

POWER GRID VERIFICATION

TABLE OF CONTENTS

Introduction	1
Power grid IR drop and ground bounce	2
Static power grid analysis	2
Dynamic power grid analysis	3
Electromigration	3
Issues in the design of power distribution systems	4
Improving full-chip power integrity and reliability	8
Reducing voltage drop	8
Reducing electromigration problems	9
Summary	11

TABLE OF FIGURES

Figure 1	Routing through a block	4
Figure 2	Routing around the blocks	4
Figure 3	Vias in a mesh array methodology	5
Figure 4	Electromigration in the power grid	6
Figure 5	Electromigration in via arrays	7
Figure 6	Electromigration risk at the full-chip level	9
Figure 7	Power grid before changes	9
Figure 8	Power grid after changes	10
Figure 9	New electromigration problems in unexpected regions after changes	10

INTRODUCTION

IC power distribution systems are designed to provide needed voltages and currents to the transistors that perform the logic functions of a chip. Based on a survey of over 206 tapeouts, targeting process technology of 0.13 micron or greater, more than 50% of tapeouts will fail if the power distribution system is not validated beforehand. At and below 0.13 micron technologies, IC designers can no longer take the risk of assuming that their VDD and VSS grids have been designed correctly, and they must perform detailed analysis to understand how robust their power distribution methodology really is.

The performance of a design's power networks, or grid, has a direct impact on its performance. Voltage (IR) drops on VDD nets and ground bounce on VSS nets affect a design's overall timing and functionality, and if ignored, will cause silicon failure. High currents in the power grids also induce electromigration (EM) effects, where the metal lines of a power grid begin to wear out during a chip's lifetime. These effects cause expensive field failures and major product liability issues.

While a clean power grid design should always be the goal of any design team, timing analysis should be performed with the effects of IR drop and ground bounce included once the power grid is routed—it is always possible that a designer can simply live with the amount of IR drop if it does not cause timing failure.

A complete picture of power grid robustness can only be obtained when effects such as IR drop, ground bounce, and EM are accurately computed and analyzed. These are full-chip issues that must be addressed by verification tools that have the capacity and performance required to analyze detailed representations of the chip in a reasonable amount of time. By analyzing and verifying the power grids at the full-chip level, designs can be taped-out with an increased expectation of first silicon success. In this white paper, the key issues of power grid design—IR drop, ground bounce, and EM—are described, along with the related analysis issues. Methodologies used to identify IR drop violations and potential EM violations in the power grids are also described, along with approaches that reduce the severity of these problems. Designing with these issues in mind and performing full-chip power grid verification enables designers to address what would otherwise be an intractable problem.

POWER GRID IR DROP AND GROUND BOUNCE

IR drop on a VDD grid is caused because the current demanded by the transistors or gates of a design flows from the VDD I/O pins (or bump bonds in the case of flip-chips) through the RC network of the power grid, and leads to a decreased VDD voltage at the devices. Ground bounce is a similar phenomenon, where the current flows back to the VSS pins, and the RC network causes the VSS voltage at the devices to rise. The risk of a design suffering from IR drop or ground bounce increases with shrinking process technology and with next-generation designs. With every technology shrink, the current demand per unit area of a design is increased, basically because of shrinks in the gate oxide thickness. This, combined with the fact that next-generation designs contain more transistors or gates, adds additional stress to power grid design and typically results in a power grid that contains an increasing number of parasitic RC values to be analyzed. For example, a leading-edge 90nm design containing over 10 million gates and processed with eight layers of metal could produce a VDD grid approaching 1 billion RC elements.

There are two approaches typically used for power grid analysis, static and dynamic. A static analysis solves Ohm's and Kirchoff's laws for a given power network but ignores localized switching effects on the power grid. A dynamic approach performs comprehensive dynamic circuit simulation of the power grid network, which includes localized switching effects. Both approaches have unique value and challenges.

