

Managing Power and Performance for System-on-Chip Designs using Voltage Islands

David E. Lackey, Paul S. Zuchowski, Thomas R. Bednar,
Douglas W. Stout, Scott W. Gould, John M. Cohn

IBM Microelectronics Division
Essex Junction, Vermont 05452, USA

Abstract

This paper discusses Voltage Islands, a system architecture and chip implementation methodology, that can be used to dramatically reduce active and static power consumption for System-on-Chip (SoC) designs. As technology scales for increased circuit density and performance, the need to reduce power consumption increases in significance as designers strive to utilize the advancing silicon capabilities. The consumer product market further drives the need to minimize chip power consumption.

Effective use of Voltage Islands for meeting SoC power and performance requirements, while meeting Time to Market (TAT) demands, requires novel approaches throughout the design flow as well as special circuit components and chip powering structures. This paper outlines methods being used today to design Voltage Islands in a rapid-TAT product development environment, and discusses the need for industry EDA advances to create an industry-wide Voltage Island design capability.

Background

The total power consumed by conventional CMOS circuitry is composed of two sources. The first is active power, this power represents the power consumed by the intended work of the circuit to switch states and thus execute logic functions. Active power is primarily composed of the power associated with the charging or discharging of the capacitance of the switching nodes [1]. The magnitude of this power is given by equation 1.

$$P_{\text{active}} = C * V_{\text{dd}}^2 * F \quad (1)$$

As silicon technology scales, the capacitance per unit area and the frequency of operation increase by 30% each. Assuming perfect scaling, these increases are exactly offset by a corresponding 30% decrease in V_{dd} , and the power per unit area remains constant. However, frequency of operation has increased at a faster rate than the scaling of the silicon process technology [2,3]. This has led to an increase in power density in each technology generation. This increasing power density drives, in turn, the need for more expensive packaging and complex cooling solutions, and decreases reliability due to increased temperatures.

In addition to active power, there are components of leakage power, the most dominant of which is the sub-threshold current of the transistors in the circuit [4]. As silicon technologies advance, smaller geometries become possible, enabling improvements of device structures including lower transistor oxide thickness (T_{ox}), which in turn increases transistor performance. To maintain circuit reliability, V_{dd} must be lowered as T_{ox} is reduced [5]. As V_{dd} is reduced, the transistor threshold voltage (V_{t}) must be reduced in order to maintain or improve circuit performance, despite the drop in V_{dd} . This decrease in V_{t} and T_{ox} then drives significant increases in leakage power, which has previously been negligible. As silicon technologies move into the 0.13 μm and 90nm lithography generations, leakage currents become as important as active power in many applications.

The combination of increasing active power density and increasing leakage currents has created a power management problem in the semiconductor industry, as depicted in Figure 1:




Fig.1. Active/Leakage Power Density by L_{poly} Width [6]
(reproduced by permission)

Voltage Islands Overview

The power challenges posed by advanced technologies force system designers to make choices concerning device structures and voltage levels for the functions they are designing. In

previous generations, large functional blocks were not integrated on the same chip, so these choices could be made independently for each block. High levels of integration supported by SoC-enabling technology drive single chip implementations, where traditional approaches to power distribution and performance optimization fail to provide the flexibility of voltage and technology optimization of the previously disintegrated solution.

Voltage Island Scenarios and Examples

The concept of “Voltage Islands” restores the concept of individual voltage optimization of functional blocks to SoC design. Individual functional blocks of the SoC design can have power characteristics unique from the rest of the design, and can be optimized accordingly. There are numerous scenarios where Voltage Islands can provide design leverage.

One scenario, shown in Figure 2, involves identifying the minimum voltage required for each island to achieve it’s required performance. Often, the most performance-critical element of the design, such as a processor core, requires the highest voltage level supported by the technology in order to maximize it’s performance. Other functions which coexist on the SoC, such as memories or control logic, may not require this level of voltage, saving significant active power if they can be run at lower voltages. In addition, voltage flexibility allows pre-designed standard elements from other applications to be reused in a new SoC application. Further, some functions, such as embedded Analog cores, are specified at very specific voltages, and can be more easily accommodated in mixed voltage systems.




