NEWS

A Semiconductor Giant Ramps Up Its R&D

As computer chips get harder to shrink down and speed up, Intel is embracing research and a new culture of openness to show the way forward

HILLSBORO, OREGON—In the Oz-like realm of computer chip giant Intel, Gerald Marcyk is the man behind the curtain. Marcyk heads Intel's research on silicon transistors and other gizmos that make up computer chips. It's his job to ensure that these chips—and hence our desktops, laptops, and personal digital assistants—continue to get faster, smaller, and more powerful, trends that have held steady for 35 years and that customers have come to expect. In June, Marcyk reported a key step toward that goal. Intel researchers, he announced, had created the world's smallest and fastest silicon transistors, each capable of winking on and off 1.5 trillion times per second. If all goes as planned, Marcyk's blindingly fast devices will be powering computers in just 6 years.

Good news, to be sure. But the announcement was spectacular for another reason—that it took place at all. Marcyk's job doesn't normally involve much fanfare, and for years Intel brass has worked hard to keep it that way. The semiconductor giant has a reputation for eschewing long-term research in favor of product-oriented development work. And where research has been needed, Intel has preferred to keep the results quiet, unlike other corporate research powerhouses such as IBM and Lucent Technologies' Bell Labs. Now that is changing. Looking to establish a new reputation as a research leader, Intel officials have played up the tiny transistors and other recent research achievements to media and industry analysts.

The shift reflects how the complex science of semiconductors is forcing major changes on the industry giant. It's getting harder and more expensive to shrink transistors and pack ever more of them onto chips. That's compelled Intel to multiply its R&D funding nearly 10-fold since 1990, to an estimated $4.1 billion this year. In recent years the company has launched new internal labs devoted to improving the layout of devices on computer chips and developing novel software applications in an effort to sell more computers powered with Intel chips. And this spring, here in the western outskirts of Portland, the company opened its first facility dedicated to silicon research.

Intel isn't only spending more; it's talking more as well. To better compete for everything from research talent to the confidence of customers, Intel is beginning to lift the curtain on its research. It is publishing more (see graph on p. 2023) and touting the results. "Intel is trying to make its advances more public," says Chris Murray, a chemist at IBM's Thomas J. Watson Research Center in Yorktown Heights, New York. "That was not true before. Intel tended to keep its cards very close to the vest."

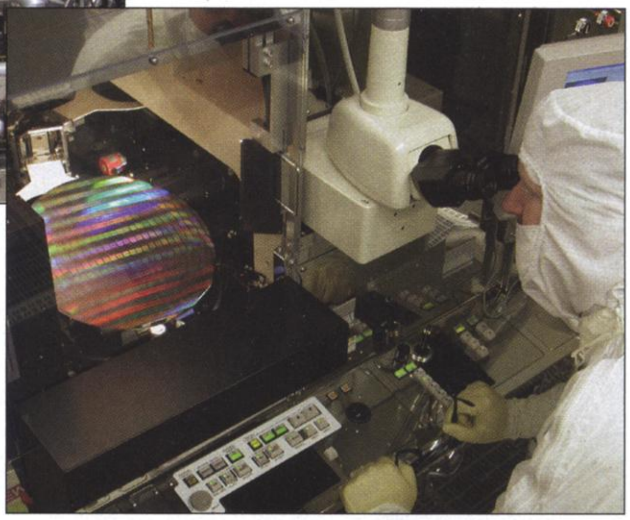

Futuristic. At Intel's new D1C plant (*above*), a sealed monorail (*right*) whisks chips through the facility.

Many computer science researchers applaud Intel's new approach. "This is vital to the industry," says Randy Isaac, who heads science and technology at IBM's Watson labs. Sandip Tiwari, who heads the Cornell Nanofabrication Facility in Ithaca, New York, agrees: "It's critical for the future. Coming to the limits [of shrinking chips] means we need to be open about research to be open to new ideas."

Growing research

Intel was anything but a research powerhouse when Gordon Moore and Bob Noyce founded the company in 1968. Before launching Intel, Moore led research at Fairchild Semiconductor, where he quickly soured on company-sponsored long-term research when he saw that many technologies invented in Fairchild's research labs failed to make it into products. At Intel, Moore and Noyce opted against building a central research lab. Instead, needed research would take place alongside teams developing better and faster computer chips, a strategy the company still follows today.

The result was a far cry from the open-ended studies of the cosmic microwave background and superconductors that won fame and Nobel prizes for Bell Labs and IBM. Intel researchers haven't netted any Nobels, but it's hard to argue that the company's strategy hasn't been a success. Today, the chipmaking giant employs 70,000 people and commands 85% of the worldwide market for microprocessors, the chips that serve as the brains of computers. Last year, the company's revenue topped $33 billion.

