## 13.5 Every Product Has a Unique PDN Requirement

One of the greatest sources of confusion in PDN design is created by taking the PDN design features of one product and blindingly applying them to another product.

---

**TIP** Unlike the design of signal paths, where the design rules in one product can often be applied to other products of similar bandwidth, the behavior of the PDN depends on the interactions of all of its parts, and the goals and constraints vary widely from product to product.

---

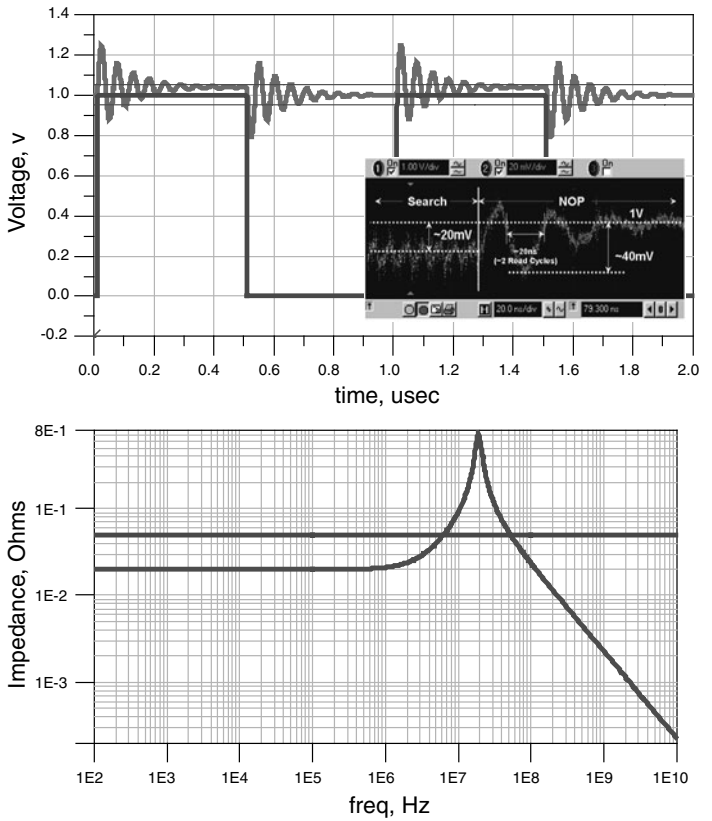


**Figure 13-9** When the impedance exceeds the target spec (bottom) and the current's peak frequency hits this impedance peak, excess ripple can result (top). The square wave is the current draw through the chip. The ringing wave is the voltage on the supply rail. Inset is an example of the measured voltage noise on a PDN showing the typical ringing response of a peak in the impedance profile.

The PDN is one giant net, rather than a large number of individual nets, with only a small amount of local coupling between them. In this respect, the PDN net is like an ecosystem of interconnects. While it may be possible to suggest an optimized design from parts of it, the most cost-effective designs are based on optimizing the entire ecology of all the elements, across the entire frequency range.

The voltage level of the rails can vary from 5 volts to less than 1 v depending on the chip type and technology node. The ripple spec may be as large as 10% in some devices, or as low as 0.5% in others, such as phase locked loop (PLL) supplies or the analog to digital converter (ADC) reference voltage rails.

The current draw from chips can vary from more than 200 A in high-end graphics chips and processors, to as low as 1 mA for some low power micro controllers.

This means that the target impedance values can vary from below 1 mOhm for high-end chips to more than 100 Ohms. This is five orders of magnitude in impedance.

There could be as many as 10 different voltage rails in some designs, many sharing the same layer, while other designs may have just one voltage and ground. Some of the planes may be solid; others may be irregularly shaped and full of clearance holes.

**TIP** This wide variety of applications and board constraints means no one solution is going to fit all. Instead, each design must be treated as a custom design.

It is dangerous to blindly apply the specific features of one design to another design. However a general strategy can be followed to arrive at an acceptable impedance profile.

## 13.6 Engineering the PDN

It is remarkable that so complex a structure as the PDN interconnects can be partitioned in the frequency domain into just five simple regions. Figure 13-10 diagrams these five regions, based on what frequency range they can influence.

At the lowest frequency, the VRM dominates the impedance the chip sees looking into the PDN. Of course, the series resistance of the interconnects can also set a limit on the lowest impedance of the PDN, if it is larger than the VRM impedance. The VRM performance dominates from DC to about 10 kHz.

The next higher frequency regime, roughly in the 10 kHz to 100 kHz range, is dominated by the bulk decoupling capacitors. These are typically electrolytic and tantalum capacitors that provide a low impedance beyond where the VRM can go.

