
Security, trust, and QoS in next-generation control and communication for large power systems

Carl H. Hauser*, David E. Bakken,
Ioanna Dionysiou, K. Harald Gjermundrød,
Venkata S. Irava, Joel Helkey and Anjan Bose

School of Electrical Engineering and Computer Science
Washington State University
P.O. Box 642752, Pullman, Washington, 99164, USA
E-mail: hauser@eecs.wsu.edu
E-mail: bakken@eecs.wsu.edu
E-mail: dionysiou.i@intercollege.ac.cy
E-mail: harald@cs.ucey.ac.cy
E-mail: virava@microsoft.com
E-mail: jhelkey@eecs.wsu.edu
E-mail: bose@eecs.wsu.edu

*Corresponding author

Abstract: The present communication architecture supporting control of the electric power grid makes it difficult to use the wealth of data collected at high rates in substations, retarding their use in new applications for controlling the grid. A flexible, real-time data network would make it possible to use these data for many more control and protection applications, potentially increasing the grid's reliability and increasing its operating efficiency. Applications that could use these data include: decentralised load frequency control; closed-loop voltage control; transient and small-signal stabilisation; and special protection schemes using data gathered over a wide area. Such applications and the flexibility of the underlying communication network imply greater data sharing between utilities, leading to new performance, availability and reliability requirements. This paper examines the security, trust and Quality of Service (QoS) requirements imposed by these applications and shows how they are met by mechanisms included in the GridStat middleware framework that we are developing.

Keywords: critical infrastructures; electric power grid; Quality of Service; QoS; real-time data communication networks; middleware; power control applications.

Reference to this paper should be made as follows: Hauser, C.H., Bakken, D.E., Dionysiou, I., Gjermundrød, K.H., Irava, V.S., Helkey, J. and Bose, A. (2008) 'Security, trust, and QoS in next-generation control and communication for large power systems', *Int. J. Critical Infrastructures*, Vol. 4, Nos. 1/2, pp.3–16.

Biographical notes: 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. Prior to joining WSU, he worked at Xerox Palo Alto Research Center and IBM Research for a total of more than 20 years. He received his BS degree in Computer Science from Washington State University in 1975 and his MS and PhD degrees in Computer Science from Cornell University in 1977 and 1980, respectively.

David E. Bakken received his BS degrees in Mathematics and Computer Science from Washington State University in 1985, and the MS and PhD in Computer Science from The University of Arizona in 1990 and 1994, respectively. He is currently an Associate Professor at Washington State University. He has worked at BBN Technologies (1994–1999) and Boeing (1985–1988). His research interests include middleware frameworks for wide-area systems and fault-tolerant computing.

Ioanna Dionysiou received her BS, MS, and PhD degrees in Computer Science from Washington State University in 1997, 2000 and 2006, respectively. She is currently an Assistant Professor at Intercollege, Cyprus. Her research interests include security and trust in distributed computing systems with emphasis on the publish-subscribe communication paradigm.

K. Harald Gjermundrød received his BS, MS and PhD degrees in Computer Science from Washington State University in 1999, 2001 and 2006, respectively. He is currently a Postdoctoral Associate at the High-Performance Computing Systems Laboratory at the University of Cyprus. His research interests include distributed computing systems, middleware and grid computing.

Venkata S. Irava is a Software Design Engineer in test in the Mobile & Embedded Devices/Communication Sector group at Microsoft. His research interests include multicast routing, mobile networking and distributed computing systems. He received his BE degree from Osmania University, India, in 1999, MS degree from Utah State University in 2002, and PhD degree from Washington State University in 2006.

Joel Helkey is an MS candidate in Computer Science in the School of Electrical Engineering and Computer Science at Washington State University.

Anjan Bose received his Btech (Hons) from the Indian Institute of Technology, Kharagpur, his MS from the University of California, Berkeley, and his PhD from Iowa State University. His career spans power utilities, suppliers and academia. His methods and software for operation and control of the power grid are used in control centres around the world. He is currently the Distinguished Professor in Power Engineering in the School of Electrical Engineering and Computer Science at Washington State University. He is a member of the US National Academy of Engineering and of the US National Research Council's Study Committee on Improving Cybersecurity Research.

