September 2008
Issue 1

SPECIAL TECHNOLOGY REPORT

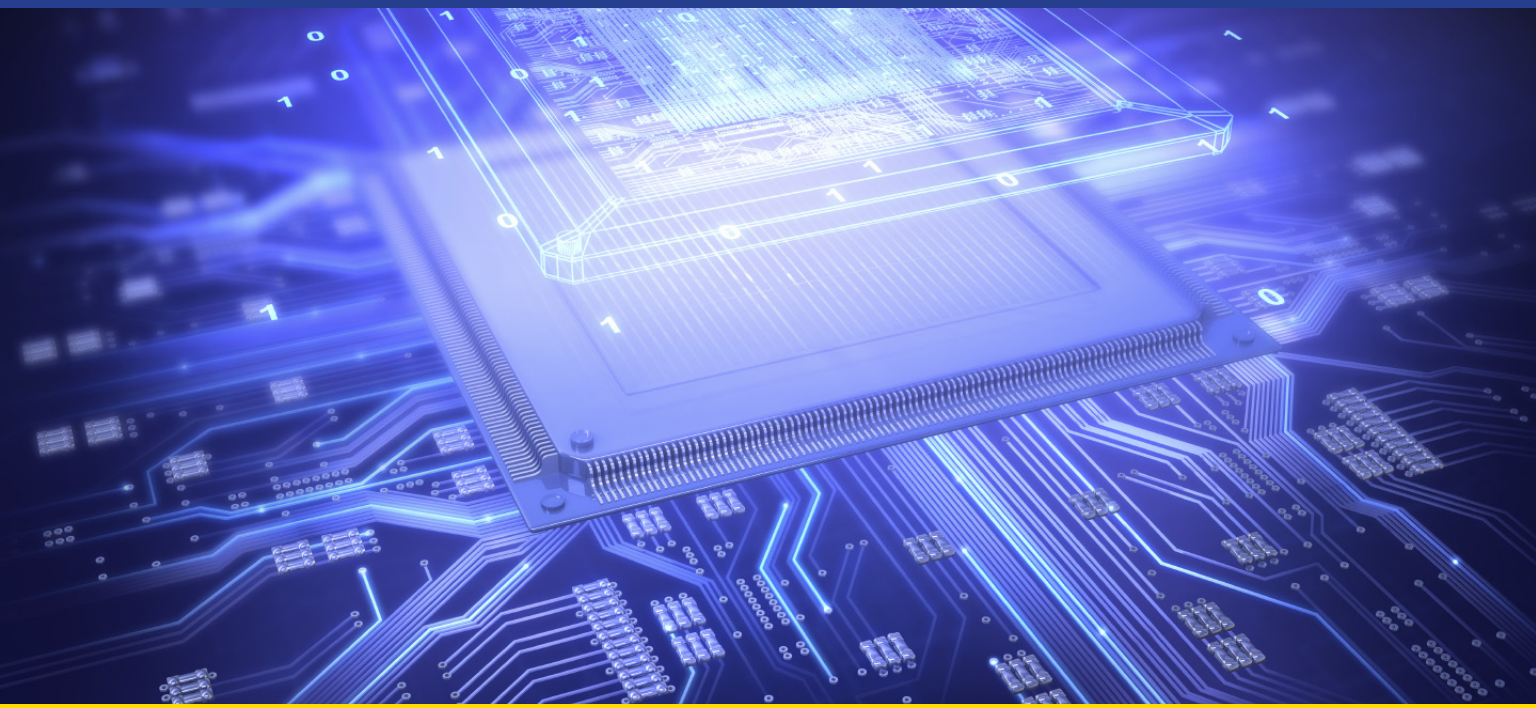

Low Power Design

AUTOMATING LOW POWER DESIGN – A PROGRESS REPORT

Richard Goering

A great deal of attention has been paid to low-power IC design, and for good reason – power consumption has become a critical, if not the most critical, issue in system-on-chip (SoC) design. While progress has been made, designers note that much more remains to be done to automate the low-power design and verification flow, especially at higher levels of abstraction.

This report identifies the bottlenecks that designers are finding in the low-power design flow, including the lack of automation at the electronic system level (ESL), and talks about some of their successes. Drawing primarily upon user experiences, it discusses the following topics:

- ESL to GDSII challenges
- Planning and implementation of power management techniques
- Architectural issues, including memories and embedded software
- ESL low-power flow – capabilities and limitations
- RTL through physical design
- Challenges of low-power verification and test

OVERVIEW: ESL TO GDSII CHALLENGES

The biggest low-power design bottleneck, many design managers say, is at the architectural or electronic system level (ESL), where the potential power savings are the greatest and where the methodology and tool support is weakest. Design teams are struggling with spreadsheets and back-of-envelope calculations to make power-critical decisions about hardware/software partitioning, processors, memories, clocking, and use of multiple power domains. Power modeling, especially at the system level, lacks a standardized approach and may provide insufficient information.

In the RTL-to-GDSII implementation flow, things look better – many IC designs have successfully incorporated power management

techniques such as clock gating, power gating, multi-Vdd (multiple supply voltages), dynamic frequency and voltage scaling, and body bias. The implementation of these techniques is not fully automated, and evaluating potential tradeoffs is not easy, but there are some power analysis and optimization tools that can help.

EDA vendors, meanwhile, have created two formats – Common Power Format (CPF) and Unified Power Format (UPF) – that allow users to express power intent in a consistent way throughout the RTL-to-GDSII flow. Both are being successfully deployed, and they're easing low-power design and verification challenges, but designers are caught in the middle of a standards war between the two competing formats.

If we're talking about RTL-to-GDSII digital hardware logic design, low-power design is doing fairly well. But beyond that realm, several huge gaps remain. Embedded software and operating systems can have a huge power impact, but little help is available to evaluate it. Memory is a major consumer of dynamic and leakage power, but is easily overlooked. Analog/mixed-signal blocks need to be power-aware, but tools and methodologies are lacking. And power management techniques, especially those that involve multiple power domains, profoundly impact and complicate functional verification. "If you don't factor in power throughout the whole development cycle, it is just not going to be optimized," said Ron Tessitore, vice president of advanced application platforms at Qualcomm. "If you don't architect it properly, you are not going to be able to achieve the power target. Technology selection is really important, and so is how you develop the software and write the algorithms. Finally, there's the optimization cycle, including test and refinement."

Asked about the biggest gap in the low-power design flow, Tessitore responded as did several other design managers questioned for this report: "The architectural, electronic system level, so you can make the right tradeoffs up-front for your low-power architecture." A second major gap, he noted, is the lack of support for power-aware software development.

"Today, decisions are made at RTL or sometimes below," said Philippe Magarshack, group vice president and central CAD and design solutions manager at STMicroelectronics. "We want to go above that to an architectural-level definition."

"Assessing power at the architectural level is still an ad-hoc approach, based on spreadsheets where IP [intellectual property] level numbers are captured and manipulated based on application use cases," said Herve Menager, architect for design methodology and flows at NXP Semiconductor. "Productivity improvements in implementation allow more trials of alternative architectures. So now the bottleneck appears at the architectural decision level."

Designers struggle with an apparent conflict between high-level decision making and accuracy. "The higher the abstraction level, the greater the leverage to save power," said Jürgen Karmann, senior staff engineer for design methodology at Infineon. "On the other hand, power analysis is most accurate on the final physical layout."

There are other concerns. David Hui, AMD fellow, said that the biggest low-power design challenge today is accuracy. This shows up in two ways, he said; the difficulty of analyzing power in the face of leakage variations, and the challenge of developing accurate stimuli that represent real-world operations.

Designer priorities may differ, but there is no controversy over the need for power management techniques at 90 nm and below. According to the 2005 International Technology Roadmap for Semiconductors (ITRS), the trend towards higher power density is far outstripping the ability to meet power design requirements (see Figure 1). The growth of leakage (static) power as process nodes shrink is a major reason for that.

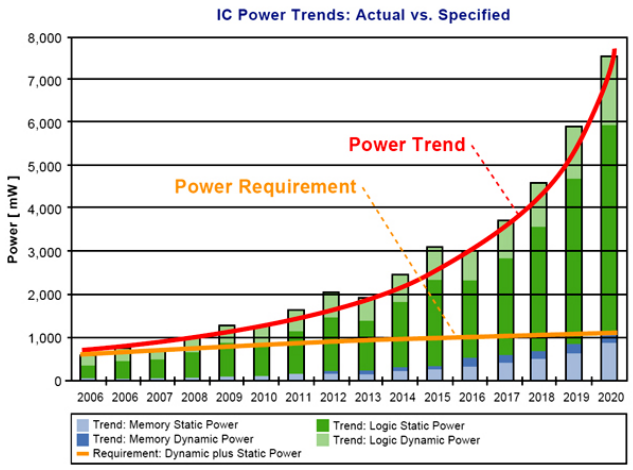

Figure 1 – Power trends are outstripping power requirements. Better design techniques are needed to close the gap. (Source: Silicon Integration Initiative (Si2), derived from ITRS 2005 Power Consumption Trends for SoC-PE).

