

# Choices for tomorrow's chips

Can the information technology revolution of the past two decades repeat itself? Much depends on how the next generation of electronic chips is made

**Alec Broers**

IN THE 1950s, an electronic circuit that could store a single "bit" of information cost more than £1. Today, a penny will buy 5000 of them. This dramatic reduction in cost led to the information technology revolution of the past two decades. Could it all happen again?

Three developments made cheap chips possible: smaller circuits, fewer electronic components per circuit and bigger silicon wafers. Semiconductor engineers are still moving as fast as ever towards further miniaturisation, for the simple reason that the number of circuits per wafer is increasing very much faster than the cost of making the wafers.

Few researchers doubt that circuits will carry on shrinking, but manufacturers face some stark choices about the technology that will bring this about. Some favour extending today's techniques for making chips, based on optical lithography; others are looking at alternatives, based on X-rays or on electron beams. The stakes are high—chip makers expect the market for the next generation of chips to be worth thousands of millions of pounds.


Circuits have been shrinking steadily for 25 years. Until the 1960s, transistors used to control the flow of current were individual components connected by wires and laid out on a circuit board. These were replaced by integrated circuits in which transistors and all their connections were formed within the surface layer of a sliver or "chip" of silicon. The size of a circuit shrank by 100 times in this process. Then integrated circuits themselves began to shrink.

## Smaller and smaller

A typical transistor is made up of two electron-rich regions with an electron-poor layer sandwiched between them. The flow of current between the two regions is controlled by a "gate" electrode placed over the electron-poor layer. By the 1970s, designers were predicting that the trend towards miniaturisation would come to a stop when the distance between the regions became as small as 0.5 micrometres. At this point, the electron-rich regions would inevitably overlap and control of current by the gate would be lost. Fortunately, electronics engineers were able to perfect processes for making transistors from thinner and thinner layers, proving the designers wrong.


Researchers at IBM's laboratories in Yorktown Heights, New York, have already shown that transistors with gate lengths as small as 0.1 micrometres work in test circuits. But they cannot be mass-produced. A major reason is that many of them are made by the expensive technique of electron beam lithography.

What manufacturers are looking for is an economical way of mass-producing chips carrying the largest possible number of features. This is particularly important for memory chips—the more features there are on the chip,



Brian Cronin





**Figure 1** Most chips are made by optical lithography. A series of patterns is transferred photographically onto a silicon wafer, elements of the integrated circuit are built into the silicon by doping, and the wafer is cut into chips. As circuits shrink, the physics of the optical system is limiting how much detail can be added to each chip

the more data it can store. Planners in some parts of the semiconductor industry say that by 1999 they will have moved a long way on from today's 4-megabit chips, to test versions of 1-gigabit (1000 megabit) chips which incorporate 0.1 micrometre features. Even by 1996, they say, features as small as 0.25 micrometres will be needed for the next generation of chips—under development by all the leading companies in the field—that can store 256 megabits of information.

Integrated circuits are made using techniques derived from lithography. In this process a series of patterns is transferred photographically onto a large slice, or wafer, of single-crystal silicon, which will later be cut into individual chips. The wafer is first heated so that a layer of silicon oxide—glass—forms on the surface. Next, it is coated with a thin, light-sensitive layer, called the resist. The "negative" of the required pattern is made in metal on a plate called a mask that is transparent to light. Light is passed through the mask, and transfers the image onto the wafer. Resist is then dissolved away from the exposed area. Elements of the circuit are built into the revealed silicon, by

doping it to make electron-rich or electron-poor regions, through holes in the glass layer. Wires that connect individual transistors are then formed on top of the silicon. There are several layers of wires, separated by insulating layers with holes in them so that the layers of wires can be connected.

To build a complete circuit may take 20 lithographic steps, each involving a new layer of resist, to expose separately the parts of the wafer which require different types of doping. The limit to how much detail can be put into the circuit—that is, how small the circuit can get—depends on the smoothness of the resist, the precision of the processing steps, and most importantly, on the size of the image.

Most chips are produced by optical projection: that is, the image of the circuit on the mask is projected onto the wafer using a projection camera, which works on the same principle as a photographic enlarger, but in reverse. Some cameras reduce the image, while others project it at full size. The most common type of optical projection camera operates by a step-and-repeat method. The exposed area is a square or a rectangle the size of a single chip, typically 10 to 20 square millimetres. The wafer is moved in steps between exposure sites, hence the name. In all optical cameras, extremely complex lenses are needed to project the finest possible lines in the same degree of detail over the whole area of the chip.

