Toward Self-Healing Energy Infrastructure Systems

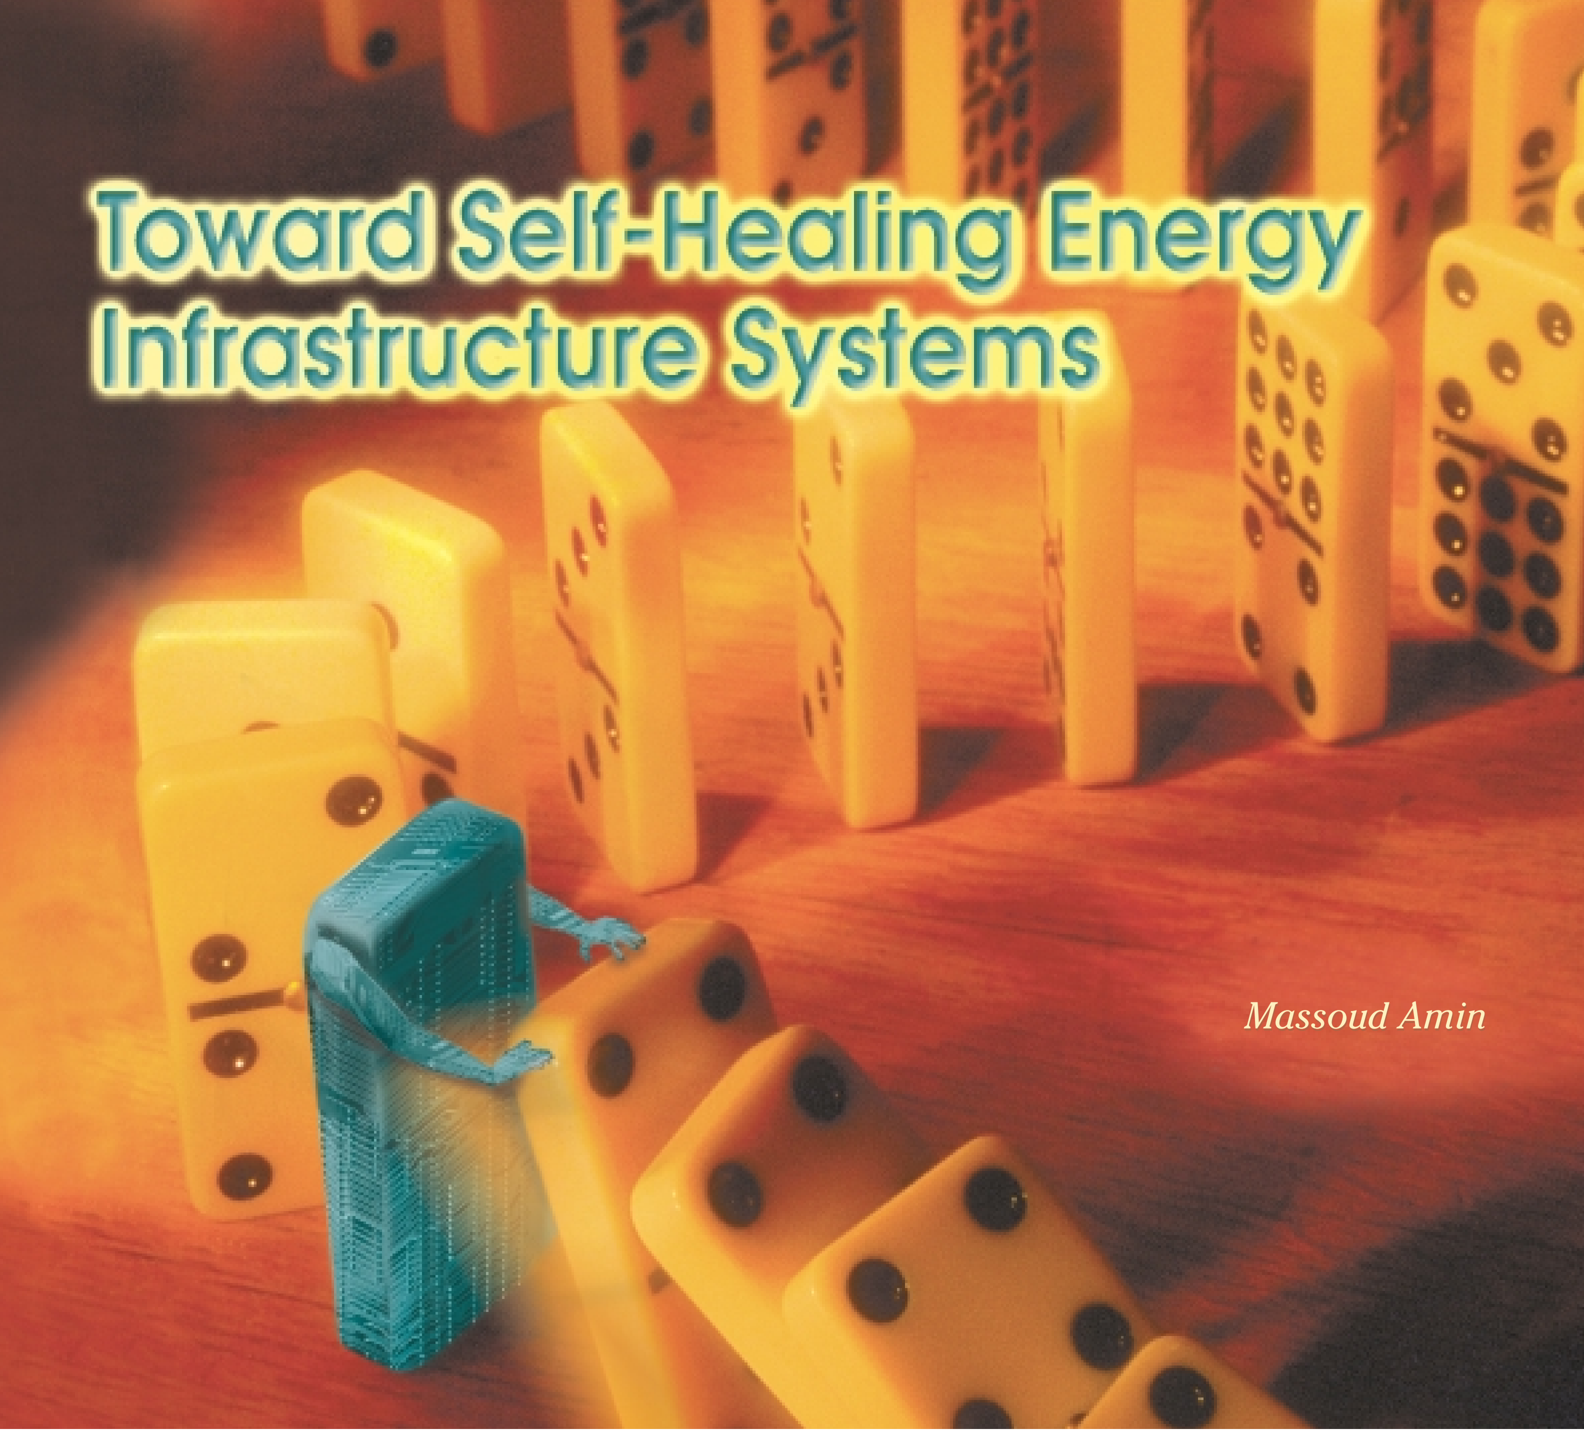

Massoud Amin

©Image Bank, Gary S. and Vivian Chapman

Virtually every crucial economic and social function depends on the secure, reliable operation of energy, telecommunications, transportation, financial, and other infrastructures. Indeed, these infrastructures have provided much of the good life that the more developed countries enjoy. However, with increased benefit has come increased risk. As they have grown more complex to handle a variety of demands, these infrastructures have become more interdependent.

The goal of the initiative is to develop new tools and techniques that will enable large national infrastructures to self-heal in response to threats, material failures, and other destabilizers

This strong interdependence means that an action in one part of one infrastructure network can rapidly create global effects by cascading throughout the same network and even into other networks. The potential for widespread disturbances is very high. Moreover, interdependence is only one of several characteristics that challenge the control and reliable operation of these networks. These characteristics, in turn, present

---

M. Amin is with EPRI, Palo Alto, California, USA.

ISSN 0895-0156/01/$10.00©2001 IEEE

unique challenges in modeling, prediction, simulation, cause-and-effect relationships, analysis, optimization, and control. What set of theories can capture a mix of dynamic, interactive, and often nonlinear entities with unscheduled discontinuities? Deregulation and economic factors and policies and human performance also affect these networks.

The Complex Interactive Networks/Systems Initiative (CIN/SI) is a joint program by Electric Power Research Institute (EPRI) and the U.S. Department of Defense (DOD) that is addressing many of these issues. The goal of the 5-year, $30 million effort, which is part of the Government-Industry Collaborative University Research (GICUR) program, is to develop new tools and techniques that will enable large national infrastructures to self-heal in response to threats, material failures, and other destabilizers. Of particular interest is how to model enterprises at the appropriate level of complexity in critical infrastructure systems.

