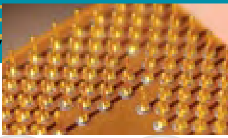

Innovative band engineering techniques target CMOS gate leakage dilemma

By Robert J. Mears, PhD

For Single Print Only

Semiconductor performance has gained enormous benefits from continuing advances in transistor scaling. While drive current enhancement techniques such as strain engineering and hybrid orientation have extended Moore's Law, one of the biggest barriers to device scaling beyond 45 nm is gate leakage, which accounts for significant static power dissipation in high-performance devices. A new platform addresses this pressing manufacturing and design challenge by reengineering silicon.

In the past 40 years, the semiconductor industry has reaped the benefits of silicon scaling to create a market where enormous and predictable advances in transistor performance are occurring concurrently with enormous and predictable reductions in transistor cost, size, and power consumption. Abiding by Moore's Law, the industry has continually doubled semiconductor performance approximately every two years while driving significant advances in areas such as portable electronics by lowering power requirements to extend battery life. Bulk silicon lies at the heart of this transformational industry and today accounts for up to 95 percent of the $250 billion semiconductor market.

Silicon's electrical properties, high-quality native oxide, low-defect density, and ability to scale to large wafer sizes have combined to make Complementary Metal Oxide Semiconductor (CMOS) silicon the industry's dominant design medium. Continual advances in the process roadmap have not come easily, however. As semiconductor manufacturers have migrated to finer process lithographies, silicon's fundamental physical properties and native oxide have been put to the test.

Perhaps the most imposing problem facing circuit designers today is current leakage, especially transistor gate leakage. As the industry moves from 90 nm manufacturing process geometries to 65 nm and below, serious degradation of IC power efficiency is becoming a dominant consideration. At 90 nm, static power consumption from leakage mechanisms, particularly gate leakage across increasingly thin gate oxides, accounts for more than 40 percent of the power dissipation of a microprocessor. At 65 nm, the problem worsens to about 60 percent. By the time the industry reaches 45 nm, the benefits of scaling will fail to boost performance in the face of imposing leakage, power dissipation, and

Perhaps the most imposing problem facing circuit designers today is current leakage, especially transistor gate leakage.

 ©2007 OpenSystems Publishing Not for distribution

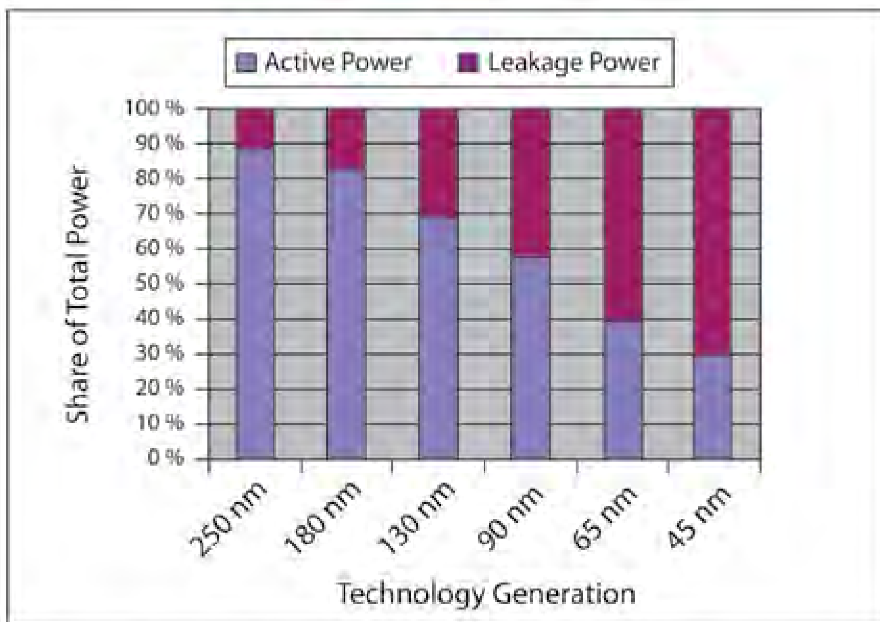

*Source: IC Insights, Intel; 45 nm extrapolation based on internal estimates

Figure 1

tunneling effects. Figure 1 illustrates how leakage is becoming a more significant concern than active power for MPUs.

No new materials

The semiconductor industry has attempted to address these scaling issues by investing in a number of innovative process improvements. The challenge has been determining how to extend the capabilities of silicon and lower power dissipation without requiring major changes in proven semiconductor materials.

