

# Designing the Next Generation of Real-Time Control, Communication, and Computations for Large Power Systems

---

KEVIN TOMSOVIC, SENIOR MEMBER, IEEE, DAVID E. BAKKEN, MEMBER, IEEE,  
VAITHIANATHAN VENKATASUBRAMANIAN, MEMBER, IEEE, AND ANJAN BOSE, FELLOW, IEEE

## *Invited Paper*

*The power grid is not only a network interconnecting generators and loads through a transmission and distribution system, but is overlaid with a communication and control system that enables economic and secure operation. This multilayered infrastructure has evolved over many decades utilizing new technologies as they have appeared. This evolution has been slow and incremental, as the operation of the power system consisting of vertically integrated utilities has, until recently, changed very little. The monitoring of the grid is still done by a hierarchical design with polling for data at scanning rates in seconds that reflects the conceptual design of the 1960s. This design was adequate for vertically integrated utilities with limited feedback and wide-area controls; however, the thesis of this paper is that the changing environment, in both policy and technology, requires a new look at the operation of the power grid and a complete redesign of the control, communication and computation infrastructure. We provide several example novel control and communication regimes for such a new infrastructure.*

**Keywords**—Distributed control, power systems, quality of service (QoS), Supervisory Control and Data Acquisition (SCADA), wide-area control.

Manuscript received May 2, 2002; revised October 25, 2003. This work was supported in part by the Consortium for Electric Reliability Technology Solutions (CERTS), funded by the Assistant Secretary of Energy Efficiency and Renewable Energy, Office of Distributed Energy and Electricity Reliability, Transmission Reliability Program of the U.S. Department of Energy under Interagency Agreement DE-AI-99EE35075 with the National Science Foundation (NSF); in part by the National Institute of Standards and Technology's (NIST's) Critical Infrastructure Protection Program and NSF's Embedded and Hybrid Systems Program under Grant CCR-0326006; in part by the Electric Power Research Institute/Department of Defense (EPRI/DOD) Complex Interactive Networks Initiative; and in part by the Bonneville Power Administration.

The authors are with the School of Electrical Engineering and Computer Science, Washington State University, Pullman, WA 99164 USA (e-mail: tomsovic@eeecs.wsu.edu; bakken@eeecs.wsu.edu; mani@eeecs.wsu.edu; bose@wsu.edu).

Digital Object Identifier 10.1109/JPROC.2005.847249

## I. INTRODUCTION

The power system networks in developed countries are extremely large and complex interconnected systems. The Eastern interconnection in North America is the largest in geographic area (stretching from the east coast to almost the Rocky Mountains), total installed generation capacity, and total miles of transmission lines. In these networks, all the generators, with the exception of a small number of renewable units, rotate synchronously in dynamic equilibrium, precisely matching generation to demand. This balance must be maintained as loads fluctuate and as disturbances such as equipment outages arise. The multilayered infrastructure to enable economic and reliable operation of such systems has evolved over many decades, utilizing new technologies as they have appeared. This evolution has been slow and incremental, partly due to the great capital expense of making changes to such a large system, but also since the operation of the power system stemming from vertically integrated utilities has, until recently, changed very little.

Large power systems exhibit a variety of dynamical behaviors ranging from very slow, minutes to hours, to extremely fast phenomena, on the order of milliseconds. Separate controllers, generally designed independently, have been developed to address the various phenomena. Control actions include both the discrete, such as the opening of a circuit breaker, and the continuous, such as the excitation control in synchronous generators. Still, except for the very slowest of controllers and a few specialized schemes, engineers have designed systems largely through local decisions based on local measurements. On the other hand, as is all too apparent during blackouts, power system disturbances may have system-wide impacts. At present, during extreme events, the burden falls on operators but, as transcripts from recent blackouts show, they do not have "situation awareness" outside their own company [1].

The tremendous cost and performance breakthroughs in the last decade in terms of computers, their networking infrastructure, and middleware technologies have created new opportunities for designing wider area controls to help overcome these limitations. If the only motivation were to take advantage of new communication technology to improve controls, then perhaps there would be no urgency to redesign the system and power control could continue to evolve slowly as it has in the past; however, in addition to the new technology concerns, the evolving regulatory environment is fundamentally altering the operation of the system. These factors include:

- the increasing number of participants whose different goals in the operation of the power grid require access to different sets of real-time (or historical) data;
- the emerging power markets which produce a completely new set of real-time data;
- the increase, by several orders of magnitude, in transactions between buyers and sellers that not only requires the tracking of more data but also produces completely unanticipated loading of the grid;
- the operation of the power grid in unanticipated operating conditions that requires more sophisticated monitoring by the regional transmission operator/independent system operator (RTO/ISO) to guarantee secure operation, where the RTO/ISO has to not only guarantee operational security but fair transmission access rules that are transparent;
- the recent blackouts and concerns with terrorism that have made it clear that security must be guaranteed against both natural (acts of God) contingencies and malicious acts;
- the new-generation technologies that are changing the generation mix to one with a large number of smaller units (such as wind units and microturbines) operating in concert with the traditional large generation plants.

In this paper, we propose a new communication and control infrastructure to address these concerns. We develop a comprehensive system that is flexible enough to manage new requirements as they evolve. The system described here not only takes advantage of new broadband, faster communication but also emerging middleware systems that use distributed computation to make such systems flexible, efficient and secure. We demonstrate how such a system can allow the implementation of improved controls that will allow the grid to be operated closer to its limits. For example, we show how the next generation of fast, system-wide controls can be added to the slow, system-wide controls [like automatic generation control (AGC)] or fast, local controls (like protection). Further, this new conceptual design is not dependent on future technology but can be implemented with available technology, such as satellite and optical communications, substation automation, voltage and transient stability computation, and microprocessor protection systems.

## II. EXISTING POWER CONTROL SYSTEMS

Given the complexity of the power system and its long history, one may not be surprised that a patchwork of controls

has developed over time. This section surveys these controls and identifies the technology or regulatory environment that led to the engineering decision. A more detailed discussion of this can be found in [2].

### A. Local Controls

Local controls are those that use inputs from within the same substation in which the control outputs are applied (i.e., no distance communication connection is required). The inputs are usually analog values (voltage, current, power, reactive power) and the control can be digital (opening/closing circuit breakers) or settings that are either discrete like transformer taps or continuous like the voltage regulator of a generator.

1) *Power System Protection:* The most fundamental control that is needed to operate the power system stems from safety concerns and the practical necessity of protecting electrical equipment from high currents that arise from faults. The simple fuse has evolved to today's sophisticated microprocessor based relays that control circuit breakers. Even the isolation of equipment from faults requires sophisticated equipment to sense the currents, recognize the faulted conditions, and coordinate the different protective devices. Relays may operate for a number of other reasons beyond simply high currents, including excessive frequency deviation, low or high voltages, or indications of instability. Thus, although not traditionally thought of in this way, protection is the front line of control. Modern protective systems have the ability to respond to a fault within two or three cycles (about 50 ms in a 60-Hz system) using microprocessor-based relays that perform sophisticated computations. Protection is the fastest control but accordingly has been the most localized of the controls.

2) *Generator Governor Control:* The power output of a generator is controlled by a governor that senses the mechanical shaft speed and adjusts the mechanical power input to the generator so that the generator can maintain the desired electrical output.

3) *Voltage Control:* Voltage control has traditionally been a local feedback control. Generator excitation can be controlled to maintain a set terminal voltage by producing or absorb reactive power. Within the transmission network, voltage can also be maintained by the use of tap-changing transformers or switched capacitors/reactors to provide/absorb reactive power. These are generally slower controls that either follow fixed time schedules or are switched at preset levels. More recently, power electronics have been used as a fast control of the production/absorption of reactive power.

4) *Power Flow Control:* Traditionally, the only way to control the power flow on an ac transmission line was with a phase-shifting transformer. This was a slow control with the same speed as tap-changing transformer control of voltage. Power flow on dc transmission lines could, however, be controlled very fast by changing the inverter/converter settings. More recently, power electronics have been used to develop fast controllers for ac transmission lines.




Fig. 1. Existing SCADA infrastructure.

