

Theory into Profit: Microsoft Invests in Mathematicians

Allyn Jackson

At a time when many companies have scaled back or eliminated their research efforts, Microsoft Corporation is betting millions on the notion that funding its own research will pay off. And it can afford the gamble: the company's revenues reached \$11.35 billion in 1997. Although Microsoft will not disclose the budget for its think tank, Microsoft Research, it says it plans to spend \$2.6 billion per year in research and development overall. The overwhelming dominance Microsoft enjoys in the world of personal computing software has brought a variety of reactions, ranging from the wisecracking—a CD-ROM satire of Windows features Bill Gates as a cyberpet who must be fed money to stay alive—to the dead serious—the Justice Department has brought an antitrust suit against Microsoft over the company's tactics for pushing its Internet browser. As a result, some have seen the growth of Microsoft Research as a publicity stunt designed to show that Microsoft is a good guy when it comes to funding research. Others have watched with envy as Microsoft snapped up some of the brightest stars in software research and development. Still others have applauded the company for providing a new haven for research when academia and the government are cutting back.

All of this might have passed unnoticed among mathematicians, who tend to be more accustomed to TeX and Unix than Word and Windows. But last July the mathematical community was startled to learn that Michael Freedman had become the first

Allyn Jackson is senior writer and deputy editor of the Notices. Her e-mail address is axj@ams.org.

Fields Medalist ever to take a job in industry: he joined the Theory Group at Microsoft Research. The group is headed by statistical physicists Christian Borgs and Jennifer Chayes and also includes graph theorist Jeong Han Kim. According to Phillip Griffiths, director of the Institute for Advanced Study in Princeton, the formation of the Theory Group "is one of the potentially most interesting developments related to the mathematical community in a long time."

A Research Lab is Born

Microsoft Research is the brainchild of the firm's chief technology officer, Nathan Myhrvold. He is someone who appreciates curiosity-driven research: he has a Ph.D. in physics (with a specialty in theoretical and mathematical physics) from Princeton University and as a postdoc studied cosmology with Stephen Hawking in Cambridge. Even after he came to Microsoft in 1986, Myhrvold remained interested in questions lying well outside the interests of the company. Last year he wrote a paper with a paleontologist in which they presented computer simulations showing that the tail of a large dinosaur might have been capable of cracking like an enormous bullwhip. When Myhrvold started Microsoft Research in 1991, he wasn't worried by the fact that no other software company did its own research. "We decided that we were of a size and our industry was in a state where it really made sense for us to put our own effort into inventing our future," he says.

Since then Microsoft Research has assembled a powerhouse staff. It now employs about two hundred fifty researchers, up from one hundred fifty

a year ago, and is slated to grow to six hundred by the year 2000. The top positions are filled by people with high quality credentials as well as hands-on experience in software development. The management bureaucracy is thin, comprising just three people: Richard Rashid, a computer scientist from Carnegie Mellon University who is Microsoft's vice president for research; Daniel Ling, an engineer from IBM who holds seven patents and is director of Microsoft Research; and James Kajiya, a graphics pioneer from Caltech who is assistant director. Microsoft Research comprises twelve groups whose areas range from speech and vision technology to cryptography to databases. Graphics is a special strength: In addition to Kajiya, one finds at Microsoft such stars as James Blinn, who rose to fame in the 1970s for his NASA animations of the moons of Jupiter and rings of Saturn, and Alvy Ray Smith, one of the founders of Pixar, which created *Toy Story*, the first entirely computer-generated film. Other luminaries include the creator of the VAX operating system, C. Gordon Bell; intentional programming guru Charles Simonyi; and laser printer inventor Charles Starkweather.

Microsoft Research has been the subject of a good deal of media attention, including a cover story in *Fortune* magazine late last year. The media often compare Microsoft Research to such legendary think tanks as Bell Laboratories and IBM Research. However, the comparison is premature. In their heydays Bell Labs and IBM Research spanned a much wider range of fundamental research than one finds at Microsoft. Bell Labs is famous not only for the invention of the transistor but also for developing the Unix programming language; IBM's innovations range from the hard disk drive to the programs that allowed Deep Blue to beat chess champion Gary Kasparov. The division of Bell Labs into AT&T Research and Lucent Technologies and the cutbacks at IBM Research have brought a more applied orientation to what they do today. Nevertheless, Andrew Odlyzko, head of the Mathematics and Cryptography Research Department at AT&T Labs, says, "if you look at how much unfettered research [Microsoft has], I would say it has still quite a bit less than" at AT&T, Lucent, or IBM.