STATIC POWER GRID ANALYSIS

The static power grid analysis approach was created to provide comprehensive coverage without the requirement of extensive circuit simulations. Typically, most static approaches are based on similar concepts:

1. The parasitic resistance of the power grid is extracted
2. A resistor matrix of the power grid is built
3. An average current for each transistor or gate connected to the power grid is calculated
4. The average currents are distributed around the resistance matrix, based on the physical location of the transistor or gate
5. At every VDD I/O pin, a source of VDD is applied to the matrix
6. A static matrix solve is then used to calculate the currents and IR drops throughout the resistance matrix

A static approach approximates the effects of dynamic switching on the power grid by making the assumption that de-coupling capacitances between VDD and VSS smooth out the dynamic peaks of IR drop or ground bounce.

The main value of the static approach is its simplicity and comprehensive coverage. Since only parasitic resistance of the power grid is required the extraction task is minimized, and since every transistor or gate provides an average loading to the power grid the solution provides comprehensive coverage of the power grid.

The main challenge of the static approach is accuracy. Local dynamic effects are not accounted for, neither are package inductance effects (Ldi/dt), both of which may result in optimistic IR drop or ground bounce results if there is insufficient de-coupling capacitance on the power grid.

DYNAMIC POWER GRID ANALYSIS

A dynamic power grid analysis requires that both resistance and capacitance of the power grid are extracted, and that a dynamic circuit simulation of the resultant RC matrix is completed. Typically, the steps to complete a dynamic power grid analysis are:

1. The parasitic resistance and capacitance of the power grid is extracted
2. The parasitic resistance and capacitance of the signal nets is extracted
3. The design netlist is extracted
4. A circuit netlist is created from the extracted parasitics and netlists
5. A circuit simulation is executed, based on a suite of simulation vectors, which simulates the transistors or gates dynamically switching and the effect of this switching on the power grid

The main value of the dynamic approach is its accuracy. Since the results are based on circuit simulation, the IR drop and ground bounce results can be extremely accurate and take into account localized dynamic and package inductance effects.

The challenges of the dynamic approach are significant.

- The parasitic extraction demands are high because you need to extract resistance and capacitance for the power grids and (as a minimum) the capacitance for the signal nets.
- The circuit simulation can contain a huge number of elements to be simulated, which strains the capacity of the circuit simulation engine.
- The vector set that is used to stimulate the simulation plays a dominant role in determining the quality of the output, if a comprehensive suite of vectors is not used, then the results will be questionable because sections of the power grid may not have been simulated.
- Finally, given the number of elements associated with a single power grid, a power grid analysis solution based on comprehensive dynamic simulation will not easily scale as design sizes continue to grow.

Many power grid analysis solutions that promote a dynamic approach must often resort to RC reduction techniques to manage the size of the data to be simulated; however, this directly conflicts with the main value of the dynamic approach, the accuracy of the results. RC reduction of the power grid can cause inaccuracies to creep into the analysis, and can hide real EM problems.

ELECTROMIGRATION

Electromigration is another important issue in the design of deep submicron power grids. High current densities and narrow line widths cause EM, and failures due to EM can be catastrophic. Failure typically occurs in a customer's hands, when the chip is already installed on a board in a system, which may result in a design recall.

While EM can cause both open and short circuits in the power grid, the most common effect is to cause increased resistance in power grid paths, which leads to increased IR drop or ground bounce and impacts chip timing. This results in a design that originally worked at its specification, but fails to operate to specification some time later in its life.

ISSUES IN THE DESIGN OF POWER DISTRIBUTION SYSTEMS

Designing a power distribution system requires the consideration of both electromigration (EM) and IR drop in a full-chip context. For example, consider the two blocks below in Figure 1. If power distribution for Block A is examined in isolation, the additional loading due to the presence of Block B is not taken into account. If power is routed through Block A to Block B, a larger IR drop will occur in Block B since power is also being consumed by Block A before it reaches Block B. As more and more blocks are added, the complex interactions between the blocks determine the actual IR drops.

The placement of these blocks is typically based on the timing requirements of a system rather than on IR drop, or else placement is based on the size and shape of blocks at the floorplanning stage. Therefore, sizing the buses properly to minimize IR drop while satisfying the required timing and area constraints is a design challenge that can only be met using full-chip analysis.