Network Reliability and Vulnerability

From a broader historical perspective, reliable networks of energy, transportation, and communication constitute the foundation of all prospering societies. For example, the U.S. electric power grid has evolved over the last hundred years, and it now underlies every aspect of our economy and society; it has been hailed by the National Academy of Engineering as the twentieth century's engineering innovation most beneficial to our civilization. The role of electric power has grown steadily in both scope and importance during this time, and electricity is increasingly recognized as a key to societal progress throughout the world, driving economic prosperity and security and improving the quality of life. Many readers of this magazine who were born before the 1950s or born in developing countries can attest to the critical importance of electricity as a truly enabling force that powers progress and transforms societies.

The Internet, computer networks, and our digital economy have increased the demand for reliable and disturbance-free electricity; and banking and finance depend on the robustness of electric power, cable, and wireless telecommunications. Transportation systems, including military and commercial aircraft, land vehicles, and sea vessels, depend on communication and energy networks. Links between the power grid and telecommunications and between electric power and oil, water, and gas pipelines continue to be a linchpin of energy supply networks.

In the coming decades, electricity's share of total energy is expected to continue to grow, as more efficient and intelligent processes are introduced into this network. For example, controllers based on power electronics, combined with wide-area sensing and management systems, have the potential to improve situational awareness, precision, reliability, and robustness of this continental-scale system. It is envisioned that the electric power grid will move from an electromechanically controlled system into an electronically controlled network in the next 2 decades.