What are designers worried about?

Everyone talks about "low power," but what do they really mean by that term? Why do we care about power, and what do we want to do about it? The answers vary greatly from one application to another. In some cases, the concern is maximizing battery life; in others, it's fitting into a cost-effective package, or reducing heat, or reducing the energy consumption of server farms. Dynamic power is the greater concern in some SoC designs, while leakage power dominates in others.

"You'll get me on a soap box here," said Stephen Padnos, methodology architect at wireless chip design firm Atheros. "Right now, power is sort of a buzz word. Everyone says 'low power' but they don't explain what it really means. You have to sit down and really understand what they mean by 'low power' instead of opening up some textbook and randomly throwing things at the problem."

With leakage power increasing faster than dynamic power (Figure 1), you might think that everyone would be most concerned about leakage. Not so. "In general, I think optimizing leakage is conceptually straightforward," Padnos said. "I think dynamic is a harder problem." One reason is the difficulty of identifying good vectors for dynamic power analysis.

When it comes to mobile multimedia devices, said STMicroelectronics' Magarshack, "the name of the game is optimizing the dynamic power of applications that may run either separately or concur-

rently." Optimizing dynamic power has "definitely" become more important than optimizing leakage power, "and unfortunately for the designer, the optimization of dynamic power is actually more difficult than static," he said.

At AMD, where Hui works with GPUs and CPUs, designers are concerned about maximizing battery life for notebook computers and containing thermal design power for desktop devices. One of the most vexing challenges is predicting and controlling the impact of process variations on leakage. For fast silicon, he noted, the leakage variation from typical to worst-case could be a factor of 3, 4 or 5.

Your specific power concerns are very much dependent on your application, noted Ameesh Desai, senior director for design tools and methodology at LSI Corp. "If a chip is going into some kind of mobile application, then leakage becomes very important. However, if it is a wired application, and it is in a high-performance storage or networking area, leakage may not be that significant. Dynamic power might be dominant," he said.

The net result: "There is no silver bullet. There is no one solution that will solve all problems. You have to look at a whole suite of different capabilities. There is a ton of techniques you can apply."

PLANNING AND IMPLEMENTING POWER MANAGEMENT TECHNIQUES

There are, indeed, a number of power management techniques that can be applied to reduce dynamic or leakage power (see Table 1). Most are implemented at the register-transfer (RT) or gate level, but some design teams start planning these techniques at the system level, and others would like to move up in abstraction.

Clock gating, for instance, is a widely-used technique that reduces dynamic power by shutting off clocks to blocks that are not in use. Multi-Vdd techniques reduce dynamic power by setting up "voltage islands" or power domains that run at different supply voltages. Dynamic frequency and voltage scaling (DFVS) and adaptive voltage scaling (AVS) are techniques that reduce voltage, or frequency, on the fly.

POWER REDUCTION TECHNIQUES

The following table gives a synopsis of the most common techniques and their uses. *Table 1*

Technique	Description
Active body bias (ABB) or substrate biasing	Applies active back bias voltage to the wells of N-MOS and P-MOS transistors to set the threshold voltages and control leakage. Compensates for process-, age- and temperature-related variations.
Adaptive voltage scaling (AVS)	Operates different blocks at variable supply voltages. Uses in-block monitors to determine frequency requirements, and adjusts voltage on-the-fly to satisfy them.
Clock gating and clock tree gating	Disables blocks or clock tree parts not in use.
Dynamic voltage scaling (DVS)	Operates different blocks at variable supply voltages. Uses look-up tables to adjust voltage on-the-fly to satisfy varying performance requirements. Signals that cross voltage domain boundaries are level-shifted.
Dynamic voltage and frequency scaling (DVFS)	Operates different blocks at variable supply voltages and frequencies. Uses look-up tables to adjust voltage and frequency on-the-fly to satisfy varying performance requirements. Signals that cross voltage domain boundaries are level-shifted.
Logic restructuring	A gate-level technique that optimizes dynamic power. Brings high-activity circuitry forward in the cone, and pushes low-activity circuitry to the back.
Multiple supply voltages (MSV)	Operates different blocks at different, fixed supply voltages. Also known as voltage islands. Signals that cross voltage domain boundaries are level-shifted.
Multi-threshold (V_{TH}) libraries	Contain gates using transistors with different switching voltage thresholds. Lower thresholds are faster and have higher leakage; higher thresholds have lower leakage, but are slower.
Power gating or Power Shut-Off (PSO)	Turns off supply voltage to blocks not in use. Significantly reduces – but does not eliminate – leakage. Block outputs float.
Power gating with retention	Stores system state prior to power-down. Avoids complete reset at power-up, which reduces power-up/reset delay and power consumption.
Transistor sizing	Upsizing increases performance and reduces dynamic power; downsizing reduces leakage power.
State retention power gating (SRPG)	Stores the system state in local registers. When on standby or idling, gates the clock, and the register saves the data. State retention registers use both a continuous power supply and a switchable supply. Other logic is powered only by the switchable supply, and can be powered down.
Save and restore power gating (S&RPG)	As SRPG, but uses a memory array.

References

How power design techniques impact SoC verification by Bhanu Kapoor (Mimasic), Arturo Salz (Synopsys), and Shankar Hemmady (Synopsys)

A Practical Guide to Low Power Design by Cadence Design Systems

When power consumption is a key consideration, describing low-power design intent with Accellera's Unified Power Format (UPF), improves the way complex integrated circuits can be designed, verified and implemented. The open UPF standard permits all EDA tool providers to implement advanced tool features that enable the design of modern low-power ICs. Starting at the Register Transfer Level (RTL) and progressing into the detailed levels of implementation and verification, UPF facilitates an interoperable, multi-vendor tool flow and ensures consistency throughout the design process. UPF is currently being ratified as IEEE Std. P1801™.

Accellera provides design and verification standards for quick availability and use in the electronics industry. The organization and its members cooperatively deliver much-needed EDA standards that lower the cost of designing commercial IC and EDA products. As a result of Accellera's partnership with the IEEE, Accellera standards are provided to the IEEE Standards Association for formalization and ongoing change control.

For further details, or to have your question addressed, please visit: http://www.accellera.org

Power shut-off or power gating switches off power to portions of the chip that are not in use. Body biasing or "back biasing" is a leakage reduction technique that uses substrate bias to raise voltage thresholds. Multi-Vt designs use cells with different voltage thresholds to trade off leakage versus power – the lower the voltage threshold, the higher the performance and leakage.

The impact of these techniques on dynamic power, active leakage power, and standby leakage power is shown in Figure 2.

Technique		Power Being Managed		
		Standby Leakage	Active Leakage	Dynamic
Power Gating	PG			
Retention with PG	RPG			
Multiple Supply Voltages	MSV			
Dynamic Voltage Scaling	DVS			
Adaptive Voltage Scaling	AVS			
Multi-Threshold CMOS	MTCMOS			
Adaptive Body-Biasing	ABB			

Primary Secondary Verification Impact

Figure 2 – Different power management techniques have differing impacts on leakage and dynamic power. (Source: Mimasic)

Who uses what? Cadence Design Systems surveyed over 3,000 engineers who registered to read its Practical Guide to Low-Power Design, and the results are shown in Figure 3. Clock gating is the most widely used technique, followed by multi-Vt libraries and power shutoff. Substrate biasing and DVFS are emerging techniques.

Many SoC designs today use multiple techniques. At the June 2008 Design Automation Conference, Magarshack described a 65 nm STMicroelectronics chip designed in 2007 that used clock gating, multiple voltage thresholds, multiple supply domains, DVFS, and state retention. He described a current chip project that uses nine power grids, has separate power supplies for memory arrays, and uses source biasing and body biasing.

More recently, Magarshack commented that body biasing is becoming a "mainstream" technique for leakage reduction at 45 nm and below. Magarshack noted, however, that "pretty much all" power management techniques are decided at the register-transfer level today. "We definitely want to bring that level of decision-making up to the architectural level," he said.

