# the first Nanochips

**As scientists and engineers continue to push back the limits of chipmaking technology, they have quietly entered into the nanometer realm**

By G. Dan Hutcheson

## For most people,

the notion of harnessing nanotechnology for electronic circuitry suggests something wildly futuristic. In fact, if you have used a personal computer made in the past few years, your work was most likely processed by semiconductors built with nanometer-scale features. These immensely sophisticated microchips—or rather, nanochips—are now manufactured by the millions, yet the scientists and engineers responsible for their development receive little recognition. You might say that these people are the Rodney Dangerfields of nanotechnology. So here I would like to trumpet their accomplishments and explain how their efforts have maintained the steady advance in circuit performance to which consumers have grown accustomed.

The recent strides are certainly impressive, but, you might ask, is semiconductor manufacture really nanotechnology? Indeed it is. After all, the most widely accepted definition of that word applies to something with dimensions smaller than 100 nanometers, and the first transistor gates under this mark went into production in 2000. Integrated circuits coming to market now have gates that are a scant 50 nanometers wide. That's 50 billionths of a meter, about a thousandth the width of a human hair.

Having such minuscule components conveniently allows one to stuff a lot into a compact package, but saving space per se is not the impetus behind the push for extreme miniaturization. The reason to make things small is that it lowers the

