Agents offer a way to buy and sell in complex systems and give stakeholders a more realistic look at that system.

Massoud Amin and Dan Ballard

Defining New Markets for Intelligent Agents

Agents are no strangers to the electronic marketplace. Indeed, the Internet version of an agent, the omnipresent "bot," may be one reason IDC expects e-commerce to top $1 trillion by 2003. Most agent applications are fairly straightforward: access a Web site, fetch material—in short, perform a simple fixed mission. Others do more personalized tasks such as filtering e-mail or updating legacy systems.

From a programming viewpoint, agents are simply active objects that have been defined to simulate parts of a model. Agent-based modeling and simulation then become a natural extension of the object-oriented paradigm. Simulations of events that involve these kinds of agents (known as actors or demons) have assisted human decision making for decades in batch manufacturing, transportation, and logistics, for example.

But work in complex adaptive systems (CAS) may be defining a new kind of agent—one that can actually evolve over time in response to its environment. This is no small undertaking. Each agent changes in a way that adapts to its environment—even while its environment is itself changing because of external forces and changes in other agents. This puts a new spin on the agent community, or agency. Agents live much the way that the human population does, finding ways to combat the stress from their environment and either thriving and evolving (including reproducing), or shrinking and dying. The "How Adaptive Agents Work" sidebar tells more about this process.

Inside

How Adaptive Agents Work

What Is an "Intelligent" Agent?

Resources

This means we may finally be able to understand the real dynamics in a business enterprise, financial market, and even the economy itself—all of which have eluded modeling with traditional methods because the mathematics just can't handle complexity on that scale. The beginnings of these adaptive systems are already evident in more advanced agents, which can do simple negotiations on a user's behalf to secure goods and services in an auction, for example.

The challenge now is to see how agents bargain and learn in a more complex environment. The Electric Power Research Institute (EPRI), for example, has funded research into agent-based auctioning as a way to address the fierce competition for resources. As electric power marketers become available, wholesale electric customers are learning to shop around for the best suppliers. Like agents that represent individual human users, the agents bargaining on behalf of these suppliers and wholesalers decide things like how much to buy, which agent to buy from, how much to pay, and how to manage the exchange of power and money. There is also concern that the entire market not be harmed by the sale. Thus, looking at how agents complete their transactions and how they learn from them provides insight into the dynamics of electrical power supply and demand.

How Adaptive Agents Work

An adaptive agent has a range of reasoning capabilities. It can create new patterns, as opposed to sorting through predetermined patterns to find the optimal response. Adaptive agents can be passive—responding to environmental changes without attempting to change the environment—or active—exerting some influence on their environment to improve their ability to adapt. Thus, an active adaptive agent essentially conducts experiments and learns from them.

Individual agents must be able to respond to environmental conditions and to other agents in a way that enhances their survival or meets other goals. To learn a strategy that increases its "fitness," the agent must gather and store enough information to adequately forecast and deal with changes that occur within a single generation. The population then adapts through the diversity of its individuals. Some individuals will always survive, and their individual actions will benefit the population's goals. Thus, the population evolves over many generations, surviving as a recognizable organization.

The agent community evolves as individual agents change their parameters. These individual changes, in turn, cause other agents to change their actions and decisions. Agents in effect "tinker" with the system's rules and structure. Those subjected to increased stress (resource shortages, environmental pressures, financial losses) increase their level of tinkering until some develop strategies that relieve the stress. These agents succeed; they grow, reproduce, and increase their profits. Agents that fail to relieve their stress shrink, die, are replaced, or are bought out.

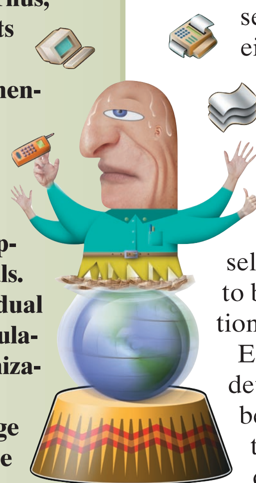

AGENT-BASED AUCTIONING

In an auction setting, intelligent agents must be able to bargain. There is essentially no middleman acting as auctioneer. Thus, the agents must represent the goods and services to be traded, understand the goals the stakeholder wants to achieve from the negotiation, and use a negotiation strategy that is at least as good as what a "qualified person" would use in the same situation. The market then becomes a mechanism that determines how stakeholders can exchange bids, offers, and other messages to implement a trade.