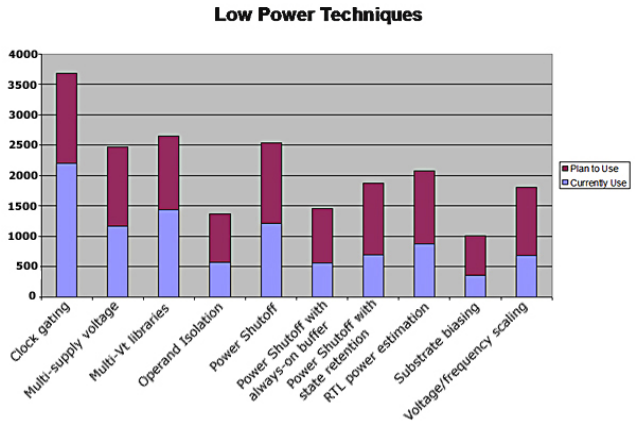

Figure 3 – Clock gating and multi-Vt libraries are widely used power management techniques. (Source: Cadence Design Systems).*

** Note that "plan to use" respondents include a number of "currently use" respondents. Thus the graph cannot be used to derive adoption growth metrics.*

LSI sees a lot of promise in adaptive voltage scaling, said Desai, because it allows designers to control voltage in order to reduce power. But it does require the installation of sensors to detect power levels, and as such requires "up front" planning at the architectural level, he said. For multi-voltage design, Desai said, figuring out which blocks can be shut down in what modes should be the "architect's job."

Most power management techniques have tradeoffs. They may require additional area, result in performance degradation, or raise signal integrity concerns. As described later in this report, they also pose a number of functional verification challenges.

"I tend to think that too many power domains is not a good thing. You get to the point of diminishing returns," Atheros' Padnos said. "There is definitely some overhead associated with these. Generally, if you have more than you can count on one hand, you've gone overboard." Padnos quickly added, "but we have another group

which used quite a few more than that for a GPS design, and they think they got really good results that way."

Qualcomm plans what will go into voltage islands "very early in the design," Tessitore said. "We have to do that because we develop our own PMICs [power management ICs]," he said. He noted, however, that "I don't believe we have any automation that determines where to cut the voltage islands."

It's best to start power management early, said Lawrence Vivolo, director of low-power solutions marketing at Synopsys. "While clock gating at the RT level can be very effective, switching off a complete block at the system level or regulating the voltage of a processor depending on its load can have a much more profound impact," he said. "It is important to apply low power techniques as early as possible."

System-level power management

One company that takes a high-level approach to power management is Texas Instruments. Richard Scales is manager of the system modeling and analysis team in TI's SoC architecture group. His team recommends to the power architecture team the power management features that are most likely to achieve the power objectives for OMAP devices, where they are required, and what level of granularity will be used. The actual design work is done by different technology groups.

Scales' team uses clock gating and power gating. The team also deploys voltage domains, and tests the SoC at different operating points (OPPs). The 100% OPP is the fastest point at which a module can execute. A 50% OPP, for example, operates at 50% of the clock frequency, but at a lower supply voltage. One area of the chip may need 100% OPP, but another area may need only 25%. That drives the rationale to have different voltages, he said.

"We implement multiple OPPs using dynamic voltage frequency scaling," Scales said. "We have an adaptive voltage scaling technology called SmartReflex. SmartReflex scales the voltage around a predefined OPP voltage value. It maintains the circuit's timing by adjusting the supply voltage on-the-fly as the silicon temperature rises or falls." Scales explained that SmartReflex uses a circuit in the silicon that constantly monitors and measures the current across a known resistive load. SmartReflex (see Figure 4) effectively compensates for process variation, resulting in a higher yield, he said.

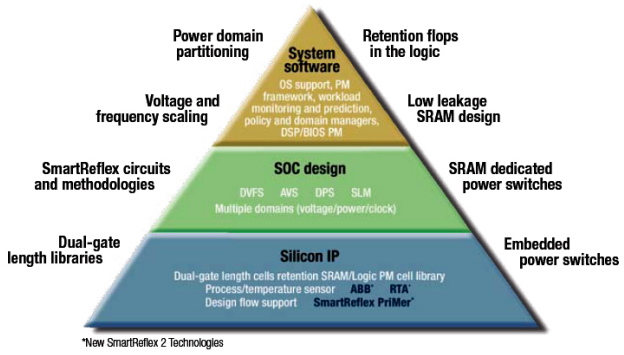

Figure 4 – TI's SmartReflex uses a combination of power management techniques. (Source: Texas Instruments SmartReflex whitepaper).

Scales' team also uses dynamic power switching (DPS), which is a form of dynamic power gating. "Take for example a use case that requires 40% loading on a particular module," Scales said. "There are two options. It can run at OPP 50%, which is enough for 40% loading and would allow reduced supply voltage. The other option is to run at 100% OPP, but for only 40% of the time." Scales said TI uses a hardware-controlled DPS technique that offers fine-grained control and has a context save-and-restore capability.

LOW-POWER ARCHITECTURAL ISSUES: PROCESSORS, MEMORIES, SOFTWARE

The use of power management techniques such as multiple power domains is just one issue that needs consideration at the architectural level. Other key issues include the number of processors involved, hardware/software partitioning, what memory architecture to use, and which process technology to select. All have profound impacts in power consumption.

"Once you have a design at RTL, 80 to 90 percent of the power budget is set in stone," noted Craig Cochran, vice-president of marketing and business development at ChipVision Design Systems. "The greatest opportunity for affecting the power is selecting an architecture that's power efficient."

Multicore architectures are one response to power challenges. The basic idea is that by using multiple processor cores, you can significantly boost performance with only a minimal increase in power consumption. But multicore SoCs are not appropriate for all applications, and they come with significant design and programming challenges that are outside the scope of this report.

Jan Rabaey, professor of electrical engineering and computer science at the University of California at Berkeley, observed that multicore devices achieve more performance for the same amount of power, but do not reduce the amount of energy required per operation. The next step, Rabaey believes, will be the use of dedicated processors for specific functions. The future SoC, he said, "is going to look like a multicore, with a variety of processors and accelerator blocks, including mixed-signal," he said.

The importance of software

Meanwhile, hardware/software partitioning allows some tradeoffs with respect to power. According to LSI's Desai, putting functionality in software may save power, because you don't have to add more hardware to the chip. "If you do something in software, you may potentially suffer in terms of performance, but you may reduce the power," he said. How is this kind of decision made? "Right now it's back of the envelope and spreadsheets," he said.

Qualcomm's Tessitore, however, said that putting functionality into dedicated hardware actually saves power. The tradeoff, then, is between the flexibility allowed by software and the potential power savings of dedicated hardware. Qualcomm developed its own architectural tool, called Visa, to help with hardware/software partitioning.

Embedded software is "a big piece" of the power consumption picture, Tessitore said. "You can have all these facilities on the chip, but if the software is not shutting down the clock, and it is not leveraging voltage or frequency scaling, you're just doing dumb

things," he said. "As the guys develop software, we tell them that they've got to hit power targets."

Power-aware software development "is another area where there aren't good tools," Tessitore said. He said that Qualcomm tried to find a fast software simulator that would factor in power, and eventually had to develop its own. "This is definitely an area that the industry needs to put more focus on," he said.

"If software needs are not accounted for in the hardware, then the optimal solution may not be achievable," said Stephen Olsen, product marketing manager for Mentor Graphics' embedded systems division. A case in point – if module or subsystem clocks are turned off entirely, the system cannot adjust to a changing hardware environment, because signals will not propagate up to the software to be processed. A better approach, he said, might be to reduce the frequency of the clock drastically when the device is no longer needed.

Power-aware software development is a tough problem, Rabaey said. "People have looked at compilers for a long time, and it turns out there are certain things you can do that help you optimize energy, but very little." However, operating systems provide a lot of opportunity for power savings, Rabaey said, provided they have an awareness of the voltage and frequency environment in which they're executing.

Memories and process technology

With all the focus on power management for digital logic, it's easy to forget the importance of memories – and unwise. In an SoC, noted STMicroelectronics' Magarshack, one-third of the power consumption may be in SRAM, one-third in the clock tree, and one-third in the remaining random logic. Memory architecture is thus a key part of STMicroelectronics' power management strategy.

In ST's designs, Magarshack said, processor cores come with on-chip L1 and L2 caches. "The SRAM architecture is tailored to the memory cache of the processor," he said. "We try to introduce some pipeline stages or try to introduce some banking. Some of the banks are in retention mode while the bank that's being accessed is powered up to produce the data that is required for the processor." He said that ST also uses power switches and power domains to minimize memory power consumption.

"We look at whether we want to implement one big memory, or a couple of smaller memories," said Atheros' Padnos. "You might be able to split a memory, where you have a small memory that has to be on all the time, together with a bigger memory that you can turn on and off depending on the mode of operation. The small memory is in a 'keep alive' state and you turn the big memory on when you do things that are more computationally intensive."

Qualcomm is experiencing a significant power savings by using low-power DDR (LPDDR1) memories in its mobile devices, Tessitore said. But there are limitations – LPDDR1 is currently limited to 200 MHz, and LPDDR2, which promises lower voltages and higher speeds, is not yet shipping.

