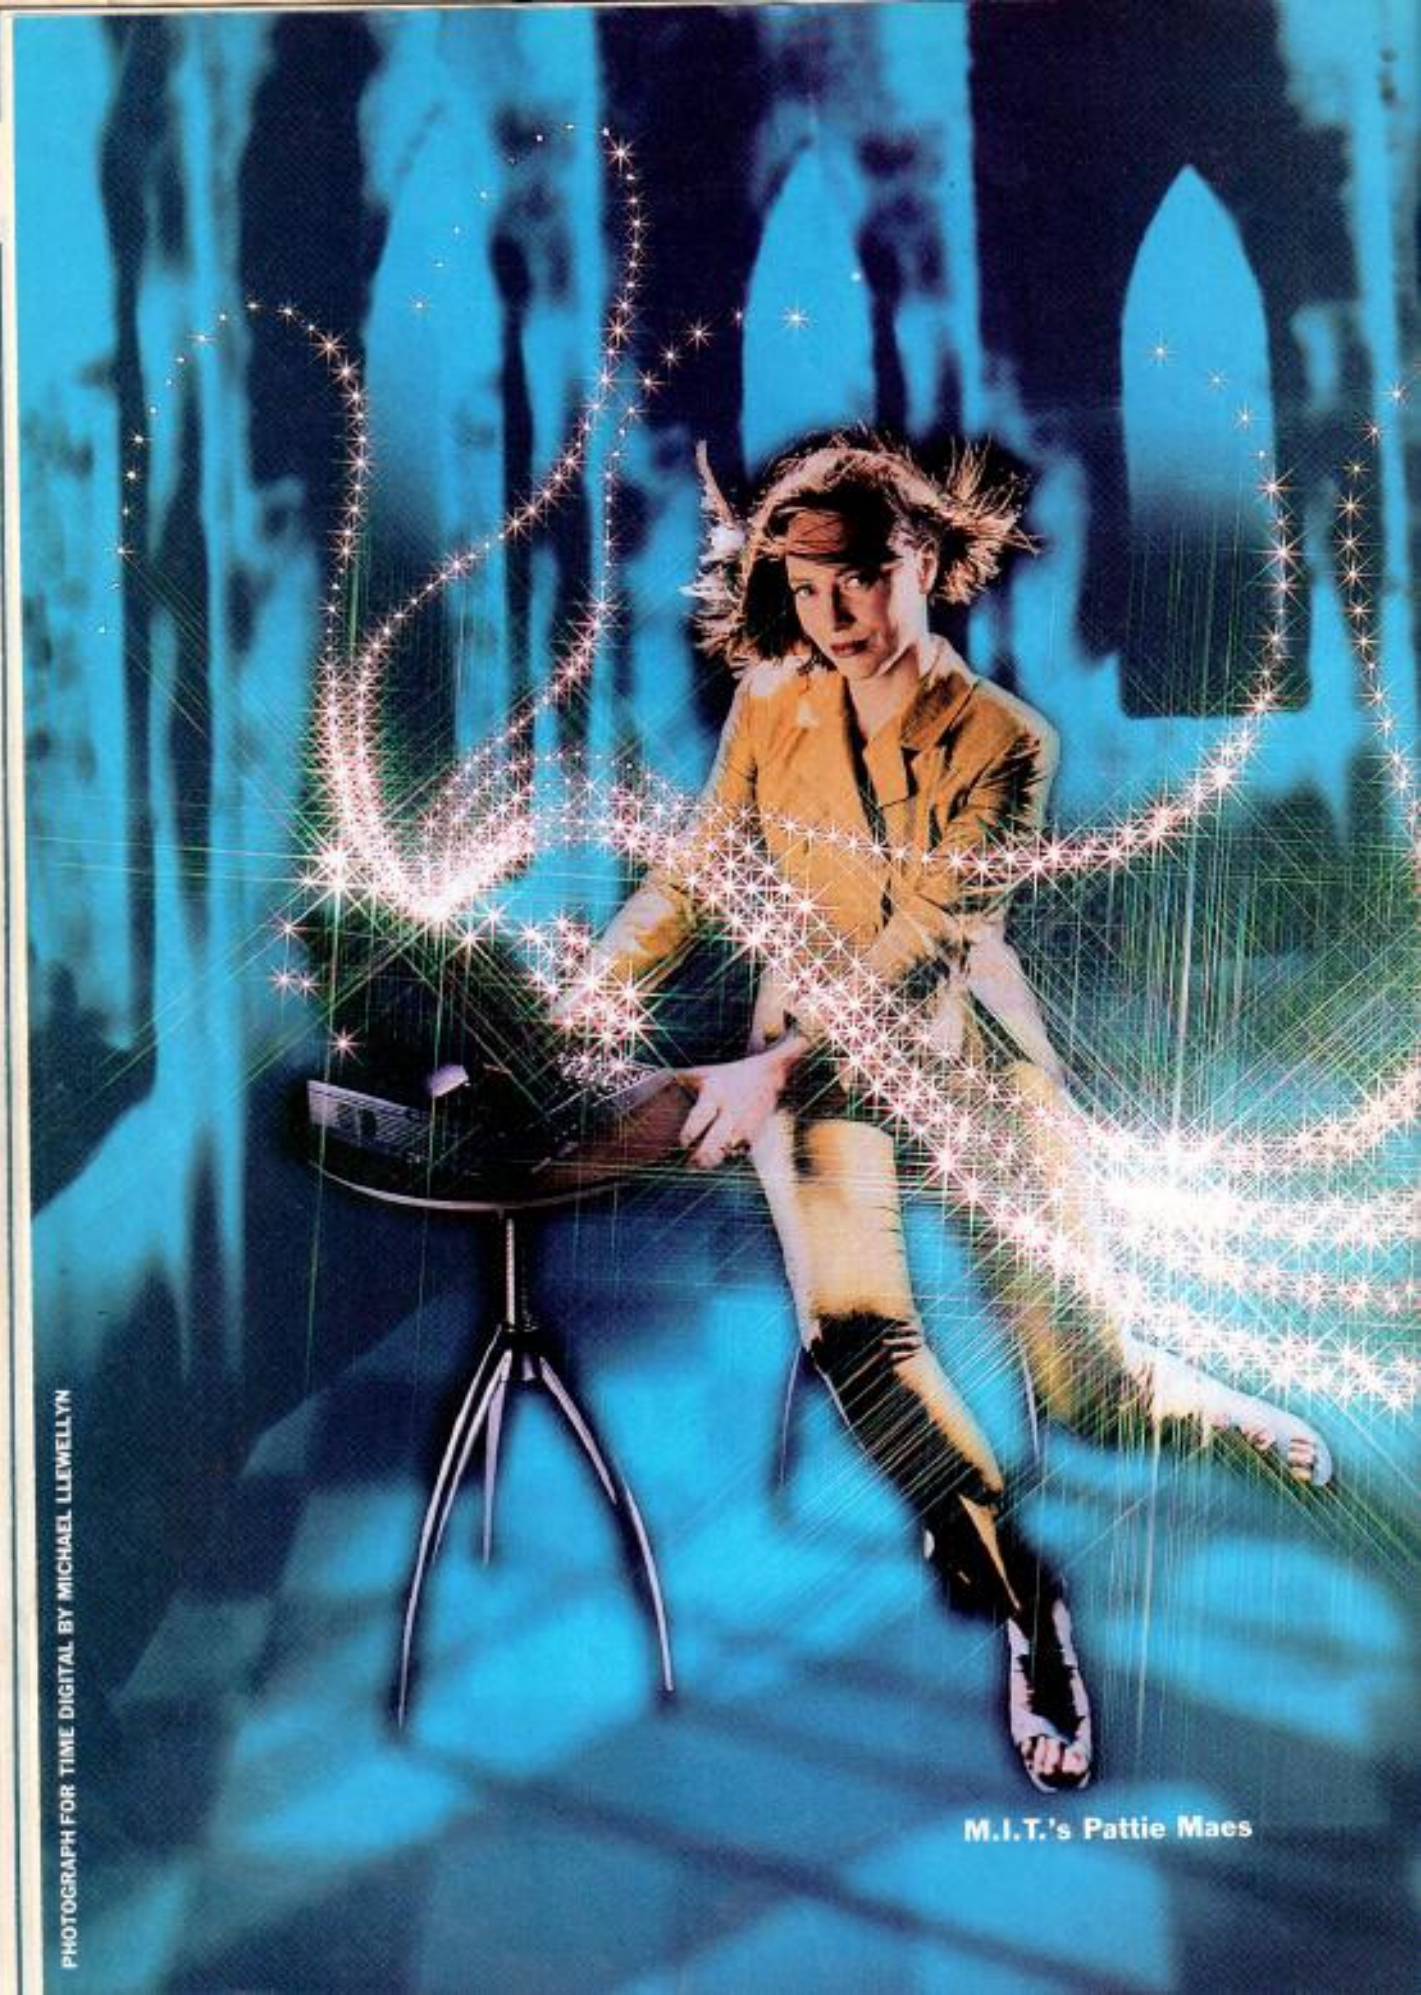

PHOTOGRAPH FOR TIME DIGITAL BY MICHAEL LLEWELLYN

M.I.T.'s Pattie Maes

Intelligent agents are changing cyberspace for good. The key ingredients: new technology, clever programming and a bit of