A PROTOTYPE AUCTIONING AGENCY

In one application of agent-based auctioning, Reticular Systems Inc. and Alternative Energy Systems Consulting Inc. explored the possibilities and ultimate limits of competitively and cooperatively scheduling electric power distribution. The project, sponsored by EPRI, found that intelligent agents could dynamically form their own marketplaces to meet their individual needs—without relying on a central authority.

Part of the project involved building a demonstration agent-based power market, or agency, using Reticular's AgentBuilder integrated toolkit. The agency consists of 10 buyer/seller agents (BSAs) and an experiment control agent.

Buyer/seller agents

As their name implies, BSAs buy and sell electric power, representing either power producers (generators) or power consumers (loads). BSAs are the agency's primary participants, doing all the buying and selling. A BSA can take on several roles. At any time, the BSA can be either a buyer or seller. It can also conduct an auction to buy or sell or participate in an auction as a bidder.

Each BSA exhibits unique behavior, determined by its own economic and behavioral model (programmed into the agent). Each agent's behavior is dictated by three components: the strategy used to buy or sell, the market supply and demand characteristics, and the power amount required.

Buying/selling strategy. Each stakeholder has different business conditions and considerations, which translate into the buying and selling strategies for the stakeholder's agents. If a stakeholder suffers a generator failure, for example, its agent will need to aggressively bid for electricity. On the other hand, if all its equipment is online and working, its agent should be quite sensitive to price. Figure 1 shows some possible buying and selling strategies.

Market supply and demand. Each BSA follows the stakeholder's supply or demand requirements and operating policies for satisfying those requirements. The supply component is the desired selling price for a quantity of electricity to be delivered at a given date and time. The demand component is the required price for buying a given quantity of electricity at a given date and time. The operating policies dictate how much each agent is willing to spend or charge for power. Supply and demand determine the amount of power each agent will be willing to buy or charge. As prices go up, agents will be willing to consume less and supply more. As prices fall, agents will want to consume more and sell less.

Operational requirements. Each agent represents its stakeholder's unique power requirements, which vary over time. Although market price dictates the agent's buying/selling behavior to some extent, agents must still ensure that they buy and sell enough power to sustain operations.

Experiment control agent

The experiment control agent provides a GUI, which the operator uses to control the simulated auction session and view the results (see Figure 2). The operator sets up the BSA agents and controls the auction's overall timing. The experiment control agent does not participate in the buying and selling of electric power. Its primary purpose is to allow an auction to be conducted as a demonstration rather than in real time (which would take one week). The experiment control agent sends simulated times to the buyers and sellers and collects auction results for logging and display.

Figure 1. Possible buying and selling strategies.

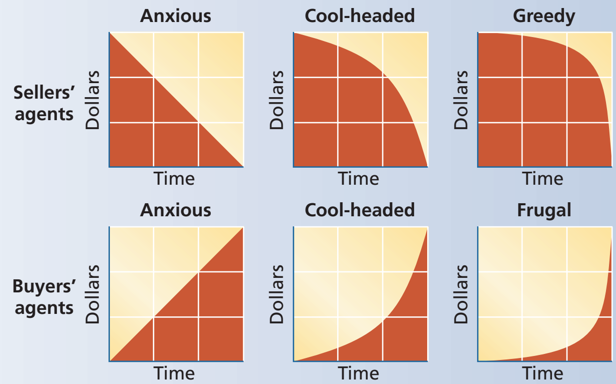

As market conditions vary, agents will use different strategies. For example, when the market demand is high and supply is low, an agent acting as a buyer will be "anxious." Each strategy is represented by a curve that corresponds to transaction time and cost of the power.

Figure 2. How operators track agents in an electrical power auction.

