A Complete Design for Power Methodology and Flow for Large ASICs

[Extended Abstract]

Raymond X. Nijssen
Magma Design Automation
2 Results Way
Cupertino, CA 95014, USA
raymond@magma-da.com

Ed P. Huijbregts
Magma Design Automation
Den Dolech 2
5612 AZ Eindhoven, The Netherlands
ed@magma-da.com

Categories and Subject Descriptors

B.7.2 [**Integrated Circuits**]: Placement and routing; B.7.1 [**Types and Design Styles**]: VLSI; B.8.0 [**Performance and Reliability**]: General; J.6 [**Computer-Aided Engineering**]: Computer-aided design (CAD)

General Terms

Algorithms Design

ABSTRACT

Few aspects in ASIC design automation have become so pervasive as power. Design constraints due to power are being imposed throughout the entire design methodology and flow to control cost, reliability and performance of the products. With every future integrated circuit fabrication technology these constraints will be tightened further and new constraints are emerging. For large designs in particular this requires a new methodology and flow. We present how this trend impacts physical design, and we convey the point that the solution must be more automated, comprehensive and integrated.

1. INTRODUCTION

Until recently fabrication technology related techniques like voltage scaling and low-k insulation have proven their great effectiveness. For various reasons these techniques do not continue to scale nearly as well beyond today's fabrication technology. In addition, further miniaturization of feature sizes will increase leakage power so much that it may soon exceed power consumption due to switching.

Design power closure and circuit power integrity have been predicted to become one of the main drains of engineering resources and therefore total time to market. Especially in large designs, analysis and correction involving evaluation and intervention by the designer are becoming impracticable due to the scale of these tasks as well as the complexities of certain techniques unique to large designs.

Much further automation of these tasks is therefore warranted, and above all the paradigm of correctness by construction has become imperative. Existing algorithms need to be enhanced to satisfy the new and tighter constraints if possible, and new algorithms may be necessary.

Several of these methods and constraints pose new challenges in the area of physical design.

Yet, importantly, few can be viewed in isolation, most are intertwined with other steps in the flow outside physical design, and the importance of these efforts must be evaluated in their full context. No longer can a point-tool based solution suffice, the correct by construction techniques and methodologies require interaction with various algorithms, transformations and analyses throughout the flow. Completeness of the methodology in all of its aspects is a must.

In this extended abstract we summarize how techniques and methods referred to as design for power translate into constraints and objectives in the ASIC design flow with an emphasis on physical design.

There are many different considerations and issues related to power consumption in ASICs. In different application domains their importance can differ greatly.

Temperature Power dissipation due to switching and leakage increases device operating temperature. In large designs this requires expensive packages and external cooling. Operating temperature limits diminish performance.

Power Supply In battery powered, mobile applications low power consumption is a key feature. In environments where ample power is traditionally available the cost of power supply has become a competitive disadvantage.

Reliability Higher currents in wires increase the likelihood of circuit failure due to electro-migration.

Performance Wider variations in temperature and device supply voltage have lead to padding of device characteristics with margins widened to cover worst case scenarios at the expense of performance. For signal wire reliability high drive strength drivers must be avoided, which adversely affects circuit performance. Devices forced to operate at lower supply voltages are slower.

Note that power related effects can be divided into local and global effects. Each category has different properties

ISPD'03, April 6–9, 2003, Monterey, California, USA.

and remedies. These effects are however becoming more interrelated, and since in practice the distinction is therefore hard to make we will not try to make it here. Clearly as these effects are becoming first order more effort is spent on reducing their common cause, namely power consumption.

2. POWER REDUCTION TECHNIQUES

There are many techniques to reduce local and global power consumption of a design. Voltage scaling, enabled by new fabrication technologies, and architecture level design, eg. decide on clock frequency, functional architecture, switching activity, power/area/timing trade-offs, etc., are beyond the scope of this article, but are indirectly responsible for some of the following topics.

More recently the following techniques within this scope have emerged or are being used more aggressively:

Clock gating This well known technique reduces switching of sequential elements whose state does not change. It increases congestion due to more signal routing.

Multi-Vt High threshold transistors are slower but have less leakage than low threshold transistors which are faster. These properties are used by limiting cells with low threshold transistors to timing critical paths, thereby reducing overall leakage power.

Low-Power optimization Slowing down non-critical paths allows for smaller cells, reducing power consumption. Power consumption-wise cells perform best in a certain range of output load relative to their size. This in turn also requires careful management of output load.

