atomera

Architecting Faster, Cooler Less-Expensive IoT SoCs

The Internet of Things (IoT) is an extraordinarily competitive, crowded market. Any participant needs every possible advantage in order to stand out and succeed. Serious players are betting serious money on their ideas by developing ASICs and systems-on-a-chip (SoCs) for their new IoT devices. But these chips must meet stringent requirements in order to be competitive.

> They must be low-cost.
> They must consume as little energy as possible.
> They will be mixed-signal SoCs.

It's easy to think that the hottest, newest chips should be developed on the most advanced process nodes available at the time. But for IoT SoCs, advanced nodes aren't always the obvious choice. Costs can actually go up – especially when transitioning to FinFET technology. Leakage current can be higher – sometimes much higher. This drains the battery even when nothing is happening – a problem for devices that will spend a good portion of their lives idle. And analog circuits don't migrate very easily from one process node to another.

So, rather than moving to aggressive nodes, there are times when it's better to remain on a well-established legacy node for as long as possible. But, to make this work for increasingly demanding SoCs, there must be a way to extend the performance of the existing node; it must be possible to reduce leakage on that node; and there must be an opportunity to shrink circuits and reduce design pessimism to optimize the cost of the SoC. There is a way to accomplish all of these goals on a legacy technology node, thereby saving the cost and effort of qualifying a new process. The approach, which Atomera calls MST, can result in faster, lower-power, lower-cost SoCs for the exploding IoT market. While useful for extending the life of legacy nodes, it also provides similar benefits to any process node.

A Layer of Oxygen

The key to MST lies in forming a layer – or a few layers – of oxygen within the silicon crystal. This is done epitaxially so that the integrity of the crystal remains intact. Electrons are constrained in these layers and move more easily than they would through the plain crystal; the improved mobility increases performance.

In addition, current in standard silicon is prone to leaking into the bulk, drifting down instead of moving along the channel. With an MST approach, the horizontal oxygen layers keep most current from moving into the bulk from the channel, reducing static leakage.

