

The Gigachips are Coming!

As the end of the millennium approaches, semiconductors are nearing their own milestone—the creation of memories that pack a billion bits in one chip. In the 1994 “National Technology Roadmap for Semiconductors” produced by the Semiconductor Industry Association (San Jose, CA), the first gigabit memory chips were predicted to debut in the year 2001. Halfway to that date, there is no reason to think the prediction needs to be substantially altered.

However, a number of challenges must be overcome before the first gigachip makes its appearance. The most important of these involve difficulties in lithography, dielectric materials, and on-chip wiring.

Lithographic limits

“For lithography, which is one of the most important tools or processes for making smaller and smaller transistors, we need light sources of shorter and shorter wavelengths,” pointed out Tak H. Ning, a fellow at IBM’s Thomas J. Watson Research Center in Yorktown Heights, New York. “Light sources of 248 nm are used in production today, and light sources of 193 nm are in development.” Gigabit chips depend on those 193-nm light sources. To make such chips, the light sources must be developed, and new imaging optics and photoresists must also be put into production.

Current refractive-imaging optics suffer from problems when exposed to deep ultraviolet light. Below about 200 nm, the optics begin to heavily absorb ultraviolet light. The repeated—nearly continuous—exposure required in cost-effective manufacturing means that the optics components will heat up, thereby distorting the light path, and the materials will also age rapidly because of the heavy absorption.

“If one continues with the optical kinds of approaches, we will have to go to reflective types of systems,” commented Sunlin Chou, an Intel (Santa Clara, CA) vice president who is in charge of component technology development.

The same absorption problem appears in


the photoresist. Because the goal of the photoresist is absorbing light, this does not seem like a problem at first glance, but the photoresist is too absorptive. At such short wavelengths of light, the top layer of resist soaks up all the energy and the bottom layer remains unexposed. One solution that has been suggested is to image only the top layer with the 193-nm light and then use other techniques to transfer that image on down through the photoresist.

Fortunately for semiconductor manufacturers, only the most critical layers, such as transistor definition or interconnect, require such short wavelengths. These critical layers constitute about one-third of the total number of layers needed. All the other layers will be able to use what will then be standard 248-nm lithography.

Controlling capacitance

The dielectric material issues are really capacitance-driven. In most circuit elements, less capacitance translates to more speed, but one very important element requires more capacitance.

As features shrink, the distance between on-chip wiring decreases both horizontally and vertically. In addition, the number of interconnect layers on a gigachip is expected to be five or six versus the four to five used today. As a result, on-chip wiring capacitance will be higher for gigachips than for today’s integrated circuits. That increasing capacitance leads to a decrease in performance, because it increases the RC (the product of resistance and capacitance) time constant. Finding materials that have lower dielectric constants, and thereby less capacitance per unit area, may be a solution. There are a number of materials under investigation, but none has been successfully integrated into existing manufacturing processes. Silicon dioxide is the current material of choice because it is easy to work with, withstands high temperatures well, and fits mechanically with silicon (see figure, top right). Any replacement has to have the same characteristics.



Texas Instruments

Ultra-high-vacuum surface science system is used to grow and characterize thin silicon oxide layer on silicon, in a search for lower capacitance.

Memory chips, on the other hand, need more capacitance, not less. Nearly all memory chips sold are DRAMs (dynamic random-access memories) with data stored as charge on a capacitor. The size of capacitor plates and the distance between them determine the amount of charge that can be stored. As feature sizes shrink, capacitance goes down because the area of the capacitor plates is decreasing faster than the distance between them.

If an energetic particle strikes a memory cell, it can alter the charge on the capacitor enough to make a stored 1 become a stored 0, and vice versa—a so-called soft error. Such soft errors were a problem with memory chips years ago because of alpha particles. That problem was corrected by a polymer coating on the chip that stopped the alpha particles. In gigabit memories, soft errors can be caused by neutrons, and no economical method stops them. Consequently, a minimum capacitance of about 25 femtofarads per memory cell must be achieved to prevent soft errors.

“You need to maintain the capacitance, which means either you need to make the cell-capacitor-insulator thickness be thinner and thinner, or you just create a much deeper and deeper trench, or build a skyscraper-type structure on top of the wafer,” said Yoshio Nishi, senior vice president of R&D for the semiconductor group of Texas Instruments (headquartered in Dallas).

Other possible solutions involve higher capacitance dielectric materials, such as

barium–strontium–titanate, for use in memory cells; but these materials are not as compatible with the existing manufacturing processes as silicon dioxide. Putting a structure on top of a wafer makes it difficult to run interconnects over it, so the most likely solution may be to build deeper and deeper trench capacitors.

Wiring

The increasing difficulty in connecting millions and millions of transistors together creates an on-chip wiring problem. The traditional solution to building integrated circuits that are more complicated has been to use more layers of interconnect. However, shrinking dimensions in both width and

height of the interconnect, together with the growing size of chips, means that wiring resistance is growing with each generation. The rule of thumb is that wiring resistance goes up by a factor of two with each generation. Thus, the other half of the RC equation also contributes to a drag on circuit speeds.

“Wiring problems are being addressed by developing wires of lower resistivity,” noted Ning. “Thus, copper will replace aluminum someday, perhaps soon.”

Another solution is to use larger wires, but that undermines the effort to continually reduce the size of chips from one generation to the next. Another technique, advocated by Intel, is to shrink wiring less in the vertical dimension—but that can lead to manufacturing problems. Siemens Aktiengesellschaft (headquartered in Berlin and Munich, Germany) is working to produce vertically integrated circuits by stacking one wafer on top of another and then wiring the two together.

Beyond lithography, dielectrics, and wiring, manufacturing gigachips faces other potential problems. For instance, just when gigabit memories should be introduced, wafer sizes may increase by 50%, and such changes have usually been fraught with manufacturing difficulties.

The right solutions

If solutions can be bought, they will be. The semiconductor companies spend billions of dollars a year on R&D—allotting a huge investment in basic-technology research. There are signs, however, that the cost of next-generation research is too high for one company, no matter how large. For example, IBM, Siemens, and Toshiba Corp. (headquartered in Tokyo, Japan) have a joint research program under way that is aimed at producing gigabit memories.

Beyond the gigabit generation, there are indications that silicon’s era as the semiconductor of choice may end. However, such predictions have been made in the past. For instance, there were dire warnings about the inability to push silicon below the 1- μm feature size. So far, those predictions of doom were proved wrong. 