

Agents of Cooperation

Orchestrating the actions of mobile snippets of smart software

By IVARS PETERSON

No traffic jams. No parking hassles. No jostling crowds. No inattentive, harried sales clerks.

Instead, you get expert help in choosing the perfect gift and finding it at the lowest possible price—all provided in the comfort of your own home.

That's one vision of shopping on the World Wide Web. Indeed, you can already take advantage of services that scour the Web to give you a list of on-line vendors and the prices they charge for a specific product, whether a video game, book, or Beanie Baby.

These services, sometimes called shopbots, automate comparison shopping. They represent a recent commercial application of research on a broad class of computer programs known as intelligent agents.


A shopbot program typically resides at a Web site. It sends out inquiries on behalf of a customer to other Web sites, collects the relevant information, and compiles a report tailored to the customer's requirements.

Other software agents automatically retrieve airline flight information, filter electronic mail to highlight messages important to the user, and alert investors to significant changes in selected stock prices.

It's like having a digital secretary to take care of some of your needs and an electronic adviser to help you make decisions.

There's much more ahead if projects now at the research stage come to fruition.

Vinton G. Cerf of MCI WorldCom predicts that, within 25 years, businesses and consumers will routinely use mobile intelligent agents to perform a wide variety of



tasks, often without the users being aware of the agents' activities. These pieces of software would roam the Internet to perform their duties, jumping from computer to computer.

Imagine a seething swarm of billions of such agents, finding and processing information, passing it on to people or other agents, negotiating with each other, and coordinating their activities. They could even alter their behavior as they learn and adapt to the ever changing landscape of cyberspace.

With access to such information resources, financial programs could tell consumers when interest rates have fallen enough to warrant refinancing a home loan. Huge, complex telecommunications networks could monitor, repair, and maintain themselves, becoming largely independent of human supervision.

Getting there, however, won't be easy. It may not even be desirable. Nonetheless, many research groups throughout the world are now exploring the possibilities.

One pioneer in the development of software agents is Pattie Maes of the Media Laboratory at the Massachusetts Institute of Technology.

In the early 1990s, she and her coworkers designed an agent that could learn to sort electronic mail according to priority by analyzing a user's reading habits. It could also figure out how to fulfill common requests, automatically replying when authorized to do so by the user.

Another Maes project, called Firefly, produced an agent that recommended movies, books, and music by quizzing visitors about their personal preferences, then comparing the answers to those of other people with similar tastes.

Allowing agents to act on their own is a key requirement. Maes and her colleagues define an autonomous agent as a "computational system" that can inhabit a complex, constantly changing environment, sense what is going on, and act independently to accomplish a specified set of tasks or achieve certain goals.

The underlying software technology is an offshoot of research in artificial intelligence. A software agent often makes inferences on the basis of a set of rules specifying its actions in a variety of situations.

Nearly a decade ago, Japanese researchers investigated putting together chunks of software that could work together as an autonomous agent. One product of that research was a talking head on a computer screen that could smile, frown, or look perplexed—using the appropriate facial expression as it replied to a spoken question. That effort was viewed as a step toward creating a "computer buddy" who could chat with its owner, joke around, or

offer advice on a variety of matters (SN: 10/9/93, p. 229).

It also provided a way to give software agents in a computer environment lifelike representations—or characters—that people would find meaningful or entertaining (SN: 12/19&26/92, p. 440).

In recent years, attention has focused on allowing autonomous agents to learn and modify their own behavior. Such capabilities would make them more versatile and robust, enhancing their value, many researchers contend.

One way to do that is by incorporating artificial neural networks. Each new piece of information that an agent encounters alters its programmed web of linked nodes, making some connections stronger and others weaker. Those changes, in turn, affect the way it processes subsequent data.

Researchers have used neural networks successfully in a wide variety of computer programs, including one that taught itself winning strategies in the game of backgammon by playing thousands of games and learning from the results (SN: 8/2/97, p. 76). Intelligent agents increasingly incorporate that kind of technology.

One capability agents may need to learn early on is how to negotiate.

In booking a hotel room or managing your stock portfolio, your agent may very well run into another agent at the other end of the line operating on behalf of the other party involved in the transaction.

A much higher level of sophistication comes with agents that cooperate, says Katia Sycara of Carnegie Mellon University in Pittsburgh. If their goals aren't compatible, the agents can even antagonize one another.

Sycara and her collaborators have been developing a system that acts as a matchmaker or broker between different types of agents. Called RETSINA, it coordinates intelligent agents that gather, filter, and integrate information on the Internet.