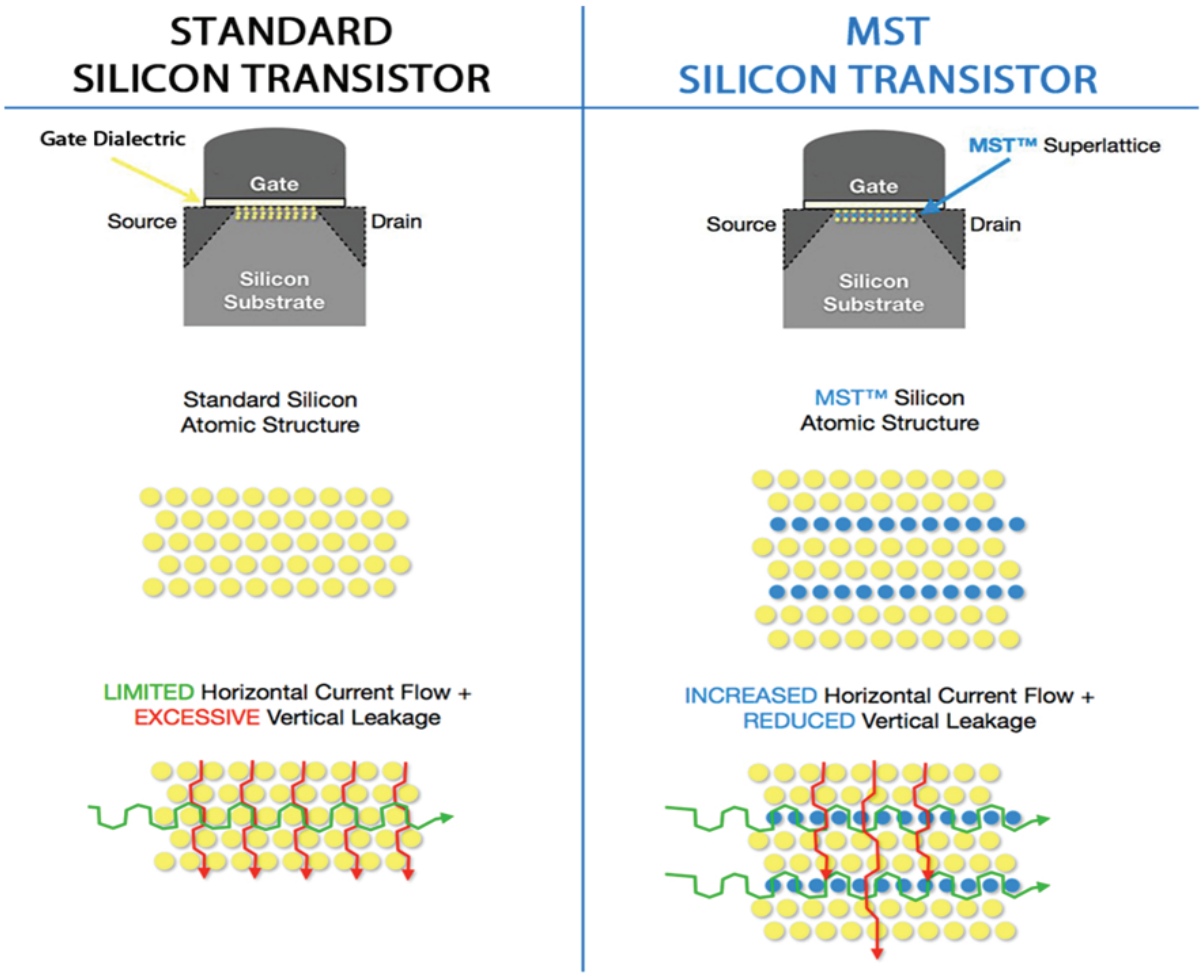

These are the most obvious benefits of MST. But there's another benefit that's just as important, even if it isn't as evident at first glance: MST helps reduce process variability.

Reducing Random Dopant Fluctuation

When doping silicon, dopant atoms are inserted using ion implantation. Even at moderate process nodes, the number of atoms implanted will vary from wafer to wafer and even from die to die. Given further processing – especially high-heat steps like oxidation, those atoms will start to migrate, spreading the distribution.

Some of those atoms will migrate towards the surface of the die, where the main channel conduction takes place. In so doing, they affect the threshold voltage VT and the drive current Idrive. So we have variation upon variation: some random percentage of a varying number of originally implanted atoms will affect these parameters, meaning that each die may be affected in a slightly different way.

Given this level of variation, designers must be pessimistic about the characteristics of the process, and so they will end up overdesigning the chip. This typically increases circuit sizes, driving up the cost of the die. Just as they keep current from drifting down to the bulk, the oxygen layers used in MST technology block the upward migration of dopant atoms after implantation. This helps to keep the dopant profile intact through subsequent processing steps, and it therefore stabilizes VT and Idrive. The parametric variation is typically reduced by around 40% – 50% for the smallest devices. The following Pelgrom plots illustrate the improvement.

