WHITE PAPER

A Holistic Approach to Energy-Efficient System-on-Chip (SoC) Design

Author

Piyush Sancheti
Senior Director, System Architects Group, Synopsys

Introduction

It takes a great deal of energy to power the modern world, and demand grows every day. This is especially true for electronics, where ever increasing automation and more intelligent devices incessantly demand more power. Many applications that use chips face a variety of pressures for reduced power consumption and better energy efficiency. In response, the semiconductor and electronic design automation (EDA) industries have developed a wide range of techniques to meet these requirements. This white paper provides some background on the problem, describes some of the available technologies to help and presents a holistic solution for energy efficiency through the system-on-chip (SoC) design flow, from architecture to signoff.

Drivers of Energy Demand

In the simplest terms, more electronic devices require more power to operate. In the last several decades, almost every aspect of daily life has evolved to rely more on electronic gadgets backed by massive data and storage centers. Naturally, demand is growing, but the percentage of energy consumption due to electronics is also growing. Figure 1 shows the growth in worldwide electricity demand due to information and communications technology (ICT), from the smallest consumer devices to the largest compute servers and networking equipment. Consumption due just to ICT applications is projected to nearly triple from 2022 to 2030. By the end of that period, ICT is expected to consume more than 20 percent of all electricity produced for all purposes.

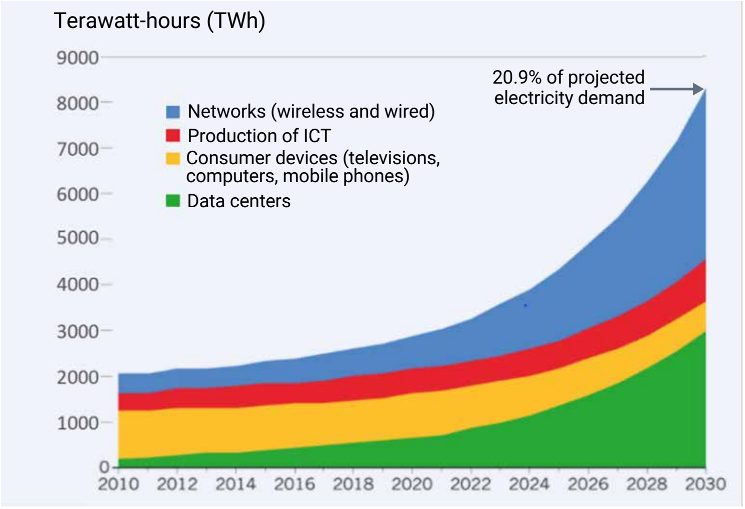

Source: www.nature.com

Figure 1: Growth in electricity demand for ICT applications

Every additional watt consumed must come from somewhere, and the cost to keep up with this demand is huge in terms of both dollars and societal impact. There is clear motivation for chip designers and manufacturers to make their devices as energy efficient as possible while still meeting the key requirements for each application. Low power consumption can be an important product differentiator, especially for devices that run on batteries. Historically, most of the work on improving efficiency and lowering power requirements has focused on portable consumer devices. Users want as much functionality as they can get while having reasonable battery life, and so the power performance tradeoffs made both during the design process and on the fly during device operation are critical.

Heat dissipation is also a major challenge for small form factors such as smartphones, tablets, laptops and wearable technology. Consumers must be able to hold these devices without discomfort, but spreading out the heat to avoid hot spots adds cost and bulk to the products. The only solution is to manage power consumption at the source. It is worth noting that small devices do not necessarily contain small chips. While many Internet of Things (IoT) applications are relatively simple, some of the largest and most complex SoCs in the world are used in smartphones. Autonomous vehicles are not exactly small applications, but they are still consumer devices and also require sophisticated SoCs, sensors and other electronic components.

