


The Near Impossibility of Making a Microchip

Somebody had to invent machinery that could reliably and affordably lay down patterns with details as small as 80-millionths of an inch.

BY DANIEL P. BURBANK

Ln the late 1950s Jack Kilby at Texas Instruments in Dallas, and Robert Noyce and Gordon Moore, at Fairchild Semiconductor near San Jose, California, independently came up with ways to cram many tiny transistors and resistors onto a small sliver of silicon. By adding microscopic wires to interconnect groups of adjacent components, they made the first integrated circuits, which are now known informally as chips. The potential range of applications was vast, limited only by the power of 1950s imaginations to conceive of refrigerator-size computers shrunk to the size of a postage stamp. As with most technological breakthroughs, however, a brilliant idea was only the first step.

The earliest chips of any complexity were logic circuits that went into the guidance systems of intercontinental ballistic missiles. For this application, mass production was not necessary, and cost was of secondary importance. For chips to become as ubiquitous as they are today, though, the technology had to be made not only feasible



The microchip seen here is actually smaller than your thumbnail—as shown on the last page.

Amazingly, fabricators can (and must) mass-produce such chips while replicating every detail flawlessly.

but economical. Achieving this demanded innovations at every step of the manufacturing process.

To visualize the Kilby/Noyce/Moore process, imagine a manic real estate developer building an entire town at once. First hundreds of bulldozers dig all the basements; then a fleet of cement trucks pours all the foundations; then an army of carpenters completes all the floors, then the walls, then the upper stories and roofs; and finally all the streets are paved. In this analogy the individual buildings correspond to individual transistors and resistors, the streets to connections between them, and the blocks of buildings and neighborhoods to functional circuits.

In real life, of course, this method would make no sense for building a town. But for fabricating integrated circuits, which are built from successive layers of silicon compounds and aluminum laid down with microscopic precision, it's the key to mass

COURTESY OF INTEL MUSEUM

When its first projection device needed 16 lenses, Perkin-Elmer decided to try a new approach

production. Many things had to be worked out along the way: isolating and purifying the different substances, finding ways of depositing them, making them adhere to the silicon wafer and to one another, and dozens of additional details. One of the most important was ensuring that the minute and complicated patterns for each layer, as drawn up by the circuit designer, were accurately reproduced on the surface of the chip.

Chip fabricators achieve this precision with the help of light-sensitive compounds and masks. For example, let's say it's time to deposit the thin lines of aluminum that connect the components with one another. Fabricators start by laying down a solid layer of aluminum and covering it with a solid layer of photoresist, a plastic resin that hardens when exposed to ultraviolet light. A glass mask with a negative image of the interconnect pattern is placed on top of the photoresist, image side down. Then a bright ultraviolet light is trained upon the mask, "exposing" the photoresist like film in a camera. Afterward a solvent dissolves the unhardened photoresist, revealing the underlying portions of aluminum. These are removed with

an acid bath, leaving the mask pattern in aluminum topped with exposed photoresist. That hardened photoresist is removed with another solvent, leaving just the network of aluminum lines.

Different combinations of chemicals are used on different layers, but the basic procedure is the same: Deposit a solid film; cover it with photoresist; imprint a pattern on the photoresist by shining intense

ultraviolet light through a mask; dissolve the unhardened photoresist; etch away some or all of the revealed film; and then dissolve the remaining photoresist. The earliest integrated circuits usually required about eight masking steps.

Experience made the masking procedure work reasonably well, and by the mid-1960s many chip makers were building their own machines to do the job for about \$10,000 apiece. The trouble was that the patterns being reproduced on the chips had details as small as 10 wavelengths of ultraviolet light. So the mask had to be clamped tightly to the wafer to prevent blurring at the edges of the image. Despite elaborate procedures aimed at maintaining pristine "clean rooms," tiny particles of dirt or dust would get stuck between the mask and the wafer. By the time eight steps had been completed, a high percentage of chips were defective.

Worse, particles often damaged the mask, and even after they had been removed, the defect remained to be replicated on all later images. The smaller and simpler a circuit was, the greater chance it had of getting through the process intact. The dismal economics of defects



Pioneers of the Micalign include Peter Moller, fourth from left; John Bossung, fifth from left; Harold Hemstreet, third from right; Abe Offner, second from right; and Jere Buckley, far right, with their creation in 1977.

