


Taking the Leap to 7nm: Risk or Reward?

By **Arvind Vel**,
Director, Application
Engineering,
ANSYS


Benefiting from the advances in semiconductor process node improvements is no easy feat. A host of design challenges and risks accompany every new generation of process nodes. The 7nm node is by far the most advanced available today from all leading foundries. This process node fundamentally challenges the basic assumptions of margin-based design. At the same time, it also pushes the boundaries of power and performance benefits that are available to designers. With the 7nm node being the industry driver for the near future, designers are fundamentally rethinking their design implementation strategies using upfront simulations.




▲ Accurate chip–package thermal analysis using ANSYS RedHawk-CTA

Power and performance have always been competing vectors in advanced system-on-chip (SoC) semiconductor design. The benefit of using a 7nm process node is the ability to operate at a much lower supply voltage (sub 500mV) without compromising performance, primarily due to the increased drive current per unit transistor. In addition, the leakage characteristics for these nodes are much lower than planar CMOS transistors. These defining characteristics of the 7nm process node make it a great target for advanced SoC designs for integrated device manufacturers (IDMs) as well as fabless semiconductor companies.


“With the 7nm node being the industry driver for the near future, designers are fundamentally rethinking their design implementation strategies using upfront simulations.”



- Lower noise margins**
- Lower voltage drop immunity
 - Validate all operating modes



- Lower gate breakdown voltage**
- Narrow ESD design window
 - Mandates simulation-based checks



- Higher I, lower EM thresholds**
- Complex EM rules
 - Power and signal EM analysis needed
 - Multimode EM sign-off needed

▲ Key reliability issues at 7nm

Historically, the move to a smaller process node is accompanied by a decrease in cost per transistor in a unit area. In the 7nm process node, the geometry scaling trend has plateaued, and the complexity of lithography masks has increased, making the cost per unit transistor significantly higher than that of previous generations.

Fabless semiconductor companies need higher margin products or significantly larger volumes to offset the cost involved in 7nm design. Additionally, the performance of the chip must add significant value to the product itself. Under these circumstances, the cost of failure for a 7nm chip is extremely high, and fabless companies are always weighing the risk-versus-reward balance in moving to this new process node.

CHALLENGES AND RISKS

Power noise closure remains one of the biggest challenges for 7nm designs. To leverage the quadratic scaling of supply voltage on dynamic power, design teams push the boundaries to operate as low as 470 mV. At the same time, the threshold voltage (VT) has remained constant over the past few process nodes. This combination of decreased supply voltage and nearly unchanged threshold voltage has led to a rapid decrease in

operating noise margin for the 7nm node. Traditional margin-based approaches for voltage drop sign-off quickly break down under these circumstances.

Supply noise mitigation requires a ground-up design philosophy at 7nm. Selecting the right logic library, power grid architecture with metal stack, clocking scheme and the appropriate IC package, have a profound effect on the noise immunity of the design. The challenges in having proper simulation coverage to capture noise-induced failures are more insidious in nature. Starting with a robust power grid to satisfy all margin requirements can be a tempting mistake. This can have consequences downstream on timing and routing closure, therefore impacting die size. Using a simulation-backed power grid architecture that covers all operating modes for local and global switching

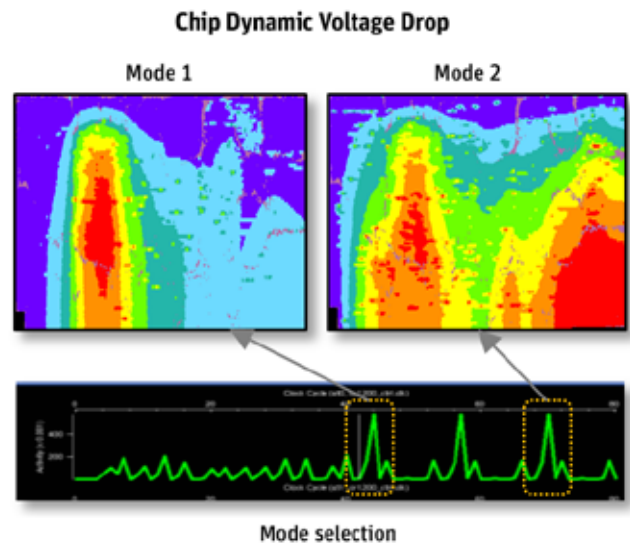
“With the 7nm node being the industry driver for the near future, designers are fundamentally rethinking their design implementation strategies using upfront simulations.”