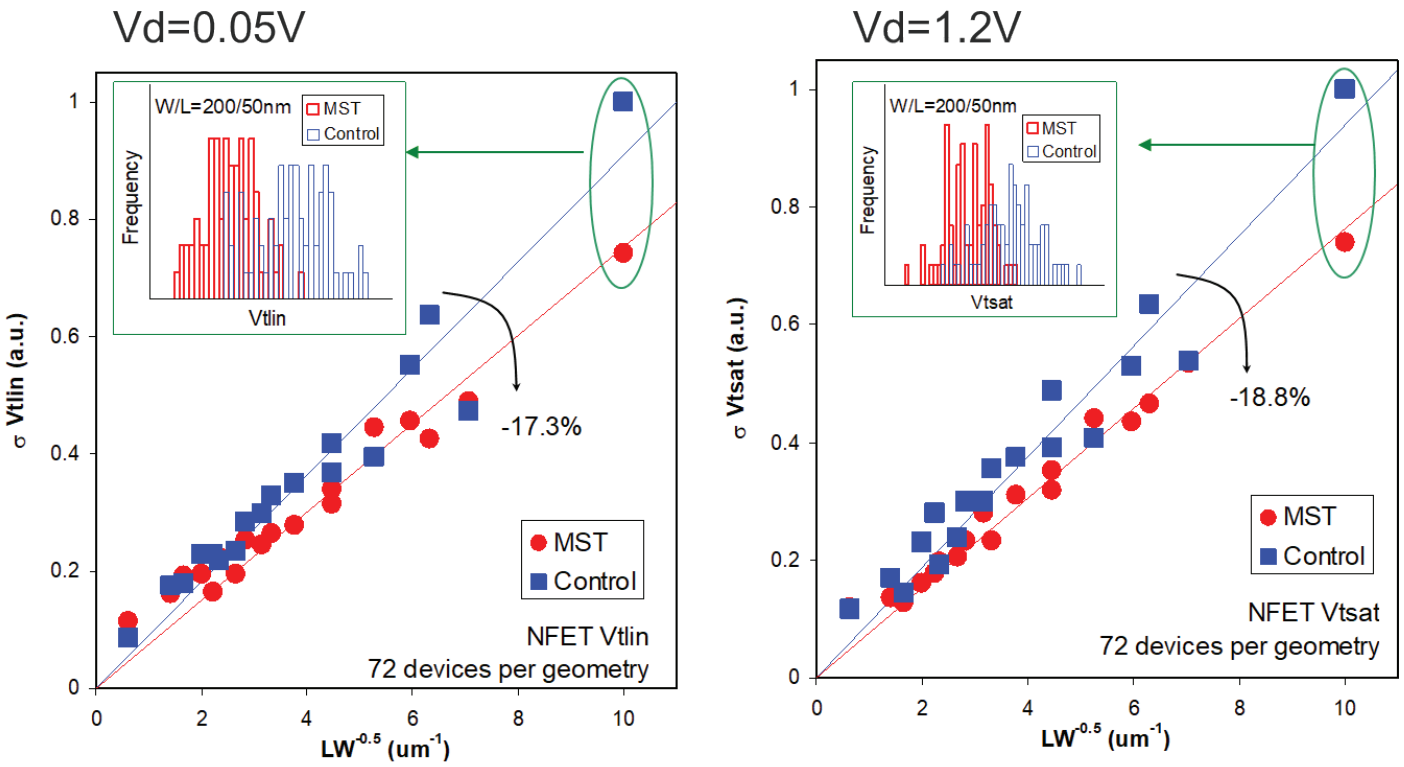

(Source: Atomera)

Lower variation can be particularly helpful for memories. With SRAMs, one can lower the operational VDDmin, reducing energy consumption. In a study of MST, yields were improved by 10% at VDD = 1.0 V, while the VDDmin at the 80% yield point was lowered by 150 mV for a 0.25-Mb array. With a 64-Mb memory having 1/64kb repair, post-repair yields increased by 20% at VDD = 1.0 V. In addition, VDD could be lowered to 0.825 V with 100% yield*. Due to the improved degradation over time, designs required less guardband, tightening up the circuits.

With DRAMs, around half of the chip real estate is consumed by the circuits that sense and control the memory array. The reduced variation possible using MST leads to increased sense-amp sensitivity; this means that smaller sense amps can be used, resulting in a lower-cost DRAM die. Alternatively, the increased sensitivity can be used to enable lower-voltage designs and, hence, lower-power operation.

Meanwhile, typical planar CMOS processes use a halo implant (also called a pocket implant) near the source and drain of a transistor to reduce leakage current. But this non-uniform doping along the channel results in increased 1/f noise (or flicker noise), limiting the performance of the chip and exacerbating the need for design pessimism.

Because the added oxygen layers reduce leakage, the amount of halo doping needed is also reduced, lowering the 1/f noise. The analog circuitry needed for radios and sensor-signal conditioning are particularly susceptible to noise, so any reduction improves performance. In addition, lower variation and lower noise mean that designers can design with less pessimism; it's no longer necessary to overdesign the chip.

Furthermore, lower 1/f noise leads to significantly lower energy consumption for the same signal-to-noise ratio. So designers can be more aggressive in getting the performance they need using smaller circuits. The area savings can significantly increase the competitiveness of the resulting chip while maintaining performance and lowering power.

Getting More out of Older Process Nodes

The MST approach makes it possible to choose any process node for an IoT SoC – in particular, legacy nodes that, without MST, would not provide the needed performance. Legacy nodes have the benefit of leveraging less-expensive fabrication equipment and saving the cost of qualifying a completely new technology. And those analog circuits that have been carefully tuned on the older node can still be used, saving months or even years of work.

With MST technology, any technology node can be made to perform faster and with less leakage. And, because natural variation in the process can be tightened up, designers don't have to hold back when designing for performance.

The MST approach gives IoT SoC designers greater flexibility on choosing the technology node that best meets their needs, generating circuits that are faster, cooler, less expensive, less noisy, and more reliable.

* The Impact of Oxygen Insertion Technology on SRAM Yield Performance, A. Marshall et al, 2017 IEEE Electron Devices Technology and Manufacturing Conference Proceedings of Technical Papers

Scott Bibaud, Atomera CEO