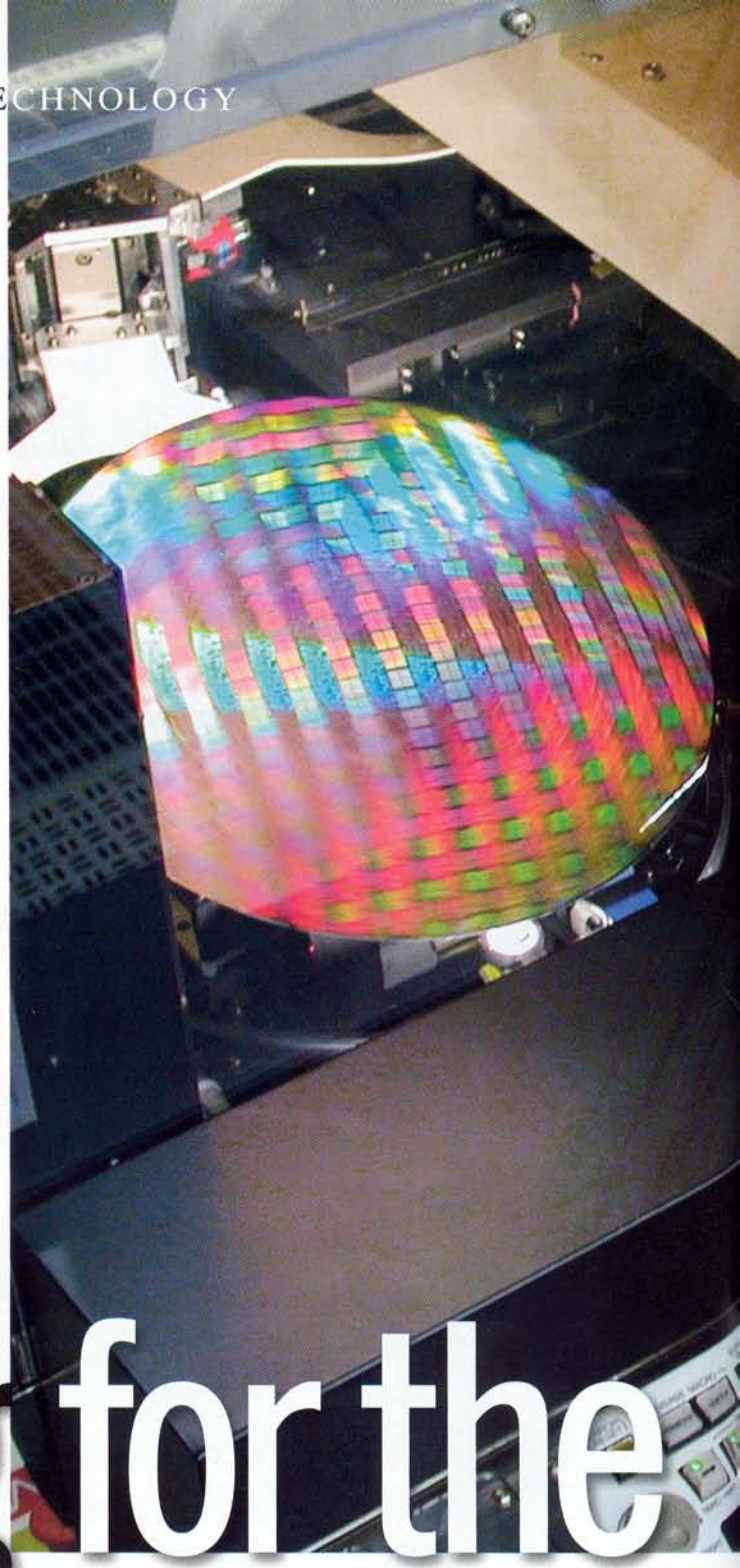

Building for the

Never mind that sales are off by 30%. Chipmakers are racing ahead with snazzy new technology.

by Stuart F. Brown

The headlong rush of semiconductor miniaturization, it seems, waits for no one. Just because chipmakers are staring at woefully thin order books doesn't mean they can stop following Moore's law, the productivity miracle that has kept the number of transistors on a chip doubling every two years or so. "The pace of technology doesn't change, even during recessions, and if you want to be in the business, you have to continue on through the downturns," says Andy Bryant, chief financial officer of industry giant Intel, whose cofounder Gordon Moore identified the law that thousands of engineers around the world are laboring to obey for another decade or more. The engineers think they can do it—good news for computer users—but for chipmakers the challenge has never been more severe.