The highest frequency impedance is set by the on-die capacitance. This is the only feature in the PDN that the chip sees in the GHz regime. It generally has the lowest loop inductance associated with it and offers the lowest impedance at the highest frequency of any element in the PDN.

**TIP** Every chip interfaces to the board it is mounted to through some mounting inductance. Usually, this is dominated by the package, the board vias, and the spreading inductance of the via contacts into the power and ground planes of the board.

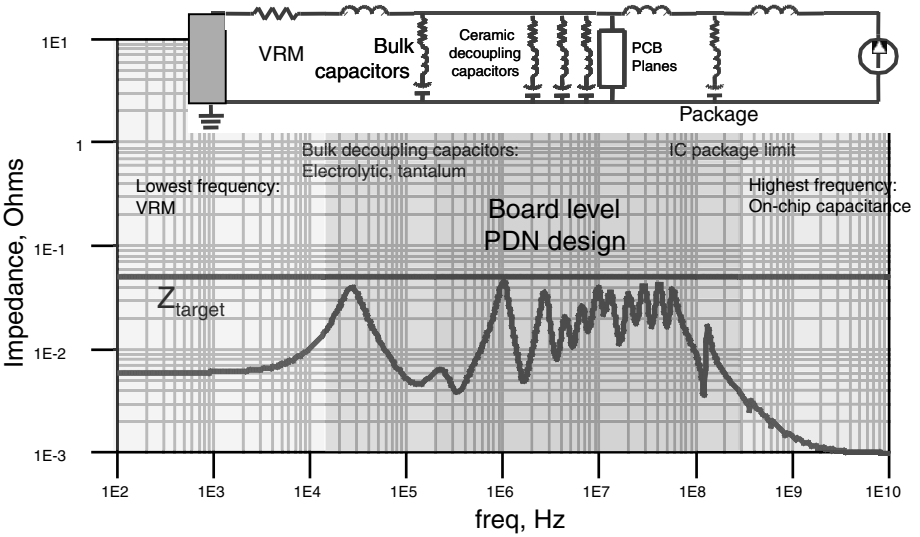


**Figure 13-10** The five parts of the PDN separated by the frequency range they influence.

The PDN interconnects in the package are generally inductive. This means that at high frequency, they will act as a high impedance path. Even if the board was designed with an impedance of a dead short, the chip would be looking at this short through the chip attach and package attach inductance, and would see an impedance limited by these inductances.

The series inductance of the package's PDN will always limit the highest frequency at which the chip will see the board level PDN. This acts as a high-frequency limit to the board level PDN design. This means that above this package-limited frequency, the impedance the chip sees will be determined by the on-die capacitance and any capacitance in the package. This limit is generally in the 100 MHz range. Above this frequency, the impedance the chip sees is all about the package and the chip.

**TIP** The frequency region that board level PDN design can influence is roughly from the 100 kHz range up to about 100 MHz. This is where the planes of the board and the multilayer ceramic chip capacitors (MLCC) can play a role.

These capacitors, typically in sizes of 60 mils × 30 mils and referred to as 0603 or 40 mils × 20 mils and referred to as 0402, are called *chip capacitors* because they look like small "chips" of something on a board. Figure 13-11 is a

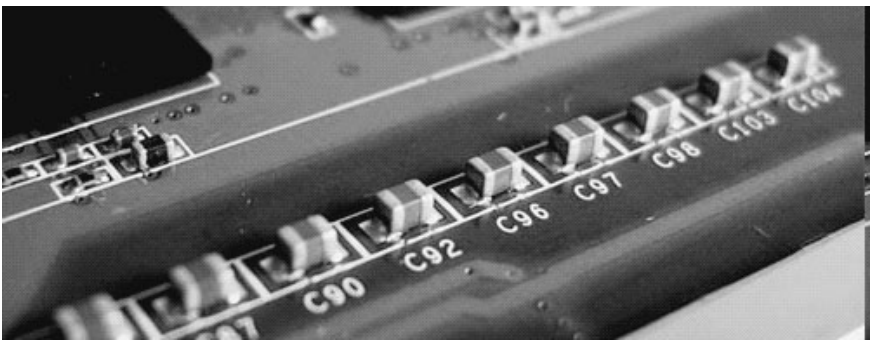


**Figure 13-11** Close-up of typical 0402 MLCC capacitors mounted to a small memory board.

close-up of some typical multilayer ceramic chip (MLCC) capacitors on a small memory board.