Ironically, in an era that has seen research at central corporate labs reined in and focused more on helping a company's bottom line, Intel is slowly moving in the other direction. The company still favors decentralized research linked to developing new products. But according to numerous company officials, long-term research is increasingly gaining favor.

Just 5 years ago, getting research money "was a struggle," says Marcyk. To a large degree that's because the path to improving chips was clear. Semiconductor companies use a technique called lithography—shining light through stencils to activate chemical etchants—to pattern devices on chips. And for decades they have managed to increase the density of these patterns by using shorter and shorter wavelengths of light, a process akin to drawing thinner and thinner lines with an ever sharper pencil. But current lithographic methods are fast approaching their endpoint, as engineers will soon run out of materials transparent to ultrashort wavelengths of light (*Science*, 3 August, p. 785).

"The industry is getting to a point where the future directions are less clear," Isaac

CREDITS: INTEL

says. That's forcing Intel to stoke its research engine and be more open about its work. "We realized a few years ago that we had to do research ourselves on new materials and processes," says Marcyk. "This is a relatively new behavior for us."

The trouble was that as chipmaking grew increasingly complex and sophisticated, Intel found itself developing a reputation for playing it safe and staying off technology's leading edge, says Manny Vara, who handles public relations concerning research for the company. Four years ago, for example, IBM announced that it was beginning to make chips with transistors connected by copper wires instead of the traditional aluminum variety. Copper conducts electrons faster than aluminum, so it has the potential to speed chips. But it can also kill semiconductor devices if not handled with extreme care. News organizations around the globe hailed IBM's effort, saying it would pave the way for chips with record-breaking speed.

Intel researchers had already evaluated the technology and concluded that it wasn't needed for the time being. Copper wires allow electrons to course through the chips like high-speed sports cars, explains Justin Rattner, who heads Intel's microprocessor research in Hillsboro. Three years ago transistors were still relatively slow. Even if electrons sped between them through copper wires, they'd still stall out at the transistors—the chip equivalent of stoplights—and reach the finish line at the same time. "We concluded the wiring delays were not the speed limitations," says Rattner. Intel officials decided that, rather than reengineer their manufacturing plants to work with copper, they would hold the technology for use in chips made with smaller and faster transistors.

"The astonishing thing was that Intel continued to deliver by far the fastest processors despite the fact we weren't using copper," says Vara. But reporters and industry analysts questioned over and over whether Intel's leadership was slipping. Now, Rattner says, the same thing is happening with silicon on insulator, a souped-up substrate on which chips sit, which is being backed by IBM, Sun Microsystems, Advanced Micro Devices, and Texas Instruments. Again, Intel researchers view the gains as meager and not worth higher manufacturing costs. "It looks like Intel is going against the grain or is antiresearch. But as is the case with copper, we are making very educated decisions," says Rattner.

To counteract this reputation, Vara and others have encouraged Intel's business leaders to unveil the work being carried out by the company's 6000 researchers in 80 labs around the globe. Vara says that they agreed reluctantly and have started slowly revealing selected results, often fuzzing the details of how they were accomplished. IBM's Isaac argues that the shift is important. As the path to improving semiconductors grows increasingly foggy, it's critical for Intel to let other companies know what its strategy is beyond the next quarter. "To get others to follow, you have to tell people what you're doing," says Isaac. "And if you can't get them to follow, you're not a leader, you're a loner."

Equally important, IBM's Murray and others say, is publicity's role in attracting new research talent and ideas to the company. It's increasingly crucial for electronics companies to forge tight relationships with university researchers, says Murray. Although Intel spends about $50 million a year to support some 300 academic projects such as quantum computing, talking up the corporation's research drives home the message that this is a place where scientists' hard work will be valued. Whether it is for enticing researchers, customers, or other industry customers to follow the company's lead, "research is becoming more of a strategic tool than it has been in the past," says Murray.

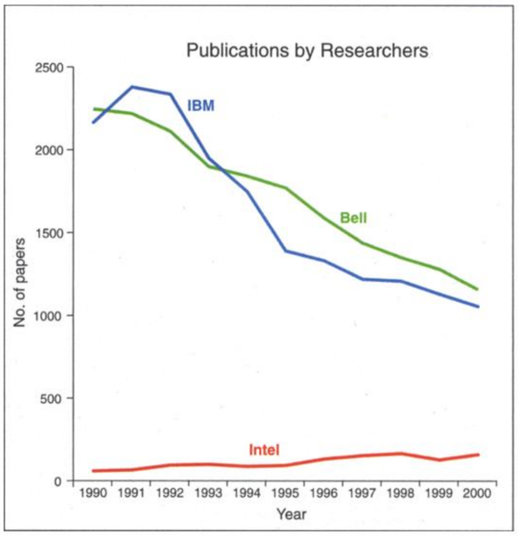

Paper chase. Intel researchers' small but growing record of publications bucks the trend of its more visible rivals.

