




DSPStore
global - interactive - database

e-Newsletter
written for engineers by engineers

This month's DSPStore Sponsor is ...

 **Next Generation Networks & Basestations USA**
Dallas, Tx
3rd- 5th Dec '08 ... defining next generation networks, optimizing this generation

[DSPStore Article Search](#)

Power vs. Performance: The Ultimate DSP Design Challenge

By Doug Morrissey, Vice President and CTO, Octasic, Inc.

For years, Digital Signal Processor (DSP) designers have tackled the daunting task of providing high-performance chips, in a small footprint, with no compromise to the flexibility and software programmability.

As new applications have evolved at a staggeringly high speed, the DSPs provided must match that speed with power, performance and shelf-life endurance to meet today's challenges and be ready to embrace tomorrow's applications. These high-performance multi-core DSPs are increasingly used in telecommunications access, edge and infrastructure equipment to process voice, video and radio signals.

Until now, telecommunications equipment manufacturers used dedicated ASICs or DSP-ASIC combinations to achieve their goals. Today, these new DSPs can replace those cumbersome solutions, and if powerful enough, they can provide the flexibility that older solutions cannot. These flexible solutions are a boon for access and infrastructure equipment that must last many years in network deployments. Given the extended service life of these types of devices and applications, the keys to success are flexibility, adaptability and field-programmability.

In today's technology sphere, ASICs are not as flexible or field-programmable as a DSP, but DSPs can be power hungry, putting gateway designers in a conundrum. But, there's hope: the new generation of multi-core DSPs can be high-performance and power-efficient. The technology exists to achieve this, but first, the issue of power dissipation, the "power crisis", must be overcome.

The Power Crisis

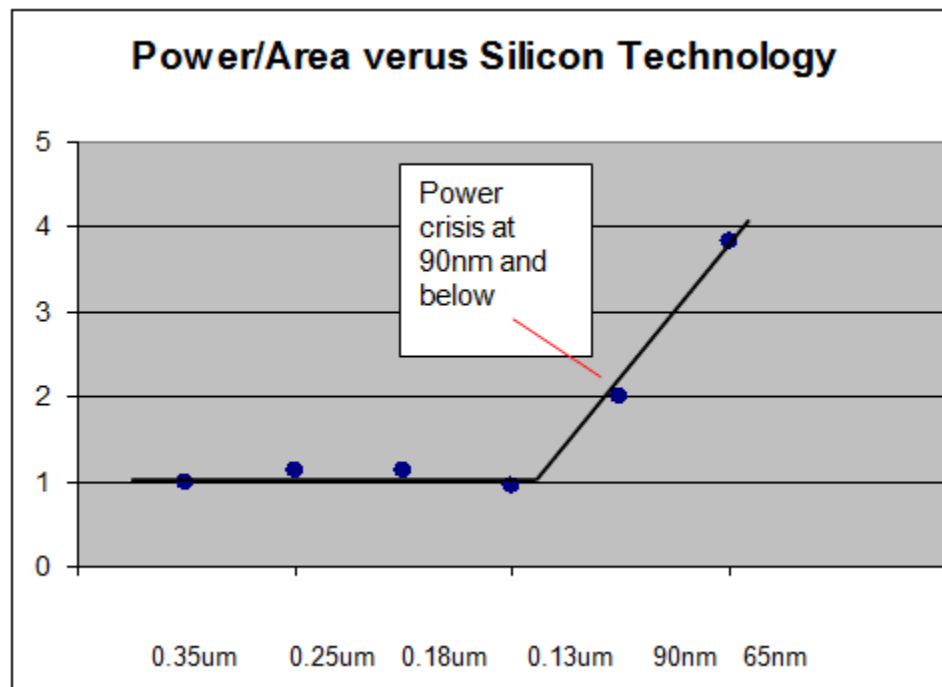
Today's chips suffer power dissipation from two sources: static phenomena in the form of leakage; and dynamic phenomena in the form of switching

operations. This power dissipation is most notable in CMOS technology in the 90nm geometry and below. New generations of DSP design, however, can not only alleviate and circumvent that power crisis, but can actually accelerate processing power in infrastructure, access, and edge equipment while containing power consumption and heat dissipation.

Some of the key metrics that define power consumption within a specific CMOS technology are:

- Supply voltage
- Gate switching speed
- Gate input capacitance
- Gate power consumption
- Dynamic energy consumed per MAC operation.

Studies have concluded that power density, or power per unit area, for an equivalent function (e.g. a MAC unit) up to and including 0.13 μ m geometry, is fairly stable. However at 90nm this parameter suddenly rises.



