

Power Distribution in High-Performance Design

Michael Benoit
Simplex Solutions
Sunnyvale, CA
(408) 617-6106

Sandy Taylor
CMOS Solutions
Cupertino, CA
(408) 617-6175

David Overhauser
Simplex Solutions
Sunnyvale, CA
(408) 617-6135

Steffen Rochel
Simplex Solutions
Sunnyvale, CA
(408) 617-6138

1. ABSTRACT

Power distribution design in high-performance chips is a task that is not eased through the application of power reduction techniques. Although the average power of a high-performance design can be reduced, the peak to average power current ratio of blocks increases as a result, aggravating the challenges faced prior to average power reduction. This paper discusses the power distribution design challenge: to reliably deliver a predictable voltage to all transistors under all operating conditions. Steps in power estimation, approaches to power distribution implementation, and verification of power distribution are reviewed. The myths versus reality of power distribution design in high-performance chips are provided.

1.1 Keywords

Power grid, Power distribution, IR drop

2. INTRODUCTION

The power distribution design challenge in ICs is simply stated as to reliably deliver a predictable voltage to all transistors under all operating conditions. Reliable power distribution includes considerations of chip lifetime IR drop, and noise. The additional conflicting constraint is the minimization of area dedicated to power routing. Meeting the challenge and constraint is significantly more difficult in high-performance design, where high-performance is defined as either high frequency, high power, or low power design.

The problems faced by power distribution designers

continuously increase in complexity due to a combination of factors: increasing design size, shrinking supply voltage, and shrinking line widths. Successful chip design is described as obtaining a balance on a butte, where the higher the performance of the chip the smaller the area of success on top of the butte, illustrated in Figure 1. Microprocessor designers face the most difficult challenges, ASIC designers strive for more conservative design that yields a larger successful design space that can be achieved.

The goal of design is to implement a design that lands on the success plateau. As technologies evolve the butte is pushed in from the sides. On one side you have the process characteristics pushing to reduce the size of the success space. On the other side you have physical issues serving to reduce the success space. The physical issues include shrinking supply voltage, power, temperature, IR drop, noise, and reliability concerns.

Process variations, CAD tools, and methodologies used in design and verification have inherent uncertainty associated with them. This is described as an uncertainty in your final location on the success plateau, illustrated by the thick lines in Figure 1. As you move to higher performance design you must make changes in your tools and methodologies to reduce the uncertainty in your design. Reducing these uncertainties increases your performance predictability at the expense of design time. The minimum uncertainty achievable is that of process variations. If your uncertainty is greater than the size of the success plateau, your design has a probability of failing or delivering low performance yield.




Figure 1. Design success space reduces with process.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ISLPED98, Monterey, CA, USA
© 1998 ACM 1-58113-059-7/98/0008..\$5.00

Low performance or ASIC design methodologies seek to avoid the challenges of power distribution design through overly conservative design at the expense of chip area or routing efficiency. This strategy has worked well for designs at 0.5 μ m and above. However, characteristics of ultra-deep sub-micron designs are beginning to expose flaws in this methodology that high-performance designers have dealt with for years. The physical design problems eroding the successful design space impacts ASIC design as ASICs move into advanced technologies. As a result, ASIC methodologies will undergo some changes to increase their performance predictability as well.

The power distribution problem is placed in perspective by examining the power density of a chip. A 20 Watt chip that is 2cm on each side has a power density of 50kW/m². This power consumption is dominated by I/Os and clocks. Considering that most power is consumed around the edges of the clock indicates enormous power cycling, or dI/dt variations, on the power pins. This power cycling frequency must be examined in conjunction with package resonant frequencies to prevent resonance due to single or multiple cycle behaviors.

In this paper the challenges of high-performance power distribution design will be discussed. Section 3 discusses the challenges introduced by deep sub-micron design. Section 4 discusses power planning and estimation. Section 5 discusses power distribution implementation. Section 6 discusses power distribution verification. Section 7 discusses some common myths about power distribution design and verification.

3. CHALLENGES INTRODUCED BY DEEP SUBMICRON DESIGN

As process technologies continuously evolve, new characteristics of these technologies force changes in the design process. Technology scaling yields narrower wires, lower supply voltages, higher frequencies, larger designs, and increased total power dissipation. These are the factors that magnify the challenges in power distribution.

Lower supply voltages result in smaller noise margins in power distribution. This means IR and LdI/dt noise that was acceptable in the past now exceeds margins. As a result the power distribution system must reduce effective resistances and inductances from the pins to the transistor loads on the chip. The power supply LdI/dt must also be reduced for the power supplies although this is complicated by increasing design frequencies.