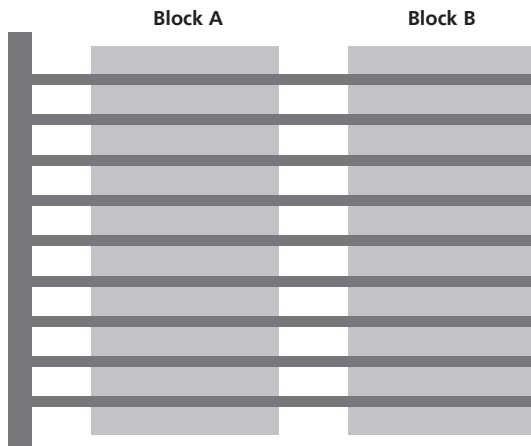


Figure 1: Routing through a block

Since the total IR drop is based on the resistance seen from the pin to the block, one could route around the block and feed power to each block separately, as shown in Figure 2. Ideally, the main trunks should be large enough to handle all the current flowing through separate branches. In this case, the T-junctions have a high current density and may be prone to EM problems. It is important in this type of grid to examine the current density at all junctions, especially the corner providing large amounts of current to each block, to ensure that EM problems do not exist. The same argument holds for every block routed in this manner. Again, power grid design is truly full-chip when IR drop and EM issues are considered.




Figure 2: Routing around the blocks

Although routing power this way is easier to control and maintain, it also requires more area to implement. The large metal trunks of power have to be sized to handle all the current for each block. This requirement forces designers to set aside area for power busing that takes away from the available routing area.

Another approach to minimizing IR drop, depicted in Figure 3, is to have a solid grid of Metal 4 and Metal 5 and use a via array to connect the two layers, effectively tying the whole grid to VDD. While this solves the problem at higher levels, it simply shifts the problem down to the lower levels of metal. What about Metal 3 and Metal 2? Are they wide enough to handle the current levels they will sustain in terms of IR drop and EM?

Depending on the methodology, lower levels are often left floating until final assembly. Low resistance, high current paths can often be created by random placement of lower blocks. In fact, when you design the logic circuitry in the block, it is not clear where Metal 3 will tap to Metal 4, so you cannot predict the current flow. And if you cannot predict it, you must analyze it.




Figure 3: Vias in a mesh array methodology

Part of the grid may have to be removed to route some signals, as shown in Figure 3. Which straps can be removed without introducing problems? If you arbitrarily pick one that is conducting a large amount of current, the excess current must flow in adjacent straps which may push the current density in them beyond acceptable levels. Clearly, such decisions cannot be made without determining the current levels in the straps and then picking ones that have lower current levels. The complexity of the problem requires a set of power grid analysis tools. These examples illustrate that design decisions must be made with a global perspective in mind.

IR drop is a dynamic phenomenon due primarily to simultaneous switching events in a chip such as clocks, bus drivers, and memory decoder drivers. As large drivers begin to switch, the simultaneous demand for current from the power grid stresses the grid. In a static context, IR drops are highest near the center of a design and lowest near VDD connections to the power supply. However, during dynamic operation, these simultaneous switching events can cause severe IR drops anywhere on the chip, and these are the ones that must be identified. These events, usually well known, can be triggered with typically fewer than 100 vectors.

The effect of IR drop on chip performance is significant. IR drop compromises the voltage noise margins of logic gates, due not only to IR drops in the power grid during the rising edge of a signal, but also to the increase in voltage in the ground grid because of the same phenomenon during the falling edge. Once the noise margins drop below the budgeted amount, typically 10%, the design is not guaranteed to operate properly.

Over the years, supply voltage has been shrinking as device dimensions are scaled to avoid transistor punch-through conditions, hot-electron effects, and device breakdown. This has resulted in smaller and smaller noise margins. With IR drop, the margins are reduced even further which makes it even more difficult to manage a multimillion-transistor design.

IR drop on a power grid primarily affects timing. IR drop compromises the drive capability of the gates and increases the overall delay. Typically, a 5% drop in supply voltage can affect delay by 15% or more. Delay in a clock buffer has been known to increase by more than 100% due to IR drop. Such an increase in delay is critical when you are managing clock skews in the range of 100 picoseconds. Imagine the effect of this type of unexpected delay along centrally located critical paths. Then path delay is no longer predictable and, in fact, the critical path may be somewhere else in the design due to IR drop. This means that the performance or functionality of the design is unpredictable. Ideally, timing calculations should take worst-case IR drop into account to improve accuracy.

