Microsoft Researches Its Future

The software giant is betting $200 million a year that research in esoteric areas of physics and mathematics will yield breakthroughs in personal computing

SEATTLE—Boot up a copy of Microsoft's Office 97, and an animated cat—or, if you prefer, a paper clip or a cartoon of Albert Einstein—strolls across your screen and settles down. Before long, the cat is kibitzing: "It looks like you're writing a letter. Would you like help?" it asks, in a cartoon-style balloon. The cute critter may seem like just another gimmick, but it's one of the first commercial products to come out of Microsoft's growing research department—and a sign of change in the way the world's largest personal-computing software company prepares for the future.

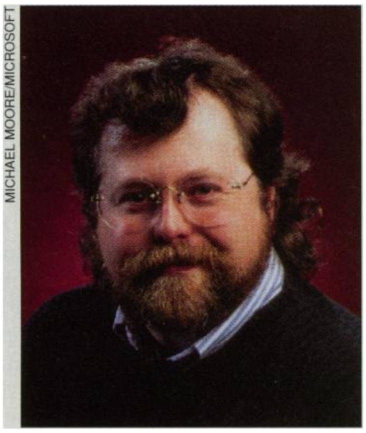

MICHAEL MOORE/MICROSOFT

"We can take [a new] idea and put it in the hands of 100 million people."

—Nathan Myhrvold

In the past, Microsoft has been slow to catch on to the hottest trends in personal computing. The first operating system with windows and icons was invented at Xerox Palo Alto Research Center (PARC) in California a decade before Microsoft released its first version of Windows—and these were standard features of Apple computers long before Microsoft caught up. Netscape enjoyed a head start of at least 2 years over Microsoft's Internet Explorer in the well-publicized "browser wars." Although Microsoft has been remarkably successful at what chair Bill Gates calls "embracing and extending" the ideas of others, it would like to be first off the blocks when the Next Great Thing comes along. To improve its chances, it is assembling a star-studded research department in the mold of such famous corporate research laboratories as IBM, Xerox PARC, and Bell Laboratories.

"Microsoft Research is unique in that it is the first software-centered lab set up by a software company," says Nathan Myhrvold, the chief technology officer of Microsoft. Founded in 1991, the laboratory already rates seventh among all academic and corporate research labs in computer science, according to an informal survey of 200 computer scientists conducted by *Business Week*. Myhrvold estimates the division's budget at about 2% of Microsoft's revenues, or about $200 million. It has grown to include 250 scientists in over 20 research groups, and this year it added an $80 million branch in Cambridge, England (*Science*, 20 June 1997, p. 1783). Its growth has placed Microsoft on the short list of companies to recognize that research "makes both business sense and research sense," says John Seely Brown, director of the 250-member Xerox PARC. "It is a wonderful event, and I highly applaud it."

Observers credit the sandy-haired, multitalented Myhrvold for getting Microsoft Research off to a flying start. "Nathan's a wonderfully charismatic and brilliant guy," says Brown. Myhrvold has authentic whiz-kid credentials: After earning a Ph.D. in physics from Princeton University at age 23, he studied briefly under Stephen Hawking at Cambridge University before launching a software company called Dynamical Systems. The company was bought out by Microsoft in 1986; 5 years later, Myhrvold himself proposed the new research division to Bill Gates.

Although Microsoft's high pressure and burnout rate are legendary, the atmosphere at Microsoft Research is low-key, academic, and collaborative. Small, oddly shaped rooms and open doors make it as easy to have a conversation in the doorway or hallway as in the office. Ties and business cards are out; Dilbert cartoons and whiteboards are in. Within broad limits, scientists are free to work on any project they want to. According to physicist Jennifer Chayes of the Theory Group, who was formerly at the University of California, Los Angeles, "I felt unquestionably more pressure to do applied research for my grant [at UCLA] than I do here. … Here, if I work on an application it's because I'm really excited about it."

MICHAEL MOORE/MICROSOFT

The combination of benevolent oversight and a long time frame has led Microsoft Research into areas of basic science that seem far removed from the personal computing world of today—quantum field theory, decision theory, and statistical physics. The ultimate goal is a computer (or operating system) that will be as easy to interact with as a human—just by talking. It will not require special words, special commands, or even special input devices, such as a keyboard or mouse. This is what Brown calls "radical research"—more cross-disciplinary and technology-oriented than "basic" research but more visionary than "applied" research. Although the humanlike computer, for example, may sound like science fiction, Myhrvold is banking on technology's version of the trickle-down theory: If the company works toward that distant goal, he says, then shorter term benefits will drop out.