For starters, business is in the dumps. "The chip downturn is certainly the worst in 15 years," says Linley Gwennap, principal analyst at the Linley Group, a technology-analysis firm in Mountain View, Calif. According to the Semiconductor Industry Association, in 2001 worldwide semiconductor revenues tumbled 30%, to $140 billion—by far the steepest annual decline ever.

At the same time chipmakers are struggling mightily to accommodate an unprecedented triple wave of new manufacturing technology now sweeping through the cutting-edge fabs that make the highest-performance and most profitable chips. There is the shift to dinner-plate-sized silicon wafers from salad-plate-sized ones; the replacement of aluminum with copper in the wiring that connects the millions of transistors on a chip; and the latest round of shrinkage in the size of the chip's microcircuits.

All this fancy new technology has boosted the cost of a new fab to about $2.5 billion, making it painful or impossible for all but the biggest players to keep up with the state of the art. As a result, many semiconductor manufacturers are forming alliances to spread the costs. And large merchant foundries like Taiwan Semiconductor Manufacturing Co. in Hsin-Chu, Taiwan, have become an enduring feature of the industrial landscape, producing made-to-order chips from designs supplied by smaller players that can't afford their own plants, as well as by customers in the computer, communications, and consumer-electronics industries.

The new wafers alone would have been a big bite for the fabs to chew. Nearly a foot in diameter (300 millimeters, to be precise), they instantly render obsolete the expensive equipment designed to make chips on 200mm wafers. The 300mm manufacturing technology is so expensive that it is out of reach for companies with annual sales of less than $5 billion or so. But it can give a huge producer like Intel, which makes about a billion chips a year, a surprisingly handsome return on investment and the prospect of rocketing away from competitors—whenever the recovery finally arrives.

next chip boom

New foot-wide wafers like Intel's yield more than twice as many chips as the smaller kind.

The big wafers' appeal has everything to do with geometry. As a circle increases in diameter, a smaller proportion of its area borders its perimeter, where rectangular chip dies can't fit. Thus less of a 300mm wafer's costly silicon real estate is wasted. In all, a 300mm wafer can yield 2.4 times more chips than a 200mm wafer, boosting a fab's output while cutting its cost per chip.

Fabs running 300mm wafers look even more space-age than their predecessors, if that's possible. The supersized wafers are so costly to put through the 350 to 500 steps in the manufacturing process that chipmakers have invested in a raft of new equipment to cut the defect rate. In older plants, for instance, wafers were cured in big ovens in which the temperature could sometimes be uneven, causing imperfections. The new fabs have small ovens that cure wafers one at a time instead of in batches. The big wafers are heavy enough, too, that "pods" full of them can't be handled safely and reliably by workers. So the new fabs have overhead monorail transporters that move the pods through the building and nifty little elevator systems that gently lower caddies of brittle wafers to the machines that process them.

The shift to copper circuitry is largely a product of the proliferation of portable gadgets like laptops and cellphones. Makers of such devices are always pushing for ever higher perfor-

mance and longer battery life while keeping the temperature down, a special concern in gadgets too small to contain a fan. Making circuits out of copper helps because it has less electrical resistance than aluminum, so electricity flows through it more quickly and without generating as much heat.

Learning to use copper took years of research and the development of sophisticated new processing equipment. "Copper is not duck soup. Otherwise the industry would have been using it a long time ago," says Ashok Sinha, a senior vice president at Applied Materials in Santa Clara, Calif., the largest of the fab equipment makers and one that has been in the thick of the chipmaking transformation. The $7-billion-a-year company builds most of the major types of machinery used in fabs, including systems that deposit, etch, cook, polish, and inspect the many millions of ultraprecise features on a chip.

A state-of-the-art chip is a complex microstructure, with eight or nine layers of wiring atop a base incorporating hordes of transistors. Switching to copper created tough problems whose solutions often begat other problems. For one, copper is much more difficult to etch than aluminum. The metal also tends to diffuse into adjacent silicon material and "poison" it, much as leaded gasoline can destroy a car's catalytic converter.