1 Introduction

New approaches to controlling the power grid are receiving increased attention in the last few years. Constrained investment in transmission infrastructure and highly visible outage events along with new monitoring and control technologies are producing pressure for a fresh look at the way information about the power grid's operation is

collected and distributed (US-Canada Power System Outage Task Force, 2004; EPRI, 2003; Wu *et al.*, 2005; Tomsovic *et al.*, 2005; Kezunovic *et al.*, 2005). GridStat, a new, flexible approach to providing communication support for electric power grid operations has been previously described by Hauser *et al.* (2005). The goal is to take advantage of modern computer networking and distributed systems knowledge to provide a communication infrastructure that can serve a multitude of communication needs for a wide range of protection and control applications in the power grid. This new approach to power grid communication offers challenges in Quality of Service (QoS), security and trust that do not arise in the power grid's existing, rather fragmented and inflexible cyber infrastructure.

GridStat's flexible architecture is intended to carry communication between substations and control centres that today is typically carried on SCADA systems. Other intended uses include gathering and disseminating Phasor Measurement Unit (PMU) data streams used in novel applications such as detection and remediation of under-damped small-signal instability, disturbance localisation, Remedial Action Schemes (RAS) and Special Protection Schemes (SPS). Meliopoulos *et al.* (2006) suggest that other potential applications include dissemination of substation status based on substation models that use redundant information collected in a single substation to derive the overall status of buses and transmission lines at that substation and other wide-area monitoring and control functions that may be invented in the future.

GridStat is based on a publish-subscribe (pub-sub) distributed system model. Devices in substations periodically publish status and analogue measurements, called *status variables* in the architecture; control centres and devices in other substations subscribe to a selected set of status variables. A publisher may produce data at a higher rate than a subscriber cares to receive them in which case the network will filter the data stream down to the required rate, reducing demands placed on network and subscriber resources. The network supports multiple subscribers to each status variable's stream of data using multicast techniques.

Taylor *et al.* (2003) describe security vulnerabilities that are frequently found in conventional communication systems based on SCADA and point-to-point communication. A pub-sub architecture for power grid communication potentially poses quite different QoS and security requirements than those that arise in conventional power grid communication: the new architecture will support a vastly richer set of interactions between power grid entities than is typical with today's architectures.

Because the new communication system enables many more interactions between many more participants, it has security requirements beyond the conventional Confidentiality, Integrity and Availability (CIA) properties provided by conventional security systems. For example, message integrity and confidentiality services have nothing to say about the quality of the data contained in a message. Nor does a confidentiality service protect against disclosure of a message by an intended recipient. As the community of participants in the power grid's operations grows, properties such as these, which involve the behaviour of participants, not just their identity, become increasingly difficult to deal with. Blaze *et al.* (1996) and Grandison and Sloman (2000) identify these properties as *trust properties*. *Trust* is defined as quantified belief by a *trustor* in the competence and dependability of a *trustee* to perform a specific action within a specific context. *Trust management* is the process by which a trustor develops and maintains its trust beliefs (Dionysiou, 2006).

It is perhaps not surprising that security, trust, and QoS requirements interact in interesting ways. The remainder of this paper exposes common themes in the security, QoS, and trust requirements imposed on the architecture by some of its intended applications and considers their consequences for the implementation of the architecture.

2 Application requirements for security, QoS and trust

In present practices, different communication needs in the power grid are met using quite different, and separate, technologies. A flexible communication infrastructure, if it is to fulfil its promise, will have to meet the security, QoS and trust needs of many kinds of applications. Some of the intended current and envisioned future applications are described below along with their currently understood security, QoS, and trust requirements.

2.1 *Communication between substations and control centres*

Communication between substations and control centres today is the domain of SCADA systems. Equipment at substations, including so-called Intelligent Electronic Devices (IEDs), reports status and measurements, such as voltage, current and power transfer, to a SCADA Remote Terminal Unit (RTU) in the substation. A SCADA master at the control centre polls the RTU every few seconds to retrieve the measurements. An Energy Management System (EMS) at the control centre displays the measurements and status to operators. The operators, in turn, issue commands to the substation equipment (*e.g.*, to open or close a breaker, or change transformer tap).