Smart Magic

By JULIAN DIBBELL

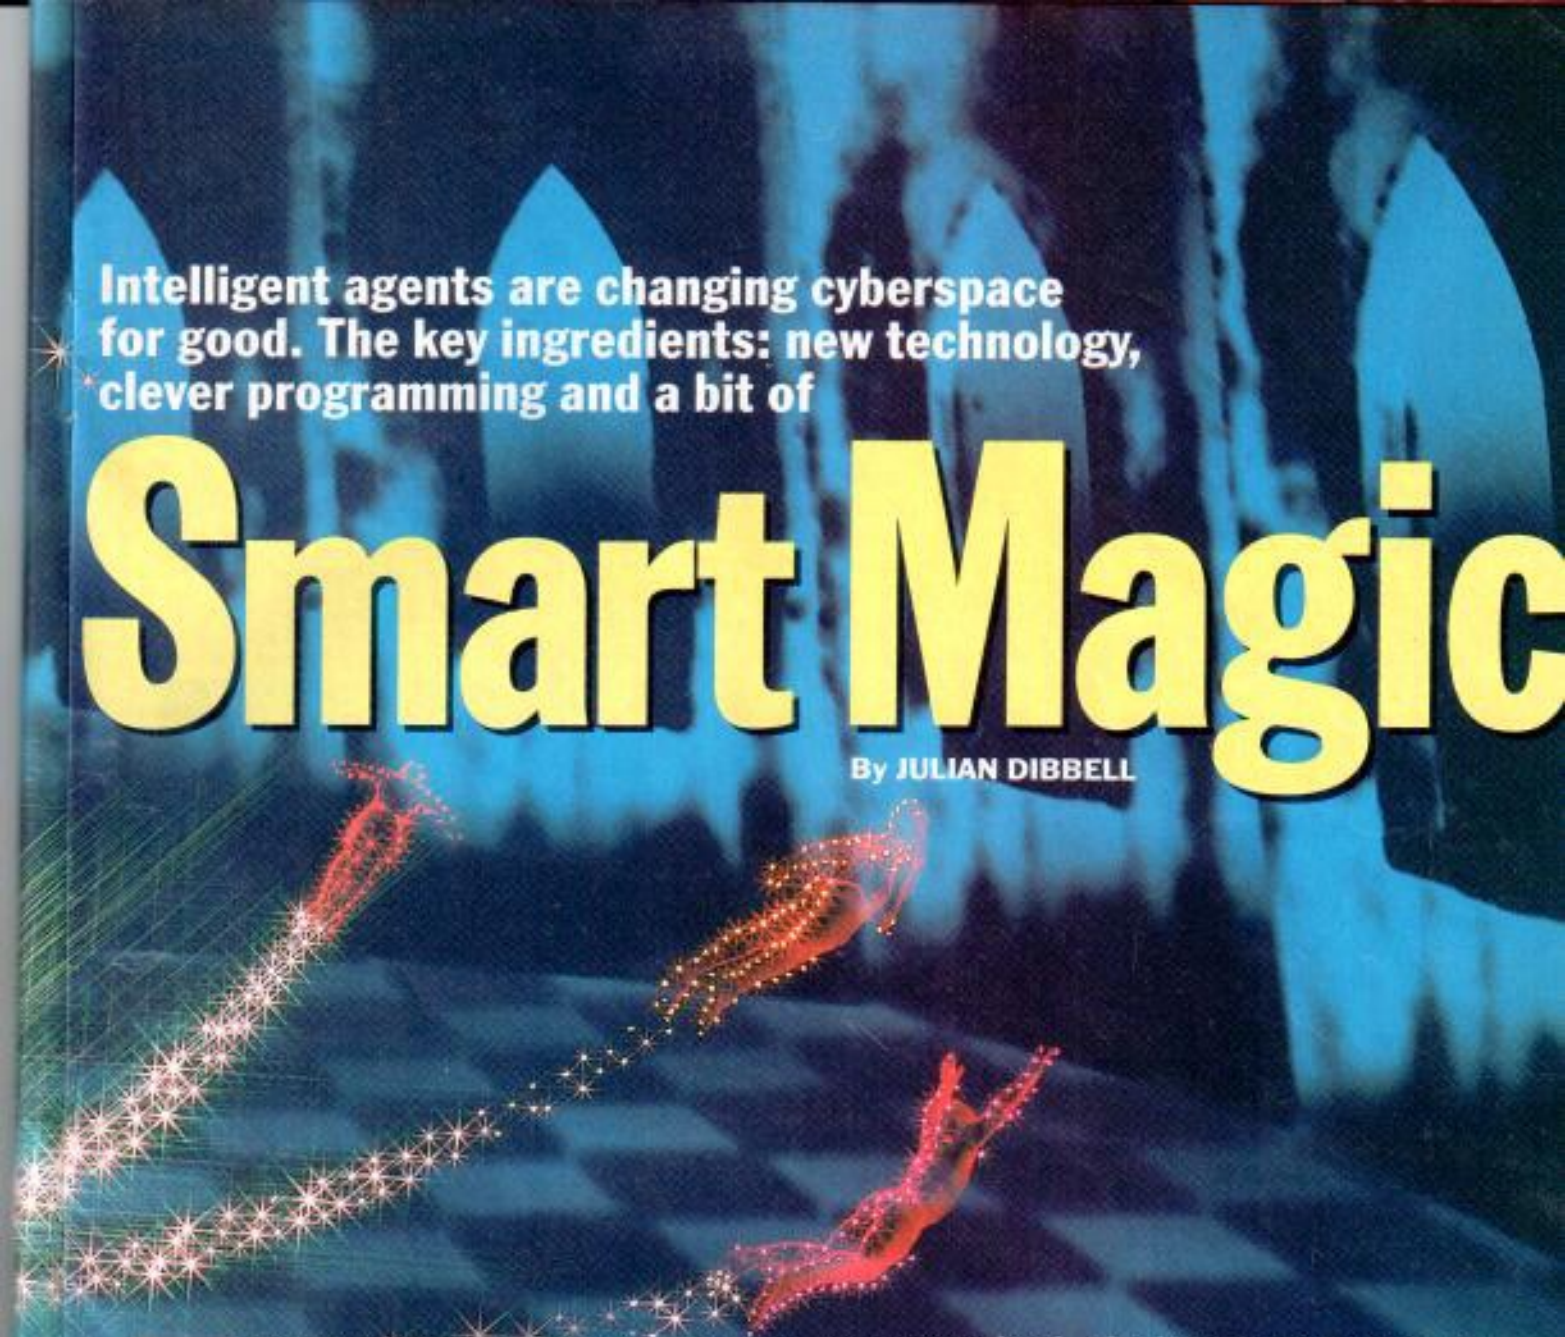

THE DREAM GOES SOMETHING LIKE THIS. You're pitching a deal in the office of a potential client, and it's turning out to be a harder sell than you had hoped. You're stressed. Your pulse quickens. Your wristwatch takes note. That's right, your wristwatch: inside it there's a tiny computer running a powerful piece of all-purpose software that watches over you every second of the day. Alerted to your rising pulse rate by a tiny biosensor built into the watchband, the program—an "intelligent agent" in geekspeak—makes wireless contact with a digital camera attached to your eyeglasses. Matching it against a database of familiar images, the agent instantly "recognizes" the face across the desk and correlates it with your stress level. The logical conclusion: you need help.

In a flash the agent makes a cellular connection to the Net and sends a dozen copies of itself out to search for information relevant to your pitch. After a few minutes you hear its digitized voice murmur discreetly in your wireless earphone. "Thought this might interest you," the agent says, as a news item—invisible to your bargaining partner—appears in the lower right-hand corner of your glasses. The spot price of one of your key raw materials has just plummeted in a certain Southeast Asian market, you are informed, allowing you to shave your bid by 5 or 6 points and clinch a deal that only moments before looked like a goner. Once again, you think to

MARK RICHARDS—CONTACT FOR TIME DIGITAL

General Magic's co-founder, CEO and chief evangelist, Marc Porat

yourself, your trusty digital companion has saved the day ... and then you wake up.

An impossible dream? Maybe not. Though the wristwatch you're wearing in 1996 knows about as much about your vital signs as Boris Yeltsin does, the dream of intelligent agents—software that "understands" what your interests are and can look after them unattended—is creeping closer to reality. The recent explosion of the Internet as a tool of our social, cultural and financial lives has significantly heightened the incentive to develop and refine such electronic servants; suddenly, information overload has practical implications. Many of us now get more E-mail than letters. But how on earth do we sort it all? The Web offers deep wells of information. But how much of it can we trust, and can we get it when we need it? "Humans can't live in cyberspace all the time," explains Paul Saffo, director of the Institute for the Future in Menlo Park, California. "They increasingly are going to need alter egos that [can]."