Fig.2. Timing-critical Voltage Island

Another scenario, shown in Figure 3, facilitates power savings in applications more sensitive to standby power, such as battery powered functions. Commonly, complex SoC designs consist of a number of diverse functions, few of which are active at any given time. Methods such as clock gating can be used to limit the active power from these idling functions, but the leakage (or standby) power remains, and can be significant in high performance technologies. If the power supplies for these functions are partitioned into islands, the function can be completely powered off, eliminating both active and standby components of power. Leveraging this concept requires that power management be built into the architecture and logic

design of the SoC, to handle power sequencing and communication issues.




Fig. 3. Voltage Island for Power-Sequencing

We will exhibit the above scenarios through two example designs. The first design is limited by the active power it may consume due to the thermal budget of the system. The second is a design whose power consumption is dominated by leakage, but requires only moderate performance when “on.”

Our first example is a 90nm chip which uses 144 high speed serial links, 10M logic gates and 20M bits of SRAM to translate between two different link protocols. The links run at a speed of 5 GHz. The rest of the chip runs at a more modest 500 MHz. In order to run at the required performance, the links require an operating voltage of 1.2V. The rest of the design, however, could meet timing at 1.0V, but without new techniques, the entire design is forced to operate at 1.2V due to the requirement of the high speed links. The active power of the design is estimated at 36.6W and the standby power is estimated at 0.5W. By placing the links and their associated high speed logic in one Voltage Island, and separating the rest of the logic into its own Voltage Island, this power can be reduced to 30.7W and the standby power is reduced to 0.25W. In this scenario, the links are supplied a continuous 1.2V supply while the rest of the logic utilizes a 1.0V supply. Due to the V_{dd}^2 term in the active power equation, the power of the design is greatly reduced.

Our second example is a 90nm chip that is used in a battery-operated application with a 0.9V power supply. It contains 5.5M gates of logic and 12M bits of SRAM. The device has three modes of operation: sleep, hibernate and active. The device is in sleep mode for 98% of the time. During hibernate mode, a portion of the design is required to run at 150 MHz. Hibernate mode occurs 2% of the time. The remainder of the design is associated with a high speed processor that runs at 400 MHz, but this portion of the design is only active 0.1% of the time. In this simple example, the standby current of the design prior to the use of Voltage Islands is estimated to be 17.41 mA (15.7 mW) at all times. During hibernate mode, the active power is estimated at 0.11W, and the total power (active +

leakage) is estimated at 0.12W. During active mode, the total power is estimated at 6.3W which is dominated by the active power.

To optimize the power of this design, we will utilize both cell-level and block-level power management techniques. We will first split the design into three Voltage Islands. The first island contains the circuitry that “listens” for the signal to wake up the rest of the chip. This logic consists of 50K gates of logic which must always be powered up. The second island, which is turned on only during hibernate mode, contains 500K gates and 2M bits of SRAM. The third island contains 5M gates of logic and 10M bits of SRAM. Using Voltage Islands, the design can be optimized to draw only 0.03 mA during standby. The circuitry that is “always on” is maintained by its own power supply and the voltage of this island is lowered to 0.7V, the minimum required for correct operation of the circuits. The rest of the chip circuitry is placed in two islands whose power supplies are switched totally off during standby. During hibernate mode, the second island is switched on. The estimated power during this mode is now 0.1W. The power in active mode is still 6.3W.

Through cell-level techniques, we can further reduce the power consumed in the standby mode. By converting the circuits in the “always on” island to high threshold transistors, the standby current can be further reduced to 0.26 μ W. Active power can also be reduced using cell level techniques. By selectively inserting lower threshold transistors for only the 10% most timing critical logic, the voltage can be reduced from 0.9V to 0.8V. This allows us to reduce the total power of the active mode to 5W. Although a larger percentage of this power is leakage or standby power, this increase is more than offset by the reduction in voltage and the V_{dd}^2 term in the active power equation.

Through these two design examples, the benefits of Voltage Islands are clearly seen.

System-level Power Management Approaches

Voltage Island techniques do not replace all other methods of power management, in fact Voltage islands concepts can complement and amplify the effectiveness of other techniques.

Clock gating can provide as much as a 20-30% power savings for high performance functions containing components that are inactive a significant portion of the overall time [7]. Clock gating can continue to be used for shorter duration “nap states” within a Voltage Island which can also be powered off for longer duration “sleep states”.