Even this very simplified description reveals fundamental requirements for the performance of the control centre:

- Control centre displays for operators must accurately reflect the system state so that control decisions are appropriate.
- Substation equipment must carry out legitimate commands, and only legitimate commands, within specified time delays.
- Certain control decisions and state information are commercially or otherwise sensitive and are not to be revealed to unauthorised parties.

On the other hand, the two-second (or even four-second) SCADA polling cycle and the almost entirely manual control system suggest that sub-second latency is *not* a requirement for these communications in today's systems and today's applications. For the future, though, the power industry recognises the need for better communication and the increased visibility into their operations that it can provide.

Historically, the obstacles that the power industry has addressed in meeting the first two requirements have been primarily ones associated with the reliability of substation equipment and communication systems. Problems include sensor failures, communication link failures, misconfiguration of substation equipment and databases that describe it, *etc.* Solutions include redundancy of communication links and use of software technology called *state estimators* to form an accurate picture of the system state based on partially missing or incorrect information.

These solutions primarily address threats to the requirements posed by unreliability of equipment and errors by human beings. In recent years, malicious interference is increasingly recognised as a threat to achieving the requirements. To date, the security approaches taken to prevent malicious interference are based on attempting to guarantee that the SCADA system is a closed, isolated system and on that basis making assumptions such as:

- Physical access to substation equipment is limited to authorised parties.
- Physical access to the communication hardware and links is limited to authorised parties.
- Authorised parties are always trustworthy.
- SCADA networks are not interconnected with networks to which unauthorised parties have access.
- Because links are used exclusively for SCADA communication there are no issues associated with allocation of bandwidth for different purposes.

Taylor *et al.* (2003) found that there is considerable accumulated evidence that existing SCADA systems do not satisfy these assumptions leading to the conclusion that the electric power infrastructure is threatened by malicious attack delivered through its control system. In the last decade the extent to which the closed system assumptions were not being met has been recognised and considerable effort by utilities has gone into remedying the situation. New operational and auditing practices have improved the situation. More recently, products providing ‘bump-in-the-wire’ encryption have become economically feasible and have seen increasing deployment.

The trend towards operating transmission systems closer to their limits, towards outsourcing of key maintenance and configuration operations, and towards separation of the businesses of generation and transmission would be better served by the more flexible communication infrastructure discussed here. However, the security assumptions of existing SCADA systems are in conflict with the evolving need for communication that crosses organisational and geographic boundaries. For the new communication architecture it is appropriate, therefore, to start not by imposing closed system assumptions but rather by looking at the required operational behaviour and developing the security, QoS and trust requirements from there.

Requirement 1 The control centre displays accurately reflect the system state (so that control decisions are appropriate).

To meet this requirement the system must ensure the integrity of data that is collected at sensors and delivered to control centres and it must ensure that sensor readings are delivered within a specified time interval from when they are collected. Meeting this requirement involves both end-to-end integrity and quality of service. Many organisations, such as the US-Canada Power System Outage Task Force (2004), are recognising that solving the situation awareness problem for operators requires that control centres have up-to-date information about the state of facilities in surrounding control areas. Operators must be convinced of the trustworthiness of the neighbouring information if they are to take actions based upon it. The operators’

trustworthiness assessment may change as evidence accumulates that a party is reporting consistently and truthfully, based on other party's reports, or that the party is reporting inaccurate information.

Public-key-based signing techniques are a potential solution to the integrity problem, but there are concerns about the computational ability of low-end and legacy sensor devices to carry out the public-key algorithms, and about the latency costs associated with the algorithms. These problems are compounded by the need to share collected data between multiple control centres and the desire to aggregate sensor measurements in substations.

Ensuring end-to-end latency of sensor data requires allocation of network resources to different source-destination flows and ensuring that the network capacity is sufficient to meet all the requirements. It is likely that the set of flows and their latency requirements will change as the operational state of the power grid or communication infrastructure changes, so the network must adapt in the face of contingencies such as link failures and cyber-attacks. The existing practice of using state estimators to augment sensor data and to protect against erroneous data provides some resilience against failure to receive information in a timely way.

The importance of trust and trust management often goes unrecognised, but a critical infrastructure's communication system should help establish appropriate trust in the information it delivers. The greater the organisational and physical distance between the communicating entities the more important this becomes. For example, in the August 2003 blackout some operators were reluctant to believe that transmission lines had tripped out based on reports from neighbouring operators, because that information conflicted with information reported by their own system (US-Canada Power System Outage Task Force, 2004).