## 13.7 The VRM

The low frequency impedance is set by the voltage regulator module. Regardless of the type of regulator, all VRMs have an output impedance profile. This can easily be measured using a two-port impedance analyzer.

An example of the measured impedance profile of a typical VRM is shown in Figure 13-12. In this example, the impedance looking into the output leads of the VRM was measured when the regulator was turned off and when it was turned on and providing regulation. In addition, the impedance of a simple two-capacitor model is superimposed.

This illustrates that when the regulator is off, the behavior seen at the output leads is almost exactly as predicted by a two-capacitor model, each capacitor being modeled as an RLC circuit.

This behavior corresponds to the two bulk decoupling capacitors associated with the VRM. The 910 uF capacitor is an electrolytic capacitor, while the 34 uF capacitor is a tantalum capacitor. This impedance profile is for the passive network of the leads and the two capacitors.

When the regulator is turned on, its output impedance drops by orders of magnitude at low frequency. This is exactly what is expected from a regulator. The output voltage is kept constant independent of the current load. A large change in current produces a small change in voltage, the behavior of a low impedance. However, we see that in the actual behavior of the VRM, this low impedance is maintained from DC only up to the kHz range.

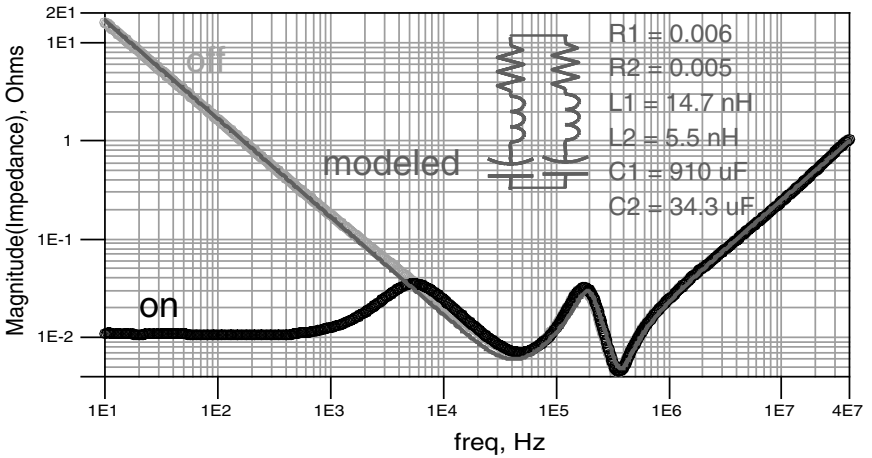


**Figure 13-12** Measured impedance profile, from 10 Hz to 40 MHz, using an Ultimetrix Impedance Analyzer for a typical VRM, showing the impedance when on, when turned off, and the modeled impedance based on a two-capacitor model.

Above about 1 kHz, the impedance is seen to increase, until it matches the impedance of the bulk capacitor at about 4 kHz, at which frequency the impedance is brought down by the passive capacitor network on the regulator. Above about 1 kHz, the output impedance of the VRM is completely due to the passive capacitors, and the active regulation plays no role at all. Whether the regulator is on or off, the impedance is the same.

This is a slight exaggeration since the regulator actually fights with the capacitance of the passive network and when the regulator is turned on its impedance is actually higher than if it were literally turned off.

**TIP** The output impedance of most VRMs is low up to the kHz regime. Beyond this, what brings the impedance down is the bulk capacitors associated with the regulator.

The total amount of capacitance needed on a board in the form of electrolytic or tantalum capacitors can be estimated based on achieving the target impedance at the frequency where the VRM is no longer able to maintain the low impedance.

The capacitance is chosen so that its impedance at 1 kHz is less than the target impedance. The minimum capacitance needed is given by:

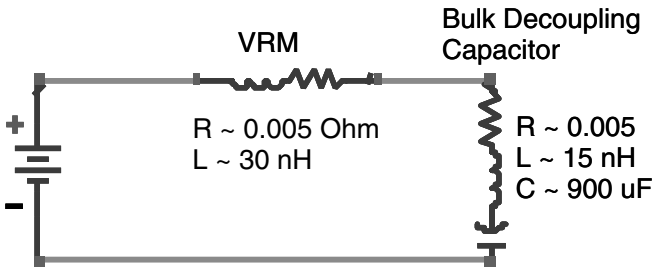


**Figure 13-13** Typical equivalent circuit model of a VRM and bulk decoupling capacitor with typical values for each element.

$$C_{bulk} < \frac{1}{Z_{target} \times 2\pi \times 1\ kHz} = \frac{160\ uF}{Z_{target}} \quad (13\text{-}14)$$

where:

$C_{bulk}$ = the minimum bulk capacitance needed, in uF

$Z_{target}$ = the target PDN impedance in Ohms

1 kHz = the frequency at which the VRM is no longer able to provide low impedance

For example, if the target impedance is 0.1 Ohm, the minimum bulk capacitance needed is about 1600 uF. Of course, this is only a rough estimate but a good starting place. When it comes to establishing the actual target values of the capacitance, the interactions of the VRM effective inductance and capacitor's capacitance must be taken into account with a SPICE simulation.