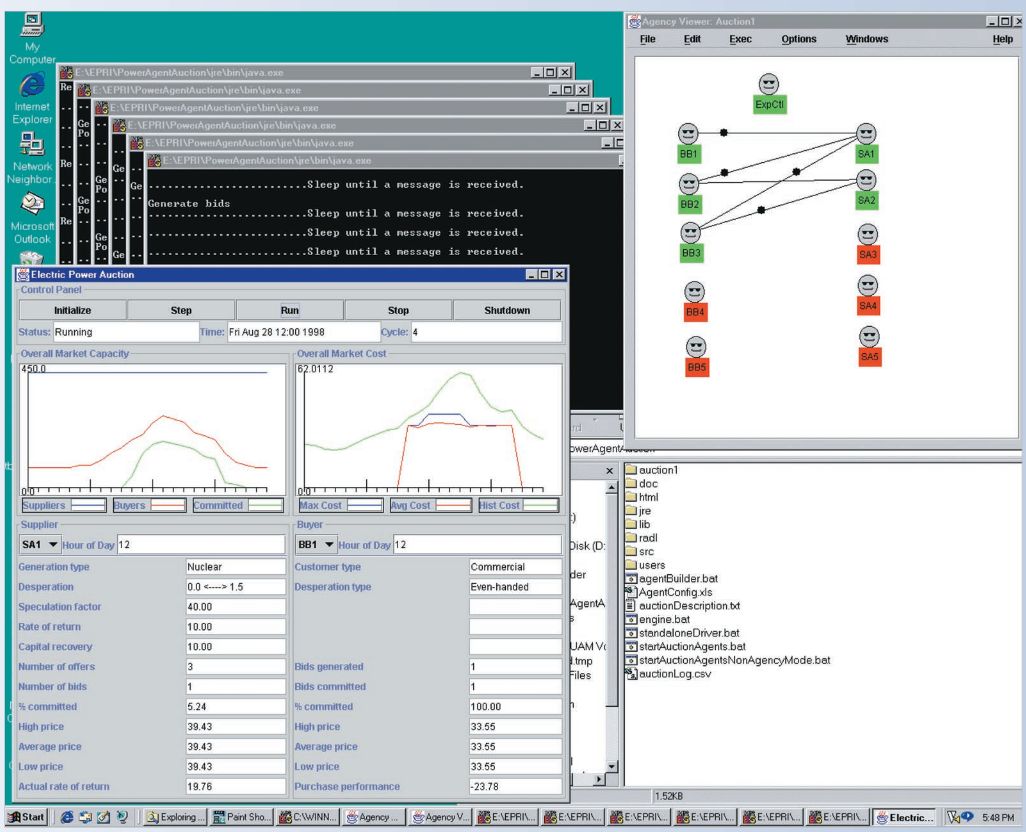

The black panels show the activity of each agent as it participates in the auction. The AgencyViewer (upper right) shows the flow of messages between agents as they bargain for power. The control panel (lower left) shows what is happening in the overall marketplace as the individual agents buy and sell electricity.

What Is an "Intelligent" Agent?

Agents have evolved in a variety of disciplines—artificial intelligence, robotics, information retrieval, and so on—making it hard to get consensus on what they are. Most researchers agree, however, that a truly intelligent agent has these attributes:

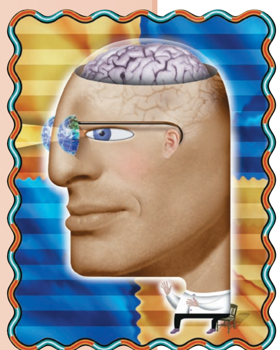

- ***Reactivity.*** **It can sense the environment and act responsively.**
- ***Autonomy.*** **It does not need human intervention.**
- ***Social and collaborative behavior.*** **It can work with other agents or humans toward a common goal.**
- ***Inferential capability.*** **It proactively seeks to meet its goals, including analyzing its environment.**
- ***Temporal continuity.*** **Its identity and state persist over long periods.**
- ***Adaptivity.*** **It can learn and improve with experience.**

The more advanced agents may also have other attributes, such as mobility (it can migrate from one host platform to another, either by directing itself or following a preprogrammed schedule) and personality (manifesting some human qualities, such as cooperation for the greater good, caution, and even greed).

The demonstration market is necessarily on a much smaller scale than a production (real-world) auction system would be. In a production system, there would be no experiment to control, but some of the same functionality would be needed to monitor and control market activities. A facilitator agent would track all participating BSAs in a production system. Each BSA would communicate with the facilitator and indicate whether it wished to buy or sell electricity and whether it wanted to be a bidder or auctioneer. The facilitator thus acts as a broker or matchmaker, connecting buying agents with selling agents and auctioneers with bidders. In the demonstration market, instead of a facilitator agent, each BSA has a priori knowledge of the other agents and where they are.

The production system would also have a grid interface agent. This agent provides a way to conveniently simulate power supply and demand, which, in turn, dictates how the BSAs will behave. The grid interface agent would also collect data from all market sessions and display it through a GUI.