5) *Power System Stabilizers (PSSs)*: Power systems exhibit undamped oscillations under certain operating conditions. Adding a PSS to a generator can help damp such oscillations (these oscillations typically range from 0.7 to 1.4 Hz.). The PSS uses local measurements to provide a supplement to the generator voltage control system.

### B. Wide-Area Controls

In this section, wide-area controls have included any control that requires some communication link to either gather the input or to send out the control signal. Obviously, some wide-area controls require more communications than others.

1) *Frequency Control*: Perhaps the earliest example of wide-area feedback control in the power system is frequency control. Frequency will deviate from the nominal setting whenever there is any imbalance between generation and load as the imbalance will be drawn from the kinetic energy stored in the rotating masses of the generators. The primary control loop to ensure balance is the local governor control (Section II-A2). In order to maintain system frequency at the nominal value and schedules between control areas, a secondary control loop, AGC, coordinates the individual raise and lower signals to the generators. The control center gathers the relevant frequency and power flow information and sends the appropriate set point adjustments for each of the units on AGC. This control varies widely throughout the world with some regions employing a tightly regulated control, e.g., North America, with commands sent typically every 2–4 s, while other regions allow much greater deviations in frequency and prefer less frequent adjustments, e.g., Scandinavia.

2) *Secondary (Regional) Voltage Control*: Some countries have implemented, or begun to implement, secondary voltage control schemes that are analogous to AGC control for frequency. These secondary schemes typically monitor voltage at key buses and then schedule voltage over a wider region to maintain the key bus voltage. These voltage set-points are then sent to the local voltage controllers (Section II-A3). This type of voltage can also be supplemented by a tertiary scheme of scheduling voltages throughout the system, typically to optimize security or minimize losses.

3) *Special Protection Schemes (SPSs)/Remedial Action Schemes (RASs)*: The one exception to local protection schemes has been highly specialized schemes, called RASs or SPSs, which respond by tripping equipment that are remote to some fault or other event. Such a scheme can be as simple as just a pilot-wire trip but are increasingly becoming

elaborate schemes that trip generators, loads or even transmission lines when a combination of events occur. The use of SPS/RAS is increasing in stability limited systems where the availability of such schemes allow power transfers that would engender instability for certain faults unless such a scheme exists. These schemes are generally hardwired and inflexible, relying on dedicated communication links and extensive offline calculations of settings.

### C. Communications for Wide-Area Controls

As the above discussion indicates, most of the power system controls are local and this is particularly so for the faster controls. The only common wide-area feedback control system is AGC, which resides at the control center. The regional voltage control (used in Europe) and the SPS/RAS control have dedicated communication and computer systems for each scheme.

1) *Control Centers*: The predominant function of the control center is an interface between the operators and the power system. Data acquisition allows the operators to monitor the condition of the system and implement supervisory (manual) controls, such as the opening and closing of circuit breakers, switching in and out capacitors, and so on. Such a Supervisory Control and Data Acquisition (SCADA) system gathers system-wide data using a communication system shown in Fig. 1. Such a “star” network gathers data sequentially from each substation and generating station with a scan rate in the 2–10-s range. The fastest scan rates (2–4 s) are used to collect the data needed for AGC (tie-line power flows and generator outputs), and the control signals to the generators are sent out at that rate. As mentioned before, this is the main wide-area control in use today and is a relatively slow control.

The control center also does other tasks, like that of scheduling power transactions, that are managed by the operators. The energy management system (EMS) software in the control center provides a number of computational tools to assist the operators in reaching their decisions, but very little, if any, of this is implemented as a closed-loop control.

The SCADA software communicates directly to the remote terminal units (RTUs) in each substation. Each RTU gathers this data from the various sensors in the substations that measure voltage, current, power, reactive power, as well as circuit breaker positions (open/close). The gathering of data within a substation has in recent years become more sophisticated especially at the higher voltages. Each substation has its own SCADA system and some of the local controls—like transformer tap control, switching shunt

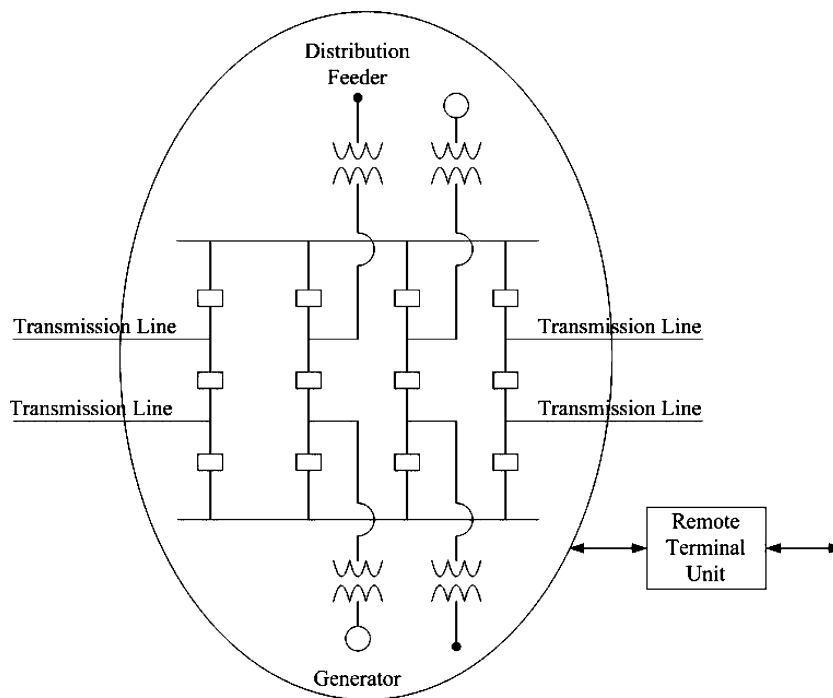


Fig. 2. Substation automation system.




Fig. 3. SPS/RAS communication links.

capacitors/reactors, etc.—are managed by it. These newer substation automation systems, shown in Fig. 2, gather data at much faster rates (milliseconds) and can record very fast sequence of events, e.g., the openings and reclosures of breakers during a fault.

It should be noted here that although data is often collected at these very fast rates at substations the present communication architecture between the RTUs and control center cannot support the gathering of all this fast data at the control center. Thus, the control center communication system cannot support fast wide-area controls unless the architecture is upgraded to handle faster communications.

2) *SPS/RAS*: Each special protection scheme today uses dedicated fast communication lines and a dedicated computer (microprocessor), as shown in Fig. 3. Once installed, these are quite inflexible except for changing the settings like trip times. Also, as SPS/RASs proliferate—the western U.S. grid now has over a hundred such schemes—their coordination becomes very difficult and the burden of offline studies to coordinate the settings becomes onerous and error-prone.

#### D. Communications Architecture of the Future

Data is being gathered at high scan rates within substations and fast local controls using digital (microprocessor) methods are commonplace. Some limited but fast wide-area

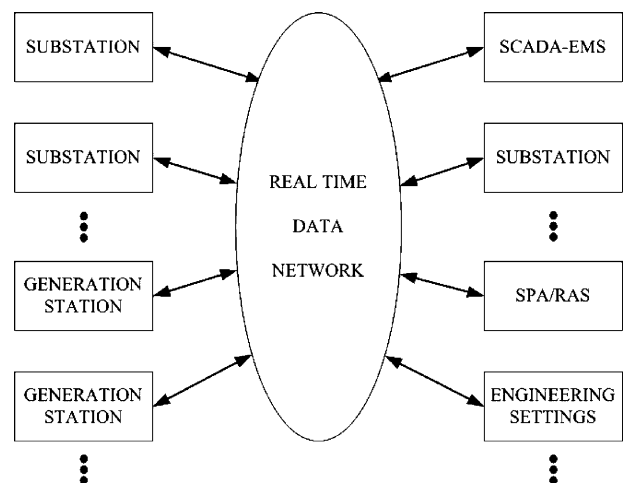


Fig. 4. Conceptual network for communication of all power grid data.

controls are already being used in the form of an SPS. However, they use dedicated communication lines rather than a network and are then inflexible as far as their control scheme is concerned.

