

Advanced ICs portend radical changes in system design

Designing systems with the next decade's increasingly complex chips will require entirely new methodologies and techniques. This article is the first in a 5-part series that will explore the changes already taking place in system design and extrapolate those changes through the year 2000.

Steven H Leibson, *Regional Editor*

Electronics occupies a unique engineering niche. Next year's jets won't fly twice as high or use half as much fuel, and except for the pinstriping, next year's cars won't differ markedly from this year's. But semiconductor technology, the driving force behind most of today's electronic design, rides upon a wave of ever-increasing capabilities.

Today, just 40 years after the invention of the transistor and 15 years after the appearance of the first 4- and 8-bit μ Ps, manufacturers have found a commercially practical way to put 1 million transistors on a single chip. Some industry experts predict that by the year 2000, engineers will be working with billion-transistor chips, which represent a thousandfold increase over today's device density. Recent achievements in creating extremely small semiconductor devices indicate that these predictions will probably come true.

Although shrinking transistor geometries are already making increasingly complex systems possible, engineers are finding that the current design methods can't manage the complexity of the new chips. To develop systems with the next decade's denser chips, designers will have to take a completely new approach to system design.

Many forces are driving the development of greater and greater device densities. Data-processing applications seem to have an inexhaustible appetite for memory and processing speed. Space and military applications continue to need devices that combine increased functionality and reliability with reduced size and power requirements. Many earthbound systems, such as medical equipment, require as much reliability as do designs intended for space, and engineers currently achieve this reliability by using redundant circuits. Faster, more complex, better, and more reliable designs all require denser circuitry.

VHSIC program drives semiconductor technology

The Department of Defense initiated the VHSIC (very-high-speed IC) program in 1980 to ensure the ongoing development of faster, denser circuits for military projects. Phase I of the 9-year, billion-dollar VHSIC program prompted contractors to perfect 1.25- μ m processes, based on optical lithography, at a time when 2- or 3- μ m geometries were the state of the art. The VHSIC program requires participants to commercialize the technology they develop.

Meanwhile, many semiconductor vendors that didn't




Fig 1—Only 32 of the 40 available 4k-byte blocks must be good for Inova to create a fully functional 128k-byte static RAM on its S128K8 die. A laser isolates bad blocks in the first layer of metal interconnection before the second layer links the remaining good blocks.

participate in the VHSIC program boosted their internally funded research on process development, in part to keep up with the growing capabilities of the VHSIC contractors. Many advanced commercial parts available now employ the 1.25- μm optical-lithography processes, which have resulted in ICs that achieve unprecedented heights of complexity and speed of operation.

Some IC vendors already offer parts, such as static and dynamic RAMs, manufactured with these advanced optical-fabrication processes. For example, the S128K8 from Inova (Santa Clara, CA) is a 55-nsec, 128k-byte static RAM that packs more than 4 million transistors onto one silicon die. It employs 1.2- μm geometries as well as redundant design to provide acceptable yields (**Fig 1**). The S128K8's design places 40 4k-byte blocks (which Inova calls slices) on the chip. A fully operational RAM requires only 32 functional slices. After the company has applied the first layer of metal interconnection, it performs wafer testing to identify bad slices, and disconnects them with a laser. The second layer of metal joins the remaining good blocks into a functional device.

Advanced process technologies benefit all ICs

Memories are not the only electronic components to benefit from the advancements in semiconductor-processing technology. Both Motorola (Phoenix, AZ) and LSI Logic (Milpitas, CA) will start producing gate arrays built with 1- μm drawn geometries this year. The largest member of LSI Logic's LCA100K family of gate arrays, the LCA100237, contains 236,880 gates. At four transistors per gate, the array holds 947,520—nearly 1 million—transistors. Motorola's HDC000 Max

family of gate arrays includes the HDC105, which contains 104,832 gates—or more than 400,000 transistors—and can operate at system frequencies exceeding 100 MHz.

Although the demise of optical lithography has long been prophesied, continual improvements allowed the technology to carry the semiconductor industry through its first 40 years. Today, optical lithography allows manufacturers to fabricate chips with much smaller geometries than many people predicted it would. However, the 0.7- to 0.8- μm processes used to build the current crop of leading-edge, commercial ICs have finally pushed conventional mercury-vapor lithography to its true limits. (After all, the wavelength of blue light is about 0.5 μm .) To achieve geometries smaller than about 0.7 μm , manufacturers must turn to exposure methods that employ shorter wavelengths.

Beyond optical lithography

In fact, the Department of Defense revamped Phase II of the VHSIC program by reducing the minimum feature size for VHSIC II chips to 0.5 μm (the program's second phase originally specified 0.8- μm geometries), thus forcing contractors to develop alternative lithographic techniques. Researchers are busily creating a bevy of exotic lithographic techniques for half-micron chips.