Qualcomm also uses stacked memory in some of its chips. For example, Qualcomm's recent Snapdragon chip stacks 32 Mbytes on top of the die, and future chips will stack 64 Mbytes. Stacking memory significantly lowers interconnect capacitance, and can cut memory power consumption by as much as 30 percent, Tessitore said. "Performance-critical things that need a lot of memory bandwidth and activity like graphics, multimedia, and modems will be placed in the stacked memory," Tessitore said. "We'll place the OS and other applications in the external memory."

Another issue that needs consideration at the system level is process technology selection. Qualcomm is currently using low power (LP) process technologies from

THE SI2 LOWER-POWER DESIGN REFERENCE FLOW

A standardized, low-power reference flow could give design teams a leg up on what is today a "make it up as you go along" type of process. That's what the Silicon Integration Initiative (Si2) Low Power Coalition (LPC) flow working group, led by IBM's Nagu Dhanwada, is trying to accomplish. AMD's David Hui is vice chair of the working group.

The working group has defined and set forth requirements for three development phases – ESL, design, and implementation. "We did focus on the ESL side of things," Dhanwada said. "There are still discussions going on about how to come up with power models, and how you can plug power models into a TLM-like simulation to get some power statistics." He said LPC is launching a modeling subgroup to look further into some of these questions.

In a presentation given at the 2008 Design Automation Conference, Hui provided some further information about the overall design flow envisioned by the working group, as well as some current thinking about ESL power modeling and analysis. The presentation defines ESL as the phase of the design process that transforms a high-level specification into an architecture, providing a platform for system and block-level power optimization. It includes IP selection, creation of power budgets for subsequent design phases, power analysis, and firmware/hardware partitioning.

One question that the flow working group is grappling with is the role of power intent formats in the ESL flow. "As we move up, there is obviously a role for extending UPF and CPF. But conceptually, does the ESL space need to consume CPF in order to have a starting design intent, or does it just generate CPF?" Dhanwada asked.

The design phase, as currently defined by the working group, includes the detailed selection, modification, coding, and interconnection of RTL descriptions of major subsystems. Each subsystem is mapped onto one or more power domains, and rules for the treatment of inter-domain signals are defined. RTL power analysis and linting are performed, with static and dynamic power targets for each power mode.

The implementation phase involves the mapping of the RTL and power description file into physical devices that can realize the design intent. Implementation tools must ensure that all constraints are satisfied, and they may generate additional power constraints that must be satisfied by subsequent tools. Power optimization strategies include such techniques as clock gating, multi-Vt, long channel devices, device stacking, body bias, power-aware placement, and activity-aware optimization.

TSMC, Chartered and Samsung. They provide very low leakage for devices such as cell phones – at a significant performance cost. "We estimate that we trade off about 30 percent in performance by using LP technology, so you can imagine the challenge," Tessitore said. Still, he noted that the Scorpion microprocessor in the Snapdragon chip set, designed in 45 nm technology, runs at 1 GHz and uses only 350 milliwatts, and the leakage is less than 10 milliwatts.

ESL TOOLS AND FLOWS: WHAT'S MISSING, WHAT'S EMERGING

It seems clear that critical decisions concerning power management techniques, processors, memories, hardware/software partitioning, and technology selection should be made at the system level. What tools exist to help design teams with such decisions?

"There really aren't any good tools at the ESL level for doing tradeoffs," said Tessitore. That's why Qualcomm developed its own hardware/software partitioning tool. Tessitore said that today's virtual platform tools are "really good environments" for software development, and he is hoping that they can be extended to examine architectural tradeoffs for bandwidth and power.

Padnos said that Atheros is not able to do much planning or modeling at the architectural level today, but would like a way of making tradeoffs before writing RTL. What's important at the system level, he noted, is not so much absolute power values, but enough information to make tradeoffs between different architectural choices. His wish list includes a tool that would evaluate the resources needed to implement a given algorithm.

What is AMD doing today at the architectural level for power management? "Not that much," replied Hui. "In fact, there has always been trouble in that area. We have a group that works on it, but the accuracy is very questionable." While accurate ESL power prediction may not be possible, Hui said, designers can still obtain some useful information. "If you do some kind of profiling, you can always translate activity into power. If you can reduce activity you can reduce power," he noted.

Nagu Dhanwada, senior R&D engineer at IBM's EDA group, chairs a working group of the Silicon Integration Initiative (Si2) Low Power Coalition (LPC) that is aiming to define a low-power design reference flow. The toughest part should not be a surprise. "Most of the unknowns are in the ESL part of it," he said. "We have a few open questions, and we are focusing on what is the right modeling abstraction for ESL."

➔ **Main article continues on page 9**

CLOSED-LOOP VERIFICATION METHODOLOGY FOR LOW-POWER SOC DESIGN

by Neyaz Khan
Low Power Architect, Cadence Design Systems

Complete and accurate verification of low-power logic is essential to an effective low-power methodology. For maximum impact, power-related decisions and tradeoffs in a system-on-chip (SoC) design must be made at the architectural level, and the effects continuously analyzed and verified at every stage from RTL down to GDSII. This article discusses such a methodology, which has been successfully deployed across a wide customer base and varied applications.

POWER INTENT

The first step is to capture power intent in a single place and format, which can then be analyzed and processed throughout the flow. This is done using the Common Power Format (CPF) [1]. Similar to the capture of design intent in HDL, a CPF file specifies complete power intent in a single place so power intent remains consistent throughout simulation, synthesis and implementation [2], including:

- Operating voltages, power domains, voltage islands, power modes
- Power saving schemes, such as power shut-off, retention, isolation, level-shifting, dynamic voltage and frequency scaling (DVFS), multiple supply voltages (MSV)
- Library- and implementation-specific information

VERIFICATION OF THE POWER INTENT SPECIFICATION

It's not enough to merely specify the power intent in a machine-executable format. The specification itself must also be analyzed and verified for functional, electrical, and structural correctness and completeness. This essential task is performed by Cadence Conformal Low Power (CLP) tools throughout the low power flow. Such CPF quality checking ensures that the specified power intent is correct and complete.

FUNCTIONAL VERIFICATION OF POWER INTENT

Functional verification is in itself a daunting task. To add to the complexity, low-power effects must also be verified as part of functional verification. For those design units that need to be powered down during normal operation of the device to save leakage power, the effects of Power shutoff (PSO), and associated Isolation (ISO) and Retention (RTN) behaviour need to be verified thoroughly to ensure that they do not affect normal operation of the device. Further complexity is introduced by various design units operating at different voltages and with the effects of simultaneously scaling supply voltage and clock frequencies to optimize for dynamic power.

This is often a rich source of error and must be thoroughly verified for all operating modes as specified in the CPF for the design. It is not enough to run LP checks on a few selected sets of stimuli. Such checks must be performed together with and as part of the complete functional verification suite to efficiently cover every aspect of the device functionality.

SIMULATION OF LOW POWER EFFECTS

As stated, hierarchical CPF is used to capture power intent. Virtual PSO, ISO and RTN behavior are simulated and the effects observed. To make this process efficient and fast, no RTL changes are required to superimpose low-power behavior on the design. There is no need to even connect the power control signals in the RTL. This is done during synthesis, saving time for early architectural exploration to gauge the effects of low-power techniques on system performance and functionality. Architectural issues and bottlenecks can be quickly

identified along with any hardware-software and other system partitioning tradeoffs without RTL modification, simply by modifying the CPF.

ASSERTION-BASED CHECKS

The CPF specifies and defines low-power intent with respect to the supplied power control signals, while an on-chip power control unit generates and drives those signals. These control signals must be exercised in a predetermined order to ensure a correct power sequence. The CPF also specifies a legal set of power modes and their transitions, along with cycle time requirements for power domains to ramp and transition to specified voltage levels.

A number of properties can be derived from the CPF specifications which are then used to automatically create assertions. These assertions monitor and check transitions in the power control signals to verify legal and correct low-power behavior [3]. These automatically-created assertions are power-aware and remain active, but dormant for powered-down modules. They also provide low-power coverage metrics.

VERIFICATION OF NON-LINEAR EFFECTS AS PART OF DIGITAL SIMULATION

Figure 1 shows the effects of voltage ramp on low-power simulation, highlighting why they cannot be ignored in digital simulation. Non-linear effects, including rise time, fall time and transients must also be verified.

