

Smart Generation and Transmission With Coherent, Real-Time Data

David E. Bakken, Anjan Bose, Carl H. Hauser
School of Electrical Engineering and Computer Science
Washington State University
Pullman, Washington, USA
{bakken, bose, hauser}@eecs.wsu.edu

David E. Whitehead and Gregory C. Zweigle
Schweitzer Engineering Laboratories, Inc.
Pullman, Washington, USA
{davewh, gregzw}@selinc.com

Abstract—In recent years, much of the discussion involving “smart grids” has implicitly involved only the distribution side, notably advanced metering. However, today’s electric systems have many challenges that also involve the rest of the system. An enabling technology for improving the power system, which has emerged in recent years, is the ability to measure coherent, real-time data. In this paper, we describe major challenges facing electrical generation and transmission today that availability of these measurements can help address. We overview applications using coherent, real-time measurements that are in use today or proposed by researchers. Specifically, we describe, normalize, and then quantitatively compare key factors for these power applications that influence how the delivery system should be planned, implemented, and managed. These factors include whether a person or computer is in the loop and (for both inputs and outputs) latency, rate, criticality, quantity, and geographic scope. From this, we abstract the baseline communications requirements of a data delivery system supporting these applications and suggest implementation guidelines to achieve them. Finally, we overview the state of the art in the supporting computer science areas of overlay networking and distributed computing (including middleware) and analyze gaps in commercial middleware products, utility standards, and issues that limit low-level network protocols from meeting these requirements when used in isolation.

Keywords—*synchrophasor; middleware; smart grid; bulk power system*

I. INTRODUCTION

Today’s large power grids evolved during the middle of the twentieth century, as utilities integrated into larger power systems in order to improve reliability. In such a structure, an entire grid, such as the USA Eastern Grid, operates at the same frequency, and supply and demand must be balanced in real-time across the entire power system. Reliability limitations of large grid systems became apparent in part due to a large blackout in the northeastern USA and southeastern Canada in 1965. As a result of this blackout, it was realized that utilities need to have better visibility into their operations beyond what can be sensed in a control center. This led to an emphasis in implementing SCADA (Supervisory Control and Data Acquisition) in the most critical substations. Nevertheless, large grids often have many subparts with only basic communications between these subparts. Now, at the beginning of the twenty-first century, new measurement and communications technologies are creating possibilities for

monitoring and control of the power grid that were infeasible with SCADA systems.

Since the 1960s, visualization and situational awareness technologies have improved and been augmented by newer networking technologies. However, the need to continue visibility improvements has been demonstrated by recent power disturbances cascading into large blackouts, for example, the ones in the USA/Canada and Italy/Switzerland in 2003 [1]. A single event, such as a transmission line fault, can start a chain reaction an hour or two prior to an actual blackout. When situational awareness is low, and no one has the entire “big picture” of the power system, the significance of one event cannot be sufficiently understood by any person or computer in order to take action to avoid the blackout. Such informational disconnects are possible given that, for example, the power systems in North America have 3500 participants that can affect system stability [2].

The electric power system depends on control and protection schemes that prevent the system from reaching instability or collapse. Again, during the USA/Canada 2003 blackout, once the last contingency occurred, the cascading over several US states and one Canadian province happened too quickly for operator intervention. Under these circumstances, the only way to avoid such cascading is to use fast control or system protection schemes to isolate impacted areas and/or adjust some controllable values.

The continuing need to improve the communications infrastructure, increase the situational awareness of operators [3], and improve wide-area power system control is driven in part by fundamental challenges facing electric power systems: demand for high reliability, efficient operation, and the evolution toward noncarbon-based energy sources.

The reliability of the power system is dependent on its ability to deliver electric power from generation to load without disruption. Thus, the grid must be able to withstand minor and major disturbances with minimal customer impact. Interruption of electricity supply is not only inconvenient to the user but it affects the overall economy (productivity) of the region.

Reliability is enhanced by adding redundancy and providing enough margin for the power system load. On the other hand, operating the grid at lower levels than its limits, introduces inefficiency because the transmission system is not fully used. There is always some compromise between

reliability and efficiency, and within the bounds of this constraint, both of them have to be optimized.

The efficiency of a bulk power system is dependent on its ability to minimize the cost of generation and delivery, which is facilitated by the transfer of large amounts of power using the most efficient generation sources while incurring the least losses in the transmission system. Because transmission lines have limits, maximizing efficiency requires a constrained and nonlinear optimization calculation, which is done in the energy market as well as in real time.

Evolving to a noncarbon-based electrical infrastructure will require integrating large amounts of nontraditional generation sources (such as wind and solar) that behave differently from existing generation (such as nuclear, coal, and natural gas) [4], [5]. The output of wind and solar generation is difficult to control because it depends on local weather conditions. The power system will have to use these intermittent sources of generation without compromising reliability and efficiency.

Meeting these challenges involves modernizing not just the power system, but also its data delivery infrastructure. The best approach is to holistically and simultaneously consider the dynamics of power systems and their data delivery infrastructure—their steady states and those perturbed by a power contingency or failure, or a cyberattack involving the data delivery infrastructure. One key recent technology involves sensor data that are given microsecond-accurate timestamps and then delivered in real-time to give a coherent picture of a system for operators—and also for closed-loop control and broader protection.

In this paper, we offer a view of the power system involving both applied electric power engineering and computer science. We describe a wide range of applications using coherent, real-time data in order to mitigate these fundamental problems. We normalize and summarize the system communications requirements, including not just traditional quality of service (QoS) metrics, such as latency and data rate but also broader metrics, which we call “QoS+” that include geographic scope, criticality, and amount of data. Next, we describe how these QoS+ metrics must necessarily impact the power system’s data delivery system at the overlay network level, including absolute requirements and recommended implementation guidelines. As part of this discussion, we also compare how existing overlay network-level technologies, middleware, and power system communications protocols map onto those requirements and guidelines. This paper does not address the lower network layers such as the physical, link, network, and transport layers, other than showing the importance of an overlay network above them and how it enables the complete network to meet the wide-area system requirements. Finally, we overview the decade-long research into the GridStat data delivery system and the emerging NASPInet concept that GridStat has influenced.

II. SOLUTIONS FOR ENHANCING GENERATION AND TRANSMISSION BASED ON COHERENT REAL-TIME DATA

Time-synchronized measurement and control is an enabling technology to help solve power system challenges.

Devices using this technology are becoming a standard part of the power system and provide microsecond time accuracy using Global Positioning System (GPS)-based clocks. These measurements have existed for over a decade within devices such as protective relays and other IEDs, which combine precise measurements of currents and voltages with accurate time recording. Synchrophasors are a common name for these measurements. They represent both the magnitude and phase angle of voltage or current waveform at a particular time, synchronized to a common reference such as a GPS clock [6]–[9]. However, the application of time-synchronized data applies beyond voltage and current signals. Accurate time-stamping of any electric power system measurement, such as breaker status, active power, reactive power, and weather effects on renewable generation, provides benefits for reliability, efficiency, and economics.

A decade ago, time-synchronized measurements were found only in stand-alone instruments called phasor measurement units (PMUs). Today, such measurements are also collected from meters, protective relays, and fault recorders, which dramatically lowers the cost of implementing synchrophasor-based control and protection strategies. Station phasor data concentrators (PDCs), which gather time-synchronized measurements from several sources within a substation, and distributed synchrophasor control devices are important new system components, providing distributed aggregation, archiving, control, and protection functions. Furthermore, new communications architectures, which include in-network data concentration, real-time distribution, and fast fault recovery provide an infrastructure with the necessary high reliability.

This section outlines a few applications that can bring increased reliability, efficiency, and stability to the entire power system. Examples are given in several broad areas: state estimation, control, protection, situational awareness, and event analysis. For each application, we explain how it is being improved with time-synchronized measurements, and then, in anticipation of the second half of this paper, we list specific communications requirements. These requirements provide the basis for subsequent communications system analysis. The following review articles provide other synchrophasor applications [8], [10]–[13].

A. State Estimation and Direct State Calculation

Knowing the system state in real time is an important first step for reliable control of the power system [14]. In the electric grid, the state of the system is defined as the voltage magnitude and angle at every bus in the system. Schweppe introduced the first system for estimating this state [15]. In this approach, still in dominant use today, the state is estimated from voltage magnitudes (no angles) and power flow measurements using iterative, nonlinear algorithms. While this has provided many benefits during the past forty years, the model nonlinearities and nonsynchronized measurements in this traditional state estimator cause limitations in computation time, solution errors, and convergence.

Fast state calculation is increasingly important for the quick response time requirements of wide-area protection and

control loops. Because time-synchronized measurements include both angles and magnitudes, the state is directly measured [16]. No additional processing is necessary. Furthermore, if some bus locations do not have PMUs installed, the only calculation required is a linear estimator [17], which does not have the time skew, convergence, and computation time issues of the traditional state estimator.

However, the new time-synchronized measurements are not yet widely deployed to measure the angle and magnitude at every bus in the system. Therefore, the traditional nonlinear state estimation will still be necessary in some systems.

State estimators must keep track of dynamic power system topology in order to correctly estimate the system state. Using traditional methods, all measurements are not necessarily taken at the same instant in time. This problem is minimized in the existing state estimators because the measured quantities they use, voltage magnitude and power, typically change slowly with time. However, they occasionally do experience more rapid changes and the result is that operators often must “suspend belief” in these quantities for a brief time after a change occurs because information collected in this asynchronous fashion creates operator display incoherency. For example, when a breaker is opened, the current through the breaker is interrupted and should read as zero amperes. However, if the breaker-open information is received and displayed before the current information is received and displayed, then the operator sees an open breaker with a nonzero current. It takes several SCADA polls for the discrepancy to be resolved, and in the meantime, the operator is not completely sure if the information is indicating a breaker failure. With time-synchronized measurements, the precise timestamps enable aligning all measurements, including contacts, disconnect switches, and tap changer values, so accurate system states can be calculated.

As a baseline case, it is interesting to note that a very simple state measurement application is immediately available with synchrophasor measurements and does not require any special communications infrastructure. One common wiring problem that is difficult to detect during commissioning is rolled power system phases. Consider the case where the VA source is wired to the VB terminals, VB to VC, and VC to VA. Using a local PMU and simple terminal connection, a technician can immediately check for this condition because signal phases are referenced to a common time standard [18]. Thus, synchrophasor capability in IEDs instantly provides these basic power system improvements to the power engineer, even without additional software applications, energy management systems, or communications infrastructure.

Time-synchronized measurements also create the capability to measure the state within substations and then share between localized regions. Overall system state calculation is then a matter of aggregating and reconciling the local measurements. Fig. 1 shows an example of how local coherent measurements improve system reliability [19].

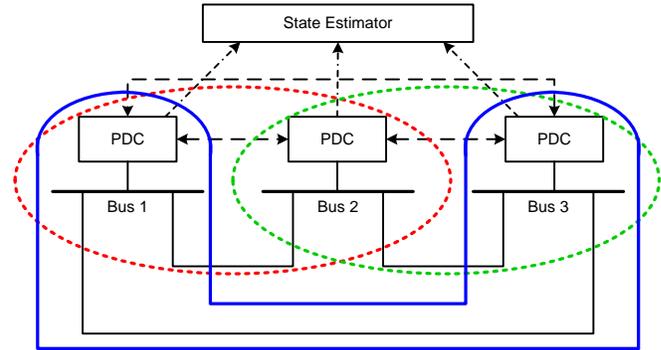


Fig. 1. Distributed Peer-to-Peer Communications Improve System Reliability

For simplicity, the PDCs in Fig. 1 are shown connected directly to the power system buses. In most systems, the PDC connects through a PMU to the bus. The PDC represents a local state calculation device; other devices are available that perform this function. Each substation PDC collects the voltages, currents, associated phase angles, and electrical topologies of the system as required by the state calculation engine in the PDC. The data are also exchanged between the PDCs so that the state is refined based on measurements from adjacent substations. The data exchange provides redundant communications paths to a state estimator to prevent lost data, should the primary communications channel be temporarily lost. If direct communications are interrupted, an adjacent PDC forwards the data.

Fig. 2 shows a two-level state estimator [20] that uses synchrophasor capabilities across a wide area. This estimator simplifies the total state estimation process by detecting and correcting topology and data errors early in the estimation process. It can also greatly lower the quantity of data sent to a control center. This particular scheme uses only two levels, but there is no inherent reason that similar techniques could not be used for more levels of a hierarchy, e.g., substation, utility subregion, utility, system operator, and transmission operator.

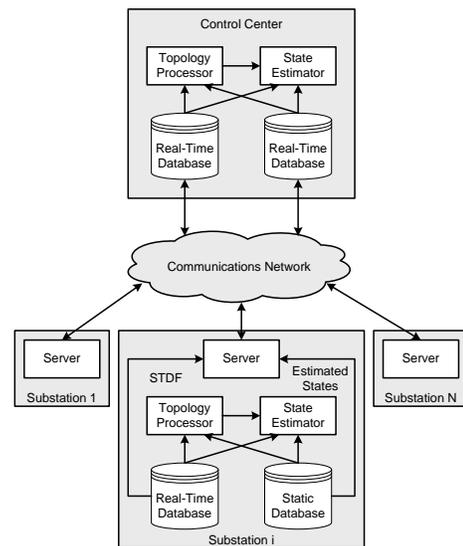


Fig. 2. Two-Level State Estimator Uses Synchrophasor Capabilities to Simplify the State Estimation Process

The communications quality of service (QoS) requirements of these state measurement and calculation schemes are dependent on the application using the state measurement, because state measurement can be the first step of wide-area protection and control schemes [21]. Today, most state estimation implementations are too slow for these schemes because of the nonlinear algorithm required when time-synchronized phase angles are not available. In these cases, the only options are either local measurements or relatively stable measurements such as power flows. With the direct state measurement made possible with synchrophasors in widely distributed power system IEDs, the measurement values can be applied immediately to the protection and control algorithms, thus enabling the use of more data that cover a wider geographic area. This enables new protection and control schemes. Meanwhile, for visualization and off-line analysis applications, the latency for these inputs, is fairly forgiving; tenths of seconds or even seconds is adequate, and a rate of a few hertz or less suffices.

B. Distributed Wide-Area Control

Power system control needs to be improved because the system is becoming more stressed each year as increased demand and supply outstrips the addition of new long-distance transmission capabilities; there are more “miles times megawatts” being travelled each year. Also, renewable energy sources are more variable, and their effect on the power system’s stability is less known than on-demand sources such as hydroelectric, coal, and nuclear with which operators and planners have greater experience. This variability can be mitigated by moving from slower operator control to use of faster algorithms with closed-loop feedback control.