The search for those alter egos is going on in computer labs around the globe, as engineers experiment with ways to locate information amid the bountiful chaos of the World Wide Web. As one would expect, commercial enterprises are in the thick of it. Three years ago, Silicon Valley start-up General Magic (backed by Apple, AT&T and other info-age giants) announced plans to create a special programming language called Telescript. The company claimed Telescript would release an army of agents that would function independently—not only of their users' supervision but of their users' machines, transporting their entire code from one computer to another, researching, negotiating, buying and selling on each user's behalf. But General Magic's dreams ran afoul of the Internet's ethos of free and open access. Its plans for a proprietary online universe lie in terminal disrepair. Regrouping, the company is shipping a new Web-centered line of Telescript development tools to anyone who asks (some 4,000 already have), and its executives rarely miss an opportunity to proselytize for their magical vision of the Net as one big, eager assistant, instead of the largely passive medium it is today.

NOT ALL MAGICIANS LIVE IN SILICON Valley, however. No less a player in the race to produce real magic is a clutch of blue-sky experimentalists at the Massachusetts Institute of Technology's Media Laboratory, known there as the Autonomous Agents Group. Led by Pattie Maes, the group's founder and principal researcher, the Media Lab team is working to create intelligent agents that are not just bits and bytes but also "living" beings.

"Ultimately, what I hope we can create is a whole new set of [evolving] life-forms on the Net," says Maes (pronounced Moss), 35, speaking with a slight Belgian accent in language more suitable to Darwin than to a programmer. "All these different species of software agents ... [would] specialize toward whatever certain users are interested in. And that's how they'd survive."

"Now there's a limit to how much damage agents can create. The moment they reach into the real world, things will get real unpredictable."

— PAUL SAFFO
Institute for the Future

The environment in which for now such creatures must survive—indeed thrive, at least in economic terms—is the World Wide Web. As General Magic's cofounder and chief evangelist Marc Porat puts it, "The Web doesn't do a thing for you unless you get out there and pump on it. If you poke it, it'll work; but the minute you stop poking, it goes fallow." And as Web surfers know all too well, finding what you need at any given moment can be a teeth-grindingly labor-intensive task; the ability to delegate the work to a Web-savvy agent would be a godsend, of course. New products with names like Surfbot and WebCompass can, for example, gather info from your favorite sites while you sleep and deliver a personalized newspaper every morning.

The Dynamic Duo Studio for Time Digital

Developers claim it won't be long before agents like these can handle more sophisticated feats—making guesses at the kinds of Web pages users might be interested in, say, instead of depending on us to tell them explicitly. But that's where the real world ends and the magic begins.

One of Maes' earliest and best-known experiments showed how simple and sensible it can be to build "natural" selection into agent systems. In a program designed to help users keep abreast of the flood of news on the Internet, Maes pitted slightly different retrieval agents against one another for the user's approval. Those that brought back the most interesting articles were permitted to pass their "genes" on to the next generation, some making precise copies of themselves, some introducing small mutations, others reproducing "sexually" by swapping patches of code with fellow agents. The net result: a self-adapting, personalized news feed that proved as flexible and fine-tuned as evolution itself.

Less lofty aims inspire the Kasbah project—a prototype marketplace for online bargaining agents designed by Maes' student Anthony Chavez. The program helps users find what they want in an online clas-

sified ad and then haggle over the right price. To be successful the Kasbah agent has to "understand" how to bid on products and how to balance the cost of an item against its features. That sort of glitzy feature—"Use an agent! Save money!"—may not have much philosophical *gravitas*, but it probably does lie a little closer to the heart of the Internet's current buzzworthiness. Right now, after all, the great looming question about the Net is whether it will ever become the vast engine of high-speed, hyperefficient electronic commerce it sometimes seems destined to be. Agents like Kasbah's could turn out to be a critical piece of that puzzle.

There is one project in Maes' stable, though, that is turning out to be something of a puzzle itself. It's called firefly, and if you've got a Web connection, it's worth a look (*www.ffly.com*). Produced by Agents Inc., a Cambridge, Massachusetts, company Maes established with some former students, firefly offers an uncannily helpful music-recommendation service. When you first log in, it invites you to rate records and artists you're already familiar with, but then it quickly analyzes those ratings and pinpoints

your musical taste on a detailed statistical map of firefly users' preferences. By peeking at the ratings lists of your nearest "neighbors" on that map, firefly can make suggestions of almost eerie perceptiveness: Like Joni Mitchell? You may love Suzanne Vega.

