Power integrity for PCB designs

by Patrick Carrier, Mentor Graphics, USA

Not that long ago, most ICs required just a single 5 V DC power supply, and all a PCB designer needed to be concerned with were a signal layer for 5 V DC and a single ground layer, and it was easy to deliver sufficient clean power to all components. Now, it's not uncommon to see multiple voltages, and as low as 0,9 V DC.

Fig. 1 shows a typical integrated circuit (IC) with multiple voltages and the complex power distribution networks (PDN) required to supply those voltages and grounds. Several factors contribute to the increase in printed circuit board (PCB) design complexity for the CAD designers as well as the EE who must analyse these PDNs: the growing number of different voltages on a design, the smaller voltage magnitudes (resulting in very tight tolerances), increasing power demands from ICs, and the ever-present issues of increasing operating frequency and the need to reduce product cost (you can no longer afford to devote a PCB layer to every unique voltage and ground).

Signal integrity issues have been around for a while – brought about by faster and faster edge rates. Both signal and power integrity problems manifest themselves as data errors, and it is common for power integrity issues to be misdiagnosed as signal integrity problems. As such, an understanding of power delivery issues, as well as a proactive approach in dealing with these issues early in a design process, is essential to the success of a printed circuit board design.

Challenge for the designer

The basic challenge for the CAD designer is how to create PDNs that feed from the voltage regulator modules (VRMs) that supply those voltages to all the ICs that require those voltage (or ground) on the fewest number of PCB layers. The good news is that defining a PDN is highly automated by functionality found in more advanced PCB layout systems like Expedition Enterprise from Mentor Graphics. Simply outline the PDN shape and the software floods the area. The bad news is that these PDNs must typically be "jigsawed" into the fewest number of layers. Some designs are known to contain over 30 PDNs creating a very challenging situation for the designer.

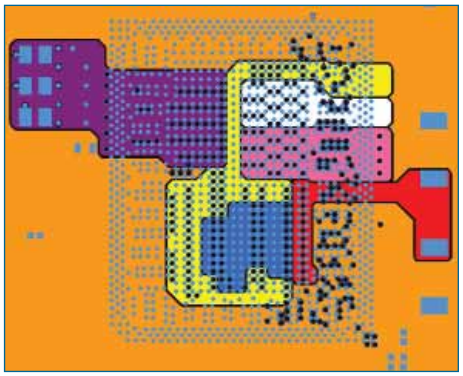

Fig. 1: This single IC requires six voltages and each must be fed by a separate PDN.

But even after the PDN layout is complete, the CAD designer is not out of the picture. They must either perform the power integrity analysis themselves or work closely with the supplier to ensure that they are delivering sufficient and clean power to all of the ICs. Corrections may require changes to the PDNs and/or the placement of capacitors on the PCB.

Minimise impedance

From an electrical point of view, the goal for designing a PDN for a PCB is simple: minimise the impedance between power and ground for your frequency range of interest. Successful execution of that design goal is not quite as simple. With multiple voltage rails and a limited number of planes available to carry those voltages, in addition to ever-shrinking real estate available for capacitors, the task becomes quite complex.

The PDN consists of a combination of the DC-to-DC converter VRMs supplying voltage to the rail, the decoupling capacitors tied between power and ground, the planes and/or traces carrying power, on-chip decoupling, and the pins and vias connecting all these elements together. The VRM is very effective in providing a low-impedance path between power and ground up to around 1 MHz. For the remainder of frequencies, the low impedance path between power and ground must be provided by the board and chip capacitances.

The board and chip capacitances combine in parallel, but are each limited in effectiveness by their parasitic inductances and resistances. Without parasitics, all the capacitances would combine to make one large capacitance, which would equate to a lower and lower impedance with increasing frequency. Unfortunately, though, each capacitance is only effective in a given frequency range, as determined by its parasitic inductance. For instance, very large electrolytic capacitors reach a low impedance at a lower frequency than smaller surface-mount caps, but because their parasitic inductance is also larger, an electrolytic capacitor's impedance will start to rise at a lower frequency than smaller surface-mount caps. Another example is the inherent capacitance between planes on the board. The parasitic inductance of the planes is extremely low, making it an effective capacitor even at higher frequencies.