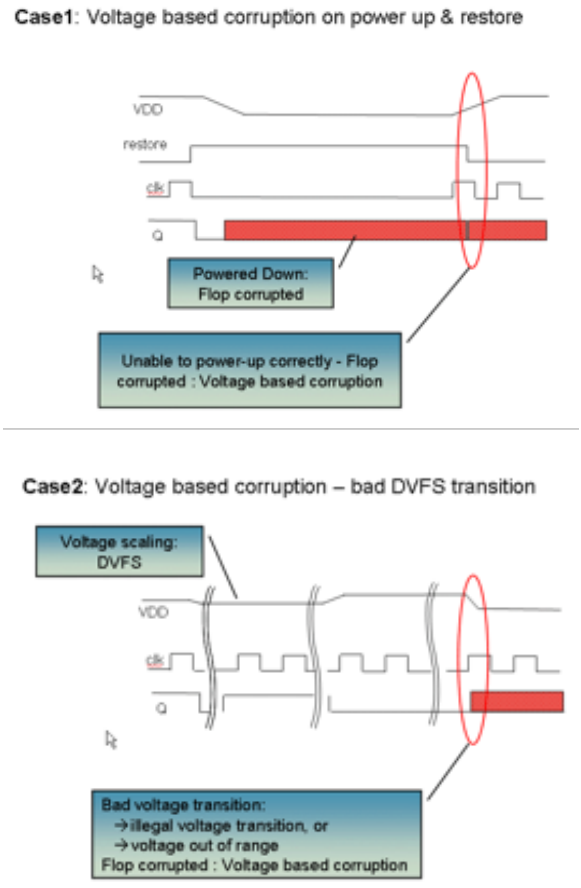

Figure 1: Effect of voltage ramp time on LP simulation

STRUCTURAL, ELECTRICAL AND EQUIVALENCE CHECKS FOR CLOSED LOOP VERIFICATION

Functional, structural and electrical checks are performed throughout the low-power flow, thus enabling a closed loop flow:

- CPF quality checks are run to ensure golden CPF.
- Low power simulation validates the golden RTL and golden CPF, which ensures that both design and power intent are correct.
- Synthesis and implementation build the design based on design and power intent. If Cadence tools are used, the golden CPF enables automation and low-power implementation which is correct-by-construction.
- Logical and physical netlists can be verified against both the golden RTL and golden CPF independently of the synthesis and implementation tools used. Checks include power-aware equivalence checks, formal and structural checks, and power intent electrical checks.

POWER CLOSURE

Unless individual tests are written for every possible mode and mode transition, randomization is used to create tests, employing a methodology such as Open Verification Methodology (OVM). As in any randomized verification environment, metrics for low-power coverage are collected during low-power operation to gauge the quality and completeness of the low-power verification effort. An executable verification plan and cover groups for collection of low-power coverage are automatically created from CPF. Power-aware assertions are also automatically created and provide additional coverage metrics.

Power closure is achieved after all metrics for both low-power functional and structural, equivalence and electrical checks have been fully met. Bucket analysis is also performed on low-power coverage data for hole analysis together with checks for illegal power modes and mode transitions.

References

[1] CPF1.0 Specification: Si2: http://www.si2.org/?page=811
[2] A Practical Guide to Low Power Design: http://www.powerforward.org/lp_guide
[3] Neyaz Khan, William Winkeler. Power Assertions and Coverage for improving quality of Low Power Verification and Closure of Power Intent: DVCON, 19-21 Feb 2008. San Jose, CA

For further details, or to have your question addressed, please send email to: holden@cadence.com

➔ continued from page 7

Creating power models at ESL, or any level of abstraction, is a big part of the problem. Dhanwada said that creating power models today "is an ad-hoc process. There's nothing standard." With ESL, he noted, models are generally transaction-level or higher and are not synthesizable. One problem is mapping between high-level events and lower-level implementation features such as clock gating. "The transition is not automatic. If we could infer structure from high-level models, part of the mapping problem would be solved," he said.

Hui is vice-chair of the LPC flow working group. "We have been discussing how we can do power modeling for ESL, and it's a very complicated issue," he said. "A low-level model, such as a flip-flop, is very easy. But when you start talking about modeling a microprocessor, then the mode that it's operating in and how you program the microprocessor impacts your power a lot. We need to be able to represent that. You need to have some kind of personality to program into the model."

Today, Hui said, AMD uses the Liberty library format for modeling, but it's "crude." One big shortcoming is that it doesn't allow the designer to model leakage power over temperature. To create models, "we basically use a spreadsheet today, but in the future we would like to come up with a more complex modeling actually based on temperature."

TI's Scales has a suggestion. "I'd like to see OSCI [Open SystemC Initiative] tackle power modeling," he said. "The challenge is that the framework rules for power management are different for each vendor. Having a common, configurable power model would be very, very difficult. But if we had a generic use-case based modeling approach, I think that could be reused."

Vendors respond to ESL demands

EDA vendors acknowledge the need for better low-power ESL tool support – and they say they're on the move. "There is a clear need for solutions in that [ESL] space, given that early decisions have the biggest impact on power consumption," said Synopsys' Vivolo. "According to anecdotal data, the impact of low power techniques at the ES level can be orders of magnitude higher than at the RT [register transfer] level."

So what are EDA vendors doing to support power-aware design at the architectural level? Neil Hand, director of Cadence Design Systems' services marketing organization, pointed to the way that EDA technologies have traditionally evolved. First, he said, comes implementation, followed by analysis and finally automation. What we're seeing at the architectural level, he said, are technologies that can help with implementation and some of the analysis, with automation to follow later.

Cadence, for example, offers several products that can help with power-aware ESL design. First, Hand said, the InCyte planning tools recently acquired from Chip Estimate can provide a system-level power analysis. Secondly, Cadence's emulation tools are often used to evaluate hardware/software tradeoffs and their impact on power. Third, Cadence's new high level synthesis tool, C-to-Silicon Compiler, includes RTL Compiler, which can implement power management techniques using the Common Power Format (CPF).

Mike McNamara, vice president and general manager of Cadence's C-to-Silicon Compiler project, said that the high level synthesis tool has a "lot of potential" for helping designers implement such techniques as multiple voltages and power shutoff, but "is not as automated as we'd like." At present, users can specify only area and timing as cost functions, but Cadence plans to add an ability to specify power as a cost function, he said.

RETHINKING LOW-POWER DESIGN

To design the next generation of low-power SoCs, we need to change our thinking, according to Jan Rabaey, professor of electrical engineering and computer science at the University of California at Berkeley. In an upcoming book entitled "Low Power Design Essentials," Rabaey describes two conceptual shifts that he thinks will be necessary to cope with future IC power challenges.

The first concept is called "always optimal" design. The basic idea is to monitor every device on the die, looking at criteria such as performance, leakage, voltage, and switching activity. Based on that information, a scheduler or operating system can schedule operations, re-adjust thresholds, and change various operating parameters (Figure 1). "You're doing a dynamic tracking of what's going on and automatically adjusting it in the most energy-efficient type of fashion," Rabaey said.

Rabaey believes that always-optimal design can be handled at the architectural level, and will be totally transparent to the application developer. It will require efficient optimization tools. It will also involve some area tradeoffs, but "area has become less of a factor than it used to be," Rabaey said.

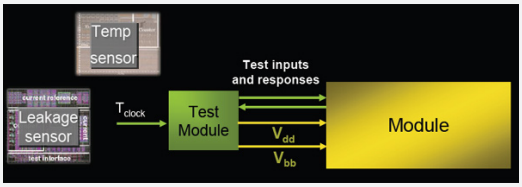

Figure 1 – Always-optimal systems are adaptively biased to adjust to operating, manufacturing and environmental conditions. (Source: Jan Rabaey, OpenAccess Conference keynote, April 2008).

The second new approach is called "better than worst case" design, or to use another term, "aggressive deployment." Today, many chips are designed to meet worst-case operating conditions. "Setting your design parameters in such a way that your design is always going to be correct and always going to be meeting specs is a really bad idea," Rabaey said. "If you look at statistics, it turns out you're hitting the worst case only rarely."

If we can do enough statistical analysis, Rabaey believes, we can design for the "medium case" rather than the worst case, and then use error-correction techniques when something goes wrong. You may occasionally miss a timing edge, for example, but you build in a mechanism to detect and correct it. Rabaey said this approach to design can result in power savings from 30 to 50 percent.

But the key message, Rabaey said, is that it's time to move away from low-power design as a set of ad-hoc techniques. "The methodology is now evolved enough that you can take a structured approach to low power," he said.

Vivolo said that Synopsys offers the ability to "instrument" SystemC transaction-level model (TLM 2.0) based virtual platforms with power-related information. How? Early estimates from RTL simulation can be fed back into system-level virtual platform models in Synopsys' Innovator environment. Also, Innovator offers a graphical user interface that allows users to enter information such as average power consumption for different power modes. The virtual platform created with Innovator can then be used to develop power management software, and can provide feedback to architects who are making decisions such as the partitioning of voltage islands.

In April 2008, ChipVision rolled out PowerOpt, a SystemC synthesis tool that generates RTL for low-power ICs. It takes in SystemC code and characterized power models, optimizes blocks for power, area and latency, and generates synthesizable RTL code. According to Cochran, customers have seen 75 to 80 percent reductions in chip power compared to high-level synthesis tools that are not power-aware. "PowerOpt does a system-level analysis using real data traffic, allowing very high accuracy estimates of what your downstream power analysis will be," Cochran said. "You can make some pretty smart tradeoffs in architecture with the power analysis."

