" Words of Power : Reusable, Holistic, Scalable approach to Multi voltage design"

Herve Menager – NXP Semiconductors

For 65nm and below, with more and more computing power integrated on portable devices, low power is crucial. Power dissipated in a cmos circuit originates from two main sources. The dynamic part is consumed when switching states. The static part is due to leakage current. With latest process technology nodes, channel length, oxide thickness and voltage threshold are smaller, thus increasing leakage power to the point where it becomes as important as active power.
Control of both leakage and dynamic power consumption is now mandatory. Traditional approaches to low-power design are being revisited and new design techniques are necessary.
Designers have to look beyond minimizing leakage current through the use of multiple threshold voltages or stopping the switching power by gating the clocks.
Leakage can be addressed by suppressing current, thus switching the voltage to a section of the logic when it is not functionally needed. To meet both chip performance requirements and power goals, new techniques consist in using voltage islands or voltage domains. The most performance critical block of the design would dynamically get the highest voltage supply, when other less critical parts can be run on lower supply, thus saving dynamic power.
But when traditional approaches have been mostly solved within the scope a single step in the design flow with point tools, new ones are much more disruptive to the entire design implementation process.
This talk presents the multiple design implications and challenges associated with designing SoC with multiple supplies or blocks being totally off during standby. They span from design architecture and partitioning, to IP (level shifters, Power switches, state retention cells) with new library elements requirements, with numerous design tools and modeling issues to address (On or off-chip Power sources, Interface logic design, Interface logic checking, Design for Test, Floorplanning and power distribution, Placement of interface logic, Timing analysis, IR drop analysis, LVS, etc)

I. New design components are required

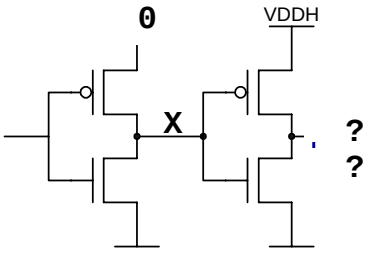

A block in power down mode cannot be directly connected to an always-on block, as its outputs will be temporary in unknown state as they gradually come to a logic 0 state. This unknown state can then propagate into the always-on section of the circuit, creating meta-stability. In addition, floating inputs will also generate short circuit current. Designing with switched blocks requires special care at the interface where signals are assumed as “unknown”. Simple nand/and or nor/or gates can be used to force a known state on the signal. Implementation choices might also require these cells to have a special layout.

In a design with blocks at different supplies, gates in a logical path spanning across voltage boundaries will operate at different voltages. This can create a problem in static CMOS. When exchanging signals between blocks at different VDD, it is necessary to insure full switching of the transistors. Going from a low vdd to a higher vdd, the output

of the VDDL gate cannot be raised higher than VDDL. When connected to a VDDH gate, the Pmos will never completely cut-off. It is thus not possible to directly connect VDDL and VDDH cells. Static current level converters are used to block this current. A key characteristic of the level shifter is that it is a std cell operating with two voltage supplies, thus creating a constraint for the layout implementation as well as for the logical connection.

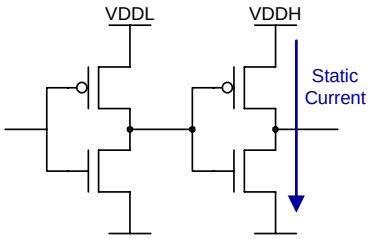

To shut-off a voltage to part of the circuit in order to minimize leakage, various implementations are possible ranging from the use of dedicated libraries where every cell is switch able individually to a more coarse approach where dedicate switches are added between a always on power net and a switched power net.

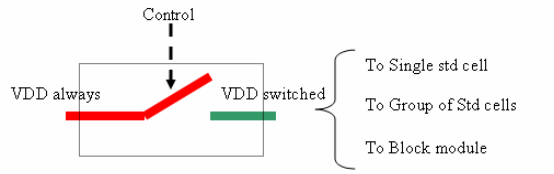

Each one of these techniques requires an additional floorplanning effort together with additional complexity for proper power distribution and analysis across the SoC.

II. New design components hamper the scalability and throughput time of the implementation

Experience has shown that EDA tools have mostly build patches to enable low power design with these new design components. Such ad hoc pragmatic approach however lacks fundamental holistic view and usually affect throughput time of the implementation. We have in some cases experienced 2X productivity drop for the back-end implementation phase. In addition to lack of tool functionality in various areas, this productivity penalty is due to the lack of scalability of the proposed approach to implement voltage islands.

- Addition of interface logic, whether it is isolation gates for power switching or level shifters for voltage scaling, does introduce additional verification challenges. Checks need to be run to verify proper isolation, proper connectivity to the right power domains, proper partitioning of the netlist, proper behavior of the interface and more.
- Similarly, a key characteristic of the level shifter is that it is a std cell operating with two voltage supplies, thus creating a constraint for the layout implementation.
- Always-on logic resulting from buffering of control logic for retention or global nets in power down blocks requires special care for proper connection of their supplies.
- Voltage islands and on-chip switches create a challenge for power distribution and limit the floorplan alternatives and flexibility. More effort is necessary for connecting power sources to the voltage domains.

