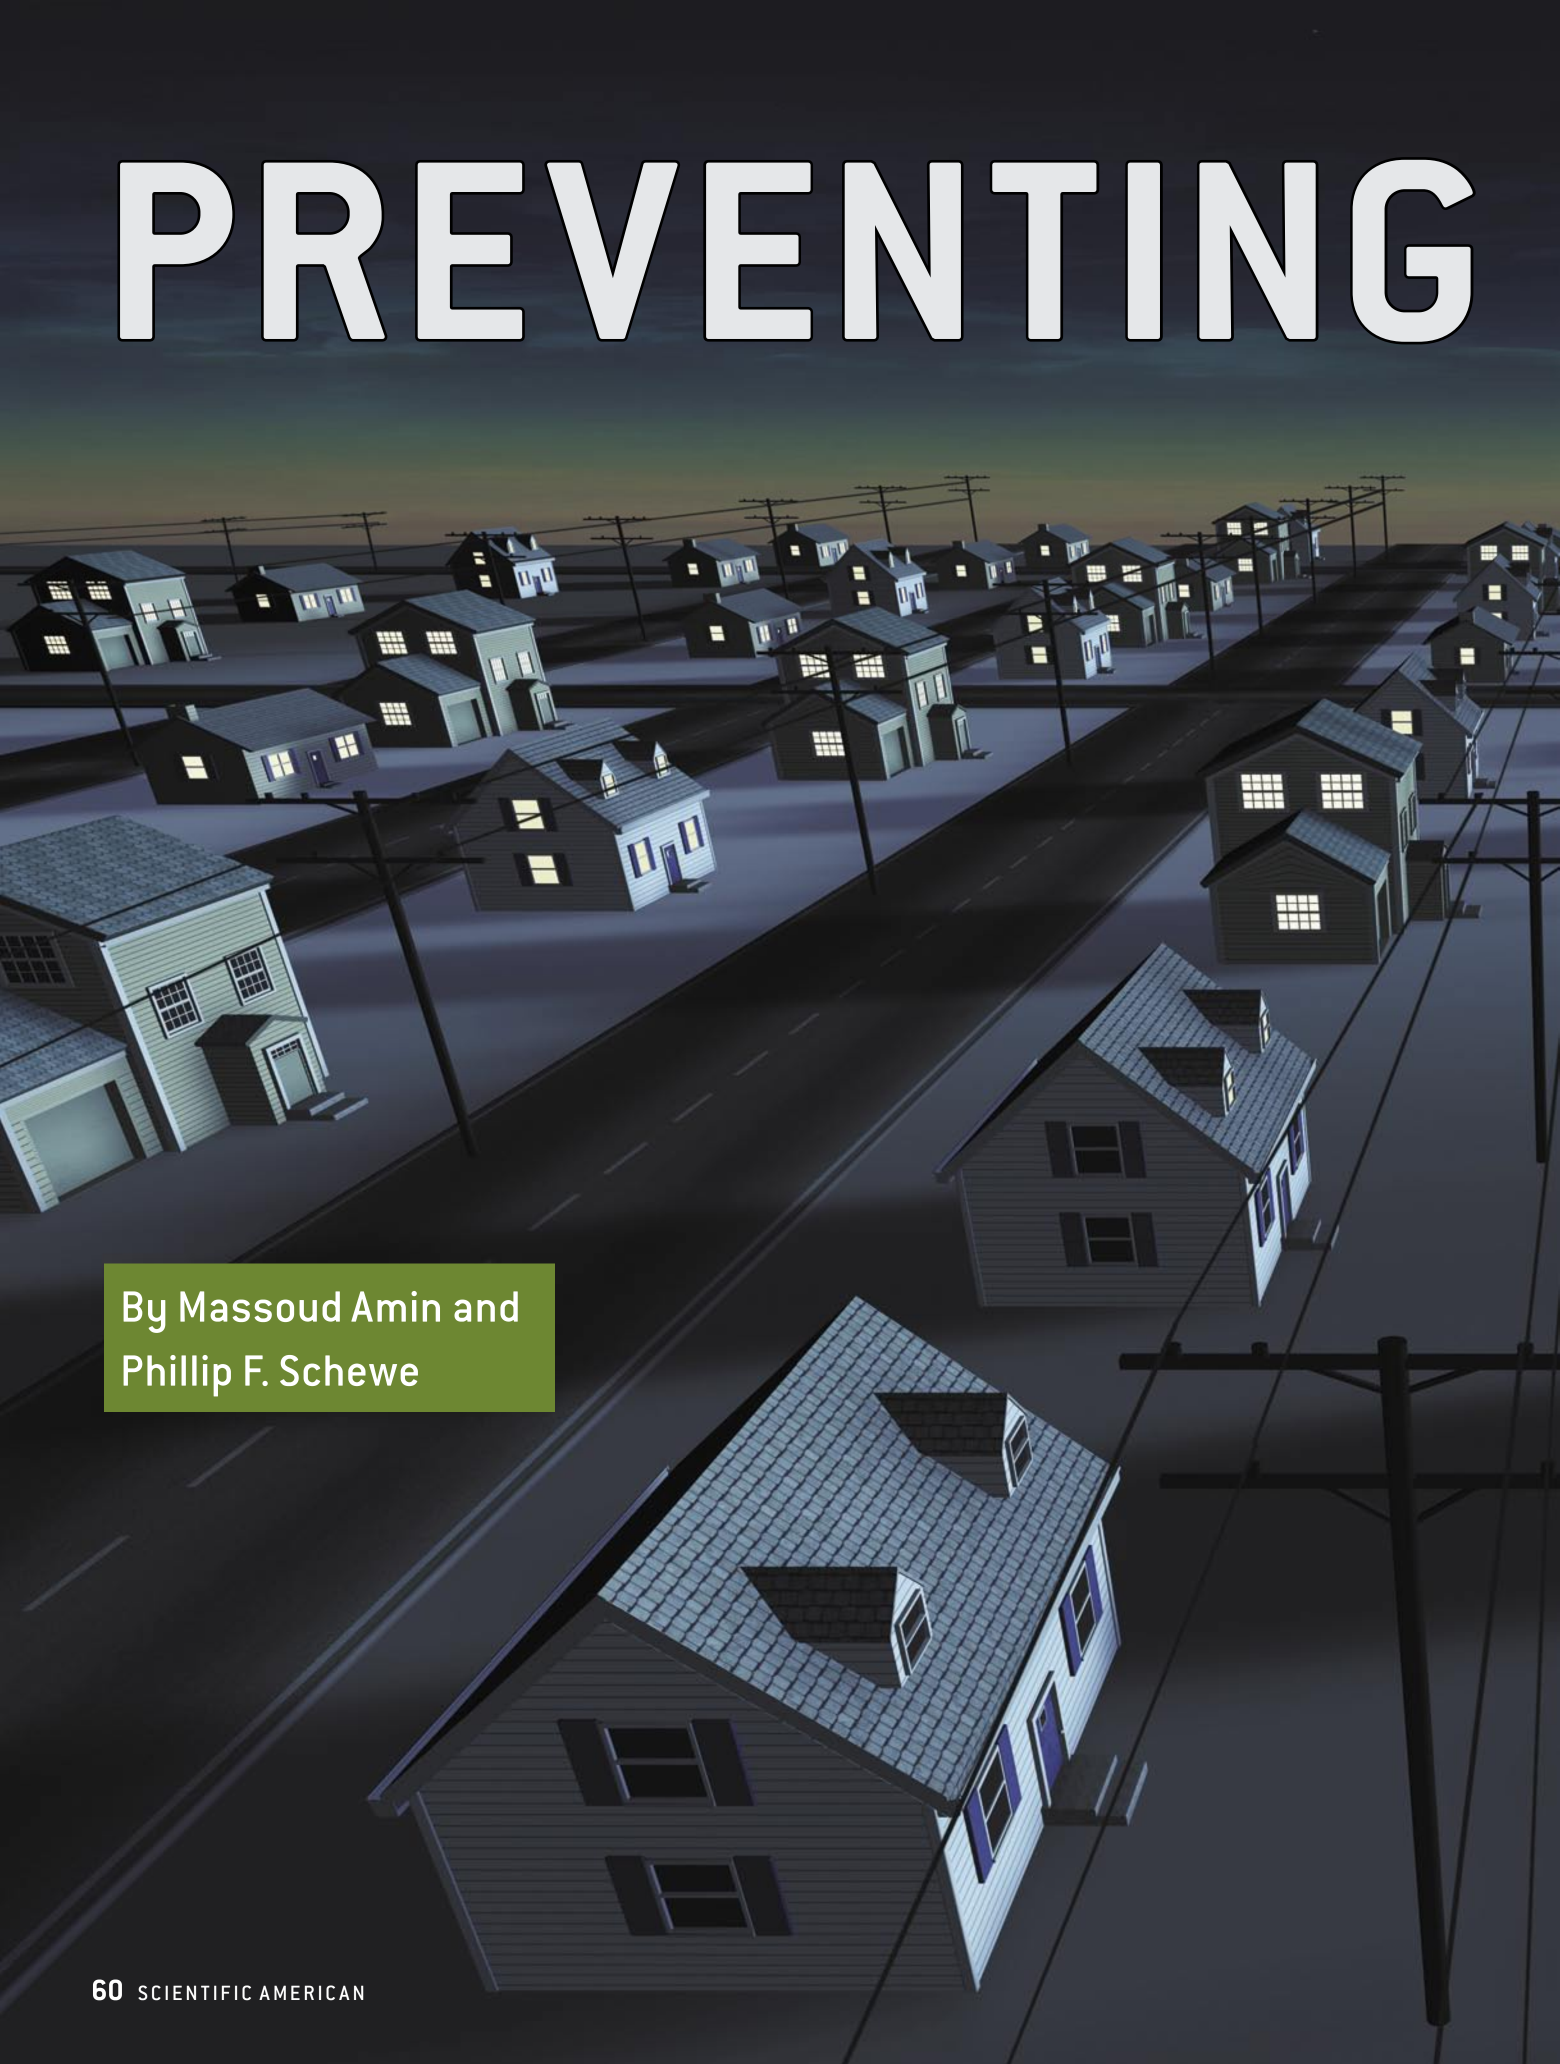

PREVENTING

By Massoud Amin and
Phillip F. Schewe

BLACKOUTS

A smarter power grid that automatically responds to problems could reduce the rising number of debilitating blackouts

RANDY LYHUS

August 14, 2003, was a typical warm day in the Midwest. But shortly after 2:00 P.M. several power lines in northern Ohio, sagging under the high current they were carrying, brushed against some overgrown trees and shut down. Such a disturbance usually sets off alarms in a local utility's control room, where human operators work with controllers in neighboring regions to reroute power flows around the injury site.

On this day, however, the alarm software failed, leaving local operators unaware of the problem. Other controllers who were relaying, or "wheeling," large amounts of power hundreds of miles across Ohio, Michigan, the northeastern U.S. and Ontario, Canada, were oblivious, too. Transmission lines surrounding the failure spot, already fully taxed, were forced to shoulder more than their safe quota of electricity.