The low-frequency model of a VRM can be easily approximated by a simple RL model with a voltage source. The equivalent circuit model of the VRM and bulk decoupling capacitor is shown in Figure 13-13.

This circuit can be used to optimize the capacitor value to keep the impedance below the target value at low frequency.

## 13.8 Simulating Impedance with SPICE

Simulating the impedance profile of different circuit models is essential in PDN design. Luckily most of the simple circuits that need to be analyzed can be simulated with free versions of SPICE that can be downloaded off the internet.

**TIP** The secret to using SPICE to perform impedance simulations is to build an impedance analyzer as a SPICE circuit. This is done with a single element in SPICE, a constant current AC current source.

This element is defined as a constant current sine wave source, outputting a sine wave of current with a constant amplitude. The output voltage of this element will be whatever it needs to be to always output a constant amplitude sine wave of current. The frequency of the current is set by the frequency of the frequency domain simulation. An example of a SPICE impedance analyzer is shown in Figure 13-14.

The amplitude of the constant current source is set to 1 A with a phase of 0. The voltage across the constant current source will depend on the impedance of whatever is connected across the leads. This voltage generated will be given by:

$$V = I(f) \times Z(f) = 1 \times Z(f) = Z(f) \tag{13-15}$$

where:

V = the voltage generated across the current source in volts

I(f) = the current from the source, a constant 1 A amplitude sine wave

Z(f) = the impedance of the device connected across the current source, in Ohms

We set the current amplitude to be exactly 1 A. This means the voltage generated across the current source is numerically the impedance in Ohms. The impedance of the circuit connected may vary with frequency. As the 1 A constant amplitude

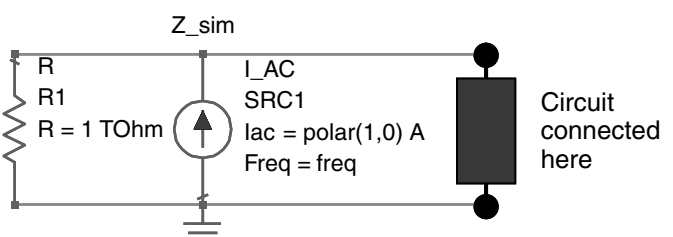


**Figure 13-14** A SPICE impedance analyzer consisting of a constant current AC sine wave source.

sine wave flows through it, a voltage will be generated that is numerically equal to the impedance. The phase of the voltage will even track the phase of the impedance.

A large shunt resistor, in this case 1 TOhm, is connected across the current source. This is to keep SPICE from halting due to an error. SPICE wants to see a DC path to ground for all nodes. Without the resistor, an open across the constant current source could result in an infinite voltage, causing an error.

With this circuit the impedance of any circuit model can be simulated. It's actually the voltage across the current source that is simulated, but this is equal to the impedance of the circuit. The impedance of the two-capacitor model in the VRM was simulated using this SPICE impedance analyzer.

## 13.9 On-die Capacitance

The impedance at the highest frequency is established by the on-die decoupling capacitance. This arises from three general sources: the capacitance between the power and ground rail metallization, the gate capacitance from all the p and n junctions, and any added capacitance.

The largest component is from the gate capacitance distributed over the die. A typical CMOS circuit found by the millions on most chips, and for some chips by the billions, is shown in Figure 13-15. At any one time, one of the gates is on and the other is off.