Requirement 2 Substation equipment carries out legitimate, and only legitimate, commands within specified time delays.

As with Requirement 1, Requirement 2 poses the need for integrity of messages along with a latency requirement. The latency requirements are perhaps more stringent because while the control centre might be able to substitute information derived from other sensors for missing inputs, it may be quite difficult to substitute different actions when a command is not carried out. Because authority to control a device is limited to a single control centre, trust is not so much of an issue for this requirement.

Requirement 3 Certain control decisions and state information are commercially or otherwise sensitive and are not to be revealed to unauthorised parties.

Requirement 3 is a confidentiality requirement. Cryptographic techniques provide a foundation on which to build a solution. More difficult problems likely lurk in techniques for distinguishing between authorised and unauthorised parties across business entity boundaries and for dynamically adapting the sets of unauthorised and authorised parties for particular information as the operating status of the power grid changes. This requirement also reveals the other side of the trust problem: disclosure to certain parties may be desirable only if their trustworthiness can be established. From the information provider's perspective, knowing the identity of the recipient is not enough – the provider also wants to know whether the recipient can be trusted.

Communication between substations and control centres is only one aspect of the power grid's cyber communication needs. While it reveals most of the qualitative aspects of the requirements and has been treated in depth, the applications that follow provide additional insight into the scope of the QoS, security, and trust needs in a flexible infrastructure.

2.2 Communication between control centres

In present-day, systems state information is shared between the EMSs of neighbouring control areas using the Inter-Control Center Protocol (ICCP). ICCP allows a utility to give access to specific items in its EMS database to authenticated EMSs at other utilities. Connections can be made over private networks or virtual private networks. The shared items may either be pulled by the receiving EMS or pushed by the sending EMS. Latencies of several seconds are acceptable.

Today, operators also rely on telephone communication with operators in nearby areas to develop their overall awareness of the operating state of the grid. In an emergency situation, such as occurred leading up to the August 2003 US blackout, operators may find it difficult to acquire enough reliable information over the telephone: it has limited bandwidth and a two-way conversation may not focus the operator's attention in the best possible place for acquiring needed information.

2.3 Collection and dissemination of phasor data

Synchronous PMUs measure current and voltage at precisely determined times, many times per second. There is considerable interest in the power industry towards applying these synchronous measurements to new control solutions to stabilise the grid and allow it to be operated more efficiently. Cai *et al.* (2005) discuss the Wide Area Monitoring System (WAMS) in the Western USA as the first widespread deployment of PMUs. The Eastern Interconnect Phasor Project (EIPP) is currently deploying PMUs in the Eastern grid. It is setting up infrastructure for collecting the data, but is only in the early stages of considering the data's use for monitoring and control applications.

In current practice as described by Cai *et al.* (2005) and Carroll (2005) an independent network is set up to carry PMU measurements to devices called Phasor Data Concentrators (PDCs) where the full data streams from many devices are time-synchronised, stored for future reference, and forwarded to applications and to Super PDCs.

2.4 Wide-area control and monitoring

Another use of relays in today's grid is to protect the grid against transient instability, a phenomenon that occurs when the grid topology suddenly changes due to unexpected loss of a large transmission line, generator, or load. When a large change occurs, power flows over a wide geographic area may begin to oscillate. If unchecked, the oscillations can damage generation and transmission equipment. To prevent damage, protective devices associated with major assets will disconnect them from the grid. Unfortunately, losing too many generators or transmission lines in this way causes break-ups of the grid and power outages.

To prevent outages due to this phenomenon, power system engineers have created Special Protection Schemes (SPS) that attempt to prevent the oscillations from growing, instead of waiting for the oscillations to force critical equipment to trip offline. A special protection scheme involves taking a specific action, such as tripping a generator, soon after a specific triggering event that may occur hundreds of kilometres away. Because the damaging oscillations have low frequency (often substantially below 1 Hz) and the concern is preventing the build-up in amplitude of oscillations over several cycles, the communication timing requirements for special protection schemes are less stringent than those for local protective relaying. Support for SPS is a reasonable goal for a wide-area power grid communication architecture.