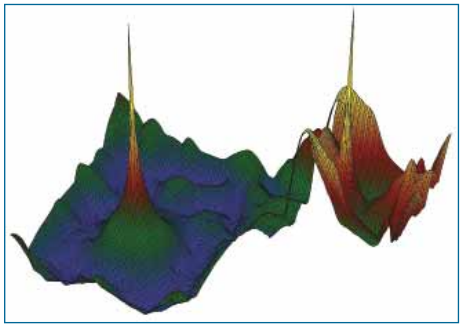

Fig. 2: Results of a pre-route noise analysis of a power distribution network in HyperLynx power integrity.

Because the amount of capacitance between planes is typically limited by their area and spacing, the plane capacitance does not equate to a low impedance until higher frequencies. As such, each of the board capacitances is inherently effective for a certain frequency range, and must all work together to provide a low impedance between power and ground across the entire frequency range.

One factor which limits the effectiveness of the board capacitances is the inductance of the

chip package. This additional inductance adds to the parasitics of the board capacitances, making them ineffective above about 1 GHz. Above 1 GHz, the on-chip capacitance (not limited by the inductance of the package) provides the low-impedance path between power and ground. As such, board decoupling is typically analysed between about 1 MHz and 1 GHz, and board PDN design is focused at minimising the impedance between these frequencies.

In order to make capacitors effective over the largest frequency range possible, choosing the largest capacitance value possible for a given parasitic inductance is the ultimate goal. Parasitics for decoupling capacitors consist of both inherent parasitics as well as mounting parasitics. The inherent parasitics, effective series resistance and effective series inductance, are properties of the capacitors themselves. The mounting of the capacitors can add significant inductance and resistance, and minimising those mounted parasitics will maximise the effective frequency range of the capacitor. The most effective means of minimising mounting parasitics is by minimising the loop area of the connection of the capacitor between power and ground. This means placing mounting vias as close together as possible, and placing the capacitor as close as possible to power and ground.

Maintaining a low PDN impedance introduces a number of benefits. The most direct benefit is minimising the amount of voltage ripple at the IC supply pins. In fact, the tolerable amount of voltage ripple at the IC supply pins is what determines the maximum allowable PDN impedance. The PDN impedance requirement may be calculated using the following formula:

$$Z_{pdn} = (\%\ ripple \times V) / I_{max}$$

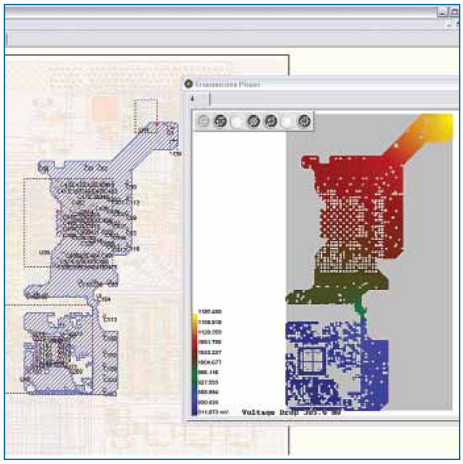

Fig. 3: Plot of DC voltage, the results of a power plane analysis.

