

Electromigration wreaks havoc on IC design

JIM LLOYD AND DAVID OVERHAUSER, SIMPLEX SOLUTIONS INC

Current IC-design practices make perfect conditions for electromigration, which causes broken connections. A thorough understanding of the problem and its prevention help prolong an IC's life.

When electrons flow through wires on a chip, they collide with metal atoms, producing a force on the atoms that causes the wires to break (see **box** “The basics of electromigration”). The immediate fix for this problem is to make the metal wires thicker. The conditions that cause electromigration (EM)—long, narrow, aluminum wires; high current densities; logic hazards; and high operating frequencies—are becoming worse, and they occur on both power grids and signal lines.

EM can only get worse as wires become narrower and operating frequencies increase. Until new metallization schemes erase EM, designers need a thorough understanding of why EM occurs and how it affects full-chip power and signal routing. Designers also need to be aware of EM's causes and how to manage EM using modern verification tools. More than ever, IC designers need tools that can find and help fix EM problems before they become problems in silicon. Only then can designers ensure reliability while squeezing as much performance as possible from fabrication.

How EM happens


Resistance does not exist in a perfect crystal lattice; electrons move without interaction with the metal atoms. However, a perfect lattice cannot exist above absolute zero because of missing atoms (“vacancies”), impurities, boundaries between crystals of different orientations (“grain boundaries”), and other regions of imperfection. The interaction between electrons and crystals results in momentum exchange. The magnitude of the momentum exchange is proportional to the current density and to the distortion in the lattice at any point. The distortion is greatest when a vacancy is nearby or in the region of a metal grain boundary. These areas are also where diffusion resulting in EM occurs.

When current passes through a conductor, the interaction of the electrons with the lattice produces thermal energy. This energy, “Joule heating,” can increase the temperature of a wire enough to cause EM. If the current is low, the heat effectively conducts away, but the temperature must still increase even if the increase is undetectable. If the current density approaches $10^6\text{A}/\text{cm}^2$, Joule heating can produce enough energy to make the conductor lines heat appreciably. This heating does not initially appear to be a problem; current densities are almost always lower than this number because of EM-induced limitations.

THE BASICS OF ELECTROMIGRATION

Electromigration (EM), the mass transportation of a metal resulting from the momentum transfer between conducting electrons and diffusing metal atoms, exists wherever current flows through metal wires. Discovered more than 100 years ago, EM first showed up in ICs as early as 1966 and has been a persistent problem since the early days of IC manufacturing.

Electromigration is one of the major causes of interconnect failure in VLSI ICs. EM causes both open and short circuits in interconnect wiring as shown below.



One of the major causes of interconnect failure in VLSI circuits, EM can cause both open (top) and short (bottom) circuits.

ELECTROMIGRATION

However, Joule heating results from rms current, not the average current as in EM. For a narrow pulse, the rms current can be much higher than the average current. The average current can be well within any guidelines you may set for EM considerations, yet significant Joule heating can result. This problem can be worse on upper level metallization, in which heat must conduct through several layers of interleaved dielectric, which is a poor thermal conductor.

Temperature gradients: the real culprit

The problem with Joule heating is not the modest temperature increase, but the resulting temperature gradients. Typically, at the current densities of modern circuitry, temperature increases occur measuring a few degrees to a few tens of degrees Celsius. This range produces temperature profiles that decay within a few micrometers, so that you can find temperature gradients of 10^4 to 10^5 °C/cm. Because EM is thermally activated, the temperature gradients produce flux divergences that approach those found in absolute divergences, such as at contacts or at microstructural features.


Thus, the rms-current-density limit should be about 2×10^6 A/cm² for lower level lines and about half this number for upper level lines. Unfortunately, you cannot accurately estimate the reliability of metal lines in the presence of temperature gradients. Temperature gradients can vary tremendously throughout a structure, depending on subtleties of the geometry and on the use of the underlying silicon devices. A conservative approach is the only way to address these issues: Limit the rms-current density to the suggested levels.