The obvious evolution of this communication system is the high-speed network that will connect all the fast scanned data from all substations (Fig. 4). This data can then be made available to the control center to do the traditional control



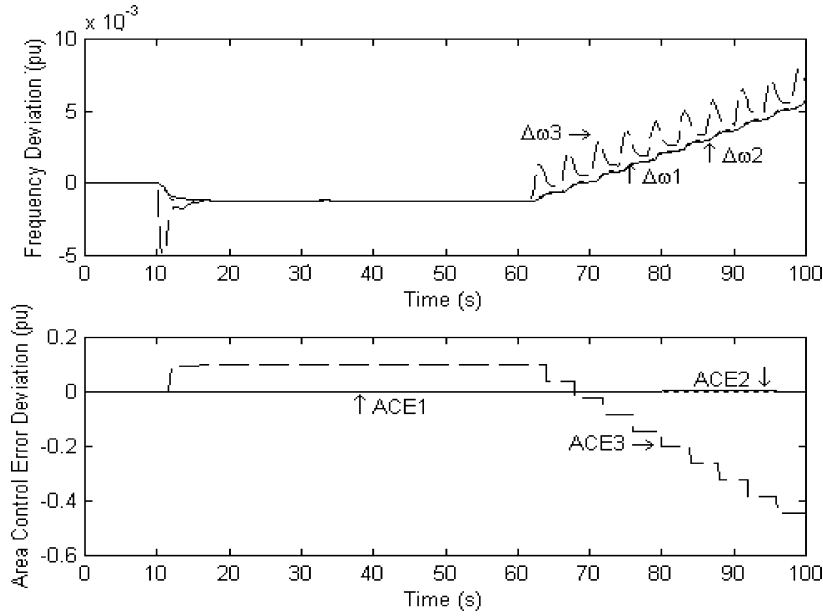


Fig. 6. Response with conventional AGC.

tain AGC cycle, say, due to a measurement failure, then one cannot calculate ACE. A practical industry approach is to pause AGC, and if this continues for too long, AGC is suspended, and alarms are issued. During the pause or suspension period, the set points remain unchanged.

To understand the impact of distributed communications on AGC, consider the problem of time delays in a closed-loop control. Assuming a linear *time-delay system* [10], let

$$\dot{x}(t) = Ax(t) + \sum_{i=1}^{n_d} A_{di}x(t - \tau_i) \quad (2)$$

where  $x(t)$  is the state at time  $t$ ,  $\tau_i$  is some fixed delay, and  $A$  and  $A_{di}$  represent the dynamics of the current and previous states, respectively. A one-area linearized AGC model could be expressed as

$$\dot{x}(t) = Ax(t) + Bu(t) + F\Delta P_d \quad (3)$$

$$y = Cx(t). \quad (4)$$

Here,  $u(t)$  is the control at the governor terminal, without this control, (3) is the same as the standard models in [8] and  $\Delta P_d$  is the load deviation. Also, note the structure of  $C$ ; the output simply tracks the frequency deviation and integral control deviation. Considering delay, the linearized model of the governor valve position would be

$$\Delta \dot{P}_v(t) = -\frac{\Delta f(t)}{RT_g} - \frac{\Delta P_v(t)}{T_g} - \frac{\Delta E(t - \tau)}{T_g}. \quad (5)$$

Notations are standard and can be found, for example, in [8]. The governor receives a delayed ACE signal, and the third term represents the impact of the former state. Thus, (3) becomes

$$\dot{x}(t) = Ax(t) + A_d x(t - \tau) + Bu(t) + F\Delta P_d \quad (6)$$

with  $F$  and  $A_d$  as defined in [11]. Notice the proposed framework is particularly appropriate for the deregulated environment where control services may be provided through bilateral contracts. With the previous states included in the controller

$$u(t) = Kx(t) + K_d x(t - \tau). \quad (7)$$

One approach to controlling time delay systems is the use of linear matrix inequalities (LMI) techniques. The stability criteria can be either delay independent or delay dependent (i.e., stability is only assured below a particular threshold) [12], [13]. In practice, delay-independent systems are not particularly meaningful, as alternative strategies will apply when communication fails or delays exceed some time threshold, e.g., use of the pause counter in AGC or simple voice communications. Still, they do provide some guideline for control design. Control gains can be chosen so as to satisfy a set of arithmetic Riccati equations [14], [15].

In the following, the control method is demonstrated on a three-area model with standard simplified models for the prime mover and governor are used. For brevity, the modeling details are left to [11]. Both fixed and random communication delays of up to 50 s are simulated with different delays in each control area.

Fig. 6 illustrates the response for a conventional AGC control. The result is a loss of power balance as the delayed signals are received. The frequency would continue to deviate if AGC was continued. This result is within expectations, since the set point is not updated with the correct and timely data. All units suffer from the delays. The point here is that if a critical channel is out of service for a certain time, it will degrade the conventional AGC system performance seriously. In this case, it is assumed that complete state information for all generators is available at each unit, and this is used in the LMI design. While it is not a practical design for such state information to be known globally, this case demonstrates an extreme example of the proposed control and communication




Fig. 7. Response using time-sensitive controls.

methodology. The system performance with the LMI control is shown in Fig. 7. The control works satisfactorily and when the slowest of the delayed signals begin to arrive at 60 s, only minor adjustments are needed.

The point in this example is not necessarily that this is the methodology one must follow to design the controller but that new approaches for bilateral and multilateral load following (balancing) services are feasible if a flexible and reliable communication network, as in Fig. 4, is available.

### B. Closed-Loop Voltage Control

There exists a rich history of research on automatic scheduling and coordination of voltage control devices [16]–[21]. In some European countries, including France, voltage control is organized as a hierarchical three-level regional structure, which was briefly discussed in Section II-B2. The automated secondary voltage control has been implemented in the French electric grid for more than a decade, and the approach is based on dividing the large power system into several voltage control areas. This approach generally is not suited to North American electric power grids, where the network tends to be tightly coupled. As the system loads increase, the interactions between power system components also become more significant. Therefore, it is difficult to divide a modern power system into decoupled voltage control areas. We have proposed a new approach for automatic voltage control in [22] and [23] that was motivated toward implementation in the transmission network operated by the Bonneville Power Administration (BPA) in the Pacific Northwest.

The controller coordinates the switching of discrete devices such as capacitor banks, reactor banks, and tap-changing transformers as well as continuous control of generator high side voltage settings. The objectives of the controller are: 1) to maintain the bus voltages across the network within specified voltage limits; 2) to minimize the number of switchings; 3) to increase the voltage

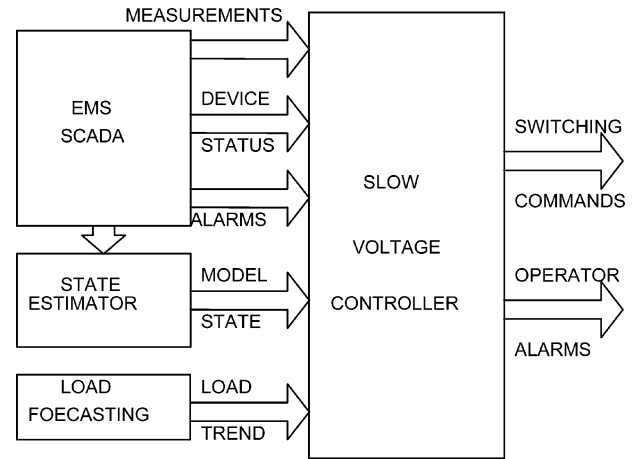


Fig. 8. Structure of the automatic voltage controller.

control reserves by keeping maximum number of devices offline; 4) to mitigate circular reactive power flows; and 5) to monitor and improve voltage security. At present, the commands for switching the voltage devices are issued by the system operators manually through the EMS supervisory control system. The proposed controller aims to automate the switching commands. By supervising and managing the reactive power resources continuously, it would be especially effective in preventing large scale events such as the 2 July 1996 western blackout and the 14 August 2003 blackout, which are related to reactive power deficiencies in some parts of the power network.