This means that the gate capacitance of one of the gates, either the p channel or the n channel is connected between the power and ground rails on the die. The capacitance per area associated with the gate is simply approximated by:

$$\frac{C}{A} = \frac{8.85 \times 10^{-12} \frac{F}{m} \times Dk}{h} \tag{13-16}$$

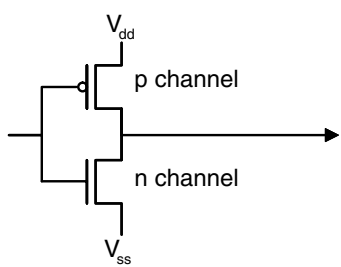


**Figure 13-15** Typical CMOS circuit model for the transistors on a chip.

where:

C/A = the capacitance per area in $F/m^2$

Dk = dielectric constant of the oxide ~ 3.9 for $SiO_2$.

h = dielectric thickness in m

In general, the shorter the channel length, the thinner the gate oxide. As a rough rule of thumb, the gate oxide thickness is about 2 nm per 100 nm of channel length. However, below about 100 nm channel length, the scaling of h flattens out, due to higher leakage currents, but then the dielectric constant is increased with the use of "high Dk" gate insulator materials. This keeps the rule of thumb a good approximation even below 100 nm channel lengths.

For the 130 nm channel length node, the capacitance per area is about:

$$\frac{C}{A} = \frac{8.85 \times 10^{-12}\frac{F}{m} \times 3.9}{0.02 \times 130 \times 10^{-9}} = \frac{8.85 \times 10^{-12}\frac{F}{m} \times 3.9}{2.6 \times 10^{-9}} = 1.3\frac{uF}{cm^2} \tag{13-17}$$

Of course, not all the die is gate area. If we assume that 10% of the surface of the die is gate capacitance, then we see that as a rough rule of thumb the on-die decoupling capacitance on a 130 nm technology chip due to its p and n junctions is about:

$$\frac{C}{A} = 130\frac{nF}{cm^2} \tag{13-18}$$

As the technology node advances and the channel length decreases, the gate capacitance per area will increase, but the total gate area on a die will stay about the same. This means the capacitance per unit of die surface area will increase inversely with the technology node.

The capacitance of 65 nm chips is about 260 $nF/cm^2$. This estimate suggests that for a die that could be 2 cm × 2 cm, close to the largest mask size in volume production, at 65 nm channel length, the on-die decoupling capacitance could easily be in excess of 1,000 nF.

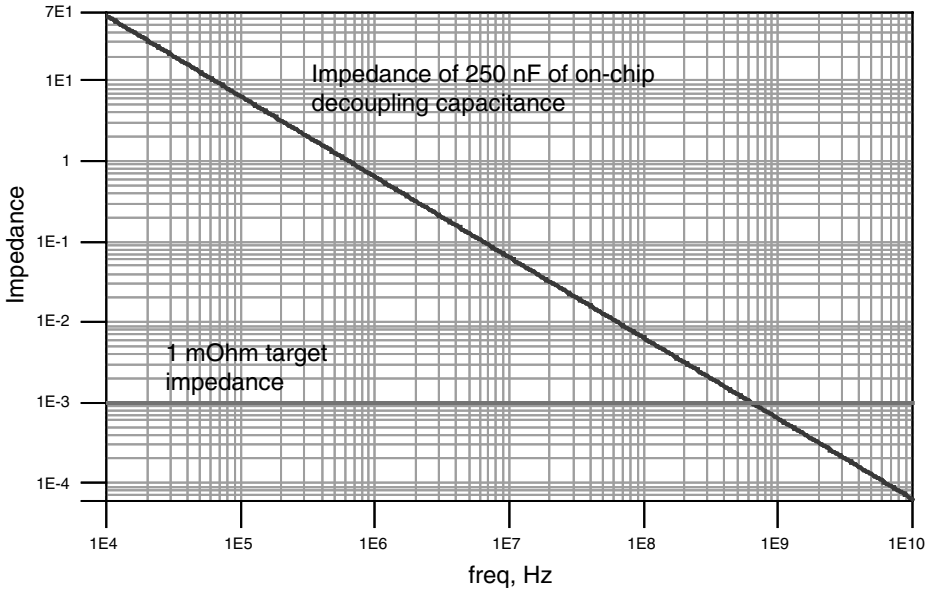


**Figure 13-16** Impedance provided by 250 nF of on-chip decoupling capacitance, typical of a 65 nm, 1 cm on a side die.

A typical chip, only 1 cm × 1 cm, typical of many embedded processors, would have as much as 260 nF of capacitance. If the gate utilization on the die were larger, the on-die capacitance could be higher as well.

**TIP** At high frequency, it is the on-die capacitance that provides the low impedance.

The impedance profile of a capacitance that is 250 nF is shown in Figure 13-16. In this example, the on-die capacitance provides an impedance below 1 mOhm, at frequencies above 800 MHz. All high-frequency decoupling is provided by this mechanism.

If the target impedance were 10 mOhms, the on-die capacitance would provide significant decoupling for frequencies above about 100 MHz.