IT'S AN ELEGANT HACK, AS THE TECHIES say, and downright useful too. But whether it is in fact an agent, as Maes claims, is another question, and one that's got leading researchers in the field scratching their heads. Even Maes' Media Lab colleague Leonard Foner (known for a much cited 1993 paper titled, "What's an Agent, Anyway?") has his doubts. On the one hand, he says, the intimate knowledge of user interests that firefly acquires does give it the personalizing ability he expects from an agent. But in the equally important aspect of autonomy, he feels, firefly comes up short. "You go to it with a question and it gives you an answer, and that's that," says Foner. "It really doesn't take independent action."

But even those two notions—independent action and specialized knowledge—may not be enough to define what an agent is. After all, an alarm clock rings independently in the morning, and you could argue that it has the specialized knowledge that you want to wake up. A stretch? Maybe, but it's not too much for the dozen companies now marketing as agent-based technology what really amounts to the software equivalent of alarm clocks—automatic file-maintenance programs, hard-drive backup schedulers, that sort of thing. So common has the ploy become, in fact, that experts long ago stopped paying it much attention. But firefly's relative sophistication commands respect—and so does Maes' reputation. Thus the system's claims to be an agent haven't so much provoked dismissal as forced researchers to re-examine their own usage of the term.

Don Norman, for one, an Apple vice president, recently told *Internet World* magazine that partly as a result of his quandary over whether to call firefly an agent or not, he had come to the conclusion that "maybe we should just never use the word anymore." The term agent has come to mean too many things to too many people, he believes, noting also that one hallmark of the kinship between software agents and artificial intelligence is that both fields aim at a maddeningly moving target. "Whenever any problem seems solved by the people in artificial intelligence, then somehow it no longer is artificial intelligence," said Norman. Defining an agent may prove no easier.

For her own part, though, Pattie Maes seems to hold with the growing number of

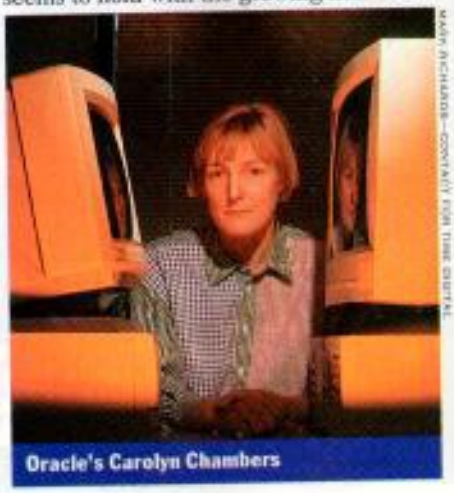
MARK RICHARDS—CONTACT FOR TIME DIGITAL

Oracle's Carolyn Chambers

people who insist that intelligent agents can and should be defined less in terms of their technological features than as a function of the mind-sets both programmers and users bring to them: if you think you are programming an agent, you are.

On the user's end, the governing philosophy might be said to borrow a page from theorist Daniel C. Dennett, who

