

Obstacles to Self-Healing Reliable Complex Control Systems

Dr. Paul Oman, Mr. Jeff Roberts, and Dr. Edmund Schweitzer
Schweitzer Engineering Laboratories, Inc.
Pullman, WA 99163-5603
(contact: paulom@selinc.com)

Submitted to the
Fourth Information Survivability Workshop
October 15–17, 2001
Vancouver, B.C., Canada

I. Real-Time Complex Control Systems

The North American electric power grid is a highly complex real-time system exposed to natural disasters, sabotage and nuisance attack, in addition to the rapid system dynamics and demand swings inherent in providing electric power across large spatial areas. Setting aside disasters and attacks for the moment, the power grid can be modeled as a complex, real-time network of dynamic load and generation balances characterized by two types of stability. *Steady-state stability* describes the nominal balancing of relatively minor disturbances in load and generation fluctuations caused by normal start-up and shut-down events associated with the millions of appliances and equipment attached to the system. Constraints on operating parameters – nominally voltage levels, current magnitudes and power flows – are used to achieve steady-state equilibrium where the generation input is matched to system losses and electrical outputs. On the other hand, *transient stability* describes the power systems ability (or inability) to absorb major disturbances and return to a relatively balanced steady-state. Load shedding, generation shedding, and regional islanding are all means to dampen the wild system oscillations evident during stages of transient stability. These two levels of stability are evident in most complex real-time control systems.

From this description we see how relatively minor disturbances, failures, and physical phenomena (e.g., nominal weather) are absorbed through steady-state balancing events controlled by protective devices such as protective relays, breakers, shunt reactors, and Flexible AC Transmission System (FACTS) devices. Whereas, natural disasters, sabotage, and attack (both physical and cyber attacks) force the system into drastic transient stability corrective actions like power outages, emergency generation, and inter-tie disconnections. Fortunately, the problem of building a self-healing, ultra-reliable real-time power control system is not unlike the process of designing and implementing a survivable, reliable wide-area computer network. It is a bit more exacting because of the time constraints inherent in critical real-time systems (protective reactions take just a few milliseconds), but the principles and technology are similar.

This paper discusses obstacles to implementing self-healing, ultra-reliable real-time complex control systems for use in critical infrastructures such as electricity and other types of power distribution (e.g., natural gas), water and sewage systems, and transportation control networks. Although we focus on electric power systems, our analysis is relevant to all complex real-time control systems, especially those used for infrastructures and utilities.

II. The Need for Self-Healing Ultra-Reliable Control Systems

Several years ago Grudin and Roytelman [1] predicted that the North American power grid would be increasingly stressed by (1) larger-than-ever bulk transfers over wider areas, and (2) unanticipated operating conditions. Those predictions came true in Fall 2000 and were evident in the subsequent attempts to transfer power into electricity-starved California in Spring 2001. Ironically, just prior to the California energy crisis, papers by Jones [2] and Amin [3] described the fragility of the power system and, in essence, predicted the service disruptions. The authors called for self-healing, fault tolerant controls and improved quality of service for infrastructure control networks.

Our digital society is demanding more reliable electric power and, in ultra-precision industries like micro-circuitry, a higher quality of power. Despite the recent California instability, the North American electric power system is 99.9% reliable, with customers experiencing an average of only 8 hours of downtime per year. However, it has been argued that to serve the increasing needs of research and technology manufacturing, reliability must improve from 99.9% to 99.99999999%, representing less than 32 seconds of downtime per year [3]. While this may be unnecessary for all customers, it points out the need for improved qualities of service from electric power providers.

Postmortem analyses of real-time control system failures show time periods in which self-healing or self-adjusting activities could occur in order to maintain stability. For example, an analysis of the 1996 West Coast cascading blackout shows several opportunities for corrective load switching to maintain steady-state stability. Further, an automated warning system operating within a 5-6 minute window could have initiated load shedding and/or given sufficient notification for operators to bring auxiliary generation on line and thereby initiate transient stability corrections without the need for widespread outages [4]. These actions would have prevented the cascading outages and subsequent islanding that affected 7.5M customers throughout the West, and cost an estimated \$1.5B dollars in damages and lost service revenues [2,3,4]. However, there are numerous barriers, both technical and socio-political, that stand in the way of implementing such a self-healing wide-area control system:

- Absence of a high-speed wide-area communications infrastructure
- Fragility of the Internet and other telecommunications infrastructures
- Lack of network quality of service guarantees and service agreements
- Immaturity, fragility, and lack of interoperability in network trust frameworks
- Lack of requirements for reporting below-threshold disturbance anomalies
- A myriad of utility protocols with minimal interoperability
- Socio-economic and political resistance to regulatory controls

III. Obstacles to Self-Healing Ultra-Reliable Real-Time Control Systems

Optimistically, we would assume that the same technologies for mitigating risk and implementing survivability in computer networks could be used for control and protection in electric power systems. Unfortunately, this is only partially true. The reliability demands and time-critical nature of real-time control systems place additional burdens on quality of service guarantees, high-speed authentication, and trusted communications. We now elaborate on these obstacles.