failures is mandatory. To boost sign-off confidence, hundreds of logically consistent scenarios need to be profiled, and the ones with the best coverage metrics should be simulated. Coverage metrics can involve switching behavior, peak currents, effective resistance, timing slack or a slew of other parameters. This is no easy task given the complexity of today’s SoC designs. Platforms such as ANSYS RedHawk-SC can profile hundreds of scenarios using a vector-based or vector-less approach to capture design weaknesses.


Reliability simulation for electromigration (EM) and thermal analysis is another challenging area for 7nm design sign-off. The FinFET device architecture fundamentally limits the thermal conduction pathway from the fin structure to the silicon substrate for every transistor. At the same time, the vertical thermal coupling between the base layers and metal routing has also increased due to higher metal densities at 7nm. These thermal characteristics lead to localized self-heating for both devices and metals, accelerate lifetime degradation and impact device performance.

The amount of time designers spend on fixing EM violations has been creeping up with every new process node — 7nm is no exception. This is mainly due to the margin-based EM sign-off approach using the worst-case temperature. Overdesign by using wider metals and more vias to solve EM violations has an unintended consequence of poor routability. Using ANSYS RedHawk-CTA for thermal-aware statistical EM sign-off, designers can reclaim valuable chip area and time spent on fixing false EM violations. This can significantly reduce the risks for schedule slips.

Device aging and variability is yet another area of importance for the 7nm process node. Understanding the effects of negative bias temperature instability (NBTI) and hot carrier injection (HCI) on device aging and performance is an important check for all FinFET nodes. Studies have shown a strong relationship between signal probabilities and aging-based performance degradation. Capturing this behavior requires detailed modeling of NBTI-aware libraries and signal probabilities across multiple workloads at the SoC level. ANSYS Path FX and ANSYS Variance FX can perform variability-aware and



▲ Multimode, multicycle vector coverage using ANSYS RedHawk-SC




▲ ANSYS SeaScape platform for chip–package–system convergence

aging-aware timing simulations to increase sign-off confidence for the 7nm node.

MODELING USABLE ACCURACY

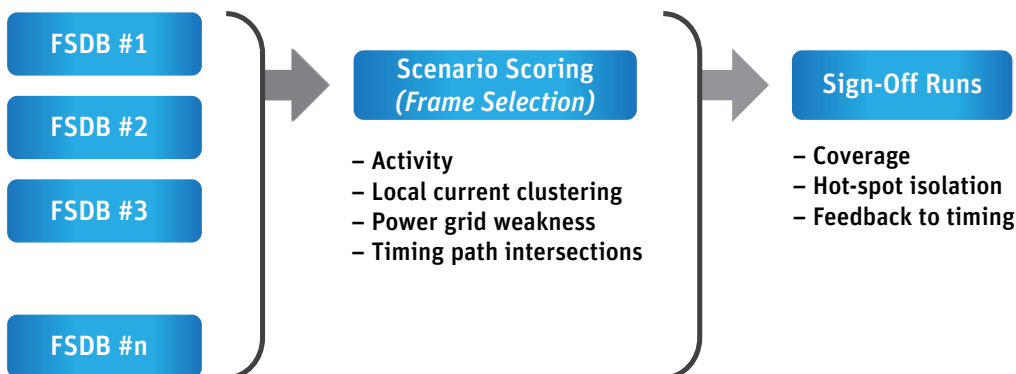
With diminished noise margins at 7nm, accuracy in modeling foundation IP and logic libraries cannot be compromised. Standard cells should be characterized at a wide range of voltage levels to capture the ultra-low-voltage effects on current consumption. Multi-bit flip-flops (MBFFs) should have accurate bit-wise currents with support for spatial allocation of current sinks. Larger analog IPs and memories should have accurate transistor-level detail for proper current distribution. Accurate channel modeling of the on-chip power delivery network along with the package and board model is also mandatory for 7nm sign-off. On-chip extraction tools should have support for multi-color-aware double and triple patterning rules and be foundry certified. Additionally, package modeling should support full-wave accuracy at a bump resolution to capture spatial variabilities. Platforms such as ANSYS RedHawk