One of the key differences is that those three labs cover a broad range of research in hardware, systems, and software, while Microsoft Research addresses only software. One would not find researchers at Microsoft producing a new surgical method for hip replacement, as IBM researchers recently did. The emphasis at Microsoft Research has been in areas such as computer graphics, databases, programming languages, operating systems, and networking, which have clear and immediate relevance to the company's main business. Odlyzko drily opines that Microsoft needs to do research in these areas "because they need to be able to build even more bloated software, loaded with novel

features, to be able to persuade people to upgrade. This is their business plan.... It's just a no-brainer." However, with the formation of the Theory Group the goals of Microsoft Research seem to be expanding. "They are hiring people with much more freedom than was customary at Microsoft," Odlyzko remarks, "which is certainly very good."

Spanning the Spectrum from Research to Products

The kind of person who fits perfectly into the culture at Microsoft Research is someone like Eric Horvitz, who works in the Decision Theory and Adaptive Systems Group. An expert in the use of Bayesian models for decision making, he writes theoretical papers in the subject and serves as the editor in this area for the *Journal of the Association for Computing Machinery*. But Horvitz is also very interested in Microsoft products: In a blaze of graphics that spill over two of the three workstations in his office, Horvitz explains how his work in decision theory formed the basis for the "Office Assistant" feature of the Office 97 software package, which uses Bayesian principles to decide when to offer help to the user and to guess at what kind of help might be needed. "I like getting my stuff into real-world products," Horvitz says. "But I also happen to like doing clean, easy-to-read papers with good results in them that can apply to anything." And he does mean anything: In addition to having a Ph.D., Horvitz earned an M.D. degree and has developed related methods to help trauma surgeons make good decisions under time and resource constraints.

Not everyone at Microsoft Research has to span the spectrum from research to product development, as Horvitz manages to do. What is more important is that the organization as a whole span that spectrum; the mathematicians in the Theory Group are at the far theoretical end. They were hired to do mathematics, which means that, like mathematicians everywhere, they write research papers, talk to colleagues, read journals, and attend conferences. They were not hired to solve somebody else's ugly integrals. As Kajiya puts it, "No research group is a service organization for the other groups." Chayes says her mathematical colleagues sometimes ask if Microsoft plans to bring in a lot of good mathematicians only to make them stop thinking about mathematics and start thinking about Microsoft products. "But why would we want to do that?" she asks. Microsoft hires plenty of good product developers, Kajiya points out. "We don't have to turn brilliant mathematicians into mediocre product developers."

While its distant connection to Microsoft products gives the Theory Group a lot of freedom, there are some risks. According to Griffiths one challenge facing the Theory Group is to be sure it cultivates that connection. "It's important for a number of



Jennifer Chayes



Christian Borgs

reasons, not the least of which is to have support in the rest of the company," he says. "Not just in principle, but when push comes to shove and resource allocations are made, research should be seen as important." The Theory Group has to stay plugged into academic research while also staying plugged into the business of Microsoft. Says Griffiths, "This is what they have to balance."

The Microsoft Milieu

Microsoft makes its home in the Seattle suburb of Redmond. Seen at a distance, the company grounds bear some resemblance to a modern college campus. Up close one realizes that most colleges couldn't afford this kind of meticulous landscaping. And there are too few battered old Volvos and too many sparkling new BMWs for this to be academia. The glossy modern buildings, as well as several bustling construction sites, connote high-tech wealth. The moneyed atmosphere continues inside Building 9, home of Microsoft Research, where original works by well-known artists such as Richard Diebenkorn hang on the walls. According to Chayes, the artwork also serves a practical purpose as signposts to help the staff find their way through the building. For here the luxury slows down, giving way to a maze of corridors with rows of offices that come in a standard size (tiny) and are generally crammed with computer terminals, books, gadgets, and other paraphernalia. Microsoft Research isn't so far from academia after all. In fact, Chayes drove a group to lunch in her well-worn Honda.