After reviewing the range of new devices available or under development in the power generation and delivery areas, if most of these devices are developed and are used, the main challenge facing the power engineering,

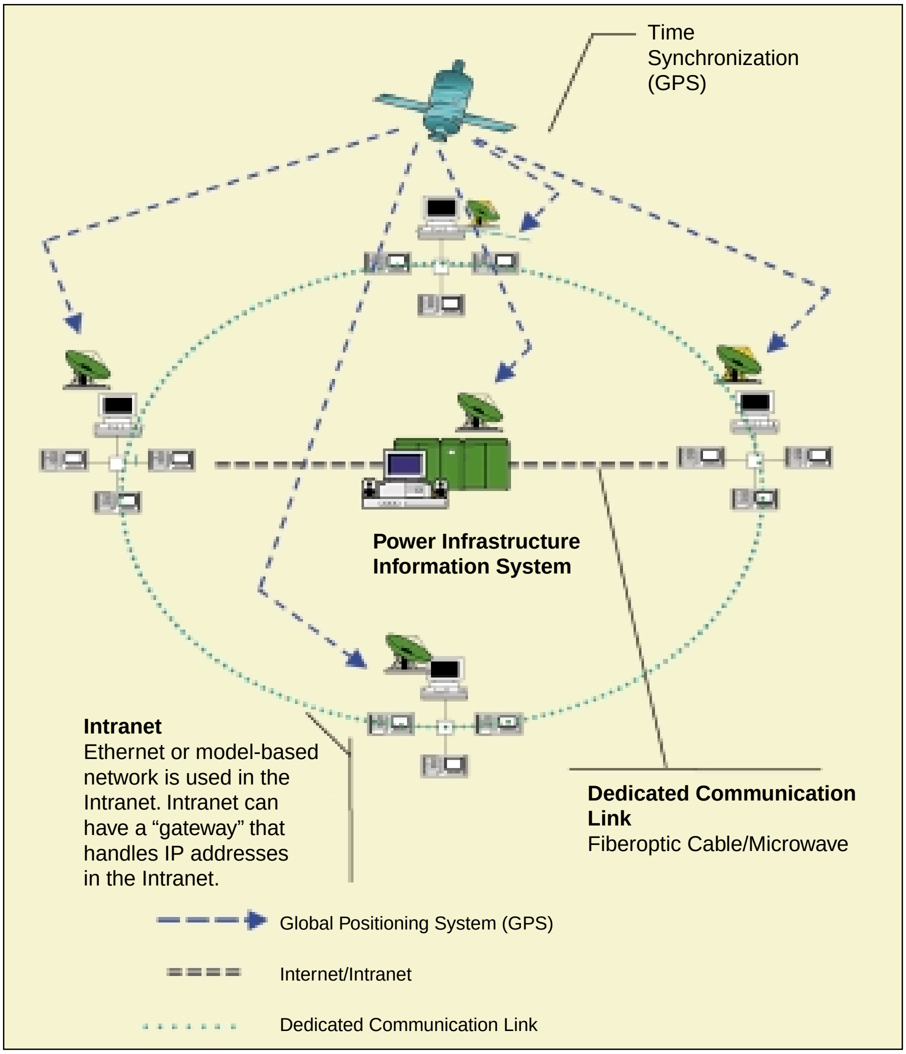

Figure 1. Telecommunications network and the electric power infrastructure (image courtesy of C.-C. Liu of the University of Washington)

computer, and control communities is how to make the entire system of the future work. The overall systems' control and robust operation will remain a major challenge. In view of this, several very timely issues have emerged, these include systematic analysis of placement of these and other devices on the network, cost/benefit and impact analyses of interaction among deployed devices, and the potential for causing unpredicted stability problems.

Another critically important dimension is the effect of deregulation and economic factors on a particular infrastructure. The electric power grid was historically operated by separate utilities, each independent in its own control area and regulated by local bodies, to deliver bulk power from generation to load areas reliably and economically. As a noncompetitive, regulated monopoly, emphasis was on reliability (and security) at the expense of economy. However, this infrastructure, faced with deregulation and coupled with interdependencies with other critical infrastructures and increased demand for high-quality and reliable electricity for our digital economy is becoming more and more stressed.

CIN/SI Objectives

Through a highly competitive source selection process, CIN/SI, which began in early 1999, has funded 6 consortia, including 28 universities and 2 utilities. Tennessee Valley Authority and Commonwealth Edison Co. are providing real-world power grid data, staff expertise, and test and demonstration sites for new modeling, measurement, control, and management tools.

The objective of CIN/SI is to significantly and strategically advance the robustness, reliability, and efficiency of the interdependent energy, communications, financial, and transportation infrastructures. Part of that work must determine if there is a unifying paradigm for simulating, analyzing, and optimizing time-critical operations.

> ***In essence, the components become context-dependent intelligent robots that cooperate to ensure successful overall operation and act independently to ensure adequate individual performance***

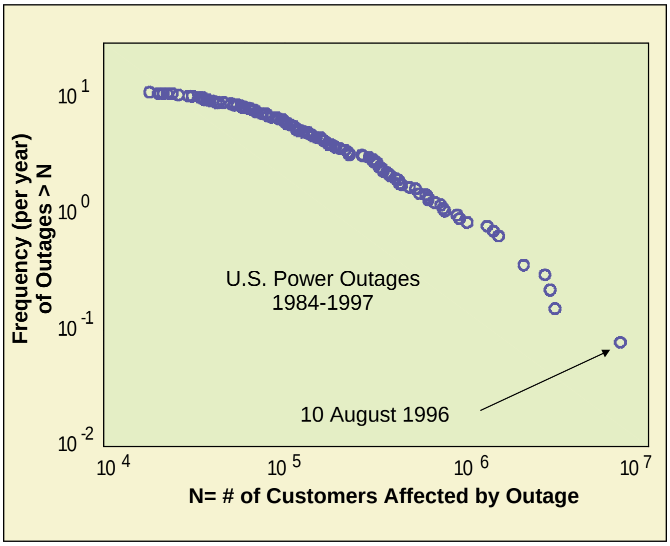

Figure 2. Major U.S. electric power outages (data from NERC, http://www.nerc.com/dawg; log-log plot courtesy of John Doyle, California Inst. of Technology)

