

## Tackling power/performance trade-offs with silicon channel engineering

Solid State Technology, July 1, 2007

[http://sst.pennnet.com/articles/article\\_display.cfm?Section=ARTCL&C=Feat&ARTICLE\\_ID=296979&KEYWORDS=Robert%20Mears&p=5](http://sst.pennnet.com/articles/article_display.cfm?Section=ARTCL&C=Feat&ARTICLE_ID=296979&KEYWORDS=Robert%20Mears&p=5)

Since its earliest days in the 1950s, the semiconductor industry has been dependent on silicon technology to meet the price, performance, power, and package demands that have enabled rapid advances in the electronics arena. Even today, nonsilicon devices account for <5% of the \$250 billion chip market. Silicon's electrical properties, high-quality native oxide, low levels of defects, and ability to scale to large wafer sizes have all combined to make it-and in particular CMOS silicon-the industry's dominant design medium.

The silicon industry is seeing an increasing drive to low-power mobile devices, while consumers continue to demand higher performance. Because of this, more applications are migrated to satisfy mobile device requirements. There is also fierce competition in the server market over how the trade-off between power and performance is achieved. Historically, the scaling of transistor geometries has enabled enormous performance gains while holding power in check. But today, the industry faces a potentially run-away power problem. For example, the *Wall Street Journal* [1] declared that server power costs were now outstripping hardware expenditure, while the Internet has driven a doubling of servers worldwide since 2000. With this, the associated power has now increased almost four-fold.


### Challenges

As semiconductor manufacturers scale their processes to achieve higher performance and smaller geometries, the fundamental physical properties of silicon and its native oxide are being put to the test. For example, performance expectations at the 130nm and 90nm process nodes led to the introduction of techniques such as strained silicon and silicon-on-insulator (SOI). However, as the industry moves from 90nm linewidths to 65nm, 45nm, and beyond, power dissipation and operating efficiency are becoming of paramount concern to chip designers.

Consider that at 90nm, static power consumption from various leakage mechanisms, particularly gate leakage across increasingly thin oxide layers, accounts for nearly half the power dissipation in leading-edge processors. At the 65nm node, the problem worsens to ~60%. It is more than a little ironic that CMOS, which was introduced to mitigate power concerns, is now hampered by static leakage. In fact, as device dimensions enter the atomic realm, leakage is but one of a number of challenges facing designers. Yet time and time again, the silicon industry has responded to such challenges, and broken through barriers that many believed to be impenetrable.

## Power/performance tradeoffs

Safely navigating the power-performance trade-off is surely every designer's mantra. Such is the magnitude of the challenge that almost every available technique for power management is likely to find some application. At the architectural level, we have already seen a shift to achieving performance increases via multicore designs, rather than by relying on clock-rate increases alone. And at a circuit level, we are seeing transistor-intensive approaches to power management, such as sleep transistors and an increasing use of body-biasing. But the real target in raising the power-performance stakes is at the material level.



[Click here to enlarge image](#)

*Figure 1. Mobility in a) conventional silicon and b) MST-engineered silicon.*

The Semiconductor Industry Association (SIA) has charted the trends of the industry since 1982. The challenge of gate leakage arising from gate dielectrics  $< \sim 10$  atomic layers thick has long been recognized. Quantum tunneling mechanisms scale exponentially with reducing dielectric thickness, which in turn, makes the conventional scaling methods that have been followed for generations of devices increasingly difficult to implement without incurring enormous static power overhead.

For the 65nm node, designers were forced to depart from the conventional scaling of dielectric thickness and were fortunate that strained silicon techniques were available to deliver higher drive current. Unfortunately, it is becoming clear that strain techniques are increasingly difficult to implement at smaller geometries, and furthermore, that the current enhancement available from strain techniques is saturating, particularly for nMOS devices. High- $k$  dielectrics and metal gates do permit further scaling by increasing gate capacitance and hence current drive for the same physical dielectric thickness, thereby reducing gate leakage.

Unfortunately, high- $k$  dielectrics tend to reduce rather than enhance the achievable drive current. The reverse trade-off to achieve higher performance through higher drive current is more complicated depending on the interplay between interconnect capacitance and gate capacitance. Furthermore, any increase in gate capacitance also leads to an increase in dynamic power. Reliability issues have also dogged early attempts at implementation, but such is the magnitude of the power problem that these techniques are now firmly on the roll-out plans of the likes of Intel and IBM.

Metal gates and high- $k$  dielectrics represent one of the most radical changes in the basic transistor architecture seen in decades, as they involve completely new materials in the fabrication process. To introduce these new materials, the silicon channel and the immediate silicon-silica interface at the bottom of the gate dielectric are left untouched in order to maintain transistor integrity and reliability. In fact, since pMOS strain techniques frequently involve source/drain refill with silicon germanium, the silicon channel is arguably the one remaining “untouchable core” of silicon transistor technology.

Is it possible to maintain a silicon channel while engineering its electronic properties so as to increase performance and reduce leakage? This is the question that we have been posing and answering with more than five years of intense research and development.


### **Silicon channel engineering**

The challenge calls for a basic re-engineering of silicon using a new technique, band engineering, and the application of the quantum mechanical principles that underlie the electronic properties of materials. In other words, since quantum mechanics lies at the heart of the leakage problem, why not make use of quantum mechanics to engineer a solution?

