

Planning power distribution in UDSM designs

Power distribution systems in ICs are designed to provide needed voltages and currents to the transistors that perform the logic functions. You normally assume the supply voltages to be constant across a chip, and expect them to operate reliably over the chip's lifetime. However, in ultra-deep-submicron (UDSM) designs, the VDD and VSS grids fluctuate in value during chip operation due to the increased resistance of metal lines, high current levels, and package pin inductances. Furthermore, narrow line-widths reduces long-term reliability of a chip.

Power systems have become so complex that you can no longer design them using intuition and "back-of-the-envelope" calculations. You can confidently tape out your design only after analyzing and verifying the power system at the full-chip level.

Voltage drop

You would typically compute resistance along a conductor by simply counting the number of squares along the line and multiplying by the

sheet resistivity, which is usually provided in ohms per square. At the UDSM level, line-widths shrink, increasing the number of squares and causing an increase in the total resistance along the line.

Similarly, you would compute the approximate capacitance along the line by multiplying the area of the line with the capacitance per unit area. However, both the capacitance and resistance are distributed along the metal line

Solving power-grid design problems at ultra-deep-submicron levels requires more than just traditional methods of rerouting and strapping metal layers. Only a full-chip analysis will help you avoid IR drops and electromigration.

By RESVE SALEH

VP of Engineering
Simplex Solutions Inc.

By MICHAEL BENOIT

Director of Customer Applications
Simplex Solutions Inc.




Figure 2: Depending on the methodology, lower levels of metal are often left floating until final assembly.

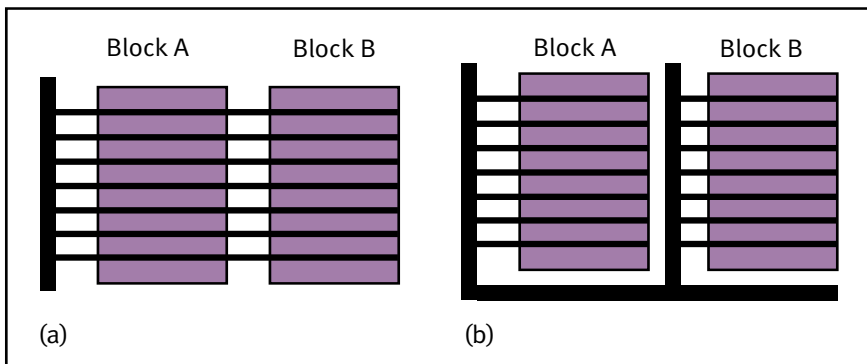


Figure 1: Rerouting or using the mesh array methodology may not reduce the IR drop. Such intuitive "quick fixes" often shift problems somewhere else within the design.

so that the RC values derived this way are already inherently accurate.

The effect of increased resistance is that supply voltage varies during normal circuit operation. The current flowing through the resistance in the power grid causes IR drops that depend on the placement of blocks, their interaction, current levels, and resistance levels. In UDSM designs, lower supply voltages yield smaller noise margins and IR drop is a first-order effect that cannot be ignored.

Another source of voltage drop in the power supply is the package pin inductance—typically around 10nH to 20nH. $L di/dt$ creates a voltage drop across an inductor. The value of di/dt has continued to in-

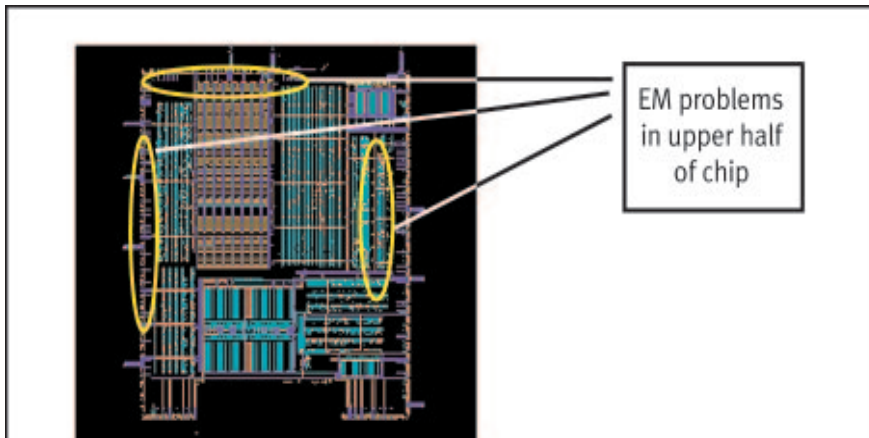


Figure 3: The most difficult aspect of power-grid design with respect to EM is that no one block can be isolated from another.

crease over the years. And the supply voltage has been decreasing from 5V to 3.3V to 2.5V and recently as low as 2V. These effects combine to a point where the Ldi/dt drop can significantly contribute to an overall voltage drop in the power grid.

Ground bounce and EM

Validating that the ground voltage does not rise above a 10-percent noise budget is as important as ensuring that VDD does not drop below a 10-percent budget. To measure ground bounce, you need to model the substrate as a distributed RC network in parallel with the metal routing for the ground grid. This significantly increases network complexity, especially when pin inductance or a more complicated pin model is included.

You can obtain a limited form of ground bounce by modeling substrate contacts as individual ideal capacitances, but these values are difficult to obtain. And if you ignore the substrate in ground-bounce analysis, you would notice that the behavior observed during analysis is worse than the actual ground bounce.

