Reed Electronics Group

MICROPROCESSOR REPORT

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

INTEL GETS EXTREME IN 2009

Extreme Ultraviolet Lithography Scheduled for Mass Production

By Tom R. Halfhill {2/10/03-02}

Intel claims it will be the first company to mass-produce microprocessors using extreme ultraviolet (EUV) lithography, a revolutionary new photomask technology. Pilot production is scheduled to begin with the 45nm fabrication process in 2007–2008, using tools and techniques now being refined. Mass production is scheduled to debut with the 32nm fabrication process in 2009.

EUV lithography will enable a significant leap forward in the circuit density of chips, because the shorter-wavelength light allows stepper tools to draw features at least 10 times smaller than is possible with today's deep-ultraviolet (DUV) lithography. (See *MPR 6/19/00-01*, "Extreme Lithography.") Intel is betting heavily on EUV to keep Moore's law from expiring for at least another decade. Without EUV, the famous "law," which observes that circuit density doubles every two years, would soon be unsustainable.

Already, some circuit features on today's most advanced microprocessors are actually smaller than the wavelengths of light that draw the features. That sounds impossible, but it's feasible with various resolution-enhancement techniques that compensate for the diffraction effects of the lithographic process. For example, a technique known as optical proximity correction uses photomasks with deliberately distorted patterns to create fuzzily etched but functional circuits with the desired patterns. Another technique shifts the phase of the patterns by using 3D structures on the masks.

These and other work-arounds greatly increase the complexity and expense of the photomasks, which is why a

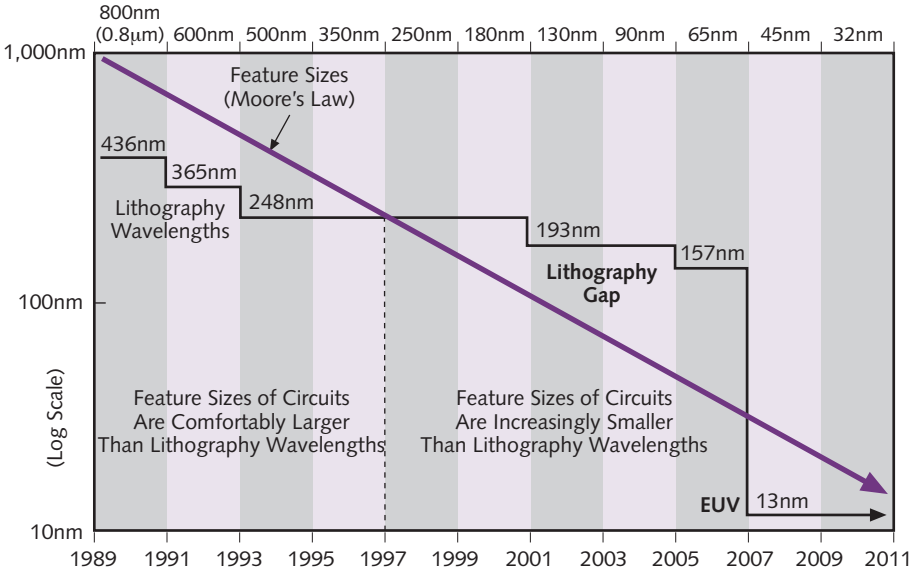

Figure 1. Improvements in lithographic imaging have been incremental for almost two decades, opening a gap between the progress of Moore's law and the resolution of circuit-imaging tools. EUV is a breakthrough that reduces the wavelength of the light from 157nm to about 13nm, which will eliminate the gap for several years.

single mask set for a microprocessor fabricated in a leading-edge process can cost a million dollars. EUV will make such clever trickery unnecessary, at least for several years after it debuts.

The Death of the Red Mask

Long-wavelength light (toward the reddish end of the spectrum) has been dead for many years in chipmaking lithography. To draw the increasingly tiny circuit features made possible by other advances in fabrication technology, chipmakers have steadily retooled their fabs with lithographic imaging systems that use shorter-wavelength light (toward the violet end of the spectrum). Today's DUV lithography uses light well below the visible spectrum. EUV is near the spectral region of soft x-rays.

By projecting the light through several photomasks, corresponding to the semiconductor, insulator, and metal layers in a chip, the lithographic optics and stepper tools repeatedly draw the die images on a wafer. Chemical solvents wash away unwanted material in the resist layers, leaving behind the finished chip images for dicing and packaging. Figure 1 shows how feature sizes have been shrinking faster than lithography wavelengths since 1989 and the dramatic difference EUV will make.