In the past, wide metal lines and low current densities, along with special processing, helped minimize the effects of EM. But today's technology—wires narrower than 0.5 μm, operating frequencies continuously increasing beyond 100 MHz, and no proportional decrease in power dissipation—has increased the potential for EM problems. The only solution is to use design tools and methodologies that minimize the potential for EM and find the problems when they occur.

Design issues related to EM fall into power-grid and signal-line categories. In a hierarchical design methodology, designers construct blocks with little or no knowledge of the block's surrounding, full-chip environment. This design methodology creates two problems. First, the designer responsible for the global power distribution has little information about the implementation of the local power grids within the individual blocks, making it impossible to predict at the top level how power will route throughout the chip. Second, block designers have no knowledge of surrounding blocks, so they don't know the quality of the power provided to their block or whether their block also serves as power routing to an adjacent block. This information is unavailable until designers assemble all blocks and perform full-chip power analysis.

A common EM design problem occurs when you place an embedded logic block near the edge of a chip. Consider a power bus that routes around a logic block to transmit power into the center of a chip. Unknown to the chip designer, the

FIGURE 1



A design that results in power routing through rather than around a logic block generates many EM violations, which appear in orange and yellow superimposed on the image of the power grid.

power wires within the logic block orient such that the effective width of the power wires inside the block are greater than the width of the power bus around the block. This design results in power that routes through rather than around the block, generating many EM violations inside the logic block (Figure 1).

Designing signal wires to minimize EM is a new and rapidly increasing challenge as wire width decreases. Although routing tools can adjust wire width to minimize interconnect delay, they don't adjust wire widths based on EM concerns. In addition, these tools don't know enough about circuit behavior to ascertain which wires may have EM problems. Because most designers lack the tools to adequately analyze power or signal EM, they rely on ad hoc methods to prevent signal EM and often don't find the problems until after they manufacture the chip. Ad hoc methods can yield designs that consume much more area than necessary because of routing.

Find EM before it's too late

To find EM problems, you need to extract and analyze full-chip data. Verification tools, such as Simplex Inc's Fire & Ice and Thunder & Lightning, extract and analyze both power-grid and signal EM. Full-chip EM analysis begins with extracting the parasitics associated with both signal and power networks from the layout. To support EM analysis and

give meaningful feedback to designers, this extraction must provide wire-segment size, layout layer, and layout location. The key to finding all EM problems is to avoid the use of reduced data because you might eliminate information critical to the power grid and signal EM analysis.

EM analysis also requires current-flow data to derive wire current densities. Power-grid current behaves as a pulsed dc current, whereas signal-wire current behaves as both pulsed dc and ac currents. Analyzing the transistor netlist in static or dynamic mode using test vectors determines the current's behavior characteristics. Simulating the circuit with many vectors or using average circuit behavior in a static analysis yields the power-grid current characteristic.

Finding signal problems requires two analyses: EM and Joule heating. EM analysis requires average-current data on signal lines; Joule-heating analysis requires rms current information. Performing the signal analysis net by net while simulating the charging and discharging for all possible paths determines the worst-case average and rms current for each wire segment. Then, current-density computations and the application of failure models to each wire segment can point out the problem areas. For example, **Figure 2** shows the locations of signal EM in a chip after applying a Joule-heating failure model.

How to manage EM

To manage EM, designers first need to understand the effects of average vs rms current density and the effects of line widths and vias on EM. It also helps to understand the current flow in each design, how to limit current density in metal lines and contacts, and how to minimize EM without overdesigning.

Designers often misunderstand the relationship between wire width and susceptibility to EM. A simple approach for managing EM is to reduce average current density by increasing the width of the metal lines. But blindly increasing the width of narrow wires may actually increase susceptibility to EM because of the material properties of aluminum wires.

Vias conduct current from layer to layer through metal lines that are thinner than the interconnect wires, causing a local increase in current density. The number of vias, their location, and their structure can impact current density enough to cause EM. Unless the verification tools analyze these effects, you may not find via-caused EM until after manufacturing.