Part of CIN/SI's work draws from ideas in statistical physics, complex adaptive systems (CAS), discrete-event dynamical systems, and hybrid, layered networks. Due to particular emphases of this magazine, the focus of this article is mainly on development of intelligent software agents to address these challenges. This area is a subset of the overall efforts; readers interested in more details on other areas, are referred to seven EPRI reports (TP-144660 through TP-114666) and the EPRI Web site (http://www.epri.com/targetST.asp?program=83). CAS researchers view the complex system as a collection of individual intelligent agents that adapt to events and surroundings, acting both competitively and cooperatively for the good of the entire system. By simulating agent-based models, stakeholders can better grasp the true dynamics of complex intercomponent and intersystem actions. As models become progressively more realistic, designers can map each system component to an adaptive agent. The adaptive agents would then manage the system using multilevel distributed control. Through its environmental sensor, each agent would receive continuous messages from other agents. If agents sense any anomalies in their surroundings, they can work together, essentially reconfiguring the system, to keep the problem local. Thus, the agents would prevent the cascading effect, the main source of vulnerability in critical infrastructure systems.

Complexity of Infrastructure Networks

Infrastructure networks have several common characteristics that make them difficult to control or to operate reliably and efficiently.

- Billions of distributed heterogeneous components are tightly interconnected. The scale is massive. For example, the time scale can range from milliseconds for one task to hours and even years for another; spatial scales can span a city or a continent.
- Attacks and disturbances can lead to widespread failure almost instantaneously.
- A variety of participants (owners, operators, sellers, buyers, customers, data and information providers, and users) interact at many points.

- The number of possible interactions increases dramatically as participants are added. No single centralized entity can evaluate, monitor, and manage all the interactions in real-time.
- The relationships and interdependencies are too complex for conventional mathematical theories and control methods. Infrastructures that interact with their users and other networks (for example, an automatic switching system for telephone calls) create additional complexity because the interaction of their elements further increases the number of possible outcomes.

In addition to these shared characteristics, each infrastructure has specific objectives that pose formidable challenges. To get an idea of the complexity involved, it helps to understand some of the issues facing the power grid, which underlies almost every other infrastructure and is vital to almost every aspect of daily living.

Power Grid

The North American power grid evolved over the past 100 years without a conscious awareness of how its evolution would affect its operation under deregulation, the digital economy, and interaction with other infrastructures. Widespread outages and huge price spikes during the past 4 years have raised public concern about grid reliability at the national level. The potential for larger scale and more frequent power disruptions is considered higher now than at any time since the great Northeast blackout in 1965. The ramifications of network failure have never been greater, as the transportation, telecommunications, oil, water and gas pipelines, banking and finance, and other infrastructures depend more and more on the power grid to energize and control their operations.

Table 1. Actions and operations within the power grid, with a time-scale variance from microseconds to years

Action or Operation	Timeframe
Wave effects (fast dynamics, such as lightning causing surges or overvoltages)	Microseconds to milliseconds
Switching overvoltages	Milliseconds
Fault protection	100 milliseconds or a few cycles
Electromagnetic effects in machine windings	Milliseconds to seconds
Stability	60 cycles or 1 second
Stability Augmentation	Seconds
Electromechanical effects of oscillations in motors and generators	Milliseconds to minutes
Tie-line load frequency control	1 to 10 seconds; ongoing
Economic load dispatch	10 seconds to 1 hour; ongoing
Thermodynamic changes from boiler control action (slow dynamics)	Seconds to hours
System structure monitoring (what is energized and what is not)	Steady state; on-going
System state measurement and estimation	Steady state; on-going
System security monitoring	Steady state; on-going
Load management, load forecasting, and generation scheduling	1 hour to 1 day or more, ongoing
Maintenance scheduling	Months to 1 year, ongoing
Expansion planning	Years, ongoing
Power plant site selection, design, construction, environmental impact, etc.	10 years or longer

The power grid is a sprawling network with many operational levels involving a range of energy sources (nuclear, fossil fuel, and renewable resources) with many interaction points (operators, power consumers and producers, and several layers including power plants, control centers, and transmission, distribution, and corporate networks). Additional complexity developed because the interaction of these elements further increased the number of possible outcomes. Because of competition and deregulation in recent years, multiple producers now share the delivery network. Demand is already out-pacing available resources in several regions: During the past decade, actual demand increased some 35%; capacity increased only 18%, because it is becoming increasingly harder for power generators and delivery entities to get permits and ensure that their return on investment is acceptable. Thus, the complex systems that relieve bottlenecks and clear disturbances during peak demand are now closer to the edge and at greater risk of serious disruption.

Another contributor to complexity is that digital users require a much higher quality of electricity. Some experts indicate that reliability will need to go from 99.9% (roughly 8 hours of power loss per year) to 99.99999999% reliability (32 seconds of power loss per year). The industry will also need new equipment to protect against sags and disruptions.

Finally, the time and operational scales at which the infrastructure operates are an important part of complexity. As Table 1 shows, the time scale for various power grid control and operation tasks can be anywhere from microseconds to a decade, which greatly complicates modeling, analysis, simulation, control, and operations tasks.