Larger systems, from desktops to servers, offer more options for heat dissipation and thermal management. Advanced chip packaging, large heat sinks, fans and even liquid cooling are feasible for many designs, but they add significant cost. Ultra-low power is not required but power consumption must be managed. Studies have shown that the cost to power a data center server over its useful lifetime exceeds the initial outlay for the hardware itself. There is clear financial motivation to make the SoCs in these systems energy efficient as well. Many buyers also want to reduce their carbon footprint whenever possible, whether due to personal concern over global climate change or to "green laws" that limit power consumption for data centers and even individual classes of machines. New applications, especially those powered by artificial intelligence (AI) and machine learning (ML), constantly demand more power. A better energy aware design process is required.

End-to-End Energy-Efficient Design

For optimal results, energy efficiency must be addressed at each stage of SoC design. Over the years, a wide range of techniques has been developed to lower and manage power consumption. Figure 2 shows the most important of these techniques, which start at the physical level. Although the word "silicon" is often used generically as a term for semiconductors, other materials have been used, and power considerations may be a factor in the selection. For example, gallium arsenide (GaAs) was once a significant competitor to silicon for some high performance applications, but silicon's higher thermal conductivity means that it can dissipate heat better and therefore handle greater power without overheating.

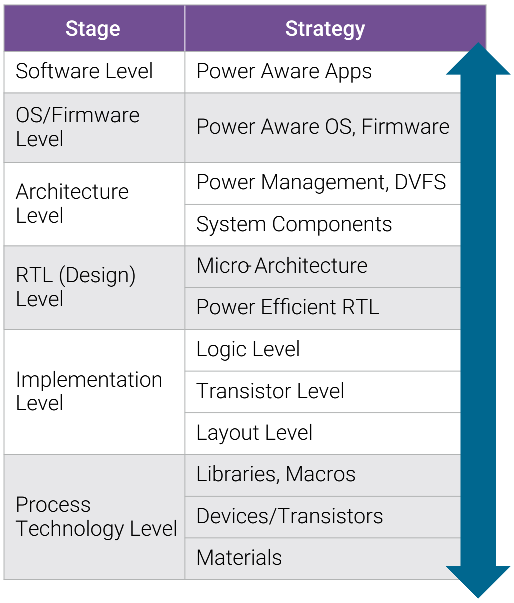

Stage	Strategy
Software Level	Power Aware Apps
OS/Firmware Level	Power Aware OS, Firmware
Architecture Level	Power Management, DVFS
	System Components
RTL (Design) Level	Micro-Architecture
	Power Efficient RTL
Implementation Level	Logic Level
	Transistor Level
	Layout Level
Process Technology Level	Libraries, Macros
	Devices/Transistors
	Materials

Figure 2: End-to-end energy-efficient design flow

Above the base materials, the structure of the transistors and other devices also has a strong effect on energy efficiency. From the earliest days of chip development, engineers made power performance area (PPA) tradeoffs by choosing the transistors with the characteristics that best matched the design targets. The dominance of fin field-effect transistor (FinFET) devices in chips of 16nm or smaller process nodes is a good example. As designers moved to deep submicron geometries, they found that leakage current went from being an annoyance to a major factor in power consumption. This issue has been ameliorated by FinFETs, which have better control over the device channel.

The majority of SoC design is performed not at the transistor level, but rather with register transfer level (RTL)—or higher level—code that is synthesized using cell libraries of common functions. Ever since the introduction of standard cells and gate arrays decades ago, libraries have contained multiple cells with equivalent functionality but different PPA characteristics. At a minimum, many libraries contained "low power" versions of many cells and designers could choose these for functionality that was not on a critical path. With the ability of logic synthesis tools to rapidly try different mappings to the cells, the more options in the library the easier it is to meet PPA targets with no need for designer intervention. Since physical effects become dominant at advanced nodes, all analysis and optimization technologies need to be physically aware and rely on high quality libraries and intellectual property (IP). Power consumption impacts the power integrity and thermal characteristics of the chip, so they must be addressed during physical implementation (synthesis and layout) and confirmed during signoff.