To make matters worse, utilities were not generating enough "reactive power"—an attribute of the magnetic and electric fields that move current along a wire. Without sufficient reactive power to support the suddenly shifting flows, overburdened lines in Ohio cut out by 4:05 P.M. In response, a power plant shut down, destabilizing the system's equilibrium. More lines and more plants dropped out. The cascade continued, faster than operators could track with the decades-old monitoring equipment that dots most of the North American power grid, and certainly much faster than they could control. Within eight minutes 50 million people across eight states and two Canadian provinces had been blacked out. The event was the largest power loss in North American history.

The 2003 disaster was a harbinger, too. Within two months, major blackouts occurred in the U.K., Denmark, Sweden and Italy. In September 2003 some 57 million Italians were left in the dark because of complications in transmitting power from France into Switzerland and then into Italy. In the U.S., the annual number of outages affecting 50,000 or more customers has risen for more than a decade.

In addition to inconvenience, blackouts are causing major economic losses. The troubles will get worse until the entire transmission system that moves power from generating plants to neighborhood substations is overhauled. More high-voltage lines must be built to catch up with the rising demand imposed by ever more air conditioners, computers and rechargeable gadgets.

But perhaps even more important, the power grid must be made smarter. Most of the equipment that minds the flow of electricity dates back to the 1970s. This control system is not good enough to track disturbances in real time—as they happen—or to respond automatically to isolate problems before they snowball. Every node in the power grid should be awake, responsive and in communication with every other node. Furthermore, the information that operators receive at central control stations is sparse and at least 30 seconds old, making it impossible for them to react fast enough to stop the large cascades that do start. A self-healing smart grid—one that is aware of nascent trouble and can reconfigure itself to resolve the problem—could reduce blackouts dramatically, as well as contain the chaos that could be triggered by terrorist sabotage. It would also allow more efficient wheeling of power, saving utilities and their customers millions of dollars during routine operation. The technology to build this smart grid largely exists, and recent demonstration projects are proving its worth.

Overwhelmed by Progress

THE TRANSMISSION SYSTEM has become vulnerable to blackouts because of a century-long effort to reduce power losses. As power moves through a wire, some of it is wasted in the form of heat. The loss is proportional to the amount of current being carried, so utilities keep the current low and compensate by raising the voltage. They have also built progressively longer, higher-voltage lines to more efficiently deliver power from generation plants to customers located far away. These high-voltage lines also allow neighboring utilities to link their grids, thereby helping one another sustain a critical balance between generation supply and customer demand.

THE PROBLEM: TOO MANY PLAYERS, TOO LITTLE INVESTMENT

Restructuring of the utility industry (*below*) because of deregulation is a major reason why U.S. blackouts are increasing (*opposite page, bottom graph*). With a single company no longer in charge of the action in a given region, the power grid is not being upgraded or expanded to keep pace with growing demand (*opposite page, top panels*).

DEREGULATION HAS FRAGMENTED CONTROL OF THE POWER SYSTEM

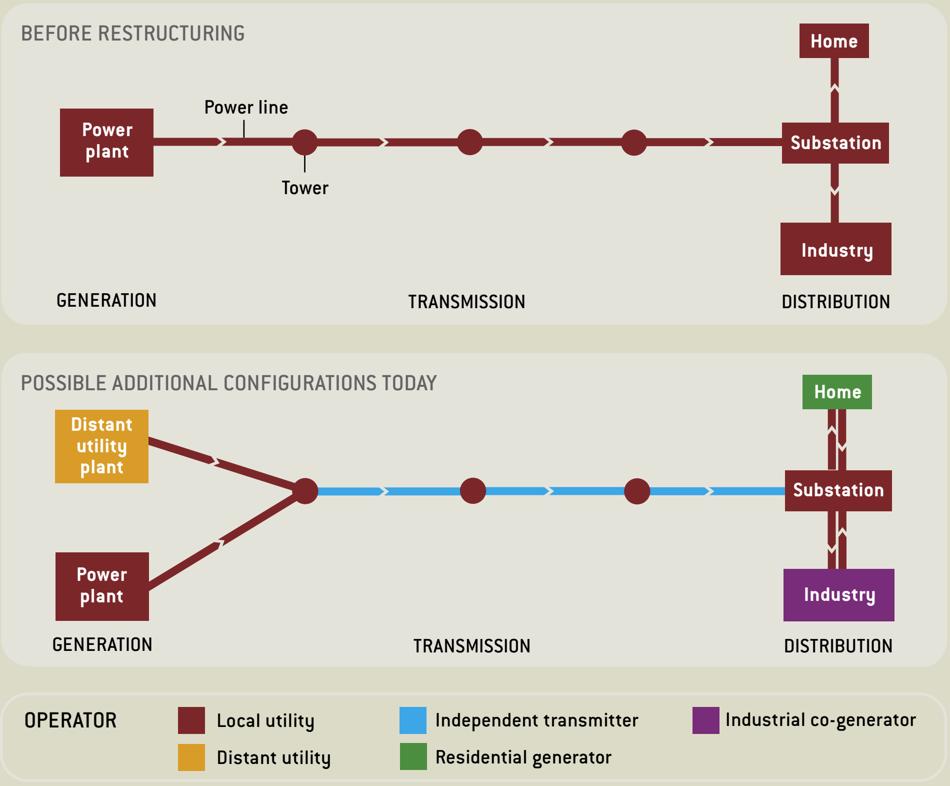

Such interconnectedness entails certain dangers, however, including the possibility that a shutdown in one sector could rapidly propagate to others. A huge 1965 blackout in the Northeast prompted utilities to create the North American Electric Reliability Council—now called the North American Electric Reliability Corporation (NERC)—to coordinate efforts to improve system reliability. Similar bodies, such as Europe's Union for the Coordination of Transmission of Electricity, exist around the world.