Anticipating the need for EUV lithography, several companies formed the EUV-LLC industry consortium in 1997 to share the burden of developing the basic technology. Intel was a charter member of the consortium, which is based in Livermore, California. Other major members are AMD, IBM, Infineon, Micron, and Motorola. For five years, they have been working with lithography tool vendors, U.S. national laboratories, and universities to create the techniques all the companies will use in their fabs. That initial research phase, known as the "pre-competitive" phase, is now ending. The next phase is actual tool development. Each company in the consortium will decide for itself exactly when and how to implement EUV.

Intel is determined to be the first company to use EUV in mass production. Intel's aggressive strategy is to prolong Moore's law—named after one of Intel's founders—by introducing a new process generation every two years. The schedule rarely slips. The current roadmap calls for 90nm lithography to debut later this year, followed by 65nm in 2005, 45nm in 2007, and 32nm with EUV in 2009. (See *MPR 9/3/02-01*, "Intel Adopts Strained Silicon.") Some EUV production may begin at the 45nm node during 2007–2008, depending on tool availability.

Intel's Competitors Stay in the Race

IBM Microelectronics is Intel's ablest rival in fabrication technology. Although IBM is a fellow member of the EUV-LLC consortium, it hasn't committed to a date for using EUV lithography in mass production. IBM says there are too many uncertainties to set a firm date: when the exposure tools for generating EUV light will become widely available; how soon manufacturers can make sufficient quantities of defect-free EUV mask blanks; and whether or not 157nm-wavelength DUV light (expected to debut in 2005) will have a longer lifespan than scientists currently anticipate. Owing to these and other factors, IBM says only that it plans to have the capability to use EUV at the same time Intel does.

It's difficult to predict when production quantities of EUV mask blanks will become available, because they have almost nothing in common with conventional mask blanks. Blanks are the starting point for masks before they incorporate the image for the mask layer. Conventional photomasks are transmissive, like 35mm color slides; lithographic exposure tools project the mask's image onto the wafer through an optical-glass lens. In contrast, EUV masks are reflective, because the exposure tool's glass lenses, and even the air between the lenses, would absorb EUV light. An EUV blank consists of about 80 layers of perfectly flat defect-free substrates, topped off by a highly reflective surface layer. EUV blanks are much more difficult to manufacture than optical mask blanks.

Although IBM made its first 45nm EUV gate masks for SRAM chips in 2001, the technology is far from perfected, notes Patricia Tivnan, manager of mask technology and operations for IBM Microelectronics. To refine the process, IBM recently began working with the University of Albany and another industry consortium, International Sematech, at a new research and development center called International Sematech North in upstate New York. Sematech North is concentrating on methods for creating defect-free EUV blanks and will share the research among its members. (At this writing, Intel isn't a member.)

AMD, Intel's closest competitor for x86 processors, doesn't currently have EUV on its official technology roadmap. However, AMD (with Intel and Motorola) was a charter member of the EUV-LLC consortium in 1997 and is actively working on EUV technology.

A likely partner for AMD is IBM, because the two companies recently formed a fabrication-technology alliance. (See *MPR 1/27/03-01*, "AMD and IBM to Develop Fab Technology.") Although their official roadmap currently stops at the 45nm node in 2007, no one should be surprised if AMD and IBM extend their relationship to develop a 32nm process with EUV, says Craig Sander, AMD's vice president of process technology development. Such a deal could give AMD the capability to deploy EUV at about the same time that Intel does.

Ye Olde Masque Shoppe

One reason Intel believes it can beat its competitors to EUV is that it still makes all its own masks for leading-edge chips—an expensive undertaking. (Intel outsources only a few mask sets for less-important products.) These days, it's more common for a semiconductor company to share a mask shop with a partner or to subcontract all or part of its mask production to a merchant mask shop. Among the few other U.S. companies that have captive mask shops are IBM and Micron.

These companies believe their mask shops provide a significant competitive advantage for two reasons: they can exercise more control over the shop's adoption of new technology, and a captive shop can shave days or weeks off the turnaround time for a new mask set.

Intel says its captive shop can make the first three layers of a current-generation 130nm mask (enough to start producing wafers) in five days, compared with seven to twelve days for a merchant mask shop. Saving a few days may not sound important, but it can have an enormous effect on revenues from popular chips. International Sematech estimates that each extra day of mask production costs about $2.5 million in lost revenue for leading-edge DRAM chips. Intel won't disclose exact figures, but it says the lost revenues would be even greater for leading-edge PC processors.