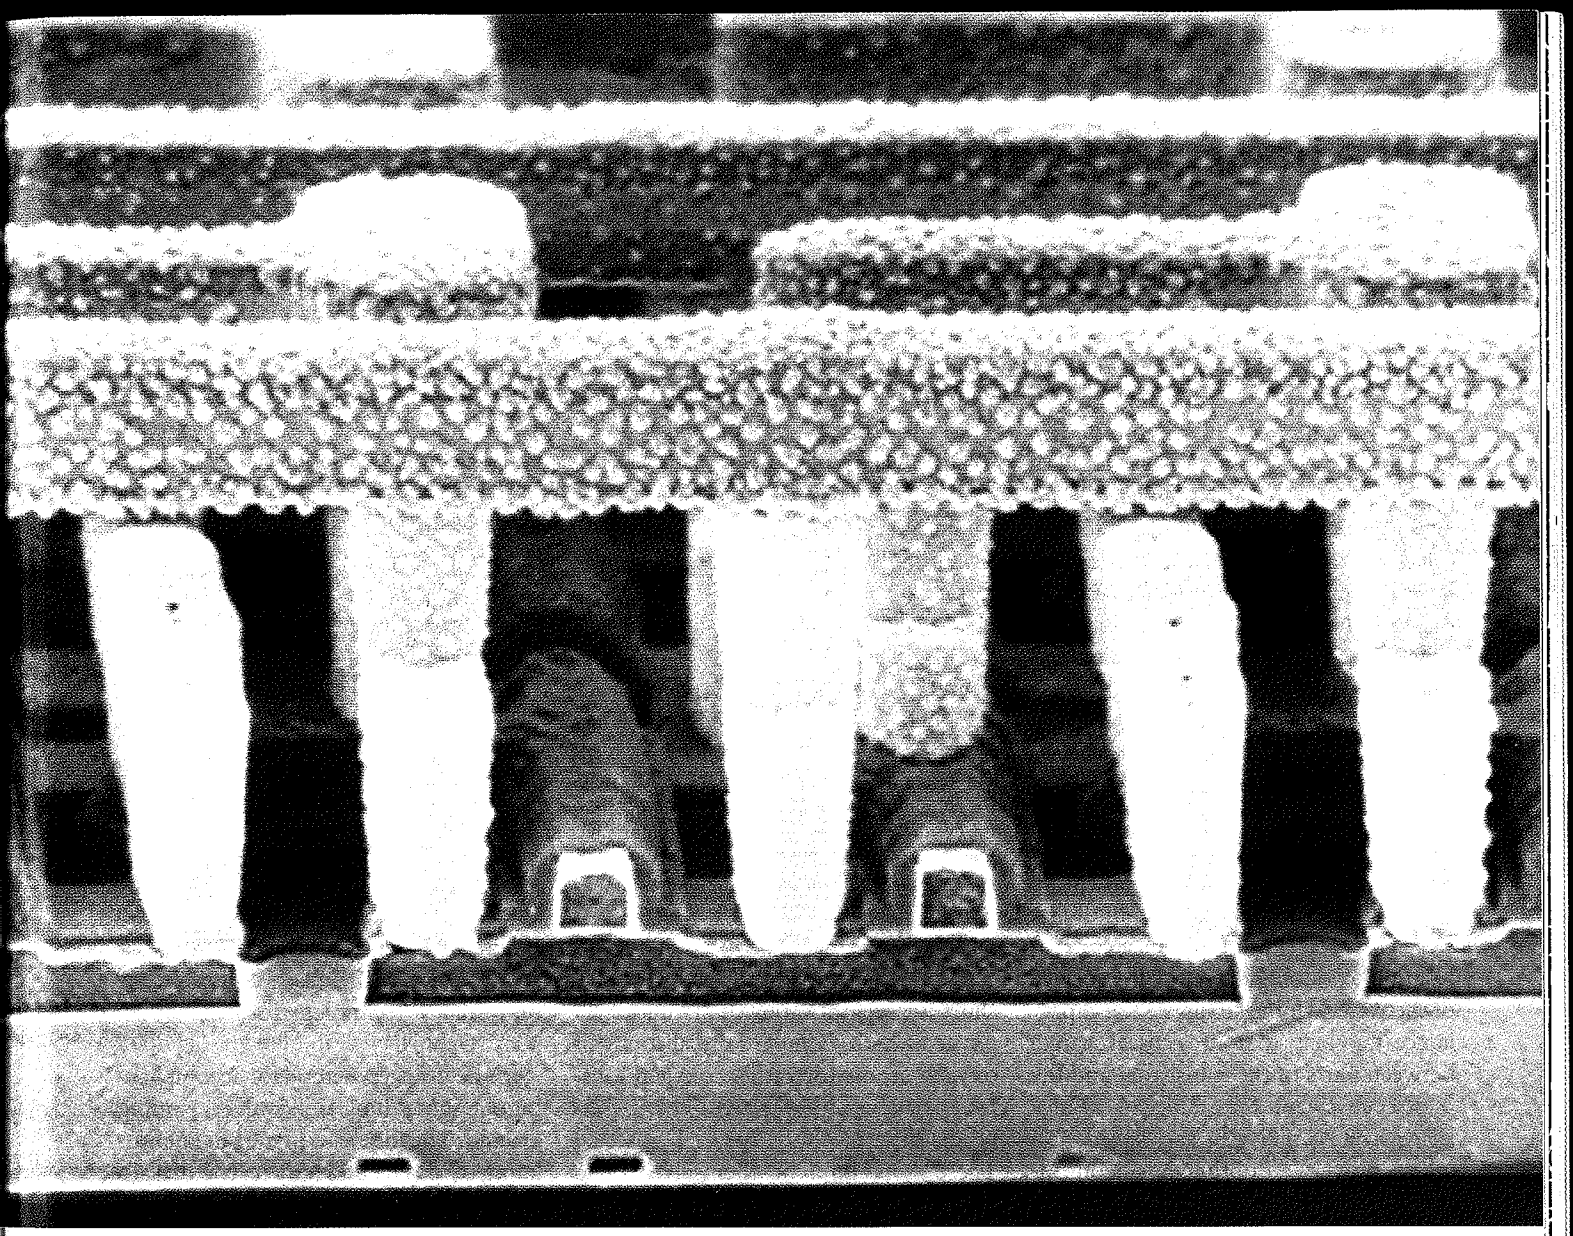

INNOVATION BOOSTS the performance of microprocessors, providing techniques that go beyond shrinking the size of transistors. The chip shown here, enlarged some 50,000 times, improves speed and saves power by placing silicon for the transistors (*light blue*) above a layer of oxide (*green*).

unit cost for each transistor. As a bonus, this overall miniaturization shrinks the size of the gates, which are the parts of the transistors that switch between blocking electric current and allowing it to pass. The more narrow the gates, the faster the transistors can turn on and off, thereby raising the speed limits for the circuits using them. So as microprocessors gain more transistors, they also gain more speed.

The desire for boosting the number of transistors on a chip and for running it faster explains why the semiconductor industry, just as it crossed into the new millennium, shifted from manufacturing microchips to making nanochips. How it quietly passed this milestone, and how it continues to advance, is an amazing story of people overcoming some of the greatest engineering challenges of our time—challenges every bit as formidable as those encountered in building the first atomic bomb or sending a person to the moon.

## Straining to Accelerate

THE BEST WAY to get a flavor for the technical innovations that helped to usher in the current era of nanochips is to survey improvements that have been made in each of the stages required to manufacture a modern semiconductor—say, the microprocessor that powers the computer on which I typed this text. That chip, a Pentium 4, contains some 42 million transistors intricately wired together. How in the world was this marvel of engineering constructed? Let us survey the steps.

Before the chipmaking process even begins, one needs to obtain a large crystal of pure silicon. The traditional method for doing so is to grow it from a small seed crystal that is immersed in a batch of molten silicon. This process yields a cylindrical ingot—a massive gem-quality crystal—from which many thin wafers are then cut.

It turns out that such single-crystal ingots are no longer good enough for the job: they have too many "defects," dislocations in the atomic lattice that hamper the silicon's ability to conduct and otherwise cause trouble during chip manufacture. So chipmakers now routinely deposit a thin, defect-free layer of single-crystal silicon on top of each wafer by exposing it to a gas containing silicon. This technique improves the speed of the transistors, but engineers have been pushing hard to do even better using something called silicon-on-insulator technology, which involves putting a thin layer of insulating oxide slightly below the surface of the wafer. Doing so lowers the capacitance (the ability to store electrical charge) between parts of the transistors and the underlying silicon substrate, capacitance that would otherwise sap speed and waste power. Adopting a silicon-on-insulator geometry can boost the rate at which the transistors can be made to switch on and off (or, alternatively, reduce the power needed) by up to 30 percent. The gain is equivalent to what one gets in moving one generation ahead in feature size.

IBM pioneered this technology and has been selling integrated circuits made with it for the past five years. The process IBM developed, dubbed SIMOX, short for *s*eparation by *im*plantation of *ox*ygen, was to bombard the silicon with oxygen atoms (or rather, oxygen ions, which have electrical charge and can thus be readily accelerated to high speeds). These ions implant themselves deep down, relatively speaking, where they combine with atoms in the wafer and form a layer of silicon dioxide. One difficulty with this approach is that the passage of oxygen ions through the silicon creates many defects, so the surface has to be carefully heated afterward to mend disruptions to the crystal lattice. The greater problem is that oxygen implantation is inherently slow, which makes it costly. Hence, IBM reserved its silicon-on-insulator technology for its most expensive chips.