Lifting the curtain

This year, Intel has allotted $4.1 billion to R&D—equivalent to roughly one-fifth of the budget for all biomedical research supported by the National Institutes of Health. Company officials won't specify how they split their spending between research and development. But most of the money, they say, is slated for places like D1C, the company's latest chip-development plant in Hillsboro, which is preparing next-generation chip technologies for manufacturing. Chips at this plant will bear features as small as 130 nanometers across—about half of the size of components on current chips. They will be the first from Intel laced with copper wiring and will be forged from 300-millimeter-diameter wafers instead of today's 200-millimeter ones—another industry first that should make them cheaper to produce.

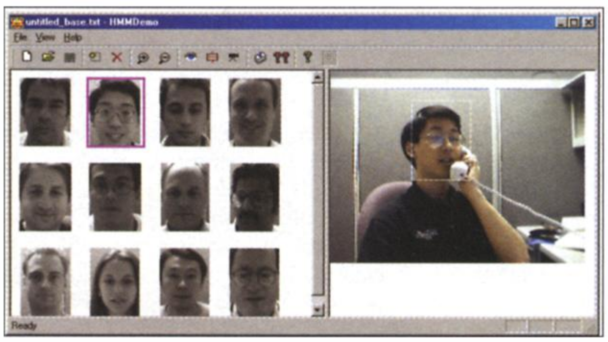

Coming soon? Intel's current projects include software designed to recognize human faces.

CREDIT: (BOTTOM) INTEL

The futuristic environment of D1C is as impressive as the chips it will turn out. Like all chip foundries, its core is a clean room dotted with stations for patterning and testing the chips. But compared with most chip clean-room facilities—which require engineers to wear special suits and to breathe into scubalike respirators to prevent even the tiniest contaminants from entering the air—D1C is relatively lax. When the chips aren't inside patterning and testing machines, they're ferried about in sealed cartridges and whisked through the plant on a robot-controlled monorail. In 2002, Intel officials expect to convert the development plant to fulltime manufacturing.

Another sizable chunk of cash goes next door to RP1, the new silicon research plant. Here, Marcyk and his crew spend their days inventing and perfecting schemes for making the generations of chips beyond those that D1C will manufacture. In addition to making tiny transistors, the RP1 crew is working to develop more potent insulators, which help confine electrons to regions where they are supposed to be on the chips. Current chips are already approaching the limits of today's insulator, silicon dioxide. So researchers are evaluating promising new compounds such as zirconium dioxide.

Marcyk's team is also looking to make high-quality lenses from calcium fluoride, a temperamental material that's transparent to photons at 157 nanometers, which makes it the leading candidate to replace current quartz optics in future lithography machines. And RP1 researchers are starting early work on replacing some of a chip's internal wires with "optical interconnects" that would use photons to speed data transfer onto and off chips. All these efforts, Marcyk says, are geared to keeping increases in computing power marching along. "It's our job to make Gordon Moore look like a genius," he says, referring to Moore's oft-quoted 1965 prediction that the number of transistors on chips would double every 2 years.

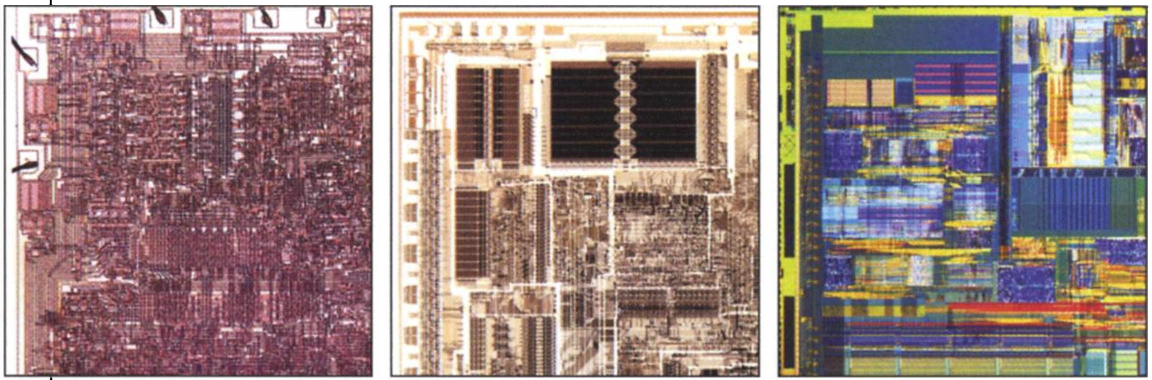

Small wonders. Silicon chips from 1971, 1982, and 2000 (the Intel 4004, 286, and Pentium 4, respectively) show the progress that researchers have made in miniaturizing electronic circuitry.