Because a SPS provides protection against only a single, predefined (and hard-coded) event at high cost and complexity, recent research by Taylor *et al.* (2005) has been investigating approaches that allow reaction based on sensing the response of the power grid to arbitrary disturbances, followed by corrective action such as capacitor bank switching, generator tripping, or load shedding. Related ideas in Quintero and Venkatasubramanian (2005) include detection and mitigation of small-signal instability, and fault location by detecting and triangulating based on frequency disturbances (Zhong *et al.*, 2005). These schemes detect anomalous behaviour of the grid by monitoring the outputs of several PMUs (or frequency measurement devices in the case of Zhong *et al.*, 2005) and deriving indications of problems such as low voltage, incorrect frequency, or oscillations. Based on the location of the problem, specific discontinuous or continuous control actions can be taken. In the Western US power grid, it is estimated that control action needs to be taken within about 1 second of the occurrence of the disturbance to suppress transient instability.

2.5 *Communication for local protective relaying*

Protective relays monitor transmission line conditions. When faults occur they are responsible for tripping circuit breakers that de-energise the line (or other equipment). Relays may also attempt to re-close breakers; if, as often occurs, the fault was transient the equipment can be put back in service in a few seconds. The measurements for protective relaying are usually made only at the substation where the breaker is located or in a neighbouring substation. The communication systems associated with relaying are, today, entirely separate from the systems used for control centres and PMU data collection. They have very stringent time requirements (only a few milliseconds latency). Local protective relaying is likely to be one of the last applications to adopt a common communication infrastructure; indeed, due to its great importance for protection of life and property and its stringent timing requirements it may never do so.

Relay settings (tripping rules) are usually statically determined from off-line analysis and simulation of the associated equipment. Better communication of overall grid conditions has the potential to allow settings to be less conservative in some operational circumstances, in turn allowing greater use of available transmission resources. Such an application would have QoS and security requirements more similar to control centre requirements than to those for protective relaying.

2.6 Summary of applications' requirements

Message CIA show up as requirements across the spectrum of applications. Hard latency requirements range from a few milliseconds for protective relaying to a few tens of milliseconds for automated dynamic stabilisation applications and a few seconds for conventional control centre applications. Applications that use inputs from a large, varying set of suppliers or make their outputs available to a large, varying set of consumers are most in need of trust management support. The communication architecture, itself, suggests that the trust management system must include support for indirect interactions – chains of processing that involve multiple participants (including the communication infrastructure).

3 Implications of the requirements

The requirements identified above have several implications for QoS, security and trust and their interactions which we now elaborate.

The communication infrastructure should support CIA properties for all messages as needed by the applications. The major challenges to doing this are:

- providing these properties using computationally underpowered devices
- establishing an authentication infrastructure that efficiently supports the operational needs of the utility industry while providing authentication for people working in many different businesses and organisations in many different roles *and* for thousands of different devices
- establishing and maintaining policies that allow appropriate monitoring and control to occur.

While we do not propose that power grid communication be implemented using the *public* internet, the cost of *internet technology* makes its use in *private* networks very attractive. We assume, for the time being, a private cyber infrastructure for the power grid based on internet technology and we look specifically at the security implications of the power grid control requirements for such a network.

The hard real-time latency requirements imply that network resources must be actively managed: best effort service is not good enough. Both Zhang (1995) and Goyal *et al.* (1996) found that first-come first-serve is not adequate as a scheduling discipline in nodes of the forwarding network if end-to-end latency requirements must be met: a guaranteed-rate packet scheduling discipline is required. For example, consider a hypothetical substation producing 60 different status and analogue values once per second and 3 PMU measurements 60 times per second. Following Tomsovic *et al.* (2005) the PMU measurements might be used in a remote action scheme that helps protect the power grid and therefore have an end-to-end latency requirement of 10 msec. This latency must be allocated between transmission, propagation, and queuing delays of multiple hops between the producer and consumer of the data. It is not unreasonable, therefore, to set a 1 msec queuing and transmission latency requirement at the originating station. Assume that the substation is connected to the data network using data link of about T1 speed and that, with overhead, each of the 63 published values requires about