The overall structure of the controller is shown in Fig. 8. It uses the power-flow model of the system from the state estimator as the internal computational engine. The effects of the switching of voltage control devices are computed by solving the power-flow problem adaptively in local areas around each device. The switching decision is formulated as a nonlinear discrete optimization problem that can be solved very fast, and is aimed toward online implementation. Short-

term load forecasting estimates from AGC are used for computing “look-ahead” power-flow scenarios at, say, 10 and 30 min ahead of present time, and the switching decisions can thus be made predictive to minimize the control actions under load fluctuations.

The voltage controller has been developed in close collaboration with the BPA system operators and engineers. A prototype version of the controller, described below, is at present being implemented by National Systems Research Co. (NSR) in the BPA grid [24]. The prototype implementation has established the feasibility of the algorithm computations for online implementation in the large electric power system.

The controller as shown in Fig. 8 is essentially a feedback control that uses the state estimation model together with the current values of bus voltages from SCADA to compute the recommended control actions. In the closed-loop implementation, it is aimed toward operating say once every minute or so, and hence is denoted as a “slow voltage controller.” The current prototype of the controller at BPA is an open-loop implementation that only recommends the control switching actions to an operator who makes the actual decision on the switching commands through SCADA. It is intended for verifying the validity and reliability of the automatic controller.

In the next few paragraphs we describe the structure of the optimization problem resulting in (8) and the flow of the voltage controller (see Fig. 10 below).

Let us consider a single power-flow scenario first. In the following discussion, the index  $i$  denotes the index for the different control variables. We will associate a cost with each control device as a combination of three terms: 1) switching costs  $C_i$  to reflect on the relative preference of the device and when it was last switched; 2) a voltage penalty term to quantify the remaining voltage violations after the device has been switched; and 3) other control objectives such as a penalty for circular VAR flow to mitigate the effects of loop reactive flows.

The switching costs  $C_i$  are set in such a way that it would increase significantly after each switching and it would decrease slowly afterwards during each control iteration. This is meant to prevent repeated switching of the same device that would be undesirable. Also, control preferences can be handled by adjusting the relative switching costs of  $C_i$  among the different types of devices. For instance, at BPA, the switching of capacitor/reactor banks is preferred over the tap changing of transformers, since transformers are more expensive to maintain. Accordingly, the transformer tap changes have a higher  $C_i$  compared with capacitor/reactor banks. Moreover, it is better to switch out devices whenever possible rather than switching in devices to correct a problem. For instance, it is better to switch out reactor banks first to address low voltages before switching in capacitor banks. Therefore, a switching-in action has a higher penalty  $C_i$  compared with a switching-out action.

Next, we will discuss the computation of the voltage penalty term. The switching effects of each control action are calculated from the large-scale state estimation model by local power flow computations around each control device.




Fig. 9. Voltage penalty function.

It has been long known in power system computations [25], [26] that switching effects of line outages and reactive power devices have a limited effect in a geographical sense. The localized echelon approach in [20] is also based on this assumption. We use a similar approach in the controller by: 1) constructing a local area around each control device using the concept of electrical distance developed by EDF in [21] and 2) by carrying out full nonlinear power-flow solutions within the local areas of each device. We want to emphasize that the local areas are constructed around each device purely for computing the switching effects and they will be overlapping among many devices. The optimization itself is carried out for the large system. The local area approach enables the computation to be very fast even for large-scale models as demonstrated by the feasibility of the prototype controller at BPA.

The objective of the optimization would be to find a switching action that corrects all the voltage violation alarms. However, typically, not all of them can be corrected by a single switching action, and there would be no feasible solution for the optimization. Therefore, we define a voltage penalty function as shown in Fig. 9, and the voltage penalty term for any device is the summation of voltage penalties for each bus in its local area.

The third term in the cost function involves the circular VAR penalty. We have proposed fast computational procedures for evaluating the existence of any circular VAR loops in a large power system model in [22] and [23], and the algorithms have also been successfully implemented in the prototype controller at BPA. It is a graph theory based depth first search algorithm [22].

Summarizing the discussion so far (see Fig. 10), the optimization can be formulated as follows:

$$\begin{aligned} & \text{Min} \sum_{i=1}^n |k_i| C_i \\ & + \sum_{m=1}^M p_m [F_m(k_1, \dots, k_n) + \dots + \lambda g_{\text{cir},m}(k_1, \dots, k_n)] \\ & \text{s.t.} \sum_{i=1}^n |k_i| \leq N_{\text{sw}} \quad k_i \in \{-1, 0, 1\}. \end{aligned} \quad (8)$$

Here, the optimization control variables are  $k_i$ , which represent a switching out ( $k_i = -1$ ), switching in ( $k_i = 1$ ), or no change ( $k_i = 0$ ), respectively. For tap-changing transformers, these actions correspond to switching up ( $k_i = -1$ ), switching down ( $k_i = 1$ ), or no change ( $k_i = 0$ ), respectively.  $N_{\text{sw}}$  denotes the maximum number of simultaneous switching actions permissible per control iteration, and it is



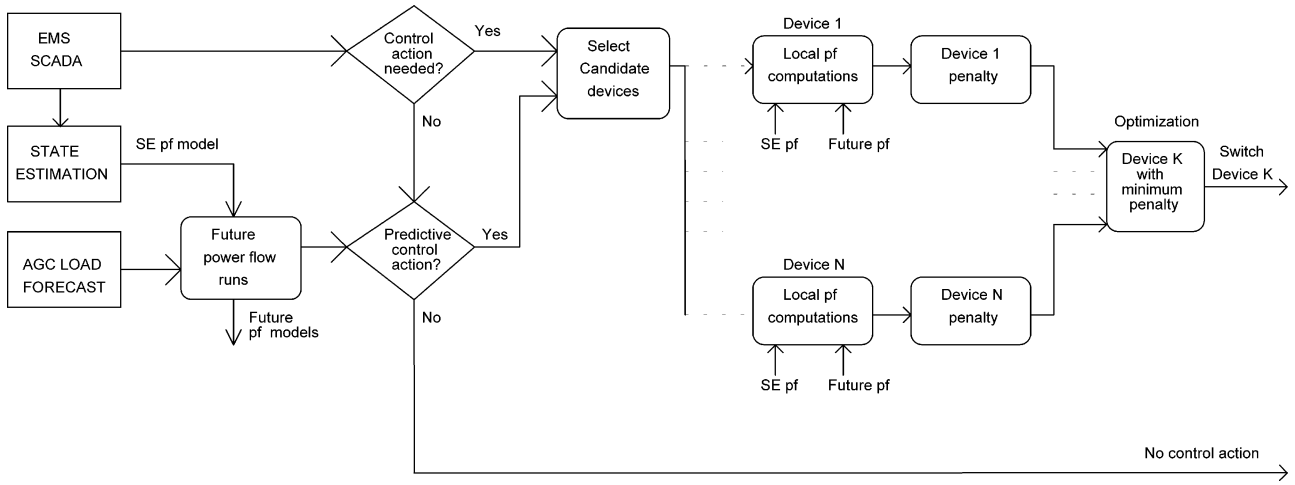


Fig. 10. Flow diagram for the automatic voltage controller.

assumed to be one for fast implementation in the BPA prototype. The variable  $C_i$  represents the switching costs of device  $i$  as discussed above.

As formulated here, the optimization is carried out in a robust sense over  $M$  different power-flow scenarios from the load forecasting computations, which are assigned relative probabilities  $p_m$ . The term  $F_m$  is for the voltage penalty calculated after the switching actions  $k_1$  through  $k_n$  for the  $m$ th power flow scenario. The final term  $g_{\text{cir},m}$  denotes the circular VAR penalty that is calculated after the switching actions  $k_1$  through  $k_n$  for the  $m$ th power flow case.

Generators can also be included by formulating the generator high side bus voltage set points in the form of certain discrete settings. The details can be seen in [22]. The controller as proposed was tested extensively on the western Oregon subsystem in the Pacific Northwest in [22]. In the prototype implementation in progress at BPA at present, the controller has been implemented on the entire BPA transmission grid [24], which establishes the feasibility of the controller for the large-scale power system.

