EVALUATION OF FIRM REAL-TIME SYNCHROPHASOR SYSTEMS

The primary traffic sources are the PMUs that produce a single message for each cycle, in synchrony with each other. However, unless the PMUs are connected to the local PDC by a LAN, the PDC will be the first point of contact with the communication infrastructure. To simplify discussion we will focus on the case of PMUs being directly connected to the PDC, as shown in figure 1.

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. Ideally the PDC should be designed to operate and the line rate of the communication link. Once the burst has been sent, there will be no further transmissions until the next cycle. 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. However, as all the messages are created at the same time, and are all bound for the same destination, bursts originating from different PDCs could interact with each other as they pass through the forwarding devices along paths that fan-in to the SPDC, as shown below in figure 4.

For economy of space, and to simplify discussion, the system shown in figure 4 is of a very small scale grid, although the overall link distances considered (71km, 111km and 121km) are not untypical for grids discussed in the literature. A more realistic scale would be in the order of 15 to 25 forwarding devices with paths involving between 3 to 20 hops and a median of between 7 and 10 hops [7]. The propagation delays shown on the links are based on the widely quoted figure for optical fiber transmission of 5 μ s/km.

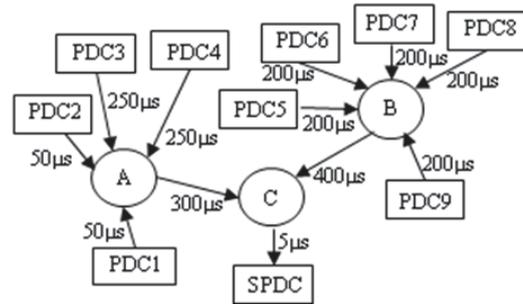


Fig. 4. PMU Message Fan-In to SPDC

Due to the variation in distance and the number of PDCs feeding in to each forwarding node and other factors, there may be a significant difference in the latency between different PDCs and the SPDC. However, since the SPDC has to wait for the whole batch of message to arrive before beginning to process them the maximum delay from the entire set is the important one. 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 in isolation. In turn this evaluation will advise the setting of the SPDCs waiting time. From the perspective of the network the PDC is just a source that periodically creates short burst of packets. Evaluating network latency involves working down through the levels of the fan-in, starting at the PDC level and evaluating the effects that each set of PDCs has on the forwarding node to which it is connected. Using propagation delay and burst lengths as parameters, the convergence and overlapping of bursts arriving at the queue need to be calculated and the effects on the queue, i.e. maximum queue length and busy periods, evaluated. Following this, the output process from the queue, which is related to the queue's busy periods, needs to be derived for use in the next stage down. Once the processing has been completed for all forwarding nodes at that level, the set of derived output process can then be used to continue the process at the next level down, and so on until the SPDC is reached. In the case of deterministic output from the PDCs evaluation can be readily achieved using a combination of arithmetic and basic network calculus [8].

Using the deterministic assumption, we carried out this process to evaluate the model shown in figure 4. The settings were as follows: propagation delays are as shown on the links in figure 4; Each PDC serving 20 PMUs each of which produces 1 message; message length 1000bits; link rates 500Mb/s, for all links; line rate switching devices with queuing at the output link only, and a fixed forwarding latency of 10 μ s per device; and line rate transmission from the PDCs. A link rate of 500Mb/s is relevant, as this is the data rate supported by the fiber optic carriers contained within an OPGW (Optical Ground Wire Systems) cable [9]. These cables provide both protection and communications and can be installed on high voltage pylons. Therefore, it is most likely that they will be used for communication at this level in the grid. The results of this evaluation are shown below in table 1.

TABLE I. LATENCY VALUES IN MICRO SECONDS

PDC No.	1,2	3,4	5,6,7,8,9
Latency[max]	455	685	855
Queue[max]	40	40	160
Queue[total]	40	70	190
Latency[prop]	355	555	605
Latency[fwd]	20	20	20
PMU tx. time	40	40	40

These values are relevant only to this one particular simple example and do not represent a typical case. However, they do serve as a scaled down example that shows the range of difference between the latency of individual paths. Even in this simple example we can see that the difference between the maximum and minimum latencies is quite significant. Also, as we would expect, Propagation delay is by far the dominant factor. However, whilst in the worst case queuing latency accounts for only about 21% of the total delay, it is not insignificant.