As an example of a wide-area control solution, consider how Southern California Edison has applied synchrophasors for wide-area dynamic voltage control [22]. The purpose of the system is to measure and control a voltage that is hundreds of miles away from the control location. The control location consists of a static VAR compensator (SVC). The SVC adjusts its local voltage, and the remote voltage signal ensures that the voltage at a remote location stays within its required limits.

The total measurement and communications latency requirement for this system is one second. This requirement was not possible to achieve with the existing SCADA system because of the slow and irregular update rate—most systems are as slow as one update every several seconds. However, the requirement was easily achieved by collecting readily available streaming, uniformly sampled, time-synchronized phasor measurements that update up to 60 times per second. A generic architecture for a wide-area control regime is depicted in Fig. 3 [23].

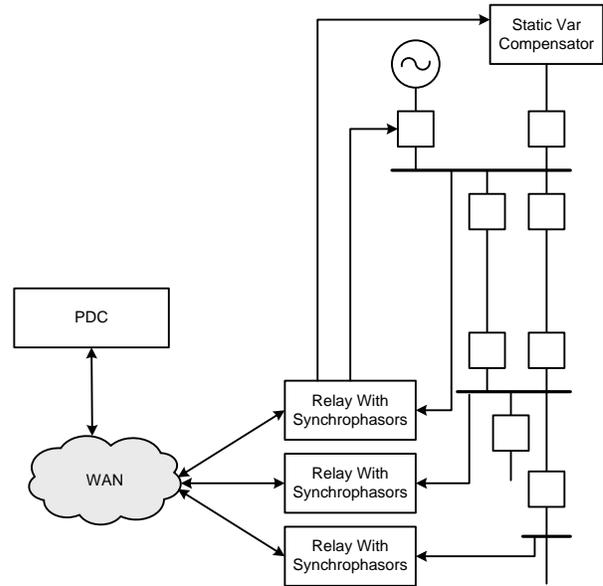


Fig. 3. Distributed Wide-Area Control Solution

Another control application is based on measuring power system off-nominal inter-area oscillations. These oscillations, often called power system modes, are caused by interactions between various mechanical and control systems coupled through long distance power lines. System disturbances, such as generation shedding or line tripping, can excite these oscillations, which may become more pronounced when wind generation is added to the power system [24]. When oscillations are well damped, the system returns to a stable state after the disturbance; however, negatively damped oscillations result in instability. Clearly the power system is not intentionally designed to trigger unstable operation conditions; it is designed with large stability margins. The system topology, however, can change in unexpected ways during a disturbance, which can lead to an unstable system. Because of the power system size, it is difficult to predict all possible topologies, parameters, and associated modes. However, the uniform sampling rate of synchrophasor measuring devices enables the ability to directly calculate the frequency, magnitude, and damping factor of each power system mode in real-time. Then, if an oscillation is detected that is not sufficiently damped, an automated control loop takes action, e.g., sheds load, to bring the system back to a stable equilibrium [25]. Another approach uses a power oscillation controller, which damps oscillations with existing control devices, such as flexible ac transmission system components or a power system stabilizer [26]–[28]. With wide-area time-synchronized information, the oscillation controller inputs are not constrained by geography and the best measurement locations can be selected for the controller input signals.

The output of wide-area control is often an actuation signal. With time-synchronized devices, it is now possible to synchronize the precise timing of these control signals. Synchronizing the control action can reduce variability and help keep the system stable. Fig. 4 shows an example of a system, which demonstrates this synchronized control method [29].

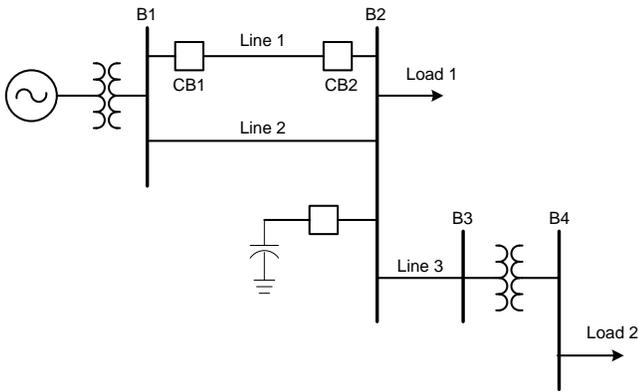


Fig. 4. Power System Model Demonstrates Synchronous Distributed Control

In Fig. 4, Lines 1 and 2 are part of the transmission network. Line 3 connects the transmission and distribution networks. Bus B4 is a distribution bus. The transformer between buses B3 and B4 is a mechanical on-load tap change transformer. As an example of the benefits of synchronized control signals, consider the case of a line removal. To remove a line from service, first an operator sends a command to open breakers CB1 and CB2; this causes a decrease in the voltage at Bus B2 due to the increased impedance from the generator through the remaining Line 2. As a result of this voltage decline, the controller at the transformer between B3 and B4 taps the transformer to restore the distribution voltage to its target levels. If the transmission voltage at Bus B2 decays to a value below the desired minimum, the operators may insert the parallel capacitor into the system. This raises the transmission bus voltage but then requires the transformer to tap back down in order to avoid exceeding the distribution bus target voltage levels. Fig. 5 illustrates the system response to these changes. These sequential operations result in unnecessary stress on power system components, and they can also contribute to a more broadly cascading event if they happen at an inopportune time.

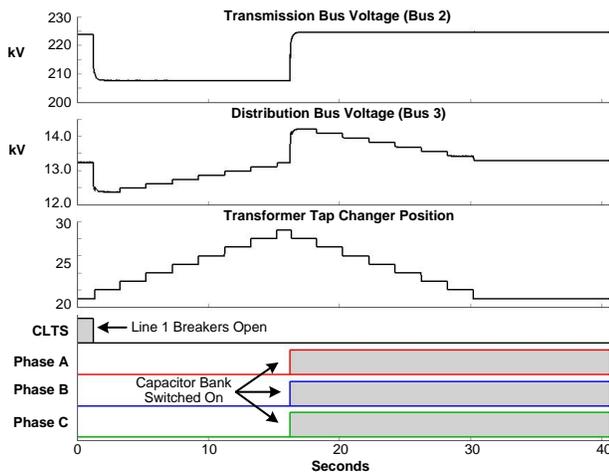


Fig. 5. Sequential Operations Disturb System Voltages and Currents and Place Unnecessary Stress on the Power System

Using time-synchronized measurements, these sequential operations are each synchronized to execute at exactly the same moment. First, the operator selects an appropriate set of

commands (or recipe) to accomplish all of the desired changes. The commands are then sent to a coordinator (such as a PDC or automation controller) at each involved substation. The PDCs send appropriate subsets of the command list to IEDs and confirm that they are in states appropriate for carrying out the commands. After receiving confirmation from each IED that the sequences of commands are ready to run, the PDC indicates to the operator that the system is ready for initiation. The operator validates that all components are ready, no cybersecurity alarms have been received, and the change is still desired. The operator then arms the system and sends the start time to the PDC. The PDC and IEDs execute each command at a preprogrammed instant in time. Fig. 6 shows a reduction in transients, which improves reliability and leaves additional margin for the dynamics of renewable energy sources.

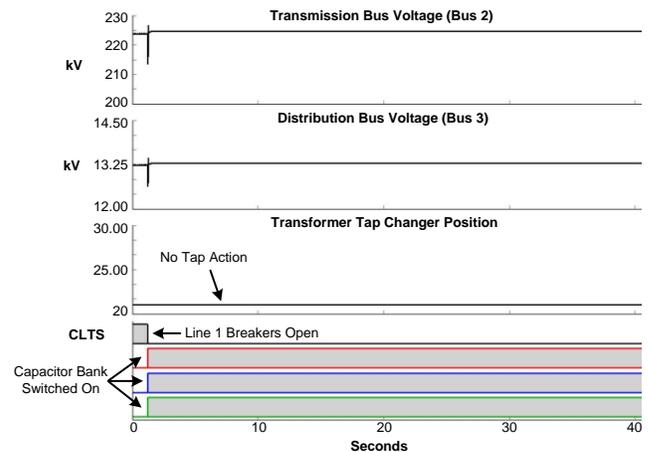


Fig. 6. Time-Synchronized Changes Cause Minimal Disturbance to System Voltages and Currents

The synchronized control system enables verification that the system operations are intended (that is, not due to a cybersecurity breach) and are suitable for the given system state. The distributed synchrophasor control device requests control validation from the system operations center or source of the synchronized commands. A local logic engine can analyze the requested operation and determine, for example, if opening a circuit breaker will result in a stability problem such as an unacceptable voltage drop or collapse. The system can also alarm to alert the operator when a new series of controls is initiated. Only after validating the commands will the operator arm the system to execute at the desired time.

Now, consider the communications system performance requirements for these and other control-loop schemes based on time-synchronized phasors. The allowable latencies for control inputs vary from roughly 100 ms to a few seconds, depending on the application. Various schemes have been proposed to compensate for excessive latencies [30]. The required data rate for inputs varies depending on the application. Voltage control inputs can be as slow as 1 sample per second, while oscillation control may require 60 samples per second. The data delivery reliability is critical. However, the reality of less than 100 percent message delivery is tolerable when mitigated by anticipation and compensation in the control algorithm. If input data are missing, there is

typically no need to retransmit because it is better to have the most recent data than to act based on historical data. The geography of inputs can vary widely depending on the control scheme.

The output is a control signal that is sent to power system devices such as a voltage regulator, reactive power controller, or load controller. The latency requirements are similar to those for the inputs. The rate of sending control output signals may be slower than the inputs, because in certain applications, a control signal is only needed when it changes. However, for continuous control outputs, a lower output rate makes the control loop slower, so increasing the rate offers benefits in some configurations. The quantity of the output signals is not large, though it obviously goes up with increased output rates. The geographic scope of the outputs is similar to the inputs.

C. Protection

Wide-area system protection applications are another class of applications where the implementation is facilitated by ability to communicate synchrophasor data across the grid. A system integrity protection scheme (SIPS), also known as a remedial action scheme (RAS), provides the next level of protection after the relays, which respond to local power system emergencies [31]. The purpose of local protection is to quickly remove the disturbance and minimize equipment damage. The purpose of system integrity protection is to ensure that the complete power system remains in a viable state after the local protection has operated. Wide-area distributed signals improve the stability of these schemes. Transient stability is a problem with many power systems in which the transfer limit on some transmission corridors is affected by the fact that short circuits make the system unstable. Actions of various kinds—shedding load and/or generation—are used to mitigate these instabilities, thus allowing higher limits on the transmission corridor. The main difficulty is that the instability occurs quite fast and any action must take place within 50 to 250 ms to maintain stability. Such fast actions are not possible without having a reliable, high bandwidth, low latency communications system.

One class of SIPS is contingency-based, where the scheme responds after a predefined event occurs, such as a topology change due to a breaker opening. Another SIPS methodology is based on analog quantities such as power flow, where the scheme responds if the power exceeds or drops below a threshold. In this case the system may shed load if the generation is unable to supply the required power. In both cases, the SIPS executes a preplanned mitigation strategy developed by a previous analysis of the power system configuration and performance.

Many applications of wide-area system protection can move toward directly measuring the system state and acting based on that information [21], [32]. A specific example of a synchrophasor-based SIPS is the Comisión Federal de Electricidad (CFE; México) automatic generation shedding scheme. A simplified diagram of the scheme is shown in Fig. 7 [33].

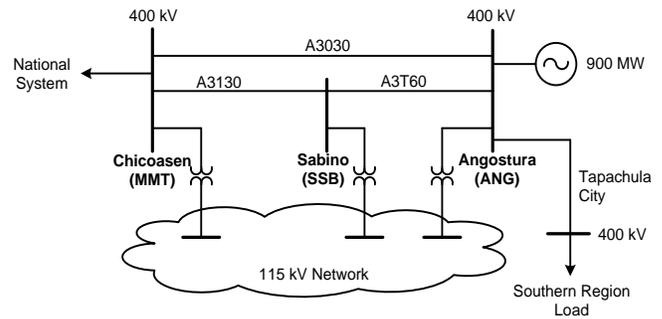


Fig. 7. CFE Automatic Generation Shedding Scheme Uses Synchrophasors to Prevent Instability

If the transmission lines between the generation at Angostura and the load at Chicoasen are lost, the system can become unstable. Initially it might seem that the easiest solution is contingency-based by monitoring the circuit breakers connected to the line. However, the resulting scheme becomes complicated because of the many circuit breakers that must be monitored. A simpler solution uses relays with time-synchronized phasors at each end of the line to measure the angle difference and compare the difference against a threshold.

Modeling using a Real Time Digital Simulator (RTDS) shows that losing both transmission lines results in an initial angle difference of 14 degrees; enough to cause system instability. Fig. 8 shows how the angle increases without constraint after the loss of both lines. A single-line fault results in a difference of less than 7 degrees and does not cause stability problems. As a result of these studies, a threshold difference of 10 degrees between measurements at Angostura and Chicoasen was selected for the synchrophasor-based scheme.

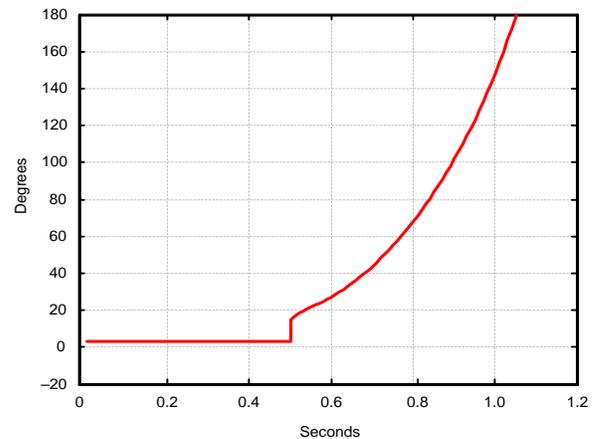


Fig. 8. Angular Difference for a Double Contingency Condition

Relays with PMU capabilities were placed at Angostura and Chicoasen. Each relay measures its local bus voltage angle. The Chicoasen relay sends its synchrophasor data to the Angostura relay. The Angostura relay time-aligns and then compares its local angle with the remote phase angle from Chicoasen. If the Angostura relay detects that the angle difference exceeds the maximum threshold of 10 degrees, it will trip generation to prevent system instability.

Fig. 9 shows the result of the synchrophasor system responding to a double-line loss and tripping the generation after 100 ms. The system remains stable.