The latest technologies have increasingly narrow wires and higher resistance contacts and vias. The resulting increased wire resistance receives a great deal of attention for signal lines, but also has a high impact on power distribution design.

Reducing power grid resistance to reduce IR drop is in conflict with the trend to narrow wires. Thus, in order to compensate for the narrow wires, relative widths of power wires must increase, consuming a higher portion of the routing area. This yields a common trend today to have metal layers dedicated to power distribution. In addition to widening wires, contact clusters must increase in size because of contact resistance. In high frequency designs skin effects are serving to reduce the effective width of power wires as well. Thus effective wider wires must be achieved using more frequent narrow wires.

Increasing design sizes also impacts power distribution because the distance from the power pins to the transistors is increasing (along with the associated resistance). This requires either increasingly wider power wires or flip-chip packages that distribute power pins throughout the design area.

The net result of these deep sub-micron trends is to increase the challenge of power distribution design. To achieve reliable power distribution the designer is forced to loosen design constraints and also become an integral part of design decisions throughout the chip.

4. Power Planning and Estimation

Chip power distribution design generally takes place at two levels: global and block. Global grid design focuses on the top-level power distribution by concentrating on the average current required by the design. The goal in global grid design is to guarantee a specified quality of power to block connections throughout the chip. Global grid design also focuses on reliability because of the large currents transmitted [1]. Block power design focuses on distributing the peak block power currents throughout the global distribution system to prevent IR drop. If block power grid design cannot distribute the current load sufficiently across the global grid so that the global grid sees a somewhat uniformly distributed load, the global grid may require modifications to support the concentrated peak load.

Prior to emphasis on low-power design techniques, such as gated clocks, the ratio of peak block current to average block current was typically 2 for logic/control blocks. The inclusion of gated clocks can take this ratio to 10-20, and can even create higher peaks than prior to clock gate insertion. Therefore block power grids must be designed for the peak currents which may occur when the block is active. The block grid must distribute the peak current over many connections to the global grid in order to prevent global grid design from being tied to peak rather than average currents. If peaky block currents propagate to the global grid and pins, they tend to compound the LdI/dt noise of the package. Empty areas of a chip are therefore commonly filled with decoupling capacitances to filter

these peaks in supply current before they propagate to the package.

Global power distribution planning and estimation is performed in conjunction with floorplanning, package selection, and pin assignment. Initial insight into global power issues comes from experience with previous designs and experiments with new technologies. Block capacitance per unit area estimations are made by scaling a block from a prior design. Given capacitance estimations, power supply voltage, chip frequency, and relative activity scaling for various block types, estimates for various blocks' peak and average current consumption are made. These estimates are made for dynamic logic blocks, static logic blocks, and other special blocks such as RAMs, ROMs, and I/Os.

These power estimates are used to design the global power distribution. The estimates may be off from the eventual actual power requirements of the completed blocks, but in a planning stage the best estimates are used. Of particular interest at this stage is the location of critical blocks, such as clocks and caches so adequate design of low resistance power routes to these blocks can be performed.

Power grid planning does not continue into great depth because all estimations performed early are the best guess as to power current requirements and additional planning introduces little value. Additional detail for blocks is not available until blocks are designed. But by that time the global grid implementation must be nearly complete. Conservative implementation is the key once early estimation and planning is performed. If the power distribution system were "pure", meaning no modification would occur late in the design process and power estimates were accurate, then additional early planning and more aggressive design could take place. However, since many local modifications to the grid are made late in the design cycle in order to assist in signal routing, aggressive approaches to power grid design are uncommon.

5. Power Distribution Implementation

Given the chip floorplan and block power estimates a conservative power distribution strategy is constructed. Power distribution in high-performance designs is split into global power routing and block power routing levels. Global power routing provides distribution of average current loading throughout the chip as well as additional routing for those portions of the design with special requirements. Block power routing provides distribution of local power consumption to the global grid in such a way that peak block current loads appear as distributed average loads at the interface between the block and the global grid.

Global routing includes meshed grids on several routing layers. This meshed grid requires additional routing space, and perhaps dedicated routing layers, but provides quality

power throughout the chip, distributes current among pins, and prevents power distribution problems from arising later in the design cycle. The global grid may include interdigitated fingers for two purposes: to help uncorrelate block currents at the global grid interface and to increase coupling capacitances between power nets. Coupling on power nets is maximized to provide the maximum level of local filtering of power consumption spikes.