Energy efficiency also places requirements on the RTL design and the gate level netlist generated from it by logic synthesis. This is where the designer defines the micro-architecture, which must take into account the power consumption goals. Common examples include the ability to shut down portions of the SoC not currently active, put them in standby mode, or control operation on the fly with dynamic voltage and frequency scaling (DVFS). It is important to find the right micro-architecture so that it does not become a bottleneck during later stages of implementation when there are fewer degrees of freedom. Advanced techniques such as DVFS require support at the macro-architectural level, so overall system power management must be part of the high level chip modeling. The SoC architects must define the power control structures and provide hooks so that they can be manipulated by production software running on the end system.

Software is the final part of the solution. Although some power management decisions can be made purely at the hardware level, for most modern SoCs the power aware firmware, operating system (OS) and applications (apps) provide most of the control. For example, an operating system knows all the apps and tasks running or scheduled to run and therefore can make intelligent decisions on powering off or slowing down portions of the chip where top performance is not needed. Some SoCs have so much functionality that the entire chip cannot run at full speed simultaneously without causing thermal damage. The operating system must take this into account when choosing how to manage the power. The programs used to test dies and chips on the manufacturing floor must also be power aware to avoid overheating.

The techniques shown in Figure 2 are applied throughout the SoC project, as shown in the timeline of Figure 3. A low power design flow must use actual software workloads with fast profiling to drive the activity (vectors) for power exploration, analysis and optimization. Typical software workloads are in the billions of cycles (generating terabytes of data), so a successive refinement approach is used to narrow down the key windows of interest for each stage of the process to make the data size more manageable. The result is that the activity derived from the software workload drives the power exploration, analysis and optimization at each stage of the flow, from architecture and RTL to implementation and signoff.

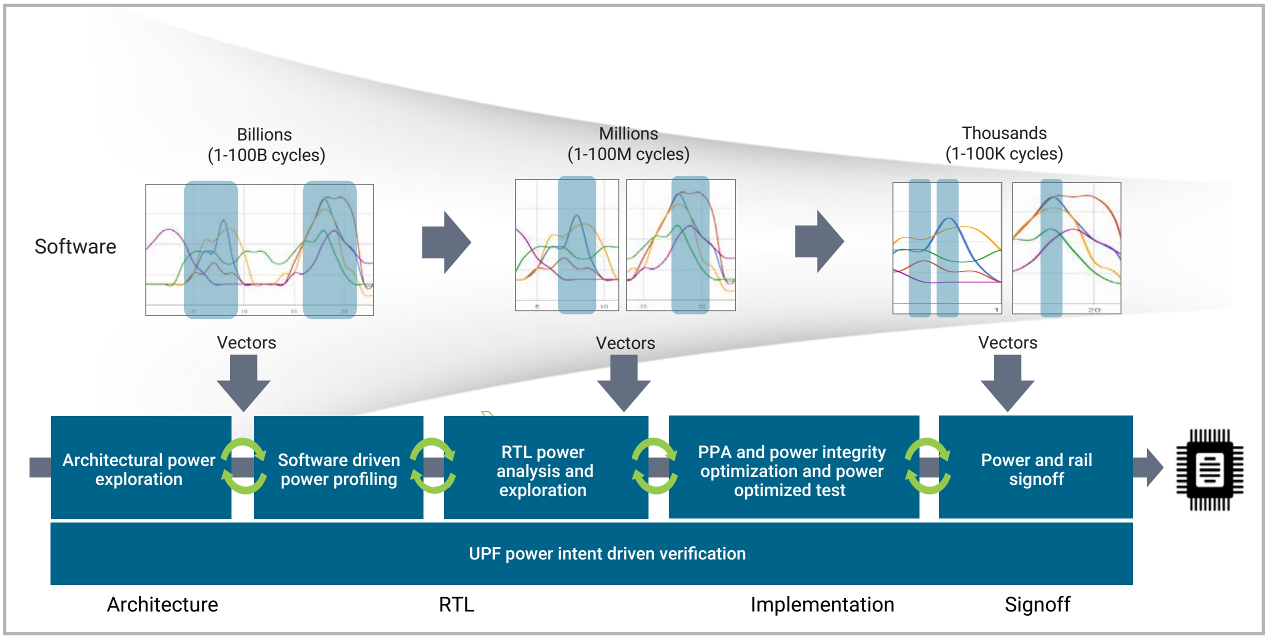
Billions
(1-100B cycles)
Millions
(1-100M cycles)
Thousands
(1-100K cycles)
Software
Vectors
Vectors
Vectors
Architectural power exploration
Software driven power profiling
RTL power analysis and exploration
PPA and power integrity optimization and power optimized test
Power and rail signoff
UPF power intent driven verification
Architecture
RTL
Implementation
Signoff