Why, then, had the U.S. grid become vulnerable enough to fail massively in 2003? One big reason is that investment in upgrading the transmission system has been lacking. Sharply rising fuel prices in the 1970s and a growing disenchantment with nuclear power prompted Congress to pass legislation intended to allow market competition to drive efficiency improvements. Subsequent laws

Overview/*Smart Grid*

- Demand for electricity has increased steadily for decades, yet transmission lines that transport power from generation plants to customers have not been added or upgraded at the same pace. As a result, the grid has become overloaded, making it more prone to blackouts, which have risen in number and severity and cost the U.S. more than $70 billion in annual economic losses.
- Even with more lines, a self-healing smart grid that can sense local problems early, and automatically fix or isolate them before they snowball, is needed to prevent the cascading power failures that cause blackouts.
- Digital controllers and real-time communications devices must be placed on every transmission line, substation, power plant and utility operations center.
- The operations centers also need updated computers and software that enable human controllers to manually take over the automated smart grid if a blackout does somehow begin. And controllers require better training to know how to react quickly.

JEN CHRISTIANSEN; SOURCES: UNIVERSITY OF MINNESOTA, EDISON ELECTRIC INSTITUTE, NORTH AMERICAN ELECTRIC RELIABILITY CORP., U.S. DEPARTMENT OF ENERGY

MONEY SPENT ON U.S. TRANSMISSION LINES IS DECLINING

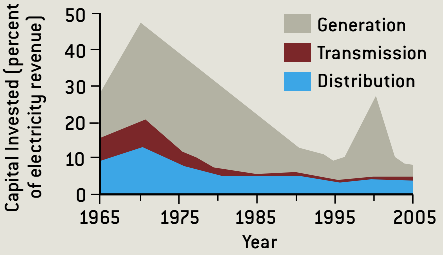

TRANSMISSION CAPACITY IS FALLING BEHIND GROWTH IN DEMAND

20%
INCREASE IN
DEMAND
FOR ELECTRICITY
(1999–2009)

7%
INCREASE IN
TRANSMISSION
CAPACITY
(1999–2009)

RESULT: LARGE BLACKOUTS ARE GROWING IN NUMBER AND SEVERITY

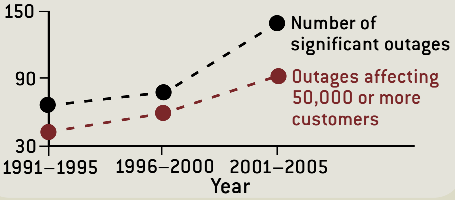

have instigated a sweeping change in the industry that has come to be called restructuring. Before restructuring began in earnest in the 1990s, most utilities conducted all three principal functions in their region: generating power with large plants, transmitting it over high-voltage lines to substations, then distributing it from there to customers over lower-voltage lines. Today many independent producers sell power near and far over transmission lines they do not own. At the same time, utilities have been selling off parts of their companies, encouraged by the Federal Energy Regulatory Commission to further promote competition. Gradually the transmission business has become a confusing mixture of regulated and unregulated services, with various companies controlling fragmented pieces.

Investors have found generation, now largely deregulated, to be attractive. But because the transmission system has been only partially deregulated, uncertainty over its fate makes investors wary. (Deregulation of distribution is still in its infancy.) Meanwhile, even though wheeling occurred in the past, since the 1990s much larger amounts of power have been moved over great distances. As a result, massive transfers are flowing over transmission lines built mostly by utilities for local use decades ago.

Proposed federal legislation might encourage more investment, but even if transmission capacity is added, blackouts will still occur. The entire power grid has to be refurbished, because the existing control technology—the key to quickly sensing a small line failure or the possibility of a large instability—is antiquated. To remain reliable, the grid will have to operate more like a fighter plane, flown in large part by autonomous systems that human controllers can take over if needed to avert disaster.

A Need for Speed

MODERN WARPLANES are so packed with sophisticated gear that pilots rely on a network of sensors and automatic controls that quickly gather information and act accordingly. Fortunately, the software and hardware innovations required to fly the power grid in a similar fashion and to instantly reroute power flows and shut down generation plants are at hand.

Reconfiguring a widely interconnected system is a daunting challenge, though. Most power plants and transmission lines are overseen by a supervisory control and data acquisition (SCADA) system. This system of simple sensors and controllers provides three critical functions—data acquisition, control of power plants, and alarm display—and allows operators who sit at central control stations to perform certain tasks, such as opening or closing a circuit breaker. SCADA monitors the switches, transformers and pieces of small hardware, known as programmable logic controllers and remote terminal units, that are installed at power plants, substations, and the intersections of transmission and distribution lines. The system sends information or alarms back to operators over telecommunications channels.

SCADA technology goes back 40 years, however. Much of it is too slow for today's challenges and does not sense or control nearly enough of the components around the grid. And although it enables some coordination of transmission among utilities, that process is extremely sluggish, much of it still based on telephone calls between human operators at the utility control centers, especially during emergencies. What is more, most programmable logic controllers and remote terminal units were developed before industry-wide standards for interoperability were established; hence, neighboring utilities often use incompatible control protocols. Utilities are operating ever closer to the edge of the stability envelope using 1960s-era controls.

The Self-Healing Smart Grid

THE RESULT is that no single operator or utility can stabilize or isolate a transmission failure. Managing a modern grid in real time requires much more automatic monitoring and far greater interaction among human operators, computer systems, communications networks and data-gathering sensors that need to be deployed everywhere in power plants and substations. Reliable operation also requires multiple, high-data-rate, two-way communications links among all these nodes, which do not exist today, plus powerful computing facilities at the control center. And intelligent processors—able to automatically reconfigure power flows when precursors to blackouts are sensed—must be distributed across the network.