The highest overall total latency in this example relates to the set of PMUs 5-7, and this value would be the one used to determine the setting for the SPDC's Waiting-Time parameter. The total latencies for the other two sets of PMUs (1-2 and 3-4) are both significantly less than this value, i.e. 400 μ s and 170 μ s respectively. Given that the QoS target is based on the delay value of the path with the overall maximum latency, then, in this particular case, it is implicit that both of these latency values could be increased to some degree without this QoS parameter being affected. Provided, however, that any such increase does not result in an adverse effect on the existing maximum delay path

We believe that such significant differences between the individual delay paths could be exploited during the evaluation and design stages. For example, the evaluation presented above is based on a Preemptive Strict-Priority Queuing discipline being applied in all forwarding devices. Whilst this particular queuing discipline provides the highest degree of class isolation to real-time class of traffic, it can have an adverse effect on lower priority classes by interrupting the transmission of their packets. Non-Preemptive Strict-Priority Queuing does not allow the transmission of packets to be interrupted and therefore is less harsh on lower priority traffic, although in certain circumstances it can lead to additional delays for real-time messages due to the residual service time of a lower class packet that may already be in transmission. However, these additional delays would most likely be acceptable in the cases presented above. Finally, Rate Based-Priority Queuing offers even better relative-fairness to the lower priority classes but results in a lower forwarding rate for the real-time class. However, this forwarding rate can be set to achieve a particular delay constraint using an appropriate evaluation technique. These possibilities suggest that a three stage evaluation approach could be beneficial. The first stage could be based on further developments of the process we have used to evaluate the example given above, and would provide results in a similar form to those presented in table 1. The second stage would

involve identifying those paths that could tolerate a significant increase in latency, then evaluating the effects of substituting Non-Preemptive Strict-Priority Queuing with either Preemptive Strict-Priority Queuing or Rate Based-Priority Queuing on particular individual forwarding devices. Finally, the third stage would carry out a re-evaluation of the complete system as modified according to the finding of stage 2. The points raised in this paragraph will be taken into consideration during our future work.

In general, we expect that the significance of queuing delay and difference between maximum and minimum path latency may both increase in line with the scale of the system. However, as yet we have not had the opportunity to test the effects of scale. Although the evaluation process for the deterministic case is relatively straight forward, it is somewhat laborious, and time consuming. To overcome this problem we have been developing a utility that provides a degree of automation. A prototype of this utility has very recently been completed and is currently being tested. Furthermore, at this stage our models can only identify the set of PDCs to which the worst case latency applies. Extension of the model into the application level devices will be required to identify the individual PMU source involved. This will be more important in future work in which a probabilistic interpretation of the delay bound, as defined by equation in section 3, will be applied.

We do not consider the process that we have used in this preliminary evaluation 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 part of its analysis, that can also serve as a starting point toward the development of a more comprehensive solutions. Currently, the evaluation process is based on a single layer Message Flow/Queuing Network Model (QNM) abstraction, with the delay effects of communication layering being subsumed into the elements of this single layer abstraction. In the next stage of our work we will develop the flow model further so that it can fully capture the delay path as it passes up, down and across the communication layers. For accurate evaluation all potential sources of delay that are above a certain level of significance will need to be investigated and generally every process along the delay path may need to be considered in greater detail.

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 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.

Currently, we are developing performance models for PDCs with the aim of capturing a number of possible alternative processing structures. In conjunction with this we are investigating the application of a convolutional approach for deriving the distribution of burst duration from the distribution of inter-packet transmission intervals produced by the PDC. Although the convolution process requires the assumption of mutual independence, which may not always be valid in all cases, biased convolution approaches have been developed and used for evaluating Probabilistic Hard Real-Time System [10]. This part of the investigation is in very early stages and development of the PDC model is still ongoing.