When she moved to Microsoft in January 1997, Chayes left behind a position as professor of mathematics at the University of California, Los Angeles. Christian Borgs, who is married to Chayes,

was a professor of physics at the University of Leipzig in Germany, where he headed a group in statistical physics. In trying to solve their intercontinental two-body problem, they considered an offer from AT&T Research. But when Nathan Myhrvold, an old friend of Chayes from their days as graduate students in Princeton, asked her and Borgs to launch the Theory Group at Microsoft Research, they jumped at the opportunity. Joining the group shortly thereafter was Jeong Han

Kim, who had been at AT&T Research for four years and was chafing a bit under what he felt was pressure to shift his research into more applied directions. The group got an enormous boost in prestige when Michael Freedman decided in July 1997 to leave UC San Diego and come to Microsoft. (Borgs, Chayes, and Freedman have not yet resigned their academic positions, and during the 1997-98 academic year Freedman was spending only one week a month at Microsoft. He plans to be there full-time by next fall.)

Why couldn't Microsoft rely on academia to do the theoretical research it needed? Kajiya says that Microsoft is in fact stepping up its support of university research; a prime example is the \$80 million computer science research center it has started at the University of Cambridge¹. "But it's just not enough," he says. "If there isn't someone on the industry side who is capable of appreciating the result, then technology transfer may be impeded." In addition, Borgs points out, academia can be slow to respond to new opportunities in research, especially in interdisciplinary areas. As Chayes puts it, "I could not go to UCLA and say, 'These areas of research are very important, and if we bring these twelve people together, we can really have an impact. So can you please give me twelve positions?' That wouldn't go over well with my colleagues."

But at Microsoft Research, this is exactly what she and Borgs are able to do: They are aiming for a steady state of twelve mathematicians on staff who, among other things, will make a headlong as-

¹The April 24, 1998, issue of The Chronicle of Higher Education contains a special collection of articles about the influence of Microsoft on college and university campuses.

sault on one of the central problems in theoretical computer science, the question of whether P equals NP (that is, whether certain algorithms that run in exponential time can be improved to run in polynomial time). Chayes says that the Theory Group is not as concerned about hiring in certain specified areas of mathematics as about getting excellent researchers with enthusiasm for fundamental problems in computer science. Joining the group at any one time will be eight to ten visitors, and perhaps four or five one- to two-year postdocs. Already the Theory Group has had a total of fifty visitors over the past year, for stays of anywhere from a half a day to six months. Sometimes long-term visitors are brought in to give courses on particular topics. When Borgs and Chayes wanted to learn more about the critical behavior of high-dimensional models in statistical physics, they bought out the teaching of Gordon Slade of McMaster University, who is an expert in the subject, and he gave them a four-month course on it. Slade also brought along a postdoc, Remco van der Hofstad, who received his Ph.D. in probability in Holland about a year ago.

With its visitors and postdocs, the Theory Group resembles a mini-mathematics institute. But unlike federally funded institutes that are accountable to the public, the Theory Group is accountable to the Microsoft shareholders. Therefore the group does not need the kind of formal scientific advisory boards one finds at publicly funded institutes (though it gets advice informally from many people in the mathematical sciences community). It also has a great deal more flexibility in how it spends money. In fact, according to Borgs and Chayes, the Theory Group spends as much as it wants on whatever it needs for its work. "We have no budget," Chayes says simply. At the beginning of the fiscal year, the Theory Group is asked to provide an estimate of big-ticket items; last year they bought some expensive computers and \$50,000 worth of subscriptions to mathematics journals. Aside from such estimates, says Borgs, "from our point of view, it's essentially an unlimited pot. Whenever we need something, we just ask for it." Visitors are well paid, receiving a consulting fee of \$300 per day, in addition to money to cover travel



Photograph courtesy Dept. of Mathematics, University of California, San Diego.

Michael Freedman



Photograph by Michael Moore. © 1998 Microsoft Corp.

Jeong Han Kim

expenses. Theory Group members will not comment on their salaries, except to say that they are comparable to "healthy" academic salaries, but not "superstar" academic salaries. They also will not reveal the total cost of their group, but it is clear that they could easily run up a bill of over a million dollars a year.

While Borgs and Chayes are clearly thrilled with the opportunities they have at Microsoft, there are some tradeoffs. They made a big leap in leaving academia: she had tenure at UCLA, and he had a lifetime job as a civil servant in Germany. Making the move from industry back to academia is not always easy. For Freedman the risk is smaller, because a Fields Medal is pretty good insurance that he could return to academia if and when he wants to. All Microsoft employees are "at will employees", meaning that they could be fired at any time. This is a big change from the security that academic tenure offers. Microsoft does not provide pension plans of the kind one finds in academia, though the legendary Microsoft stock options could provide a comfortable retirement. As in academia, the staff at Microsoft Research are free to publish their work wherever they like. Decisions about when a patent is needed to protect intellectual property are made by the researchers themselves.