Global power distribution mesh parameters are specified in terms of width and pitch, such as 25 μ m wide routes at 150 μ m spacing. Specific width and pitch are functions of the mesh size, chip power consumption, IR drop margins, and chip frequency (skin effects).

Block power distribution is handled in one of two ways. In regular arrayed structures the power grid is an integral portion of the array design. This insures adequate power distribution within the array. Driver locations within the array, such as in a datapath, include additional power bussing to adequately handle expected power peaks. The power distribution in these regular arrays are characterized carefully because their normal and worst-case conditions are understood very well by the designer.

The second style of power distribution design is for control logic blocks. These blocks present the more difficult challenge because they are typically implemented using automatic design tools and therefore their behavior is less predictable. Block power grid design tools must have some understanding of power intensive regions or else they are likely to undersign the power routing to critical areas.

Hard IP cores, in which block layouts are provided to customers, include conservative power grids to provide robustness in the connection to global power grids. However, since they have no knowledge of their environment when they are designed, they still require adequate verification after insertion into the chip. Hard IP cores should never be used to distribute power.

One of the most common and risky situations that occurs regularly in the design of high-performance chips is the removal of small global routing wire segments in order to complete the routing of block signals. Expecting this is why global power grid designers will over design their power distribution and also why global meshes are preferred. The cutting of power grid segments introduces risk to the power grid resistance goals from the transistors to the pins, that results in additional need for verification.

6. Power Distribution Verification

Power distribution verification has grown significantly in its application in the past several years due to the increasing number of design failures caused by inadequate power distribution implementation. Power distribution problems have even plagued some ASIC designs.

Traditional design relied on conservative practices to prevent power distribution problems, and verification may have been as simple as visual inspection. Because deep sub-micron is physical, designs cannot afford to be so conservative in design or casual in verification.

Designers must know not only how much power is consumed in a design, but where it is consumed. Place and route tools commonly used today have no support for power distribution analysis or reliability analysis. Therefore manual intervention may be required to resolve problems in power distribution. Methodologies that separate the logical design from the physical implementation can be dangerous in this respect because designers have no understanding where functionality is located after a block has been automatically designed.

Four factors during the design stage can introduce power distribution problems in conservatively designed grids. First, process modifications can change the characteristics of a grid, increasing resistances. Second, signal routing complexity today generally results in segments of the power grid being sacrificed in order to achieve routing completion. Third, partitioning of power distribution into global and block grids commonly results in errors in the attachment of the grids. Fourth, final placement of large drivers may be in areas with insufficient power routing. All these factors impacting the grid require verification.

Adequate discussion of power distribution verification requires more space than available here [2], so only basics are discussed. A variety of aspects of power distribution design can be verified using commercial CAD tools today. The approaches available provide a range of tradeoffs between verification time and vector-dependent detail.

The most common and critical errors in power grid implementation can be found using static analysis of grids. Static power grid analysis does not require simulation vectors to find problems, but rather assumes some form of activity distribution across a design to qualitatively analyze the power grid. This form of analysis finds significant IR drop problems, unexpected current paths, and visibility into electromigration problems. An activity distribution of the transistor loads on a power grid can be obtained in a variety of ways: assuming all transistors are ON and scaled, using chip frequency and gate switching density to compute average currents, or simulation using vectors to compute average or peak currents.

Although static analysis is not guaranteed to find the worst-case currents and IR drop on a power grid, experience has shown that most problem areas in a worst-case sense exhibited their effect in static analysis as well. Proper application of static analysis can be extremely effective in verifying power distribution. Many examples of problems found using static power grid analysis may be found in [1] and [2].

Dynamic power grid verification uses simulation vectors to activate various portions of a circuit in order to compute instantaneous power grid behavior. The concern in this case is twofold. First, instantaneous IR drop and Ldi/dt noise happens in a small portion of a clock cycle, so dynamic analysis to detect these conditions must have greater resolution than a single clock cycle. These hot spots caused by blind placement algorithms (focused on timing only) are common. Second, finding worst-case vectors for IR drop can be very difficult.

Worst-case IR drop vectors can be independent of worst-case vectors for total power consumption. IR drop is a local power distribution phenomenon. The greatest IR drop on a chip can, however be related to certain predictable situations. Clock switching and synchronous bus driving are two specific situations that create strong localized power consumption. Therefore test vectors for these specific conditions can be generated. Otherwise, mechanisms to construct a worst-case IR drop test vector using data gathered from many test vectors should be applied [3].

Proper and thorough power distribution verification requires a combination of static and dynamic techniques in order to find grid weak spots and vector-dependent failure mechanisms.