In Figure 4 below, a portion of a design is shown with two metal lines connected by a narrow strap of metal. The metal lines must be wide enough to carry the average current needed to feed the circuitry connected to it. If the lines are too narrow, EM or IR drop may occur.

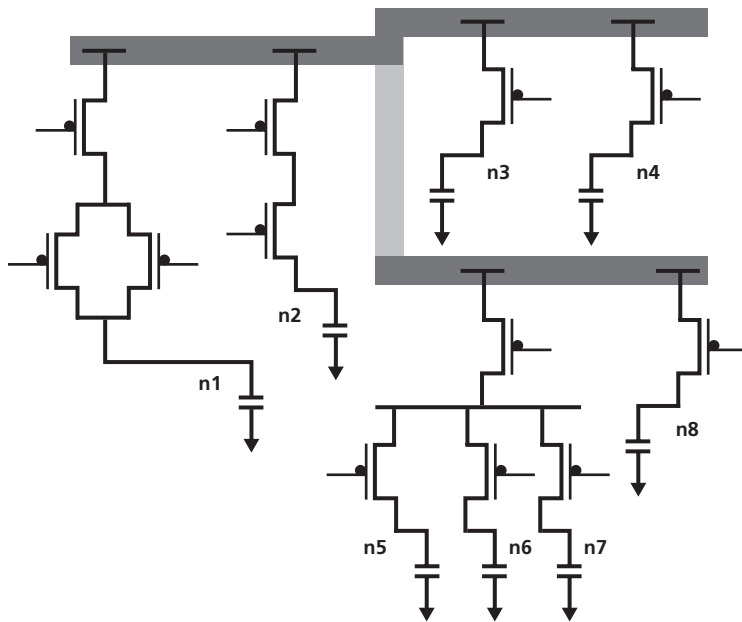


Figure 4: Electromigration in the power grid

Since large currents flow in the periphery of a design, EM problems are usually observed in the outer regions of a chip. However, vias scattered all over the design may also be prone to EM problems. Furthermore, the lower levels of metal connected to devices are usually narrower and may cause EM problems depending on the current levels. Therefore, it is important to look for EM across the entire chip rather than just specific regions.

Finding all areas susceptible to EM prohibits any use of data reduction. You must include all the detailed extracted resistance data—otherwise, you may lose useful information. For example, a via cluster that has been reduced to one via resistor may mask a potential electromigration failure, and an EM analysis tool would miss the problem.

In Figure 5 below, current flows from Metal 5 to Metal 4 through a via array. Crowding occurs as the current "hugs the curve" going from one level to the other. Some of the vias in the center of the layout have been tagged as ones that may suffer from electromigration (EM). If the 16 vias in the array were collapsed into one via, this region would not be flagged as having a problem. In reality, the nine indicated vias in the cluster may fail due to the high current density in the narrow dimension of the cluster. Any extraction and analysis for EM must have unreduced data to provide useful feedback.

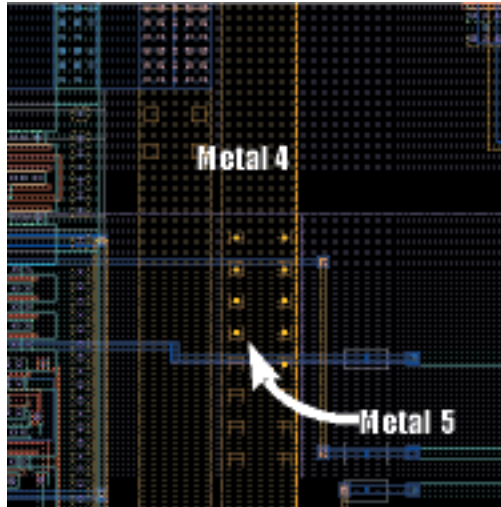


Figure 5: Electromigration in via arrays

Electromigration in the power grid is a DC phenomenon due to the average current flow in metal lines and vias. Design guidelines for EM are based on average current levels which, in turn, depend on signal line capacitance. Therefore, obtaining an accurate EM prediction requires the use of accurate capacitance information. Furthermore, since metal lines vary in height and material properties at different levels in the design, each metal layer has different failure criteria. To identify all potential areas of EM problems across a chip, the only solution is to perform full-chip analysis.