Take that cat—or Office Assistant, as it is formally known. Behind that reassuring figure, says developer Eric Horvitz of the Decision Theory Group, lies a Bayesian inference engine that "generates a probability distribution over the problem a user might be having before a query is input." Thus, Office Assistant can evaluate the probability that a user is writing a letter (a probability that goes way up when it encounters the words "Dear John") and anticipate that the user might want to know about Office 97's letter formats. The cartoon figure, for its part, comes from research by the User Interface Group on how humans interpret gestures and expressions—research that will be needed by that far-future computer to make sense of the vagaries of human communication.

By some measures, Microsoft Research is just getting started. For example, the 199 patents it received in 1997 are dwarfed by the 1724 patents awarded to industry leader

"I felt unquestionably more pressure to do applied research for my grant than I do here."

—Jennifer Chayes

Solving 'Hard' Problems—or Dodging Them

Theoretical mathematics and physics may seem far removed from the challenge of building better software. But in an example of Microsoft Research's willingness to gamble on basic research (see main text), one of its newest research groups is looking to those disciplines for approaches to some of computer science's toughest problems.

Computer scientists have found that certain types of problems, called NP-hard, are especially intractable, because the time it takes to solve the most difficult examples of an NP-hard problem seems to grow exponentially as the amount of input data increases. These problems range from the utterly pure (telling mathematical knots apart) to the highly practical (scheduling airline flights so that two planes don't need the same runway at the same time). Not only are all NP-hard problems equally refractory, but they are actually equivalent: An algorithm to solve one of them in subexponential time could be converted into an algorithm to solve any of them.

Because of this equivalence, researchers can work on any NP-hard problem they want to. Michael Freedman, a topologist who is the newest member of the Theory Group, studies the Jones polynomial—a tool for telling knots apart. As in any NP-hard problem, it takes exponentially longer to compute the Jones polynomial of a knot as the number of overpasses and underpasses in the knot increases. But physicist Ed Witten of the Institute for Advanced Study in Princeton, New Jersey, showed in 1988 that the Jones polynomial could be directly measured for knots in a topological quantum field theory—one of a group of theories describing the microscopic structure of the vacuum. Thus, in some as yet unimaginable "quantum field computer"—a device that would, in effect, compute with the quantum vacuum—Freedman believes it may be possible to "enter the den of the exponential dragon" and slay it.

His colleagues, statistical physicists Jennifer Chayes and Christian Borgs, are instead looking for ways to detour around the dragon's lair. Their favorite NP-hard problem, called three-satisfiability or 3-SAT, is a type of logical problem similar to the popular mystery game "Clue," but with clues written by an incompetent author. The goal is not necessarily to find the murderer(s) but to determine whether any of the suspects could possibly satisfy all the clues. More suspects make it more likely that someone fits the description of the murderer; if there are too many clues, it is likely that no one does. At either extreme, the problem is easy to solve. In between, when the ratio of clues to suspects is about 4.2, a "phase transition" occurs, and the two possibilities are about equally likely (*Science*, 27 May 1994, pp. 1249 and 1297). At that point, determining whether the murderer can be identified rises to a peak of difficulty.

Chayes and Borgs are now trying to determine how to identify similar phase transitions in other NP-hard problems—such as scheduling airplanes. It might then be possible to predict where the critical regimes are and avoid them, and abstruse physics might have a very practical payoff for software design. **–D.M.**

IBM. But the company has already made its mark in one way: its ability to snare eye-catching talent. The "look who's at Microsoft now" club includes Michael Freedman, formerly of UC San Diego, who is a MacArthur Fellow and a winner of the Fields Medal (considered the mathematical equivalent of the Nobel Prize) for his work on computational complexity. Other stellar recruits include James Blinn, a MacArthur Fellow for his work at the California Institute of Technology in educational animation, and Gary Starkweather, the inventor of the laser printer, snared from Xerox PARC. Myhrvold explains the quality and diversity of the staff this way: "We're willing to believe in a discipline and get a collection of people together that wouldn't have happened otherwise. We can create a hell of an interesting environment."