- Communication between voltage islands may create logical path spanning power domains boundaries. This creates challenges for sign-off by increasing the number of corners and modes hence the number of STA runs.

It is not sufficient for tool vendors to address these areas by simply providing the basic low level hooks in their tool infrastructure. Tools understanding the same power design intent with the highest possible level of abstraction are needed to compensate the throughput time overhead introduced by designing with multiples supplies.

III. POWER CONNECTIVITY BECOMES PART OF THE FUNCTIONAL DESCRIPTION

In a SoC, power and ground nets have traditionally been defined and implemented outside of the scope of the logical design description. Logical views for basic library elements, as well as HDL descriptions did not have implicit representation of these nets as they would not have any functional impact. As a result, they were usually handled at a late stage during physical implementation and usually needed special handling through global connection in the back-end phase. The integrity of their implementation was usually performed by re defining the global power nets at LVS stage.

With power islands being turned on and off to minimize leakage current, power nets have now become partly functional as the behavior of the SoC now depends on the state of these nets. The number of voltage islands has increased the complexity of the description of the power architecture. Power and Ground nets now need a standardized placeholder that can be used as golden reference across all design steps requiring such power intent specification.

For IP reuse purpose, it is important that the specification of the power network intent remains separated from the logic functionality specification as it can change from one implementation to another.

IV. METHOD FOR SIMULATING POWER MODES

Power and Ground nets are now part of the functional behavior and need to be simulated! Traditionally in digital design verification, only functional signals are simulated, not power and ground.

An increasing number of designs are implemented with multiple islands in which temporarily inactive blocks can be temporarily powered down without affecting the functionality of the rest of the design. Powering down a block can lead to propagation of unknown signals to the rest of the design. Such meta-stability can be prevented by proper insertion of isolation logic between the blocks. This interface logic needs to be verified. Beyond the necessary common placeholder to capture the power design intent, tool support is needed for the functional simulation of the power modes and the verification of the correctness of the clamping values. It is important that such approach allows exploration of alternatives functional partitioning to power islands without rewrite of the RTL code.

V. IP MODELING AND HIERARCHICAL USE MODEL

With the design specification now consisting of a {Power Intent, Functional Specification} pair, it is important to define a hierarchical precedence mechanism in

order to reuse existing IP specification or constrain an IP implementation. Several obvious use cases have to be accounted for:

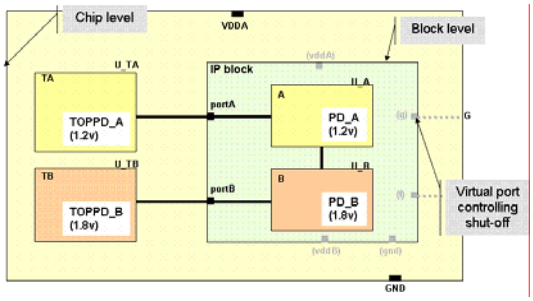

- Bottom-up reuse: Power design intent has been developed together with an IP. It should be reusable for the integration of this IP.
- Top-down constrain of lower level IP implementation: Chip level power design intent is created. Low level blocks should have their power design intent derived from this chip level description.
- IP implementation with visibility of the context of its instantiation: IP implementation is done with the knowledge of the power domains at its boundaries.

The requirements in nature have similarity with the timing constraints, timing models and timing budgeting.

VI. CONSISTENT POWER ARCHITECTURE SPECIFICATION

Designing with Voltage Islands is not entirely new. However, it is now becoming mainstream and where proprietary solutions were created, commercial tooling needs to substitute. Past SoC designs have clearly identified limitations associated with the design of Multi Supply Voltages SoCs:

- No placeholder for the power and ground nets and to describe power spec and constraints
- No possibility to verify power modes and power sequences in functional simulation.
- No reusability of IPs with multiple power domains into SoCs.
- Tremendous increase of implementation throughput time due to lack of automation.
- Recurrent specification of the same power intent for each tool in the design flow.

Beyond the discussion and controversy related to formats, NXP semiconductors has relentlessly been providing requirements and driving tool vendor‘s additional functionality required for designing with Voltage Islands. Tool vendors have now started to provide support for features built on a single description within their toolset offering. So far CPF (Common Power Format) has been qualified as a good basis for insuring consistency of the power specification, allowing portability across tools, being a golden reference for the power specification, allowing the potential specification of several “power architecture” to co-exist with the same RTL

A comparison with UPF has shown that, as they both came to the conclusion that RTL was not the placeholder for low power specification, a fundamental difference originally though to be a technical motivation for divergence has disappeared.

A larger scope towards the system level and additional capabilities allowing IP reuse with complex architecture are two obvious opportunities for these two industry standards to converge.