### C. Small-Signal Stabilization Controller

When all the eigenvalues of the linearized system Jacobian evaluated at the current system equilibrium point have negative real parts, the system is said to be small-signal stable [8]. That is, all the modes of the linearized system must have positive damping for small-signal stability. Under stressed power system operating conditions, if any of the eigenvalues cross over into the right half complex plane, a small-signal instability results, typically leading to diverging or sustained oscillations.

A classical example of the small-signal instability was the 10 August 1996 western blackout [27] when the 0.25-Hz WSCC interarea mode became unstable, leading to diverging oscillations and the large scale blackout. Such oscillatory instabilities typically develop over several minutes before resulting in the eventual serious consequences such as the system separation. Therefore, with the existing technology of wide-area measurements and the emerging technology of wide-area control architectures, it is possible to develop

small-signal stabilization controllers, which can detect the evolution of the small-signal instability and take counteractions to stabilize the unstable modes.

We are working on the design of such a wide-area small-signal controller at Washington State University (WSU), Pullman, and the proposed configuration is shown in Fig. 11. The recent doctoral dissertation [28] at WSU provides a full design for the small-signal stability real-time controller. We discuss the salient features of the controller in this section. The controller as shown includes two sets of components: 1) a central controller and 2) fast control devices such as static VAR compensators (SVCs) and thyristor controlled series compensation (TCSC) devices.

The main difference in our approach from previous research is that our controller is exclusively designed as a safety-net-type real-time control. The proposed controller will initiate specific control actions only when it detects the emergence of poorly damped or negatively damped oscillatory modes in the power system being monitored. The controller is designed for real-time implementation in the large power systems by utilizing the wide-area measurements together with the heuristic control rules developed from offline studies. When the oscillatory instability is detected, the central controller switches the operation of a specific SVC from the normal voltage regulation mode into an aggressive power system damping control (PSDC) mode for damping out the oscillations.

The dissertation [28] proposes the following simple rule for the SVC PSDC design. The SVC is chosen so that it lies on the sending or the receiving end of a critical intertie line that is participating in the oscillatory mode. Then, the active MW power flow on the critical intertie line is used as the control input for the SVC, and the PSDC is designed as a phase lag design when the SVC is at the sending end, while it is designed as a phase lead design when it is at the receiving end of the active power flow on the tie line. This work shows the effectiveness of this rule above by analyzing the nature of the interarea mode in the Kundur two-area system, and the rule is also verified in a WECC large-scale example, namely, a validated model of the 10 August 1996 western blackout from [27].




Fig. 11. Real-time small-signal stability centralized control strategy.

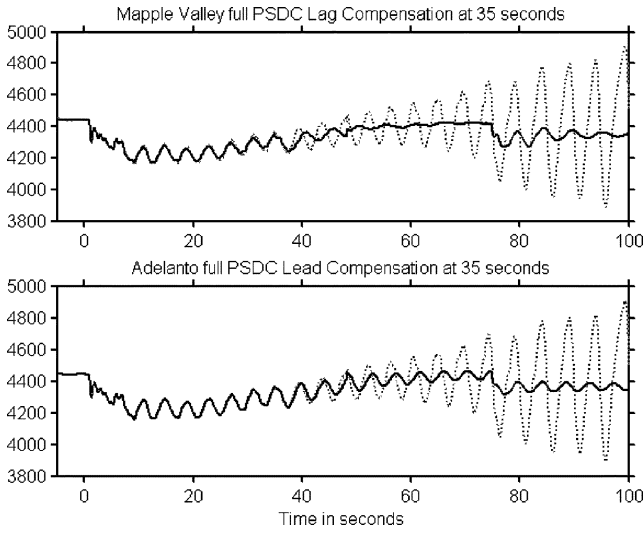
Specifically, the simulations in [28] show that the oscillatory WECC instability of 10 August 1996 could possibly have been prevented with the implementation of the proposed wide-area control. The rules for reliable oscillation detection were tested in [28] using the actual wide-area measurement data that was recorded during the 10 August 1996 disturbance. Based on this data, if the threshold for controller action had been set at say  $-1\%$  damping for any oscillatory mode, then the controller would have reliably concluded the emergence of the instability for the 0.25-Hz WECC inter-area mode at about 34 s after the initiating disturbance of Ross–Lexington line outage.

The active tie-line flow on the California–Oregon ac inertia line (COI) is then used as the control input. Two SVCs which were in operation during the 1996 event—namely, Maple Valley, located near Seattle, WA, and Adelanto, located near Los Angeles, CA—were then studied as the potential SVC controls for damping the oscillations. As per the rule stated above, Maple Valley PSDC is designed as a phase lag control with the COI MW flow as the PSDC input signal, since Maple Valley is located on the sending end of the COI MW flow. On the other hand, Adelanto is located on the receiving end on the receiving end of the COI MW flow, and hence Adelanto PSDC is designed as a phase lag control with the COI MW flow as the PSDC input. Simulations in [28] show that the oscillatory instability of 10 August 1996 could have been prevented by triggering *either* of the two SVC devices into the PSDC operation as defined above. The two SVCs were operating in the normal voltage regula-

tion mode during the 10 August 1996 event, and they experienced the diverging oscillations like the rest of the WECC system. In the simulations, the SVCs are assumed to change to the PSDC mode at 35 s after the initial Ross–Lexington outage, which is a full 1 s after the central controller reliably detects the emergence of the oscillatory instability of the 0.25-Hz mode at 34 s using the  $-1\%$  threshold for the 0.25-Hz mode.

Fig. 12 shows the comparison of the COI MW flows in the validated model with and without either of the two SVC PSDC controls. That is, the system experiences an oscillatory instability with diverging oscillations when the SVC is in the voltage regulation as was the case during the 10 August 1996 event. However, the oscillations remain well-damped throughout the disturbance in the simulation, when the SVC is switched into the PSDC control mode at 35 s after the Ross–Lexington outage. The switching sequence of line trippings and generation outages are identical in the two simulations, and they correspond to the actual sequence of events during the 10 August 1996 event. The simulations illustrate the potential usefulness of the proposed control methodology.

It is recommended that every power system should have a few SVCs marked for operating in the PSDC mode under emergency conditions whenever they are required to do so by the small-signal stability controller. It is important to develop other safety-net-type wide-area controls for targeting phenomena such as transient instability and fast voltage collapse phenomena.



**Fig. 12.** Simulated COI behavior for an SVC triggered at 35 s to full PSDC at Maple Valley with COI phase lag compensation at Adelanto with COI phase lead compensation.

#### D. Other Stability Controls

The first two examples for frequency and voltage control are relatively slow in that control actions are required in the range of several seconds. In the case of the third example on small-signal stability, a few seconds are available for control intervention. For such oscillatory stability we are also developing control methods that can:

- adjust parameters on PSS based on local and wider area measurements [29];
- switch capacitors, shed load, or trip generation to improve system damping.

The availability of fast communication and computations also allows more sophisticated control schemes for power systems, especially for systems that are transient stability or voltage stability limited. The fast control needed to stabilize power systems today is limited to local controls (protection) or the hard-wired RAS/SPS. Although the fast calculation of needed controls to stabilize a transient instability case within the milliseconds required for intervention remains a difficult task, a coordinated design of stability control can remove the limitations of hard-wiring these SPS/RAS and the need for large offline studies to fix their parameters. We are developing schemes that can in real time:

- quickly determine the amount and location for generator rejection and load tripping to stabilize the system when the dynamic security analysis detects instability for a particular contingency;
- determine which available RAS/SPS scheme ought to be activated to stabilize such an instability.

A more exciting possibility brought about by the availability of fast data acquisition by a wide-area communication network is the elimination of the hard-wired RAS/SPS schemes that are proliferating to take care of various operating conditions. If such a flexible communication network is available, a new RAS/SPS scheme can be implemented by software alone. Such a soft-wired RAS/SPS scheme will

make it possible to easily bring into play new controllers and phase out old controllers without having to do physical installations or removals. Finally, the objective for all of these schemes is to allow the operation closer to stability limits without sacrificing reliability.