Figure 3: Low power design flow

Managing energy efficiency across the complete low power design flow has been easier since the introduction of the Unified Power Format (UPF) standard (IEEE 1801-2018). UPF specifies many aspects of the power control network for the SoC, including:

- Power supply nets and power switches
- Power/voltage domains
- Level shifters and isolation cells between these domains
- Power states and transitions between these states
- Memory retention when power is turned off to a portion of the chip

The UPF file is a specification of power intent. Design tools read this file and use its contents to guide implementation through logic synthesis, placement and routing. Architectural tools use UPF to reflect power management in virtual models and help make the macro-architectural tradeoffs. Many verification tools (simulation, emulation, formal analysis and signoff checks) also take power structures into account.

Synopsys End-to-End Low Power Solution

Architects, designers and verification engineers looking for an end-to-end energy-efficient SoC development flow have a solution available today, proven on many chip projects. The Synopsys low power solution provides software driven power verification, exploration, analysis and optimization from architecture to signoff. Figure 4 overlays the flow in Figure 3 with the Synopsys products that contribute to the industry's most complete and most effective solution. These include:

- Platform Architect™ for architecture exploration and early performance power tradeoffs using pre-RTL architecture models and software workloads
- ZeBu® Empower for power emulation with the capacity and performance for profiling software workloads to identify key windows of interest for further analysis and exploration
- SpyGlass® Power for RTL power exploration with fast turnaround time during initial stages of RTL development
- PrimePower RTL with RTL Architect for RTL power exploration with high accuracy as the RTL matures
- Fusion Compiler™ for RTL to GDSII implementation with the best PPA results
- Fusion of PrimePower and RedHawk™ (from Ansys®) signoff engines for fastest convergence and best quality of results (QoR)
- PrimePower and Ansys RedHawk for golden signoff
- TestMAX™ for power optimized automatic test pattern generation (ATPG) to ensure power is managed properly when chips are tested