In conjunction with TRW (Redondo Beach, CA), Motorola developed a 0.5- μm CMOS process for the VHSIC Phase II program. The process uses direct-write electron-beam (e-beam) lithography for two mask steps (**Fig 2**). Motorola built a 1k-bit static RAM as the first test vehicle for this process, and ICs from the



Fig 2—The junction widths as small as 0.5 μm .

device's
Motorola
ber of de
increases
reduced b
feature si

TRW p
further a
Using the
measuring
the compa
mately 34
will build

Enginee
sign pract
doesn't ex
have 34.7 m
extremely
wafers mal
pany incorp
tensive red
cumvent d
power is a
circuitry id
a fully oper
blocks.

An on-chip
The super
spare parts
fails during




Fig 2—The 0.5- μm CMOS process developed by Motorola in conjunction with TRW for a VHSIC Phase II contract produces line widths as small as 0.25 μm .

device's initial manufacturing run worked. By itself, Motorola's 0.5- μm process dramatically boosts the number of devices you can put on an IC. Circuit density increases by a factor of 14.7, although feature sizes are reduced by only a factor of 2.5 in comparison with the feature size obtained from the 1.25- μm process (**Fig 3**).

TRW plans to develop this 0.5- μm process much further as its part of the VHSIC Phase II contract. Using the 0.5- μm design rules on a piece of silicon measuring approximately 2 $\times$ 3 in. on a side will allow the company to build a "superchip" containing approximately 34.7 million transistors (**Fig 4**). TRW says it will build the first superchips before 1989.

Engineers at TRW abandoned conventional IC-design practices for the superchip project. The company doesn't expect to manufacture many superchips that have 34.7 million perfect transistors because of the ICs' extremely large size; defect densities on the raw silicon wafers make such an event unlikely. Instead, the company incorporated systems-level features, including extensive redundancy, in the superchip's design to circumvent defective regions on the device. Each time power is applied to the superchip, built-in self-test circuitry identifies the working modules and constructs a fully operational device from the properly functioning blocks.

An on-chip tool box

The superchip also carries the equivalent of tools and spare parts to repair itself. If a module on the superchip fails during operation, the self-test circuitry on the




Fig 3—Shrinking minimum feature sizes and design rules on an IC from 1.25 to 0.5 μm increases circuit density by a factor of 14.7. (Photo courtesy Motorola Inc)

device can repair the damage by switching in a spare block. The company estimates that the superchip's self-healing capabilities give the component an expected life span of 50 years on Earth-orbital platforms. System specifications and applications often change over such a long lifetime, and engineers can also use the superchip's software-configurable architecture to build systems that reconfigure themselves for new applications.

Although direct-write e-beam lithography is the process of choice for today's experiments with half-micron chips, most IC vendors agree that production lines for 0.5- μm devices will probably use other lithographic processes. Even with the new photoresists developed for e-beam processes over the past few years, e-beam writers still take enormous amounts of time to draw the tiny lines on an IC.

GCA Corp (Andover, MA) has already shipped a wafer stepper that uses an excimer laser as the illumination source. The wafer stepper can fabricate 0.5- μm circuits. Other stepper vendors are close behind GCA. Some IC manufacturers are considering x-ray lithographies for half-micron (and smaller) devices. Further, one of the goals of the Department of Defense's MMIC (Monolithic Microwave IC, pronounced "mimic") program is to make ICs with 0.25- μm geometries manufacturable. The MMIC program is expected to bear fruit in the early 1990s. Semiconductor manufacturers show no signs of abandoning their quest for ever-smaller geometries and denser ICs.

As researchers continue to shrink device geometries,



Fig 4—Using a piece of silicon about the size of a credit card (shown in mock-up form on the right), TRW's "superchip" contains approximately 34.7 million transistors, dwarfing a conventional IC. The superchip incorporates redundancy and built-in self-test/self-configuration circuits that allow the device to circumvent inoperative circuitry.

the mathematical models describing transistor operation start to fall apart. Scientists at IBM's T J Watson Research Laboratory in Yorktown Heights, NY, have fabricated ICs containing NMOS transistors built with 0.07- μm geometries in an attempt to discover whether such small devices would operate as transistors (Fig 5). About 75% of the structures on the test chips were operational.

Because of their tiny feature sizes, the 0.07- μm transistors operate from a 1V power supply and are cooled by liquid nitrogen to combat thermal noise. The thin gate oxide—it's less than 50 Å (or fewer than 20 atoms) thick—necessitated the low power-supply voltage in order to prevent destruction of the oxide by high electric-field stresses. Although IBM's researchers designed their devices for cryogenic cooling, they see no fundamental reason why FETs with 0.1- μm gate lengths can't operate at room temperature as well.