Black's law is used to predict the mean-time-to-failure (MTTF) of a metal line using the average current density, J , seen by the line. The more accurate the average information, the better the estimate of the MTTF. To obtain this information, you need to use a large number of vectors to exercise the design. The average current in every metal line must be measured and then divided by the width and thickness of the line. This is clearly impossible to do on a fabricated chip, and prohibitive to do using circuit simulation.

An alternative to expensive transistor-level simulation is to obtain average currents from activity information, in the form of toggle data, using a gate-level or higher-level tool. Toggle data is simply the number of times a gate switches high or low during a simulation of thousands of clock cycles. If the toggle data is divided by the number of clock cycles, the activity information is obtained. For example, the core of a memory circuit may have an activity of 0.02% while a data path may be closer to 5%. These factors can be converted into average current information for the transistors connected to the power grid.

You must also determine the average flow of current in the entire power grid to assess reliability risks of a given design. It is not sufficient to determine the average behavior of a block taken in isolation, because the block may only be exercised periodically in a full-chip context. Furthermore, changes to the power grid in one section tend to have a global impact. Data reduction cannot be used either since some of the real EM problems may be masked by the reduction itself. Therefore, an accurate picture of EM risk cannot be obtained unless the entire chip is verified as a single entity. Any tool used for this purpose must have the capacity to analyze multi-million resistor grids.

IMPROVING FULL-CHIP POWER INTEGRITY AND RELIABILITY

The problems described above must be identified and fixed before going to silicon since they are very expensive to debug after fabrication. Verification tools exist for this purpose. Clearly iterations through a verification loop are preferable to more expensive iterations through a fab-find-and-fix loop.

When you look at methods of performing full-chip verification, it is clear that an engineering solution must be developed. A dynamic simulation with power grids containing many millions of resistors and a similar number of capacitors, is prohibitively expensive. Any simulation approach will suffer from severe capacity limits unless severe data reduction techniques are utilized, which minimizes the overall value of this approach.

As mentioned above, block-based methods by themselves do not suffice since power distribution planning is a full-chip issue. Given the scope, trying to find the IR drops and EM risks is certainly a daunting problem. But if power grid verification is the goal, excellent approaches exist to address the problem.

REDUCING IR DROP

As described earlier, IR drops in the power grid can be caused by two different type of phenomena—IR and Ldi/dt. Reducing the impact of IR drop in a power distribution system can be accomplished in several ways:

WIDEN THE POWER ROUTES

The simplest approach is to widen the lines that experience the largest voltage drops since increasing the width decreases the resistance (hence the IR drop). However, this may not always be possible due to constraints in the routing area. Via arrays should also be maximized wherever possible, since the resistance associated with a single via can have a significant impact on IR drop.

MAXIMIZE THE USE OF DECOUPLING CAPACITANCE

One effective approach is to use decoupling capacitors between power and ground, which can deliver the additional current needed by the power distribution system. These decoupling caps are usually scattered throughout the power grid, in any available space, and enable the use of the static approach to power grid analysis. Ldi/dt effects can be mitigated by placing large capacitances near the pins.

STAGGER THE SWITCHING

Since IR drop is due primarily to simultaneous switching events, another (more difficult) approach is to stagger the gates that are switching together such that they switch at slightly different times—at least enough to keep the problem within the noise budget. Alternatively, you could reduce the buffer size, but this may not be possible if the design fails to meet performance requirements with smaller devices. Device switching can be staggered to reduce the peak demands of current by introducing delays on the signals driving the gates.