Chipmakers tried to solve the diffusion problem by applying an insulating layer of silicon dioxide, or glass, to the copper wires. Alas, glass has a relatively high capacitance—the tendency to soak up electrons, slowing current flow. So chipmakers switched to so-called low-K materials—insulators with lower capacitance than silicon dioxide's. But these materials are also softer than silicon dioxide, and without great care, too much of the coating can wash away during polishing, resulting in bad chips.

To help chipmakers master such quirks and get up and running quickly, Applied Materials developed so-called multimachine modules that take some of the trickiness out of the manufacturing process. "Time to market is everything, and the window of opportunity for a chip design is very small," says executive vice president David Wang. "Within

Sandia Labs' prototype prints chips using extreme-ultraviolet light.

Taking light to the limits of the extreme

Chipmakers don't have to look very far ahead to see an obstacle that could derail Moore's law. A few generations from now the teensy features on a chip will be smaller than the wavelength of the light currently used to produce them. Thus the photolithography machines, or steppers, won't be able to keep pace in the miniaturization race.

Today's steppers beam deep-ultraviolet light through sets of stencil-like masks that project precise images of the chip's intended circuitry onto silicon wafers treated with photosensitive chemicals. When line widths get down to about 50 nanometers—minuscule enough to achieve a truly awesome milestone, the billion-transistor chip—steppers will need a light source with a shorter wavelength.

A consortium of chip companies spearheaded by Intel is banking on extreme-ultraviolet light for the next leap forward. The consortium, EUV Lithography, which includes IBM, Motorola, and others, is tapping teams of bright types from traditionally rivalrous federal weapons labs in Livermore, Calif. Their aim: to overcome the technical hurdles to a commercially usable EUV stepper, which are far from trivial.

Consider this unhelpful attribute of extreme-ultraviolet light: It is completely absorbed by most materials, including air and glass. The work-around for the air problem was to enclose parts of the machine in a vacuum chamber. The glass problem posed a bigger brain teaser. Optical devices like steppers are full of glass lenses, and masks are made from transparent quartz patterned with chromium metal. The answer was to shunt around the EUV light using precisely curved mirrors with special multilayer coatings that reflect it. To beam the desired circuitry pattern onto a photosensitive wafer, the team replaced traditional transmissive masks, which would gobble up all the EUV light, with reflective ones. These are made by patterning a precise image of the chip features onto an absorptive background on an extremely flat piece of glass. The mask then projects an EUV image of the chip features onto the silicon wafer.

Developing the light source itself was the final challenge. A powerful laser beam focuses on a supersonic jet of xenon gas, which forms a plasma from which EUV light can be filtered and put to work. The hard part is getting the thing to shine brightly enough to produce 80 300mm wafers per hour, which is the pace of work on a busy fab line.

The consortium's proof-of-concept EUV stepper—called the Experimental Test Stand—has the shiny, exotic appearance of a Star Wars weapons experiment. Indeed, Charles Gwyn of Intel, the program's director, compares the task of making a chip with features smaller than 70 nanometers to printing perfectly aligned, successive images the size of a quarter down here on Earth from an orbiting space shuttle. "The devil is in the little details in making such a complex system work right," he says. Intel expects delivery of the first prototype in 2005.

four or five months, the price of a new device drops by half. The fab has got to be productive right away to enjoy a premium price for even a short while."

As in almost any type of manufacturing, the goal of a process module is to reduce variability from one product to the next, and ideally to make no bad products at all. All of the machines in a module are controlled by a single suite of software. The module also incorporates sensors that gauge dimensions and other parameters while a process is under way. Thus a machine that applies a coating to a wafer can immediately report that the wafer varies slightly in thickness from one side to the other. The module's polishing machine can then adjust its efforts in order to "planarize" the wafer to near-perfect flatness.

Things were a lot different in the early days of chipmaking, recalls Dan Hutcheson, CEO of the chip-manufacturing research firm VLSI Research in San Jose, Calif. "In the 1960s, a chipmaker bought stand-alone machines that were still pretty researchy. They figured out how to use them on their own, and they lived with the 10% to 20% yields of good chips that came along with this," he says. "Today chipmakers have to come out of the box with yields of 80%, or they're not ready for the market." That's why Applied Materials has invested several hundred million dollars in its process-integration center.