It is also important to note that the key elements and principles of operation for interconnected power systems were established prior to the emergence of extensive computer and communication networks. Computation is now heavily used in planning, design, simulation, and optimization at all levels of the power network, and computers are widely used for fast local control of equipment as well as to process large amounts of sensor data from the field. Coordination across the network happens on slower timescales. Some coordination occurs under computer control, but much of it is still based on telephone calls between system operators at the utility control centers (even, or especially, during an emergency). There is not yet a significant and intimate interaction of an extensive computer/communication network layer with the primary physical layer in the operation and control of a power system. However, economic restructuring and increasingly powerful sensing, computation, and control possibilities are changing the context in which power systems are operated and studied (e.g., see George Verghese's presentation at http://discuss.santafe.edu/dynamics/). A deeper understanding of power systems as complex interacting networks is likely to play an important role in the future.

> ***CAS researchers view the complex system as a collection of intelligent agents that adapt to events and surroundings, acting both competitively and cooperatively for the good of the entire system***

Telecommunications

The globalization of our economy is built on telecommunication networks, including fixed networks, (public switched telephone and data networks), wireless (cellular, PCS, wireless ATM), and computers (Internet and millions of computers in public and private use). These networks are growing rapidly and require secure, reliable, high-quality power supplies.

This telecommunication infrastructure, like the power grid, is becoming overburdened. The satellite network, just one segment of the infrastructure, is a good example. The satellite network has three main layers:

- Low-earth orbit (LEO), at 200 to 2,000 km (little LEOs at 750 to 1,500 km), operating at VHF, UHF below 500 MHz, with low complexity
- Medium-earth orbit (MEO), at 2,000 to 20,000 km (big LEOs and MEOs at 750 to 11,000 km), operating at Land S microwave (1.6 and 2.5 GHz), with high to very high complexity
- Geosynchronous orbit (GEO), at 36,000 km, operating at K microwave (19 and 29 GHz), with variable low to high complexity.

Some of the most familiar services are detailed earth imaging, remote monitoring of dispersed locations, and highly accurate location and tracking using the continuous signals of the global positioning system (GPS). Satellite-based business and personal voice and data services are now available throughout much of the world.

The Internet is rapidly expanding the range of applications for satellite-based data communications; two of the most popular applications are accessing the Internet itself and connecting remote sites to corporate networks. Some satellite systems, including those of satellite TV providers, let users browse Web pages and download data, at 400 kbps, through a 21-inch (53-cm) roof-mounted dish receiver connected to a personal computer with an interface card. This capability could become a valuable tool for expanding an enterprise network to remote offices around the world.

Some utilities are diversifying their businesses by investing in telecommunications and creating innovative communications networks that cope with industry trends toward distributed resources, two-way customer communications, and business expansion, as well as addressing the measurement of complex and data-intensive energy systems via wide-area monitoring and control. Figure 1 shows a possible scenario. Challenges include how to handle network disruptions and delays and manage orbits from the satellite. A major source of complexity is the interdependence of the telecommunication networks and the power grid.

Figure 1 shows how the telecommunications network and the electric power grid are becoming increasingly interdependent. Issues range from the highest command and control level (the power infrastructure information system) to the individual power stations and substations at the middle level, and then to the devices and power equipment at the lowest level. The Internet/intranet (solid blue line) connecting the middle level stations is an Ethernet or model-based network with individual gateways. The dedicated communications link is a fiberoptic cable or microwave system. The GPS handles time synchronization. In this scenario, satellite technology is used for a range of utility and business applications including direct-to-home interactive services and wide-area monitoring and control.

Cost of Cascading Failures

The occurrence of several cascading failures in the past 35 years has helped focus attention on the need to understand the complex phenomena associated with these interconnected systems. According to data from the North American Electric Reliability Council

(NERC), outages from 1984 to the present affect nearly 700,000 customers annually, or 7 million per decade (Figure 2). Many of the outages were exacerbated by cascading effects.

Perhaps the most famous recent example is the August 1996 blackout in the western North American grid. On 10 August 1996, faults in Oregon at the Keeler-Allston 500 kV line and the Ross-Lexington 230 kV line resulted in excess load, which led to the tripping of generators at McNary Dam, causing 500 MW oscillations, which led to separation of the North-South Pacific intertie near the California-Oregon border. This led to islanding and blackouts in 11 U.S. states and 2 Canadian provinces and was estimated to cost $1.5 billion to $2 billion and included all aspects of interconnected infrastructures and even the environment. Among several analyses that followed, some researchers have shown that dropping (shedding) approximately 0.4% of the total network load for 30 minutes would have prevented the cascading effects of the August 1996 blackout.

Past disturbances in both the power grid and telecommunications infrastructures provide some idea of how cascading failures work. In some cases, the local disturbance affected geographically distant areas. In others, a failure in one infrastructure led to breakdowns in other infrastructures. Because these and other infrastructures support critical services and supply critical goods, disturbances can have serious economic, health, and security impacts. Therefore, there is a need to develop an ability for these infrastructures to self-heal and self-organize at the local level in order to mitigate the effects of such disturbances.

In most critical infrastructure networks, systems are spread across vast distances, are nonlinear, and are highly interactive. In any situation subject to rapid changes (from natural disasters, purposeful attack, or unusually high demands), completely centralized control requires multiple, high-data-rate, two-way communication links, a powerful central computing facility, and an elaborate operations control center. But centralized control may not be practical in this setting, because a failure in one part of the network can spread unpredictably almost instantaneously, including to the centralized control elements. Thus, centralized control is likely to suffer from the very problem it is supposed to fix. A pertinent question is how to manage and robustly operate these systems that have hierarchical layers and are distributed at each layer. An alternative strategy is to have some way to intervene locally (where the disturbance originates) and stop problems from propagating through the network.