When Myhrvold decided that the area of Bayesian statistics sounded promising, for example, he says he "set out to track down the best guys in the field." David Heckerman, a decision theorist, recalls what happened then. "My dissertation [on Bayesian networks] won the Association of Computing Machinery doctoral dissertation award. It caught Nathan's eye." At first, Heckerman says, he told Myhrvold, "There's no way you're going to get me up here, but I have this company [Knowledge Industries] with Eric [Horvitz] and Jack [Breese], and they might want to join." Several visits later, Myhrvold persuaded Heckerman to come—and Horvitz and Breese as well. They now form one-third of Microsoft's Decision Theory Group, which created Office Assistant.

Another expanding group is the Theory Group, which recently brought Freedman aboard. Freedman says that Chayes, the co-founder of the group, came up with the name "Theory" as a euphemism for "Mathematics." But "Theory" may be the only term broad enough to encompass the group's interests. Freedman started out as a topologist, but his career took a detour into the theory of computational complexity by way of knot theory. Chayes and her husband Christian Borgs, on the other hand, are statistical physicists. Yet all three of them are bringing their insights to bear on solving "NP-hard problems," a hard nut at the center of computer science that bears on challenges from scheduling airlines to cryptography (see sidebar).

Talking picture. Comic Chat—a product of basic research at Microsoft—records an Internet chat by assigning cartoon figures to participants.

From concept to product

Such cross-disciplinary stimulation helps attract researchers to Microsoft—as does the company's ability to make young researchers into millionaires through stock options. Researchers there say they are also exhilarated by the speed at which a new idea can be transformed into software. As Myhrvold says, "If someone has a new idea, we can take that idea and put it in the hands of 100 million people."

One example is Comic Chat—an idea born out of the graduate dissertation of David Kurlander of the User Interface Group. Kurlander proposed that the history of an Internet chat session could be portrayed in a comic strip, making it easy for newcomers to the session to see what had gone on before their arrival. The idea went from concept to product in just 9 months and is now a regular feature of Microsoft's Internet Explorer.

Ideas reach Microsoft's product division through official channels—a Technology

Applications office—and direct contacts with researchers. Indeed, Kurlander was so committed to Comic Chat that he had himself transferred to a product group to see it through. Myhrvold adds that the proximity of research and product divisions on the same campus has helped produce a culture that encourages researchers to ponder the possible commercial applications of their work, no matter how abstract it may seem.

Yet researchers at Microsoft say they face few restrictions on their freedom to publish results. Unlike many other companies, including IBM and Xerox, Microsoft has no review process for papers and no intellectual property department. Researchers are simply expected to file their own patents when necessary. The benefits of this nonpolicy, in Myhrvold's view, are the automatic quality control that comes from peer review and the freedom to exchange ideas with outside researchers. The Theory Group, for example, plans to host regular visitors from academia for periods of anywhere from 1 day to 12 months. And Microsoft researchers often serve as de facto advisers for graduate students at the University of Washington. "People worry, 'That means we're going to lose some ideas,'" says Myhrvold. "Well, I've found that people who are too afraid of losing ideas are people that don't have very many."

Research leaders elsewhere might not be that sanguine. But Microsoft does not seem to be alone in its strategy of keeping the research focused on a problem relevant to the business—such as creating a computer with a completely intuitive interface—while letting people take any approach to it they want. Andrea Califano, the manager of computational biology at IBM Research, says the atmosphere there has changed "quite significantly" since the company's financial crisis in the early '90s, from a pseudoacademic environment where publishing papers was the only requirement to a more technology-driven model. Still, Califano says, "at least 50% to 60% of the work that we do would be very basic science." Michael Garey, director of mathematical research at Lucent Technologies Bell Laboratories, complains about the "misperception" that his company now does only applied research. "It hasn't gotten any less fundamental. We think in terms of building an intellectual foundation for a technological area we see as important to the company down the road."

Now that Microsoft is a convert to that notion, Myhrvold says it's time for other technology companies to take a longer view as well. "Most have lots of what I call 'r&D'—little r, big D. Or even no r, in the sense that they do no pure research."