So how does one re-engineer silicon? There are two parameters: the atoms themselves and the lattice arrangement. Since the premise is to continue to use silicon, the only thing left to modify is the lattice. Herein lies the subtlety. At first, one would think that the silicon lattice would need to change in the plane of the device—something that is impossible on an atomic scale using lithography and, as we have seen, is limited to the very small macroscopic changes that are introduced by strain techniques. Fortunately, it turns out that breaking the periodicity in the vertical direction is sufficient to manipulate silicon’s electronic bands in a planar silicon device. Rather than using lithography, this is something that is much more simply achieved via a modified conventional epitaxial process step.

Thanks to the advances that have been made in silicon epitaxy in recent years, it is now possible to build up silicon, atom layer by atom layer. At the same time, quantum mechanical simulation tools have advanced to the point where realistic numbers of atoms can be modeled to determine band-structure differences. For example, by intermittent insertion of a “nonsemiconductor layer” during silicon epitaxy, the resultant layer, though fully single-crystal, has a strongly anisotropic band structure. In the vertical direction, the effective mass can be increased substantially (that is, the band curvature is reduced) by up to an order of magnitude. Thus, electron flow in the vertical direction is impeded as the superlattice layer blocks the vertical conduction between the gate and the channel to reduce gate leakage. In contrast, the same technique reduces the effective mass (curvature increased) in the plane of the device, because electron density is more evenly distributed. That distribution supports the creation of channels for electrons to travel parallel to the surface more freely, which improves electron mobility for both nMOS and pMOS devices. In a sense, the re-engineered silicon is akin to a silicon laminate, with increased

mobility in the plane and reduced (gate) leakage vertically. The technology has been launched as MEARS Silicon Technology (MST) (**Fig. 1**).



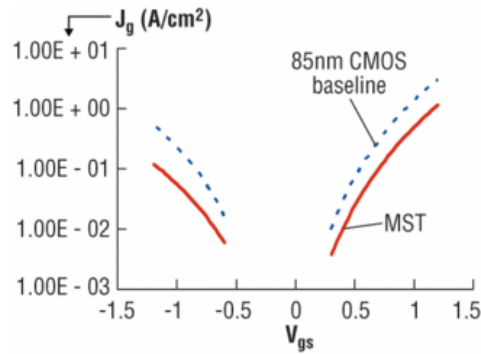
[Click here to enlarge image](#)

*Figure 2. Electron density stratification as a result of insertion of a nonsemiconductor layer during silicon epitaxy.*

*Figure 2*-point) from a quantum mechanical simulation of a small cross-section of the channel region from one such MST design. It can be seen that the conventional silicon orbitals are “squashed” on either side of the nonsemiconductor layer. This leads to a greater connectivity in the plane of the device, which illustrates the increased ease of electron transport in the plane. In a sense, the connected orbitals in the plane of the device form electron “pipes,” much as in carbon nanotubes, but in this case engineered through epitaxy in a conventional planar silicon device. Conversely, the greater separation of the orbitals in the vertical direction results in reduced transport and hence reduced current flow. Put another way, the insertion of the nonsemiconductor layer creates stratification in the electron density, thereby impeding vertical current flow (gate leakage current). The technology benefits arise from the inherent properties of the channel rather than external factors such as applied stress.

### **Impact on electrical properties**

The ability to manipulate the electrical properties of silicon has far-reaching impact. An increase in in-plane electron transport leads to increased drive current and improved transistor performance for microprocessors, DRAM, SRAM, flash, and other memory ICs, and RF and mixed-signal devices. Drive current enhancement in excess of 30% has been observed in nMOS and pMOS devices, and is expected to be additive to other enhancement techniques such as strain engineering. At the same time, gate leakage can be simultaneously reduced by as much as 80% by impeding (gate) current flow in the vertical direction (**Fig. 3**). The technology lends itself to further enhancement, as the epitaxial layer design is optimized for specific applications.



[Click here to enlarge image](#)

**Figure 3.** Gate leakage reduction for nMOS and pMOS with the MST Platform. (Data from ATDF)

MST has been demonstrated on hundreds of wafers run through International Sematech's ATDF facility. The technology can be directly inserted into existing CMOS manufacturing processes and involves no new materials in the fab. By manipulating the electrical properties of silicon without changing the material itself, and in particular without changing the silicon-silica channel-dielectric interface, the technology is fully compatible with semiconductor manufacturers' baseline processes. Potential also exists to further enhance future devices on the *International Technology Roadmap for Semiconductors* using MST. For example, wrap-around epitaxy on a silicon fin can be used to produce a channel-engineered FinFET.

## Conclusion

Quantum mechanical techniques based on band engineering are complementary to existing silicon optimization approaches. Although the band structure can appear quite complicated, it is essentially a function of two things: the arrangement (spacing, periodicity) of the atoms—referred to as the lattice—and the electronic nature of the individual atoms themselves. Both within and beyond CMOS, there are other areas of device design that would benefit from a band-engineered silicon solution. Increasing the band-gap increases the breakdown voltage without the need to move to more expensive materials such as GaAs. Even more intriguing is the prospect for band-engineered direct bandgap silicon that would enable optoelectronic functionality.

## Acknowledgment

MST (MEARS Silicon Technology) is a trademark of MEARS Technologies.

## Reference

1. J. Carlton, "Slowing the Data Center Power Drain," *The Wall Street Journal*, December 21, 2006.

**Robert J. Mears** is president and CTO of *MEARS Technologies*, 1100 Winter St., Waltham, MA 02451, United States; ph 617/219-0600, e-mail [robert.mears@mearstechnologies.com](mailto:robert.mears@mearstechnologies.com).

*Solid State Technology* July, 2007

**Author(s):** R.J. Mears