The mathematical models for transistor operation suggested that devices scaled to such small geometries wouldn't function. However, IBM's transistors not only worked, but also exhibited excellent transconductance characteristics. The company claims that its e-beam lithographic process can write patterns as small as 0.02 to 0.05 μm onto a silicon wafer. IBM continues to conduct research on such small structures.

The device density in the TRW/Motorola CMOS process increased by a factor of 14.7 when geometries




Fig 5—Built with a minimum feature size of 0.07 μm , (which is about 10 times smaller than the geometries of today's commercial ICs), this NMOS FET developed at IBM's T J Watson Research Center exhibits excellent transconductance characteristics, proving that such small devices are indeed feasible.

shrank from 1.25 to 0.5 μm (a linear scaling of 2.5). A further reduction in linear feature size from 0.5 to 0.07 μm (a reduction of a little more than 2.5²) could therefore produce an additional density increase of about 216 (14.7²). Apply that increase to the TRW superchip's transistor count of 34.7 million and you arrive at a phenomenal 7.5 billion transistors on one 2×3-in. piece of silicon!

Giga-scale integration

So, by the end of the 90s, the combination of sub-micron geometries, expanding silicon die size, and redundant IC design promise to put billions and billions of transistors on a chip. However, few engineers are thinking about such giga-scale integration (GSI) today. Such tremendous ICs won't find their way into every system—many designers will probably make do with a few million transistors per chip—but you should consider now how GSI will affect your designs and your design methods over the next decade.

According to Mel Thomsen, associate director of

FREE
160-page catalog

Just in time delivery

on over 100,000
electronic hardware parts.

We've been cutting our customers' inventory costs with "JUST IN TIME DELIVERY" for ten years. We stock over 80% of our items for immediate shipment to your scheduled delivery dates. Non-catalog or out-of-stock parts are shipped within a few weeks.

We manufacture our own parts to ensure consistent high quality. When you need special hardware, we custom-manufacture to your blueprint in any quantity, from small lot to high volume.

For quality, selection, competitive pricing, and "JUST IN TIME DELIVERY," depend on RAF. Write or call today for our free catalog, or a quotation.

RAF

Electronic Hardware, Inc.


95 Silvermine Road • Seymour, CT 06483
(203) 888-2133 • FAX (203) 888-9860




CIRCLE NO 12

Data Line EMI Filter

- Virtually eliminates line conducted EMI from data lines
- 12 dBm attenuation from 3 to 500 MHz
- Frequencies below 100 kHz pass unattenuated
- Available in 8-line or custom versions
- Low cost, standard DIP package



Coilcraft
800/322-COIL • Cary, IL 60013



**The future of
system design**

Dataquest's (San Jose, CA) Semiconductor Industry Service, 1 billion transistors on a chip will let you build 128M bytes of RAM, 1000 DEC VAX superminicomputer CPUs, 20 Cray 2 supercomputer CPUs, 10 complete VAX systems with memory, or 1/40 of a complete Cray 2.

That is, you could put such systems on silicon if you had the proper tools. Most system designers don't have the tools necessary to build such incredibly complex ICs. TRW uses a combination of purchased and internally developed tools for the superchip project. Because even a GSI-level device couldn't incorporate all of the circuitry of today's most complex systems, such as the Cray 2, it's likely that at least some future systems will incorporate several GSI devices. Current CAE tools simply can't handle systems of that complexity, and today's design methodologies aren't capable of addressing such a monumental task. In addition, the test philosophies currently employed by most electronics companies virtually guarantee that systems built from such sophisticated components will be untestable, as will the components themselves. Today's chip- and system-level packaging schemes seem equally unable to cope with GSI devices.

All these inadequacies place electronics designers at a crossroads. They must choose whether to keep today's design methods, which will limit them to today's level of system complexity, or adopt new tools and technologies to progress to the next decade's level of system complexity. Current system-design techniques just will not work with tomorrow's IC complexity. Companies that plan to design future systems with the techniques they use now risk being left in the dust over the next decade.

The tools and methodologies that will allow engineers to create electronic systems from extremely complex components are already starting to appear in prototypical form. Although they're not yet ready to tackle GSI components or systems that incorporate them, these developments suggest the system-design trends of the 1990s. The remaining articles in this series will explore these emerging methods, technologies, and trends.

EDN

Reference

1. Meindl, James D, "Chips for Advanced Computing," *Scientific American*, October 1987, pg 78.

Article Interest Quotient (Circle One)
High 497 Medium 498 Low 499