One piece of the work addresses the problem of integrating information coming from diverse sources. In addition, the researchers have developed a language called LARKS that allows agents to communicate with each other—a prerequisite to cooperation and negotiation.

Coordinating the activities of a collection of agents spread across a computer network offers several advantages, says Mark Klein of the Center for Coordination Science at MIT's Sloan School of Management. Efficient coordination would simplify the acquisition of information, improve performance, and make the system more tolerant of component failure, he contends.

Researchers are still striving to find coordination mechanisms that are efficient, reliable, and work with many agents, Klein says.

So far, agents typically have been confined to a single computer or a limited network of computers all using the same operating system. Installed at a host computer, they perform their duties by sending messages over the Internet and ordering the transfer of information, as needed, from place to place.

Now, researchers are designing agents that can travel around networks of dissimilar computers. Whenever necessary, such a mobile agent sends itself to another computer, where it continues its assigned task.

Mobile agents have the advantages of reducing communication requirements and operating more flexibly than current agents, Klein says.

A data-processing agent, for example, could move to a computer where the data it needs are stored, sift through the information, and come back with a compact result. That saves having to move the data from computer to computer.

Mobile agents offer an attractive answer for managing commercial telecommunications networks, says Patricia Morreale of the Stevens Institute of Technology in Hoboken, N.J.

Instead of merely sending messages identifying a trouble spot, for example, a network maintenance agent could travel to the computer where the problem appears most severe to get a firsthand look at the situation. It could then respond, perhaps by deploying other specialized agents to do its bidding.

Mobile agents can react much more quickly than their static counterparts in situations where instantaneous action is often required, Morreale remarks. Moreover, because a mobile agent can focus on the trouble spot, the system is much less likely to be swamped by extraneous alarms and messages when an emergency occurs.

Another possibility, say researchers at the Hewlett-Packard Laboratory in Bristol, England, is the use of a swarm of navigation agents to wander telephone networks to find speedy paths for telephone calls from one switching center to another. By leaving a kind of digital "scent" trail, agents following the shortest, quickest paths at any given moment could signal the best available routes for telephone calls.

To enhance their capabilities further, autonomous, mobile agents could be allowed to evolve and develop new behaviors to cope with new situations.

Some researchers, however, are nervous about what might happen when mobile, autonomous intelligent agents begin working together in an open computer network like the Internet, and individual agents evolve into new forms. Agents could become more akin to artificial life than mere pieces of software.

In 1991, Thomas S. Ray, now at the Uni-

versity of Oklahoma in Norman, created an electronic ecosystem he named *Tierra* (SN: 8/10/91, p. 88). The organisms that populated his world consisted of strings of computer instructions, expressed as ones and zeros. Various rules specified how these strings interact with each other, depending on their places in a computer's memory.

By allowing mutations that modify instruction strings and other types of interactions, Ray found that he could obtain all sorts of surprising behavior, including rampant propagation and diversification among *Tierra's* digital denizens.

Intelligent agents have far greater capabilities than Ray's simple automatons though they don't yet mutate and evolve in a Darwinian sense. What might happen in an information-driven, agent-mediated economy is one of the main issues now being studied at the recently formed Institute for Advanced Commerce at the IBM Thomas J. Watson Research Center in Yorktown Heights, N.Y.

Computer simulations of bustling digital marketplaces populated by various types of agents provide insights into the kinds of collective behavior that could emerge.

In one scenario, agents allow vendors to change prices and product offerings in response to demand and the competition much faster than a person ever could. Simulations show that such an economy quickly degenerates into vicious price wars and market crashes, while customers interested in the less popular products are poorly served.

By studying the conditions that lead to undesirable collective behavior, IBM researchers hope to understand better how to prevent these phenomena from occurring in the real world as agents race onto the Internet.

"These insights can then help us design rules of behavior and interaction that we can incorporate into agents and agent markets," says IBM's Jeffrey O. Kephart.

It may be useful, for example, to make agents more aware of the likely outcomes of their decisions and so to instill some common sense. That remains a tough problem for artificial intelligence researchers, however.

Compared with the intelligent agents being studied at IBM, today's shopbots are decidedly primitive. Extremely myopic, they typically search a limited number of Web sites, sometimes only those of vendors who have paid for the privilege of being included.

You don't get advice on the ideal gift. There's no guarantee that the final list really includes the lowest available price.

That may change very quickly, however. Each week brings new agents onto the scene as one startup company after another tests its wares on the Internet. The fittest survive—at least for a while. □