400 bit times to transmit (250 μ sec). With an unmanaged network a PMU measurement might be delayed as much as 15.5 msec (worst case analysis). Power grid measurements become more effective when they are taken synchronously across the network, so there is actually a high probability of these measurements arriving at the queue at the same time. In simulations of this scenario the 60/second measurements frequently suffer delays in excess of 1 msec. Interestingly, better real-time performance of the substation systems leads to closer arrival times at the queue and worse queuing delays. With earliest-due-date scheduling for the outgoing link all 3 PMU measurements can be scheduled to meet a 1 msec deadline. Note that in this scenario the T1 link is busy only 6% of the time, yet the latency for the PMU measurements is unacceptably large. This suggests that solving latency problems with ever-higher bandwidth connections is an expensive proposition. Furthermore, without admission control, the system remains susceptible to overload that could disrupt the real-time control of the power grid.

A corollary of these observations is that the public internet, lacking both admission control and guaranteed-latency delivery, cannot completely supplant private data networks for the power grid's control communication infrastructure. Policy mechanisms will be needed as part of the communication infrastructure to ensure that subscriptions needed for safe operation of the power grid are always admitted. Also, determining which subscriptions are most needed and which are merely nice to have will be a part of the engineering process in designing control systems in the future. The question of which subscriptions are most needed is compounded when one considers that the answer may change as power flows change.

Even on a private IP network the standard TCP byte-stream protocol is inappropriate for most of the applications described here. Although TCP provides reliable delivery, its delivery model – do not deliver byte n to the application until byte $n-1$ has been delivered – and its retransmission, window management, and timeout strategies geared to congestion control render it unsuitable for real-time control.

The design of the communication infrastructure must also consider malicious threats to meeting QoS requirements. Denial of service through consumption of network resources is a threat that must be countered. Even though use of the public internet has already been shown to be inappropriate, the control network design should suppose that at least from time to time, if not permanently, the two networks become connected. It must be assumed therefore that packet injections are possible so denial of service attacks cannot be prevented, but must be thwarted. Security mechanisms will have to be leveraged to prevent unwanted traffic from consuming network resources.


4 GridStat

For the past five years, the GridStat project, described in detail by Gjermundrød (2006), has been designing a flexible power grid communication architecture and creating a publish-subscribe middleware framework instantiation of it. GridStat was designed to provide flexible communications with managed QoS. Flexibility in communications means that the system will provide legitimate subscribers with the status data they request without any hard-coding of subscribers, publishers, routes the data will take, *etc.* These aspects are all handled by the GridStat software. This allows a wide variety of control and protection schemes to be deployed, including ones designed in the future, while still

using the same communications infrastructure. QoS in communications means that the status data will be delivered in a timely manner over redundant paths which are managed by GridStat. GridStat's delivery infrastructure was designed to be able to be ported to a wide variety of underlying network technologies.

GridStat's architecture is illustrated in Figure 1. GridStat middleware partitions the network functions into a management plane, consisting of *QoS Brokers*, and a data plane consisting of *status routers*.

Figure 1 GridStat architecture



The QoS Brokers allocate resources for subscriptions, create multicast routes to be loaded into the status routers, perform admission control, and monitor the data plane, adapting it to changing circumstances. For example, a QoS Broker can change the set of active subscriptions in response to a change in the operational status of the power grid; or it can reroute subscriptions around a failed status router or communication link.

The status routers implement a multicast-based, real-time, publish-subscribe model for providers and consumers of operational data: breaker status, bus voltage, line power transfer, *etc.* Unlike IP routers (Network layer 3), status routers are session-aware (subscriptions) and can even incorporate application-level functionality, for example, by performing data aggregation using GridStat's *condensation function* mechanism. A condensation function allows a status router to combine multiple input streams into a single output stream using a built-in or user-supplied combination rule.

Status routers incorporate a mechanism for pre-configuring sets of subscriptions (subscription modes) corresponding to different operating conditions of the power grid and making smooth transitions between them.

In the GridStat prototype the QoS Brokers demonstrate disjoint-path, multicast routing, hierarchical network management and network monitoring. To increase delivery reliability, a subscription request can specify that data for the subscription be delivered over multiple, disjoint paths. Irava and Hauser (2005) and Irava (2006) discuss research on routing heuristics for efficient, delay-constrained, disjoint-path multicast. The status routers incorporate the ability to perform multicast forwarding, subscription modes, and

condensation functions. Research based on Zhang (1995) and Goyal *et al.* (1996) is currently under way to add guaranteed-rate forwarding to the status routers to achieve guaranteed end-to-end latency.