#### IV. PROPOSED COMMUNICATION INFRASTRUCTURE

We have designed and built a prototype of a new communications framework called GridStat for delivering *status information* and data representing dynamic operational phenomena, such as voltage, current, breaker status, phasor measurement data, and command decisions to be delivered [30]. It provides the communication services depicted in Fig. 4. GridStat allows *status variables*, named program language variables with status information, to be programmed at a high level that is easy for programmers not expert in network or distributed computing to use and reuse. It then delivers the status variables with the required timeliness, redundancy, and security. GridStat is implemented as *middleware*, a high-level software functioning above the operating system that makes distributed computing systems much easier to program by providing higher level abstractions for programmers to use. Middleware also allows for interoperability across different operating systems, network technologies, programming languages, computer architectures, and even across different vendors' middleware frameworks. GridStat is a specialization of a middleware architecture called publish-subscribe, with a management plane and with optimizations to take advantage of the semantics of status variables.

##### A. GridStat Overview

Fig. 13 gives a high-level overview of GridStat. *Subscribers* are programs that receive status variables. *Publishers* are programs or intelligent devices that create status variables, and then call a GridStat application programming interface (API) to create the status variable and then each time an update needs to be delivered. Publishers do not need to be aware of particular subscribers, and *vice versa*. Still, publishers and subscribers can be aware of each other indirectly if they require. Ongoing research is developing a trust management framework, which allows constraints on what entities may receive or forward a particular status variable. The cloud in Fig. 13 depicts an opaque network that the status variables are routed through.

Status variables need to be delivered with specified QoS requirements, including timeliness, redundant paths, and computer security. To facilitate this, above the publisher-subscriber data plane is a *QoS Management* plane. This takes QoS parameters both when a publisher creates a status variable and when a subscriber subscribes to it. It then exerts control over the network to ensure that status variables are delivered with the requested QoS.

GridStat is designed to be easy for subscribers to program. For example, subscribers are provided with a local object that caches the latest value of the status variable, so programmers do not have to deal with callbacks or interrupts (though they

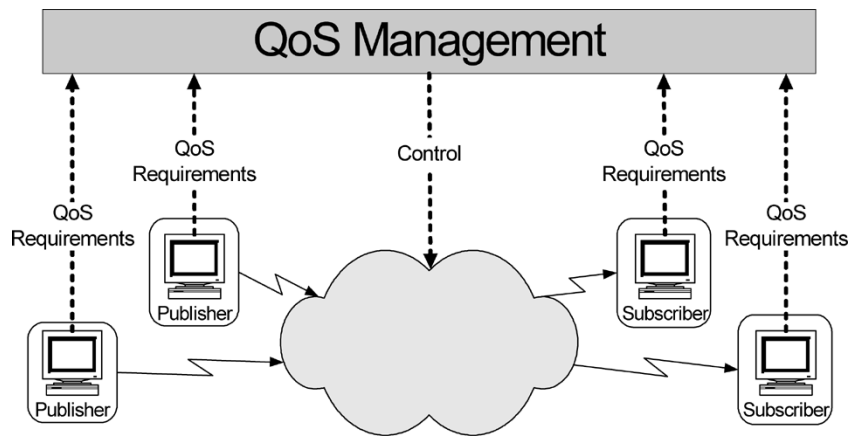


Fig. 13. Overview of GridStat.

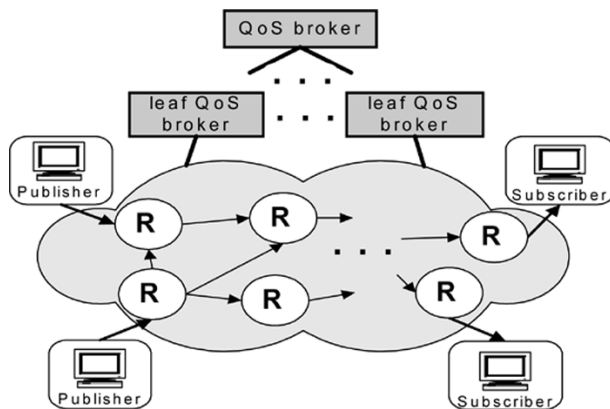


Fig. 14. Status routers and QoS broker hierarchy.

have the option of getting a callback if their timeliness or redundancy requirements are violated). GridStat also provides its interfaces using Common Object Request Broker Architecture (CORBA), a widely utilized middleware standard.

### B. Detailed Architecture

Fig. 14 provides more detail of GridStat's architecture. For GridStat, the network consists of a series of *status routers*, computer nodes that forward status events, much like IP routers in today's computer networks. Status routers filter status updates based on the rate and priority of subscribers that are "downstream." Status routers today are in software, just like initial IP routers were, but work is underway implementing them on very fast network processors. It is unnecessary for the QoS management plane to be put in hardware, because it is not on the critical path for fast delivery of status data. It is also undesirable to put it into hardware, since its functionality would very likely be augmented on a regular basis. The data path is much simpler and would likely change more slowly, hence its suitability to hardware implementation once more fully researched and defined.

GridStat's QoS management plane is not a single, centralized module. Rather, it is a hierarchy of modules called *QoS brokers*. A QoS broker contains and enforces policies on the entities that should be allowed to subscribe to status

variables within its domain of control. QoS brokers are arranged in a hierarchy: the lowest one, called a *leaf broker*, will typically be a substation; say, a particular utility's control center might be at the third or fourth level; and the ISO or RTO would be at the top. A nonleaf broker must approve subscriptions between its children's domains. GridStat has domains and subdomains, for reasons similar to the Internet's domain name system, to allow for local autonomy, but in GridStat's case via local policies. Additionally, this supports a divide-and-conquer approach to the routing task.

### C. Routing in GridStat

Fig. 15 details the routing process within GridStat's architecture. Most status routers connect only to other status routers in the same domain. An *edge status router* has additional functionality to allow programs to publish and subscribe to status variables. A *border status router* interfaces with a status router outside its own domain, namely, another border status router. The QoS broker that is the lowest common ancestor up the broker hierarchy controls its allocation.

When a subscriber subscribes to a status variable, a route is chosen from the publisher to the subscriber through status routers, using information on node and link delays. This route does not change during normal operation, unlike IP routing, for performance reasons. If a subscriber requests redundant paths, then one or more disjoint paths are also selected. Status variables are then delivered in parallel over these disjoint routes. This compensates for permanent and transient network and QoS router failures and also can provide for lower delivery latency (i.e., for the first copy of the status update to arrive) in the face of overload or cyberattack.

### D. Implementation Status

GridStat has been under development for almost three years. An initial prototype was demonstrated at NIST in 2002, and a pilot demonstration project is underway at Avista, an electricity and gas utility in five western U.S. states, to evaluate GridStat's capabilities. The current GridStat prototype provides mechanisms to create and deliver status information, filter it at status routers based on the