Flying the grid begins with a different kind of system design. Recent research from a variety of fields, including nonlinear dynamical systems, artificial intelligence, game theory and software engineering, has led to a general theory of how to design complex systems that adapt to changing conditions. Mathematical and computational techniques developed for this young discipline are providing new tools for grid engineers. Industry working groups, including a group run by one of us (Amin) while at the Electric Power Research Institute (EPRI) in Palo Alto, Calif., have proposed complex adaptive systems for large regional power grids. Several utilities have now deployed, at a demonstration scale, smart remote terminal units and programmable controllers that can autonomously execute simple processes without first checking with a human controller, or that can be reprogrammed at a distance by operators. Much wider implementation is needed.

A self-healing smart grid can best be built if its architects try to fulfill three primary objectives. The most fundamental is real-time monitoring and reaction. An array of sensors would monitor electrical parameters such as voltage and current, as well as the condition of critical components. These measurements would enable the system to constantly tune itself to an optimal state.

The second goal is anticipation. The system must constantly look for potential problems that could trigger larger disturbances, such as a transformer that is overheating. Computers would assess trouble signs and possible consequences. They would then identify corrective actions, simulate the effectiveness of each action, and present the most useful responses to human operators, who could then quickly implement corrective action by dispatching the grid's many automated control features. The industry calls this capability fast look-ahead simulation.

The third objective is isolation. If failures were to occur, the whole network would break into isolated "islands," each of which must fend for itself. Each island would reorganize its power plants and transmission flows as best it could. Although this might cause voltage fluctuations or even small outages, it would prevent the cascades that cause major blackouts. As line crews repaired the failures, human controllers would prepare each island to smoothly rejoin the larger grid. The controllers and their computers would function as a distributed network, communicating via microwaves, optical fibers or the power lines themselves. As soon as power flows were restored, the system would again start to self-optimize.

THE SOLUTION: A SMART GRID THAT HEALS ITSELF

Imagine that a thunderstorm knocks out power lines L5 and L6. This occurrence would typically cause a chain reaction of line faults that would black out Area 1. But a smart grid would isolate and correct the problem as depicted below. The action begins as a look-ahead computer at the control center simulates corrective actions in less than half a second and sends instructions to control computers around the grid.

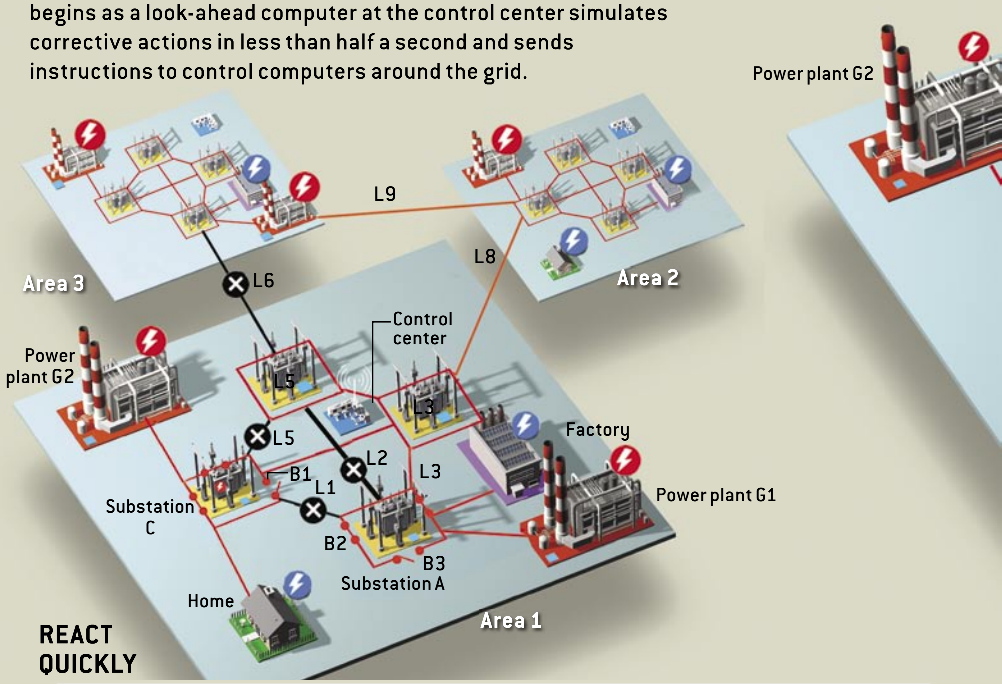

REACT QUICKLY

0.04 second later
The loss of L5 and L6 causes a fault in line L1. Control computers tell circuit breakers B1 and B2 to open to isolate the fault, but B2 becomes stuck in the closed position.

0.1 second
Power generator G1 automatically accelerates to meet demand from the loss of G2 caused by problems on lines L5 and L1. G1 also accelerates to attempt to keep line voltage throughout Area 1 at the required 60 hertz (cycles per second).

0.4 second
The control computer-simulator in substation A tells breaker B3 to open to protect the substation against damage from excessive current flow through it. B3 opens, shutting down line L2. G1 accelerates further to compensate.

0.5 second
The control center shuts down generator G1 to prevent damage to it from excessive acceleration.

DON FOLEY; SOURCE: ELECTRIC POWER RESEARCH INSTITUTE

To transform our current infrastructure into this kind of self-healing smart grid, several technologies must be deployed and integrated. The first step is to build a processor into each switch, circuit breaker, transformer and bus bar—the huge conductors carrying electricity away from generators. Each transmission line should then be fitted with a processor that can communicate with the other processors, all of which

THE AUTHORS

MASSOUD AMIN and *PHILLIP F. SCHEWE* have extolled the virtues of a smart grid for years. Amin is professor of electrical and computer engineering at the University of Minnesota and is director of the school's Center for the Development of Technological Leadership. When he was with the Electric Power Research Institute in Palo Alto, Calif., he led the development of more than 20 advanced technologies and devised the foundations of a "self-healing grid," a term he coined. Schewe is chief science writer at the American Institute of Physics and edits the institute's weekly newsletter, *Physics News Update*.