The use of multi-threshold libraries is becoming a common method for trading-off active and standby power for a function. Low threshold devices provide a performance advantage over higher threshold transistors, particularly at lower voltage. Using Low-Vt transistors can allow timing closure at a lower voltage level, which can be a great savings for overall active power. Low device thresholds also imply higher levels of leakage current, however, which can be detrimental to standby power sensitive applications. For this reason, logic libraries

utilizing high threshold transistors can be used in logic paths without critical timing. The higher voltage, required for these circuits to meet performance goals, can be justified by the reduction in standby power. In an SoC with varied performance and power requirements, these device and library options can be intermixed to optimize the diverse functions.

Voltage Island architecture methods can enhance the usefulness of multi-threshold design techniques. An island can be created to run an active power sensitive block with Low Vt's, at a lower voltage than the rest of the design. In addition, this leaky, low-Vt block can be shut off completely during sleep modes to eliminate standby power. Similarly, functions which are “always on” can be held at a higher voltage to accommodate less “leaky” high-Vt transistors, or be powered from a separate, back-up supply.




Fig. 4. Multi-level Voltage Islands Architecture

As shown in Figure 4, Voltage Islands can be used at different levels of the design hierarchy to amplify their effectiveness. A block which can be powered off could exist within a larger block which is running at a unique voltage, for example. Constructing a Voltage Island capability with a fine hierarchical granularity can enable a large variety of useful permutations.

Design Implications of Voltage Islands

An SoC architecture based on Voltage Islands requires additional design components to ensure reliable communications across island boundaries, distribute and manage power, and save and restore logic state during power-off and on.

Multiple Power Sources

A Voltage Island represents a level of hierarchy with unique powering, and exists within a parent block which constitutes the physical region in which the island is placed. An island's parent block may be the top level of a chip design or even another island at the next highest level of chip hierarchy. As shown in Figure 5, the circuits within a Voltage Island are primarily powered from the island voltage, called VDDI

(VDD-island or VDD-inside), while the circuits in the parent terrain are powered from a supply voltage called VDDO (VDD-outside). With deeper hierarchy, the VDDO of one island may be equivalent to the VDDI of a parent island in which it is contained.




Fig.5. Voltage Island Powering and Switching Control

Voltage Island Boundary Requirements

The relationship between the voltages (VDDI and VDDO) of an island and its parent block may vary considerably depending on how Voltage Islands are being employed. For example, a dynamically powered island might have VDDI greater than VDDO when operating at maximum performance, VDDI less than VDDO when operating at reduced performance or to preserve states, and VDDI=0 volts when fully powered down for standby current control.

These methods of voltage variation present a real problem for traditional, static CMOS logic gates. When such a gate operates at a voltage sufficiently lower than the gate it drives, signal margins and performance will degrade, and the driven circuit will consume significantly higher power. Further increases in the voltage difference will eventually result in unreliable signal switching. Clearly, additional circuitry is necessary to handle the differences in both magnitude and timing that can occur between VDDI and VDDO at island boundaries.

Circuits called Voltage Island Receivers perform this function for signals going from the parent block into the island, while Voltage Island Driver cells perform the equivalent for signals from island to parent block. These Drivers and Receivers must provide reliable voltage level shifting from VDDI and VDDO for a wide range of operating voltages, and do so with minimal impact to signal delay or duty cycle.

In applications where VDDI or VDDO are allowed to assume voltage values below those necessary to support reliable signal switching, the Voltage Island boundary also requires functions to disable communications across island boundaries and provide reliably controlled states (eg. logic 0, logic 1, or hold last active state) to downstream logic. Such an operation,

known as fencing, prevents the undesired propagation of unknown (X) states by powered-off logic.

Many possibilities exist for powering Voltage Islands, driving further requirements for special design components. VDDI or VDDO may be supplied directly from a unique, non-switched power distribution. One or both may be the output of an on-chip voltage regulator, whose voltage value may be fixed or programmable. Finally, VDDI or VDDO may be a switched version of some other voltage supply, controlled by one or more PFET or NFET switches. A given SoC design may use one or more of these approaches depending upon the product design objectives.