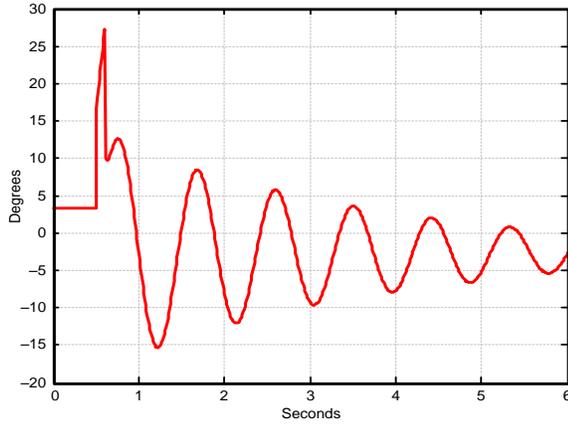


Fig. 9. Angular Difference for a Double Contingency at 500 ms and Tripping of the SIPS 100 ms Later

Synchrophasor technology has also been applied in islanding control [34] and anti-islanding applications. Presently the IEEE 1547 Standard, “Interconnecting Distributed Resources With Electric Power Systems,” [35] specifies that a distributed generation source must disconnect from a locally islanded system within two seconds. Such a requirement is important for safety reasons, quality of power, and out-of-phase reclosing avoidance. One approach to anti-islanding uses local voltage or frequency information to determine if the frequency or voltage magnitude is outside thresholds set by planning and engineering. However, if the power mismatch between the islanded source and the local load is small, it is difficult to detect an island and respond quickly using voltage and frequency information. Breaker status is another source of information that indicates when the system is islanded, but this approach can require many communications channels, causing overall poor reliability [36]. Having the inverter continuously attempt to shift its local frequency is another method that is used to indicate an island. This method becomes less effective for high photovoltaic (PV) penetration levels [37].

A synchrophasor-based anti-islanding system helps alleviate disadvantages of existing approaches by making the implementation simpler. Furthermore, as the density of renewable energy sources increases, forced islanding reduces power system reliability. IEEE 1547 also requires disconnecting for sagging voltage under high demand. With a small amount of generation, this requirement is reasonable, but disconnecting a large number of solar generators can cause the low-voltage condition to accelerate. In the future, it will be important to keep these sources online during certain power system events because the large quantity of generated power can help keep the system stable. Synchrophasors enable a wide-area view of the system and therefore enable solutions that can keep distributed generation online during transient conditions.

Using synchrophasor technology, islanding control for a PV system is set up as shown in Fig. 10 [29]. The relays include PMU capabilities and are connected by a wireless link.

The solar PV panel is connected through a breaker to the distribution power system and then to the bulk power system. Both relays acquire voltage phasor measurements locally. Relay 1 then sends the synchrophasor values to Relay 2 at a rate of 60 messages per second. Relay 2 receives the remote synchrophasor values and calculates the angle differences between the remote and local values.

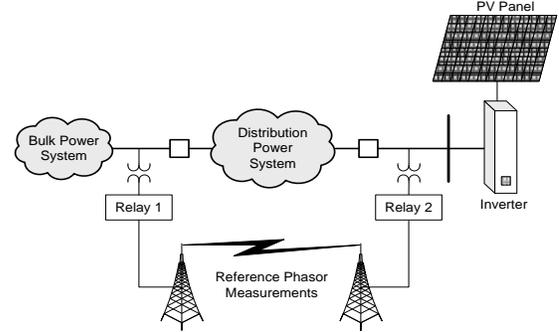


Fig. 10. Anti-Islanding Scheme Using Relays, Synchrophasors, and an Inverter

The angle difference between the relays is defined as δ_k in (1). The rate of change of δ_k is the relative slip frequency, S_k in (2), where MRATE is the synchrophasor message rate. The change of slip frequency with respect to time, measures the acceleration between the two terminals. This value is defined as A_k in (3).

$$\delta_k = \angle V_k^{(1)} - \angle V_k^{(2)} \quad (1)$$

$$S_k = \delta_k - \delta_{k-1} \text{ MRATE} \quad (2)$$

$$A_k = S_k - S_{k-1} \text{ MRATE} \quad (3)$$

Combining slip (S_k) and acceleration (A_k) results in the island detection phase diagram shown in Fig. 11 [38]. In steady state, the slip and acceleration are at the origin. When an island condition occurs, slip and acceleration are possible, and either can push the phase into the Islanded region of the phase diagram. Normally, the system is indicated as connected when the slip (2) and acceleration (3) are within the Connected region of the diagram.

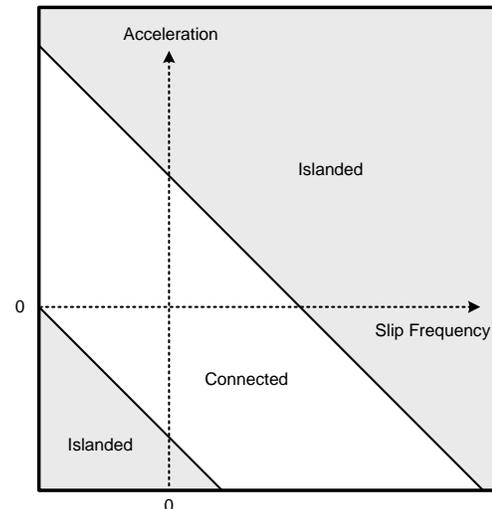


Fig. 11. Islanding and Connectedness Based on Acceleration and Slip Frequency

The communications requirements for wide-area system protection are challenging. The input data rate is the highest of all the applications considered, and the latency must be very low. The criticality of its inputs (and outputs) is extremely high. For example, a system integrity protection system might be installed in order to transfer more energy over a line than it can handle under all contingencies. Therefore, if a contingency happens, it will have to respond by curtailing generation or load. If the protection scheme fails to operate, the contingency can cascade into a blackout [39]. As a specific example of communications requirements, the CFE system required a data exchange of 20 messages per second in order to meet the operating time of 100 ms. This message rate was met by using a 19,200 baud fiber-optic serial connection between the relays.

Outputs from a wide-area protection algorithm are a condition-based control signal to initiate any of a number of actions to compensate for the contingency, e.g., tripping a breaker, generator, or load. The outputs should be delivered with very low latencies. The criticality of the control actions is high, though the quantity is low. The output control signals sometimes are delivered over less distance than the inputs when the logic is located close to the power system element that it is controlling.

D. Wide-Area Situational Awareness

Operator displays are the primary window by which engineers monitor the operational state of the electric power system. Most existing operator displays update slowly based on data collected from a SCADA system every few seconds. These data are insufficient to reveal some crucial dynamic phenomena, such as oscillations, that can indicate progress toward undesirable operating conditions. With so much new renewable generation being connected to the power system, it is difficult to analyze the power system in sufficient detail to predict some of these oscillations, so detecting them when they occur is crucial. Many oscillations have such a high frequency that they are not detectable with the slowly updating SCADA data. Presenting operators with results of analysis based on synchrophasor measurements made at much higher rates offers a remedy for this. Many systems have been described in the literature [40]–[45]. New tools based on wide-area information are becoming available to help operators determine abnormal conditions and either assist in selecting the appropriate response or automatically perform a control action. These tools include dynamic security assessment (DSA) [46], mode meters [47], and voltage collapse detection [48], [49].

The communications latency constraints for wide-area visualization are not strict because the data arrive coherently; updates every few seconds are sufficient and displays can lag by a few seconds. However, the latency becomes more important when data are used for more than visualization, such as in security or oscillation monitoring applications that predict whether the power system is moving into an unstable state. The quantity of data gathered with synchrophasor measurements is large because of the high sampling rate and increased number of measurement points across an entire utility or ISO (wide area). Note the difference to existing

SCADA systems. The existing SCADA systems update every few seconds with a single instance of measurements from strategic substations, so the sample rate is equal to the update rate. The wide-area scheme based on synchrophasors also might update every few seconds but with a sequence of measurements. The sample rate of the measurements might be 30, 60, or 120 samples per second, arriving as a set of values at the slower message rate.

For visualization, it is not always critical that every measurement arrive. If there is a gap in communications, the systems are designed to buffer and retransmit critical information during a fault or other problem. This kind of data transfer is different from some uses of sensor updates, where each update may be critical to deliver.

E. Post-Event Analysis

A system disturbance in the power grid can lead to an outage at some scale. Utilities are required to save key sensor data in a database so regulatory authorities, such as NERC in North America, can ascertain the root cause of the problem. Post-event data transfer, then, involves transferring key related database entries for an event. The data messages of the transfer need not have any kind of latency guarantees, because the post-event analysis will be conducted offline. However, it is important to be able to transfer a reasonable amount of event data within a few hours or at most a few days. If the size of the required dataset is too large, and the communications system is not adequately designed, it may not be possible to do this without interfering with important real-time data. One post-event application is model validation. Gathering of archived data has no sub-second data delivery requirements for a given sensor update, but it is important that the data transfers in this class happen in a reasonably predictable amount of time and that all of the data be transferred, requiring a reliable data transfer mechanism.

III. POWER APPLICATION REQUIREMENTS MAPPED TO DATA DELIVERY SERVICE REQUIREMENTS

The power applications described in the previous section have a wide range of data delivery requirements in many dimensions. In this section, we summarize the requirements of communicating synchrophasor data to show the breadth of the requirement space and introduce the idea of a wide-area measurement system for data delivery (henceforth WAMS-DD). WAMS-DD middleware lies between the lower network layers and the power system applications.

A. Normalizing WAMS-DD QoS+ Parameters

Table 1 presents WAMS-DD requirements in a qualitative form, normalized to indicate the level of difficulty, where 5 means most difficult and 1 means least challenging to provide. This methodology enables comparison of different properties that have very different ranges, to get a sense of the wide ranges of difficulty or easiness involved for different power applications. It is important to note that a given application will not have all of its values in the same row; some requirements will be quite stringent (e.g., ultra-low latency) while others may be more forgiving (e.g., low volume of traffic) for a given application.

TABLE 1: NORMALIZED VALUES OF QoS+ PARAMETERS

Difficulty (5 hardest)	Latency (ms)	Rate (Hz)	Criticality	Quantity	Geography	Deadline (for bulk traffic)
5	5–20	240–720+	Ultra	Very High	Across grid or multiple ISOs	<5 seconds
4	20–50	120–240	High	High	Within an ISO/ RTO	1 minute
3	50–100	30–120	Medium	Medium	Between a few utilities	1 hour
2	100–1000	1–30	Low	Low	Within a single utility	1 day
1	>1000	<1	Very Low	Very Low (serial)	Within a substation	>1 day

Included in Table 1 are representative values of the following data delivery requirements.

Latency: What latency is required for the delivery?

Rate: At what message rate does/should the input be delivered, both now and in the future?

Criticality: How critical is this input [50]? i.e., what is the severity of the consequences if data are not delivered for a short period of time?

Quantity: How much information needs to be delivered?

Geography: How far do the data have to travel?

Deadline: For bulk data transfer (defined shortly), when does the transfer have to be completed?

As noted earlier, some of these parameters are called *quality of service (QoS)* by networking researchers. We denote this entire collection, then, as *QoS+* to indicate that it includes other information needed in communications system design. *QoS+* also refers to cybersecurity issues, though these are beyond the scope of this paper.

The WAMS-DD requirements include the entire communications system. Lower-layer protocols, such as the physical layer, network layer, and transport layer, must be selected so that they meet the WAMS-DD requirements. However, these lower-layer protocols will also include other requirements and capabilities that are unique to them. Addressing their functionality is outside the scope of this paper.

B. Comparing WAMS-DD Parameters for Selected Power Applications

We now use the numerical difficulty values identified in Table 1 to summarize the *QoS+* requirements for the power applications described in Section II. This is depicted in Table 2. The columns of this table are the different applications. The rows are the *QoS+* attributes of the application’s data delivery requirements along with three other kinds of information about the application:

Loop Entity: Where does the application’s output go, a person (**P**); or a computer (**C**)?

Inputs and Outputs: What kind of data delivery is the input or output, streaming sensor updates (**SS**), condition-based (**Co**), i.e., aperiodic events triggered by some condition, or bulk data transfer (**Bu**)?

Note that the inputs and outputs for a given application can be different. For example, an application can take in **SS** updates but only emit an output when those inputs show a certain condition (**Co**). Also note that **Co** and **Bu** inputs and

outputs do not have a delivery rate and that a **Bu** input or output does not have a required latency, which in this table represents a per-message guarantee. **Bu** inputs and outputs are also the only kinds that have a “soft” deadline.

NASPinet Class: What service class is this kind of traffic, see Section IV D 4.

It is crucial to note that the requirements of even this small set of applications have great diversity. This means that the data delivery requirements are very broad, and many different kinds of traffic have to be managed in order for each application to receive its required delivery guarantees. This is exactly the opposite of “one size fits all” regarding data delivery. Further, we note that power system dynamics can be affected by data delivery dynamics [51], [52].

We now examine what these data delivery requirements are in greater detail, along with issues involved with implementing them.

IV. COHERENT REAL-TIME DATA DELIVERY ENABLING THESE APPLICATIONS

Data delivery in the power system today can be improved by reducing the use of hard-coded protocols, developing more reusable systems, and providing real end-to-end performance guarantees. For example, in protection applications, over-provisioning provides low latencies and high availability in the steady state but not necessarily in the face of IT failures, bugs in software or hardware that cause spurious traffic, or cyberattacks. As more applications that can exploit coherent, real-time data delivery emerge, such as those outlined in Section II, using isolated networks may soon become unsustainable, as will designing a new communications system for each new application or application family.

Fortunately, the state of the art in distributed computing, real-time systems, and fault-tolerant computing does support providing strong guarantees with data delivered to many applications. If designed, implemented, and validated correctly, a state-of-the-art data delivery system can greatly lower the barrier to enter (in both time and money) and enable deployment of new power applications by simplifying the process of adding new sensors. If designed incorrectly, it will be difficult to maintain in the future because it will not be able to keep up with increasing demands. Further, these data delivery systems will have a long life, and no single network-level mechanism (for multicast, security, or *QoS*) can be assumed to be everywhere.