The AgencyViewer, part of Reticular's AgentBuilder toolkit (http://www.agentbuilder.com), lets users view communication messages among agents. As Figure 2 shows, lines are drawn between the icons to represent each communication message sent between agents. When the operator starts AgencyViewer, a window with 11 agent icons will appear: five bidder/buyers (BB1-BB5), five sellers/auctioneers (SA1-SA5), and the experiment control agent.

The auctioning process

Auctioning begins when the operator uses the experiment control agent to assign each agent a role. The agents then begin to buy and sell electricity in accordance with the economic needs and their individual buying and selling strategies, which the stakeholders define and set according to their business needs.

An agent initiates an auction by advertising that it is going to conduct an auction and specifying the auction time. All bidders interested in participating can then join the auction and submit bids. Each auctioning agent is in a sense conducting an auction that is selling/buying 24 items: power supply or demand for one particular hour in one particular day.

The auctioning agent offers a quantity of electricity at an asking price for delivery at some hour on some day. It also specifies the time of the next auction round to sell the power not sold in the current round and the price adjustment that will be applied to the remaining power in the next round.

Each auctioning agent also maintains a reserve price—the minimum/maximum price the agent will accept before refusing any further transactions at a lower/higher price. The reserve price is known only to the auctioning agent and is not available to bidding agents. An auctioning agent may decide to reintroduce this block of capacity back into the marketplace at some later time.

Agents that want to buy/sell power from/to the auctioning agent specify the quantity that they are willing to buy/sell at the specified price. The auctioneer then forms a contract with the bidder to deliver that power (subject to collision constraints). Agents are encouraged to buy during a particular round because they cannot be sure that they can get the power they require in a later round if other agents purchase it in this round.

Although bidder agents may wish to defer contracting with the auctioneer, they do so at the risk of being unable to acquire the amount of power they require. For the demonstration, the initial round is for delivery of power at the specified hour seven days in the future. The next round is for delivery of power at that same hour six days in the future. If bidding continues to the last round, this round is for the sale of power to be delivered on the next day. Since bidders don't know when the reserve price will be reached and the remaining power withdrawn from the market, they are encouraged to bid early in the round.

SCALING UP COMPLEXITY

The demonstration market shows that agents, acting independently, are an effective way to study the many issues that affect the power market as it struggles with adapting to changes caused by deregulation. But it cannot capture network behavior on a very large scale. To gain insights into something as large as the entire North American power grid, for example, we must view the entire grid as a complex adaptive system that can exhibit global change almost instantaneously from actions taken in only one part of it.

Viewing the grid in this way means coming up with a new kind of simulation model that looks at the grid from the bottom up—as a collection of interacting, adaptive agents. Using this model, EPRI hopes to make the grid self-healing—grid components could actually reconfigure to respond to material failures, threats, or other destabilizers.

Figure 3. Agent structure for power market simulation.

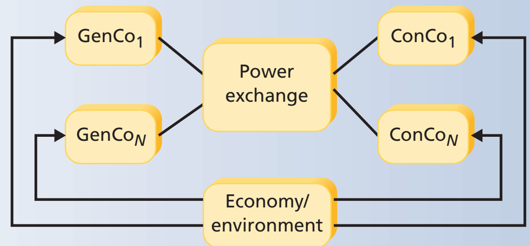

Each agent in the simulation functions independently. The agents interact with each other by sending and receiving messages, electric power, and money. Corporate agents—power generators (GenCo) and consumers (ConCo)—reflect their stakeholders' business objectives in their internal dynamics in the simulation. Noncorporate agents, such as the power exchange, residential loads, and the economy/environment have less complex internal dynamics.

SEPIA

Over the past few years, EPRI members funded and sponsored a research effort to develop this new model, and the first phase is complete. SEPIA (Simulator for Electrical Power Industry Agents) is a multiple adaptive agent model of the grid and of the industrial organizations that own parts of it or are connected to it. Developed by Honeywell Technology Center in conjunction with the University of Minnesota, SEPIA provides a way for EPRI member companies to conduct computational experiments for any kind of scenario. This helps them see the power market's true dynamics, which in turn gives them strategic insight into market share.

SEPIA is an object-oriented Windows application with plug-and-play agent architecture. Users can readily adapt simulations to a parallel computing environment, including multiprocessor PCs and PC networks. SEPIA agents are autonomous modules that encapsulate specific behaviors. They are implemented as independent ActiveX applications that intercommunicate by messages sent through the SEPIA agent bus. The messaging mechanism can handle a variety of communication needs: electric power transmission, information flows among corporate agents, and money transfers.