Until achieving 0.13 μ m technology, DSP designs would benefit from both an increase in performance and a reduction in power, which allowed more circuits to be packed onto a single die. This resulted mainly from the reduction in geometry coupled with a reduction in voltage. Once the technology hit 90 nm, all bets were off.

It's now a matter of trading performance for functionality, an equipment manufacturer's nightmare: more circuits in a single package with lower

performance; or fewer circuits with less functionality.

As the “power crisis” continues, designers have been increasing power consumption to achieve benefits in both performance and functionality. But a new risk occurs: heat dissipation limits have been reached and the resulting problems can already be documented with the latest generation of generic multi-core DSPs on the market today.

Zero-Sum Game: Static Power-Performance

Because performance is the main goal in infrastructure, access and edge applications, zero standby power is usually not a concern for designers. As a result, a general-purpose silicon process has been the preferred approach to optimize performance, rather than a low-leakage silicon option. Choosing low-leakage silicon would provide lower standby power, but also lower speed and performance.

This requires a choice of transistors.

In battery-operated devices, High-Voltage Threshold (HVT) would be optimal, but in infrastructure applications, a Standard-Voltage Threshold (SVT) technology is preferred.

Take, for example, a design using HVT logic operating with a supply voltage of 1.2V, which continually draws 20mW of leakage power. While operating at maximum capacity, it consumes 1W of dynamic power.

An identical design using SVT logic operating at 1.0V delivers roughly the same performance and draws 5 times more leakage power (100mW), but dynamically consumes only 694mW ($1.0^2 / 1.2^2 = 0.694$).

The higher-leakage SVT design will therefore only consume a total of 790mW of power compared to 1.02W for the HVT design. That is a 23% power savings.

Comparison of Power Consumption of HVT and SVT Designs

Power Consumption	HVT design ($V_{DD} = 1.2V$)	SVT design ($V_{DD} = 1.0V$)
Leakage Power	0.02W	0.10W
Dynamic Power	1W	0.69W
Total Power	1.02W	0.79W
Performance	Both designs deliver the same performance.	

Although it is counter-intuitive, this example illustrates that the use of higher-leakage SVT logic versus low-leakage HVT logic can save power overall, due to the high amount of switching activity in the circuit. This is particularly useful for Multiply-and-Accumulate (MAC) circuits but is counterproductive when used on circuits with low activity factors like RAMs and test circuits. As a result, SVT logic is the way to go for infrastructure, "always-on" equipment.

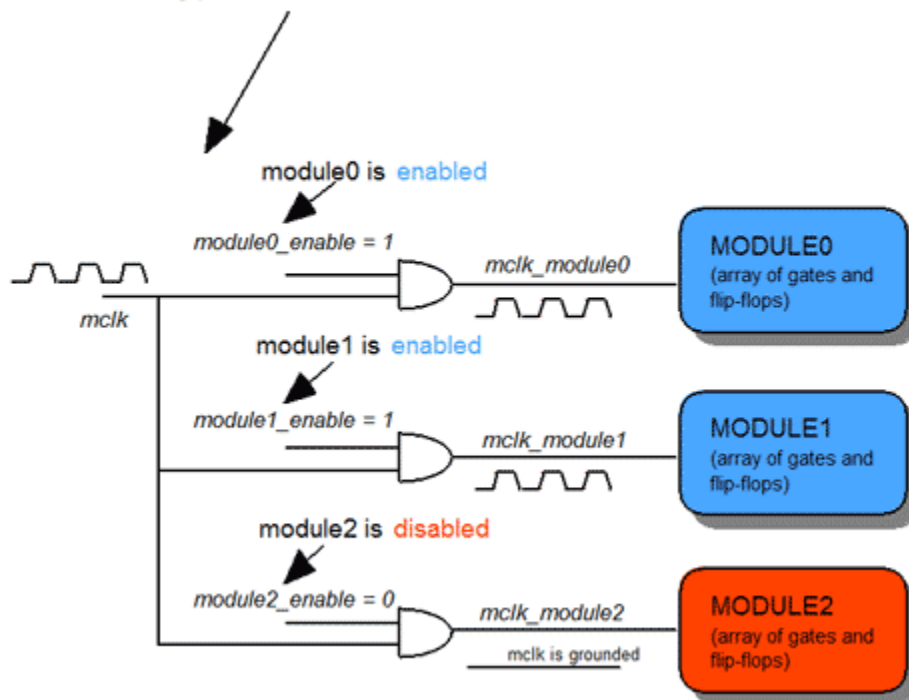
Being Dynamic: Power-Performance Optimization

Both clock trees and logic transitions cause dynamic power consumption, and must be dealt with in the new generation of multi-core DSPs. Power-performance metrics can be substantially improved by aggressively optimizing the design of these two power culprits.