EDA vendors also recognize the need for ESL power models. "From what we have seen, most companies are using estimations or manual modeling for creating power models," said Glenn Perry, general manager of Mentor's ESL and HDL Design division. "This results in an error-prone, inefficient process." Perry said Mentor offers a methodology that enables power models to be annotated into standard SystemC TLM 2.0 models.

ChipVision is actively working with the Si2 LPC towards a power modeling standard, said Thomas Blaesi, ChipVision president and CEO. "We believe it is definitely necessary to have a power modeling standard, because SoC blocks come from different sources and you want to be able to use the same model," he said.

Meanwhile, ChipVision offers P-SAM, a combination of products and services that can build SystemC TLM power models. The models are available at a variety of abstraction levels, and can be used in any SystemC system-level simulation, Cochran said.

Successful ESL flows

Challenges notwithstanding, STMicroelectronics is actively working on an ESL flow for low-power design. Magarshack said the company is using high-level, power-aware synthesis from Mentor Graphics and Synfora, and is working with Atrenta to develop a high-level power estimation tool (see Figure 5). He said the approach allows ST to quickly compare different architectural solutions.

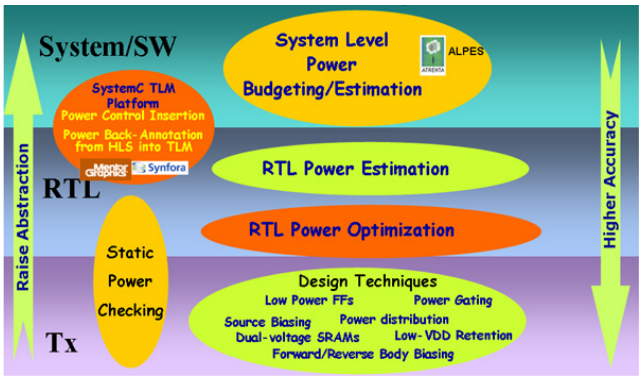

Figure 5 – STMicroelectronics uses a variety of power management techniques, starting at a high level of abstraction (Source: STMicroelectronics).

At TI, Scales' group starts its work with a marketing requirement document that defines use cases. "An OMAP device has about 10 use cases that are considered to be very important," he said. "Those use cases help us define the device's functionality."

"We have internally-developed SystemC-based virtual platforms that are built for various 'what if' analyses," Scales said. "We use those primarily for chip-level interconnect performance analysis. For example, do we have enough bandwidth for these particular traffic scenarios, or does bus arbitration ever block a requestor for too long?"

"At the same time we do some really high level analysis of the power structure. From the use case performance analysis, we understand the data flow and control flow at a high level. We then analyze the use cases and their flows to identify the interdependencies that help us to determine the design's partition into power and voltage domains."

Scales said his team also runs a rapid, power "what if" analysis, using a power management (PM) model together with an XML configuration model that reconfigures the PM automatically. "The PM models the chip's controller blocks. It also models the power structure of the entire chip – the hierarchy of voltage, power and clock domains; the clock tree; and the power source tree. Using the PM and XML models, we can rapidly analyze multiple use cases on different power structure architectures. For example, I may have four modules that are sometimes executing at the same time. Do I save more power by using one power domain for all four modules, or by allocating the modules to two domains?"

➔ **Main article continues on page 12**

STRAIGHT TALK ABOUT LOW POWER DESIGN

by Stephen Bailey
Mentor Graphics director of product marketing, Design Verification Technology

Fortunately, designing for low power is one of those challenges that can be tackled with the divide and conquer approach engineers thrive on. So let's do a little divide and conquer discussion here to examine some useful low power know-how. Specifically, let's look at architectural or system design, functional verification and design implementation.

ARCHITECTURAL OR SYSTEM DESIGN

There are two key elements that are required for an architectural level flow:

- Accurate power models and a proven platform that support those models.
- Tools and methodologies for optimizing power at the architectural and system design level.

Most companies use estimations or manual modeling for creating power models. This is an error-prone, inefficient process. Mentor offers a methodology, solution, and automation for the fast, accurate creation of power models that can plug into any transaction level modeling (TLM) platform.

FUNCTIONAL VERIFICATION

Early verification significantly improves productivity and finds power-related design errors early when it is easier and cheaper to resolve them. Functional verification of low power designs requires a power-aware simulator that:

- Understands Unified Power Format (UPF) — the power design intent.
- Can model the logical implications of the power state of the design — corruption when powered down and conditional corruption in bias modes.
- Can insert power-related functionality as required by the power intent — switches, isolation and retention behavior.
- Facilitates debugging through the automatic generation of assertions.
- Facilitates coverage-driven verification through automatic generation of coverage metrics and data.
- Supports hardware/software co-verification of the power management functionality.

DESIGN IMPLEMENTATION

A common technique to reduce dynamic power is the use of multiple voltage islands (domains), which allows some blocks to use lower supply voltages than others, or to be completely shut off for certain modes of operation. These techniques require new thinking in physical design in order to:

- Correctly place and route across multiple domains and honor the multi-voltage domain specifications.
- Optimize routing by analyzing timing, power, and signal integrity across all modes and corners concurrently.

The core architectures of incumbent place and route solutions in the market are at least 10 to 15 years old and were intended to handle at most one or two scenarios. A multi-voltage multi-mode multi-corner (MCMM) implementation system for low power designs must be built from the ground up to address new MCMM requirements.

Clock trees have a big impact on dynamic power. Low power SoCs can have dozens of clock trees operating at different rates, times, and modes. Clock tree synthesis that is concurrently optimized across all modes and corners reduces the number of clock buffers required and can significantly reduce the total power consumption.

Low power design can be a formidable task, but with Mentor's design, verification and implementation tools and methodologies, it's one that can be tackled with routine success

For more information on Mentor's technologies for low power design go to: http://www.mentor.com/lowpower

For further details, or to have your question addressed, please visit: http://www.mentor.com/products/product_maincontact.cfm

➔ **continued from page 11**

The power estimates themselves are calculated by a TI-internal tool. In addition to the XML configuration file, Scales' team has done a lot of work on internal tooling and infrastructure to increase automation. As a result, the team can handle complex interdependencies that are difficult – or impossible – to capture and validate in an Excel spreadsheet.

"A few years ago, only one person could manage use case entry into a spreadsheet because it involves so much knowledge of the SoC dependencies," Scales said. "Today, using the PM anyone can enter a use case – and they don't have to know anything about the actual hardware and its interdependencies."

Infineon is bridging the gap between high-level design and lower-level accuracy for its automotive ICs by building what Karmann calls "layout prototypes" that help predict power consumption. These are created with a fast logic synthesis run, followed by a "run time optimized" placement and routing. "Quality and reliability is a must in automotive design, so all power-related decisions done at the spreadsheet level must be confirmed by power analysis during layout," he said.

RTL ON DOWN: STILL SOME GAPS

Support for low-power design is far more established at the register-transfer level, gate level, and physical design level than it is for ESL design. This is where most power management techniques are implemented, often with the support of analysis and optimization tools. Providers such as Cadence Design Systems and Synopsys offer low-power design flows that include RTL synthesis, verification, and placement and routing, and Mentor Graphics targets low-power design with its Olympus-SoC netlist-to-GDSII suite.

Vivolo noted that multi-Vt and dynamic and leakage optimization support has long been available in Synopsys' synthesis, placement and routing products. Synopsys power optimization and analysis capabilities include automated power management cell inference, activity-based power-aware placement, automated power gate insertion, automated multi-supply rail connection, multi-mode/multi-corner optimization, power network synthesis, and power network analysis.

Cadence introduced a comprehensive low-power solution last year. According to Hand, Cadence's Incisive verification tools and Encounter IC implementation tools support "all" power management techniques from RTL to signoff, using CPF as a common format for conveying power intent.

"It seems to me we went from no support for multiple supply voltages in the implementation tools two years ago to a reasonable level of automation today for low power," said NXP's Menager. "As a result, techniques have become mainstream, and low power implementation has been enabled."

But designers still note gaps. "Automation for power optimization is addressed in place and route, but it has only just started at RTL," said Infineon's Karmann. "In my opinion, activity optimization has the highest leverage to save power, but the lowest degree of automation." What's needed, Karmann said, is a "consistent RTL-to-layout activity optimization flow instead of point solutions."

Atheros primarily uses a Cadence flow. "The pieces are all there," Padnos said, "but they are not always as mature as I would like." Padnos said his IC implementation tools do a good job of providing power reports, but he added that "they need to get better in dynamic and rush current analysis. In doing multiple power domains the tools are pretty good at inserting switches, and not all that good at helping you make sure you've got the right number of switches and that you've sequenced them on and off correctly."