Absence of a High-Speed Wide-Area Communications Infrastructure: As a result of the 1996 West Coast blackouts, the Western Systems Coordinating Council adopted an operating policy prohibiting operators from putting the electric power system into unanalyzed conditions for which there exist no contingency procedures [4]. But in the absence of a wide-area communication infrastructure, control decisions are localized and can only be shared in an ad-hoc manner. Grudin and Roytelman point out that a wide-area centralized control system could have diagnosed the 1996 power disturbance and may have reduced the magnitude of the outage. They describe a high-speed protection-level communications infrastructure and response system similar to that discussed by other researchers [1,5,6].

Fragility of the Internet and Other Telecommunications Infrastructures: As an alternative to a separate protection-level communications structure, several utilities have experimented with the Internet for access to real-time control and metering data. Experience shows, however, that while the Internet is sufficient for observation and maintenance planning, it is unsuitable for real-time protection because of its non-deterministic delivery and dynamic routing. Further, Internet E-commerce applications have experienced hacking, sniffing, spoofing, and deliberate overloading (e.g., denial of service attacks), which exposed the weaknesses of the Internet for critical applications [7,8]. Other telecommunications systems include the Public Switched Telephone Networks (PSTN), and leased lines forming Asynchronous Transfer Mode (ATM) networks, Frame-Relay Permanent Virtual Circuits (PVCs), and Frame-Relay Switched Virtual Circuits (SVCs). The ATM and PVC solutions have reliability and quality of service suitable for critical applications, but PSTN and SVC solutions have reliability and quality of service

concerns, respectively, that create questions about their use in real-time applications. Thus, ATMs and PVCs are the most reliable communications pathways outside of wholly-owned dedicated lines, but leased lines are still vulnerable to electronic intrusions.

Lack of Network Quality of Service and Service Agreements: The original IEEE 802.x Ethernet specifications called for multicast token-ring or token-bus data distribution that were unsecure and inefficient for time-critical applications. It was not until the IEEE 802.10 specification that point-to-point addressing and security considerations were implemented in Ethernet communications. As a result, many utility networks were implemented with little or no communication service guarantee. There are two mechanisms for ensuring quality of service guarantees over a network: (1) Leased resources sufficient to handle the maximum load, and (2) Packet prioritization that ensures high-priority packets are delivered at near minimum times. Implementing packet prioritization on proprietary networks has been done for many years, but on Ethernet networks this is still a research topic. Several companies have implemented Ethernet TCP packets over leased ATM or PVC lines. Fortunately, this provides quality of service guarantees suitable for time-critical applications. Unfortunately, the end-to-end TCP flow-control necessary for quality of service implementation can interfere with ATM and Frame-Relay packet construction, causing indeterminate degradation in service quality [9]. Further work is needed to better define quality of service mechanisms within ATM and Frame-Relay packets.

Immaturity, Fragility, and Lack of Interoperability in Network Trust Frameworks: Trusted communication between sender and receiver is crucial when sending control information across spatial, economic, and governing boundaries. There are three frameworks for establishing trusted, secure interconnections across public lines: (1) IP Security (IPSec), (2) Public Key Infrastructure (PKI), and (3) an informal “web of trust.” IPSec is an effort of the Internet Engineering Task Force to add security to the Ethernet TCP/IP layers. The disadvantage to this approach is the additional burden for the authentication and encryption/decryption of the data packet. The advantages are that it allows both protected and unprotected communications, and that COTS hardware/software solutions are available (but not always interoperable). PKI is an attempt to create a world-wide infrastructure for secure communications based on asymmetric public-key cryptography. Users have both private and public crypto-keys; their public keys are distributed via a trusted-third party Certificate Authority. The system works but is cumbersome and fragile – and therefore suspect for critical applications [10]. As an alternative to PKI, the Pretty Good Privacy group has implemented an informal “web of trust” where trusted users vouch for and include others in formalized lists of who to trust [11]. The system works for non-critical applications, but it is doubtful if it would suffice for critical control and communications. Other trust frameworks are being explored, but by-and-large these efforts are focused on E-commerce and are not sufficiently robust for critical infrastructure control and protection [11].