TABLE 2: DIVERSITY OF DATA DELIVERY OF SELECTED POWER APPLICATIONS

		Direct State Measurement	Operator Displays	Catch Up for Operator Displays	Distributed Wide-Area Control	System Protection With Time-Synchronized Data	Anti-Islanding	Post-Event Analysis	Research
Loop Entity		P/C	P	P	C	C	C	P	P
Inputs	Kind	SS	SS	Co	SS	SS	SS	Co	Co
	Lat.	1–5	1	1	2–4	4–5	4–5	1	1–5
	Rate	1–5	2–3	1	2–4	4–5	3–5	1	1–5
	Crit.	1–5	2–4	1–2	5	5	4–5	1–5	1–5
	Quan.	1–5	3–5	3–4	3–5	2–4	1–3	5	1–5
	Geog.	1–5	5	5	1–5	1–4	1–2	3–5	3–5
	Dline	—	—	5	—	—	—	2–3	1
Outputs	Kind	SS	SS	Bu	SS/Co	Co	Co	Bu	Bu
	Lat.	1–5	1	1	2–4	5	3–5	—	—
	Rate	1–5	1	1–2	2–4	—	—	—	—
	Crit.	1–5	2–4	1–2	5	5	5	1–2	1
	Quan.	1–5	3–5	3–4	1–3	1–3	1–2	5	5
	Geog.	1–5	1	1	1–4	1–3	1–2	5	5
	Dline	—	—	5	—	—	—	2–3	1
NASPInet Class		B	D	—	B	A	A	C	E

It is crucial, therefore, that data delivery systems between the mission-critical peer-to-peer automatic protection and control systems, and the power grid’s operations IT backbone, have interoperability between different kinds of network mechanisms providing the same property, such as delay guarantees [53].

In this section, we examine how a WAMS-DD will be an enabling technology for the new and emerging power system. We first overview the performance and reliability requirements that a WAMS-DD must meet. We then present implementation guidelines, based on best practices in other industries and in the field of distributed computing systems, for achieving these requirements. Next we compare how existing technologies meet these delivery requirements and design guidelines when used in isolation without additional overlay networks. This includes technologies and standards at the network layers (and below), the middleware layer(s), and related ones from the power industry. We also discuss relevant research and development for wide-area middleware. After this, we discuss the emerging NASPInet effort and the GridStat data delivery middleware. Finally, we conclude this section with a brief discussion of pertinent cybersecurity issues for next-generation data delivery services for the electric power grid.

Note that the following analysis focuses on coherent but asynchronous data delivery for operations, but the emerging communications infrastructure will provide the additional benefit of distributing the time signal required for time-synchronized measurements and control. Typically time is received via GPS and distributed over a separate physical

network using protocols such as IRIG. This results in a physical cable connection to the measuring devices such as PMUs. Combining time distribution with the communications network provides advantages such as simplicity and reliability [54]. Furthermore, for many applications, such as a control scheme or system protection scheme, which use separate mission-critical peer-to-peer communications, operation can proceed even if global time is lost, as long as they maintain a local coherent time signal. The communications infrastructure can provide this locally common time signal when the primary GPS signal is unavailable.

A. System Model

Fig. 12 depicts the architecture of a WAMS-DD. Application programs or firmware that emit a stream of updates are called publishers, which are denoted as Pub_1 through Pub_N in the diagram; Pub_1 , for example, outputs updates to variables X and Y. Applications that receive these updates are called subscribers, which are denoted as Sub_1 through Sub_N . In the diagram, Sub_1 subscribes to Y from Pub_1 and to W from Pub_N .

In the usual case in publish-subscribe (pub-sub) systems, neither publisher nor subscriber needs to know about the other; they are decoupled such that they only know about the variable they publish or subscribe to and how to contact the delivery system. In cases where the subscriber requires confirmation that the update came from its legitimate publisher—which may be common with a WAMS-DD—data integrity techniques from the computer security field can be used by the data delivery system.

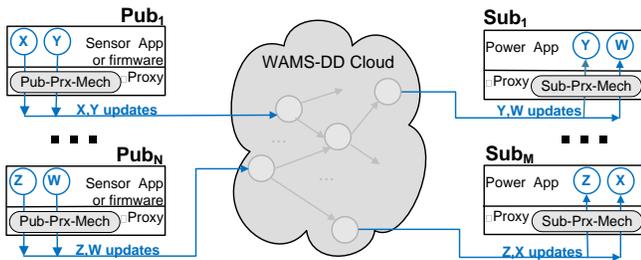


Fig. 12. Architecture and System Model of a WAMS-DD

Creating a pub-sub delivery path requires two steps. Publishers register their variables with the delivery system (only once per variable, not once per subscriber), and subscribers request a subscription to a given variable. For both publishers and subscribers, the delivery system returns a handle to a piece of code called a proxy, which is generated at compile time by the data delivery middleware. This proxy contains logic provided by the data delivery service, which, besides doing the usual middleware proxy activities such as packaging of the parameters into a message, is also a place where data delivery mechanisms may reside. In Fig. 12, we denote a publisher-side proxy as Pub-Prx-Mech and the subscriber-side proxy as Sub-Prx-Mech.

After the variable is registered and subscribed to, updates to variables flow from publishers to subscribers, as shown in blue in Fig. 12. To do this, they traverse what we call the WAMS-DD Cloud. This is opaque because, as shown later in this section, it can be implemented in different ways resulting in different tradeoffs. For the purposes of our system model, the WAMS-DD Cloud consists of a graph where the edges are network links and the nodes contain forwarding mechanisms that can forward a message on its way toward a subscriber.

Updates from a publisher of a sensor variable thus traverse one or more paths to be delivered to a given subscriber. Along a given path, an update may be delayed, so that its required delivery latency cannot be met, or the update may be dropped due to failures in a network link or forwarding node or due to a cyberattack. However, the probabilities of an update not meeting its delivery requirements can be held extremely low by carefully designing the WAMS-DD and by allocating multiple paths for important updates. That is, a WAMS-DD can be constructed so that the on-time delivery probability is very high, so long as its design constraints are met. Informally, these include forwarding capacity per node, maximum link traffic, number and kind of benign failures, and cyberattacks, etc.

We now overview the delivery requirements in Section IV B; then in Section IV C, we describe implementation guidelines that can be used to meet these delivery requirements with extremely high probabilities. These probabilities offer the potential to practice using dual isolated networks for critical protection applications, while at the same time supporting many more application families with thousands of update flows. However, such delivery technologies clearly need to be proven in the field before any migration to them can begin to be contemplated.

B. Delivery Requirements for a WAMS-DD

The following delivery requirements (DRs) must be met by a WAMS-DD [55], [56]; these do not include the details of cybersecurity-related requirements. These DRs are in addition to the requirements of other network layers, such as the physical, link, network, and transport layers, which are outside the scope of this paper.

Requirement 1. *Hard, end-to-end (E2E) guarantees must be provided over an entire grid because protection and control applications depend on the data delivery. The guarantees must be deterministic: met unless the system's design criteria have been violated (e.g., traffic amount, number of failures, and severity of cyberattack).*

Requirement 2. *WAMS-DDs must have a long lifetime and thus must be designed with future-proofing in mind. This is crucial in order to amortize costs over many projects, utilities, grids, etc. The goal of NASPInet, for example, is to last at least 30 years.*

Requirement 3. *Multicast (one-to-many) is the normal mode of communications, not point-to-point. Increasingly, a given sensor value is needed by multiple power applications.*

Requirement 4. *End-to-end guarantees must be provided for a wide range of QoS+. Data delivery for the power system is not "one size fits all" [50], as shown in Section II and Table 2. For example, to provide very low latencies, very high rates, and very high criticality/availability to all applications would be prohibitively expensive. Fortunately, many applications do not require these stringent guarantees, but their less stringent requirements must nevertheless be met.*

Examples of the wide ranges that must be provided follow (summarized partly from Table 1 and Table 2):

- A. Latency and Rate: Ten milliseconds or less, up to seconds (or hours or days for bulk transfer traffic), .001 Hz to 720 Hz or more.
- B. Criticality/Availability: IntelliGrid [50] recommends five levels of availability of data, from ultra to medium.
- C. Cybersecurity: Support a range of tradeoffs of encryption strength compared to delay induced and resources consumed.

Requirement 5. *Some merging and future SIPS, transient stability, and control applications require ultra-low latencies and one-way delivery on the order of a half or full power cycle (8–16 ms in the US) over hundreds of miles [31]. Thus, any forwarding protocols should not add more than a millisecond or two of latency (through all forwarding hops) on top of the speed of light in the underlying communications medium.*

These latencies must be provided in a way that:

- A. Is predictable and guaranteed for each update message, not a much weaker aggregate guarantee over longer periods of time, applications, and

locations such as is provided by multiprotocol label switching (MPLS) technology [57].

Each sensor update needs to arrive within its required guaranteed deadline.

B. Tolerates nonmalicious failures in the WAMS-DD infrastructure.

No system can tolerate unlimited kinds and numbers of failures. However, much like the power system must continue in the face of one or more known contingencies, the IT infrastructure on which it increasingly depends must still provide these hard, end-to-end guarantees in the face of failures (up to design limits).

C. Tolerates malicious cyberattacks.

Power systems are known to be subjects of extensive study and probing by multiple organizations that have significant information warfare capabilities, including nation states, terrorist organizations, and organized crime. A WAMS-DD must adapt and continue to deliver data despite cyberattacks of a designed severity (a bar that should be increasable over the life of the system). Note that a bug in hardware or software that generates spurious traffic can have an effect similar to that of a cyberattack.

Requirement 6. *Extremely high throughput is required.*

Today’s synchrophasor applications are generally limited to 30 or 60 Hz in the USA, in part because the communications systems they use are not designed to support higher rates. To not provide much higher sustainable throughput would greatly limit the number of new applications that can help the power system’s stability. Indeed, not just synchrophasors but digital fault recorders (DFRs) and IEDs in substations provide a wealth of data. It is quite conceivable and likely that “If you build it, they will come” and there will be many thousands of synchrophasors, relays, DFRs, and other sources of sensor updates across a grid. These devices can output at 720 Hz and sample at 8 kHz, but their full output is not always used remotely due to communications limitations. If key relay or DFR data could be delivered from a set of devices across a grid at 720 Hz, many new opportunities would open up for transient protection without using expensive dedicated networks or “drilling down” into the root causes of an ongoing power contingency using additional contingency-specific data.

We are not aware of any commercial or military market for a wide-area data delivery infrastructure that has the stringent requirements of a WAMS-DD including ability to enforce complete perimeter control, ability to know the vast majority of the traffic ahead of time, and other factors incorporated into the implementation guidelines described next in this paper. The reason is quite simple; electric power is the only market with such stringent requirements. However, these requirements are achievable using state-of-the-art distributed

real-time embedded computing [58], [59], as long as a careful end-to-end analysis is done [60] and the core data delivery mechanisms are not saddled with unnecessary features. Much broader reliability has been explored in the fault-tolerant distributed computing community, from where appropriate lessons, both good and bad, should be heeded [61], [62].

C. Implementation Guidelines for a WAMS-DD

The requirements outlined in the previous subsection were kept to a minimum. In order to achieve them, a number of implementation guidelines (IGs) are enumerated and explained in this section, many of which are quite different from what is provided in today’s best-effort Internet and what has been the standard practice in networking and distributed computing research.

Some of the IGs below (e.g., IG4 and IG5) are actually deemed requirements for NASPInet [56], but we describe them here as IGs because it is possible to build a WAMS-DD without them. These IGs are drawn from a number of sources, including our knowledge of what the state of the art in distributed computing has demonstrated is feasible, best practices in other industries, and decades of experience gained in Defense Advanced Research Projects Agency (DARPA) wide-area application and middleware projects [63]–[67], [58], [59].

We also note that the scope of these IGs involves only the data delivery system for a WAMS-DD. It does not include the supporting services that will be required for configuration, security, path allocation, resource management, etc. It is important to avoid hard-coding these tools in a WAMS-DD, but rather allow them to be specified in a high-level policy language (or at least a database) [66]–[68]. For an example of a hierarchical version of such services (a “management plane”), see [69], [70].

Guideline 1. *Avoid post-error recovery mechanisms.*

Traditional protocols for the Internet in general and reliable multicast protocols from the fault-tolerant computing research community use post-error recovery [62]. In these protocols the receiver either sends a positive acknowledgement (ACK) when it receives a message, or it sends a negative acknowledgment (NACK) when it concludes that the message will not arrive. However, this can add considerable latency when a message¹ gets dropped; three one-way latencies are required plus a relatively large timeout. This violates DR1, DR5A, and DR5B.

The better alternative is to send sensor updates proactively over multiple disjoint paths, each of which would meet the latency and rate requirements [71], [72]. Indeed, if multiple independent messages, each going over a QoS-managed path, cannot meet the delivery deadline, then sending ACKs or NACKs is very unlikely to help and the resulting additional network traffic may make things worse.

¹ We use the term “message” rather than “packet,” because in many cases we are describing middleware-layer mechanisms above the network and transport layers.

Note that the guideline to avoid post-error correction is only for data with guarantees on a per-message basis. Bulk data transfer is similar to a remote file transfer and will almost certainly employ post-error correction. However, bulk data transfer mechanisms must be different from the ones that have to provide per-message guarantees and isolated from the hard real-time mechanisms.

Guideline 2. *Optimize for rate-based sensors.* A WAMS-DD can be made with higher throughput and robustness if not over-engineered. General-purpose pub-sub systems offer a wide range of traffic types, because they are designed to support a wide range of applications [73]. However, in a WAMS-DD, the vast majority of the traffic will be rate-based.

Guideline 3. *Provide per-subscriber QoS+.* It is crucial that different subscribers to the same sensor variable be able to have different guarantees in terms of latency, rate, and criticality/availability. If not, then a lot of bandwidth will be wasted; all subscribers will have to receive the most stringent QoS+ required by any of its subscribers.

Guideline 4. *Provide efficient multicast.* In order to achieve the highest throughput possible, it is imperative to avoid unnecessary network traffic. Thus, never send an update over a link more than once. Also, as a sensor update is forwarded through the network, the update message should be dropped if it is not needed downstream in the multicast tree (e.g., by subscribers who require it at a lower rate than other subscribers). This can be implemented using a rate down-sampling mechanism as is done in GridStat [71], [72].

These first four guidelines add up to a need for multicast routing heuristics that provide multiple disjoint paths to each subscriber, with each path meeting the subscriber's latency requirement. A family of heuristics developed for this multicast routing problem [74], [75] confirms the feasibility of the approach at the anticipated scale (see IG10 in Table 3) if routing decisions are made statically (see IG11 in Table 3). GridStat's route selection mechanisms are based on these four guidelines.