Microsoft wants to give the Theory Group plenty of freedom, but it still needs to be sure that its investment is worthwhile. How will it make such an evaluation? Largely on the quality of the group's research, as measured in the usual ways such things are measured in academia and by research funding agencies. According to Chayes, right now the Theory Group evaluates itself by writing reports akin to progress reports on research grants. One important factor will be the level of interactions

with other groups in Microsoft Research, which so far has been high. Might there be some kind of review in, say, five years, to assess how well the Theory Group has done? Myhrvold says no. “If three years from now people in that group all want to do something else, why should I make them wait five? Conversely, if they’re going great and the question never comes up, well, terrific, more power to them.”

“Managing research is always tricky,” Myhrvold continues. The researchers “know more about their specialty than anybody else, so how are you going to measure them on that?” In such an endeavor one has to embrace a certain amount of uncertainty, and Microsoft is willing to do so. Even if the Theory Group were to succeed in its audacious quest to crack the problem of whether P equals NP , it is not clear what the payoff would be to Microsoft. “Broadly speaking, the theoretical implications of work like that could affect lots of our products,” Myhrvold says. “And we could make lots of money on it. Now, exactly how is hard to say, because we don’t have the result yet. And in any piece of research it’s possible that you won’t succeed. It’s possible we’ll succeed on a theoretical basis, but won’t succeed in transitioning to products. But that’s the risk you have to take.”

P, NP, and All That

Michael Freedman is known for being an enthusiastic outdoorsman, and at forty-seven he looks the part, tanned and wiry. Enamored of big ideas, he is not one to sweat the small stuff. While the merits of stock options versus pension funds can elicit serious discussion from Borgs and Chayes, the subject only kindles Freedman’s irreverence. “I used to think I was much more likely to be buried in an avalanche than to retire,” he says. “Mount Rainier is very near Seattle, so that’s a draw. It might reduce the risk of retiring.” Like retirement, departure from academia was something he hadn’t thought about. He is not leaving San Diego out of any dissatisfaction with the university. However, he finds the research environment at Microsoft more enticing than that in academia. At most university mathematics talks, he says, there is a “polite distance” between the audience and the speaker. The audience doesn’t expect to learn much, but only to get a flavor of what the talk is about. “It’s a little like watching television,” he says ruefully. By contrast, when he went to speak at Microsoft, many people came up afterward, immediately wanting to nail down the details. “And I found that really exciting,” he said. “They wanted to know if this was going to work or not, that day.”

It might seem strange that Microsoft would want to hire Freedman, who is best known for his solution of the four-dimensional Poincaré conjecture and his classification of simply connected four manifolds, feats that won him a Fields Medal

in 1986. In recent years his research interests have widened to include theoretical computer science and, in particular, the question of whether P equals NP . What follows is a brief description of some of the ideas that he and the other members of the Theory Group are working on.

A problem is in the class P (standing for “deterministic polynomial time”) if there exists an algorithm that solves the problem within an amount of time that is bounded by a polynomial function of the size of the problem (i.e., the number of input bits). The algorithm is assumed to run on a computer based on the Turing machine model. The class NP (standing for “nondeterministic polynomial time”) is defined as the class of problems for which it takes polynomial time to check whether a proposed solution actually is a solution. The naïve approach of trying out all possible solutions to an NP problem could take exponential time, since there are potentially many possible solutions to a typical NP problem. The question of whether or not there exist more clever algorithms that can find solutions in polynomial time—that is, whether $P=NP$ —is the most fundamental problem in theoretical computer science. In 1972 Richard Karp showed that a host of hard-to-solve problems in computer science were in NP . Around that same time, Stephen Cook showed that problems in NP are all equivalent, in the sense that an algorithm that would solve one of them could be used to solve any of them.

In 1988 Edward Witten produced a physics interpretation of the knot polynomials discovered by Vaughan Jones, leading to what is now known as Jones-Witten theory. That same year, in a development that received less notice, computer scientists connected the Jones polynomial to the world of NP problems. The connection is rather simple: Take a fifth root of unity, $e^{2\pi i/5}$, and plug it into the Jones polynomials for links to produce a function that takes links to the set of algebraic integers. As Freedman puts it, “If you had an oracle that could tell you, when you inputted the link, what that algebraic integer was, then modulo that oracle, P would equal NP .” This calculation is hard because the time required to carry it out rises exponentially with the complexity of the links.