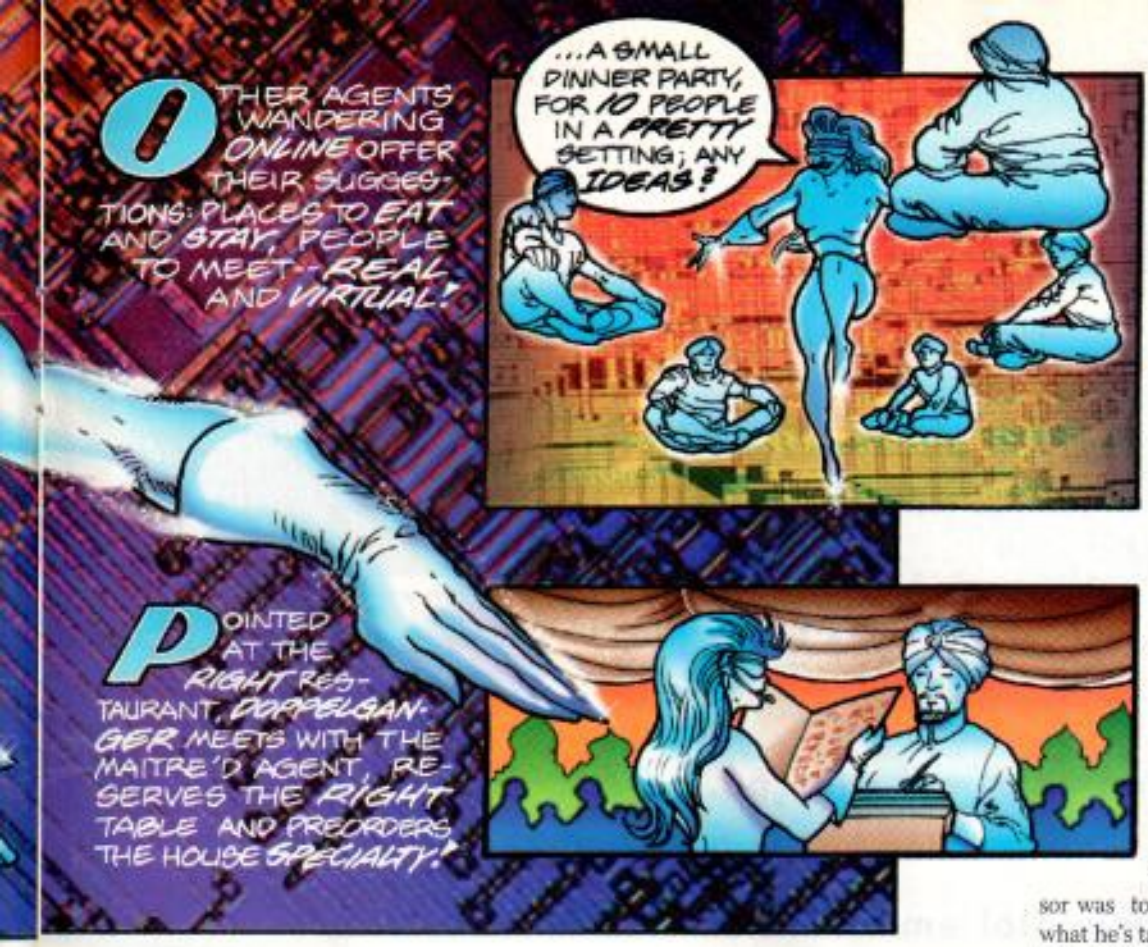

makes a case for the usefulness of attributing "belief, desires and rational acumen" to otherwise complicated systems such as people, animals and computers. In other words, to the extent you're capable of sharing an engineer's understanding of firefly as a "mechanism for vector-based matching of user interests in a centralized system," then that's exactly what firefly is. If, on the other hand, you find it more helpful to think of firefly as your very own digital proxy, consulting with other people's proxies 24 hours a day in its unflagging eagerness to sniff out cool records for you, then by all means go ahead and call it an agent.

THAT'S WHAT THE FOLKS AT BIG software companies are starting to do, cranking out new projects with the word agents in them as fast as possible. Oracle has worked for years to figure out the fastest way to ship search requests through giant databases of text and photos. Oracle Mobile Agents, an intelligent networking product, delivers the results of an agent's search expeditiously to you, whenever and wherever you need it. "If I want to keep track of competitors or important mail, I want some sort of process or intelligent agent going out, looking for that information, gathering it up and sending it to me the next time I use my laptop," says Carolyn Chambers, who runs the Mobile Systems Group. And the real magic, she says, is that the agent does this by itself, autonomously.

IBM has found similar promise in agent technology. Among its newest products is Alter Ego, a program designed to manage and handle E-mail. The system starts off by employing rules that the user defines—"Notify me of all mail from my boss," for instance—and then slowly learns what the user needs and when. Next up is adaptive behavior, which will let the program learn how you read mail and speed-sort the most important letters to the top of the stack. "E-mail is ubiquitous, and everyone gets overloaded," explains Don Gilbert, who manages IBM's intelligent-agent research group. But turning over your E-mail account to a computer program requires more than great technology: "You have to trust your agents," Gilbert says.

That may be the toughest leap of all. Looked at from just about any angle, the mix of software agents and the Internet poses risks that no amount of tinkering can ever fully defuse. Cryptographic security measures may help keep the mobile programs of Telescript and similar schemes from getting out of hand as they flit unsupervised from computer to computer. But once those measures fail, not much remains to distinguish such an agent from a garden-variety computer virus. Similarly, a sophisticated battery of crypto-locks will have to be applied to even the smartest personal programs before they can be trusted to head out onto the Net with the wealth of confidential data they will need to do their job. Even then there won't be any ironclad guarantee that these programs won't give up information they shouldn't to a slickly programmed con agent.