The importance of clocking

IC design tools have good support for multiple voltage thresholds, Padnos said, and synthesis tools provide automatic clock gating. "Something that's relatively new, and has received a lot of attention from the big vendors in the last couple of years, is clock tree synthesis. I think the tools have improved a lot in that area," he said. To supplement the capabilities offered by RTL synthesis, Atheros uses a standalone clock tree synthesis tool from Azuro.

Why is clocking so important? Depending on your design, the clock tree can consume as much as 30 percent of the chip's total dynamic power, Padnos noted. And there are other impacts too – for example, a poorly designed clock tree can add so many wires that chip congestion becomes unworkable. "Everything kind of gets bad if you don't do a good job on clock tree synthesis," he said.

Clock gating has been around for at least 10 years, but there's still room for improvement, AMD's Hui said. "One of the big problems with clock gating is that a lot of the time, designers are working with very small pieces of code," he said. Seeing a bigger picture, he noted, could provide additional clock gating opportunities. New tools are adding more visibility to clock gating, he noted.

Qualcomm uses extensive clock gating, possibly requiring gating for 90 percent of the clocks, Tessitore said. Why? "It is typical for just the clocking on the chip to consume 30 to 40 percent of the power." Tessitore also noted that some clock tree synthesis tools do a great job of minimizing skew, but "then they really load up the clock tree with buffers and inverters and get carried away. They are supposed to go back and optimize but they really don't."

STMicroelectronics discovered that relaxing maximum clock skew can achieve dynamic power savings, Magarshack said. At the recent

Design Automation Conference, he discussed how ST was able to reduce clock tree power by up to 40 percent with the help of Mentor's Olympus-SoC suite. Most existing clock tree synthesis tools, he noted, typically aim for zero skew, or as close to zero as possible. The result? Clocks "end up with very sharp slopes that overlap with each other. You have a high peak current and IR drop, and you have sharp rising or falling edges."

Synopsys' Eclypse low-power solution, introduced Feburary 2008, adds a new capability to the DC Ultra synthesis tool that lets users make tradeoffs to get the best possible combination of low skew and low power. The tool can use automatic clock tree rebalancing to add or remove levels of clock gating. IC Compiler, meanwhile, added algorithms to optimize clock tree placement.

Impact of variability

Another implementation-level issue that concerns Magarshack is modeling the impact of process variability on power. "You have to mitigate all your constraints," he noted. "You have to design for chips that will leak the most, and those that will be the fastest chips. You still have to meet your leakage budget and you have to have the same design still work when you are in a slow process corner. You have all these contradicting constraints that have to be managed."

Since variability is increasing, he said, designers must ensure that the design will work within a wide process corner. And it isn't only a matter of process corners – voltage and temperature corners must be considered too. To model random variability, ST is beginning to use statistical timing analysis. For systematic variations, the company is using modeling.

What about analog/mixed-signal circuitry? "As a percentage of total power on the chip, analog has become more important," said Padnos. Unfortunately, he noted, "in some cases there is not a lot you can do – you are still going to burn a lot of power. You can look at the efficiency of your amplifiers or oscillators or things like that." Other than Spice, he noted, there are few tools that help analog designers with power analysis. Sometimes, he noted, it makes sense to put more functionality into the digital domain, where more power management techniques are available.

"Accurate analysis and verification of power-sensitive designs have increased the need for transistor-level analysis of logic blocks and on-chip memory, as well as integrated analog/mixed-signal circuitry," said Synopsys' Vivolo. He said Synopsys' HSIM suite provides simulation and analysis tools that help with power management.

Hand said that Cadence's InCyte chip estimation tool can capture analog power data, but he said there isn't yet any automation of analog power optimization. He noted that some customers get a "huge" power savings by moving functionality from analog to digital circuitry.

Formats add productivity – and controversy

One development that has greatly eased RTL-to-GDSII low-power design and verification challenges is the emergence of standard formats for expressing power intent. There are two. The Common Power Format (CPF), initially developed by Cadence, is currently maintained by the Si2 Low Power Coalition (LPC). The Unified Power Format (UPF), backed by Synopsys, Mentor Graphics, and Magma Design Automation, is now the IEEE P1801 effort.

Both formats make it possible to express power intent and constraints to both design and verification tools throughout the RTL-to-GDSII flow (Figure 6). This includes features difficult or impossible to describe in HDLs, such as multiple power domains, retention, isolation, and level shifting. Last year Cadence announced a low-power design flow with broad CPF support across a range of IC implementation and verification tools. In early 2008 Synopsys, Mentor, and Magma all announced support for UPF. Some of the smaller EDA vendors support both – an example is ChipVision, whose PowerOpt SystemC low-power synthesis tool outputs constraints in both CPF and UPF.

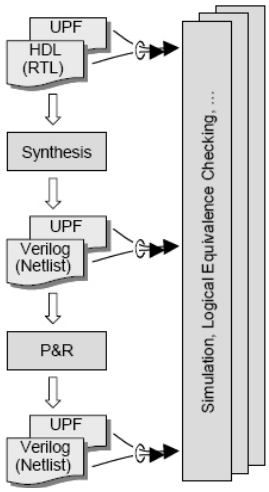

Figure 6 – The Unified Power Format (UPF) conveys power design intent to synthesis, simulation, and placement and routing tools. (Source: Accellera UPF 1.0 description).

Design teams are increasingly making use of these formats. Mentor's Bailey noted that UPF 1.0 adoption has grown faster than originally anticipated – so much so that the IEEE P1801 group recently recognized that it could not ignore backward compatibility considerations in drafting UPF 2.0. Hand said that a "huge number" of Cadence customers are using CPF and that "thousands of designs" have employed it.

"UPF is a great initiative to trigger power-aware RTL design and verification," said Infineon's Karmann. "UPF enables a consistent power intent configuration along the design flow."

The problem is that there are two of them, leading Atheros' Padnos to decry "that whole CPF/UPF nonsense" as one of the biggest challenges in the low-power design flow. Like many companies, Atheros has to use both. The company primarily uses Cadence tools, which take CPF, but must use UPF for other tools, making it necessary to go back and forth between two similar formats with different semantics.

Few people are as knowledgeable about both formats as Gary Delp, distinguished engineer at LSI Corp. As an architect for the Si2 LPC, he's heavily involved with CPF. He is also co-chair of the IEEE P1801 working group. Delp said that the use of these formats is still "fairly limited," with CPF implementation around a bit longer than UPF, but UPF adoption is increasing as more tools become available.

While clock gating can be expressed adequately in HDLs today, CPF and UPF are needed to fully express intent and constraints for power gating and multiple power domains, Delp said. "With different domains, you need level shifting," he said. "Being able to designate where the level shifters need to be inserted into the design, and what states the design might operate under, helps you determine which of those level shifters needs to be put into place and what can be optimized," he said. Thus, he said, the formats help designers model how different power domains will interact, and implement their design intent "in a way that is consistent with what has been modeled."

Delp said that LSI uses UPF today to drive simulations, and to work with customers who are handing off designs. Compared to CPF, he said, UPF has more "breadth" in what it can describe in terms of power verification and implementation, particularly in areas such as back biasing and retention. On the other hand, Delp said LSI also uses CPF in conjunction with Cadence tools to check power implementation, making it possible to "check to make sure that when you say things are in this power domain, they really are powered by the things you said they were powered by."

Today, both UPF and CPF are used primarily in the RTL-to-GDSII design flow, while both have aspirations to move upwards. Delp said that a UPF 2.0 draft now under development will add features to support ESL design, including the notion of "supply sets," which are bundles of supply nets that can be associated with specific states or power domains. Further, he said, UPF 2.0 will have "crisper" simulation semantics and a full understanding of hierarchy.

It would be possible to write a translator from CPF to UPF, but it would be more difficult to go in the other direction because UPF provides more "flexibility in expression," Delp said. He added that it would "certainly" be technically possible to combine both into a single standard, should the EDA vendors agree to do so. He hopes that will be the case. "There is a good number of formats and there is only one number which is better – one," Delp said.

VERIFICATION AND TEST CHALLENGES

CPF and UPF help ease the challenges of low-power design verification by bringing power intent to simulators and formal verification tools. But the fact remains that power management techniques strongly impact, and complicate, functional verification.

At the 2008 Design Automation Conference, Yoshio Inoue, chief engineer in the design technology division at Renesas Technology,

➔ Main article continues on page 16

LOW POWER SYSTEM-LEVEL DESIGN

by Glenn Perry
Mentor Graphics general manager, ESL and HDL Design division

Clock gating, power gating and voltage islands are commonly used techniques for low power design tuning and are available at RTL design and post synthesis. While these approaches are valuable, they are limited by the system level architecture decisions that have gone before. In contrast, optimizing at the architectural level allows the system to be tuned in the context of the overall target application, including the interactions between the software and data that the device must process.