The status variable abstraction is provided to GridStat publishers and subscribers through middleware libraries. Recall that a status variable is a periodic sequence of time-stamped values produced by a single publisher and consumed, potentially, by many subscribers. The status variable abstraction relieves the application of some of the more difficult issues that arise in programming distributed applications. At the same time, the fact that publications are periodic makes it possible for resource scheduling in the status routers to provide guaranteed end-to-end delay.

GridStat provides two alternative APIs for subscribers to obtain status data. The *Pull-from-Cache* API lets the application obtain the value of the most recently received, timestamped, update from a cache maintained by GridStat. This lets a subscriber use the value like a local variable in its program, and be shielded from having to deal with the asynchronous arrival of update messages. The *Direct Push* API lets the subscriber receive each update. This is useful for inserting the updates into a relational database. Finally, subscribers can also optionally request a *QoS Push*, a callback that lets it know that the QoS it requested has been violated (for example, the latency is too high or the rate too low).

5 Ongoing work

A new centre-scale project was initiated in August 2005 by the US National Science Foundation (with funding from the US Departments of Energy and Homeland Security) in order to address many of the issues discussed in this paper. The centre, Trustworthy Cyber Infrastructure for the Power Grid (TCIP), involves computer scientists and power engineering researchers from the University of Illinois, Cornell University, Dartmouth College, and Washington State University (TCIP Project, 2006). TCIP researchers are conducting fundamental research which will be applied to wide-area communications infrastructures, such as GridStat, in order to better meet the needs described in this paper. Particular areas of research include: hardware, software, and firmware to provide a reliable and secure computing base that can protect existing substation devices against accidental failures and malicious attacks; providing data communications within substations and at higher levels (*e.g.*, to control centres and between them) which offers more security but also quantified tradeoffs between security, performance, and trust; secure data aggregation techniques to deal with malicious and other incorrect data inputs; and quantitative validation of the above technologies using simulation and modelling.

6 Conclusion

Security and trust requirements for the power grid's cyber infrastructure are derived from the need to monitor and control grid components by multiple interacting business entities, with both cooperative and competitive interests. Creating the power grid's cyber infrastructure as an isolated network has obvious appeal, but does not solve problems associated with insider threats, nor is complete isolation likely to be achieved and maintained at all times.

Security techniques for confidentiality, integrity and availability are important components of any solution. The need to support interactions involving devices with low computational power, and to support many business entities with many people in each business pose significant challenges to the application of well-known technologies such as public key encryption and message signing.

Thwarting denial of service attacks is a key need in achieving the required quality of service. A power grid communication infrastructure that is not intended to be open to all comers (as the public internet is) will allow controls to be placed on traffic so that the network can be protected from denial of service attacks: to work, this must be true even when cross connections to the public internet occur, as they inevitably will.

Acknowledgement

This work has been supported in part by US National Science Foundation Grants CCR-03-26006 and CNS-05-24695.