COURTESY SVG

from contact printing placed a low ceiling on practical circuit complexity, even for well-funded U.S. Air Force programs. The Air Force was looking hard for a way out from under this ceiling. The obvious solution was to separate the mask and wafer using a lens to project an image of the mask onto the wafer, as in a photographic darkroom enlarger.

Perkin-Elmer of Norwalk, Connecticut, one of the foremost optical-instrument companies in the country, had a long history of building special-purpose optical systems for scientific work and the defense industry. So it was no surprise when Perkin-Elmer was awarded an Air Force contract to build a lens-based projector in June 1967.

Within a few years the company had come up with its Microprojector, which met the Air Force specifications but was less than satisfactory for industrial use. The main problem was that lenses refract light of different wavelengths at different angles—a basic property of glass called dispersion. Designers of optical systems have learned to combine elements of different shapes, made from different types of glass, to compensate for this phenomenon. These methods involve tradeoffs affecting resolution, image size, and the acceptable range of wavelengths. Perkin-Elmer's contract required its projector to resolve details as small as 2.5 microns, or 100 millionths of an inch, at an exact 1:1 scale. That's the

equivalent of more than 300 million pixels (picture elements) on a two-inch wafer, which may contain hundreds or thousands of chips. (By comparison, a high-quality 35mm camera lens resolves about 6 million pixels in a 28-by-35-millimeter format, roughly half the area.) To achieve such high resolution, the Microprojector used 16 lens elements. All those lenses meant that it worked only in a narrow 200-angstrom sliver of the ultraviolet spectrum (which extends from 40 to 4,000 angstroms). Outside that range, dispersion would cause a faulty image. This meant that most of the light from the system's 1,000-watt mercury-vapor lamp had to be thrown away.

The use of multiple lenses has another consequence that made the Microprojector hard to operate. Before exposure, the mask and the wafer must be aligned with great precision by a technician. The Microprojector system was made to transmit ultraviolet light, and designers found it impossible to get enough visible light through the labyrinth of lenses for humans to see. Fortunately, at about that time, night-vision apparatus developed for Vietnam was declassified. Perkin-Elmer installed an image-intensification system to let operators use ultraviolet light for alignment

The night-vision fix worked, but it was a stopgap that only increased the unit's cost. "We had accomplished the contract requirements," says John Bossung,

an engineer who worked on the project for Perkin-Elmer, "but we hadn't really accomplished what the government wanted, a viable commercial product." Abe Offner, who was one of Perkin-Elmer's most formidable optical designers, agrees: "It was at the very limits of what could be manufactured at the time. No one thought that we could make it in production."






Perkin-Elmer's experience with the Microprojector convinced Harold Hemstreet, general manager of the electro-optics division, that lenses were the wrong approach. With the Micro-projector still in development, he called once again on Abe Offner, who decided to investigate reflective systems—those using mirrors instead of lenses. Reflective systems image light of all wavelengths in exactly the same way, solving the problems of alignment and wavelength range. But they create distortion of their own, which is known as aberration. For that reason, designers had previously avoided them. And to make matters worse, by this time the resolution requirements had gotten even tougher. Now the system would have to resolve details of two microns, or 80 millionths of an inch (sometimes expressed as 250 line pairs per millimeter). On the newly popular three-inch wafers, that worked out to more than a billion pixels.

The task sounded daunting until Offner realized that he didn't have

to resolve those billion pixels all at once. He needed only to achieve high precision in a small area and use it to scan the mask a little at a time, the way a photocopier scans a sheet of paper. He accomplished this by cleverly combining two spherical mirrors in such a way that there was a small ring in their image field in which the mirrors' aberrations canceled each other out.

Offner's design paired a "primary" concave spherical mirror about 10 inches across with a smaller, convex "secondary" one about 2 inches across. (See Fig. 1.) By arranging them along a common axis of curvature, he made a distortion-free ring about a millimeter wide, with a diameter of about 5 inches around the axis of symmetry. (See Fig. 2.) Light from above the mask would travel in a *W* path, with the mask at the top of one arm, the wafer at the top of the other, the concave primary mirror forming a base for the *W* and the convex secondary at the middle peak. Moreover, the magnification was inherently IX, which meant that mask features would be imaged at the same size on the wafer, and the system was telecentric, meaning that it would still work if the mask or wafer was moved slightly out of focus. This allowed much greater tolerance for mask or wafer surfaces that were not perfectly flat. And

The parts of the roof mirror are held together without glue, by intermolecular forces alone



In fig. 4 a roof mirror replaces one of the flat fold mirrors, giving the images on mask and wafers the same orientation. Fig. 5 shows how mask and wafer are held in place in a scanning carriage that rotates by means of a flexure bearing (not shown).