7. Power Distribution Design Myths

Many designers have a variety of beliefs about power distribution on chips that are not necessarily true. This is due to the fact that in most low-performance designs power distribution is not a common source of design failures. Design decisions made based on these myths may result in the failure of a design.

7.1.1 Power integrity problems are easy to find on silicon

Power integrity problems have a variety of symptoms on silicon. Some of those symptoms associated with total logic functionality failure are easy to identify. However, in many cases the symptom of a power integrity problem is a timing violation or logical error. These symptoms can therefore point a designer down the wrong path to a solution.

7.1.2 I don't need to worry about power integrity in the design of my block

One class of power grid integrity problems is result of how gates are assembled in a block, particularly data paths. The orientation of power routes at the lowest level and the distribution of large drivers on power rails are critical. A number of large drivers activating synchronously on a single power wire can result in large IR drops and subsequent timing impacts.

7.1.3 Guard banding my block is sufficient isolation of a block from global distribution issues

The design of a block's power distribution can have a large impact on global distribution issues, even if guard banding is used. For example, if a block contains many power rails crossing the block, those power rails can serve as a very effective power feed-through. A number of designs have experienced power problems in blocks because the block power distribution system unintentionally focused current through the block rather than around the block in guard bands. In addition, unnecessary guard banding increases the chip area, and consequently reduces profit.

7.1.4 Block power grid verification is sufficient in a design

This is a dangerous myth. The primary source of power distribution problems in a chip result from interactions between blocks on the chip or from inadequate connection of a block to the global grid. Adjacent blocks commonly interact to introduce IR drop or modify current flow to inject reliability problems. These issues must be verified when the chip is assembled.

7.1.5 IR drop verification is sufficient in a design

IR drop is only one of the primary sources of power grid integrity problems. Electromigration and Ldi/dt are other sources of design failure. All these issues must be included in power grid verification.

7.1.6 LVS is sufficient in verifying connectivity between my block and the global distribution

In reality insufficient connection between a block and the global power grid has been a common cause of power distribution implementation problems. A number of designs have been observed in which either a place-and-route tool has performed inadequate connectivity or via insertion scripts missed some vital power tap points. These types of failures are not found using an LVS tool because LVS only requires a single connection someplace whereas common failures occur because of insufficient connectivity.

7.1.7 Once a conservative global power distribution system is designed, I don't need to worry about it

This myth is very dangerous. Global power distribution is commonly designed in an early phase of a project before the detailed routing of signals takes place. Signal routing is a very difficult task to complete. It is therefore very common for segments of the power distribution to be sacrificed in order to complete signal routing. If the impact of this change to the power grid is not verified, excessive IR drops may occur. Traditionally, high-performance designers had no systematic method for verifying the power distribution after all routing was complete, so little

or no verification was performed. Experience with these problems has resulted in more extensive verification in designs today.

7.1.8 Flip-chip prevents power grid problems

While flip-chip packages aid in the reduction of resistance from various locations in the design to power pins, it does not resolve all problems. Flip-chip technologies have two impacts. First, their application generally results in thinner power routing in the chip, so verification is still required in order to verify proper power distribution without consuming too much routing area. Second, flip-chip introduces some thermal issues in the design due to the high currents accumulating at power pads.

7.1.9 The physical power grid implementation matches the plan

At the time the power grid plan is completed, many issues about the physical implementation are not completely resolved. Any excessive discrepancy in the physical implementation can invalidate the plan. For example: were block power estimates close enough to actual? More importantly, were process characteristics close enough to predicted? Issues such as contact resistances increasing by a few Ohms can invalidate many resistance estimations from pins to loads.

8. Summary

High-performance power distribution design issues apply to high-power, low-power, and high-frequency design styles. Low-power design techniques can actually increase peak block currents. Global and local power distribution techniques must work together to address both average and peak power conditions to reliably deliver predictable voltage to all transistors on a chip. An overview of deep sub-micron design challenges, power planning, implementation, and verification have been discussed. In addition some common myths about power distribution design have been presented.

9. REFERENCES

- [1] S. Rochel, G. Steele, J. R. Lloyd, S. Z. Hussain, and D. Overhauser, "Full-chip reliability analysis," *Proceedings, 1998 Intl. Reliability Physics Symp.*, March 1998.
- [2] G. Steele, D. Overhauser, S. Rochel, and S. Z. Hussain, "Full-chip verification methods for DSM power distribution systems," *Proceedings, 1998 DAC*, June, 1998.
- [3] *Thunder Manual*, Simplex Solutions, 1998.