Clock trees (the network of wires and buffers used to deliver a synchronous clock to flip-flops of a design) suck some energy from a chip during the flip-flop operation itself. Power is also expended in charging and discharging the often massive clock trees that span modern high speed chips. In addition, some new generation DSPs use faster clock speeds of 1GHz and higher that require larger and more power-hungry drivers. These larger drivers are required to minimize clock propagation delays through the chip die and associated skew. This results in the consumption of an even larger amount of power.

Clock Tree Gating for Reduced Power

An unused module can be disabled anytime using an enable signal. Associated logic and clock trees contained in a disabled module will therefore stop consuming power.



Equipment designers can reduce the power consumption in clock trees by combining sophisticated techniques that include:

- Individually clock-enabled flip-flops to restrict flip-flop operation to the times when clocking is required.
- Gated clock trees to dynamically prevent clocking entire circuit sections when not in use.
- Multi-cycle path designs to reduce the number of flip-flops in circuits as well as the frequency at which they are triggered.
- Combining computational circuitry whenever architecturally feasible, where a series of MAC operations can be implemented in a cascaded combinatorial circuit rather than in a synchronous feedback circuit. Borrowing on the multi-cycle path technique, this approach substantially reduces the number of flip-flops used and the frequency at which they are triggered.
- Minimizing the area used by flip-flops and circuits to have physically smaller clock trees requiring smaller drive buffers.

Finally, eliminating clock trees altogether would drastically reduce power while improving performance. Clock-less design techniques can be implemented in portions of logic circuits where the most power is consumed. Forward-thinking designers are actively pursuing these solutions. Clock-less design is the most efficient, cost-effective means of achieving the continuing battle of performance versus power.

Optimization of Logic Transitions

Logic transitions play an important role in power consumption, because overall power is consumed in charging or discharging it to its new state. A combination of sophisticated techniques can be applied to minimize the power consumed in logic transitions.

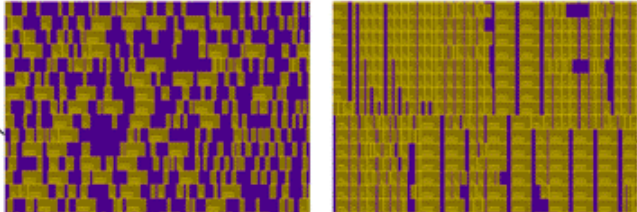
- **Optimize physical gates:** This technique can deliver the highest return in power-performance metrics, especially with smaller geometry technologies. Although very simple in principle, implementing this technique is a challenge when using today's layout tools and methodologies, originally developed to speed time-to-market, that have traded performance for higher level, more complex designs.

Ultimately, physical gates were invented to allow an abstract language (e.g. VHDL) to create silicon based on a designer's functionality goals. There are upsides and downsides to this technique. The standard methodology today removes the designer from the details of how the physical implementation will be accomplished, which can be a time-to-market benefit.

The downside is that designers of a complex chip lose control of their design, including the length of wires that can substantially increase the total capacitance of a circuit. Human designers are still better than design tools at finding the optimal placement of wires and circuits. If armed with sophisticated tools and an intimate knowledge of the design, a human brain can still exercise better judgment. The designer can also readily see where minor changes in the synthesized logic may slash interconnect wire lengths in half. In fact, anecdotal information demonstrates that physical gate technologies, with human intervention, can reduce average circuit wire lengths by a factor of up to two (compared to the same design implemented in conventional, state-of-the-art automatic back-end tools). Furthermore, circuit compaction stemming from strategic placement easily yields 90%+ silicon usage efficiency. That represents approximately a 20% increase over results with automatic back-end tools. I

In addition, the gates that drive these very short wires generally end up being smaller in size and power consumption compared to those in an automatically placed and routed design. As a result, entire circuits are smaller, operate faster, and consume a fraction of the power compared to their automatically-placed counterparts. Using only slow HVT logic elements in 90nm technology, this circuit compaction technique allows entire data-path engines to run at 1.5-2GHz while consuming up to four times less power than equivalent, conventionally-designed circuits.

Design placed by conventional back-end tools: 70-80% silicon usage efficiency



Designer Optimized circuit placement: 90% silicon usage efficiency

Advantages of Optimized Placement: Circuit Compaction and Power Reduction.
The gates are illustrated in yellow, unused silicon is shown in purple.