Multi-Power To achieve high performance in critical parts of the design a subset of the devices may operate a higher voltage. In other application domains a large part of the design may be off most of the time until powered up by an always-on block.

3. POWER CONSTRAINTS

In addition to these constraints and restrictions aimed at power reduction, other objectives arise to counteract the effects of the higher overall power consumption and/or the tightening of existing rules.

3.1 Power Domains

We employ the concept of power domains. A power domain is a group of netlist objects sharing the state of their power supply voltage(s). It specifies the set of power nets are to be used in this domain. From these definitions it follows what voltages these objects can be at, from which it follow what cell characteristics to use.

Special attention is required where different power domains interact. Level shifters are typically used to convert signal levels and protect against sneak leakage paths. With great care level shifters can be avoided in some cases, but this will become less practicable on a wider scale.

Physical design hierarchy is used to separate different power domains, so power domain hierarchy forms a strict superset of the former. Physical hierarchy is a strict superset of the logical design hierarchy in all practical cases. Thus the power domain hierarchy is a strict superset of the logical hierarchy.

For electrical consistency multi-power designs requires strict adherence to a set of rules, for example,

- A cell can only be placed in an area associated with the cell's own power domain.
- A cell can only be driven by a cell in the same domain.
- Wires belonging to nets of different power domains must be spaced at more than minimum spacing at the top level.
- Nets must not be routed through a block belonging to a different power domain.

Such rules affect almost every aspect of the design flow. We review how these rules are relevant to physical design.

3.2 Floorplanning

Multi-power methods greatly constrain the floorplan. In most of today's methodologies devices at different voltages cannot be mingled freely in the same block. This is due in part to limitations in power routing rules and power routers. This leads to power domains dictating design partitioning.

More complicated and elaborate power supply meshes have to be implemented to support power domains.

IR-drop in large designs can no longer be kept within acceptable limits using peripheral power pads. The use of flip-chip pads to bring external power supply directly to the center of the chip has recently become mainstream.

3.3 Routing

Existing power routers will have to be enhanced to support the more complicated power routing requirements.

Vast amounts of routing resources are spent on wide power meshes at the expense of routing closure. In many technologies the upper metal layers are geared toward power distribution, typically by making the wires less resistive by making them both higher and wider. Due to these different characteristics they must be avoided by signal routing.

Wide via arrays from the upper layers for power distribution to the active devices on the lower layers can turn intermediate layers into labyrinths of obstacles complicating congestion calculation.

The wires in the power mesh need to be sized differently depending on calculated local power consumption. Wide power wires must be slotted or bamboo-d so as not to violate metal isodensity rules.

In signal wires limitations on maximum via current necessitate addition of redundant vias. These consume a large amounts of local routing resources because they block at least one adjacent routing track at at least one routing layer.

Temperature differences stress via connections, causing via stacking rules to be reinstated in new technologies.

3.4 Placement

Power domain rules require that cells be placed only in certain areas, at the expense of first order placement objectives of wirelength, congestion minimization and density control. Level shifter cells have special placement requirements so that they can be conntected to two power meshes.

Creation of local hotspots of high power consumption may have to be avoided by the placer, so placers need to be able to utilize power analysis output.

3.5 Optimization

Long wire buffering at the top level of the design already likens a feasibility problem more than an optimization problem. Power domains add additional constraints to the permissible locations per repeater. Otherwise optimal topolo-

gies may be illegal because the right power supply cannot be brought to certain locations on the chip.

High drive strength buffers needed to drive long wires may not be permissible due to reliability concerns, or the net driven is going through congested areas that cannot accommodate insertion of redundant vias.

IR-drop effects are normally accounted for by varying the cell characterizations. Cell delays vary dynamically due to local power supply fluctuations, causing both setup and hold violations. Presently spatial correlations are ignored, causing over-design of the circuit. Voltage scaling has eroded supply voltage variation margins.

4. CONCLUSION

In this summary we have made the point that power related aspects affect almost every part of the design methodology and flow. We focused on physical design aspects.

Existing implementations need to be enhanced or replaced, and new algorithms are needed. The many cross-dependencies clearly necessitate a comprehensive, flow-wide view requiring a highly integrated approach.

While these factors introduce new algorithmic challenges, at the same time they in many cases increasingly diminish the relevance of work ignoring these factors.