Fig. 1 shows the basic idea behind Offner's breakthrough: Light passes through the mask, is reflected by the concave primary mirror twice and the convex secondary mirror once, and then exposes the photoresist on the wafer. Fig. 2, looking down along the axis of symmetry, shows the thin ring of near-perfect resolution and how it is used to scan the image on the mask. In fig. 3 a pair of flat fold mirrors has been added, allowing mask and wafer to be scanned in the same direction.

The company learned that to sell a tool, you have to develop expertise in its application.

finally, since the system was entirely reflective, it would work at all wavelengths.

Now Offner could project a sliver of light—carefully aligned to correspond with a segment of this distortion-free ring—and move the mask through the sliver until it had been completely covered. The resulting scan, reflected through the system onto the silicon wafer, would have excellent resolution everywhere. His plan had the disadvantage of introducing moving parts into the system, but that could be coped with. Much more important was the advantage of replacing a complicated, unwieldy system of 16 lenses with an arrangement of two mirrors, a system so simple and elegant that it is found today in optical textbooks.

Bossung built a bench-top proof-of-concept model using photographic film instead of photoresist to record the scanned image. The demonstration was enough to get another \$100,000 from the Air Force. With the design's feasibility demonstrated, Perkin-Elmer addressed the task of building a sturdy, reliable machine that could be produced at reasonable cost. In May 1971 Hemstreet put Jere Buckley, a mechanical designer, and Dave Markle, an optical-systems engineer, in charge of the project.

The most obvious approach was to configure the system just as Offner

had drawn it up and simply pull the mask and the wafer through at the same speed. That would avoid the need for any more mirrors. Unfortunately, it would also require that the mask and the wafer be driven in precise, coordinated linear fashion at exactly equal velocities but in opposite directions, across

the ends of the W of light, to submicron tolerances. A photocopier or fax machine can stretch or compress an image slightly without major consequences, but even a tiny amount of scanning distortion would have rendered a chip useless. The mechanical and servo design

required by such a scheme would have been so complicated and expensive that the scanner could never have been commercially viable.

Markle and Buckley next considered what would happen if they placed a pair of flat fold mirrors into the light path, reflecting the ends of the W outward. (See Fig. 3.) Doing so could simplify the scanning process by allowing the mask and the wafer to be held a fixed distance apart, facing each other, and moved in the same direction by a common scanning platform, or carriage. This would have the advantage of requiring only one precision servomechanism instead of two. However, it would still invert the image—that is, it would turn a "left-handed" image into a "right-handed" one. This meant that a company adopting the scanner would have to discard its existing inventory of masks.

Even worse, it meant that if the carriage was bumped perpendicular to the scan direction, the image on the mask would move one way but the image reproduced on the wafer would move the other. This would produce an unwanted jog in the imaged line. To prevent such a situation, the carriage would have to be driven in exactly linear fashion and in exactly the right direction, with any deviation causing distortion. The mechanical



In a "clean room" used for chip fabrication, elaborate precautions are taken to prevent contamination by even the minutest particles.

COURTESY OF SEMI

requirements would be simpler than in Offner's original system but still too complicated for commercial production. Markle made a breakthrough by changing one of the two flat fold mirrors into a roof mirror. (See Fig. 4.) A roof mirror consists of two flat fold mirrors joined at a right angle. Large versions of roof mirrors are often found in clothing stores, barbershops, and hair salons, where they allow customers to see themselves from any angle. Inserting a roof mirror diagonally added an extra reflection to the light path, yielding an image that was not inverted. Now customers could use up their stock of masks that had been made for contact printing. More important, the three-reflection design eliminated the requirement of a precisely linear scanning movement. As long as the mask and the wafer moved in tandem with each other, the direction of the scan did not have to be held constant.

In fact, Markle realized, the scan would not have to be linear at all. This was an unexpected benefit. Incorporating a linear translation into the system would have been difficult, requiring an expensive and high-maintenance carriage supported on an air bearing and floated over a granite slab. But the new design allowed the linear scan to be replaced by rotation around an axis, much more reliable and less complicated. (See Fig. 5.)