Agent Technology

Infrastructures are highly interconnected and interactive, making them well suited for agent technology. Indeed, infrastructure networks already use agents in the form of decision-making and control units distributed among layers throughout physical, financial, and operational subsystems (including supervision, maintenance, and management). Agents assess the situation on the basis of measurements from sensing devices and information from other entities. They influence network behavior through commands to actuating devices and other entities. The agents range in sophistication from simple threshold detectors, which choose from a few intelligent systems on the basis of a single measurement, to highly intelligent systems.

The North American power grid has thousands of such agents, and power system dynamics are extremely complex. Actions can take place in microseconds (such as a lightning strike), and the network's ability to communicate data globally is limited. For these reasons, no one can preprogram the agents with the best responses for all possible situations. Thus, each agent must make real-time decisions from local, rather than global, state information. Many agents (particularly, controllers for individual devices) are designed with relatively simple decision rules based on response thresholds that are expected to give the most appropriate responses to a collection of situations generated in offline studies.

Context-Dependent Agents

This approach does not offer sufficient reliability, however. Power grid agents have been known to take actions that drive the system into undesirable operating states. In some cases, the agents acted as programmed, but the predesigned actions were not the best responses to the actual situation, the context. In many cases, the agent could have been made aware of the context and thus would have known that the preprogrammed action was not appropriate.

Context dependence is a key difference between agents as they are currently designed and the adaptive agents that CIN/SI researchers are developing. In a context-dependent agent-based network, agents cooperate and compete with each other in their local operations while simultaneously pursuing the global goals set by a minimal supervisory function. In the power grid, for example, a network of local controllers would act as a parallel, distributed computer, communicating via microwaves, optical cables, or the power lines themselves, and intelligently limit their messages to information needed to optimize the entire grid and recover from a failure. Thus, in essence, the components become context-dependent intelligent robots that cooperate to ensure successful overall operation and act independently to ensure adequate individual performance.

Agent Evolution

The agents evolve, gradually adapting to their changing environment and improving their performance even as

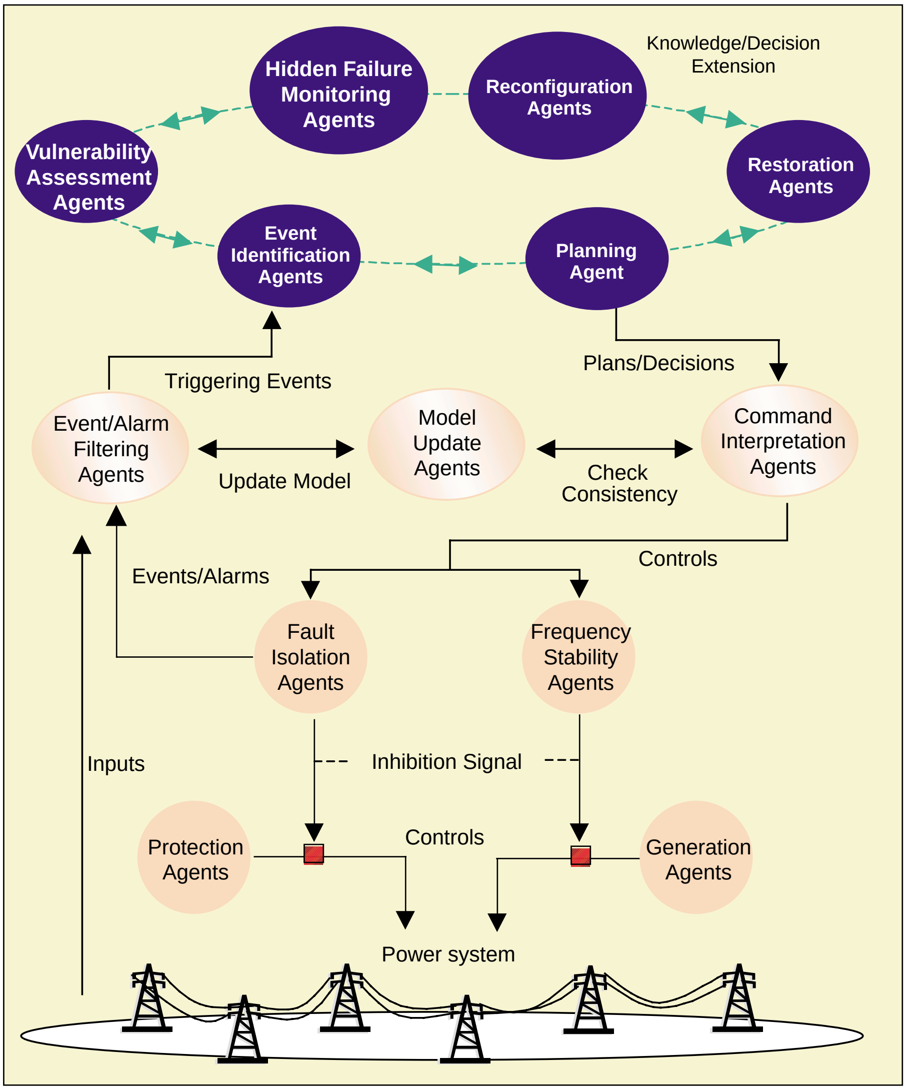

Figure 3. A multiagent system design (image courtesy of C.-C. Liu of the University of Washington)

conditions change. A single bus would strive to stay within its voltage and power flow limits while still operating in the context of the voltages and flows that power system managers and other agents impose on it. Advanced sensors, actuators, and microprocessors would be associated with generators, transformers, buses, and so on. Modelers use object-oriented methods and object hierarchies of simpler components to model more complex components, such as a generating plant or a substation, thus creating a hierarchy of adaptive agents.