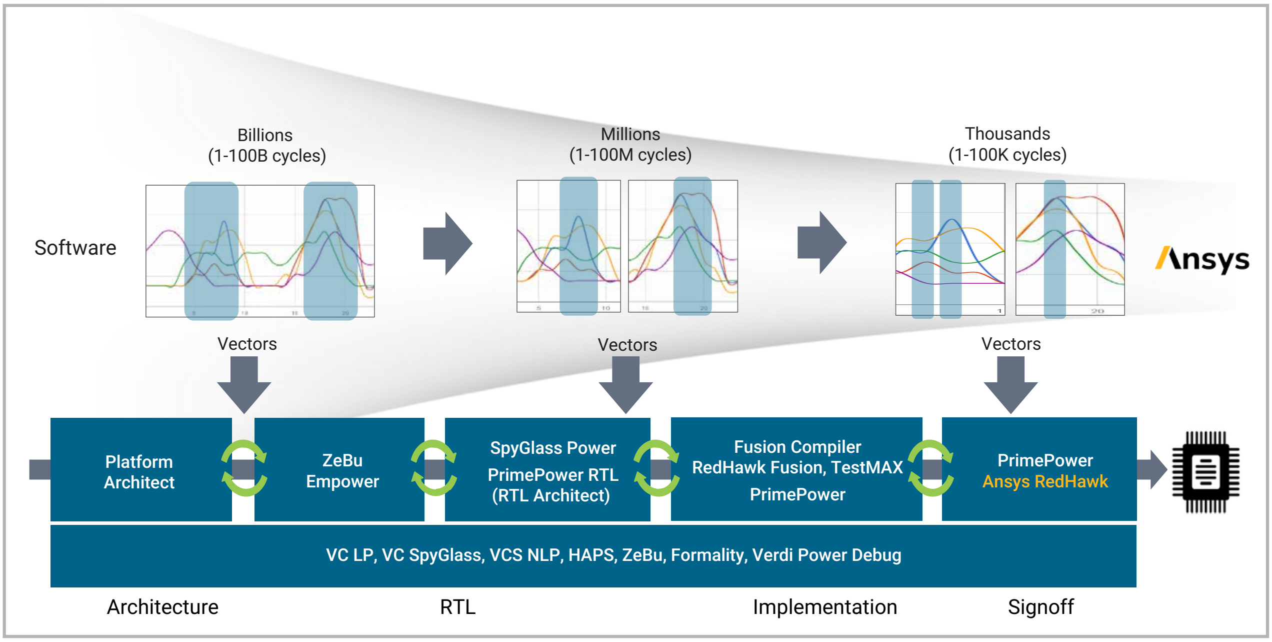

Figure 4: The Synopsys low power solution

The Synopsys development flow produces optimal energy efficient results by maximizing power reduction opportunities at each stage of design. The power performance tradeoffs are based on actual software workloads to avoid unpleasant surprises when the SoC is run in the bring-up lab. Early and accurate power analysis ensures fast and predictable convergence to the PPA goals. Figure 5 drills down into the details of the Synopsys flow, showing how the components of the solution are applied at different levels of design. During the architectural phase, Platform Architect provides the unique capability to use abstract models for exploring macro-architectural options and their power performance tradeoffs. Some IP development can proceed in parallel, but most of the RTL design occurs after the macro-architecture is settled.