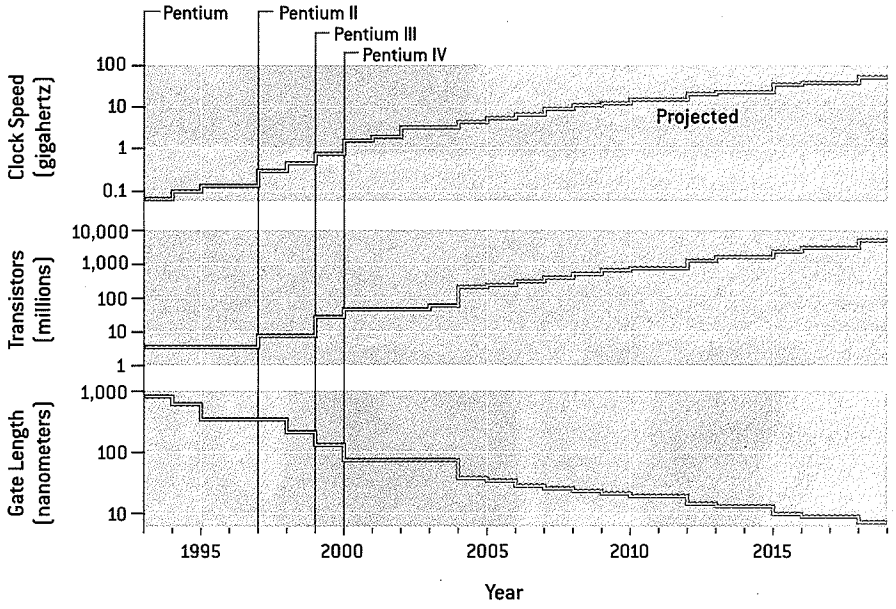


MICROPROCESSOR components have entered the nano realm during the past decade, as illustrated by the evolution of Intel's Pentium series (*blue*), which shows remarkable gains in the speed and quantity of transistors, both of which rise as the gate length of the transistors diminishes. If the semiconductor industry even comes close to matching its forecasts (*yellow*), these trends should continue.

A new, faster method for accomplishing the same thing is, however, gaining ground. The idea is to first form an insulating oxide layer directly on top of a silicon wafer. One then flips the oxidized surface over and attaches it onto another, untreated wafer. After cleverly pruning off most of the silicon above the oxide layer, one ends up with the desired arrangement: a thin stratum of silicon on top of the insulating oxide layer on top of a bulk piece of silicon, which just provides physical support.

The key was in developing a precision slicing method. The French company that did so, Soitec, aptly trademarked the name Smart Cut for this technique, which requires shooting hydrogen ions through the oxidized surface of the first wafer so that they implant themselves at a prescribed depth within the underlying silicon. (Implanting hydrogen can be done more rapidly than implanting oxygen, making this process relatively inexpensive.) Because the hydrogen ions do most of their damage right where they stop, they produce a level within the silicon that is quite fragile. So after flipping this

## Overview/*Smaller, Faster, Better Chips*

- In 2000 the semiconductor industry quietly began producing "nanochips"—chips with features measuring less than 100 nanometers (roughly one thousandth the thickness of a human hair). These devices are found in the average desktop computer today.
- Reducing the size of features boosts speed and improves the economics of manufacture by allowing more transistors (often more than 50 million) to be put on a single chip. In just a few years, a typical microprocessor will contain about 10 times that number.
- Significant gains have accrued from several new forms of technology, including improved materials and methods to correct for distortions that occur from optical diffraction when patterning the chips.

LUCY READING; SOURCES: INTEL AND INTERNATIONAL TECHNOLOGY ROADMAP FOR SEMICONDUCTORS

INTEL (*micrograph*); BRYAN CHRISTIE DESIGN (*illustration*)

treated wafer over and attaching it to a wafer of bulk silicon, one can readily cleave the top off at the weakened plane. Any residual roughness in the surface can be easily polished smooth. Even IBM now employs Smart Cut for making some of its high-performance chips, and AMD (Advanced Micro Devices in Sunnyvale, Calif.) will use it in its upcoming generation of microprocessors.

The never-ending push to boost the switching speed of transistors has also brought another very basic change to the foundations of chip manufacture, something called strained silicon. It turns out that forcing the crystal lattice of silicon to stretch slightly (by about 1 percent) increases the mobility of electrons passing through it considerably, which in turn allows the transistors built on it to operate faster. Chipmakers induce strain in silicon by bonding it to another crystalline material—in this case, a silicon-germanium blend—for which the lattice spacing is greater. Although the technical details of how this strategy is being employed remain closely held, it is well known that many manufacturers are adopting this approach. Intel, for example, is using strained silicon in an advanced version of its Pentium 4 processor called Prescott, which began selling late last year.

## Honey, I Shrunk the Features

ADVANCES IN the engineering of the silicon substrate are only part of the story: the design of the transistors constructed atop the silicon has also improved tremendously in recent years. One of the first steps in the fabrication of transistors on a digital chip is growing a thin layer of silicon dioxide on the surface of a wafer, which is done by exposing it to oxygen and water vapor, allowing the silicon, in a sense, to rust (oxidize). But unlike what happens to the steel body of an old car, the oxide does not crumble away from the surface. Instead it clings firmly, and oxygen atoms required for further oxidization must diffuse through the oxide coating to reach fresh silicon underneath. The regularity of this diffusion provides chipmakers with a way to control the thickness of the oxide layers they create.