As often happens, the best choice of bearing for the

axis of rotation turned out to be the simplest: the flexure bearing. A familiar application of flexures is the plastic cap with flip-up lid found on ketchup and shampoo bottles. The Micralign's more sophisticated high-performance flexure bearing uses thin, flexible stainless steel leaf springs arranged to allow limited rotation around one axis only. "We did look at other bearing types, a rotary air bearing in particular, but chose the flexure bearing for its simplicity, accuracy, and reliability," Buckley says. "You could quite literally throw a handful of sand into the bearing with no loss of performance."

Although the team completed the basic design of the Micralign (as it was dubbed) by November 1971, a number of hurdles lay between that design and a practical manufacturing tool. As Markle recalls, "One challenge was finding a way to efficiently and uniformly illuminate a curved slit 1 millimeter wide and about 80 millimeters long." Simply projecting an ordinary lamp through a slit-shaped hole would not do, because it would waste too much light and thus require too much time per scan. Instead, says Markle, "the answer turned out to be a curved capillary mercury arc lamp which Ray Paquette from ARC made and tested for us in about two hours after I phoned him one day." Offner designed a light source around this special lamp to collimate the ultraviolet sliver—that is, make all the light waves parallel. Using this source, a typical

scan could be performed in 10 to 12 seconds.

Other problems cropped up as the team began building prototypes. "Our first system had all the electronics in the desktop—a bad decision because the electronics kept growing and the desk didn't," says Markle. Another challenge was giving the operator a way to align the mask image with previously patterned wafer features. "This was solved by putting a dielectric coating on one of the fold mirrors for fine alignment and on the secondary [spherical] mirror for coarse alignment." This coating reflected the ultraviolet wavelengths onto the wafer for exposure but transmitted the yellow and green wavelengths through the mirror to the operator's viewing scope for alignment.

Fabrication of the roof mirror also required a major development effort. Because of the extreme precision required of the right angles, it could not be built the conventional way, by gluing two mirrors together with optical cement. Instead the bonding surfaces had to be made so smooth and clean that inter-molecular forces alone would hold the components together, essentially making the two pieces become one.

Jere Buckley designed the alignment system, of which he is still proud: "If you've ever had much exposure to optical systems, you may have noticed that they tend to be full of adjustments, many of which are cross-coupled. What I was able to do on that fold

mirror array mount was contrive a system in which there were three adjustment knobs. You turn one, and it is strictly a focus adjustment. There is no cross coupling to anything else. You turn another, and it exclusively adjusts distortion in the direction of the scan. You turn the third knob, and it adjusts distortion perpendicular or across the direction of the scan and affects nothing else. With that system you could lay out a cookbook procedure for the service engineer, where in a matter of minutes he could make all these adjustments with a minimum of confusion."

Through all the technical ups and downs, Harold Hemstreet provided the vision and management skill to keep the project on track. Buckley says, "I remember Harold sitting in his office and saying, 'Someday we're going to sell a thousand of these machines,' and everybody thought he was totally bananas." The development staff was much less sanguine than Hemstreet, especially as pressure mounted with the product launch approaching in the summer of 1973. Late one night Peter Moller, a marketing executive with Perkin-Elmer, was in the laboratory with a small group of engineers. As Buckley recalls, "The machine wasn't behaving properly, and it

wasn't really all that clear that we were on solid ground with this thing. Peter Moller said, 'I'll give you a trip to Bermuda when we sell the hundredth machine or a cup of coffee now.' We all took

The Micralign scanner made the PCs Possible

the cup of coffee. And I wasn't even a coffee drinker!"

The launch had its bumpy passages, as Moller remembers. "Texas Instruments came in, and we ran a lot of wafers for them. They looked at them through microscopes and so on, and they seemed to be fairly impressed with what we had done. . . . We then launched and went to the West Coast with what we called our 'golden wafers' to show the industry." At Raytheon "we gave them the wafer and they looked at it under the microscope. And the head of production at the time let out a loud bellow, and he said, 'This is s—!' We were absolutely stunned. We left there. It was time for lunch, and we sat around lunch and decided that he probably didn't know what he was talking about. We went into National Semiconductor, and we dealt with a fellow there who was more mask-oriented. He was mostly interested in whether [the image magnification] was 1:1. He looked at the wafer and the mask, and, boy, they were exactly 1:1. So we thought that was great. That was the first day, a 50-50 deal."