State-saving

Leakage or standby power can be reduced by lowering the voltage of functionally-inactive islands well below the level required for reliable operation. However, some subset of the logic state, prior to power-down, may need to be preserved to resume operation once the island is again powered up, at the end of the inactive period. Special state-saving latches provide a solution to this problem, eliminating the need to transfer logic states off-island and back in order to save and restore necessary logic states. Whereas a standard latch in a given island would operate from the island voltage (VDDI), a state-saving latch is a modification of the standard latch, adding both a VDDO connection and a state control input to select between normal and state-saving operation. In normal operation, the state-saving latch behaves identically to the standard latch. In state-saving operation, the latch data is preserved in a portion of the latch powered only by VDDO, and all other latch inputs (clocks, data, scan) are ignored. As long as VDDO remains active, VDDI may be powered down without concern that unreliable logic levels will effect the latch's logic state. State-saving latches are designed to consume minimal power from VDDO. The Voltage Island can be quickly returned to normal operation once VDDI is restored, via the latch's input data port.

Design Planning and Chip Implementation




Fig.6. SoC Design Flow

A traditional methodology for SoC architecture and chip implementation includes the following steps, as shown above in Figure 6:

- ✍ Functional Partitioning
- ✍ Functional Implementation (RTL)
- ✍ Synthesis
- ✍ Design for Test (DFT) logic design and transformation
- ✍ Simulation
- ✍ Timing verification for entry to Physical Design (layout)
- ✍ Floorplanning
- ✍ Physical Design and timing optimization
- ✍ Final timing verification and release

Designing for Voltage Islands involves additional operations that affect each step in the design flow, and the successful integration of the new requirements can significantly affect the TAT of the chip design project.

The remainder of this section will discuss the methodology requirements for Voltage Island design, and describe methods proposed and in use for satisfying these requirements.

Functional Partitioning

The traditional process for the partitioning of an SoC design involves division and subdivision into an n-level functional hierarchy. The resulting functional components are grouped based upon minimizing the number and timing-criticality of signals that connect different groups. The chip area of each group is maintained between minimum and maximum sizes (high performance requirements may reduce maximum size of a group, and the need to limit floorplanning complexity may in turn limit minimum group size). Recently, the EDA industry has created a new wave of tools intended to aid the designer in chip partitioning. The methods employed by these tools range from early SoC block-level planning [8], to physically-aware gate-abstraction techniques [9], to quick placement of the netlist for floorplanning insight [10].

Designing for Voltage Islands places additional requirements on the functional partitioning process. The SoC designer seeks an optimal voltage for each functional component that minimizes active power at the required performance, and the designer needs to identify components whose voltage supply can be independently sequenced. Satisfying these requirements means finding a partitioning solution that minimizes chip power within additional chip-level constraints: maximum peak power, the available voltage range of each power source, and the maximum peak and average power for each power source.

First, the operating states of each functional component are defined: operating (active), or inactive. For the inactive period, a component might be considered as a Voltage Island whose power source can be turned off (ie. a power-sequenced Voltage Island), if the duration of the inactive period exceed the minimum power on-to-off and off-to-on duration of the power-switching circuitry. Next, the power requirements of

state-saving circuits needed within the Voltage Island must be satisfied by the alternate source and accompanying power supply structure. If both these conditions are satisfied (required duration of the inactive state, and limit to the alternate source requirements of the state-saving circuitry), then one might proceed with a power-sequenced Voltage Island. If not, one might consider other methods of power reduction for inactive logic, such as gating the clock [7].

Next, the performance requirement of each functional component is considered, by specifying the voltage range under which the component will correctly operate. This evaluation suggests a set of SoC chip-level solutions, where each solution represents a mix of functional components each at a voltage within its performance operating range and where the SoC-level voltage mix satisfies all timing requirements. Advances to Static Timing tools in the industry are necessary that will provide quick performance evaluation of these potential chip-level solutions. As detailed in the following discussion on Timing Analysis requirements, conformance to IEEE 1481 provides a key component of this requirement: the ability to evaluate timing, in a single analysis, for multiple voltages across the chip and within a single path.