–Dana Mackenzie

Dana Mackenzie is a mathematics and science writer in Santa Cruz, California.

PHYSICS

Reports Call for New Super-Accelerator

Physicists in the United States have been understandably timid about asking for a major new accelerator. The debacle of the Superconducting Super Collider (SSC), which Congress canceled in 1993 when it was already under construction, is still fresh in their minds. And it took years of negotiation to arrange a consolation prize: U.S. participation in what will be the highest energy collider ever built, the Large Hadron Collider (LHC) at CERN in Geneva, Switzerland. But the message was clear in two U.S. reports on the state of particle physics released last week: Another collider will be needed if physicists are to assemble a complete picture of the particles and forces that constitute the world.

Plans for some kind of successor or companion to the LHC—which hurls protons against protons—have been in the works since the mid-1980s. But in a once-a-decade review of the field prepared by the National Research Council (NRC) and a draft report to the Department of Energy's (DOE's) influential High Energy Physics Advisory Panel (HEPAP), physicists have made their most public plea yet for a new machine. "The LHC is set, so the stage is now open for the next [machine]," says Columbia University physicist William Willis, a member of the HEPAP panel.

On the wish list are three very different devices: a scaled-up LHC, called a Very Large Hadron Collider; a 30- to 50-kilometer-long Next Linear Collider (NLC) that would smash electrons together; or an even more fanciful device that would collide muons—the electron's short-lived, heavy brothers (*Science*, 9 January, p. 169). Physicists have been exploring all three possibilities, and some portions of the NLC have even been benchtested. Any of the three options will cost over a billion dollars and take a decade or more to plan and build.

But such a behemoth will be essential for moving beyond the existing picture of the subatomic world, says University of Chicago physicist Bruce Winstein, who chaired the NRC committee. The missing piece of the decades-old theory called the Standard Model is an account of how particles get mass. Most physicists think a hypothetical particle called the Higgs boson works behind the scenes to confer mass. Although physicists expect the Higgs to tumble out of the collisions in the LHC, sightings are expected to be rare at the energies the accelerator can achieve. And occasional glimpses won't be enough, because the Higgs is expected to lead the way beyond the Standard Model to an even more fundamental theory of particles and forces.

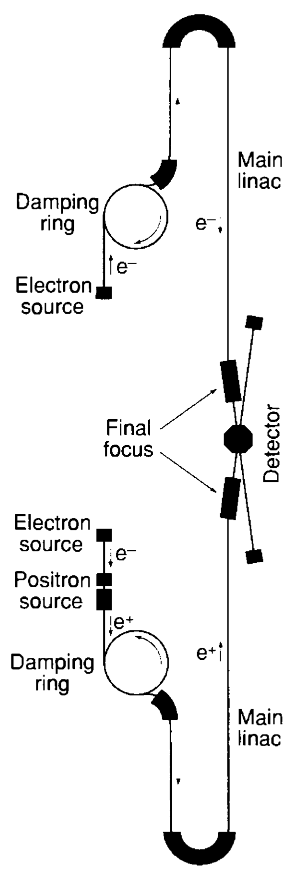

Object of desire. A concept for a future electron collider, tens of kilometers long.

SOURCE: SLAC

Finding out which, if any, of several candidate theories is right will take a more detailed investigation of the Higgs—and any other particles that turn up—than will be possible at the LHC, the reports say. The electron and muon colliders would give physicists a cleaner environment for studying the new particles, as these collisions produce less debris than do the proton collisions of the LHC. And a supersized LHC would generate higher energy collisions, allowing physicists to search for still more particles.

But the specter of the failed SSC clearly darkens the pages of both reports. The SSC cost rose billions of dollars over initial estimates, to nearly $12 billion just before its demise, and the project came up short in attracting funding from other countries. The SSC "cannot happen again," says Cornell physicist Persis Drell, who helped draft the NRC report. "If nothing else, that is branded on our foreheads."

The key to success, physicists hope, will be an affordable price tag and a global effort. Building even the NLC (the furthest along of the three options) will be "damn difficult and very, very expensive," says Donald Shapero, director of the NRC board on physics and astronomy. "This is going to be an international game from now on. No one country is going to be able to contemplate doing it alone."

Yet the NLC camp is already divided,