Analysis at the system level dramatically expands the degree of flexibility in making tradeoffs related to performance, power requirements and user experience with minimal impact to the design schedule. Some examples of enhanced exploration capabilities available at the system level are the following:

- Data representation choices can be explored and the impact fully characterized. For example, wireless applications can realize huge power savings by optimizing the hardware-intensive compression/decompression tasks vs. transmission costs.
- Data processing requirements can be quantified; decisions such as the speed and number of processors vs. additional hardware accelerators can be clearly understood.
- Sleep and power down modes can be examined and the impact on system performance and user experience can be assessed and fine tuned.
- Cache and memory management decisions can be analyzed to maximize system performance while minimizing power requirements.
- Bus structure and arbitration schemes can be implemented to minimize system bottlenecks and wasted power caused by contention and transfer delays.

These system level optimizations can easily produce power improvements that are orders of magnitude beyond what is possible at RTL design or post synthesis. However, they require investment in appropriate system-level tools and modeling to accurately perform these optimizations. For example, using high level synthesis tools to describe and synthesize an algorithm in untimed ANSI C++ is vastly more concise than RTL, thereby enabling extensive design exploration. The designer can find solutions balancing area, performance and power outside the scope of hand-coded RTL such as:

- Optimizing individual blocks for power by scaling clock frequencies, and using closed-loop power analysis for early power estimation.
- Supporting C++ bit-accurate data types, which allows sophisticated numerical refinement to reduce area and power, while remaining within error tolerances and supporting industry-standard C++ compilers.
- Sharing hardware memory and computing resources, which is much more effective in ESL synthesis than RTL synthesis to save power.
- Generating verified RTL from C++ that automatically supports clock gating and interfaces to system functions for dynamic clock and voltage management.

Solution space exploration is paramount to finding the right tradeoff when making difficult choices between power, performance and user experience. Fortunately, ESL solutions have matured to address the growing demand for system-level power optimization. With the aid of these tools, it is now possible to evaluate and optimize power, performance as well as the end-user experience without endangering tight development schedules. With increasingly powerful system level solutions, designers now have the tools needed to tackle low power with confidence.

For more information on achieving optimal designs using ESL technologies, go to www.mentor.com/esl

For further details, or to have your question addressed, please visit: http://www.mentor.com/products/product_maincontact.cfm

➔ continued from page 15

discussed the challenges of verifying a chip with over 20 power domains. He noted that the different power modes resulted in over 36,000 possible test sequences. Inoue said he is wary of new low-power design methodologies because "I'm not sure I can complete verification."

Low-power verification raises two concerns, according to Hui. One is making sure that state retention issues are handled correctly when power domains are turned off. Another concern is the time it takes to turn power domains on and off. "For example," he said, "when you turn a power domain on, how much time does it take for this power gate to actually power on the design?" That's one issue CPF doesn't handle today, he noted, even though CPF does a "pretty good job" in conjunction with the Cadence simulator.

"Low power affects verification in several ways," said Padnos. "You have to make sure level shifters are included where you need them, and that that is carried through the entire flow. And you have to make sure that, as you go in and out of different power states, the state machine handles the transitions correctly."

In an SCDsource article titled "How power design techniques impact SoC verification," consulant Bhanu Kapoor, president of Mimasic, outlined a number of verification concerns raised by power management techniques. For example, power domains should be isolated under power-down to avoid floating signals driving active regions, and voltage domain crossings must be correctly level shifted.

The big picture is that power management adds new power states that must be verified. No longer is a chip simply "on" or "off." Chips today might have 20 voltage islands that are switching off and on independently. "These are power states that didn't exist before," Kapoor said. "Now we have to ensure that the chip functions in each of these power states, and functions as we transition these states."

Testbenches, therefore, must be written so that simulation can validate the design in each of the different power states, and during the transition time from one state to another. Simulators must be "power aware" so they can recognize different power states. Finally, Kapoor noted, functional coverage needs to ensure that all the different power states are covered.

Meanwhile, formal verification is becoming more power-aware. Hui said that AMD uses Cadence's Conformal LP as its power signoff tool. In addition to equivalence checking, he said, it checks the legality of voltage islands and isolation cells and the correctness of power gating. "We run Conformal LP on the output against the original netlist, so if anything is done incorrectly by any tool, we can catch it," he said.

Mentor's Bailey said that the use of power management techniques requires simulators to understand connectivity from a power domain perspective, infer registers in the design, simulate and evaluate supply networks, and consider behavior such as power shutoff, non-operational body bias modes, and non-shifted logic voltage values. "Of course, the debug and analysis environment needs to understand the power specification," he added.

"PM [power management] techniques have changed the verification paradigm," said Synopsys' Vivolo. "A traditional 0-1 simulation doesn't understand voltage values and will miss bugs, causing silicon failures and respins." What's needed, he said, is a "voltage-aware, low power simulator," and he identified Synopsys' MVSIM as one such tool.

"The biggest bottleneck in the low power verification flow today is a lack of verification methodology for low power," said Vivolo. To fill this gap, he said, Synopsys is working with ARM, Renesas and other companies to develop a verification methodology for low-power design called VMM-LP. It includes a book and a set of base classes set for release in fall 2008.

Hand said Cadence offers support for power verification today, and also supports automatic test generation capabilities to ensure that power modes are tested. "We are adding the capability to generate necessary assertions to track different power modes," he said. Hand noted that Cadence offers low-power formal verification through both the Conformal LP equivalence checker and the Incisive Formal Verifier (IFV) model checker.

Complicating test

Low power design complicates test, said LSI's Desai, because when you place a chip on a tester and start toggling a large number of circuits, it negatively impacts power. "You may end up using more power than the circuit was designed for, because in actual function, it will never toggle that many nodes," he said. "So do you design the circuit for this level of test power, or do you pick the right vectors so the test itself is aware of the limitations of handling power?" The latter would be a better solution, but will require enhancements to automatic test pattern generation (ATPG) tools, he noted.

One development that may become a challenge for test is the use of many small power domains. "If a power gate fails, I don't know how to tell," Hui said. "A single power gate failing on one big power island is not a big deal. A power gate failing on a small island is a big issue. A lot of times, you won't be able to test that."

"A chip's power constraints can adversely affect manufacturing test unless ATPG is power aware," said Greg Aldrich, marketing director of Mentor's design for test division. Referring to Mentor's TestKompress, he said "fortunately, there is a test pattern compression solution that enables effective testing while staying within the power budget."

"Testing of designs that implement PM techniques can be impacted dramatically if care is not taken in advance," said Synopsys' Vivolo. "Scan chains, for example, can span across multiple power domains. Not only must the scan insertion program take into account voltage boundaries, it must also be aware of power tables and transitions." Vivolo also noted that a chip tester armed with the wrong test vectors "can theoretically create unrealistic test scenarios which could easily consume more power than the device would ever use in a real world environment."

Vivolo observed that Synopsys' DFTMax supports UPF today, and said the company's TetraMAX ATPG solution can account for PM techniques by creating test vectors that limit simultaneous switching throughout the test process.

CONCLUSION: FROM CHIPS TO SYSTEMS

It's clear that power management impacts the entire SoC design and verification flow. The greatest potential gains are at the application and architectural levels, where design teams make crucial decisions about memory, processors, hardware/software partitioning, and the use of multiple power domains – but there's little automation in terms of tool support. There's reasonably good support in the RTL-to-GDSII flow for techniques such as clock gating and multiple voltage islands, but that's the low-hanging fruit. The challenge now is to reach towards the top of the tree.

It's beyond the scope of this report, but what really matters is the total power consumption of the end product that's delivered to the consumer. That comes not only from the SoC, but also the SoC's packaging, the board on which it sits, and all the other ICs and discrete components on that and other boards, analog as well as digital. It also concerns the mechanical enclosure, the power supply, and the absence or presence of cooling elements such as fans. Embedded software applications and operating systems must also support power management. Ultimately, power analysis and optimization needs to move up to the full application and system level and take all these factors into account.

-- Additional reporting by Bill Murray

RELATED ARTICLES

Low power flows still need work, speakers say
Unified Power Format moves towards standardization
Understanding RTL power reduction techniques
Analyzing IC power at the electronic system level
Keynote: System-level approaches scale the power wall
Synopsys offers low-power IC design suite
'Off by design' architectures curb energy waste
It's time to shift the low-power debate
ChipVision offers ESL synthesis for low-power ICs
Resonant clocking research promises IC energy savings
Designers cite low-power challenges, solutions
Optimize IC power by understanding circuit activity
How power design techniques impact SoC verification
How high-level modeling speeds low-power design
TI's Richard Scales use cases guide power optimization