For example, the thin oxide layers required to insulate the gates of today's tiny transistors can be made by allowing oxygen to diffuse for only a short time. The problem is that the gate oxide, which in modern chips is just several atoms thick, is becoming too slim to lay down reliably. One fix, of course, is to make this layer thicker. The rub here is that as the thickness of the oxide increases, the capacitance of the gate decreases. You might ask: Isn't that a good thing? Isn't capacitance bad? Often capacitance is indeed something to be avoided, but the gate of a transistor operates by inducing electrical charge in the silicon below it, which provides a channel for current to flow. If the capacitance of the gate is too low, not enough charge will be present in this channel for it to conduct.

The solution is to use something other than the usual silicon dioxide to insulate the gate. In particular, semiconductor manufacturers have been looking hard at what are known as high-K (high-dielectric-constant) materials, such as hafnium oxide and strontium titanate, ones that allow the oxide layer to be made thicker, and thus more robust, without compromising the ability of the gate to act as a tiny electrical switch.

Placing a high-K insulator on top of silicon is, however, not nearly as straightforward as just allowing it to oxidize. The task is best accomplished with a technique called atomic-layer deposition, which employs a gas made of small molecules that naturally stick to the surface but do not bond to one another. A single-molecule-thick film can be laid down simply by exposing the wafer to this gas long enough so that every spot becomes covered. Treatment with a second gas, one that reacts with the first to form the material in the coating, creates the molecule-thin veneer. Repeated applications of these two gases, one after the next, deposit layer over layer of this substance until the desired thickness is built up.

After the gate insulator is put in place, parts of it must be selectively removed to achieve the appropriate pattern on the wafer. The procedure for doing so (lithography) constitutes a key part of the

## FIELD-EFFECT TRANSISTOR

**THE FUNDAMENTAL BUILDING BLOCK of a microprocessor is the field-effect transistor, which acts as a simple switch. The proper voltage applied to the gate electrode induces charge along the channel, which then carries current between the source and the drain, turning the switch on. With sufficiently small gates, these transistors can switch on and off billions of times each second.**

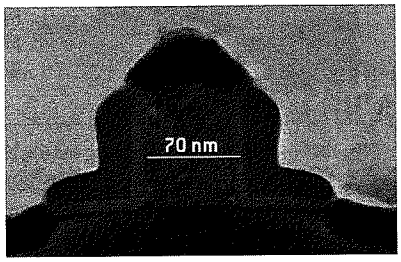


FIRST-GENERATION nanometer-scale transistor, produced by Intel.

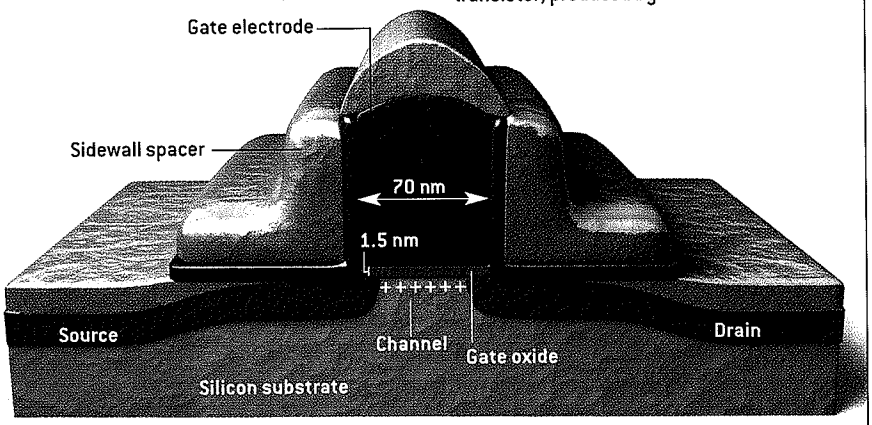


INTEL (micrograph); BRYAN CHRISTIE DESIGN (illustration)

# BASIC CHIPMAKING PROCESS

A CIRCULAR WAFER of silicon about the size of a dinner plate provides the starting point for the step-by-step chipmaking process, which sculpts transistors and their interconnections. Some of the manipulations shown below are repeated many times in the course of production, to build complex structures one layer at a time.

## BASIC CHIPMAKING PROCESS

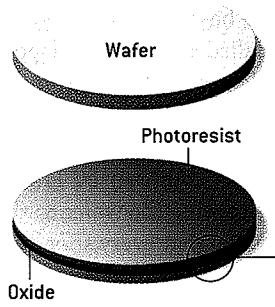


1 Steam oxidizes surface (*red layer*)

2 Photoresist (*dark blue layer*) coats oxidized wafer

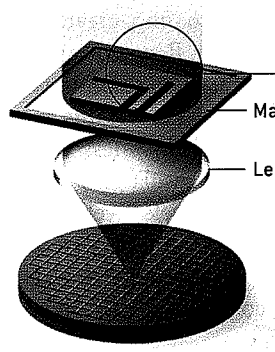


3 Lithography transfers desired pattern from mask to wafer

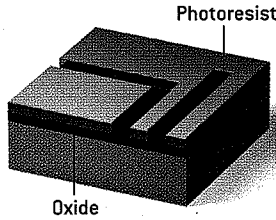


4 Chemicals and baking harden unexposed photoresist. Other parts of photoresist are removed

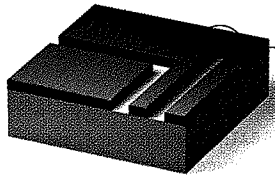

5 Chemical etching selectively strips off the oxide where no photoresist protects it. The rest of the photoresist is removed.