### How small can circuits shrink?

A fundamental limit to the detail achievable by this process is set by the wavelength of light and the size of the lenses used to project the image. If the lens is small and the wavelength of light large, the light waves are "diffracted" or diverge as they pass through the lens, rather than continuing unperturbed, just as real waves spread out after entering a narrow harbour entrance. In the case of light, such spreading blurs the image. To avoid this, engineers use the shortest possible wavelength of ultraviolet light and the biggest possible lenses—about five times as large as a standard photographic camera lens. Researchers are now trying to produce features whose size is approaching the diffraction limit of 0.15 to 0.2 micrometres for imaging with ultraviolet light.

The minimum feature size that an optical projection camera can reproduce—its "resolution"—is set approximately by the ratio of the wavelength of light to the numerical aperture, which in turn is approximately equal to the ratio of the diameter of the lens to its focal length. So, the shorter the wavelength and the bigger the lens, the smaller the size of feature. In practice, the minimum wavelength is currently 250 nanometres, which is in the deep ultraviolet. (There are lasers which are more intense, and which produce wavelengths shorter than 250 nanometres, but so far these are proving difficult to use.) The maximum numerical aperture is about 0.6, so the minimum line width is currently about 0.35 micrometres.

A drawback with the optical method is that it demands a perfectly flat sample to reach the ultimate limits of resolution. This is because as the lenses get bigger, the convergent angle of the rays at the surface becomes larger, and the depth of focus—the depth of the region over which focus is maintained—becomes shallower (see Figure 1). Until about five years ago, semiconductor process engineers thought that optical lithography would reach its limit when lines became as fine as about 0.5 micrometres. This is because the layers which make up the integrated circuit are about 1 micrometre thick, giving rise to surface roughness of the same size.

This problem was unexpectedly solved in 1986, when Fedor Coopmans and Bruno Roland of Leuven University in Belgium,



among other groups, developed a new way of processing resists. They discovered how to form the image in a thin layer on the resist surface. They could then transfer the image into the underlying resist by dry etching—a sort of sand blasting with oxygen ions. By this means, the depth of focus can be reduced to below a micrometre, and the resolution of optical lithography refined to at least 0.35 micrometres.

In the past two or three years researchers have also developed ways of getting round the classical diffraction limit. Researchers at Hitachi and other Japanese companies developed a technique called focus drilling, in which either the lens or the wafer is moved up and down while the wafer is being exposed. This gives a more uniform exposure. It is especially useful for getting rid of variations at different depths in the resist caused by interference between incoming light waves and those reflected from the surface of the wafer.

Another method, proposed by Marc Levenson of IBM's Almaden Research Laboratories in Palo Alto, California, is to use so-called phase shifting masks. With these masks selected areas are covered with layers that slow down the light wave and so shift the phase of the transmitted light in such a way as to increase the contrast for particular small features. This method is most easily used for regular patterns such as arrays of holes which are all identical in shape and size. Used together with focus drilling, it may allow researchers to produce features as small as 0.15 micrometres.

All these advances help to put more features onto each square millimetre of a chip's surface, but that only goes part of the way towards meeting the ambitions of chip designers. Their latest creations contain so many features—and therefore cover such a large area—that it will become increasingly difficult to expose an entire chip with a single image. To get all the detail in would require an impractically large and expensive lens. Instead, the

chip pattern will have to be exposed in sections and “stitched” together—with perfect seams—or perhaps built up by scanning, as is already done in the latest optical scanning cameras. This change to scanning or stitching will probably be essential for lines finer than about 0.3 micrometres, and for chips bigger than 25 square millimetres.

The dilemma facing the semiconductor industry is whether to throw all its efforts into developing optical lithography, in the ways just described, for example, or whether to switch to an alternative. The leading candidate is X-ray lithography. X-rays have a shorter wavelength than ultraviolet light, so they should give better resolution. The technique was developed in the early 1970s and first demonstrated by Henry Smith's group at MIT's Lincoln Laboratories but because of the difficulty of finding adequate sources to produce the X-rays it has been “waiting in the wings” until relatively recently.

### The search for an alternative


In this method, X-rays “shadow print” an image onto the wafer—there are no lenses involved. It should, within four or five years, be able to produce components with dimensions as small as 0.25 micrometres and eventually it will produce components as small as 0.1 micrometres. But the technology presents a formidable challenge. To be competitive, it needs an elaborate X-ray source, such as an electron synchrotron storage ring, and a mask that is fragile and difficult to make.

A third possibility engineers have considered is to write the circuits onto chips directly, using electron beams. But this is really only feasible for producing specialised chips, not for the mass production of memory chips (see Box overleaf).