Given a set of SoC chip-level partitioning solutions that satisfy timing requirements, and given the set of power-off candidates discussed earlier, one can now proceed to the final grouping of functional components into Voltage Islands. This is an optimization problem with the following considerations:

- ✍ Identifying functional components with similar inactive periods.
- ✍ Assigning functional components to possible chip-level power sources capable of providing required voltage level.
- ✍ Identifying the optimal grouping of components, based upon power sequencing (affects static power) and operating voltage (affects active power) that minimizes chip power within the limits (such as peak power) of the SoC.
- ✍ Identifying or creating, and connecting, logic signals that will be used to control power-sequencing circuitry or control clock gates.
- ✍ Connecting alternate voltage sources to latches or arrays used to save state across power sequencing.

For SoC's to achieve the substantial power optimization capability possible through the Voltage Island technique, industry EDA tools for logic partitioning need to consider the additional objectives of optimal voltage assignment and active/inactive operating state.

Timing considerations in Synthesis and Static Timing

SoC's designed using the techniques described in this paper will contain logic paths that enter or exit Voltage Islands. As a result, such paths will contain logic gates operating at a mix of voltage levels. Since the voltage range of a gate (minimum or worst-case, maximum or best-case) significantly affects its delay characteristics, paths that traverse the boundaries of Voltage Islands present a new set of challenges for the steps in

the design flow that deal with timing behavior, particularly synthesis (logic and physical) and static timing analysis.

Once functional partitioning has been completed, the SoC designer can begin synthesizing the design. Each partition should be synthesized independently, with each Voltage Island synthesized at its targeted worst-case (WC) and best-case (BC) voltage levels.

Top-level SoC optimization takes on new challenges with Voltage Islands, as paths now exist that traverse voltage levels. This challenge is answered in part by synthesis tools that can

- ✍ Operate using IEEE 1481-compliant timing models (DCL) [11], which model multiple voltages within a single model, thus allowing a single timing library to be used for an SoC with Voltage Islands.
- ✍ Set each Island (or individual gate instances) to its unique voltage range (WC/BC), and calculate the timing of paths that traverse islands based upon the voltage assigned to each gate in the path.

For synthesis tools not supporting these capabilities, top-level synthesis becomes somewhat daunting, but can be approached through reporting of delays from the separate timing of each Voltage Island, creating black boxes of each Island for synthesizing the top-level, and applying the reported delays of each Island as timing constraints at the boundaries of the black boxes.

Additionally, top-level synthesis is more than a problem of timing calculation for multiple-voltage paths; providing the correct optimization choices (timing correction) is the objective. This requires the optimization function of synthesis to determine the optimum change within either a Voltage Island or at the chip's top-level. The requirement for voltage level-shifting circuits (Level Shifters) at the boundary signals of Voltage Islands needs to also be considered by synthesis: delay of Level Shifters in timing paths, and including this delay when inserting Level Shifters and making choices for partitioning and island voltage assignments.

Static timing analysis tools have all the requirements for delay calculation that are described above for synthesis. Further, the affects of the additional power supply structures need to be considered for timing sign-off: the tolerance of each power supply needs to be independently considered in the delay calculation. This tolerance is a function of both the off-chip source, as well as the effects of the on-chip supply and distribution structures. Calculation of internal voltage drop is determined by applying electrical simulation to the on-chip distribution, loads, and logic switching behavior. The tolerance of external sources is characterized by the SoC product designer, as well as determination of whether the power supplies vary in opposite directions (which increases variation of the resulting on-chip timing between Islands) or in like directions (reduces variation). Multiple voltage sources (including off-chip variation) also need to be considered when designing and timing clock distribution structures that span islands with different power sources. The voltage differences represent differences in the calculation of clock skew, and

independent variation of voltages can further add to the clock skew between Islands.

Design for Testability and Manufacturing Test

Design For Test (DFT) and full scan for Voltage Island designs raises the question of whether power-sequencing circuitry will be held to the power-on state during test operation. In this case, DFT is greatly simplified since any scan chain may thread through any Voltage Island, and all test clock and control signals are driven by powered logic throughout test. DFT automation can be accomplished, for the most part, without consideration of whether an island is power-sequenced. A downside of this approach is that any on-chip power-switching circuitry is not tested in the power off state.

An alternative to holding all islands in the power-on state is to isolate scan and test logic, and test operation, by Voltage Island; each power-sequenced island is then tested independently. At a cost of increased DFT logic complexity and potentially increased tester time, island power-switching circuits are allowed to operate as used in the system environment; however, special circuitry may be required to observe the power-off condition in the corresponding island.