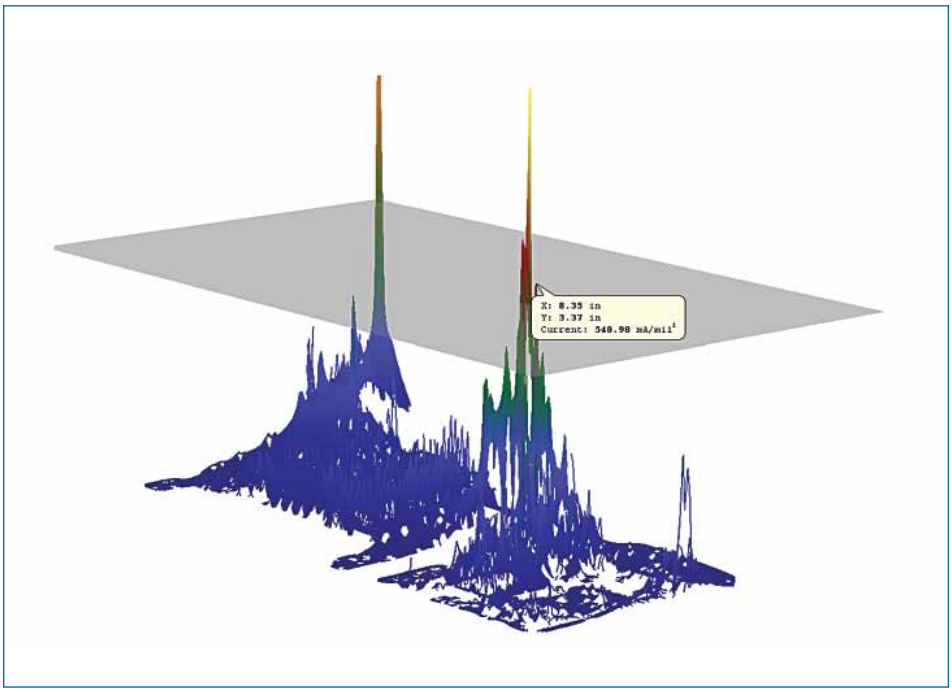

Fig. 4: Three-dimensional plot of current density on a power plane, the results of analysis performed with HyperLynx power integrity.

where V is the rail voltage and Imax is the peak current draw of the IC. As such, ICs with higher current draw and lower operating voltages require a lower PDN impedance.

Another benefit of a low PDN impedance is the reduction of noise propagation throughout the board. Power pins, signal pins, and vias can all introduce noise onto the planes which can propagate around the entire PCB and affect other circuits, as can be seen in Fig. 2. The presence of low-impedance paths between power and ground facilitate better flow of return currents. Since noise is fundamentally unresolved return currents, the presence of these paths helps eliminate the noise issue.

DC problems

At DC, the problem becomes much simpler, but failures can be much more severe. The goal is still to minimise the impedance of the PDN, but at DC the main concern is providing an adequate amount of metal to distribute up to tens-of-amps of current. With planes being divided up to accommodate the growing number of different voltages on a PCB, this can be quite a daunting task.

Analysis of plane shapes within the PDN can be very useful for identifying areas of low voltage, such as in Fig. 3. Areas of the board where voltage drops below a certain minimum threshold can lead to IC malfunction, including but not limited to, signal integrity problems and logic errors.

In addition to voltage drops, of equal concern are areas of high current density on the board. These can be caused by neck downs due to the partitioning of the plane, but are also common in dense connector and IC pinfields, where antipads eliminate a great deal of the copper on plane layers. Areas of high current density lead to voltage drop, but can also lead to board failure. A narrow neckdown can act like a fuse which ultimately leaves power disconnected.

If the current density is high enough, it can also lead to dielectric breakdown and even cause the board to catch on fire. As such, it is crucial to monitor areas on the planes where current densities may become excessive. Analysis of the PDN at DC will also produce this useful current density information, like the plot in Fig. 4. It is also important to monitor current through vias connecting planes together or connecting VRMs to planes. Vias can act as similar neck downs and areas of high current density, resulting in similar failures.

Summary

Through an understanding of proper design of a power distribution network, coupled with analysis early in the design phases, severe issues with power integrity can be avoided. This ultimately results in meeting schedule and budget goals, while designing products that are more reliable and last for many years to come.

Contact Willem Hijbeek,
Tempe Technologies, Tel 011 452-0530,
willem.hijbeek@telkomsa.net