need for a long, slow twisting would necessarily space jerks by many decades, as observed, and the inherent rapidity of loop formation would still ensure a sharp jerk. Indeed, a jerk may be a particularly prominent example of the twisting of the east-west–oriented, ring-shaped core field that presumably generates the north-south dipole field observed at the surface.

Jerks may provide a view of the inner workings of the core, but they may also tell researchers about the window through which the jerks must be viewed. The electrically conductive mantle screens out much of the magnetic activity of the core before it can reach the surface. On the basis of the nature of the 1969 jerk, the French group calculated a relatively low mantle conductivity of less than 100 ohm^{-1} $meter^{-1}$. George Backus of the University of California at San Diego has since shown through theoretical considerations that the abruptness of the jerk is not limited by mantle conductivity in any simple way. Thus, contrary to initial expectations, extracting measurements of conductivity from magnetic observations will not be easy.

Another approach would be to watch how fast the earth rotates, of all things. The fluid core can alter the rate of earth rotation through the coupling of its magnetic field to the mantle. The higher the mantle's conductivity, the stronger the coupling. Searches for a correlation between the rate of rotation and magnetic field variations continue, but no single research group has found a long-term correlation that holds up throughout a complete record. Both kinds of records have their share of imperfections, but the synthesis of magnetic variations from worldwide observatories is a major stumbling block.

Geophysicists see a satellite or a series of satellites as the only practical solution to their magnetic data problems. They had one in orbit in 1980 but only long enough to get a single picture of the magnetic field. They say they need another one soon to see how the 1980 field is changing before it becomes unrecognizable.—RICHARD A. KERR

Additional Readings

1. L. R. Alldredge, *J. Geophys. Res.* **89**, 4403 (1984).
2. J. Ducruix, V. Courtillot, J.-L. Le Mouël, *Geophys. J. R. Astron. Soc.* **61**, 73 (1980).
3. C. Gire, J.-L. Le Mouël, J. Ducruix, *Nature (London)* **307**, 349 (1984).

The Intelligence of Organizations

Humans work in organizations, and increasingly, so do computers; are there lessons to be learned?

Traditionally, artificial intelligence research (AI) has focused on human cognition, the individual human mind. More and more, however, researchers are turning their attention to a collective form of human intelligence: organizations.

This is not exactly a new idea. Carnegie-Mellon University's Herbert Simon became one of the founders of AI in the 1950's out of an interest in organizational decision-making, for which he subsequently won the Nobel Prize in Economics. But the last few years have brought an upsurge of interest in such things as parallel processors and robot assembly lines (*Science*, 10 August, p. 608), societies of machines that turn out to face all the same issues of communication, coordination, and organizational structure that human societies face. Those problems, and the organizational metaphor, were the subject of a panel discussion last month at the annual meeting of the American Association for Artificial Intelligence (AAAI)*.

A straightforward example of the organizational approach was described by Thomas W. Malone of the Massachusetts Institute of Technology (MIT), who addressed an increasingly common situation in laboratories and offices: high-performance personal computers are scattered around on desks and laboratory benches, with a lot of them sitting idle at any given time. This is obviously a waste, says Malone. However, if the computers can communicate with each other through a local area network, some of that idle power can be harnessed using a system called "Enterprise," which he has developed along with Richard Fikes and Michael Howard of the Xerox Palo Alto Research Center.

Enterprise is based on the "contract network" scheme first proposed by Reid Smith of Schlumberger-Doll Research Center. "If you have a computation-intensive task to do," Malone explains, "your computer sends out a 'request for proposals' over the network describing the task and its priority." Each of the other processors then checks its own priorities and its available data, he says, and responds with an estimate of how quickly it could finish that task. The lowest bidder gets the job.

"So by making lots of local decisions in the bidder and client machines," says Malone, "you get a globally coherent assignment of tasks without having to set up any one machine as a 'foreman.' " In fact, a mathematical analysis suggests that Enterprise will often be substantially better than having a foreman.

Victor Lesser of the University of Massachusetts, meanwhile, is concerned with how networks can cope with the uncertainties of the world.

Imagine something like an automated factory or an air traffic controller network, he says, where lots of sensors and processors are distributed over a wide area. "In classical distributed processing," he says, "each processor is assumed to produce accurate results based on correct and complete information." But in practice, doing things this way is either impossible or very costly and inefficient. "There is an enormous burden of communication and synchronization," he points out. "The processors spend most of their time waiting for someone else. Worse, as you build larger and larger systems, you can't assume that all the processors will be functional. You can't assume that you have global information. You can't assume that all the communication channels are working."

An effective way to cope with such uncertainties is an AI technique known for historical reasons as the "blackboard" architecture, which dates back to the HEARSAY speech understanding system that Lesser helped design at Carnegie-Mellon in the mid-1970's. The idea in HEARSAY was that multiple "agents" would analyze the incoming sounds from differing points of view. One would identify phonemes, for example, another would piece together words,

*The fifth annual meeting of the American Association for Artificial Intelligence, 6 to 10 August 1984, the University of Texas at Austin.

and yet another would try to predict the next word on the basis of the preceding phrase. Each one alone would reach tentative, incomplete conclusions. But together they could usually narrow down the possibilities enough to reach a satisfactory answer.