In X-ray lithography (see figure 2) the image is obtained by shining low energy, “soft” X-rays with wavelengths in the range 0.4 to 5 nanometres onto the wafer through a mask made from a silicon or silicon carbide membrane less than 5 micrometres thick: the glass plates used for optical masks are not transparent to soft X-rays. The wavelength is a compromise. Harder X-rays—with shorter wavelengths—would go through glass masks, but they would also go straight through the resist without exposing it. A gap between the mask and the wafer ensures that they do not touch and damage each other. The metal pattern on the mask that absorbs the X-rays is written using electron beams.

X-ray lithography produces sharper images in thick resist layers than optical lithography, because the effective resolution is better and the depth of focus is larger. A curious but important advantage is that exposure is relatively unaffected by small particles, such as flakes of human skin, that frequently spoil chips made by optical lithography. Because these particles are of low atomic weight, X-rays pass straight through them. The resolution of X-ray lithography is limited by the spreading of the X-rays between the mask and the surface of the sample, as a result of diffraction. For an X-ray wavelength of 1 nanometre and a gap between the mask and the wafer of 10 micrometres it should be possible to image 0.1-micrometre lines. The resolution of X-ray lithography currently used at IBM and other laboratories is 0.25 micrometres, as the gaps are typically 40 micrometres.

The disadvantages of the technique lie with the mask and the X-ray source. The metal absorber that carries the pattern must be much thicker than for an optical mask. The mask for X-ray lithography must also be the same size as the wafer, in contrast to optical lithography, where in some cases the image is reduced in the lithographic process, allowing the features on the mask to be larger and therefore easier to make. The features on the



**Figure 2** X-ray lithography avoids many of the problems associated with optical systems. But to compete on cost, the X-rays must be extremely intense. One of the latest X-ray lithography facilities uses a compact electron accelerator, incorporating superconducting magnets, to achieve this.



X-ray mask are so small that stresses in the relatively thick metal absorber can distort the image.

The biggest problem, however, is the X-ray source. With a conventional X-ray source of the type used for medical X-rays, it takes several hours to expose a wafer 200 millimetres in diameter. These sources are used in the laboratory but would not be economic for mass production. The only X-ray source fast enough to compete with optical methods of producing chips is an electron synchrotron storage ring.

Researchers in many fields, especially spectroscopy, have been using such storage rings since the late 1970s, because they are the most intense broadband radiation sources known. The storage ring is a circular vacuum tube in which electrons circulate at speeds very close to the speed of light. They are held in their orbit by magnets, and their energy is increased and maintained by radio frequency fields in the ring. As they are deflected by the magnets the electrons emit an intense beam of electromagnetic radiation which ranges from visible light to X-rays.

At first sight, the main obstacle seems to be cost. A storage ring for lithography costs as much as \$25 million. But the ring could have 20 exit ports, each producing enough radiation to expose 40 wafers per hour, equivalent to the output of a modern optical step-and-repeat camera. A step-and-repeat camera costs about \$2 million, so the ring has the same potential production capacity as \$40 million worth of step-and-repeat cameras—and gives higher resolution. The disadvantage which worries manufacturing engineers is that if the electron beam fails, all the exposure stations stop working.

One of the latest developments in storage rings is the compact electron synchrotron, which uses superconducting rather than normal magnets. This makes it smaller, so it should produce chips more cheaply than a full-sized ring, while maintaining the same output. In a forward-looking project set up by Peter Williams of Oxford Instruments, a team of researchers led by Martin Wilson built a compact storage ring. Shaped like a race-track, and called Helios, it was developed in collaboration with

## Electron beam lithography: the finer points of writing chips

AN alternative to optical lithography some researchers have considered is scanning electron beam lithography, a technique pioneered at Tübingen University and at the University of Cambridge in the late 1950s and early 1960s.

In this technique a fine electron beam writes the pattern straight into the resist, without a mask. The beam scans across the wafer rather like the spot scans across a television screen. At first sight this looks like the ideal way to write integrated circuit patterns.

Electron beam lithography has been worked on for more than 30 years and has been shown to have much higher resolution (smaller feature size) than optical lithography. Structures and devices with dimensions one-hundredth the wavelength of light (10 nanometres to 20 nanometres) have been made by many researchers, including the author at IBM and Cambridge; Michael Isaacson of Cornell University; and Colin Humphreys, then at the University of Liverpool and now at Cambridge. This resolution is well beyond the foreseeable needs of integrated circuits.

Electron beam systems have successfully made the masks for optical and X-ray lithography and a few entrepreneurial spirits have opted to use them for the manufacture of custom and prototype devices. The drawback is that they write very slowly.