Electromigration (EM) is another important issue at UDSM. High current densities and narrow linewidths cause EM. Chip speeds of $\geq 100\text{MHz}$ and geometries of $\leq 0.35\mu\text{m}$ have increased the potential for EM problems.

Failures due to EM occur in the field when the chip is in a system. Depending on the location and number of failures, the chip may begin to operate incorrectly or shut down completely.

Routing around the problem

Suppose that you have two blocks somewhere in your design (**figure 1a**). If power distribution for Block A is examined in isolation, the additional loading due to Block B is neglected. If power is routed through Block A to Block B, a large IR drop will occur in Block B since power is also being consumed by Block A

requirements rather than on IR drop. Or you would base it on the size and shape of blocks at the floorplanning stage. Therefore, sizing the buses to minimize IR drop while satisfying timing and area constraints becomes a design challenge.

Since the total IR drop is based on the resistance seen from the pin to the block, you could route around the block and feed power to each block separately (**figure 1b**). In this case, the T-junctions have a high current density and may be prone to EM problems. It is, therefore, important in this type of grid that you examine the current density at all junctions, especially at the corner providing large current to each block.

Moreover, in routing power this way requires that large metal trunks be sized to handle all the current for each block. This will force you to set aside area for power busing that takes away from the available routing area.

In another approach, you can have a solid grid of Metal 4 and Metal 5 and likewise use a via array to connect the two layers, effectively tying the whole grid to VDD. However, this remedy simply shifts the




Figure 4: Strapping reduces the depth of the IR drop valleys to acceptable levels and spreads the voltage drops over a wider area.

before it reaches Block B. As blocks are added, the interactions between the blocks determine the actual voltage drops.

You would typically place these blocks based on the system timing

problem down to the lower levels of metal (**figure 2**). 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? ▸

When you design the logic circuitry in the block, it is not clear where Metal 3 will tap to Metal 4 and you cannot predict the current flow. If you cannot predict it, you must analyze it.

You may have to remove part of the grid to route some signals. 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. You need to determine the current levels in the straps and then pick ones that have lower current levels. The complexity of the problem requires a set of power-grid analysis tools.




Figure 5: Solving IR problems alone may create EM problems and vice versa. You must consider them together during the design stage.

A dynamic phenomenon

Note that the IR drop is a dynamic phenomenon due primarily to simultaneous switching events in a chip resulting from clocks, bus drivers, and memory decoder drivers. In a static context, voltage 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 voltage drops anywhere on the chip. These events, usually well known, can be triggered with typically fewer than 100 vectors.

Supply voltages have considerably decreased with shrinking device

dimensions. This has decreased noise margins. With IR drop, the margins are reduced even further.

However, voltage 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 percent drop in supply voltage can affect delay by ≥ 15 percent.

EM in via arrays

Since large currents flow in the periphery of a design, you will usually observe EM problems 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 gen-

strains in the routing area. Another 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 that are driving the gates.

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 capacitors are usually scattered throughout the design in any available space, using transistors with their gates tied to VDD and their source-drains tied to VSS. All empty regions of the chip are filled with decoupling capacitors using the philosophy that you can never have enough. You can reduce Ldi/dt effects by placing large capacitances near the pins.

Reducing EM problems

The basic idea in all approaches that solve EM problems is to reduce the average current density seen by any metal segment. The simplest method comprises widening 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 a full-chip context, current flowing to adjacent blocks may overload the power connections in the

erally narrow and may cause EM problems depending on the current levels.

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 EM failure, and an EM analysis tool would miss the problem.

Reducing voltage drop

The simplest approach for reducing the impact of IR drop is to widen the lines that experience the largest voltage drops. However, this may not always be possible due to con-

block, and the analysis tool identifies an EM risk (**figure 3**). Recognizing these problems at the planning stage is helpful, but difficult to do. A complete picture of EM risk can only be obtained at the verification stage.

A key point is that you cannot solve IR drop and EM problems separately. Consider how to solve an IR drop problem in the chip in **figure 4a**. The figure shows a power flow diagram of the VDD grid in a multimedia chip. Different shading indicates various levels of voltage drops. The darkest areas are the lowest points (valleys) of the IR drop contours. A significant voltage 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. If we strap the upper and lower regions together in two

places, the voltage drop problem is significantly reduced (**figure 4b**).

However, if you examine the result in the context of EM, you will notice that fixing the IR drop problem has caused an EM problem in the lower portion of the design (**figure 5**) 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 is built between the two, but it was not clear exactly how current would flow and exactly where EM problems might occur.

Repairing all the areas with potential EM problems would be labor-intensive, time-consuming, and unnecessary. Since every chip has a lifetime associated with it, you can use the MTTF factor to compute a probability of failure due to EM in a given lifetime. The immediate 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 more manageable.

In the past, certain DRC and visual checks were performed on the grid to ensure compliance with the constraints imposed by these issues. Usually, over-designing was an acceptable solution. But as technology moves deeper into UDSM, this is not a viable approach.

Too much performance is sacrificed or the area penalty of over-designing leads to decreased yields. And under-designed chips often fail on the test bench or later in the field. Situations of over-design and under-design must both be identified when evaluating the integrity of a power distribution system. **ee**

You may send your comments on this article to Resve Saleh through e-mail at res@simplex.com, or Michael Benoit at mbeinoit@simplex.com; fax: 1-408-7740285.