Freedman began to wonder if there was a connection between this computer science result and Witten’s work. “I remember in the fall of 1988 being struck by this idea that if Witten is telling us that the values of the Jones polynomial are physics, and the computer scientists are telling us that the values of the Jones polynomial represent universal computation, then one should try to build a computer which measures the physics in its computational step—the computer should somehow measure, certainly not calculate, the Jones polynomial.” But he quickly found out that such a plan would be extremely difficult to carry

out. Nevertheless, he remained fascinated by the connection between physics and computation and had these ideas simmering on the “back burner” as he continued to work in topology. “But then I thought somehow it just couldn’t sit on the back burner,” he remarks. “At some point I felt that I needed to take five years and figure out whether something could be done with this.” He has been working on it ever since.

The question of whether P equals NP arises for computers based on the Turing machine model. Are there other computing architectures in which P equals NP? One possibility is known as “quantum computing”. In analogy to what happens in quantum mechanics, a quantum computer produces not one answer, but a superposition of all possible answers, with different probabilities of occurring. “That makes the game much more lively, because it means you have to be clever with your algorithms,” says Freedman. “You have to arrange for the interesting information to be reinforced and the boring answers to somehow interfere in the sense of wave mechanics.” Of course, such answers have a nonzero probability of being wrong. “From the point of view of a mathematician coming from a Platonic ideal with absolute precision, there’s a little hurdle to get over,” Freedman concedes. “That kind of small beauty mark definitely lies on quantum computation.”

Christian Borgs and Jennifer Chayes are looking at related problems but from a very different viewpoint. They became interested in computer science about three years ago when they attended a talk by Scott Kirkpatrick on work he did with Bart Selman that revealed an affinity between phase transitions in statistical mechanics and NP problems. Cook’s work in the 1970s showed that any problem in NP could be reduced to what is known as the k -satisfiability problem. Given N Boolean variables, form M clauses of length k by linking the variables within each clause with the “or” operator, and form a function F by linking the clauses with the “and” operator. The question is whether one can find values of the Boolean variables that make F a true statement. If N is large compared to M , there are many variables and few constraints, so generally it is easy to satisfy F . Conversely, if M is large compared to N , there are many constraints and few variables, so generally it is easy to show that it is impossible to satisfy F . But there is a certain value of the ratio M/N for which it is very difficult to say whether or not one can satisfy the function. Graphing M/N against the probability of solving the problem gives a curve that tends to a step function that takes values 1 and 0 as N tends to infinity.

What fascinated Borgs and Chayes was the fact that this graph resembles the phase transitions they have studied in physics. For example, if one heats water and graphs temperature against vol-

ume, one gets a graph that develops a jump discontinuity as the volume tends to infinity, with the sharp change occurring at the boiling point for water. The k -satisfiability problem has applications to problems like airline scheduling, where there are many variables and constraints to be juggled. It turns out that the most efficient scheduling scenarios occur right at the phase transition. As Borgs puts it, “economics drives it toward the phase transition.” This is an example of what is known in statistical physics as “self-organized criticality”, in which the dynamics of the system tend to drive it toward the phase transition. These critical configurations are among the most fascinating problems in statistical mechanics.

While the difficulty of solving the k -satisfiability problem is bad news for airline schedulers, it may be good news for cryptographers. Rather than avoiding the hard instances of such problems, cryptographers want to home in on them, because the sheer difficulty of solving the problems provides a means for encryption. Borgs and Chayes are working on ways to apply their understanding of critical configurations of phase transitions in physics to helping cryptographers develop new codes. It was their interactions with the cryptographers at Microsoft Research that led Borgs and Chayes to consider this new dimension on their work. The connection to cryptography “is something that we never would have realized had we not come to Microsoft,” says Chayes.

In contrast to the research of Borgs, Chayes, and Freedman, which has a distinct physics flavor, the work of Jeong Han Kim is in the areas of extremal graph theory, probabilistic methods, combinatorial optimization, and neural networks. He has worked in the area of “semi-random” methods in graph theory, which he used to settle a famous problem that had been open for sixty years. For this result he received the 1997 Fulkerson Prize, which is jointly sponsored by the AMS and the Mathematical Programming Society. Although Kim works in areas that have traditionally been close to theoretical computer science, his research is still quite distant from immediate application to something that Microsoft might sell. Asked what connection his work might have to Microsoft products, Kim describes a philosophy that might sum up the entire purpose of the Theory Group. “I can’t say that in five years I am going to invent something. But I am going to research something, and I’m going to have something....If I find the truth, then the truth eventually will be useful.”