The user interface, which is modeled after the Windows GUI, lets users specify agents and agent relationships, modify agents, and guide and monitor the simulation. Users conduct simulations by defining scenarios through drag-and-drop operations on icons representing the agents, then interconnecting the agents, and pressing a "run" button. Simulation results are shown dynamically on graphs and reports. Simulators can modify agent policies and parameters dynamically as well.

Agent structure

The SEPIA project team has implemented the agent model, simulation engine, and GUI, and defined agent base classes: generation unit, transmission system, power exchange, transmission zone, load, and corporate. Corporate agents represent either a power-consuming company (ConCo) or a power-generating company (GenCo). Figure 3 presents a simple scenario with four agent types: power-generating companies (GenCos), power-consuming entities (ConCos), a power exchange, and the economy/environment.

Power-generating companies. The structure of a power generating company will give you some idea of how complex the internal behavior of a GenCo agent is. Each power-generating company has five departments: production, marketing, procurement (purchasing), finance, and strategy. Each department, in turn, has its own market-oriented schedule and budget: Procurement has a schedule and budget for buying fuel. Production has a schedule and budget for generating power. Marketing has a schedule and budget for selling power. Finance has a schedule and budget for accounts and available credit. On top of that, each department maintains an internal (read-protected) file. Procurement has a fuel inventory file. On top of this, each department maintains an internal file; for example, marketing has a file of contracts with the power exchange. The companies receive forecasts from the economy/environment agent and quote-board prices from the power exchange agent. On the basis of these and their internal files, the strategy department evaluates the four market-oriented schedules and budgets and provides direction to the specialized departments for their revision.

Power-consuming entities. Power-consuming entities can be industrial, commercial, or residential. The commercial

Resources

Web sites

http://www.botspot.com. All about agents, new and old. News, conferences, and forums.

http://agents.www.media.mit.edu/groups/agents/. MIT Media Laboratory's Software Agents Group. Projects, news, and conferences.

http://www.agent.org. Agent Society, not-for-profit agency to further the development and use of agent technology.

http://www-cia.mty.itesm.mx/~lgarrido/Repositories/IA/agents.html. Repository of articles, books, conferences, and news on software agents.

http://www.htc.honeywell.com/projects/sepia. SEPIA project site. Visitors can view a self-running demonstration of SEPIA.

http://www.epri.com. EPRI site describes the EPRI/DoD Complex Interactive Networks Initiative. Visitors can order technical reports: "Computer Simulation of Adaptive Agents for an Electric Power Auction," 1999 (TR-106975); "Simulator for Electric Power Industry Agents (SEPIA): Complex Adaptive Strategies," 1999 (TR-112816); "Prototype Intelligent Software Agents for Trading Electricity: Competitive/ Cooperative Power Scheduling in an Electronic Marketplace," 1999 (TR-113366); and "E-Commerce Applications and Issues for the Power Industry," 2000 (TR-114659).

http://www.petzinger.com/complexity.shtml. Annotated bibliography of about 30 books with metaphors and models using complexity theory for business problems.

Books, articles, and papers

Agent Technology: Foundations, Applications, and Markets, N.R. Jennings and M.J. Wooldridge, eds., Springer-Verlag, New York, 1998: Good introduction to ideas and basic technologies involving agents.

"AgentBuilder: An Integrated Toolkit for Constructing Software Agents:" 1999 white paper at www.agentbuilder.com provides a very readable tutorial on the mental models behind agents, and how to build them using AgentBuilder.

Amin, M., "National Infrastructures as Complex Interactive Networks," *Automation, Control, and Complexity: New Developments and Directions*, Samad and Weyrauch, eds., John Wiley & Sons, 2000.

Amin, M., "Complexity and the Deregulation of the Electric Power Industry," *Proc. Third Embracing Complexity (EC3) Conf.*, 1998, pp. 101-106.

Amin, M., A.M. Wildberger, and G. McCarthy, "Self-Healing Electric Power Grid as a Complex Adaptive System," *Proc. Fourth Joint Conf. Information Sciences (JCIS'98)*, Vol. I, 1998, pp. 295-301.

Wildberger, A.M., "Complex Adaptive Systems: Concepts and Power Industry Applications," *IEEE Control Systems*, Dec. 1997, pp. 77-88.