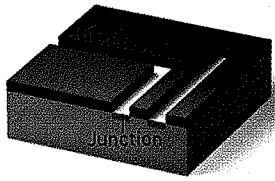


6 Ions shower etched areas, forming source and drain junctions

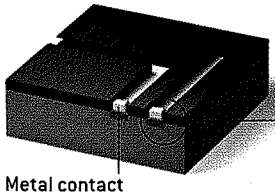


7 Metal contacts are added using lithography during later stages of fabrication

## REFINEMENTS IN CHIPMAKING

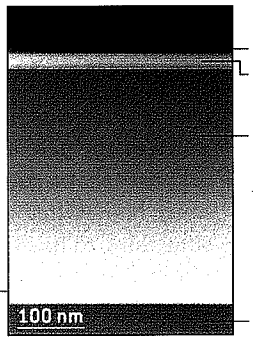


PERFORMANCE HAS IMPROVED with the growing use of wafers having a buried oxide layer or those fashioned to have a thin layer of strained silicon at the top—or by employing both techniques at once, as shown here.

**Pattern on mask**

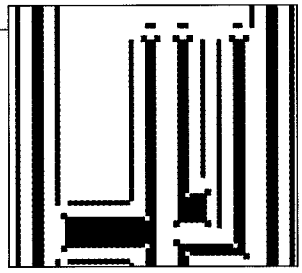

**Pattern on wafer**

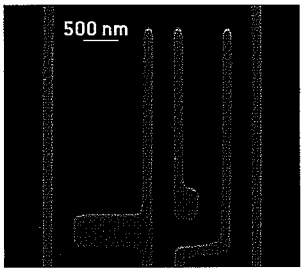


IF MODERN TECHNIQUES such as "optical proximity correction" are applied to compensate for the blurring effects of diffraction, photolithography can create features smaller than the wavelength of light used in projecting the pattern. In this example of optical proximity correction, a complicated pattern used for the mask (*left*) results in crisp features on the chip (*right*).

**Before cleaning**

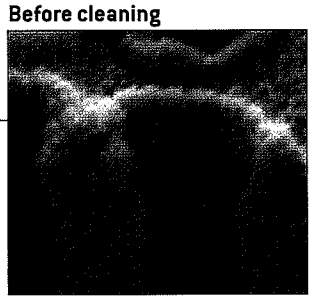

**After cleaning**

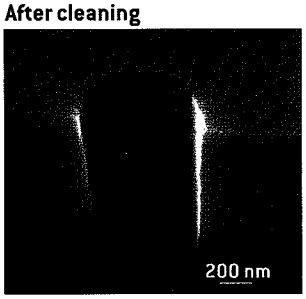


AS FEATURE SIZES SHRINK, removing photoresist and residues that remain after etching (*left*) becomes difficult. But supercritical carbon dioxide can penetrate tiny openings and dislodge particles without leaving traces of cleaning fluid behind (*right*).

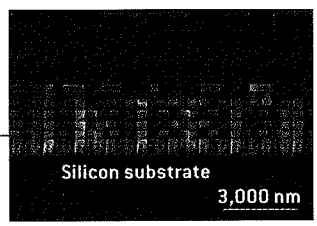


AS MANY AS EIGHT levels of wiring now connect the millions of transistors found on a typical microprocessor. Aluminum, the metal long used for this purpose, has given way to copper, which is more difficult to emplace but improves the speed and integrity of the signals carried on the wires.

BRYAN CHRISTIE DESIGN (*illustrations*); GIANNI TARASCHI *Massachusetts Institute of Technology* (*top*); IMAGES COURTESY OF ASML MASKTOOLS (*top middle*); IMAGES COURTESY OF SC FLUIDS, INC. (*bottom middle*); COURTESY OF INTERNATIONAL BUSINESS MACHINES CORPORATION. UNAUTHORIZED USE NOT PERMITTED (*bottom*)

BRYAN CHRISTIE DESIGN

technology needed to create transistors and their interconnections. Semiconductor lithography employs a photographic mask to generate a pattern of light and shadows, which is projected on a wafer after it is coated with a light-sensitive substance called photoresist. Chemical processing and baking harden the unexposed photoresist, which protects those places in shadow from later stages of chemical etching.

Practitioners once believed it impossible to use lithography to define features smaller than the wavelength of light employed, but for a few years now, 70-nanometer features have been routinely made using ultraviolet light with a wavelength of 248 nanometers. To accomplish this magic, lithography had to undergo some dramatic changes. The tools brought to bear have complicated names—optical proximity correction, phase-shifting masks, excimer lasers—but the idea behind them is simple, at least in principle. When the size of the features is smaller than the wavelength of the light, the distortions, which arise through optical diffraction, can be readily calculated and corrected for. That is, one can figure out an arrangement for that mask that, after diffraction takes place, yields the desired pattern on the silicon. For example, suppose a rectangle is needed. If the mask held a plain rectangular shape, diffraction would severely round the four corners projected on the silicon. If, however, the pattern on the mask were designed to look more like a dog bone, the result would better approximate a rectangle with sharp corners.

