



Chasing Moore's Law with 90-nm: More Than Just a Process Shrink

In the electronics industry, the term "process shrink" is often used to refer to when a semiconductor company migrates an existing design to a smaller process technology. In many cases, this is an upgrade path for reducing the cost, size, and power consumption of chips. An example of a "simple" process shrink is that of Texas Instruments' 720MHz TMS320C6416 DSP. Moving the C6416 DSP from a 130-nm CMOS process to 90-nm resulted in a price reduction of 50%.

If this were the end of the story, it would still be impressive. However, this only describes a single facet of the changes that must come when a new process technology is introduced. A process shrink alone will not enable you to keep pace with Moore's Law. To achieve the greatest gains, it is necessary to innovate in several dimensions simultaneously. To only claim the advantages of a smaller die is to ignore the new levels of SoC integration enabled by finer geometries. The real excitement starts when designers are able to integrate divergent technologies in ways never before possible.

More than Just Faster and Cheaper

As chips clock faster and faster internally, the process technology, which defines the distance between transistors, becomes a limiting factor. To achieve even greater speeds, either the distance between gates must be further reduced or new architectures developed.

With the move to 90-nm, TI was able to manufacture a production-quality C6416 DSP running at 1GHz. However, increasing the clock rate, especially at such high frequencies, requires engineering beyond just redefining an existing design with a faster clock. Running a longer pipeline 40% faster won't result in 40% more processing unless changes are made to facilitate the internal architecture to handle data and processing more efficiently. These changes build upon the new process technology, further exploiting its potential.

In anticipation of moving to 1GHz, Texas Instruments began implementing architectural changes to the C6416 DSP for the migration to 720MHz. These changes were to the low-level micro-architecture and not visible to designers other than as a resulting increase in processing efficiency; the devices are still code-compatible with all of TI's C6000 DSPs.

The micro-architectural changes increased efficiency in a variety of ways. Adding subword SIMD (single instruction multiple data) extensions to the 8-wide VLIW (very long instruction word) instructions enabled more compact code and more efficient utilization of pipeline functional units. The pipeline was also altered to allow certain complex operations requiring up to four pipelined cycles to execute in a single clock

cycle. Reducing the instruction fetch by one cycle opened a space in the pipeline for a register read stage to improve processing efficiency.

Other changes addressed optimizing utilization of the increased clock rate. Since register file size grows in proportion to the square of the number of ports, implementing register files in a dual datapath architecture results in less porting than a single register file. Register forwarding and a pipelined move path between the two paths, as well as a substantial depth of thirty-two 32-bit registers, also increased efficiency. The performance of data functional units were matched within 5% of each other along the critical speed path, meaning that the performance of one functional unit was not overallocated in comparison to the others and thus underutilized.

To get to 1 GHz from 720 MHz, additional changes were required, such as further refinements in the critical speed path and memory pipelines to improve internal processing efficiency, as well as improved clock skewing at the circuit level for signals traveling from one side of the chip to the other.

Without these and other architectural enhancements, the migration to 90-nm would not have produced improvements of such magnitude. Shrinking the process is only one aspect of realizing Moore's Law. Without innovative advancements in architectural design as well, performance gains would be much less substantial.

Designing in 90-nm

TI builds a variety of optimized process flows for each technology generation to provide the best performance for different end equipment requirements. Adjustments to transistor gate length, threshold voltage, gate oxide thickness, and bias conditions all change the performance specification of the millions of transistors on the final integrated circuit. The different flows are carefully targeted to achieve the right application balance between transistor performance and power consumption. For example, the low power process flow for applications such as mobile devices trades clock-speed performance for power efficiency.

These process flows define the design rules and constraints to which chip designers must adhere. Each architecture must be designed specifically to the appropriate process, and the process flows need to be redefined at each process node. Because TI owns its own fabs, TI is able to closely couple process and design engineering to drive process and architecture developments in tandem. As devices grow to millions of gates and silicon features shrink, the linkage between design and manufacturing becomes increasing critical. Additionally, by working hand-in-hand with customers, TI is able to direct research in directions that meet the specific needs of customers for their next generation designs.

Under the Hood

To be truly useful, process migration must be implemented transparently so that designers don't have to understand the details of how the changes are made in order to benefit from them. While the product roadmaps of most companies show process migration years down the road as *fiat accompli*, process shrinking is still an impressive technological feat. Each process node presents new manufacturing challenges that at first

sight may seem insurmountable. However, experts have finally stopped declaring that CMOS has hit the wall, so to speak, because every time this appears to be the case, industry leaders innovate their way past the current set of barriers.

90-nm CMOS process technology is no exception.

TI has greatly improved transistor performance by offering a 2.9 k (OSG) dielectric material at the interconnect level (compared to 3.6 k (FSG) used in the previous generation). Low-k materials reduce capacitance and propagation delays within the interconnect layers of a device while increasing drive current. Given that these are primary factors in transistor switching speed, low-k materials boost overall chip operating frequency and performance. Low-k dielectrics also allow metal lines to be packed closer together on a chip with less risk of electrical signal leakage.

Traditionally when TI introduces a new material or process, it first appears in the ultra-high performance process flow. Then, after TI gains experience with the new material or process, it is fanned out at the next node to the other flows. The 2.9 k OSG dielectric, for example, was introduced in the ultra-high performance process flow at 130-nm. Now it is available for DSP architectures at 90-nm.