CREDITS: INTEL

Maintaining that pace also involves plenty of work beyond the confines of RP1. In a nondescript office complex just down the road, Rattner and his colleagues in the microprocessor research group are working on designs for the high-density chips that Marcyk's team is learning to build. Topping Rattner's list are microprocessors hardwired for vision and hearing. Such chips, says Rattner, will one day allow your computer to recognize you when you walk in the room, turn itself on, and fetch your latest e-mail, access the Web, or launch other programs you tell it to launch. Literally *tell*, Rattner says, because new hardware designs are also expected to vastly improve today's rudimentary voice recognition systems. "On these foundations you can build all sorts of unique and wonderful tools," such as computers that control all the appliances in your house with simple voice commands, says Rattner.

Many of those tools will likely come from the Intel Architecture Lab (IAL), also centered in the Hillsboro complex. Here, researchers are pushing the limits in an area that few realize has long been an Intel hotbed: software. According to IAL researcher Steven Spina, Intel researchers actually did the lion's share of developing components of now-standard programs, such as RealPlayer—a music and video player—and an animation program called Shockwave that has been downloaded by more than 200 million users worldwide. Spina and Vara say that although Intel spent millions on the products, it essentially gave away the licensing rights. The idea, Vara explains, is to create must-have applications that will drive demand for new computers, most of which will presumably be powered by Intel chips. "We want to continue to grow the entire computing pie," says Vara.

Spina and his IAL colleagues have several software packages on the horizon. One converts video coverage of sporting events into three-dimensional animation that can be viewed from any angle on the field, including the perspective of specific players, referees, coaches, and even the ball. Another is a set of video games that track a player's real-life motions by camera and use them to control the action on screen. More mundanely, the researchers envisage computers that start up in seconds rather than the 1 to 2 minutes that we've come to tolerate first thing in the morning. Inventing new computer applications is a job Spina and his colleagues clearly enjoy. "This is a place where we get to be around a lot of cool stuff and get to play," Spina says.

But with the semiconductor industry in the midst of a wrenching downturn, it's not clear how long the toys—and the money to buy them—will keep coming. Also unclear is how long the corporate glasnost will continue to offer outside researchers insight into the possible road ahead and a chance for Intel's own legions of trained talent to talk in detail about their latest work.

"I don't know if this is a fluctuation or a sustained trend," says Isaac. "[Intel] still is not as open as IBM or Bell Labs would be," says Cornell's Tiwari, but for a company better known as a consumer of research than as a provider, the trend is in the right direction.

–ROBERT F. SERVICE

NEWS

Better Searching Through Science

Next-generation search tools now under development will let scientists drill ever deeper into the billion-page Web

In the beginning, the Web was without form, and void. Vast heaps of information grew upon the deep, and it was good for one's desktop. But users across the land were befuddled and could not find their way. There arose the tribes of the Yahoos, the HotBots, and the AltaVistas to bring order out of chaos. Google and CiteSeer prospered and lent guidance. But researchers and scientists, learned ones who had built the Web in their own image, yearned for something more. …

As myths go, this one may lack staying power, but there is no doubt that in some sense scientists have been victims of their own success. The real creation story is that the World Wide Web began as an information-sharing and -retrieval project at the European particle physics lab CERN, and many scientists of all fields now depend vitally on the Internet to do their jobs. It's only recently that it has evolved into a convenient way to buy stuff. And although this commercial proliferation has been good for the Web's growth, it has frustrated researchers seeking quality content and pinpoint results among the noise and spam.

Now, a handful of companies and academic researchers are working on a new breed of search engines to undo this second curse of Babel. "I think the real action is in focused and specialized search engines," says Web researcher Lee Giles of Pennsylvania State University, University Park. "This is where we're going to see the most interesting work."

The first generation of search engines was based on what computer scientists such as Andrei Broder of AltaVista like to call classic information retrieval. Stick in a key word or phrase, and the software scurries around looking for matching words in documents. The more times a word pops up, the higher the document ranks in the output results.

But ranking by hits did not say how important or authoritative or useful the pages might be. "The original idea was that people would patiently look through 10 pages of results to find what they wanted," says Monika Henzinger, director of research at Google. "But we