At the 130 nm node, manufacturers began introducing commercially viable Silicon-On-Insulator (SOI) enhancements. This technology places a thin layer of insulation such as silicon oxide between the silicon surface and the substrate. SOI lowers parasitic capacitance and junction leakage, and in the process, boosts performance by 20-35 percent. It can also be used to lower power consumption compared to devices running at the same frequency in bulk CMOS. However, manufacturers using SOI technology face significantly higher substrate costs and yield challenges as the technology scales to finer process nodes and thermal-resistance issues arise.

At the 90 nm node, manufacturers turned to strained silicon techniques to address these scaling limitations. These technologies increase or decrease the normal inter-atomic distance between silicon atoms in the wafer plane. The small distortion of the crystal lattice alters the electronic band structure of the silicon. If the inter-atomic distance is increased (tensile strain), electrons travel more freely, but holes are impeded. The reverse is true if the inter-atomic distance is decreased (compressive strain). By combining tensile and compressive strain approaches, the mobility of both electrons and holes can be significantly improved, which in turn boosts chip operating speeds.

But as channel lengths shrink, the benefits of strained silicon approaches diminish. Regardless of which strained silicon approach is used, for instance, drain current enhancement in strained Negative-channel Metal Oxide Semiconductor (NMOS) devices saturates at 25 percent. And as manufacturers move to smaller geometries and scale up stresses, the technology becomes increasingly difficult to implement.

Intel and IBM recently announced advances in high-k dielectric oxides, which deliver better insulation than conventional silicon dioxide to reduce gate leakage. But problems persist with Fermi-level pinning and the generation of higher-threshold voltages, mobility degradation, and reliability issues. Despite these efforts, scalability issues remain. Furthermore, new high-k techniques introduce exotic materials such as hafnium to the manufacturing process, adding complexity and cost to the process flow.

Reexamining band engineering

In the past few years, researchers at MEARS Technologies have developed a new platform that addresses the fundamental obstacles to CMOS scaling by significantly reducing gate leakage while simultaneously enhancing drive current. This cost-neutral process does not use new materials in the fabrication process and is compatible with existing manufacturing infrastructure, including bulk CMOS, SOI, strained silicon, and high-k approaches.

Two well-known observations drove the search for this new methodology. The first was that the electronic properties of any material are primarily governed by its electronic band structure or the relationship between the energy and the momentum (speed) of electrons and holes in the material. The band structure of a material is a function of the arrangement of the atoms, or the lattice, and the electronic nature of the atoms themselves. The curvature of a material's band minima and maxima in turn determines how easily an electron or hole is accelerated. The larger the curvature, the more easily electrons are accelerated with an applied electric field – in other words, the electrons have a reduced effective mass.

The team determined that to improve performance, the lattice of the atoms had to be modified. Normally one would assume that altering the design of the lattice in the plane of the device would improve performance. Instead, the researchers discovered they could manipulate the bands in a planar silicon device by intermittently breaking the silicon periodicity in the vertical direction (see Figure 2). Furthermore, by leveraging their ability to build up the silicon by individual atom layers and taking advantage of recent advances in manufacturing equipment, they could complete this enhancement with a conventional silicon epitaxial step.

The additional silicon layer is fully crystalline but has a strongly anisotropic electronic band structure. Silicon atoms are intermittently spaced apart slightly more in the vertical than the horizontal plane of the device. That distribution supports the creation of *channels* for electrons and holes to travel parallel to the surface more freely in the wafer plane, improving electron mobility for both NMOS and Positive-channel Metal Oxide Semiconductor (PMOS) devices. This effect is in distinct contrast to strain techniques, for example, which require separate tensile strain for NMOS and compressive strain for PMOS.

In contrast, electron flow in the vertical direction is impeded. In this direction, effective mass can be increased and band curvature reduced by up to an order of magnitude. This unique property results in a significant reduction in gate leakage even while the additional silicon layer enhances drive current in the horizontal plane of the device.

For Single Print Only

Multiple benefits

This innovative approach to reengineering silicon, known as the *MST Platform*, or simply *silicon-on-silicon*, offers semiconductor manufacturers a variety of benefits. In-plane mobility increases drive current and can be used to improve transistor performance in microprocessors, DRAM, SRAM, flash and other memory ICs, as well as RF and mixed-signal devices. Some of this increased drive current capability can also be traded for higher-threshold voltages and lower-subthreshold leakage power. At the same time, the technology reduces gate leakage by up to 80 percent, far outperforming strained silicon or SOI technologies, while remaining completely compatible with these existing industry approaches. Early tests indicate that the silicon-on-silicon platform delivers up to 30 percent enhancement in drive current.