V. CONCLUSION AND FUTURE WORK

This paper has addressed the requirements of Time-Critical Wide Area Measurement and Control systems. Systems that are intended to facilitate automatic control in future Smart Grids. It has examined their operation and analyzed their communication traffic characteristics. It has shown that these characteristics are significantly different to those of the current near real-time wide area measurement applications that provide visualization to aid manual grid control. Therefore, it recognizes that they will also have significantly different QoS requirements. In particular the delay targets are an order of magnitude less, and therefore, evaluation will need to be carried out at the micro second level. Furthermore, the communication delay bound ultimately applies to the maximum delay for a set of messages that are created simultaneously at the beginning of a periodic cycle. If the maximum delay can vary between batches, then the delay bound refers to the maximum possible delay for all batches, over the operational period of the system. However, this requirement can be based on a probabilistic interpretation. Also, due to the wide range of difference in latency that can be expected between different paths through the fan-in to the SPDC, it may be possible to concentrate effort on those paths that are nearest to the limit in terms of latency.

The paper has also discussed the problem of evaluating the performance of these time critical system in advance of their deployment, and has presents the first stage in a body of work aimed at developing models and techniques to facilitate the performance evaluation process. It has presented some preliminary results and outlined the direction for the next stage of the investigation.

We have outlined the direction of our investigation for the immediate future and once the PDC models have been developed we will use them to investigate the effects of variable bust duration. These models will also be used to support the development of a theoretical basis for PDC performance compliance testing. Our future work will also need to addresses the problems of non-line-rate forwarding devices. This will require a similar approach to that we using to evaluate PDCs except the internal queuing model will, most likely, be more complex. To develop the internal queuing

model will require detailed analysis of multiple delay paths and the numerous sources of delay contained within the device. However, such information may not be readily available, particularly in the case of propriety equipment and therefore alternative options may need to be considered. Given that one of our main aims is to develop generic and parameterized models for performance evaluation, part of our future work will be to consider the viability of using such models as an alternative to enable the manufacturer to supply the information required for performance evaluation, without having to provide details of the actual implementation. Also, as another alternative, we will address the viability of obtaining the required information through measurement, that could possibly be carried out during compliance testing.

Finally, as the investigation progresses we will need to consider the results of our work within a wider contextual framework. We believe that the paradigms of Software Defined Networking (SDN) and Software Defined Infrastructure (SDI) are appropriate areas to consider for this purpose. Mainly due to their ability to provide isolation between the function of flow and control. Also, the systems that our models are intended to support are closed loop control system that operate over a communications network. Therefore, we will also need to consider the implications of our work within the context of Networked Control Systems (NCS) research.

REFERENCES

- [1] Bakken, D.E., Bose, A., Hauser, C.H., Whitehead, D.E., Zweigle, G.C., "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] R. B. Bobba, J. Dagle, E. Heine, H. Khurana, W. H. Sanders, P. Sauer, and T. Yardley, "Enhancing Grid Measurements", IEEE power & energy magazine, PP 67-72, January/February 2012.
- [4] 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. School of Engineering, University of Aberdeen, Aberdeen AB24 3UE, UK
- [5] D. M. Laverty, R. J. Best, P. Brogan, I. A. Khatib, L. Vanfretti, and D. J. Morrow, "The OpenPMU Platform for Open-Source Phasor Measurements", IEEE Transactions on Instrumentation and Measurement, Vol. 62, No. 4, April 2013.
- [6] US Department of Energy "Factors affecting PMU installation Cost", Oct 2014. Smartgrid.Gov
- [7] P. Kansal and A. Bose, "Bandwidth and Latency Requirements for Smart Transmission Grid Applications", IEEE Transactions on Smart Grid, Vol. 3, No. 3, pp 1344-1352, September, 2012.
- [8] W.M. Sofack, M. Boyer, "Non Preemptive Static Priority with Network Calculus," ETFA 2011, pp. 1 9 Sept. 2011
- [9] 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, 10-13 June, 2013.
- [10] G. Bernat A. Colin S. M. Petters, "WCET Analysis of Probabilistic Hard Real-Time Systems", 23rd IEEE Real-Time Systems Symposium, RTSS 2002., 279 - 288