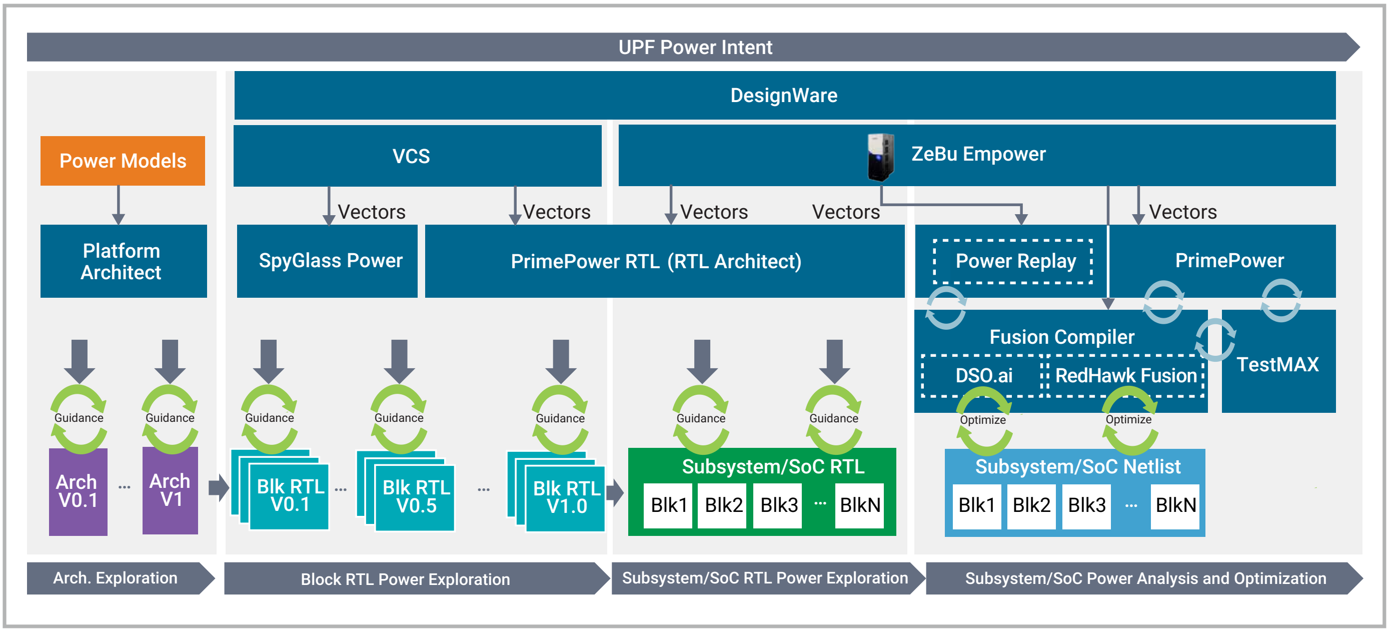

Figure 5: Software driven power exploration, analysis and optimization

During the early stages of RTL block development, as the micro-architecture is being defined, SpyGlass Power is used with vectors generated by the VCS® simulator for early power exploration. As the RTL blocks mature and move closer to implementation, PrimePower RTL provides more accurate analysis built on embedded RTL Architect physical and timing aware predictive technology and signoff PrimePower engines. Once the SoC or subsystem is ready for emulation, ZeBu Empower is used to profile software workloads to identify windows of high interest (such as peak power and regions of high average power) that can then be used for more detailed analysis in PrimePower RTL. Only ZeBu Empower has the power emulation capacity and performance to handle full SoC software workloads with billions of cycles.

As the design goes into implementation, further refinement of the activity windows drives implementation in Fusion Compiler. Its power driven RTL to GDSII flow provides optimal PPA results quickly, aided by the artificial intelligence engine DSO.ai™ (Design Space Optimization AI). DOS.ai, the industry's first autonomous application for chip design, searches for optimization targets in the exceptionally large solution spaces of SoC designs. The activity windows from ZeBu Empower also drive signoff in PrimePower. Its Power Replay capability reuses the vectors from VCS RTL simulations on the gate-level netlist after implementation. Golden power signoff with PrimePower includes key technologies for glitch analysis and debug, delay shifting for timing accuracy, and modeling for advanced process nodes. Finally, TestMax takes power considerations into account during manufacturing test generation.

In addition to the custom RTL portions of the design, every SoC makes use of commercial IP. Synopsys provides the broad DesignWare® portfolio of low power IP for processors, interface, sensors, analog/mixed-signal (AMS), memories and logic libraries. These come with predefined UPF descriptions to complement the files provided by the designers. All the libraries and IP, and all the tools shown in Figure 4, are linked together with consistent UPF support. This unified approach includes the Synopsys tools used for low power verification, shown in more detail in Figure 6. This combination provides comprehensive low power verification for block, subsystem and complete SoCs based on UPF power intent.