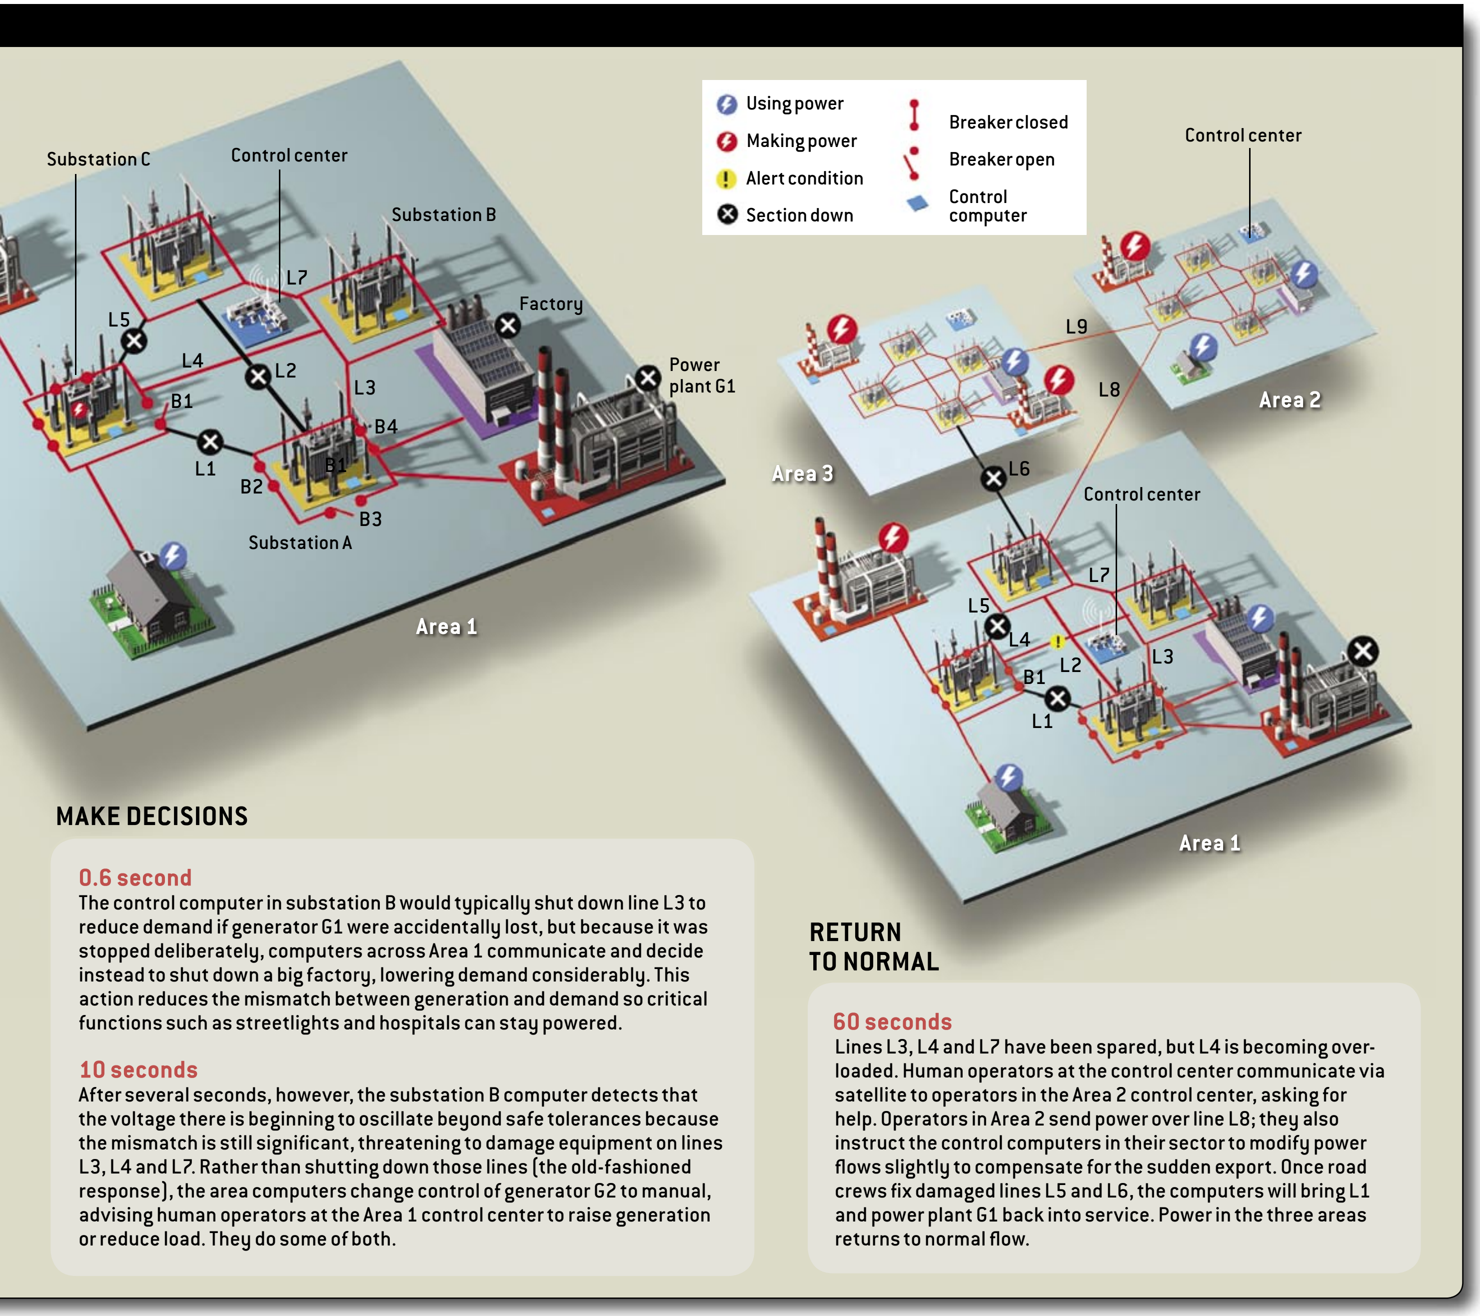

MAKE DECISIONS

0.6 second
The control computer in substation B would typically shut down line L3 to reduce demand if generator G1 were accidentally lost, but because it was stopped deliberately, computers across Area 1 communicate and decide instead to shut down a big factory, lowering demand considerably. This action reduces the mismatch between generation and demand so critical functions such as streetlights and hospitals can stay powered.

10 seconds
After several seconds, however, the substation B computer detects that the voltage there is beginning to oscillate beyond safe tolerances because the mismatch is still significant, threatening to damage equipment on lines L3, L4 and L7. Rather than shutting down those lines (the old-fashioned response), the area computers change control of generator G2 to manual, advising human operators at the Area 1 control center to raise generation or reduce load. They do some of both.

RETURN TO NORMAL

60 seconds
Lines L3, L4 and L7 have been spared, but L4 is becoming overloaded. Human operators at the control center communicate via satellite to operators in the Area 2 control center, asking for help. Operators in Area 2 send power over line L8; they also instruct the control computers in their sector to modify power flows slightly to compensate for the sudden export. Once road crews fix damaged lines L5 and L6, the computers will bring L1 and power plant G1 back into service. Power in the three areas returns to normal flow.

would track the activity of their particular piece of the puzzle by monitoring sensors built into their systems.

Once each piece of equipment is being monitored, the millions of electromechanical switches currently in use should be replaced with solid-state, power-electronic circuits, which themselves must be beefed up to handle the highest transmission voltages: 345 kilovolts and beyond. This upgrade from analog to digital devices will allow the entire network to be digitally controlled, the only way real-time self-monitoring and self-healing can be carried out.

A complete transition also requires digitization of the small, low-voltage distribution lines that feed each home and business. A key element is to replace the decades-old power meter, which relies on turning gears, with a digital meter that can not only track the current going into a building but also track current sent back out. This will allow utilities to much better assess how much power and reactive power is flowing from independent producers back into the grid. It will also allow a utility to sense very local disturbances, which can provide an earlier warning of problems that may be mounting, thereby improving look-ahead simulation. And it will allow utilities to offer customers hour-by-hour rates, including incentives to run appliances and machines during off-peak times that might vary day to day, reducing demand spikes that can destabilize a grid. Unlike a meter, this digital energy portal would allow network intelligence to flow back and forth, with consumers responding to variations in pricing. The portal is a tool for moving beyond the commodity model of electricity delivery into a new era of energy services as diverse as those in today's dynamic telecommunications market.