---

and residential entities are little more than stochastic loads that act according to the state of the economy/environment agent. The industrial entities, also called power-consuming companies, have internal departments and structures similar to those in the power-generating companies.

Power exchange. The power exchange is a special agent that has three departments: exchange management, brokerage, and finance. The basic functions of these departments are similar to those of a futures exchange. Thus, although the internal dynamics aren't as complex as those of GenCos and ConCos, the departments must track considerable activity, as Figure 4 shows.

The power-price quote board is basically a table indicating the highest bid and lowest offer for standardized quantities of power for delivery during each specified hour at each specified location. Other data such as opening price, daily high and low, open interest, and volume are also available to give traders some idea of the market's state and recent conditions.

Agent adaptation and learning

Agents in SEPIA are adaptive. To provide this capability, simulators associate general or customized adaptation modules (algorithms) with each agent. This means that a simulation can actually evolve. For example, as the competition increases, the simulation can actually produce an optimized pricing structure for a generating company in a competitive, dynamic environment.

The current version of SEPIA offers two reusable adaptation algorithms: a form of the Learning Classifier System (LCS) and a variation of the popular Q Learning algorithm. LCS uses rule representation. Agents "discover" things through a genetic algorithm and reinforced learning. Q Learning offers a formula that enables an agent to learn

Figure 4. Structure of the power exchange agent, showing agent functions (rectangles), internal data files (ellipses), and public records (diamonds).

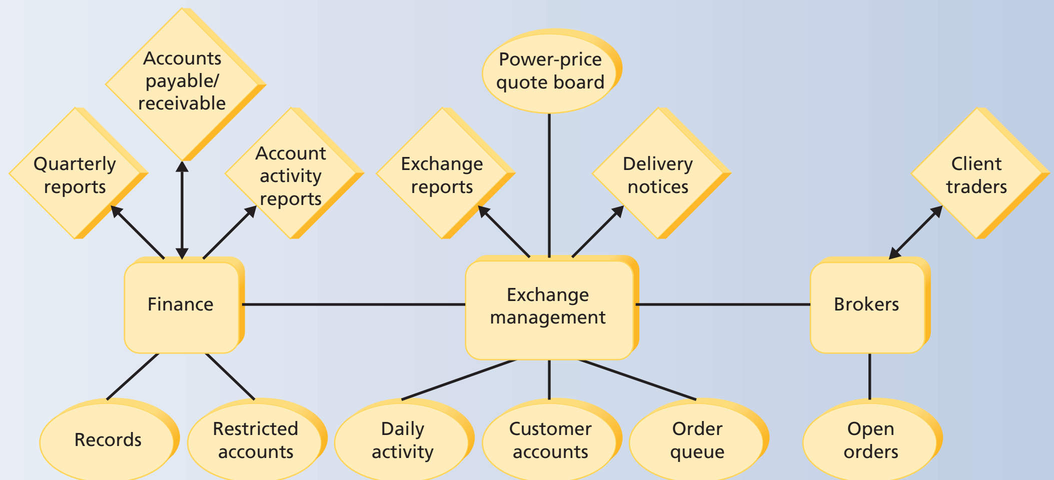

Exchange management maintains a power-price quote board, keeps track of orders placed by the brokers, executes orders when appropriate, and notifies brokers of executed orders. It also serves as a clearinghouse for the positions open when the contract expires, sending delivery notices to the appropriate counter parties.

to take an action when given an observable state. Lacking a world model to allow offline search, the Q learner uses the world as its own model. In SEPIA, the algorithm is implemented using a C++ class that accommodates any finite set of states and actions. Hence, it can be applied in many different agents or in different ways in the same agent.

The demonstration auction market proves that agents can cooperate and compete to dynamically create their own electrical marketplaces. SEPIA is essentially extending these results to a model of the entire power industry. Already, SEPIA experiments have made it possible to develop more sophisticated business scenarios. The next step is to create more realistic models—of both physical components and trade aspects—to gain more meaningful insights into market dynamics. ■

Massoud Amin *is manager of mathematics and information science at the Electrical Power Research Institute. Contact him at mamin@epri.com.*

Dan Ballard *is president and chief technical officer of Reticular Systems Inc. Contact him at ballard@reticular.com.*

An earlier version of this article was presented at the Infrastructure Protection and Emergency Management (IPEM) Symposium held as part of the 2000 Advanced Simulation Technologies Conference. The Society for Computer Simulations International sponsors ASTC.