MAXIMIZE THE POWER I/O PINS

A more aggressive solution is to use a ball-grid array, sometimes called solder bumps or C4 bumps, where the power supply connections can be at various points within the chip. This expensive solution requires placing many C4 bumps across the chip to minimize the worst-case IR drop in any location. This solution tends to push EM problems to lower levels of metal that are usually narrower. Also, this solution cannot be used in sensitive areas such as memories and dynamic logic because C4 bumps generate alpha particles that may cause logic value upsets in the sensitive nodes. Nevertheless, when used appropriately, C4 bumps can reduce IR drop. The key to design is proper placement of the C4 connections, which can only be done effectively with full-chip analysis.

REDUCING ELECTROMIGRATION PROBLEMS

Electromigration failures can be reduced in several ways. The basic idea in all approaches is to reduce the average current density seen by any metal segment. The simplest approach is to widen the metal lines. However, increasing the width beyond a certain point leads to over-design, which costs area and can reduce yields. Another approach is to change the current flow in the power grid itself by adding jumpers and straps between different points in the grid. This would reroute current around the affected areas, but such changes would require another verification pass to confirm that the problem has not simply been moved to another area of the design.

In Figure 6, note that the standard cell block on the right would not have shown any EM risk if analyzed by itself. However, in a full-chip context, current flowing to adjacent blocks overloads the power connections in the block, and the analysis tool identifies an EM risk. Recognizing these problems at the planning stage is helpful, but difficult to do. EM requires a detailed grid with unreduced data. Therefore, a complete picture of EM risk can only be obtained at the verification stage.




Figure 6: Electromigration risk at the full-chip level.

A key point made earlier is that IR drop and EM problems cannot be solved separately; they must both be considered during design. To illustrate this, consider how to solve an IR drop problem in the chip in Figure 7. The figure shows a power flow diagram of the VDD grid in a multimedia chip. Different shading indicates various levels of IR drops. The darkest areas are the lowest points (valleys) of the IR drop contours. A significant IR drop occurs in the center region of the chip because only the top portion of the power grid feeds the large drivers in the top section. The upper and lower regions of the power system are not connected.




Figure 7: Power grid before changes

If we strap the upper and lower regions together in two places, the IR drop problem is reduced significantly, as indicated in Figure 8. The depth of the IR drop valleys has been reduced to acceptable levels, and the IR drops have been spread over a wider area of the grid. The lower region is now supplying more current to the upper region and therefore a better power distribution has been obtained by adding the two straps.




Figure 8: Power grid after changes

However, when examined in the context of EM, the results show that fixing the IR drop problem has caused an EM problem in the lower portion of the design. A review of Figure 6 (before the straps were added) shows EM problems at the periphery of the chip due to the high current levels in those regions. The lower half of the chip shows no EM problems.

But in Figure 9 (after straps were added), new EM problems are evident in the lower half as indicated by the small horizontal white lines. It was clear that the lower portion would supply additional current to the upper half of the design once a bridge was built between the two. However, it was not clear exactly how current would flow and exactly where EM problems might occur.




Figure 9: New electromigration problems in unexpected regions after changes

Repairing all the areas with potential EM problems would be labor-intensive, time-consuming and, frankly, unnecessary. Since every chip has a lifetime associated with it, the MTTF factor can be used to compute a probability of failure due to EM in a given lifetime. The goal of any changes to the power grid would be to decrease the probability of failure to an acceptable level. This limits the actual number of repairs needed and makes the job manageable.

Does increasing a line width always improve EM risk? No — thin wires can have better EM characteristics than wider wires due to the physics of EM. Be aware that more is not necessarily better. Proper EM analysis accounts for this width dependence.

SUMMARY

The design of power distribution systems for all ICs is complicated by issues such as IR drop, ground bounce, and EM. In the distant past, DRC, LVS and hand calculations were performed on the power grid to ensure clean power grid designs, and over-designing the power grid was considered an acceptable solution. In today's highly competitive market, the area penalty of over-designing leads to decreased yields and non-competitive designs, and the penalty of under-design is tapeout failure, silicon respins and costly field failures—you lose in either case.

Power grid analysis is now a critical part of design verification prior to tapeout, because Murphy's Law can be directly applied to power grid design: if something can go wrong, it probably will.



Cadence Design Systems, Inc.

Corporate Headquarters
555 River Oaks Parkway
San Jose, CA 95134
800.746.6223
408.943.1234
www.cadence.com