Guideline 5. *Provide synchronized rate down-sampling.* In providing rate down-sampling, it is important to not down-sample in a way that destroys the usefulness of some data. For example, synchrophasors are used to take a direct state measurement at a given microsecond. If some subscribers require only a small fraction of the updates for a set of synchrophasor sensors, the updates that reach the subscriber at each interval must carry the same timestamp. If a subscriber only requires a tenth of the updates from two different variables, then it would not be useful to get updates {#1, #11, #21} from one synchrophasor and updates {#2, #12, #22} from another synchrophasor, because the given measurements do not correspond to the same time. They are not the same snapshot, which is the main point of synchrophasors.

Also, for applications that require reconstruction of the original signal, it is important to maintain Nyquist bandwidth filtering to avoid aliasing.

Guideline 6. *Don't depend on priority-based "guarantees."* Pub-sub delivery systems typically offer a way to specify a priority, so if the traffic gets too heavy, less important traffic can be dropped. However, this does not provide a hard end-to-end guarantee to subscribing applications, and even applications that are not of the highest criticality still need to meet their DRs. Instead of priorities, mechanisms must be used that exploit the characteristics of a WAMS-DD (as outlined in these guidelines) to provide each subscriber firm assurances that its guarantees will be met so long as the design criteria in terms of kind and numbers of failures (DR5B) and cyberattacks (DR5C) are not violated.

Guideline 7. *Provide end-to-end interoperability across different/new IT technologies (providing multicast, latency, rate, etc.).* Many WAMS-DDs will span multiple utility and network organizations. It is unlikely that the same mechanisms will be present across all these organizations. And, even if they are today, if the WAMS-DD gets locked into the lower-level application programming interfaces (APIs) and semantics of a given multicast or QoS mechanism, it will be difficult to "ride the technology curve" and use newer and better mechanisms that will inevitably become available over the long lifetime of the WAMS-DD. This is a stated goal of the GridWise community, see [76]. Fortunately, it is possible to use middleware to span these different underlying technologies in order to guarantee diversity of the underlying networks that must be spanned. Indeed, this is one of the main reasons for the development of middleware over the last three decades.

Guideline 8. *Exploit a priori knowledge of predictable traffic.* Internet routers cannot in general make assumptions or optimizations based on the characteristics of the traffic that they will be subjected to, because they are intended to be general-purpose and support a wide range of traffic types. A WAMS-DD, however, has traffic that is not just rate-based but is often known ahead of time, as is the case when an engineering survey is made of a new power application. This common case can be optimized, as described in later IGs below.

Guideline 9. *Have systematic, quick, internal instrumentation.* In order to provide end-to-end guarantees across a wide area despite failures and cyberattacks, IG8 must be exploited to provide systematic and fast instrumentation of the WAMS-DD. This allows much quicker adaptations to anomalous traffic, whether accidental or malicious in origin. Finally, this instrumentation should exploit the pervasive presence of GPS clocks in substations and in likely sites for WAMS-DD backbone mechanisms.

TABLE 3: IMPLEMENTATION GUIDELINES AND THE DELIVERY REQUIREMENTS THAT MANDATE THEM

DR1: Hard E2E WAN Guarantees	DR2: Future-Proofing	DR3: Multicast	DR4: Wide Range of QoS+...	4A: Latency and Rate	4B: Criticality/Availability	4C: Cybersecurity	DR5: Ultra-Low Latencies...	5A: Per-Update and Predictable	5B: Tolerating Failures	5C: Tolerating Cyberattacks	DR6: High Throughput	IGx Prerequisites	Summary of Implementation Guideline IGx
X								X	X				IG1: Avoid post-error recovery mechanisms
X				X				X	X	X	X		IG2: Optimize for rate-based sensors
		X									X		IG3: Provide per-subscriber QoS+
		X									X		IG4: Provide efficient multicast
											2, 3		IG5: Provide synchronized rate down-sampling
X					X			X			X		IG6: Don't depend on priority-based "guarantees"
X	X			X	X	X							IG7: Provide end-to-end interoperability across different/new IT technologies (multicast, QoS+)
X								X			X		IG8: Exploit <i>a priori</i> knowledge of traffic
X								X	X	X		8	IG9: Have systematic, quick internal instrumentation
X								X					IG10: Exploit smaller scale of the WAMS-DD
X								X				8–10	IG11: Use static, not dynamic, routing
X								X	X	X			IG12: Enforce complete perimeter control
X								X	X	X	X	12	IG13: Reject unauth. messages quickly and locally
											X	2, 8	IG14: Provide only simple subscription criteria
								X			X	2	IG15: Support transient, not persistent, delivery
								X			X		IG16: Don't over-design consistency and (re)ordering
											X	2, 8, 14–16	IG17: Minimize forwarding-time logic
X	X			X	X	X							IG18: Support multiple QoS+ mechanisms for different operating conditions
								X			X	17	IG19: Inspect only message header, not payload
X								X			X		IG20: Manage aperiodic traffic

Guideline 10. *Exploit smaller scale of the WAMS-DD.* This is crucial if the challenging delivery requirements are to be met over a wide area with reasonable cost. However, this requires rethinking the conventional wisdom in networking research and commercial middleware products.

NASPINet data buses (NnDBs), see Section IV D 4, will be orders of magnitude smaller in scale than the Internet at large², so it is feasible for the entire configuration to be stored in one location for the purposes of (mostly offline) route selection. Additionally, academic computer science researchers historically consider

² For example, in the entire USA there are approx 3500 companies that participate in the grid [2]. Therefore, in the case of a broker-based pub-sub system (defined later), the number of router-like forwarding engines that would be required for an NnDB backbone is at most 10^4 and likely only around 10^3 .

something that is $O(N^2)$ for path calculation with N routers or forwarding engines to be infeasible; see for example [70]. However, this assumption ignores two key factors for WAMS-DDs. First, N is not in the neighborhood of 10^8 as in the Internet, but rather is more likely $\sim 10^3$ at least for the next 5–10 years. Even $O(N^2)$ storage of state is feasible at this scale. Second, as a rule, power engineers do not decide that they need a given sensor's values seconds before they really need it, due in part to the fact that today's data delivery infrastructure requires them to recode hard-coded socket programs and then recompile. Rather, power engineers plan their power contingencies (and what data they will need in them) months ahead of time with detailed engineering studies, and they plan similarly for their monitoring, protection, control, and visualization needs. Thus, the routing/forwarding decisions involved in path selection

can be done offline well ahead of time, while still allowing for handling a modest number of subscription requests at runtime.

It is also feasible for router-like forwarding engines to store state for each flow. Having a router keep per-flow state has long been considered a bane to networking researchers, because it is considered to be prohibitively unscalable. However, with the much smaller scale, and the much more limited type of applications for a WAMS-DD, storing per-flow state is not only feasible but it is a requirement for providing IG3 (per-subscriber QoS+) with IG4 (efficient multicast); this is something that the GridStat project has been using for many years [69]. Recently, however, networking researchers are realizing the necessity of storing per-flow state to provide any reasonable kind of QoS [77]. Other recent efforts with roughly similar approaches include CHART [78] and PHAROS [79].

Guideline 11. *Use static, not dynamic routing and naming.* Much stronger latency guarantees can be provided when using complete knowledge of topology (IG10) coupled with static routing. Complete topology knowledge is a reasonable assumption in an NnDB, given that it will be a carefully managed critical infrastructure with complete admission control. Also, almost all of the sensors and power applications will be known well ahead of time, so optimizations for static or slowly-changing naming can potentially be useful and can be done while still providing more flexible and dynamic discovery services at a much lower volume. We note that networking and security researchers generally assume that the membership of multicast groups (or a set of subscribers) may change rapidly; see for example [70]. However, this is not the case with a WAMS-DD.

Guideline 12. *Enforce complete perimeter control.* All traffic put onto a WAMS-DD must pass admission control criteria (permissions based on rules for both cybersecurity and resource management) via a management system where the publisher registers a sensor variable at a given rate and the subscriber asks for a subscription with a given rate and end-to-end latency. This is essential to provide guarantees at a per-message granularity. It also enables quicker adaptations.

Guideline 13. *Reject unauthorized messages quickly and locally.* Messages that have gone around the admission control perimeter should be rejected as soon as possible, ideally at the next WAMS-DD forwarding engine, rather than going most or all the way across the WAMS-DD consuming resources along the way. Detection of such unauthorized packets is an indicator of anomalous traffic and hence, evidence of a failure or cyberattack that needs to be reported to the management infrastructure. When sufficient evidence over sufficient time is collected, an appropriate adaptation can occur.

Guideline 14. *Provide only simple subscription criteria.* This is exactly the opposite of what is usually done with

general purpose pub-sub systems in either academic research or commercial products. Both tend to favor complex subscription criteria, which are expensive to evaluate because each update is forwarded through the system [73]. For example, in GridStat, the subscription criteria are latency, rate, and number of paths, and the forwarding decision is completely based on rate with static routing. Note also that the lower-level ID of a sensor variable could still be looked up through a complicated discovery service. This guideline is concerned with avoiding complex forwarding logic.

Guideline 15. *Support only transient delivery, not persistent delivery.* Most pub-sub systems offer persistent delivery, whereby if an event cannot be immediately forwarded, it is stored for a period of time and then the delivery retried. This method harms throughput, however, and potentially harms the per-packet predictability because it requires storing the data. Persistent delivery may also be unnecessary in many cases when using real-time visualization, control, and protection, due to the temporal redundancy inherent in rate-based update streams. In addition, the next update will be arriving very soon, so the usefulness of a given update decays quickly. Thus, it is inadvisable to complicate delivery mechanisms to support persistent delivery, though it can be provided “on the side” by other mechanisms. Furthermore, in the power system, historian databases are already required for archiving data, so there is no reason to complicate the design or otherwise bog down the fastest and highest availability mechanisms of a WAMS-DD in order to deliver historical data³.

Guideline 16. *Don't over-design for consistency and (re)ordering.* Research in fault-tolerant multicast tends to provide different levels of ordering between updates from the same publisher or between different clients of the same server, as well as consistency levels between different replicas or caches of a server [61], [62]. There is no need for anything like this in a WAMS-DD. The only requirement for such consistency that we have found is reflected in IG5 for synchrophasors, and the only ordering of any kind is where a PDC combines updates from different PMUs into one message to pass on. When using devices containing synchrophasors that have an accurate GPS clock, the order of events is clear, and the only delivery ordering mechanism required is that which the application performs.

Guideline 17. *Minimize forwarding-time logic.* In order to provide the highest throughput, the forwarding logic that decides how a packet or update is to be forwarded should be kept as simple as possible. On the GridStat project, forwarding decisions are made based solely on the subscription rate of subscribers downstream in the multicast tree [69], [72]. Given that the traffic is rate-

³ We note that such post-event historical data can be delivered by the same network links as the fast traffic with traffic isolation mechanisms; indeed, this is one of the main traffic categories for the emerging NASPInet.

based (IG2) and known ahead of time (IG8), subscription criteria are kept simple (IG14); only transient delivery is supported (IG15), and there are no consistency semantics (IG 16). Much logic can be pushed off to subscription setup time or even offline. This reduces the logic necessary when an update arrives at a forwarding engine (or peer-to-peer middleware mechanisms at an edge) and hence, greatly increases throughput and decreases latency.

Guideline 18. *Support multiple QoS+ mechanisms for different runtime conditions.* A given mechanism that provides guarantees of latency and security, for example, will not be appropriate for all the runtime operating conditions in which a long-lived WAMS-DD may have to operate. This is because different implementations of a given QoS+ mechanism can require very different amounts of lower-level resources such as CPU, memory, and bandwidth [64]. This will be particularly important as WAMS-DD deployments span areas that cannot be controlled nearly as closely as the core backbone.

Guideline 19. *Inspect only packet header, not payload.* In order to provide the highest throughput and lowest latency, ensure that subscription criteria and consistency semantics allow a forwarding decision to be based solely on a packet header. This is not possible for pub-sub middleware that has complicated subscription topics, as is typical with commercial and research systems. For such middleware, data fields in the payload also have to be inspected.

Guideline 20. *Manage aperiodic traffic.* Any traffic that is aperiodic, i.e., not based on rate but on a condition, must be isolated from rate-based periodic traffic and managed accordingly. This can be done deterministically, for example with OSI Layer 1 optical wave division multiplexing (OWDM) hardware. Further, aperiodic traffic should be aggregated intelligently—ideally based on updateable policies rather than hard-coded settings—instead of sending all alarms/alerts to the next level up for processing.

It is important to recognize that you can't have the highest level of all the properties described in the DRs for every sensor variable. Reference [53] lists the following DR observations:

1. Different properties inherently must be traded off against others.
2. Different mechanisms for a given property are appropriate for only some of the runtime operating conditions that an application may encounter (especially a long-lived one).
3. Different mechanisms for the same nonfunctional property can have different tradeoffs of lower-level resources (CPU, bandwidth, storage).
4. Mechanisms most often can't be combined in arbitrary ways.

Even if you somehow could have them all at once, it would be prohibitively expensive. Given these realities, and the fact that application programmers rarely can be expert in dealing with the above issues, middleware with QoS+ properties

supported in a comprehensive and coherent way is a method of packaging and handling these issues and allowing reuse across application families, organizations, and even industries.

Similarly, it is important to note that meeting IG3 (and others) requires the data delivery system to be provided at the middleware layer. This is because network-level mechanisms know about packets and IP addresses, not middleware-layer sensor variables and the power applications that subscribe to their updates. There is thus no way that network-level mechanisms can provide different QoS+ guarantees to different subscribers of the same sensor variable, which is mandated by efficient multicast (IG4).

Finally, because of length constraints, it is not possible in this paper to fully discuss the cybersecurity issues that arise in a WAMS-DD. Clearly, a WAMS-DD, providing universal connectivity, creates cybersecurity challenges beyond those arising in a conventional, single-utility SCADA system. Cybersecurity also interacts with DRs and IGs. For example, techniques used for message confidentiality and authentication must not impose too much additional latency, yet the multicast requirement appears to limit use of symmetric-key cryptography for authentication. Of the traditional "CIA" cybersecurity properties (confidentiality, integrity, and availability), many power practitioners consider availability to be the most important for a WAMS-DD. See [80] for an example.

D. Analysis of Existing Technologies for a WAMS-DD

We now analyze how existing technologies and standards meet the above DRs and IGs.