This general strategy now allows transistors with 50-nanometer features to be produced using light with a wavelength of 193 nanometers. But one can push these diffraction-correction techniques only so far, which is why investigators are trying to develop the means for higher-resolution patterning. The most promising approach employs lithography, but with light of much shorter wavelength—what astronomers would call "soft" x-rays or, to keep with the preferred term in the semiconductor industry, extreme ultraviolet.

Semiconductor manufacturers face daunting challenges as they move to extreme ultraviolet lithography, which reduces the wavelengths (and thus the size of the features that can be printed) by an order of magnitude. The prototype systems built so far are configured for a 13-nanometer wavelength. They are truly marvels of engineering—on both macroscales and nanoscales.

Take, for instance, the equipment needed to project images onto wafers. Because all materials absorb strongly at extreme ultraviolet wavelengths, these cameras cannot employ lenses, which would be essentially opaque. Instead the projectors must use rather sophisticated mirrors. For the same reason, the masks must be quite different from the glass screens used in conventional lithography. Extreme ultraviolet work demands masks

## SLICING A NANOCHIP

SILICON-ON-INSULATOR technology, which has helped improve chip performance considerably, has become cheaper and easier to adopt, thanks to a technique called Smart Cut, developed by Soitec, a French company.

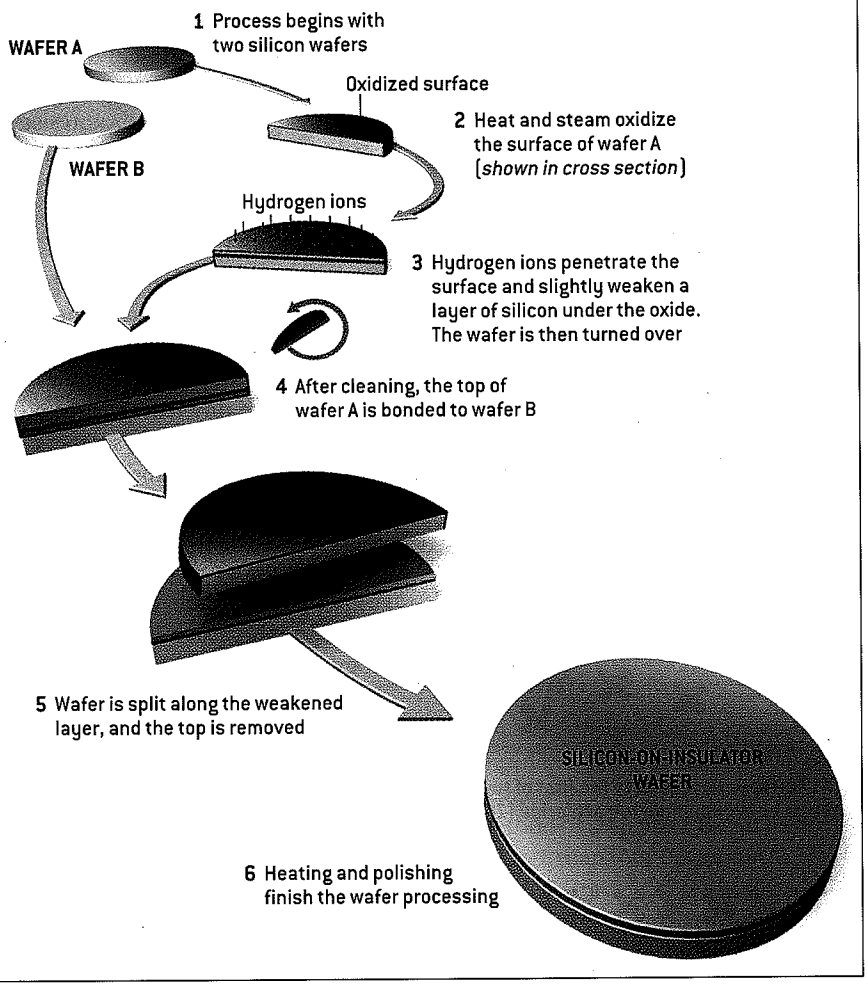



THE AUTHOR

*G. DAN HUTCHESON* is chief executive officer and president of VLSI Research Inc., a market research and economic analysis firm serving the semiconductor industry. Hutcheson, who holds a master's degree in economics from San Jose State University, has constructed various quantitative models that chipmakers can use to forecast costs and to guide them in procuring equipment. As an industry analyst, he follows the emerging technologies of semiconductor manufacture and provides summaries of the latest research advances and manufacturing trends to interested companies.


BRYAN CHRISTIE DESIGN

that absorb and reflect light. To construct them, dozens of layers of molybdenum and silicon are laid down, each just a few nanometers thick. Doing so produces a highly reflective surface onto which a patterned layer of chromium is applied to absorb light in just the appropriate places.

As with other aspects of chipmaking, these masks must be free from imperfections. But because the wavelengths are so small, probing for defects proves a considerable challenge. Scientists and engineers from industry, academe and government laboratories from across the U.S. and Europe are collaboratively seeking solutions to this and other technical hurdles that must be overcome before extreme ultraviolet lithography becomes practical. But for the time being, chipmakers must accept the limits of conventional lithography and maintain feature sizes of at least 50 nanometers or so.

Using lithography to imprint such features on a film of photoresist is only the first in a series of manipulations used to sculpt the wafer below. Process engineers must also figure out how to remove the exposed parts of the photoresist and to etch the material that is uncovered in ways that do not eat into adjacent areas. And one must be able to wash off the photoresist and the residues left over after etching—a mundane task that becomes rather complicated as the size of the features shrinks.