Introductions pushing the parameters of 90-nm process technology in the ultra-high performance process flow for Sun Microsystems include a 37-nm gate length using nickel-silicide metal gates and a strained silicon approach with ultra-shallow source/drain junctions to drive performance in both NMOS and PMOS transistors. The shorter gate length, perhaps the most aggressive in the industry, yields higher performance. Nickel-silicide lowers gate resistance while strain induced on the transistor channel increases electron mobility.

While DSPs won't see these enhancements until 65-nm (samples expected in 2005), they are a portent of performance improvements to come: TI expects its highest performance 90-nm transistor to exhibit 50% better performance than its highest performance 130nm transistor. .

Future process enhancements being explored, include a new high-k gate dielectric, Hafnium silicon oxynitride (HfSiON). As transistor dimensions shrink, a higher-k material can prevent the material from becoming so thin that leakage current goes up. HfSiON is being closely watched by the semiconductor industry as the most stable high-k gate dielectric material to date [reference 1].

Another part of the innovation equation is TI's move to 300 mm (12 ") wafers. A 300 mm wafer holds 2.4X the die of a 200 mm, further reducing manufacturing costs, which directly translate to less expensive technology. TI also moved from aluminum to copper at the 130-nm node.

By introducing copper and 300 mm wafers at the 130-nm node, TI was able to leverage the learning curve for these processes when moving to 90-nm while reducing overall risk by reducing the number of transitions made between 130-nm and 90-nm. By building on previous technology and planning ahead with architectural changes, TI has also been able to maintain code compatibility along family lines as clock rates increase. This means that designers who have been developing on the 600 MHz C6416 DSP have a head start of several years on developing for the 1 GHz device.

Excellence in chip manufacturing must also include packaging and manufacturing techniques. TI offers a variety of packages using a lead-free Nickel/Palladium (Ni/Pd) finish introduced to the IC market by TI in 1989. Lead-free ball options are also available on certain high-volume MicroStar BGA products. TI has also been a leader in the development of essential manufacturing techniques, such as optical proximity correction (OPC) to sharpen lithography imaging at advanced nodes.

TI's excellence is reflected in its SRAM technology as well. The six-transistor cell is only 0.97 micron in size. Given that memory is important to DSP applications in that overall performance is tied to how often you have to go off chip, the ability to cost-effectively increase on-chip memory translates to improved performance.

Alternatively, since memory consumes the majority of die space, a smaller memory cell either allows reduction of the die size – and thus cost – or leaves room for adding more functionality. The C6416 DSP, for example, has an integrated programmable Viterbi/Turbo Code coprocessor that is separate and parallel to the execution pipeline. Not only do communications applications gain the efficiencies of a faster clock at a smaller process node, the addition of the Viterbi/Turbo Code coprocessor compounds these gains by offloading Forward Error Correction (FEC) processing and enabling even greater channel density.

Now applications that require greater processing capacity are possible. The ability to handle almost twice as many voice channels in a basestation or on a line card reduces deployment cost resulting in a low price threshold for adopting a technology. Other applications include integrating what used to reside on several chips onto a single die. For example, integrating a digital television or cellular handset onto a single chip reduces both cost and device form-factor.

More than just a process shrink

Digital RF is only one example of how these finer geometries enable the integration of RF, analog, DSP, and memory functions into a single SoC. TI believes that CMOS, with its existing infrastructure, will continue to serve as the workhorse for cost-sensitive, high volume markets.

As illustrated above, the move to 90-nm involves much more than just a process shrink. It is a compounding of many different technology advancements. By investing heavily in developing new semiconductor manufacturing technology and advanced architectures like Digital RF, TI will continue to push the performance envelope without sacrificing cost, yield, or reliability and maintain its leadership by extending the capabilities of CMOS on all fronts:

- Process advancements such as 2.9k OSG dielectrics, 37-nm nickel-silicide gates, strained silicon, , resulting in better transistor packing, smaller transistors, and improved speed of operation.
- micro-architectural enhancements for increased efficiency
- advanced circuit design for improved process flows that provide tuned configurations for performance, density, and/or power consumption based on specific application needs
- Advanced SRAM technology providing higher densities of internal memory

- integrated programmable coprocessors to offload the DSP processor and increase overall processing capacity
- system expertise for developing advanced power efficiency strategies
- manufacturing techniques such as 300mm wafers, optical proximity correction, and lead-free packaging
- industry-leading integration of analog components for SoC to enable the next generation of consumer electronics devices

Reference 1: TI presented a paper on the electrical stability of HfSiON at IEDM 2003 in December 2003.

For more information on silicon development strategies from Texas Instruments, visit www.ti.com/research/docs/index.shtml

Glossary:

FEC: Forward error correction

FSG: Fluorinated Silicate Glass

GSM: Global System for Mobile Communications

HfSiON: Hafnium silicon oxynitride

Ni/Pd: Nickel/Palladium

NMOS: N-channel Metal Oxide Semiconductor

OPC: optical proximity correction

OSG: Organo-Silicate Glass

PMOS: P-channel Metal Oxide Semiconductor

RF: radio frequency

SiGe HBT: Silicon Germanium Heterojunction Bipolar Transistor

SIMD: single instruction, multiple data

SoC: System-on-Chip

VLIW: very long instruction word