Another question is whether the tester can supply a single common external voltage to the chip. Again, this can greatly simplify the cost of test, but at the potential expense of greater power consumption during test (if the common voltage is the highest level used amongst the Islands), and variation in logic timing behavior from that of the system environment. These issues can be alleviated by using on-chip voltage regulators to provide the unique voltages used by the various islands.

Voltage Islands can provide some advantages for burn-in testing, which is a technique used to improve long-term chip reliability. During burn-in, chips are tested under elevated voltages and temperatures to accelerate the wear-out mechanisms that would otherwise be seen over time for the chip in the field. Chip power increases dramatically under burn-in conditions. As chip power increases, the number of chips that can be burned in at one time decreases due to test equipment power supply limits. This ultimately limits throughput. Voltage Islands offer the ability to sequentially power down portions of the chip to keep the total average power load on the burn-in tester within limits. While this will increase the burn-in test time, the ability to burn-in more parts at the same time can improve total test throughput.

IDDq testing has long been a useful tool for quickly identifying bad chips by measuring quiescent current. As chip leakage power has increased, the small current differences that separate good chips from bad are becoming harder to detect. Voltage Islands allow the possibility of sequentially powering off sections of the chip. Not only does this allow the background leakage current level to be reduced, it also aids in localizing IDDq-detected faults such as power and ground shorts.

Logic Simulation

Power sequencing creates a requirement for simulation: Logic levels driven by a powered-off Voltage Island must be observed as *unknown* by observing logic that remains powered-on. This observed state is important for verifying that state-saving logic truly preserves functional values across a power-off/power-on sequence, and for verifying that any constant generators or state-holding logic (eg. fencing), that may be required for predictable operation between islands, is functioning as architected.

This capability is needed in logic simulators provided by the EDA industry. In the absence of such support, this requirement can be approached by user-defined simulation software that selectively forces the X-state at the output of powered-off logic gates. This enables the verification of logic intended to block unknown state propagation from powered-off islands.

Physical Planning and Implementation of Structures

A Voltage Island requires a minimum of one power source and complete isolation from all other supplies on the chip, in order to enable independent power sequencing. Physical planning of Voltage Islands includes determining the number of power sources that meet each island's power requirements. The following are some examples of off-chip and on-chip power source types:

- ✍ Solder bumps or C4 pads
- ✍ Wire bond pads
- ✍ Island power switches
- ✍ Voltage regulators

Voltage Islands must be floorplanned (placed) in close proximity with the chip pins corresponding to its off-chip power source. On-chip power sources must also be placed within the voltage island, or externally [12] for off-island power switches. To verify that no electromigration (EM) or IR-drop issues exist due to the relative placement of the island and its power sources, a power grid analysis must be performed [13].

Previous SoC designs have implemented Voltage Island structures using a custom core approach. Custom cores have separate power grid structures designed within them, and are isolated from the standard chip power grid. A PLL is an example of this approach, where its analog supply is isolated from the ASIC power grid. However, this custom core approach does not provide the full design flexibility required by SoC designers to leverage any functional partition as a Voltage Island. A general Voltage Island design capability requires power bus automation for connecting the power sources to all the required Voltage Island circuits, with little or no manual intervention.

Some types of Voltage Island circuits require both the internal voltage island supply (VDDI) and the external supply (VDDO) for proper operation. These include level-shifting circuits, state-saving latches, and power switches. Power-routing tools must be aware of these constraints and make efficient VDDI and VDDO connections without creating

electrical or wireability issues. [12] describes two approaches to connecting power to island power switches placed either internal to or external to the island.

Circuits within islands that require VDDO supplies may have control signals that must be driven by circuits powered outside the voltage island. One example is the "Sleep Mode" signal that controls an island's internal power switch. If that signal requires buffering within the island (for example, to avoid signal transition degradation, or to provide noise immunity), then the buffer circuits must also be powered by VDDO.

Coupled noise effects (causing timing degradation and false switching) must be accounted for in deep sub-micron SoC designs which contain Voltage Islands. For accurate coupled noise analysis, each circuit's voltage rail level and power-switching states must be considered by the coupled noise analysis tools, so that cross-coupling between signals within Voltage Islands and external signals (e.g. at the top-level) will be properly considered for their impact on timing delay and false signal switching.