In the early systems the pattern was written with a small round electron beam with a diameter about a quarter of the required linewidth. Shaped electron beams that project the equivalent of about 20 round beams at once improve the writing speed.

The first successful beam systems of this type were built by Hans Pfeiffer at IBM in about 1970 and the writing speed was later increased again with beams whose shape

could be changed as the pattern was written. Even with these variably shaped beams, the necessary exposure time remains longer than with an optical projection camera, and the systems are more expensive too.

Engineers have struggled for 25 years to reduce the cost of direct-writing with electron beams. The problem is that the number of pattern elements (pixels) per chip has increased faster than the rate at which pixels are exposed. This disadvantage rules out scanning electron lithography for mass production.


The potential of this method lies in the manufacture of custom chips—special-purpose chips used for selected applications, as opposed to the standard memory chips and microprocessors made in their millions for storing information and as the standard processors for personal computers. The custom chips are the so-called application-specific integrated circuits (ASICs). A company called European Silicon Structures of Rousset, Aix-en-Provence, founded in 1985 by Robb Wilmot and Bob Hikes, uses only electron beams to write ASICs.

It is uncertain, however, whether electron beam lithography can be made fast enough and cheap enough to remain economic for the fabrication of such chips.

Also in question is whether the technique can be justified for making devices that are needed in


a hurry. Until now, chips produced in under 10 days were expensive enough to cover the cost of direct writing with an electron beam. To make and inspect the masks for optical cameras would have taken too long. Recent advances in mask making, however, have reduced this time to a few days, shifting the balance in favour of optical methods.

One way of increasing the number of wafers processed is to increase the current density of the beam—in other words to increase the number of electrons per flash. But even with the most modern, shaped beams, taking this beyond a certain point increases interactions between electrons and blurs the beam. □



In electron beam lithography a narrow beam of electrons writes the pattern for each chip directly onto the silicon wafer. This technique gives extremely good resolution compared with optical lithography but is slow and expensive





David Parker/SPL

Drawing the line: chip manufacturers face a difficult choice

Physicists from the Synchrotron Radiation Source at Daresbury in Cheshire. The first Helios was delivered last year to IBM in New York, where it now forms part of one of the world's most advanced optical lithography systems.

No matter how the competition between optical and X-ray methods is resolved, the cost of microelectronic chips will keep falling for at least another 15 years, and perhaps 30. The result will be a repeat of the staggering increases in performance and cuts in price we have seen over the past two decades.


With optical lithography, engineers are faced with the problems of perfecting lenses, light sources and resists. X-ray lithography presents an entirely different set of problems, centred on the masks and the X-ray source. Despite its drawbacks, X-ray lithography is perhaps the most promising method for large-scale production of devices with dimensions less than 0.35 micrometres, because it produces better resolution and depth of focus than the optical method and is cheaper than scanning electron beams.

In retrospect, the best strategy for the past 20 years would have been to develop optical projection for exposing wafers and scanning electron beams for writing the masks, and not to have worked on alternatives such as X-rays. But to continue to pursue all the alternatives will be prohibitively expensive. Despite the lack of information on which to choose, decisions are already being made. The spoils are likely to go to those who invest enough in their choice to prove it was the right one. □

**Alec Broers** is Master of Churchill College and professor of electrical engineering at the University of Cambridge.

# NewScientist BINDERS

Keep the world of  
science and technology  
at your fingertips



Every week, New Scientist explores the issues and scientific developments that affect our lives. But the interest of many features extends well beyond the week of publication. If you have ever searched in vain for a mislaid back issue, you'll appreciate the practical value of the New Scientist binder. Each issue is kept to hand, accumulating week by week into a unique permanent encyclopaedia of scientific and technological developments.

Each binder holds 13 copies and costs £5.50 including VAT, postage and packing (overseas £1.00 extra). Just fill in the coupon and send it with your payment to John Denton Services, Unit 13, Thornham Grove, Stratford, London E15 1DN.

Please note that these binders are only suitable for issues of New Scientist dated from 3 February 1990.

Please allow 28 days for delivery.



## New Scientist Binder Order Form

Complete this coupon and send it with your payment to John Denton Services, Unit 13, Thornham Grove, Stratford, London E15 1DN.

ALL EUROPEAN ORDERS MUST BE PAID BY EUROCHEQUE

Please send me \_\_\_\_\_ binders at £5.50 each.

I enclose a cheque/PO for £ \_\_\_\_ made payable to New Scientist.

Name \_\_\_\_\_

Address \_\_\_\_\_

Post Code \_\_\_\_\_

☐ From time to time you may receive further information about offers, services and products that may be of particular interest to you from other organisations. If you would prefer not to receive these, please tick the box.