For lower-volume chips, Intel makes its own multiproject shuttle masks with 2–25 chip designs grouped together in a single mask set. Because all these designs are internal, Intel says it can achieve a better balance of production priorities than can an independent foundry, which typically combines chips from several different customers on one shuttle mask. (What this seems to mean, however, is that someone at Intel resolves the conflicting priorities of internal design teams by edict, whereas an independent foundry would resolve the conflicts among customers by negotiation or by the amount of money each customer paid.)

Like IBM, Intel has already produced some experimental EUV masks. To fix the inevitable manufacturing defects that mar all photomasks, Intel has been working with tool vendors to design equipment that can detect and repair defects at very small geometries. The prototype equipment can find every defect as small as 11nm high on a six-inch EUV mask blank in 20 minutes. That's comparable to finding every two-inch bump on an area of land as large as California. Other equipment repairs tiny defects on patterned masks by etching them with electron beams.

Intel: EUV May Stabilize Mask Costs

The skyrocketing costs of conventional mask sets are a growing obstacle for developers of ASICs and SoCs, especially for chips that have little chance of attaining high volumes. If the leading-edge technology of EUV lithography costs even more than optical lithography, many chips could be priced out of their markets.

However, Intel thinks EUV masks may hold the line on costs—at least in relative terms. One factor in favor of EUV is its potential to reduce the complexity of the mask images. Consider that when Intel introduces its 90nm process later