1) Technologies and Standards at the Traditional Network Layers

Traditional network protocols, including the OSI-2 "Data Link" layer (e.g., Ethernet), OSI-3 "network" layer (e.g., IP), and the OSI-4 "transport" layer (e.g., TCP, UDP, SCTP) do not provide end-to-end QoS+ guarantees or multicast [81], [82]. This is because they are at lower networking layers and end-to-end functionality is not their intended use. All of these lower-layer protocols can be part of the complete network solution that WAMS-DD sits above. Nevertheless, some systems do apply them in ways that are nearly end-to-end in scope, and therefore we now examine these protocols and extensions to them to see how they meet the requirements and guidelines if they were implemented as the end-to-end solution. We do not consider experimental or emerging network technologies such as CHART [78], PHAROS [79], and Anagram's Flow routers [77]. Such technologies may someday be helpful in providing QoS guarantees across parts of a WAMS-DD. Also, mission-critical power system applications dictated the creation of several new Ethertypes to enable design for deterministic behavior. Ethernet multicast generic object-oriented substation event (GOOSE), sampled values, and line current differential messages each have their own Ethertype and operate at the OSI-2 layer with other Ethernet frames. Their deterministic behavior lies outside the scope of this paper.

IPv6 flow labels [83] associate each "reservation" with an application-to-application network socket connection, which contains many different sensor update streams with a wide

range of required QoS+. Packets are processed in a flow-specific manner by the nodes that have been set up with a flow-specific state. The nature of the specific treatment and the methods for the flow state establishment are out of scope of the specification.

IP multicast provides efficient multicast for a single, nonreplicated flow. However, if multiple IP multicast groups are used as a replication mechanism, there is no guarantee that the corresponding multicast trees will be disjoint, which is important not only for efficient multicast (IG4) but also for providing low latencies in the face of failures (DR5B). It also does not, by itself, have other end-to-end capabilities that are necessary for a WAMS-DD.

MPLS is designed to give Internet Service Providers (ISPs) a set of management tools for bandwidth provisioning, not to provide fine-grained (per-update) QoS [84]. Its guarantees are weak compared to the needs of a critical infrastructure. For example, it gives aggregate economic guarantees over user, location, and protocol, not hard guarantees (DR1) for each update (DR5A). Further, different ISPs can implement MPLS in different ways. There are no facilities for combining flows across different ISPs, as would be required in a WAMS-DD, or for predicting the end-to-end delays.

MPLS has some fault tolerance mechanisms, such as a fast reroute feature, detour merging, and end-to-end path protection. However, these mechanisms presently provide a minimum latency of about 50 ms, which is too long for the emerging SIPS applications described earlier in this paper.

Virtual local-area networks (VLANs) and virtual private networks (VPNs) do not meet the DRs listed above because their purposes are orthogonal to the DRs. A VPN or VLAN could be part of a WAMS-DD, but VPN and VLAN technologies alone do not meet the requirements and can add to latency and decrease throughput.

Pragmatic General Multicast (PGM) is a transport-layer multicast protocol [85]. Implementation by Microsoft is known as Reliably Delivered Messages (RDM). PGM runs over a datagram multicast protocol such as IP multicast, to provide basic reliable delivery by use of negative acknowledgements (NACKs). PGM uses a rate-based transmission strategy to constrain the bandwidth consumed. However, it does not provide real-time guarantees.

Spread can be considered a high-level multicast protocol that provides a range of ordering strengths across a wide-area network (WAN) [86], [87]. It supports ordered delivery and the resulting consistency, even in the face of network partitions, and it is used largely for replicating databases. It has no real-time mechanisms.

Asynchronous Transfer Mode (ATM) and Synchronous Optical Networking (SONET) are networking technologies sometimes employed in WANs. They offer strong latency guarantees on a per-message basis. ATM does not support multicast (DR3) and multiple disjoint paths (DR4B). Given ATM's strong latency guarantees, at the right granularity, the ATM protocol can be part of a WAMS-DD.

2) Commercial middleware technologies and standards

There is a wide range of commercial, off-the-shelf (COTS) middleware frameworks providing different kinds of services

with some relevance for a WAMS-DD. We first consider middleware supporting the pub-sub paradigm. There are two distinct architectures for pub-sub middleware, each with advantages and disadvantages.

Broker-based (BB) pub-sub systems rely on an infrastructure of broker nodes to forward messages toward subscribers. The Data Delivery Plane (DDP) for a WAMS-DD, though not necessarily commercial systems, is a managed WAN because it implements IG12 (complete perimeter control).

A BB pub-sub WAMS-DD is depicted in Fig. 13. A node in the DDP is called a Forwarding Engine (FE) and is a device specialized for the particular BB pub-sub framework. We depict the mechanisms that a BB WAMS-DD system can exploit in green; these consist of the proxies and the FEs.

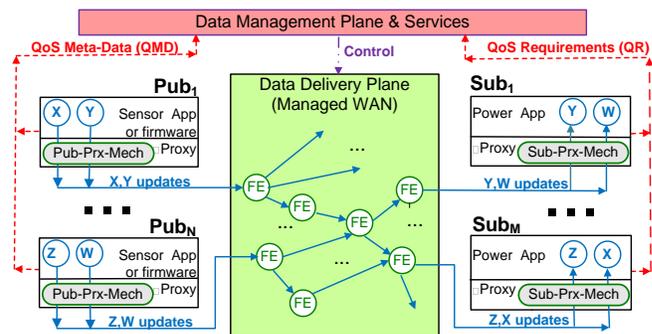


Fig. 13. Broker-Based WAMS-DD

In BB WAMS-DD systems intended for mission-critical applications, there is often a separate “plane” for managing the system⁴ and providing services. This Data Management Plane (DMP) is depicted in red in Fig. 13; it is shown here as a single entity but is often distributed. Publishers provide the DMP with basic QoS meta-data (QMD) about their publications, e.g., the rate at which they will output updates. Subscribers provide QoS requirements (QR) including rate and latency. The DMP then exerts control over the DDP (depicted in purple) in order to provide the delivery guarantees, e.g., by updating a forwarding table for an FE.

BB pub-sub systems require a broker/server infrastructure to be installed; you can’t just buy an IP router from Cisco or others. This can be a disadvantage, which often, for small and medium scales, cannot be amortized over enough applications to be justified. BB pub-sub systems have an advantage, however, in that they place intelligence inside the network, not just at the edges. This enables, for example, efficient multicast (IG4) and rate down-sampling throughout the data delivery system, not just at the edges. It also creates the potential to reject unauthorized packets at their next “hop” through the system (IG13).

Additionally, BB systems can exploit mechanisms in the graph of FEs in order to meet more of the IGs. For example, such an FE can be used to provide per-subscriber QoS+ (IG3), provide synchronized rate down-sampling (IG5), exploit *a priori* knowledge of traffic (IG8), and exploit the smaller scale

⁴ In telecommunications parlance, this is often called the “Control Plane,” hence our use of the term “plane.” Telecommunications parlance also refers to WAMS-DD as the “data plane.”

of the WAMS-DD (IG10), i.e., it can contain per-subscriber state, a forwarding table entry for every subscription for which it forwards updates. An example of a BB WAMS-DD is GridStat.

Peer-to-peer pub-sub systems place mechanisms for reliability and filtering only at the edges of an infrastructure. A canonical architecture for a peer-to-peer pub-sub configuration of a WAMS-DD is given in Fig. 14. For the DDP, peer-to-peer systems typically rely on a combination of IP multicast and Ethernet broadcast to be as efficient as possible. Note that in Fig. 14, we omit the DMP, which is often not present as a separate core entity in peer-to-peer systems; the edge mechanisms collectively implement it.

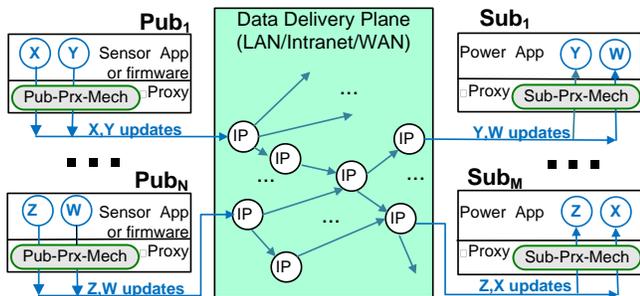


Fig. 14. Peer-to-Peer WAMS-DD

One other thing to note in Fig. 14 is that the controllable mechanisms for affecting traffic lie at the edges, in the proxies. Certainly a peer-to-peer WAMS-DD will exploit IP multicast as much as possible, but this has its limits, as described previously. Because its control mechanisms are at the edges, both QMD and QR are communicated to other proxies that collectively provide the delivery guarantees. Similarly, the only WAMS-DD-specific mechanisms are in the proxies, so control messages also must go there. In Fig. 14, to help aid understanding, the red and purple traffic lines are omitted. In practice, the red and purple traffic would be delivered via the DDP using IP routers.

Peer-to-peer pub-sub systems have an advantage in smaller and medium sized deployments, but for larger scales, the lack of mechanisms in the backbone core for rate down-sampling and fault tolerance limit their abilities to achieve extremely low latencies in the presence of failures.

A federated combination of peer-to-peer and BB pub-sub systems has the potential to offer much of the best of both worlds. Here, peer-to-peer pub-sub systems are employed near the edges, i.e., within a single utility or sometimes within an ISO. Between utilities or ISOs, BB pub-sub systems are used in order to support higher throughputs and the lowest possible latencies over distance. A federated amalgamation of peer-to-peer systems would feature a globally unique namespace for variables, and utilities and could seamlessly pass messages with standardized wire and message formats [53].

Business-to-business and web services is another middleware category called streaming queries (also known as complex event processing). It consists of a network of computer nodes that manipulate data streams through continuous queries in order to selectively propagate data, merge streams with existing data, or store data in a distributed

database. Such systems are not designed to provide hard end-to-end WAN guarantees (DR1) with per-message granularity (DR5A) while tolerating failures (DR5B). Given their intended application domain, they also do not follow most of the IGs.

More recently, a number of vendors are offering middleware based on web technologies such as HTTP, XML, and “web services” for use in the power grid. We note that scalability and throughput of such systems is difficult due to the added integration layers [88], [89].

3) Existing Power Technologies and Standards

Middleware is rarely used in today’s electric power systems, despite being considered a “best practice” in many other industries for a few decades [53]. It is not surprising, then, that there seems to be no networking technologies developed for the power grid that meet all of the DRs above. Part of this limitation is because commonly used power technologies are intended for a substation scope, with the QoS+ “mechanism” being over-provisioning of bandwidth. When moving from a LAN to a WAN environment, there are implicit design decisions that cannot be solved by layering a new “WAN-appropriate” API over existing LAN-based protocols [64]. We now overview some of the more common power protocols and standards related to communications.

OPC-UA [90] was designed for substations. It uses TCP, which was not designed for predictable latency and does not support multicast. Subscribers and publishers “ping” each other to verify if the other is up, which does not scale but ignores best practices for pub-sub systems.

The IEC 61850 communications standard was also designed primarily for applications associated with a substation automation system (SAS). It was conceived and created by protection providers in order to move data and information to, from, and among intelligent protection, control, and monitoring devices instead of legacy SCADA and RTU methods. The standard has over seven protocols designed within it that use several messaging methods mapped directly into one or more Ethernet frames. Some of the protocols use TCP methods to transport manufacturing messaging specification (MMS) messages to report data and transfer files to clients asynchronously via multiple frames. Several other protocols use layer-two methods and specially assigned Ethertypes to accomplish multicast messaging restricted to single frames for performance. Its use outside the substation is inherent in the chosen technology, and now work is being performed to map message contents to other protocols frequently used outside the substation as well as appropriate data delivery mechanisms. Its Common Information Model (CIM) can potentially be of use in a WAMS-DD, especially when the harmonization with C37.118 is completed, in particular in helping automate QoS+ settings and perhaps adaptation strategies for a wide variety of sensors and applications that use them. When IEC 61850 MMS and GOOSE APIs are successfully extended across the WAN, then IEC 61850 may well be able to successfully use a WAMS-DD transport. However, this will only be true if the WAMS-DD transport is carefully designed to support layer two multicast messaging in addition to TCP mechanisms. The DRs and IGs

closed loop teleprotection, automation, and telecontrol via multicast GOOSE require more than can be provided at the network layers. If a WAMS-DD network is deployed that does not support data link layer multicast, an overlay network of some kind will need to be created and provided. The extensions proposed in IEC 61850-90-5 to extend it to the wide area do not address underlying multicast mechanisms for a WAN. Such a multicast would need to meet the requirements in this paper, as well as many of the implementation guidelines, if low latencies and high throughput are to be achieved. Presently, IEC 61850-90-5 discusses delays of 50–500 ms, which do not support some of the more challenging applications outlined earlier in this paper.

IEEE C37.118 is a standard for synchrophasors that includes standard message formats. C37.118 is being revised to allow different data delivery mechanisms to be used. If successful, then C37.118 synchrophasor updates should easily be deliverable by any WAMS-DD transport.

MMS also does not have data delivery mechanisms. It can map onto the OSI protocol stack (which was not adopted in practice) and TCP/IP; see [81], [82].

An information architecture for the power grid is proposed in [91], which contains an analysis of 162 disturbances between 1979 and 1995. Reference [91] indicates that information systems have an impact on power grid reliability and points out major deficiencies in the current communications scheme. The paper contains proposals for different ways to structure interactions between control centers and substations, and it also contains reliability analyses of different schemes. However, it does not propose communication mechanisms and relies on off-the-shelf network technologies, which do not meet many of the DRs and IGs.

4) *NASPI*

The North American Synchrophasor Initiative (NASPI) is a government-industry consortium dedicated to effective deployment of synchrophasors in the United States. It is the only effort worldwide that is dealing with end-to-end WAMS-DD issues at a more-than-superficial level. To support the use of synchrophasors, NASPI has been developing the notion of NASPI_{net} (Nn), which has two main components, the data bus (NnDB) and the phasor gateway (NnPG). The NnPG is the edge component of Nn, interfacing the utility or ISO to the NnDB.

The NnDB is the electricity version of what is sometimes called an enterprise service bus (ESB), which provides communications services for business-to-business exchanges. NnDB satisfies the DRs described earlier in this paper. Five initial service classes have been identified for the NnDB in recognition of the fact that different kinds of traffic with different delivery requirements must be carried.

- A. Feedback Control (e.g., small signal stability)
- B. Feed-Forward Control (e.g., enhancing state estimators with synchrophasors)
- C. Post-Event (post-mortem event analysis)
- D. Visualization (for operator visibility)
- E. Research (testing or R&D)

Each class has associated qualitative requirements for such properties as low latency, availability, accuracy, time alignment, high message rate, and path redundancy. Distinguishing the classes in this way is an important first step for a WAMS-DD system. The NnDB classes also consider the lowest required latency to be 100 ms, which is insufficient for some applications. See Table 2.