Lack of Requirements for Reporting Below-Threshold Disturbance Anomalies: There are no reporting regulations for wireless communications and mobile network failures; in electric power systems FERC requires reporting outages affecting 50,000 customers for 3 hours or more; in telephone communications the FCC requires reporting failures affecting 30,000 customers for over 30 minutes. In none of those three domains are there any reporting requirements for “near-critical” conditions and “near-failure” events. Such data would clearly be useful in modeling and predicting power system disturbances, and modern digital metering equipment has the capability to record and report these conditions. However, there exists no infrastructure for collecting, analyzing, and processing data from near-critical conditions, even though those data would be quite useful for maintaining steady-state stability.

A Myriad of Utility Protocols With Minimal Interoperability: In a typical control station or substation you find a plethora of communications protocols, including obscure proprietary protocols, EIA-232, EIA-485, Ethernet, Utility Communications Architecture, Distributed Network Protocol, Modicom’s Modbus and Modbus-Plus, Profibus, Foundation Fieldbus, and ControlNet. These protocols are used to connect the protection equipment such as breakers, reclosers, relays and meters, to control equipment like Remote Terminal Units, Data Processing Units, communications controllers, local PCs, and Supervisory Control and Data Acquisition devices. The diversity and lack of interoperability in these

communication protocols create obstacles when attempting to retrieve disturbance data (failure, near-failure, critical, or near-critical) from the station.

Socio-Economic and Political Resistance to Regulatory Controls: Approaches to implementing very wide-area (e.g., regional or national) control networks range from strongly centralized to highly decentralized systems with very fast calculations, high-speed communications, uniform or homogenized information structures, a robust communications infrastructure resilient to attack and natural phenomena, guaranteed quality of service levels and service agreements, and an integrated trust framework. While there have been several calls for improved control structures and increased regulatory requirements, it is doubtful at this time whether the socio-economic and political climate would support centralized control. Despite the failed California deregulation attempt, other deregulation efforts in the Southeast U.S. are ongoing and appear successful. It remains questionable, however, if critical infrastructure control systems should be stressed and jeopardized by deregulatory machinations.

IV. Summary and Conclusion

We are approaching an era in which complex systems can be modeled in real-time using high-speed communications and data from disturbances and near-critical conditions. These models can be used to prescribe corrective actions to maintain steady-state equilibrium, or at least dampen transient-state fluctuations. Technological and socio-political barriers to the successful implementation of self-healing, ultra-reliable real-time control systems are similar to those of implementing a survivable, reliable wide-area computer network, but more exacting because of the time constraints inherent in critical real-time systems. We can learn from advances in computer networking and we can glean technologies from that domain in order to improve the survivability and reliability of complex real-time control systems.

References

1. N. Grudin and I. Roytelman, "Heading Off Emergencies in Large Electric Grids," *IEEE Spectrum*, Vol. 34(4), April 1997, pp. 42-47.
2. A. Jones, "The Challenge of Building Survivable Information-Intensive Systems," *IEEE Computer*, Vol. 33(8), August 2000, pp. 39-43.
3. M. Amin, "Toward Self-Healing Infrastructure Systems," *IEEE Computer*, Vol. 33(8), August 2000, pp. 44-53.
4. J. Daume, "1996 Western Systems Coordinating Council Power System Disturbances," Paper #1, 24th *Annual Western Protective Relay Conference*, (Oct. 21-23, Spokane, WA), 1997.
5. K. Stahlkopf and M. Wilhelm, "Tighter Controls for Busier Systems," *IEEE Spectrum*, Vol. 34(4), April 1997, pp. 48-52.
6. K. Birman, "The Next-Generation Internet: Unsafe at Any Speed?," *IEEE Computer*, Vol. 33(8), August 2000, pp. 54-60.
7. J. McHugh, "Security and Quality of Service Interaction," *Proceedings 23rd National Information Systems Security Conference*, (Oct. 16-19, Baltimore, MD), NIST, Gaithersburg, MD, 2000, p. 585.
8. P. Oman, E. Schweitzer, and J. Roberts, "Safeguarding IEDs, Substations, and SCADA Systems Against Electronic Intrusions," Paper #1, *Western Power Delivery Automation Conference*, (Apr. 10-12, Spokane, WA), 2001. Available from www.selinc.com.
9. S. Dixit and Y. Ye, "Streamlining the Internet-Fiber Connection," *IEEE Spectrum*, Vol. 38(4), Apr. 2001, pp. 52-57.
10. R. Forno and W. Feinbloom, "PKI: A Question of Trust and Value," *Communications of the ACM*, Vol. 44(6), June 2001, p. 120.
11. T. Wilkinson, D. Hearn, and S. Wiseman, "Trustworthy Access Control with Untrustworthy Web Servers," *Proceedings 15th Annual Computer Security Applications Conference*, (Dec. 6-10, Phoenix, AZ), IEEE Computer Society, Los Alamitos, CA, 1999, pp. 12-21.