References

- Blaze, M., Feigenbaum, J. and Lacy, J. (1996) 'Decentralized trust management', *Proceedings of the 1996 IEEE Symposium on Security and Privacy*, Washington, DC, USA, pp.164–173.
- Cai, J.Y., Huang, Z., Hauer, J. and Martin, K. (2005) 'Current status and experience of WAMS implementation in North America', *2005 IEEE/PES Transmission and Distribution Conference and Exhibition: Asia and Pacific*, Dalian, China, pp.1–7.
- Carroll, J. (2005) 'Eastern interconnect phasor project: TVA's super phasor data concentrator architecture', *Presentation at the Eastern Interconnect Phasor Project Meeting*, Chattanooga, TN, USA, April, http://phasors.pnl.gov/Meetings/2005%20April/presentations/TVAs%20Super%20Phasor%20Data_Carroll_EIPP.ppt (retrieved February 2006).
- Dionysiou, I. (2006) 'Dynamic and composable trust for indirect interaction', PhD dissertation, Washington State University.
- Electric Power Research Institute (EPRI) and Electricity Innovation Institute (2003) *The Integrated Energy and Communication Systems Architecture (Intelligrid Architecture)*, Palo Alto, California, USA, http://intelligrid.info/IntelliGrid_Architecture/Overview_Guidelines/Adl_Deliverables_List.htm (retrieved February 2006).
- Gjermundrød, K.H. (2006) 'Flexible QoS-managed status dissemination middleware framework for the electric power grid', PhD dissertation, Washington State University, August.
- Goyal, P., Lam, S.S. and Vin, H.M. (1996) 'Determining end-to-end delay bounds in heterogeneous networks', *ACM/Springer-Verlag Multimedia Systems Journal*, Vol. 5, No. 3, pp.157–163.
- Grandison, T. and Sloman, M. (2000) 'A survey of trust in internet applications', *IEEE Communications Surveys and Tutorials*, Vol. 3, No. 4, pp.2–16.
- Hauser, C.H., Bakken, D.E. and Bose, A. (2005) 'A failure to communicate: next generation communication requirements, technologies and architecture for the electric power grid', *IEEE Power and Energy Magazine*, Vol. 3, No. 2, pp.47–55.
- Irava, V.S. (2006) 'Low-cost delay-constrained multicast routing heuristics and their evaluation', PhD dissertation, Washington State University.
- Irava, V.S. and Hauser, C. (2005) 'Survivable low-cost low-delay multicast trees', *Proceedings of the IEEE Global Telecommunications Conference (Globecom '05)*, St. Louis, MO, USA.

- Kezunovic, M., Djoki, T. and Kosti, T. (2005) 'Automated monitoring and control using new data integration paradigm', *Proceedings of the 38th Hawaii International Conference on System Sciences (HICSS 2005)*, Hawaii, USA.
- Meliopoulos, A.P.S., Cokkinides, G.J., Galvan, F. and Fardanesh, B. (2006) 'GPS-synchronized data acquisition: technology assessment and research issues', *Proceedings of the 39th Hawaii International Conference on System Sciences (HICSS 2006)*, Hawaii, USA.
- Quintero, J. and Venkatasubramanian, V. (2005) 'A real-time wide-area control framework for mitigating small-signal instability in large electric power systems', *Proceedings of the 38th Hawaii International Conference on System Sciences (HICSS 2005)*, Hawaii, USA.
- Taylor, C., Oman, P.W. and Krings, A.W. (2003) 'Assessing power substation network security and survivability: a work in progress report', *Proceedings of the International Conference on Security and Management (SAM '03)*, Las Vegas, NV, USA.
- Taylor, C.W., Erickson, D.C., Martin, K.E., Wilson, R.E. and Venkatasubramanian, V. (2005) 'WACS-wide-area stability and voltage control system: R&D and online demonstration', *Proceedings of the IEEE*, Vol. 93, No. 5, pp.892–906.
- Tomsovic, K., Bakken, D.E., Venkatasubramanian, V. and Bose, A. (2005) 'Designing the next generation of real-time control, communication and computations for large power systems', *Proceedings of the IEEE*, Vol. 93, No. 5, pp.965–979.
- TCIP Project (2006) 'Trusted Cyber Infrastructure for the Power Grid overview', <http://www.iti.uiuc.edu/tcip/> (retrieved January 2007).
- US-Canada Power System Outage Task Force (2004) *Final Report on the August 14, 2003, Blackout in the United States and Canada: Causes and Recommendations*, US Department of Energy, Washington, DC.
- Wu, F.F., Moslehi, K. and Bose, A. (2005) 'Power system control centers: past, present and future', *Proceedings of the IEEE*, Vol. 93, No. 11, pp.1890–1908.
- Zhang, H. (1995) 'Service disciplines for guaranteed performance service in packet-switching networks', *Proceedings of the IEEE*, Vol. 83, No. 10, pp.1374–1396.
- Zhong, Z., Xu, C., Billian, B.J., Zhang, L., Tsai, S-J.S., Connors, R.W., Centeno, V.A., Phadke, A.G. and Liu, Y. (2005) 'Power system frequency monitoring network (FNET) implementation', *IEEE Transactions on Power Systems*, Vol. 20, No. 4, pp.1914–1921.