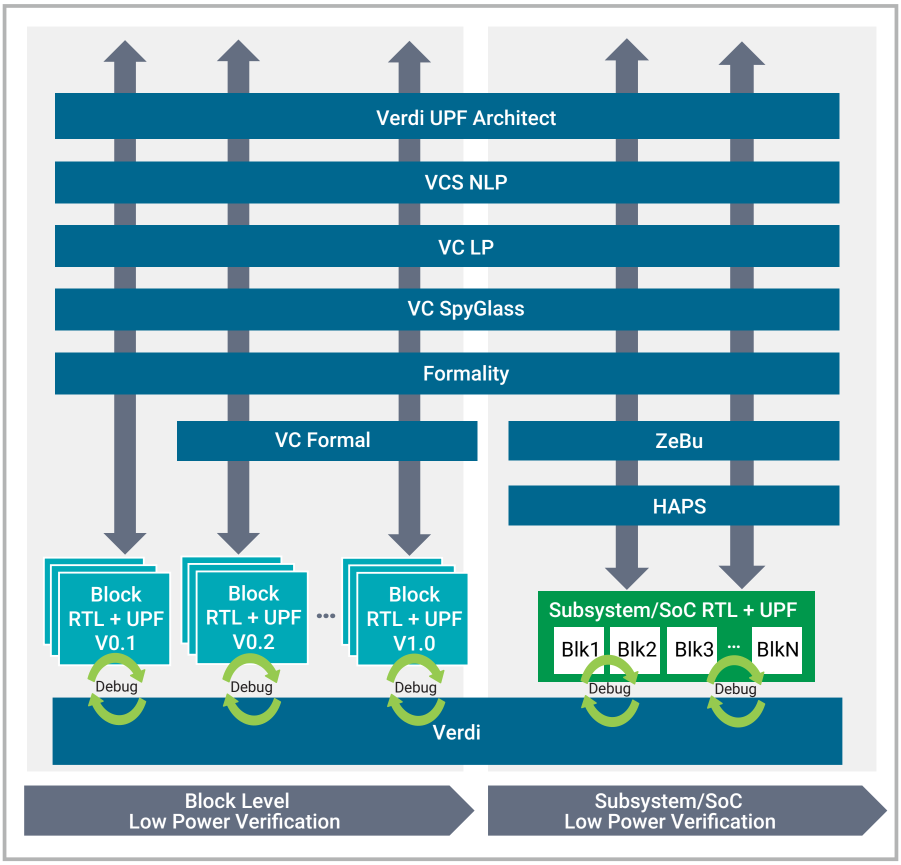

Figure 6: Synopsys low power verification

The process begins with correct-by-construction UPF generation with Verdi® UPF Architect, part of the industry leading Verdi Debug Solution. Verdi UPF Architect enables users to generate UPF files based on high-level descriptions of the power intent for the design. The UPF is loaded into VCS® Native Low Power (NLP) along with the design and the testbench. VCS NLP models the entire power network described in UPF and takes power events into account during dynamic simulation. For example, if the control signal for a power domain switches it off, then VCS LP will set all signals in that domain to unknown. The VC LP™ static low power verification solution includes over 650 checks to catch low power bugs earlier and faster than traditional methods, including:

- Power Intent Consistency Checks: Syntax and semantic checks on UPF that help validate the consistency of UPF prior to implementation
- Architectural Checks: Global checks at RTL for signals violating power architecture rules
- Structural and Power and Ground (PG) Checks: Validation of insertion and connection of isolation cells, power switches, level shifters, retention registers and always-on cells throughout the implementation flow
- Functional Checks: Checks for the correct functionality of isolation cells and power switches

These checks can be run at any stage of design from RTL to the final layout netlist. Additionally, these checks can be invoked directly from Fusion Compiler at various stages of the RTL-to-GDSII flow to ensure that the power intent of the design is preserved during implementation. The VC SpyGlass™ RTL Static Signoff platform also reads UPF so that checks for clock domain crossing (CDC) and reset domain crossing (RDC) instances are power aware. The same is true for the logical equivalence checking (LEC) performed by Formality® and the analysis performed by VC Formal™. Both the ZeBu emulation system and the HAPS® prototyping solution also take UPF into account. All these tools and technologies share Verdi as a unified debug platform, with many power aware debug features available. The result is a seamless flow from power intent specification through all phases of functional verification.

Conclusion

Many SoC applications require minimal power consumption to extend battery life or careful power management to meet regulatory and market requirements. However, the overall PPA goals cannot be comprised by focusing just on power. In response, the chip industry has developed an end-to-end design flow for energy efficiency at every stage. Early architectural exploration of power options must be followed by power aware implementation and verification, linked by a common UPF description of power intent as well as unified debug. The holistic Synopsys end-to-end low power solution is the industry's most mature and most advanced way to achieve end-to-end energy efficiency in SoC development. Every project for every application should have these capabilities.

02/17/22.CS829506549-A Holistic Approach to Energy-Efficient System-on-Chip (SoC) Design-WP.

Pub: Feb. 2022