This multiagent approach has obvious parallels to human teamwork, as in, say, a panel of consulting physicians. To see how it might work in the example of a distributed air traffic control system, says Lesser, imagine that one radar installation has only an approximate idea of an airplane's position when it hands off control to the next radar down the line. The second radar could still use that information to narrow its area of search, improve its signal-to-noise ratio, and identify the airplane faster and more accurately than it could have on its own.

However, there is one major problem to this approach, Lesser admits, and that is global coherence. "How do you get each node to do the most fruitful thing for the overall activity?" he asks. The whole idea is to keep the processors from getting bogged down in communications and coordination. "But without a centralized view," he says, "you may have nodes doing redundant processing, or *wrong* processing." To reach some acceptable compromise, Lesser and his colleague Daniel Corkill are drawing on another mainstay of human organizations: put an overall structure in place to impose long-term coherence, and then give each node a certain range of responsibility within which it can make its own decisions.

The idea of computers forming organizations raises some fundamental research questions. For example, as things are now, the organizations are specified by the programmers beforehand. Can the computers be taught to organize and reorganize themselves on their own to fit the problem at hand? Lesser has been thinking about how to do that, but finds it slow going. "You find that the question of 'What is an organization?' is very difficult to define," he says. "Part of our work is to define a language in which you can talk about organizations symbolically." Malone has also been thinking along these lines. He and several colleagues have begun to develop an analytic framework for evaluating the efficiency and flexibility of organizations, including such factors as production costs, coordination overhead, and the vulnerability of the system to isolated failures or to sudden changes in the environment.†

Another research question: How can one machine reason about another's knowledge, intentions, and beliefs? "In human communication, a lot of what I say depends on what I believe about your state of mind," says MIT's Randall Davis, who organized the AAAI panel. "For example, if I think you know about something, I won't bother to explain it to you. If I think you don't believe it, I may argue for it." Exactly the same kind of considerations come up when machines have to communicate.

Michael Genesereth of Stanford University has been looking at some of these issues by mathematically modeling groups of computers, or "agents," that interact according to rules based on game theory. "The thing that intrigues me," he told the AAAI, "are the circumstances in which cooperation will emerge spontaneously from individual agents.

The simplest case is when the agents cannot communicate with each other, he explains. As long as the agents know about each other's desires and intentions, they still end up cooperating simply because that is the way they can best achieve their individual goals. "What we've found is that rationality necessarily leads to cooperation," he says.

On the other hand, says Genesereth, things begin to get very interesting indeed when the agents *can* communicate. Sometimes they cooperate. Sometimes they establish ad hoc organizations. But sometimes they try to manipulate each other. Sometimes they withhold information. And sometimes they lie. Genesereth hopes to do a lot more work in understanding why and when this happens.—**M. MITCHELL WALDROP**

†Thomas W. Malone and Stephen A. Smith, *Tradeoffs in Designing Organizations: Implications for New Forms of Human Organizations and Computer Systems* (CISR WP #112 and SLOAN WP #1541-84, Center for Information Systems Research, Massachusetts Institute of Technology, Cambridge, March 1984).

High Spatial Resolution Ion Microprobe

With focused scanning ion beams, researchers can now make elemental maps with 400-angstrom resolution by secondary ion mass spectrometry

Photons and electrons are the mainline tools for imaging and analyzing the composition of solid samples, but ions can be used just as well. In mid-July at the 31st International Field Emission Symposium held in Paris, Riccardo Levi-Setti of the University of Chicago described the current state of the art: a scanning ion microprobe that has produced elemental maps with a spatial resolution of 400 angstroms by means of secondary ion mass spectrometry. The instrument can also make scanning electron microscope-like images at the same level of detail.

Researchers are lining up at Levi-Setti's door with proposals for collaborations. But it may not be necessary to fly to Chicago to find a high-resolution scanning ion microprobe. The flip side of imaging is pattern generation. Microelectronics researchers, especially those in Japan, are investigating ways to use scanning ion beams to draw the fine features of integrated circuits. Spurred primarily by this interest, several companies are or soon will be marketing scanning ion beam machines with promised resolutions as high as 500 angstroms.

At Chicago, Levi-Setti has been interested for several years in high-resolution ion microscopy by means of focused scanning ion beams. However, it is not images but the ability to make spatially resolved elemental analyses (microanalyses) at the submicrometer level that has researchers scrambling to use the scanning ion microprobe. Frederick Coe of the Chicago Medical School, who is collaborating with Levi-Setti on projects to study developmental processes in the skull bone (calvarium) of the fetal rat and to investigate the microstructure of kidney stones, says that the ability to map the presence of individual atomic species or especially isotopes of the same atom with submicrometer resolution is the key. "Otherwise I would just use [a scanning electron microscope]."

Microanalysis by means of the x-rays emitted when a scanning electron beam strikes a solid surface is the standard laboratory method. Although the resolution for imaging is fixed by the diameter of the electron beam and can be as small