- **Optimize routing of long signals:** When used in conjunction with other power and speed circuit elements, long signal routing can supply significant performance improvements. For example, data busses use long routes and frequently change states. Reducing the overall capacitance of such lines considerably reduces power consumption, increases speed and reduces the need for buffering. However, the designer's challenge is to place long signals with a larger spacing to reduce capacitance, while still allowing the router to close extremely dense portions of the design. Some tools and methodologies include:

- **Eliminate circuits that change states uselessly:** Disable any circuit whose changed output is not used. This can be done through the use of clock-gating.

- **Reduce the amount of high-frequency gates:** PC processor chips, like Pentium™ and others, have proven that higher functionality comes at the price of increased power consumption. This exponential increase in power consumption results from increasing a circuit's performance using one or more of the following techniques:

- The use of more complex circuits (i.e. look-ahead vs. ripple-carry adders) which consume more area and power;
- The use of larger gates, buffers, and drivers to speed transitions, which yields diminishing returns.

Often, an equivalent performance can be readily achieved using much simpler and slower circuits that operate in parallel or in a staggered multi-cycle path, with much lower power consumption. And although it may seem counter-intuitive, such circuits often occupy a smaller overall area. Indeed, even when used in parallel, their aggregate implementations are often smaller. That's because individually they use fewer and smaller gates per instance than those required for a faster, bigger, more power-hungry monolithic circuit.

- **Reduce the size of voltage transition swings:** Further reducing power consumption can be achieved by reducing the voltage transition swings with long buses and clock lines. It involves the use of balanced transmission line technology with smaller voltage swings like those used in high-performance memory design (such as differential amplifiers). These transmission lines operate with smaller voltage transitions, greatly reducing power consumption. Even though this technique usually requires intermediate voltage rails/planes within the chip, these transmission lines can change states ten times faster than conventional CMOS rail to rail circuitry, for the same power consumption, dramatically improving power-

performance metrics.

- **Scale voltage operation:** As designers spec their systems, they should incorporate some moderation. Not every element within a system needs to scream with performance, particularly if it is not one of the 10 percent of functions that could paralyze the entire system. In fact, running the other 90 percent of functions in the leanest way possible is acceptable. Therefore, designers should apply different voltage rails to different parts of circuits. For example, 10 percent of a chip's circuitry could be fed with 1.2V to run at 3GHz, another 40 percent at 1.0V to run at 1GHz, while the remaining 50 percent is fed with 0.8V to run at 400MHz. The aggregate would provide the best overall power-performance metrics achievable for that particular application.

Harnessing the Power-Performance Beast

As applications become increasingly diverse, and tools become increasingly complex, designers of telecommunications access and infrastructure equipment are scratching their heads about how best to build a high-performance product, at the right price point, with a service life that makes sense. However, the ever-increasing compartmentalization and specialization in chip design methodology keeps these techniques out of reach for many products. It is especially difficult for chips designed with state-of-the-art, back-end design tools by large teams of specialized engineers. Happily, there is a wide array of techniques to manage chips' power-performance metrics to achieve as much as a 3:1 MIPS/power ratio. These techniques range from very simple to extremely complex and offer a wide spectrum of possibilities for improvement.

Surprisingly, the most efficient techniques, such as optimized placement and routing, based on the designer's best judgment and smarts, can be relatively simple when given tools designed for the purpose.

Doug Morrissey, Vice President and Chief Technology Officer

Doug Morrissey is Vice President and Chief Technology Officer at Octasic and has over 10 years of experience in the definition and marketing of semiconductor devices. Joining Octasic in 1999, Morrissey strategically focuses on issues with regard to the technical evolution of future Octasic products within the Voice over Packet market. Prior to joining Octasic, Morrissey worked as Marketing Manager for ATM and DSL products for Agere (formerly Lucent Technologies, Microelectronics Group). Previously to that, he was Senior Systems Architect at Unisys Corporation. Morrissey holds a BSc from Rochester Institute of Technology.

About Octasic

Octasic Inc. is a global provider of media processing silicon and software solutions for the converged carrier, enterprise and end-point communication equipment markets. The company's leading quality VoIP DSP solutions are based on Opus, a unique asynchronous DSP architecture. Octasic allows next-generation equipment manufacturers to significantly reduce system costs by offering unmatched performance in

terms of density and power consumption. Founded in 1998, Octasic is a privately-held company headquartered in Montreal, Canada. For more information, please visit www.octasic.com.