The next day was less successful. A wafer they brought to Fairchild "came back with electron microscope pictures of godawful-looking edges. . . . They said this just wasn't hacking it." Another visit that afternoon left a fourth

potential customer similarly unimpressed.

As it turned out, the trouble did not lie with the Micralign itself. "By the time I got home," Moller continues, "Raytheon had decided that we probably didn't know anything about photoresist. They were basically right about that." Perkin-Elmer was not a chip fabricator, and its lack of experience in processing the demonstration wafers had nullified much of the precision engineering that went into the Micralign. Raytheon sent an experienced fabricator to show Perkin-Elmer how to process photoresist. Through this experience Moller and the company learned that in order to sell a tool, you have to develop expertise in its application.

In spite of these initial setbacks, manufacturers came to appreciate the Micralign's superior performance. The first one was sold to Texas Instruments in 1974 for \$98,000. Intel and Raytheon were also among the early purchasers. The Micralign would prove to be a high-tech cash machine for its users, but fitting it into existing production processes was much more complicated than plugging a new component into a stereo system. Fabrication technicians were accustomed to rugged, low-tech contact-printing equipment and had to be taught to respect a unit that was much less forgiving of vibration and

production-floor wear and tear. "Operators would put their feet up on the aligner during exposure," Bossung says. Process engineers then had to figure out why the images looked so much worse than they had in the demonstration. But once the problems were ironed out, cost savings were dramatic. A Texas Instruments manager said the units paid for themselves in 10 months. Perkin-Elmer turned out Micraligns as fast as it could, but new customers had to wait as long as a year.

Early marketing materials emphasized the enormous improvement in mask life. "Contact printing was godawful," Bossung says. "The emulsion [on the masks] would lift off. . . . After ten uses, the masks were useless." Buckley agrees: "Places like TI were buying masks, literally by the truckload, using them six to ten times, then putting them in the landfill." By contrast, the Micralign offered mask lifetimes of at least 100,000 exposures.

Gradually, though, the semiconductor companies began to realize that the greatest savings from projection alignment came not from reduced mask costs but from improvements in yield. No longer did particles transfer from wafer to mask, to be replicated with each subsequent use. And even if the yield for a single step is increased by only a few percentage points, by the

time the step is repeated half a dozen times, the probability of getting a good chip will be much greater.

In 1975 a report by a consulting firm summarized the savings. For small integrated circuits, such as the SN7400 TTL logic series, which could fit on a silicon chip 35 by 48 thousandths of an inch, yields improved from 75 percent with contact printing to 90 percent with the Micralign. A wafer three inches in diameter could hold 4,000 copies of this chip, so the Micralign printed 600 more good chips per wafer. Results were even more dramatic for larger chips with higher profit margins. In the mid-1970s logic circuitry for a four-function calculator could fit on a chip 140 thousandths of an inch square. Contact printing yielded a dismal 30 percent usable chips, while the Micralign yielded 65 percent.

Ultimately, the Micralign scanner made the personal computer possible by allowing the necessary microprocessors to be manufactured cheaply enough for a reasonably priced unit. In June 1978 Intel introduced the 8086 chip. A year later came its sister chip, the 8088, which was used in the first IBM PC. These chips were about one-fifth of an inch square and contained some 29,000 transistors each. With

contact printing, only 20 percent of these chips were salable. With the Micralign, the yield shot up to 60 percent.

Perkin-Elmer reached total market dominance in the early 1980s, when more than 2,000 Micraligns were in use worldwide. Later in the decade,

however, it began to encounter serious competition from firms in the United States, the Netherlands, and Japan. After some rough times Perkin-Elmer sold its Microlithography Division to Silicon Valley Group, a manufacturer of wafer-processing equipment.

Today's chip-fabrication tools can combine laser light sources and sophisticated lens systems with scanning technology and "step-and-repeat" methods to print chips one at a time across a wafer that may be 4 by 6 inches or larger.

Present-day projection aligners don't look very much like the original Micralign, but their performance continues to improve relentlessly on the upward trajectory set by the launch of that machine a quarter-century ago.

Daniel P. Burbank is a photolithography engineer with Seagate Technology, a maker of computer disk drives. He lives in Minneapolis.