One issue with the class definitions is that a customer, such as a utility, ISO, RTO, or NERC, cannot specify only what it wants from a telecom provider, e.g., a “Class A” network. This will not result in a WAMS-DD that meets the requirements across multiple traffic classes. For example, if too much traffic of “easier” classes is on the network, then you will not get Class A guarantees. Rather, to provide the DRs identified in Table 2, one needs to do resource management within the data delivery service. Network management components must account for all traffic associated with each subscription using a given level of QoS+. This is embodied in a number of IGs, including IG8 (exploit traffic knowledge), IG9 (systematic, quick, internal instrumentation), IG12 (complete perimeter control), IG13 (reject unauthorized packets quickly and locally), and IG20 (manage aperiodic traffic).

5) *GridStat*

GridStat is a data delivery service designed to support the DRs discussed in this paper. Its research results have influenced the shape of NASPI_{net} [56]. The GridStat research started in 1999 by looking at the QoS+ requirements of innovative power applications being developed by power researchers and analyzing closely what the state of the art in applied distributed computing systems could support. After significant gaps were identified, the detailed design and then programming of GridStat began in 2001.

GridStat is a BB pub-sub system that meets all of the DRs from this paper except for 5C, Tolerating Cyberattacks, which has been planned for and is near-term future research. It also implements all but three of the IGs, which have similarly been planned for and are also near-term future research. More on GridStat overall can be found in several publications:

- General details [55], [72], [92]
- QoS routing [74], [75]
- Securely upgradeable encryption and authorization infrastructures [93], [94]
- Forwarding Engines (NASPI_{net}-like routers) [95]
- Security and trust management issues for power grid WAMS-DD [96], [97]
- Advanced GridStat mechanisms [98], [95]

On a 2007-era PC, GridStat adds ~0.1 ms per overlay hop and handles ~20,000 forwards/s at each forwarding engine. On 2003-era network processor hardware, it adds ~0.01 ms/hop and scales to a few million forwards/s [99]. Using 2010 era hardware and a small cluster, these results would easily be extendable to achieve 50–100 million forwards/s while rejecting unauthenticated messages, monitoring traffic patterns, and checking for evidence of intrusions and cyberattacks. Custom hardware implementations would likely support even more throughput.

V. CONCLUSION

Creating and operating electric power grids that meet society's demands for reliability, efficiency, and integration of renewable energy sources is an ongoing challenge. Synchrophasors and other coherent, high-rate, measurements taken at hundreds or thousands of points in the grid and delivered in real-time to monitoring and control applications promise to help meet these challenges. In this paper, we have described some of the applications for these data that exist today, as well as new applications that are being investigated.

The quantities, rates, and real-time delivery requirements for these data, as well as the number of different applications that will use them, are different from the needs of current SCADA systems. Detailed consideration of the applications' QoS+ requirements for their data, exposes a need for a very flexible data delivery service to support these applications. Along with six hard requirements for the resulting WAMS-DD, we have identified twenty implementation guidelines that suggest how the requirements can be met. We find that off-the-shelf computer networking technology alone is not always sufficient for the task, unless augmented with middleware technology, such as the GridStat framework, to manage network resources and exploit them efficiently.

VI. BIOGRAPHIES



David E. Bakken is an Associate Professor of Computer Science at Washington State University (WSU), Pullman, WA. He has worked for Boeing, BBN, and has consulted to Amazon.com, Harris Corp, Real Time Innovations, and others. He received his B.S from WSU, and his M.S. and Ph.D. from The University of Arizona. Dr. Bakken is a Senior Member of the IEEE.



Anjan Bose is a Regents Professor and holds the endowed Distinguished Professor in Power Engineering at Washington State University, Pullman, WA. He received his B.Tech. (Honors) from the Indian Institute of Technology, Kharagpur, his M.S. from the University of California, Berkeley, and his Ph.D. from Iowa State University. He has worked for industry, academe and government for 40 years in power system planning, operation and control. Dr. Bose is a member of the National Academy of Engineering and the recipient of the Herman Halperin Award and the Millennium Medal from the IEEE.



Carl H. Hauser is an Associate Professor of Computer Science in the School of Electrical Engineering and Computer Science at Washington State University (WSU). His research interests include concurrent programming models and mechanisms, networking, programming language implementation, and distributed computing systems. His recent research focus has been on security aspects of emerging control systems for the smart grid. Prior to joining WSU, he worked at Xerox Palo Alto Research Center and IBM Research for a total of over 20 years after receiving his Ph.D. in Computer Science from Cornell University.



David Whitehead, P.E., is the vice president of research and development at Schweitzer Engineering Laboratories, Inc. (SEL). Prior to joining SEL, he worked for General Dynamics Electric Boat Division as a combat systems engineer. He received his B.S.E.E. from Washington State University in 1989, his M.S.E.E. from Rensselaer Polytechnic Institute in 1994, and is pursuing his Ph.D. at the University of Idaho. He is a registered professional engineer in Washington and Maryland and a senior member of the IEEE. Mr. Whitehead holds seven patents with several other patents pending. He has worked at SEL since 1994 as a hardware engineer, research engineer, and a chief engineer/assistant director and has been responsible for the design of advanced hardware, embedded firmware, and PC software.



Greg Zweigle is an engineer at Schweitzer Engineering Laboratories, Inc. He received his M.S.E.E. and M.S. in Chemistry from Washington State University. He also received a B.S. in Physics from Northwest Nazarene University. He is presently pursuing a Ph.D. focusing on energy systems. Mr. Zweigle holds seven patents and is a member of the IEEE.

VII. REFERENCES

- [1] G. Andersson, P. Donalek, R. Farmer, N. Hatziaargyriou, I. Kamwa, P. Kundur, N. Martins, J. Paserba, P. Pourbeik, J. Sanchez-Gasca, R. Schulz, A. Stankovic, C. Taylor, and V. Vittal, "Causes of the 2003 Major Grid Blackouts in North America and Europe, and Recommended Means to Improve System Dynamic Performance," *IEEE Trans. Power Systems*, vol. 20, no. 4, p. 1922, Nov. 2005.
- [2] U.S.-Canada Power System Outage Task Force, "Final Report on the August 14th, 2003 Blackout in the United States and Canada," 2004. [Online]. Available: <https://reports.energy.gov/>
- [3] *The Economist*, "Building the Energy Internet," 11 May 2004 (Technology Quarterly section).
- [4] L. Meegahapola and D. Flynn, "Impact on Transient and Frequency Stability for a Power System at Very High Wind Penetration," in *Proc. 2010 IEEE Power and Energy Society General Meeting*.
- [5] D. Gautam, V. Vittal, and T. Harbour, "Impact of Increased Penetration of DFIG-Based Wind Turbine Generators on Transient and Small Signal Stability of Power Systems," *IEEE Trans. on Power Systems*, vol. 24, no. 3, p. 1426, Aug. 2009.
- [6] A. G. Phadke and J. S. Thorp, "Synchronized Phasor Measurements and Their Applications," Springer, 2008.
- [7] K. Martin and J. Carroll, "Phasing in the Technology," *IEEE Power and Energy Mag.*, vol.6(5), pp. 24–33, Sept. 2008.
- [8] J. Thorp, A. Abur, M. Begovic, J. Giri, and R. Avila-Rosales, "Gaining a Wider Perspective," *IEEE Power and Energy Mag.*, vol. 6 no. 5, pp. 43–51, Sept. 2008.
- [9] S. Chakrabarti, E. Kyriakides, B. Tianshu, C. Deyu Cai, and V. Terzija, "Measurements Get Together," *IEEE Power and Energy Mag.*, 7(1), pp. 41–49, Jan 2009.
- [10] L. Ramesh, S. P. Chowdhury, and S. Chowdhury, "Wide Area Monitoring Protection and Control – A Comprehensive Application Review," in *Proc. 2010 Developments in Power System Protection, Managing the Change, 10th IET International Conference*.
- [11] S. H. Horowitz, D. Novosel, V. Madani, M. Adamiak, "System-Wide Protection," *Power & Energy Management*, Sept/Oct 2008.
- [12] A. G. Phadke, H. Volskis, R. Menezes de Moraes, T. Bi, R. N. Nayak, Y. K. Sehgal, S. Sen, W. Sattinger, E. Martinez, O. Samuelsson, D. Novosel, and V. Madani, "The Wide World of Wide-Area Measurement," p. 52, Sept/Oct 2008.
- [13] D. Novosel, V. Madani, B. Bhargava, K. Vu, and J. Cole, "Dawn of the Grid Synchronization: Benefits, Practical Applications, and Deployment Strategies for Wide Area Monitoring, Protection, and Control," *IEEE Power Energy Mag.*, vol. 6, no. 1, pp. 49–60, Jan. 2008.
- [14] A. Abur and A. G. Exposito, "Power System State Estimation, Theory and Implementation," New York: Marcel Dekker, 2004.
- [15] F. C. Schweppe and J. Wildes, "Power System Static-State Estimation, Part I: Exact Model," *IEEE Trans. Power Apparatus and Systems*, vol. PAS-89, pp. 120–125, Jan. 1970.
- [16] J. S. Thorp, A. G. Phadke, and K. J. Karimi, "Real Time Voltage-Phasor Measurements for Static State Estimation," *IEEE Trans. Power App. and Sys.*, vol. PAS-104, no. 11, pp. 3098–3106, Nov. 1985.
- [17] R. Zivanovic and C. Cairns, "Implementation of PMU Technology in State Estimation: An Overview," in *Proc. 1996 IEEE 4th Africon*, vol. 2, pp. 1006–1011.
- [18] E. O. Schweitzer III and D. E. Whitehead, "Real-World Synchrophasor Solutions," in *2008 35th Annual Western Protective Relay Conference Proceedings*.
- [19] E. O. Schweitzer III and D. E. Whitehead, "Real-Time Power System Control Using Synchrophasors," in *2007 34th Annual Western Protective Relay Conference Proceedings*.