The EPRI project to design a prototype smart grid, called the Complex Interactive Networks/Systems Initiative, was conducted from 1998 to 2002 and involved six university research consortia, two power companies and the U.S. Department of Defense. It kicked off several subsequent, ongoing efforts at the U.S. Department of Energy, the National Science Foundation, the DOD and EPRI itself to develop a central nervous system for the power grid. Collectively, the work shows that the grid can be operated close to the limit of stability, as long as operators constantly have detailed knowledge of what is happening everywhere. An operator would monitor how the system is changing, as well as how the weather is affecting it, and have a solid sense of how to best maintain a second-by-second balance between load (demand) and generation.

As an example, one aspect of the EPRI's Intelligrid program is to give operators greater ability to foresee large-scale instabilities. Current SCADA systems have a 30-second delay or more in assessing the isolated bits of system behavior that they can detect—analogous to flying a plane by looking into a foggy rearview mirror instead of the clear airspace ahead. At EPRI, the Fast Simulation and Modeling project is developing faster-than-real-time, look-ahead simulations to anticipate problems—analogous to a master chess player evaluating his or her options several moves ahead. This kind of grid self-modeling, or self-consciousness, would avoid disturbances by performing what-if analyses. It would also help a grid self-repair—adapt to new conditions after an outage, or an attack, the way a fighter plane reconfigures its systems to stay aloft even after being damaged.

Who Should Pay

TECHNOLOGICALLY, the self-healing smart grid is no longer a distant dream. Finding the money to build it, however, is another matter.

THE HUMAN FACTOR

If a local blackout begins to escalate beyond a smart grid's ability to automatically keep it in check, human operators in regional control rooms could attempt to cut off the chain reaction. To do so, they would need complete and up-to-the-second information about the network, consistent computer protocols, predetermined response procedures and solid training. Each of these prerequisites was lacking when the massive 2003 U.S. blackout began to snowball, as dialogue during the first minutes of the event shows (*portions appear below*). The recorded conversations, published by the North American Electric Reliability Council, took place between reliability controllers in neighboring regions who were trying to help one another balance power flows that were heading out of control.

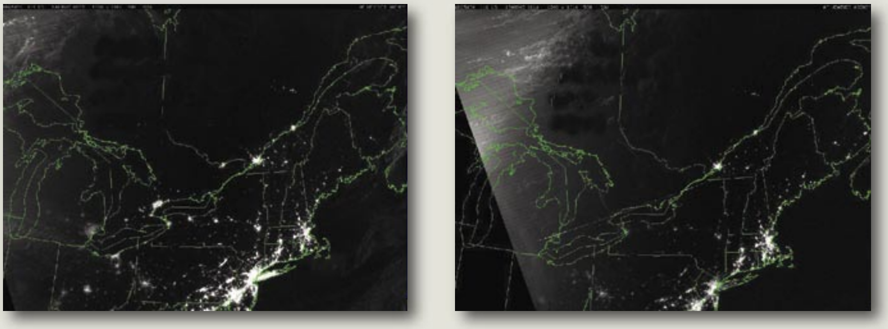

SATELLITE IMAGES show the Northeast the night before (*left*) and after the 2003 blackout (*right*).

NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION

POOR OPERATOR TRAINING, LACK OF REAL-TIME DATA

AEP operator: “What do you have on the Sammis-Star [line]?”
PJM operator: “I'm sorry? Sammis-Star, okay, I'm showing 960 on it and it's highlighted in blue.... Tell me what that means on your machine.”
AEP: “Blue? Normal.... I mean—that's what's on it?”
PJM: “960, that's what it says.”
AEP: “That circuit just tripped. South Canton-Star.”
PJM: “Did it?”
AEP: “It tripped and reclosed ...”

INCONSISTENT COMPUTER PROTOCOLS, LACK OF CONTINGENCY PLANS

PJM: “I'm still seeing flow on both those lines. Am I looking at state-estimated data?”
AEP: “Probably.”
PJM: “Yeah, it's behind, okay. You're able to see raw data?”
AEP: “Yeah; it's open. South Canton-Star is open.... We have more trouble ... more things are tripping. East Lima and New Liberty tripped out. Look at that.... Oh, my gosh, I'm in deep ...”
PJM: “You and me both, brother. What are we going to do?”

INCOMPLETE INFORMATION, LACK OF REAL-TIME DATA

PJM: “... it looks like they lost South Canton-Star 345 line. I was wondering if you could verify flows on the Sammis-Star line for me.”
MISO operator: “Well, let's see what I've got. I know that First Energy lost their Juniper line, too.”
PJM: “Did they?”
MISO: “And they recently have got that under control here.”
PJM: “And when did that trip? That might have ...”
MISO: “I don't know yet....”
PJM: “And right now I am seeing AEP systems saying Sammis to Star is at 1378....”
MISO: “Let me see. I have got to try and find it here, if it is possible.... I see South Canton-Star is open, but now we are getting data of 1199, and I am wondering if it just came after.”
PJM: “Maybe it did.”

PJM: Pennsylvania–New Jersey–Maryland
AEP: American Electric Power
MISO: Midwest Independent System operator

RESEARCHERS at the Pacific Northwest National Laboratory sit in a simulated regional control center and test prototype software that would provide human operators with real-time grid information, necessary to stop a nascent blackout before it expands.

PACIFIC NORTHWEST NATIONAL LABORATORY

The grid would be costly, though not prohibitively so given historic investments. EPRI estimates that testing and installation across the entire U.S. transmission and distribution system could run $13 billion a year for 10 years—65 percent more than the industry is currently investing annually. Other studies predict $10 billion a year for a decade or more. Money will also have to be spent to train human operators. The costs sound high, but estimates peg the economic loss from all U.S. outages at $70 to $120 billion a year. Although a big blackout occurs about once a decade, on any given day 500,000 U.S. customers are without power for two hours or more.

Unfortunately, research and development funding in the electric utility industry is at an all-time low, the lowest of any major industrial sector except for pulp and paper. Funding is a huge challenge because utilities must meet competing demands from customers and regulators while being responsive to their stakeholders, who tend to limit investments to short-term returns.

Other factors must be considered: What terrorism threat level is the industry responsible for and what should government cover? If rate increases are not palatable, then how will a utility be allowed to raise money? Improving the energy infrastructure requires long-term commitments from patient investors, and all pertinent public and private sectors must work together.

Government may be recognizing the need for action. The White House Office of Science and Technology Policy and the U.S. Department of Homeland Security recently declared a "self-healing infrastructure" as one of three strategic thrusts in their National Plan for R&D in Support of Critical Infrastructure Protection. National oversight may well be needed, because the current absence of coordinated decision making is a major obstacle. States' rights and state-level public utility commission regulations essentially kill the motivation of any utility or utility group to lead a nationwide effort. Unless collaboration can be created across all states, the forced nationalization of the industry is the only way to achieve a smart grid.

At stake is whether the country's critical infrastructures can continue to function reliably and securely. At the very least, a self-healing transmission system would minimize the impact of any kind of terrorist attempt to "take out" the power grid. Blackouts can be avoided or minimized, sabotage can be contained, outages can be reduced, and electricity can be delivered to everyone more efficiently.

Had a self-healing smart grid been in place when Ohio's local line failed in August 2003, events might have unfolded very differently. Fault anticipators located at one end of the sagging transmission line would have detected abnormal signals and redirected the power flowing through and around the line to isolate the disturbance several hours before the line would have failed. Look-ahead simulators would have identified the line as having a higher-than-normal probability of failure, and self-conscious software along the grid and in control centers would have run failure scenarios to determine the ideal corrective response. Operators would have approved and implemented the recommended changes. If the line somehow failed later anyway, the sensor network would have detected the voltage fluctuation and communicated it to processors at nearby substations. The processors would have rerouted power through other parts of the grid. The most a customer in the wider area would have seen would have been a brief flicker of the lights. Many would not have been aware of any problem at all.

MORE TO EXPLORE

Technical Analysis of the August 14, 2003, Blackout: What Happened, Why, and What Did We Learn? North American Electric Reliability Council, 2004. Available at **www.nerc.com/~filez/blackout.html**

Toward a Smart Grid: Power Delivery for the 21st Century. Massoud Amin and Bruce F. Wollenberg in *IEEE Power and Energy Magazine,* Vol. 3, No. 5, pages 34–41; September/October 2005.

The Grid: A Journey through the Heart of Our Electrified World. Phillip F. Schewe. Joseph Henry Press, 2007.