"Time to market is everything, and the window of opportunity for a chip design is very small."

The most important imperative in chip design is the drive toward miniaturization. The newest chips have circuits whose lines are a scant 0.13 microns wide, about one-thousandth of the breadth of a human hair and down from 0.18 microns in the previous generation. The appeal of tininess is simple: Finer lines mean more transistors in the same space, and the more transistors on a chip, the faster it can do its work. Intel's original Pentium chip, introduced in 1993, had three million transistors. Today's Pentium 4, which is similar in size, crams in 52 million transistors. That's why Moore's law is such a big deal.

Each successive generation of feature shrink brings new puzzles. The elements on the new 0.13 micron chips, for example, are so tiny that insulating and conductive layers are only 25 atoms thick—fiendishly difficult to apply with consistency. Applied Materials developed new control software and feedback sensors to ensure that any wafer drifting off the specs is detected and fixed right away.

So far the industry has shifted less than a tenth of its capacity to 0.13 microns. Yet the next step is already in the offing. This fall both Intel and Taiwan's TSMC will start rolling out 0.09 micron chips. Intel's version, a microprocessor code-named Prescott, will contain more than 100 million transistors. The new chips will be the first to get to wear the sexy prefix "nano" (as in nanotechnology, the science that will one day build micromachines atom by atom), because their line widths translate to 90 nanometers. The R&D people at the chip companies and their equipment suppliers are busy working on generations of shrinks beyond 90 nm too (see box), which will bring plenty more engineering excitement.

Or so the high priests of chip design believe. There are still financial questions as immense as the transistors are tiny. Even if technology can keep pace with Moore's law for another decade or more, will chipmakers be able to afford the new equipment they'd have to buy every two years? Will enough customers have any use for the ultra-mega-turbo chips the fabs would spit out?

Andy Bryant at Intel remembers running a study for Gordon Moore years ago when the company was debating whether to replace its 286 microprocessor with the more powerful 386 model. The study concluded that, based on the computer applications of the day, the upgrade wouldn't generate a sufficient return on investment. "Gordon said, 'I agree with everything you say, and I don't see why people need that much power,' " recalls Bryant. " 'But this is what we do, and if the software comes, maybe it will work.' And of course it did." Now that he's CFO, Bryant still questions the financial sense of rendering Intel's product line obsolete every two years. So far, the answer has always been the same: The faster, the better. F

FEEDBACK *sbrown@fortunemail.com. Stories from the Industrial Management & Technology section can be found at fortune.com/imt. Executives in manufacturing and research and others eligible to receive* FORTUNE*'s Industrial Edition can subscribe by calling 888-394-5472.*

THE NAM / FORTUNE MANUFACTURING INDEX

Still feeling confident

Manufacturers remained upbeat in the second quarter, though cautious in their spending and hiring plans. In a survey of 265 companies by the National Association of Manufacturers and FORTUNE, 82% of large firms said they were optimistic, up from 80% three months earlier. Among smaller manufacturers (employing fewer than 1,000) the upbeat mood eased slightly.

Sales prospects look better but not dazzling. Large manufacturers expecting sales to increase up to 5% over the next 12 months outnumber, by 55% to 39%, those betting on a bigger improvement. Still, the percentage of big firms planning to boost capital spending more than 5% has jumped from 7% to 25% since the third quarter of 2001.

Hiring plans suggest employment may improve in the second half of this year. Only 21% of large manufacturers expect a drop in payrolls over the next 12 months, while the rest plan to take on more workers. Most smaller firms also plan to hire.

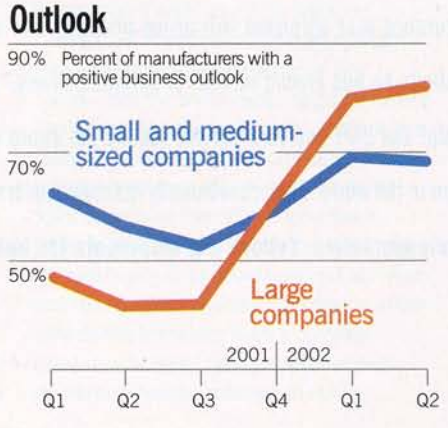

FORTUNE CHART