Histogram of peak-to-peak voltages across 26 switching combinations of 4 CPU cores



<10 hrs to evaluate more than 3,000 switching scenarios

▲ Rapid design exploration for thousands of scenarios




▲ Critical vector selection based on multiple parameters

“One benefit of using a 7nm process node is the ability to operate at a much lower supply voltage without compromising performance.”

and ANSYS RedHawk-CTA provide foundry certified extraction rules as well as the necessary accuracy for package modeling.

BREAKING THE SILOS

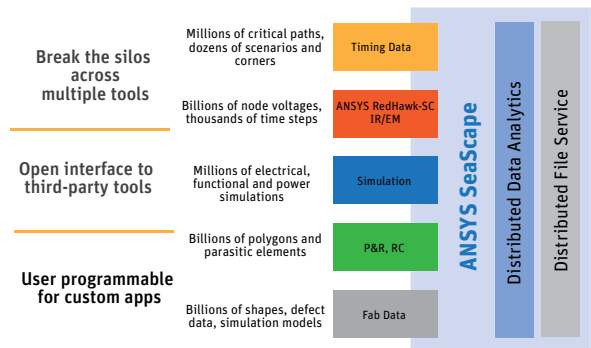
Timing, functionality, power noise and reliability closure are all different stages in the verification process of a chip and typically have a siloed approach toward sign-off. Margining and overdesigning for these verification steps have worked well at older technology nodes without much impact. However, for 16nm and 7nm nodes, the amount of time spent on these independent verification steps has gone up significantly, affecting actual project schedule and cost. For instance, the interdependency of power noise on timing closure has gone up significantly at 7nm. Similarly, the interdependency of self-heat on EM or device aging on timing closure has also increased. If design teams are not adopting new methodologies, verification silos will quickly risk schedule and cost. Platforms such as ANSYS SeaScape that can provide the rapid simulation coverage and bridge the analytics gap across all these silos will be necessary to avoid overdesign at the 7nm node.



▲ ANSYS RedHawk-SC platform

ANALYTICS TO THE RESCUE

The ability to create actionable information from vast amounts of simulation data is an important part of 7nm design closure. Traditional tools lack the ability to morph metrics from different tools into a change action. For instance, trying to avoid low-timing slack paths going through a high-dynamic voltage drop area would require the knowledge of all timing paths along with all voltage drop scenarios. Similarly, trying to find instances that have a high peak current and a high resistive path requires knowledge of the effective resistance map along with the instance load versus current models. Using a targeted design-fixing approach with these analytics is a much more efficient way to fix issues at 7nm than a shotgun approach. Design teams can fundamentally improve the way design closure is done using a purpose-built analytics platform such as ANSYS SeaScape, which is built on a scalable big-data architecture, for multiphysics optimizations.



▲ Breaking the silos using the ANSYS SeaScape platform

SUMMARY

The risk versus reward in moving to the 7nm process node can be quite daunting for IC design teams. Designers must fundamentally rethink implementation and verification methodology to tackle multiphysics problems through innovative analytical means. Breaking down the silos between power noise, timing, thermal and reliability issues is the only way forward without overdesigning. Design teams must retool and rethink design processes to meet the inevitability of process migrations. ⚠



ANSYS SeaScape — A Big Data Approach to Complex Chip Design
ansys.com/seascape-big-data