Future considerations

An industry-wide capability to design SoC's with voltage islands requires key EDA capabilities as described earlier, including:

- ✍ Automatic partitioning that considers optimizing voltage levels and potential idle states of SoC functional partitions.
- ✍ Timing calculation and optimization capabilities that consider multiple voltages (including multiple voltage paths) and differing variations amongst multiple sources.
- ✍ Consideration in DFT for multiple voltages and powering down as part of logical behavior.
- ✍ Verification of system functions that manage logic state during power sequencing.
- ✍ Automatic floorplanning and placement with constraints on proximity to power sources.
- ✍ Power bus automation of power supply connections, including VDDI/VDDO requirements of special circuits, and repowering of island signals requiring VDDO.
- ✍ Noise coupling analysis that considers multiple voltages and power sequencing conditions.

Other considerations include:

- ✍ Optimization of performance and power late in the design cycle: Establishing alternate voltage regions by chip area, and using gate placement location in these areas to selectively decrease delay (higher voltage) or decrease power (lower voltage).
- ✍ Accurate power measurements that do not require costly simulation-based switching vectors, but instead can infer much of the switching characteristics from boolean and sequential design characteristics.
- ✍ Industry IP standards for multiple voltage functional interfaces and power, clock, and reset sequencing.

Conclusion

The Voltage Island architectures and design methods described in this paper provide an opportunity to achieve a significant reduction in power consumption for SoC's, designed within rapid-TAT development environments. Issues of both active power and leakage power can be addressed effectively using these techniques.

We described a working design methodology for Voltage Islands, and identified several fundamental enhancements needed in design tools and modeling to create an industry-wide Voltage Island capability. These include functional partitioning, synthesis, timing analysis, test, simulation, and physical design.

References

- [1] Enomoto, E., "Low Power Design Technology for Digital LSIs," IEICE Transactions on Electronics v E79-C n 12, Dec. 1996, pp 1639-1649.
- [2] Pollack, F.J., "New microarchitecture challenges in the coming generations of CMOS process technologies," Proceedings of the Annual International Symposium on Microarchitecture, 1999, p 2.
- [3] Borkar, S., "Design Challenges of Technology Scaling," IEEE Micro, Vol. 19, No. 4, July-August 1999, p 23-29.
- [4] Adan, A.O. and Higashi K., "OFF-State Leakage Current Mechanisms in BulkSi and SOI MOSFETs and Their Impact on CMOS ULSIs Standby Current," IEEE Transactions on Electron Devices, Vol. 48, No. 9, Sept. 2001, pp 2050-2057.
- [5] Stathis, J.H., "Physical and Predictive Models of Ultra Thin Oxide Reliability in CMOS Devices and Circuits," IEEE 39th Annual International Reliability Physics Symposium, Orlando, Florida, 2001, pp 132-149.
- [6] Nowack, E.J., "Maintaining the Benefits of CMOS Scaling when Scaling Bogs Down," IBM Journal of Research and Development, No. 2/3 March/May 2002.
- [7] Gary, S., "Low Power Microprocessor Design," Low Power Design Methodologies, Kluwer Academic Publishers, 1996, pp 279-281.
- [8] "IC Wizard - The Hierarchical Design Planning Tool," © 2002 Monterey Design Systems, Inc., http://www.mondes.com/prod_icw.html
- [9] "TeraForm® RTL Design Planner for Deep Submicron SOCs," © 2002 Tera Systems, Inc., <http://www.terasystems.com/products/datasheet.htm>
- [10] "First Encounter," © 2002 Cadence Design Systems, Inc., http://www.cadence.com/products/first_encounter.html
- [11] 1481-1999 IEEE Standard for Integrated Circuit (IC) Delay and Power Calculation System, copyright 1999 by IEEE.
- [12] Kosonocky, S.V.; Irrmediato, M.; Cottrell, P.; Hook, T.; Mann, R.; Brown, J., "Enhanced Multi-Threshold (MTCMOS) Circuits Using Variable Well Bias", Low Power Electronics and Design, International Symposium, 2001, pp 165-169.
- [13] Nassif, S.R. and Kozhaya, J.N., "Fast Power Grid Simulation", In Proc. Design Automation Conference, Los Angeles, CA, June 2000, pp 156-161.