The problem is that, seen at the nanometer level, the tiny features put on the chip resemble tall, thin skyscrapers, separated by narrow chasms. At this scale, traditional cleaning fluids act as viscous tidal waves and could easily cause things to topple. Even if that catastrophe can be avoided, these liquids have a troubling tendency to get stuck in the nanotechnology canyons.

An ingenious solution to this problem emerged during the 1990s from work done at Los Alamos National Laboratory: supercritical fluids. The basic idea is to use carbon dioxide at elevated pressure and temperature, enough to put it above its so-called critical point. Under these conditions, $CO_2$ looks something like a liquid but retains an important property of a gas—the lack of viscosity. Supercritical carbon dioxide thus flows easily under particles and can mechanically dislodge them more effectively than can any wet chemical. (It is no coincidence that supercritical carbon dioxide has recently become a popular means to dry-clean clothes.) And mixed with the proper co-solvents, supercritical carbon dioxide can be quite effective in dissolving photoresist. What is more, once the cleaning is done, supercritical fluids are easy to remove: lowering the pressure—say, to atmospheric levels—causes them to evaporate away as a normal gas.

With the wafer cleaned and dried in this way, it is ready for the next step: adding the junctions of the transistors—tubs on either side of the gate that serve as the current "source" and "drain." Such junctions are made by infusing the silicon with trace elements that transform it from a semiconductor to a conductor. The usual tactic is to fire arsenic or boron ions into the surface of the silicon using a device called an ion implanter. Once emplaced, these ions must be "activated," that is, given the energy they need to incorporate themselves into the crystal lattice. Activation requires heating the silicon, which often has the unfortunate consequence of causing the arsenic and boron to diffuse downward.

To limit this unwanted side effect, the temperature must be raised quickly enough that only a thin layer on top heats up. Restricting the heating in this way ensures that the surface will cool rapidly on its own. Today's systems ramp up and down by thousands of degrees a second. Still, the arsenic and boron atoms diffuse too much for comfort, making the junctions thicker than desired for optimum speed. A remedy is, however, on the drawing board—laser thermal processing, which can vary the temperature at a rate of up to five *billion* degrees a second. This technology, which should soon break out of the lab and onto the factory floor, holds the promise of preventing virtually all diffusion and yielding extremely shallow junctions.

Once the transistors are completed, millions of capacitors are often added to make dynamic random-access memory, or DRAM. The capacitors used for DRAM have lately become so small that manufacturing engineers are experiencing the same kinds of problems they encounter in fashioning transistor gates. Indeed, here the problems are even more urgent, and the answer, again, appears to be atomic-layer deposition, which was adopted for the production of the latest generation of DRAM chips.

## New Meets Old

ATOMIC-LAYER DEPOSITION can also help in the next phase of chip manufacture, hooking everything together. The procedure is to first lay down an insulating layer of glass on which a pattern of lines is printed and etched. The grooves are then filled with metal to form the

ATOMIC-LAYER DEPOSITION allows chipmakers to lay down coatings that are extremely thin. Cycling through these steps repeatedly builds up the coating—one molecule of thickness at a time.

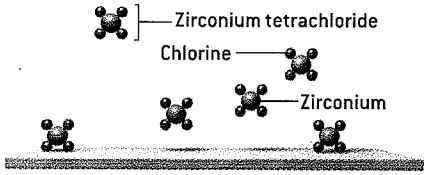


1 The surface is exposed to the first of two gases, here zirconium tetrachloride ($ZrCl_4$).

2 Molecules of $ZrCl_4$ adhere to the surface but not to one another.

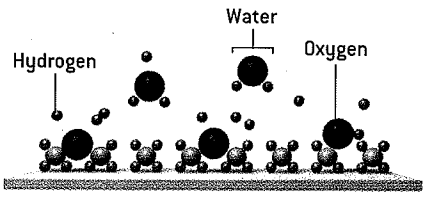


3 The coated surface is exposed to a second gas, in this case steam ($H_2O$).

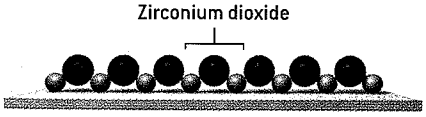


4 The $ZrCl_4$ on the surface reacts with the water ($H_2O$) to form a single-molecule-thick veneer of the desired material, zirconium dioxide ($ZrO_2$).

BRYAN CHRISTIE DESIGN

BRYAN CHRISTIE DESIGN

## EXTREME ULTRAVIOLET LITHOGRAPHY

LENSES, which are used in conventional lithography systems, would absorb the extreme ultraviolet light required for patterning features smaller than 50 nanometers. As a result, lithography systems may soon use multilayer mirrors instead of lenses to focus extreme ultraviolet radiation from a plasma and to reduce the size of the image projected from the mask. This illustration is based on one of the design concepts under consideration by the Dutch manufacturer ASML.