Understanding, analyzing, and predicting current flow can be difficult because of the high connectivity of power grids. The most difficult aspect of power-grid design is that you can't isolate one block from another. One logic block analyzed in isolation may show no EM risk. However, viewing that block in a full-chip context may show how current flowing to adjacent blocks is overloading the power connections in the analyzed block and causing EM.

Designers frequently lack information on where power is dissipating, and they typically pay little attention to "local" decisions they make during power-grid design that may impact the overall power flow. But EM risk lurks everywhere. Logic gates with high loading can be anywhere. Designers

often remove power-grid wires to complete signal routing. Large power buses designed to protect a block can attract more power-grid current because they offer a lower resistive path to the pins.

Make a reliability budget


Eliminating the last trace of EM may seem desirable, but this goal may not be worth the resulting extra design time or design area. How much time should designers spend fixing problems, and how do they prioritize the problems? When should they stop improving the design? Some design tools, such as Thunder & Lightning, can now help in "reliability budgeting"—helping designers estimate a chip's lifetime, given its existing EM violations, and, based on the estimate, determine the time to stop working on additional EM issues.

Repairing all areas of a chip that have potential EM problems is labor-intensive, time-consuming, and, probably, unnecessary. You can use an MTTF factor to compute the probability of chip failure from EM during the chip's lifetime and then modify the power grid to decrease the probability of failure to an acceptable level.

Designers typically base design rules on worst-case scenarios that assume that all wires use the highest current-density limit. This approach is confining and results in conservative designs. EM failure is statistical, and only a few circuit elements typically run at the EM limit. Most circuit elements have little to no current flowing through them.

To perform reliability budgeting, you need to know how much current is flowing through each element. Tools can compute a probability of failure in time by calculating the MTTF from each circuit element and applying extreme value "links-in-a-chain" statistics. If the probability of failure is

FIGURE 2



Tools such as Simplex Solutions Inc's Thunder & Lightning can highlight the signal-line locations that are subject to Joule-heating failures when you operate the chip at a particular frequency, in this case, 200 MHz.

ELECTROMIGRATION

unacceptable, you can eliminate identified trouble spots and design a more reliable circuit.

The role of place-and-route tools

Ideally, a place-and-route tool should be able to route a design to make it resist both signal and power-grid EM. Unfortunately, current place-and-route tools are not ideal. To ask an already-overburdened place-and-route tool to handle EM issues is probably too much to ask.

Power-grid EM issues pose problems because place-and-route tools lack three pieces of critical information. The first missing piece is the average switching information for each of the cells within the block. This switching information is useful for estimating the average current that each cell in the design dissipates. Second, these tools also lack loading information for each cell that drives an output of the block. Finally, these tools don't tell you the flow of power current through the block that results from the block's orientation with respect to power pins and other blocks on the global grid.

Today's place-and-route tools could route a block to make it resist signal EM if they could perform net simulation to compute the appropriate current data for each wire segment and if they could access clocking information, cell-driver characteristics, and information on the switching activity of every signal. Without this information, most place-and-

route tools can route a signal-EM-resistant block, but the resulting design would be too conservative to accommodate the lack of critical information. EDN

Authors' biographies

James Lloyd PhD, is president of Lloyd Technology Associates Inc and has 19 years experience in metallization reliability with IBM and Digital Equipment Corp. He has taught graduate-level courses in materials science at Stevens Institute of Technology (Hoboken, NJ) and Polytechnic Institute of New York (Brooklyn, NY), and many seminars in electromigration technology. Lloyd is the author of more than 60 papers in technical and scientific journals. Currently he is an adjunct professor at the University of Maryland—College Park, and was a Visiting Scientist at Max-Planck-Institut in Stuttgart, Germany.

David Overhauser is vice president at Simplex Solutions (Sunnyvale, CA), where he has worked for three years. He received a PhD and an MSEE from the University of Illinois—Urbana/Champaign and a BS from Purdue University (West Lafayette, IN).

VOTE

Please use the Information Retrieval Service card to rate this article (circle one):

High Interest
586

Medium Interest
587

Low Interest
588