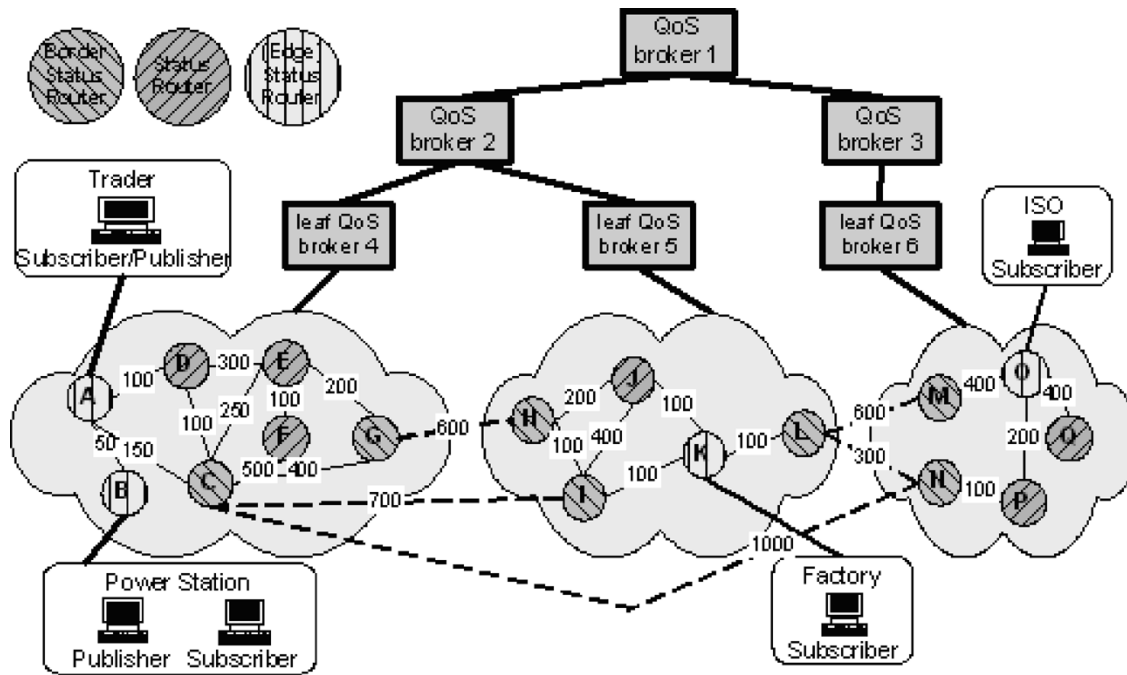


Fig. 15. QoS routing.

rates of “downstream” subscribers, and many optimizations. It includes a QoS management hierarchy where policies on allocation and security can be programmed. GridStat also includes a visualization subsystem which projects onto maps and stripcharts dynamic information about status variables, status routers, link delays, etc.

We now provide an overview of GridStat performance. The latency added by one status router’s forwarding operations is less than 1 ms in a nonoptimized Java version, and we believe that multiple such operations will be achievable in a more optimized version in C++. An implementation is underway on Intel IXP 2400 network processors. A design study and comparison with other reference implementations of forwarding protocols indicates that the latency should be significantly less than one microsecond. Allocation of paths for status delivery from a publisher to subscriber is significantly more expensive. However, it is not on the critical path of data delivery, and is almost always done offline.<sup>1</sup>

#### E. Summary of Related Computer Science Work

A closely related communications infrastructure, namely a status dissemination framework, is PASS [31]. PASS provides information about the status of communication network resources, namely, their availability (i.e., 1-b status variables). It is intended for military operations, where it has been fielded repeatedly, but optimizes heavily based on 1-b variables because bandwidth is typically very scarce in its intended environment. PASS has low-level status delivery mechanisms, but only for 1-b data, not floating point or integer data. It also does not have a management infrastructure or any adaptation, and all subscribers must

<sup>1</sup>More information on GridStat, its performance, and its status can be found at <http://gridstat.net>

have the same QoS parameters such as rate or maximum latency. It also does not support redundant paths or other capabilities necessary for a large critical infrastructure such as an electric power grid.

Sienna is an Internet-scale publish–subscribe event notification framework [32]. It has flexibility in the way publisher–subscriber graphs are organized, but does not provide any QoS properties or management infrastructure. It also only delivers generic events, not status variables, and thus cannot perform any of the optimizations and management that GridStat does on status variables.

#### F. Summary of Related Power Engineering Work

A number of distantly related efforts have come from the electric power research community and industry, but none are a usable framework with the scale and management needed for the electric power grid and other critical infrastructures. The work in [3] points out the major deficiencies in the current communication and information systems and proposes a new information architecture tailored to the electric power grid, including different kinds of redundancy. GridStat provides most or all of the mechanisms and management systems that are seen as required by their analysis.

The Electric Power Research Institute (EPRI) has also started the Infrastructure Security Initiative (ISI), which focuses in part on how to develop a secure private communication network for the power industry as an alternative to Internet-based systems [33]. It is not a deployable framework that can provide QoS guarantees, such as GridStat, or even an architecture for such a framework. NERCnet is a new communication infrastructure being developed by NERC. There is relatively little published about its details, but it seems to be a network infrastructure with no communication services such as GridStat to make it programmable and QoS

controls or management infrastructure (at least not at the middleware or application layer).

EPRI has created the Utility Communications Architecture (UCA), version 2 [34]. It offers interconnectivity between devices and interoperability between databases. Still, UCA is only designed to operate within a single substation, does not support QoS guarantees, and has no management infrastructure. As such, it is unsuitable to be deployed beyond the substation level. Further, it is a standard, and not a technology. Still, the creation of a translation layer from UCA to GridStat could combine the best of both infrastructures and be of practical use. Indeed, the UCA standard [34] notes that one of its benefits is the ability to incorporate future communications innovations and discusses possible future expansion of the UCA protocol suite to other application protocols such as CORBA. This indicates that hybrid configurations with UCA being used within a substation and GridStat managing flows of packed UCA streams between substations and above them would likely being a feasible and pragmatic combination of the two technologies.

Other related projects such as IntelliGrid and OASIS do not provide real-time status delivery mechanisms. The WAMS and EIPP projects do deliver status data via dedicated links. While they are important building blocks, they do not provide end-to-end QoS management.

## V. CONCLUSION

While the potential for improved control afforded by the new communication systems is great, the basic principles of designing a large complex system with high reliability requirements, such as the power system, will force a conservative design. A system with great flexibility and functionality as proposed here may not lend itself to analysis and performance guarantees. The overriding concern in power systems engineering has been safety and reliability of the system. With or without deregulation, that is not expected to change. An extreme implementation of the proposed approach could easily create more problems than it solves. For example, the fundamental function of relays must remain protection though in this new framework it can take on a secondary role.

In the past, transient stability analysis and control methods have been based on differential-algebraic representation of the quasi-stationary power system dynamics. In the proposed real-time control framework, wherein the measurements and control signals travel over the communication network, the mathematical representation of the future power system must include the transmission and communication networks operating in parallel, serving the respective purposes of electrical power exchange and information exchange. A comprehensive model is needed to provide a platform for developing simulation, analysis, and design tools. For example in transient simulation studies, there will be a need for including a rudimentary model of the protection, a representation of the delays in receiving measurement signals, and a data flow model that shows the data available to each controller. The

development of appropriate simulation tools is part of our ongoing research.

The system we are proposing may also require system operators to have a very different feel for the system. Today, operators working in the control center rely on experience to sense when the system is becoming overstressed and can take appropriate remedial actions. With a more distributed structure, important control decisions could be made completely outside the view of the operators and at speeds far faster than the cognitive ability to track changes as already happens today with the SCADA alarms systems. Again, new tools for operator training will be needed.

The ultimate goal for the real-time control, communication, and computation schemes proposed here is to control the dynamics directly without having to set special protection parameters. Although the available communication today is fast enough, the computation needed for such real-time control is still very complex and poorly understood. This paper has discussed the various issues for real-time stability control and the advances needed to make this feasible given the appropriate communication infrastructure.

## ACKNOWLEDGMENT

The authors would like to thank H. Gill for her useful comments on this research, H. Gjermundrød, I. Dionysiou, R. Johnston, P. Jiang, and S. Sheshadi for their work on GridStat, and S. Bhowmik, Y. Chen, R. Kavasseri, E. Nobile, J. Quintero, K. Schneider, and X. Yu for their work on the power system controls.