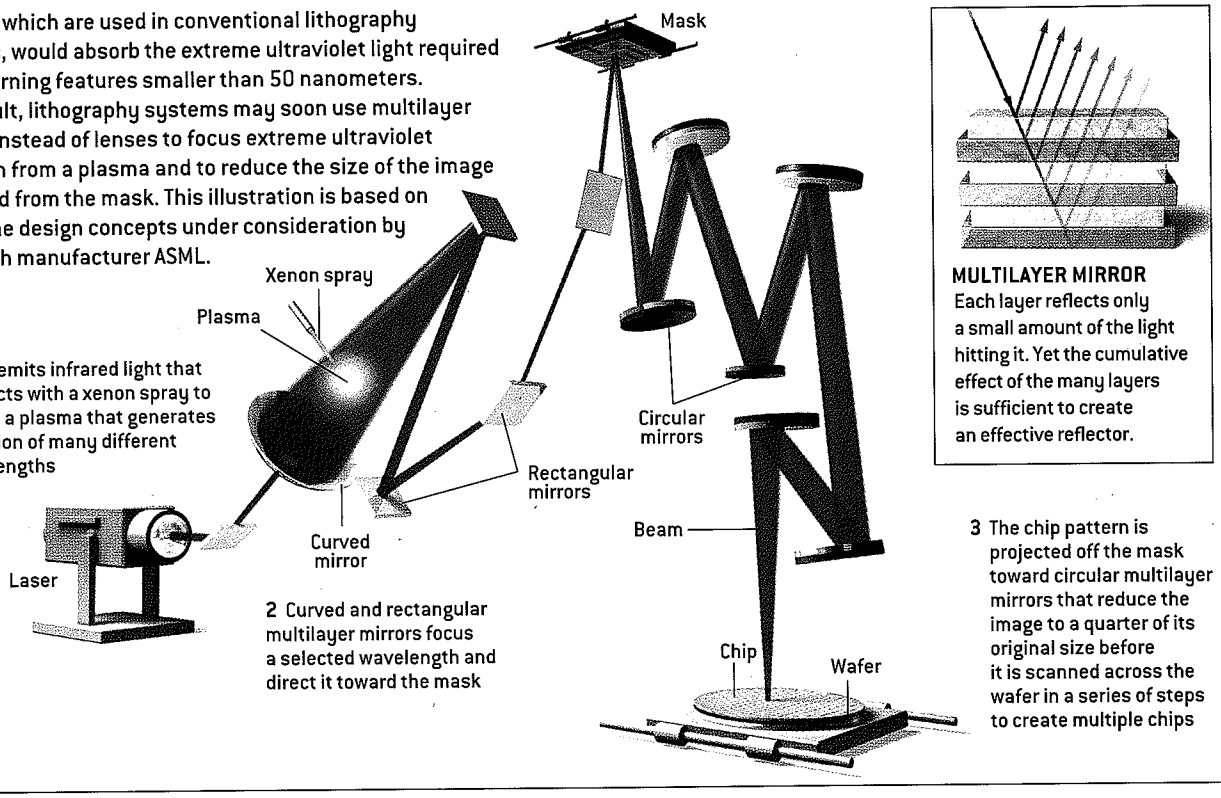




**MULTILAYER MIRROR**
Each layer reflects only a small amount of the light hitting it. Yet the cumulative effect of the many layers is sufficient to create an effective reflector.

wires. These steps are repeated to create six to eight layers of crisscrossing interconnections. Although the semiconductor industry has traditionally used aluminum for this bevy of wires, in recent years it has shifted to copper, which allows the chips to operate faster and helps to maintain signal integrity. The problem is that copper contaminates the junctions, so a thin conductive barrier (one that does not slow the chip down) needs to be placed below it. The solution was atomic-layer deposition.

The switch to copper also proved challenging for another reason: laying down copper is inherently tricky. Many high-tech approaches were attempted, but none worked well. Then, out of frustration, engineers at IBM tried an old-fashioned method: electroplating, which leaves an uneven surface and has to be followed with mechanical polishing. At the time, the thought of polishing a wafer—that is, introducing an abrasive grit—was anathema to managers in this industry, which is downright obsessed with cleanliness. Hence, the engineers who originally experimented with this approach at IBM did so without seeking permission from their supervisor. They were delighted to discover that the polishing made the wafer more amenable to lithographic patterning (because the projection equipment has a limited depth of focus), that it removed troublesome defects from the surface and that it made it easier to deposit films for subsequent processing steps.

The lesson to be learned here is that seemingly antiquated methods can be just as valuable as cutting-edge techniques. Indeed, the semiconductor industry has benefited a great deal in recent years from combinations of old and new. That it has advanced as far as it has is a testament to the ingenious ability of countless scientists and engineers to continually refine the basic method of chip manufacture, which is now more than four decades old.

Will the procedures used for fabricating electronic devices four decades down the road look anything like those currently employed? Although some futurists would argue that exotic forms of nanotechnology will revolutionize electronics by midcentury, I'm betting that the semiconductor industry remains pretty much intact, having by then carried out another dazzling series of incremental technical advances, ones that are today beyond anyone's imagination. SA

### MORE TO EXPLORE


**Toward Point One.** Gary Stix in *Scientific American*, Vol. 272, No. 2, pages 90–95; February 1995.

**Technology and Economics in the Semiconductor Industry.** G. Dan Hutcheson and Jerry D. Hutcheson in *Scientific American*, Vol. 274, No. 1, pages 54–62; January 1996.

**Handbook of Semiconductor Manufacturing Technology.** Edited by Yoshio Nishi and Robert Doering. Marcel Dekker, 2000.

2003 International Technology Roadmap for Semiconductors. Available online at **http://public.itrs.net/Files/2003ITRS/Home2003.htm**

News from International SEMATECH, a global consortium of leading semiconductor manufacturers, is available online at **www.sematech.org**