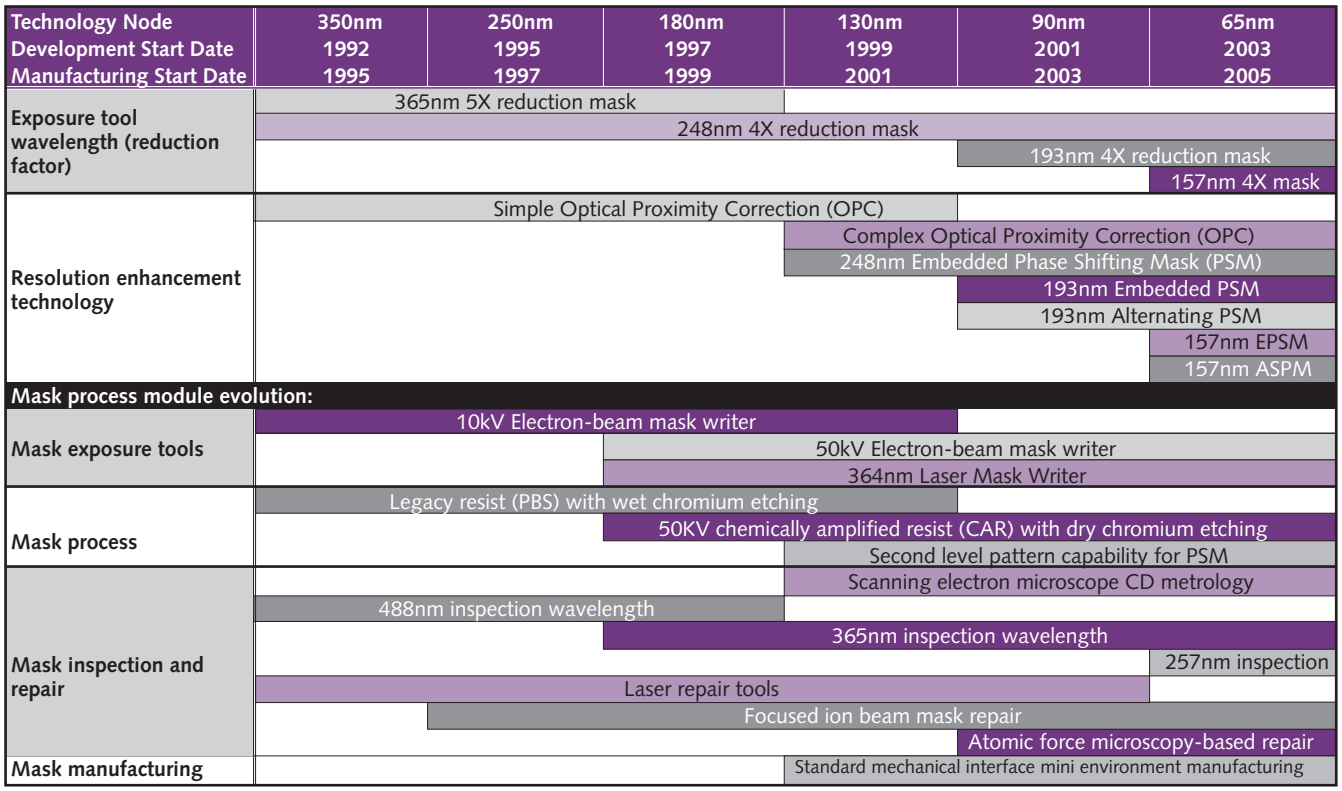

Figure 2. When feature sizes became smaller than lithography wavelengths around 1997, opening up the "lithography gap" shown in Figure 1, mask makers began resorting to many new techniques to keep Moore's law on track. Pessimists interpret the chart above as evidence that mask making will continue to become more complicated and expensive. Optimists hope EUV will simplify mask making by once again shrinking wavelengths below feature sizes.

For More Information

Intel's Web site has a page explaining the basics of EUV lithography at *www.intel.com/home/trends/future/euv_lithography.htm*. More detailed information about Intel's lithography research appears at *www.intel.com/research/silicon/lithography.htm*.

this year, a mask set for a Pentium 4-class processor will require 22–25 photolayers consisting of one trillion pixels in a 200GB file. Much of that complexity is a result of resolution-enhancement techniques to compensate for inadequate lithographic optics. EUV masks won't need those tricks, so they won't be as complex (assuming a microprocessor of comparable design).

Therefore, Intel expects EUV mask sets in 2007–2009 to cost about the same as conventional mask sets will. Although the actual cost of EUV masks will almost certainly increase with the rising complexity of microprocessor designs, Intel believes the relative cost, compared with that of conventional masks, may decline over time. Others dispute that hypothesis, arguing that the greater cost of manufacturing the many-layered EUV blanks and other factors will outweigh any savings from the reduced complexity of the mask images. Even without EUV, mask-making has become much more challenging over the past decade, as Figure 2 shows.

One factor that complicates the long-range forecast of mask costs is the unexpected longevity of existing technology. It's possible that unforeseen advances in DUV will delay the introduction of EUV. Recall that scientists once despaired of creating masks capable of resolving features smaller than one micron—an imaginary barrier shattered more than a decade ago. As recently as 2001, some critics predicted that optical lithography would hit a wall at 65nm. Yet today, every manufacturer working toward 65nm (and even 45nm) lithography plans to use DUV instead of EUV at those nodes. The 157nm-wavelength DUV technology scheduled to debut in 2005 may prolong the life of optical lithography longer than anyone anticipates. If that happens, semiconductor manufacturers will almost certainly postpone their costly retooling for EUV.

Another consideration is whether EUV is the best solution for next-generation lithography. Although EUV now appears to be the front-runner to replace DUV in large-volume fabs, there are some alternatives. X-ray lithography, electron-beam lithography, electron-beam direct writing, and other technologies all have their advantages and disadvantages compared with EUV. Some of these technologies are used today for low-volume production and will probably persist for certain kinds of masks. For example, electron-beam lithography and direct writing are too slow for complex mask layers, but they are sometimes suitable for less-complex contact and via layers.

Indeed, even after EUV or another next-generation lithography debuts, there will continue to be applications for conventional optical lithography in mask layers that don't need extremely small features, such as some metal layers. Typically, manufacturers use the most-advanced lithography at any node for only the most critical two to four mask layers. This practice will continue with EUV and will cushion the expense of retooling the fabs.

Shuttle Wafers Could Distribute Costs

When introducing any new fabrication technology, Intel has a unique advantage: industry-leading volumes of PC and server processors with comparatively high average selling prices. That advantage will allow Intel to amortize the cost of retooling its fabs and creating EUV masks by spreading the expense over tens of millions of chips.

Other chipmakers with smaller fabs, smaller volumes, and lower prices will have trouble matching Intel's economies of scale—just as they do today. If nothing else, however, the much greater circuit densities enabled by EUV lithography will allow chip makers and foundries to combine more designs on multiproject shuttle wafers, thus spreading the costs among more customers.

Today, shuttle wafers are an increasingly popular way to prototype and sample new chips, but production runs rarely combine multiple designs on a wafer. In the future, more customers may use shuttle wafers for actual production, as Intel does—if the foundry can coordinate different customers' quantity demands, an admittedly difficult juggling act. But something has to give. As EUV lithography enables smaller die sizes, and as wafers grow in diameter from 200mm to 300mm to 450mm over the next decade, fewer individual chip designs will be able to amortize the climbing costs of nonrecurring engineering and wafer manufacturing.