- [20] T. Yang, H. Sun, and A. Bose, "Two-Level PMU-Based Linear State Estimator," in *Proc. 2009 IEEE PES Power Systems Conference and Exposition*, pp. 1–6.
- [21] M. Glavic and T. Van Cutsem, "Wide-Area Detection of Voltage Instability From Synchronized Phasor Measurements. Part I: Principle," *IEEE Trans. Power Systems*, vol. 24, no. 3, pp. 1408–1416, Aug. 2009.
- [22] A. Johnson, R. Tucker, T. Tran, D. Sullivan, C. Anderson, and D. E. Whitehead, "Static Var Compensation Controlled via Synchronphasors," in *2007 34th Annual Western Protective Relay Conference Proceedings*.
- [23] E. O. Schweitzer III, D. E. Whitehead, A. Guzmán, Y. Gong, and M. Donolo, "Advanced Real-Time Synchronphasor Applications," in *2008 35th Annual Western Protective Relay Conference Proceedings*, pp. 21–23.
- [24] M. Klein, G. J. Rogers, and P. Kundur, "A Fundamental Study of Inter-Area Oscillations in Power Systems," *IEEE Trans. Power Systems*, vol. 6, no. 3, pp. 914–921, Aug. 1991.
- [25] Y. Gong and A. Guzmán, "Synchronphasor-Based Online Modal Analysis to Mitigate Power System Interarea Oscillation," in *2009 DistribuTech Conference Proceedings*.
- [26] J. Chow, J. Sanchez-Gasca, H. Ren, and S. Wang, "Power System Damping Controller Design Using Multiple Input Signals," *IEEE Control Systems Mag.*, vol. 20, no. 4, pp. 82–90, 2000.
- [27] M. Aboul-Ela, A. Sallam, J. McCalley, and A. Fouad, "Damping Controller Design for Power System Oscillations Using Global Signals," *IEEE Trans. Power Systems*, vol. 11, no. 2, pp. 767–773, 1996.
- [28] C. Lu, Y. Han, X. Wu, P. Li, J. Wu, J. Shi, "Field Experiments of Wide Area Damping Controllers for Multiple HVDC Links," in *Proc. 2008 IEEE Asia Pacific Conference on Circuits and Systems*, p. 627.
- [29] E. O. Schweitzer III, D. E. Whitehead, G. C. Zweigle, and K. G. Ravikumar, "Synchronphasor-Based Power System Protection and Control Applications," in *2009 36th Annual Western Protective Relay Conference Proceedings*.
- [30] N. R. Chaudhuri, B. Chaudhuri, S. Ray, and R. Majumder, "Wide-Area Phasor Power Oscillation Damping Controller: a New Approach to Handling Time-Varying Signal Latency," *IET Generation, Transmission and Distribution*, vol. 4, no. 5, p. 620, 2010.
- [31] S. Horowitz, D. Novosel, V. Madani, M. Adamiak, "System-Wide Protection," *IEEE Power and Energy Mag.*, vol. 6 no. 5, pp. 34–42, Sept. 2008.
- [32] A. Guzmán, A.; D. A. Tziouvaras, E. O. Schweitzer III, E.O., and K. Martin, "Local- and Wide-Area Network Protection Systems Improve Power System Reliability," in *2006 Power Systems Conference: Advanced Metering, Protection, Control, Communication, and Distributed Resources Proceedings*, p. 174.
- [33] E. Martínez, N. Juárez, A. Guzmán, G. Zweigle, and J. León, "Using Synchronized Phasor Angle Difference for Wide-Area Protection and Control," in *2006 33rd Annual Western Protective Relay Conference Proceedings*.
- [34] R. J. Best, D. J. Morrow, D. M. Laverty, P. A. Crossley, "Synchronphasor Broadcast Over Internet Protocol for Distributed Generator Synchronization," *IEEE Trans. on Power Delivery*, vol. 25, no. 4, p. 2835, Oct. 2010.
- [35] *IEEE Standard for Interconnecting Distributed Resources With Electric Power Systems*, IEEE Standard 1547™, 2003.
- [36] J. Mulhausen, J. Schaefer, M. Mynam, A. Guzmán, and M. Donolo, "Anti-Islanding Today, Successful Islanding in the Future," in *2009 36th Annual Western Protective Relay Conference Proceedings*.
- [37] R. Bhandari, S. Gonzalez, M. E. Ropp, "Investigation of Two Anti-Islanding Methods in the Multi-Inverter Case," in *Proc. 2008 IEEE Power and Energy Society General Meeting – Conversion and Delivery of Electrical Energy in the 21st Century*.
- [38] A. Guzmán, V. Mynam, and G. Zweigle, "Backup Transmission Line Protection for Ground Faults and Power Swing Detection Using Synchronphasors," in *2007 34th Annual Western Protective Relay Conference Proceedings*.
- [39] J. Å. Walseth, J. Eskedal, and Ø. Breidablik, "Analysis of Misoperations of Protection Schemes in the Nordic Grid 1st of December 2005," *Protection, Automation and Control World*, Mar. 2010.
- [40] R. Moxley P.E., C. Petras P.E., C. Anderson, and K. Fodero II, "Display and Analysis of Transcontinental Synchronphasors," *2004 Western Power Delivery and Automation Conference Proceedings*.
- [41] D. J. Trudnowski, J. W. Pierre, N. Zhou, J. F. Hauer, and M. Parashar, "Performance of Three Mode-Meter Block-Processing Algorithms for Automated Dynamic Stability Assessment," *IEEE Trans. Power Systems*, vol. 23, no. 2, May 2008.
- [42] M. Parashar and J. Mo, "Real Time Dynamics Monitoring System (RTDMS): Phasor Applications for the Control Room," in *2009 42nd Hawaii International Conference on System Sciences Proceedings*.
- [43] J. M. Ordagci, H. C. T. Santos, S. R. Morand, R. Cespedes, R. Mano, and D. Caceres, "ONS – Brasil New Control Center Architecture Conceptual Design," in *2008 Transmission and Distribution Conference and Exposition Proceedings, Latin America*.
- [44] G. Zhang, S. Lee, R. Carroll, J. Zuo, L. Beard, and Y. Liu, "Wide Area Power System Visualization Using Real-Time Synchronphasor Measurements," in *Proc. 2010 IEEE Power and Energy Society General Meeting*.
- [45] J. Wang, Y. Hu, A. Johnson, H. Tram, and R. Nasri, "System Requirements of Visualization Platform for Wide Area Situation Awareness System," in *Proc. 2010 IEEE Power and Energy Society General Meeting*.
- [46] K. Sun; S. Likhate, V. Vittal, V. S. Kolluri, and S. Mandal, "An Online Dynamic Security Assessment Scheme Using Phasor Measurements and Decision Trees," *IEEE Trans. Power Systems*, vol. 22, no. 4, p. 1935, 2007.
- [47] Z. Huang, N. Zhou, F. Tuffner, Y. Chen; D. Trudnowski, W. Mittelstadt, J. Hauer, and J. Dagle, "Improving Small Signal Stability Through Operating Point Adjustment," in *Proc. 2010 IEEE Power and Energy Society General Meeting*.
- [48] Y. Gong, N. Schulz, and A. Guzmán, "Synchronphasor-Based Real-Time Voltage Stability Index," in *Proc. 2006 Power Systems Conference and Exposition*, p. 1029.
- [49] M. Donolo, M. Venkatasubramanian, A. Guzmán, F. de Villiers, "Monitoring and Mitigating the Voltage Collapse Problem in the Natal Network," in *Proc. 2009 Power Systems Conference and Exposition*.
- [50] Electric Power Research Institute (EPRI), "The Integrated Energy and Communication Systems Architecture," vol. IV: Technical Analysis, 2004.
- [51] S. Bhowmik, K. Tomsovic, and A. Bose. "Communication Models for Third Party Load Frequency Control," *IEEE Trans. Power Systems*, vol. 19:1, pp. 543–548, Feb. 2004.
- [52] J. Nataro, P. T. Kuruganti, L. Miller, S. Mullen, and M. Shankar, "Integrated Hybrid-Simulation of Electric Power and Communications Systems," in *Proc. 2007 IEEE Power Engineering Society General Meeting*, pp. 1–8.
- [53] D. E. Bakken, R. E. Schantz, and R. D. Tucker, "Smart Grid Communications: QoS Stovepipes or QoS Interoperability," in *Proc. 2009 Grid-Interop*, GridWise Architecture Council, Denver, Colorado. [Online]. Available: <http://gridstat.net/publications/TR-GS-013.pdf>.
- [54] K. Fodero, C. Huntley, D. E. Whitehead, and B. Kaszenny, "A Novel Scheme for Wide-Area Time Synchronization," in *Proc. 2010 Developments in Power System Protection, Managing the Change, 10th IET International Conference*.
- [55] D. E. Bakken, C. H. Hauser, H. Gjermundrød, and A. Bose, "Towards More Flexible and Robust Data Delivery for Monitoring and Control of the Electric Power Grid," Washington State University, Pullman, WA, Tech. Rep. EECS-GS-009, May 2007.
- [56] North American Synchronphasor Initiative, "Quanta Statement of Work."
- [57] E. Rosen, A. Vishanathan, and R. Callon, "RFC-3031: Multiprotocol Label Switching Architecture," *The Internet Society*, 2001. [Online]. Available: <http://datatracker.ietf.org/doc/rfc3031>.
- [58] Y. Krishnamurthy, V. Kachroo, D. A. Karr, C. Rodrigues, J. P. Loyall, R. E. Schantz, and D. C. Schmidt, "Integration of QoS-Enabled Distributed Object Computing Middleware for Developing Next-Generation Distributed Applications," in *Proc. 2001 ACM SIGPLAN Workshop on Optimization of Middleware and Distributed Systems*.
- [59] R. E. Schantz and D. C. Schmidt, "Research Advances in Middleware for Distributed Systems: State of the Art," in *Proc. 2002 IFIP World Computer Congress*.

- [60] J. Saltzer, D. Reed, and D. Clark, "End-to-End Arguments in System Design," *Trans. on Computer Systems, Association of Computing Machinery*, vol. 2, no. 4, pp. 277–288, Nov. 1984.
- [61] G. V. Chockler, I. Keidar, and R. Vitenberg, "Group Communication Specifications: A Comprehensive Study," *ACM Computing Surveys*, 33(4), pp. 1–43, Dec. 2001.
- [62] X. Défago, A. Schiper, and P. Urbán, "Totally Ordered Broadcast and Multicast Algorithms: Taxonomy and Survey," 36(4), pp. 372–421, Dec. 2004.
- [63] R. E. Schantz, R. H. Thomas, and G. Bono, "The Architecture of the Cronus Distributed Operating System," in *Proc. 1986 6th International Conference on Distributed Computing Systems*, IEEE Computer Society Press, Cambridge, Massachusetts, USA, pp. 250–259, May 1986.
- [64] J. Zinky, D. E. Bakken, and R. Schantz, "Architectural Support for Quality of Service for CORBA Objects," *Theory and Practice of Object Systems*, April 1997.
- [65] J. P. Loyall, D. E. Bakken, R. E. Schantz, J. A. Zinky, D. A. Karr, R. Vanegas, and K. R. Anderson, "QoS Aspect Languages and Their Runtime Integration," *Lecture Notes in Computer Science*, vol. 1511, Springer-Verlag, in *Proc. 1998 Fourth Workshop on Languages, Compilers, and Run-Time Systems for Scalable Computers*, pp. 28–30, May 1998.
- [66] P. P. Pal, J. P. Loyall, R. E. Schantz, J. A. Zinky, R. Shapiro, and J. Megquier, "Using QDL to Specify QoS Aware Distributed (QuO) Application Configuration," in *Proc. 2000 3rd IEEE International Symposium on Object-Oriented Real-Time Distributed Computing*.
- [67] R. E. Schantz, J. P. Loyall, M. Atighetch, and P. P. Pal, "Packaging Quality of Service Control Behaviors for Reuse," in *Proc. 2002 5th IEEE International Symposium on Object-Oriented Real-Time Distributed Computing*.
- [68] D. E. Bakken, "Quality of Service Design Considerations for NASPInet," *presentation to the North American Synchrophasor Initiative (NASPI) Work Group Meeting*, Feb. 4, 2009.
- [69] K. H. Gjermundrød, I. Dionysiou, C. H. Hauser, D. E. Bakken, and A. Bose, "Flexible and Robust Status Dissemination Middleware for the Electronic Power Grid," *School of Electrical Engineering and Computer Science*, Washington State University, Pullman, WA, Tech. Rep. EECS-GS-003, Sept. 2003.
- [70] R. Bobba, E. Heine, H. Khurana, and T. Yardley, "Exploring a Tiered Architecture for NASPInet," in *Proc. 2010 IEEE Conference on Innovative Smart Grid Technologies*.
- [71] K. Tomsovic, D. E. Bakken, M. Venkatasubramanian, and A. Bose, "Designing the Next Generation of Real-Time Control, Communication and Computations for Large Power Systems," in *Proc. IEEE Special Issue on Energy Infrastructure Systems*, 93(5), May 2005.
- [72] K. H. Gjermundrød, D. E. Bakken, C. H. Hauser, and A. Bose, "GridStat: A Flexible QoS-Managed Data Dissemination Framework for the Power Grid," *IEEE Trans. Power Delivery*, 4(1), pp. 136–143, 2009.
- [73] P. T. Eugster, P. A. Felber, R. Guerraoui, and A. M. Kermarrec, "The Many Faces of Publish-Subscribe," *ACM Computing Surveys*, vol. 35, no. 2, pp. 114–131, June 2003.
- [74] V. S. Irava and C. H. Hauser, "Survivable Low-Cost Low-Delay Multicast Trees," in *Proc. 2005 IEEE Global Telecommunications Conference*.
- [75] V. S. Irava, "Low-Cost Delay-Constrained Multicast Routing Heuristics and Their Evaluation," Ph.D. dissertation, Washington State University, Pullman, WA, Aug. 2006.
- [76] GridWise Architecture Council, *Interoperability Constitution Whitepaper (v1.1)*, Dec. 2006.
- [77] L. Roberts, "A Radical New Router," in *Proc. 2009 IEEE Spectrum*, pp. 35–39.
- [78] J. Brassil, R. McGeer, R. Rajagopalan, P. Sharma, P. Yalangadula, S. Banerjee, D. P. Reed, S. J. Lee, A. Bavier, L. Peterson, S. Schwab, L. Roberts, A. Henderson, B. Khorram, S. Zhang, S. Sohn, B. Mark, J. Spies, and N. Watts, "The CHART System: A High-Performance, Fair Transport Architecture Based on Explicit-Rate Signaling," *ACM SIGOPS Operating Systems Review*, Jan. 2009.
- [79] K. Rauschenbach, R. Hain, A. Jackson, J. Jacob, W. Leland, J. Lowry, W. Milliken, P. Pal, R. Ramanathan, and C. Santivanez, "Dynamic Provisioning System for Bandwidth-Scalable Core Optical Networks," in *Proc. 2009 of MilCom*.
- [80] D. Norton, "Security Strategy for NASPI Implementation at Entergy," *presentation at 2009 NASPI Meeting*.
- [81] K. P. Birman, J. Chen, K. M. Hopkinson, R. J. Thomas, J. S. Thorp, R. van Renesse, and W. Vogels, "Overcoming Communications Challenges in Software for Monitoring and Controlling Power Systems," in *Proc. IEEE*, 9(5), May 2005.
- [82] K. Hopkinson, G. Roberts, X. Wang, and J. Thorp, "Quality of Service Considerations in Utility Communication Networks," *IEEE Trans. Power Delivery*, 24(3), July 2009.
- [83] J. Rajahalme, A. Conta, B. Carpenter, and S. Deering, "RFC3697: IPv6 Flow Label Specification," *The Internet Society*, 2004. [Online]. Available: <http://www.faqs.org/rfcs/rfc3697.html>
- [84] A. Farrel (ed), "Network Quality of Service Know It All," *Elsevier*, 2009.
- [85] PGM Reliable Transport Protocol Specification, RFC 3208, IETF, 2001.
- [86] Y. Amir, C. Danilov, and J. Stanton, "A Low Latency, Loss Tolerant Architecture and Protocol for Wide Area Group Communication," in *Proc. 2000 IEEE First International Conference on Dependable Systems and Networking*.
- [87] [Online]. Available: www.spreadconcepts.com.
- [88] K. Birman, "Like It or Not, Web Services Are Distributed Objects!" *Communications of the ACM*, 47:12, pp. 60–62.
- [89] K. Birman, "The Untrustworthy Web Services Revolution," *IEEE Computer*, 39:2, pp. 98–100, Feb. 2006.
- [90] W. Mahnke, S. Leitner, and M. Damm. *OPC Unified Architecture*, May 2009.
- [91] Z. Xie, G. Manimaran, V. Vittal, A. Phadke, and V. Centeno, "An Information Architecture for Future Power Systems and Its Reliability Analysis," *IEEE Trans. Power Systems*, 17:3, pp. 857–863, Aug. 2002.
- [92] K. H. Gjermundrød, "Flexible QoS-Managed Status Dissemination Middleware Framework for the Electric Power Grid," Ph.D. dissertation, Washington State University, Pullman, WA, Aug. 2006.
- [93] E. Solum, C. H. Hauser, R. Chakravarthy, "Modular Over-the-Wire Configurable Security for Long-Lived Critical Infrastructure Monitoring Systems," in *Proc. 2009 3rd ACM Int'l Conf. on Distributed Event-Based Systems*.
- [94] R. Chakravarthy, C. H. Hauser, and D. E. Bakken, "Long-Lived Authentication Protocols for Critical Infrastructure Process Control Systems," *Fourth IFIP WG 11.10 Int'l Conf. on Critical Infrastructure Protection*, Mar. 2010.
- [95] S. F. Abelsen, H. Gjermundrød, D. E. Bakken, and C. H. Hauser, "Adaptive Data Stream Mechanism for Control and Monitoring Applications," in *Proc. 2009 1st International Conference on Adaptive and Self-Adaptive Systems and Applications*, pp. 86–91.
- [96] I. Dionysiou, D. Frincke, C. H. Hauser, and D. E. Bakken, "An Approach to Trust Management Challenges for Critical Infrastructures," *Lecture Notes in Computer Science 5141*, Springer, Berlin, 2007.
- [97] C. H. Hauser, D. E. Bakken, I. Dionysiou, K. H. Gjermundrød, V. S. Irava, J. Helkey, and A. Bose, "Security, Trust and QoS in Next-Generation Control and Communication for Large Power Systems," *International Journal of Critical Infrastructures (Inderscience)*, 2007.
- [98] E. Viddal, D. E. Bakken, K. H. Gjermundrød, and C. H. Hauser, "Wide-Area Actuator RPC Over GridStat With Timeliness, Redundancy, and Safety," in *Proc. 2010 4th International Conference on Complex, Intelligent and Software Intensive Systems*, pp. 17–24.
- [99] K. Swenson, "Exploiting Network Processors for Low Latency, High Throughput, Rate-Based Sensor Updated Delivery," M.S. Thesis, Washington State University, Pullman, WA, 2009.