So that it is aware of context and can evolve, each agent and subagent that is represented as an autonomous active object is equipped with appropriate algorithms (intelligence). Evolution is enabled through a combination of genetic algorithms and genetic programming. Classes are treated as an analogy of biological genotypes, and objects are instantiated from them as an analogy of their phenotypes. When instantiating objects to form individual agents, operations typical of genetic algorithms, such as crossover and mutation, can select and recombine their class attributes, which define all the potential characteristics, capabilities, limitations, or strategies these agents might possess.

The physics specific to each component will determine the object-agent's allowable strategies and behaviors. Researchers can augment existing instrumentation and control capabilities and run computer experiments with hypothetical, optional capabilities to evaluate their benefit.

Figure 3 describes how the agents are organized in three layers. The reactive layer (bottom) consists of agents that perform preprogrammed self-healing actions that require an immediate response. Reactive agents, whose goal is autonomous and fast control, are in every local subsystem. The agents in the coordination layer (middle) include heuristic knowledge to identify which triggering event from the reactive layer is urgent, important, or resource consuming. These agents, whose goal is consistency, also update the system's real-world model and check if the plans (or commands) from the deliberative layer (top) represent the system's current status. If the plans do not match the real-world model, the agents in the middle layer trigger the deliberative layer to modify the plans. The deliberative layer consists of cognitive agents that have goals and explicit plans that let them achieve their goals. The goals of agents in this layer are dependability, robustness, and self-healing.

As part of the research on context-dependent network agents, investigators are developing a robust, dynamic, and real-time computing architecture that will:

- Ensure the robustness of the software infrastructure using an analytically redundant software architecture with two complementary components, a simple and highly reliable core component that guarantees the minimal essential services and a complex component that provides many desirable features, such as the ability to replace the control agents without the need to shut down and then restart the normal operations. The useful but noncritical complex component will extensively use commercially available software components to lower the cost. The reliable core will function in spite of failures in the complex component and will provide the network state information to restart the complex component should it fail.

- Provide timely and consistent contexts for distributed agents. The stochastic events arising from the dynamics of the power network drive the coordination between distributed agents. An event-driven real-time communication architecture will assemble relevant distributed agents into task-driven teams and will provide the teams with timely and consistent information to carry out coordinated actions.

Multiagent Systems

A multiagent power grid system uses two types of agents, cognitive (rational) and reactive.

Each cognitive agent has a knowledge base that comprises all the data and know-how required to carry out its task and to handle interactions with the other agents and its environment. Cognitive agents are also intentional, in that they have goals and explicit plans that let them achieve their goals.

The reactive agent, in contrast, claims that it is not necessary for agents to be individually intelligent for the system to demonstrate intelligent behavior overall. Their active agents work in a hard-wired, stimulus-response manner. The reactive agent's goals are only implicitly represented by rules (or simple logic), so it must consider each and every situation in advance. The reactive agent's advantage lies in its ability to react fast.

As Figure 3 shows, this multiagent system has three layers. The reactive layer (bottom) is in every local subsystem and performs preprogrammed self-healing actions that require an immediate response. The agents in the middle layer, the coordination layer, include heuristic knowledge to identify which triggering event from the reactive layer is urgent, important, or resource consuming. If a triggering event exceeds a threshold value, this agent will allow the event to go to the deliberative layer, which contains the cognitive agents. The agents in the deliberative layer develop plans according to their virtual models, which they keep current with information from the coordination layer. However, the virtual world model could be outdated because the agents in the deliberative layer do not always respond to the current situation.

For this reason, the agents in the coordination layer continuously compare the world models between the deliberative and reactive layers. They update the current real-world model and check if the plans (or commands) from the deliberative layer represent the system's current status. If the plans do not align with the real-world model, the agents in the coordination layer trigger the deliberative layer to modify the plans.

In addition, events from the reactive layer might contain too much detailed information for the agents in the deliberative layer. On the other hand, the plans from the deliberative layer might be too condensed for the agents in the reactive layer. There may be more than a few control signals in the reactive layer originating from the deliberative layer. The coordination layer analyzes the command and decomposes it into actual control signals. This layer might be at every local subsystem that interfaces with the reactive layer. The agents in the deliberative layer prepare higher level plans, such as vulnerability assessment and self-healing.

Modeling the power industry in this control-theory context is especially pertinent, since the current movement toward deregulation and competition will ultimately be limited only by the physics of electricity and the grid's topology. A CAS simulation will test whether any central authority is required, or even desirable, and whether free economic cooperation and competition can, by itself, optimize the efficiency and security of network operation for the benefit of all.

> *Economic restructuring and increasingly powerful sensing, computation, and control options are changing the context in which power systems are operated and studied*

Infrastructures for a Digital World

No one is outside the infrastructure, and there are clearly many opportunities for modeling and simulation, as well as for the use of computers, machine intelligence, and human performance engineering. Agent-based modeling and CAS are only a fraction of what's involved in capturing the level of complexity in infrastructure systems. Modeling complex systems is one of three main areas in CIN/SI's charter. The others are measurement, i.e., to know what is or will be happening and develop measurement techniques for visualizing and analyzing large-scale emergent behavior, and management, i.e., to develop distributed management and control systems to keep infrastructures robust and operational. Some specific areas are being investigated.

- Robust control: Extend the theory of robust control (managing the system to avoid cascading failure in the face of destabilizing influences such as enemy threats or lightning strikes) beyond the relatively narrow problem of feedback control
- Disturbance propagation: Predict and detect the onset of failures at both the local and global levels. This includes establishing thresholds for identifying when instabilities trigger failures.
- Complex systems: Develop theoretical underpinnings of complex interactive systems.