## REFERENCES

- [1] "Final report on the August 14, 2003 blackout in the United States and Canada: Causes and recommendations," US-Canada Power System Outage Task Force, 2004.
- [2] A. Bose, "Powers system stability: New opportunities for control," in *Stability and Control of Dynamical Systems with Applications*. Boston, MA: Birkhauser, 2003.
- [3] Z. Xie, G. Manimaran, V. Vittal, A. G. Phadke, and V. Centero, "An information architecture for future power systems and its reliability analysis," *IEEE Trans. Power Syst.*, vol. 17, no. 3, pp. 857–863, Aug. 2002.
- [4] E. Nobile, A. Bose, and K. Tomsovic, "Feasibility of a bilateral market for load following," *IEEE Trans. Power Syst.*, vol. 16, no. 4, pp. 782–787, Nov. 2001.
- [5] X. Yu and K. Tomsovic, "Application of linear matrix inequalities for load frequency control with communication delays," *IEEE Trans. Power Syst.*, vol. 19, no. 3, pp. 1508–1515, Aug. 2004.
- [6] E. Nobile and A. Bose, "A new scheme for voltage control in a competitive ancillary service market," presented at the Power Systems Computation Conf., Seville, Spain, 2002.
- [7] S. Bhowmik, K. Tomsovic, and A. Bose, "Communication models for third party load frequency control," *IEEE Trans. Power Syst.*, vol. 19, no. 1, pp. 543–548, Feb. 2004.
- [8] P. Kundur, *Power System Stability and Control*. New York: McGraw-Hill, 1994.
- [9] M. Illic and S. Liu, *Hierarchical Power Systems Control*. New York: Springer-Verlag, 1996.
- [10] L. Dugard and E. I. Verriest, *Stability and Control of Time-Delay Systems*. New York: Springer-Verlag, 1997.
- [11] X. Yu, "Power systems control via open communication networks," Ph.D. dissertation, Washington State Univ., Pullman, 2002.
- [12] T. Mori, "Criteria for asymptotic stability of linear time-delay systems," *IEEE Trans. Autom. Control*, vol. AC-30, no. 2, pp. 158–160, 1985.
- [13] T. Mori and H. Kokame, "Stability of  $\dot{x}(t) = Ax(t) + Bx(t - \tau)$ ," *IEEE Trans. Autom. Control*, vol. 34, no. 4, pp. 460–462, Apr. 1989.

- [14] M. S. Mahmoud and M. Zribi, " $H_\infty$  Controllers for time-delay systems using linear matrix inequalities," *J. Optim. Theory Appl.*, vol. 100, no. 1, pp. 89–112, Jan. 1999.
- [15] M. Zribi, M. S. Mahmoud, M. Karkoub, and T. T. Lie, " $H_\infty$ -Controllers for linearised time-delay power systems," *IEE Proc. Gener., Transm., Distrib.*, vol. 147, no. 5, pp. 401–408, Sep. 2000.
- [16] J. P. Paul, J. T. Leost, and J. M. Tesson, "Survey of the secondary voltage control in France: Present realization and investigations," *IEEE Trans. Power Syst.*, vol. PWRS-2, no. 2, pp. 505–511, May 1987.
- [17] S. Corsi, P. Marannino, N. Losignore, G. Moreschini, and G. Piccini, "Coordination between the reactive power scheduling function and the hierarchical voltage control of the EHV ENEL system," *IEEE Trans. Power Syst.*, vol. 10, no. 2, pp. 686–694, May 1995.
- [18] H. Vu, P. Pruvot, C. Launay, and Y. Harmand, "An improved voltage control on large scale power system," *IEEE Trans. Power Syst.*, vol. 11, no. 3, pp. 1295–1303, Aug. 1996.
- [19] A. Zobian and M. Ilic, "A steady state voltage monitoring and control algorithm using localized least square minimization of load voltage deviations," *IEEE Trans. Power Syst.*, vol. 11, no. 2, pp. 929–938, May 1996.
- [20] C. Yu, Y. T. Yoon, M. Ilic, and A. Catelli, "On-line voltage regulation: The case of New England," *IEEE Trans. Power Syst.*, vol. 14, no. 4, pp. 1477–1484, Nov. 1999.
- [21] P. Lagonotte, J. C. Sabonnadiere, J. Y. Leost, and J. P. Paul (EDF), "Structural analysis of the electrical system: Application to the secondary voltage control in France," *IEEE Trans. Power Syst.*, vol. 4, no. 2, pp. 479–484, May 1989.
- [22] Y. Chen, "Development of automatic slow voltage control for large power systems," Ph.D. dissertation, Washington State Univ., Pullman, 2001.
- [23] Y. Chen and V. Venkatasubramanian, "Automatic on-line controller for coordinated slow voltage control," *IEEE Trans. Power Syst.*, submitted for publication.
- [24] "Final report for the BPA on-line demonstration of the automatic voltage control," Ciber Inc., Vancouver, WA, 2004.
- [25] J. Zaborzsky, K. Whang, and K. Prasad, "Fast contingency evaluation using concentric relaxation," *IEEE Trans. Power App. Syst.*, vol. PAS-99, no. 1, pp. 28–36, Jan./Feb. 1980.
- [26] V. Brandwajn and M. Lauby, "Complete bounding method for AC contingency screening," *IEEE Trans. Power Syst.*, vol. 4, no. 2, pp. 724–729, May 1989.
- [27] V. Venkatasubramanian and Y. Li, "Analysis of 1996 Western American electric black-outs," presented at the Bulk Power System Dynamics and Control VI Conf., Cortina, Italy, Aug. 2004.
- [28] J. Quintero, "A real-time wide-area control for mitigating small-signal instability in large electric power systems," Ph.D. dissertation, School Elect. Eng. Comput. Sci., Washington State Univ., Pullman, 2005.
- [29] R. Kavasseri, "New methodology for design of power system stabilizers," Ph.D. dissertation, School Elect. Eng. Comput. Sci., Washington State Univ., Pullman, 2002.
- [30] C. Hauser, D. Bakken, and A. Bose, "A failure to communicate: Next-generation communication requirements, technologies, and architecture for the electric power grid," *IEEE Power Energy*, vol. 3, no. 2, pp. 47–55, Mar./Apr. 2005.
- [31] J. Zinky, L. O'Brien, D. Bakken, V. Krishnaswamy, and M. Ahamad, "PASS: A service for efficient large scale dissemination of time varying data using CORBA," in *Proc. Int. Conf. Distributed Computing Systems*, 1999, pp. 496–506.
- [32] A. Carzaniga, D. S. Rosenblum, and A. L. Wolf, "Design and evaluation of a wide-area event notification service," *ACM Trans. Comput. Syst.*, vol. 19, no. 3, pp. 332–383, Aug. 2001.
- [33] "EPRI-ISI, infrastructure security issues, project status report," Elect. Power Res. Inst. (EPRI), 2003.
- [34] "Utility communications architecture (UCA) version 2.0," IEEE-SA, Tech. Rep. 1550, 1999.

**Kevin Tomsovic** (Senior Member, IEEE) received the B.S. degree in electrical engineering from Michigan Technological University, Houghton, in 1982, and the M.S. and Ph.D. degrees in electrical engineering from the University of Washington, Seattle, in 1984 and 1987, respectively.

He held the Advanced Technology for Electrical Energy Chair at Kumamoto University in Japan from 1999 to 2000. He is currently a Professor at Washington State University, Pullman. Visiting university positions have included National Cheng Kung University, National Sun Yat-Sen University, and the Royal Institute of Technology, Stockholm. He is also currently a Program Director in the Electrical and Communications Systems division in the Engineering Directorate at the National Science Foundation.

**David E. Bakken** (Member, IEEE) received the B.S. degrees in mathematics and computer science from Washington State University, Pullman, in 1985, and the M.S. and Ph.D. degree in computer science from the University of Arizona, Tucson, in 1990 and 1994, respectively.

He has worked at BBN Technologies (1994–1999) and Boeing (1985–1988). He is currently an Associate Professor at Washington State University.

**Vaithianathan Venkatasubramanian** (Member, IEEE) received the B.E. (Honors) degree in electrical and electronics engineering from Birla Institute of Technology and Science, Pilani, India in 1986, and the M.S. and D.Sc. degrees in systems science and mathematics from Washington University, St. Louis, MO, in 1989 and 1992, respectively.

He is currently a Professor in the School of Electrical Engineering and Computer Science, Washington State University, Pullman.

**Anjan Bose** (Fellow, IEEE) received the B.Tech. (Honors) from the Indian Institute of Technology, Kharagpur, in 1967, the M.S. degree from the University of California, Berkeley, in 1968, and the Ph.D. degree from Iowa State University, Ames, in 1974.

He has worked for the Consolidated Edison Co. (1968–1970), the IBM Scientific Center (1974–1975), Clarkson University (1975–1976), Control Data Corporation (1976–1981), and Arizona State University (1981–1993). He is currently the Distinguished Professor in Power Engineering and Dean of the College of Engineering and Architecture, Washington State University, Pullman.