The MST Platform is the result of a five-year design effort conducted in partnership with International Sematech's Advanced Technology Development Facility and Lawrence Semiconductor Research Laboratory and has been tested on more than 1,000 wafers. Validated by industry-standard E-test and bench measurements, the new platform has been successfully applied to sub-90 nm processes and is compatible with industry-standard manufacturing equipment. Given its adaptability to existing CMOS infrastructure, the MST Platform is expected to be adaptable to next-generation 32 nm and 22 nm flows. MEARS Technologies' MST integration is shown in Figure 3.

From a semiconductor manufacturer's perspective, this technology offers a number of advantages. It reduces risk by relying on established CMOS manufacturing techniques and requiring no new materials or complex architectures. The technology also simplifies implementation because it can be easily inserted into existing manufacturing processes. It only requires the addition of a few simple steps to a standard CMOS flow used to construct the transistor channel. And it relies on standard equipment. In most cases, manufacturers can implement the technology in six months or less.

The MST Platform is also attractive from a cost standpoint. Given potential die real estate savings (from reducing the widths of non-minimum-geometry devices) and the technology's ability to use existing manufacturing tooling, implementing the technology is practically a cost-neutral exercise. Manufacturers using this technology simply insert the special expitially grown silicon into their standard CMOS flow as a channel replacement layer.

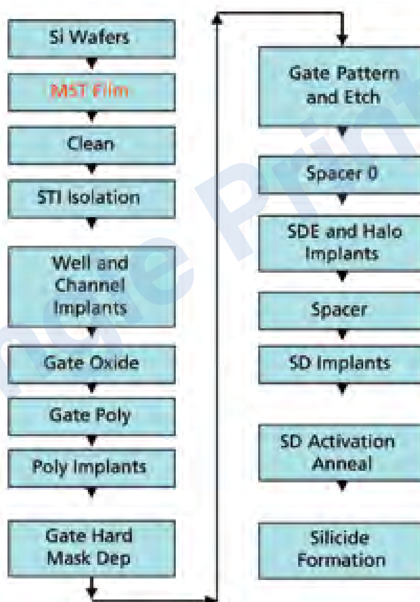

Figure 3

Perhaps the greatest cost impact of this new technology is the investments it allows semiconductor manufacturers to forgo. By increasing drive current and reducing gate leakage, the silicon-on-silicon approach allows semiconductor manufacturers to implement 65 nm and 45 nm CMOS generations without resorting to the added expense associated with alternative technologies. **ECD**

***Robert J. Mears**, PhD, is MEARS Technologies' founder, president, and CTO. Robert has 25 years of research experience in electronics and photonics and is the inventor of the erbium-doped fiber amplifier, which helped enable broadband Internet. He has authored or coauthored approximately 150 publications and numerous patents and is an Emeritus Fellow of Pembroke College, Cambridge.*

To learn more, contact Robert at:

MEARS Technologies
1100 Winter Street, Suite 4700
Waltham, MA 02451
617-219-0600
info@MEARStechnologies.com
www.MEARStechnologies.com

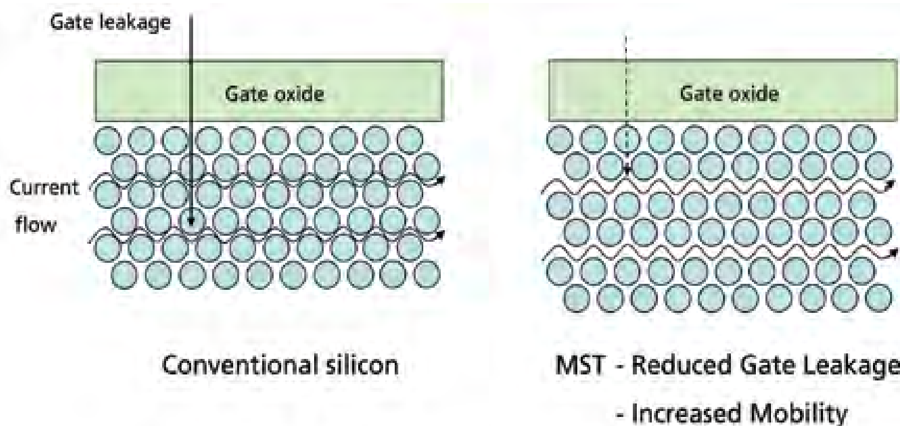

Figure 2

 ©2007 OpenSystems Publishing Not for distribution