- Dynamic interaction in interdependent layered networks: Create models that capture network layering at many levels of complexity.
- Modeling in general: Develop efficient simulation techniques and ways to create generic models. Develop a modeling framework and analytical tools to study the dynamics and failure modes in the interaction of economic markets with power and transportation systems.
- Forecasting network behavior and handling uncertainty and risk: Characterize uncertainty in large distributed networks. Stochastically analyze network performance. Investigate handling rare events through large-deviations theory.

An October 1997 report from the U.S. President's Commission on Critical Infrastructure Protection (PCCIP) cited the growing importance of infrastructure networks in many application areas. The PCCIP report and subsequent studies recognized the damaging and even dangerous ways cascading failures can affect the economy, security, and health of U.S. citizens in unpredictable ways. Indeed, even the weather can create cascading effects. In the summer of 1998, for example, temperatures were considerably above normal, the power demand increased, the transmission capacity could not meet it, and prices in the Midwest jumped from $30 to $50 per MWh to $7,000 per MWh (http://www.ferc.fed.us/electric/mastback.pdf). Similar price spikes of 100-fold have been experienced during peak demand.

CIN/SI represents a huge undertaking, and, long after the initiative is over in 2003, work will continue on the foundation it provides. The EPRI Electricity Technology Roadmap shows approximate milestones for the larger effort to resolve infrastructure vulnerability:

- By 2003, strengthen the power delivery infrastructure. Resolve electric power infrastructure vulnerability threats.
- By 2005, enable customer-managed service networks. Build an integrated services delivery network as the superhighway system for e-commerce.
- By 2010, boost economic productivity and prosperity. Create the advanced electrotechnology platforms needed to accelerate productivity growth and global competition.
- By 2015, resolve the energy/carbon conflict. Electrify the world to stimulate more efficient patterns of production and consumption.
- By 2025, manage global sustainability.

As these milestones show, CIN/SI's immediate and critical goal is to avoid widespread network failure. Although "resolve vulnerability threats" has many forms (DOD is more concerned with enemy threats, and EPRI with natural disasters and material failures), there is little difference in the effects and recovery task, whether lightning or a terrorist destroys the power pole. The milestones are ambitious: Achieving and sustaining infrastructure reliability, robustness, security, and efficiency requires strategic investments in research and development. Given economic, societal, and quality-of-life issues and the ever-increasing interactions and interdependencies among infrastructures, this objective offers exciting scientific and technological challenges.

Acknowledgments

An earlier version of this article was published in the August 2000 issue of *IEEE Computer* magazine and smaller portions in the *IEEE Control Systems* magazine, August 2000.

CIN/SI is cofunded by EPRI and DOD through the GICUR program. Within DOD, I thank Delores Etter, deputy under secretary of defense, and Robert Trew, director of research, funding the DOD part of this GICUR program in the Office of the Director, Defense Research and Engineering. Within the U.S. Army Research Office (ARO), I thank CIN/SI manager Robert Launer, ARO director Jim Chang, Richard Smith, and Mitra Dutta, as well as Robert Singleton, formerly of ARO. Several colleagues formerly with DOD were involved in the initiative's planning stages, including Arthur Diness, Laura S. Rea, George Singley, and James Garcia. During the formative stages of the initiative, Anita Jones of the University of Virginia was director of Defense Research and Engineering. Within EPRI, Gail Kendall, Martin Wildberger, Revis James, Ram Adapa, Aty Edris, Paul Grant, Hung-po Chao, Richard Lordan, Dejan Sobajic, Peter Hirsch, Steve Lee, Steve Gehl, John Stringer, and many others provided technical input and other assistance. I am also grateful for the contributions of the 108 faculty members, as well as dozens of other researchers and students in CIN/SI-funded universities.

Further Reading

M. Amin, "National infrastructures as complex interactive networks," in *Automation, Control, and Complexity: An Integrated Approach*, T. Samad and J.R. Weyrauch, Eds., New York: Wiley, 2000 ch. 14, pp. 263-286.

"Complex interactive networks/systems initiative: overview and progress report for joint EPRI and U.S. Department of Defense university research initiative," EPRI, Palo Alto, CA, Rep. TP-114660, first annual rep., 2000.

"Critical foundations: protecting America's infrastructures," U.S. President's Commission on Critical Infrastructure Protection (PCCIP), Washington, DC, Oct. 1997. Available: http://www.ciao.ncr.gov.

"Electricity technology roadmap: 1999 summary and synthesis," EPRI, Palo Alto, CA, Rep. CI-112677-V1, 1999. Available: http://www.epri.com.

J.F. Hauer and J. E. Dagle, "Review of recent reliability issues and system events," PNNL rep. PNNL-13150, 1999.

"Utility communications go into orbit," *EPRI J.*, pp. 18-27, Fall 1999.

Biography

Massoud Amin is manager of mathematics and information science at EPRI, Palo Alto, California, where he leads strategic research in complex interactive networks, including national infrastructures for energy, telecommunication, transportation, and finance. He is the author or coauthor of more than 75 research papers on the theoretical and practical aspects of online decision support, system optimization, differential game theory, and intelligent and adaptive control of uncertain and large-scale systems. He received a B.S. (cum laude) and an M.S. in electrical and computer engineering from the University of Massachusetts at Amherst and an M.S. and a Ph.D. in systems science and mathematics from Washington University, St. Louis. He is a member of Sigma Xi, Tau Beta Pi, Eta Kappa Nu, AAAS, AIAA, NY Academy of Sciences, the IEEE, SIAM, and Informs. He may be reached by e-mail, mamin@epri.com.