Or consider the possibilities that the futurist Paul Saffo raises in asserting that dirt-cheap sensors—"eyes, ears, sensory organs for our computers"—may be to the next decade what the microprocessor was to the past one. On the one hand what he's talking about is the advent of that ultimate intelligent agent we dream about—the constant companion that looks out *through* us as well as for us, monitoring our heartbeats and our daily lives, lending a hand when we least expect it. On the other, though, what he's talking about is something whose full dimensions we can scarcely imagine. "Right now," says Saffo, "there's a limit to how much damage agents can create because they can only operate in cyberspace. But the moment they have the capability to reach out into the real world, then things will get real unpredictable." Unpredictable—maybe. But also, possibly, wonderfully easier. —*With reporting by Lisa Granatstein/New York*


IBM's Don Gilbert smartens up E-mail

MARK RICHARDS—CONTACT FOR TIME DIGITAL

General Magic's co-founder, CEO and chief evangelist, Marc Porat

yourself, your trusty digital companion has saved the day ... and then you wake up.

An impossible dream? Maybe not. Though the wristwatch you're wearing in 1996 knows about as much about your vital signs as Boris Yeltsin does, the dream of intelligent agents—software that "understands" what your interests are and can look after them unattended—is creeping closer to reality. The recent explosion of the Internet as a tool of our social, cultural and financial lives has significantly heightened the incentive to develop and refine such electronic servants; suddenly, information overload has practical implications. Many of us now get more E-mail than letters. But how on earth do we sort it all? The Web offers deep wells of information. But how much of it can we trust, and can we get it when we need it? "Humans can't live in cyberspace all the time," explains Paul Saffo, director of the Institute for the Future in Menlo Park, California. "They increasingly are going to need alter egos that [can]."

The search for those alter egos is going on in computer labs around the globe, as engineers experiment with ways to locate information amid the bountiful chaos of the World Wide Web. As one would expect, commercial enterprises are in the thick of it. Three years ago, Silicon Valley start-up General Magic (backed by Apple, AT&T and other info-age giants) announced plans to create a special programming language called Telescript. The company claimed Telescript would release an army of agents that would function independently—not only of their users' supervision but of their users' machines, transporting their entire code from one computer to another, researching, negotiating, buying and selling on each user's behalf. But General Magic's dreams ran afoul of the Internet's ethos of free and open access. Its plans for a proprietary online universe lie in terminal disrepair. Regrouping, the company is shipping a new Web-centered line of Telescript development tools to anyone who asks (some 4,000 already have), and its executives rarely miss an opportunity to proselytize for their magical vision of the Net as one big, eager assistant, instead of the largely passive medium it is today.

NOT ALL MAGICIANS LIVE IN SILICON Valley, however. No less a player in the race to produce real magic is a clutch of blue-sky experimentalists at the Massachusetts Institute of Technology's Media Laboratory, known there as the Autonomous Agents Group. Led by Pattie Maes, the group's founder and principal researcher, the Media Lab team is working to create intelligent agents that are not just bits and bytes but also "living" beings.

"Ultimately, what I hope we can create is a whole new set of [evolving] life-forms on the Net," says Maes (pronounced Moss), 35, speaking with a slight Belgian accent in language more suitable to Darwin than to a programmer. "All these different species of software agents ... [would] specialize toward whatever certain users are interested in. And that's how they'd survive."

The Dynamic Duo Studio for Time Digital

"Now there's a limit to how much damage agents can create. The moment they reach into the real world, things will get real unpredictable."

— PAUL SAFFO
Institute for the Future

The environment in which for now such creatures must survive—indeed thrive, at least in economic terms—is the World Wide Web. As General Magic's cofounder and chief evangelist Marc Porat puts it, "The Web doesn't do a thing for you unless you get out there and pump on it. If you poke it, it'll work; but the minute you stop poking, it goes fallow." And as Web surfers know all too well, finding what you need at any given moment can be a teeth-grindingly labor-intensive task; the ability to delegate the work to a Web-savvy agent would be a godsend, of course. New products with names like Surfbot and WebCompass can, for example, gather info from your favorite sites while you sleep and deliver a personalized newspaper every morning.

Developers claim it won't be long before agents like these can handle more sophisticated feats—making guesses at the kinds of Web pages users might be interested in, say, instead of depending on us to tell them explicitly. But that's where the real world ends and the magic begins.

One of Maes' earliest and best-known experiments showed how simple and sensible it can be to build "natural" selection into agent systems. In a program designed to help users keep abreast of the flood of news on the Internet, Maes pitted slightly different retrieval agents against one another for the user's approval. Those that brought back the most interesting articles were permitted to pass their "genes" on to the next generation, some making precise copies of themselves, some introducing small mutations, others reproducing "sexually" by swapping patches of code with fellow agents. The net result: a self-adapting, personalized news feed that proved as flexible and fine-tuned as evolution itself.

Less lofty aims inspire the Kasbah project—a prototype marketplace for online bargaining agents designed by Maes' student Anthony Chavez. The program helps users find what they want in an online clas-