CHAPTER 13

The Power Distribution Network (PDN)

The power distribution or delivery network (PDN) consists of all those interconnects from the voltage regulator module (VRM) to the pads on the chip and the metallization on the die that locally distribute power and return current. This includes the VRM itself, the bulk decoupling capacitors, the vias, the traces, the planes on the circuit board, the additional capacitors added to the board, the solder balls or leads of the packages, the interconnects in the packages mounted to the board, the wire bonds or C4 solder balls, and the interconnects on the chips themselves.

The primary difference between the PDN and signal paths is that there is just one net for each voltage rail in the PDN. It can be a very large net that can physically span the entire board and have many components attached.

TIP As we will see, the PDN is like an ecosystem. If one small part of the PDN were to change, the entire system performance can be affected. This makes generalizations very difficult.

13.1 The Problem

An example of a motherboard with all these interconnects is shown in Figure 13-1.

First and primary is to keep a constant supply voltage on the pads of the chips, and keep it within a narrow tolerance band, typically on the order of 5%.

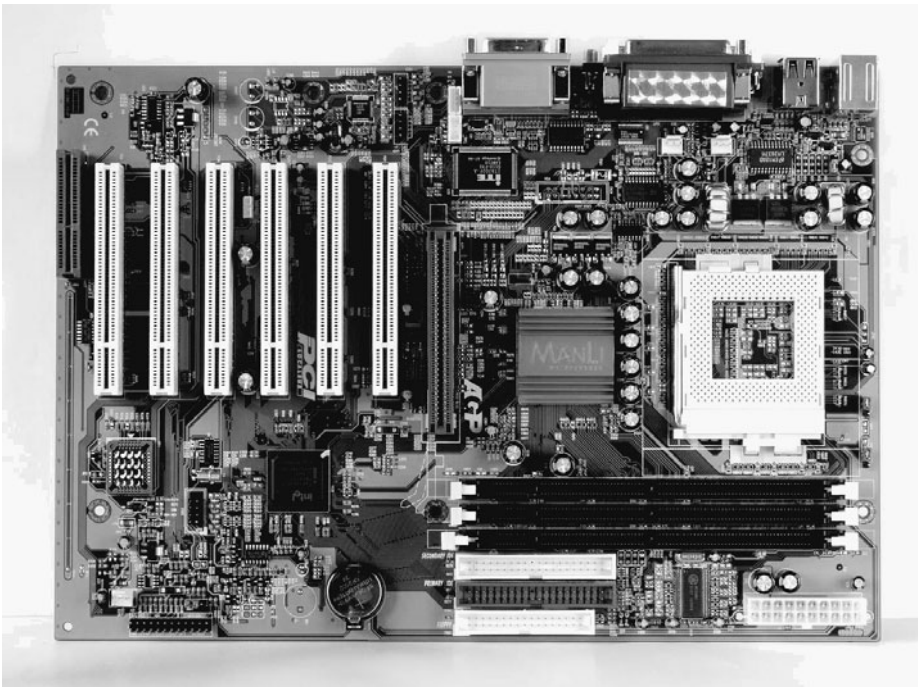

Figure 13-1 Example of typical motherboard showing all the interconnects of the PDN.

TIP The purpose of the PDN is threefold: Keep the voltage across the chip pads constant, minimize ground bounce, and minimize EMI problems.

This voltage has to be stable, within the voltage limits, from DC up to the bandwidth of the switching current, typically above 1 GHz.

Second, in most designs, the same PDN interconnects that are used to transport the power supply are also used to carry the return currents for signal lines. These interconnects must provide a low impedance return path for the signals.

The easiest way of doing this is by making the interconnects wide, so the return currents can spread out as much as they want, and by keeping the signal traces physically separated so that the return currents do not overlap. If these conditions are not met, the return current is constricted and the return currents from different signals overlap. The result is ground bounce, also called simultaneous switching noise (SSN) or just switching noise.

Finally, since the PDN interconnects are usually the largest conducting structures in a board, carry the highest currents, and sometimes carry high frequency noise, they have the potential of creating the most radiated emissions and causing failure of an EMC certification test. When done correctly, the PDN interconnects can mitigate many potential EMI problems and help prevent EMC certification test failures.

The consequence of not designing the PDN correctly is that there will be excessive noise on the voltage rails of the chips. This can cause a bit failure directly, or it can mean the clock frequency of the chip can't be met and timing errors result.

An example of the voltage noise on the pads of a processor chip is shown in Figure 13-2. In this example, the nominally constant 2.5 v rail to the core of the chip, referred to as Vdd, shows voltage noise of as much as 125 mV on some pads. As the Vdd supply drops, the propagation delay of the core gates will increase and timing problems can cause bit failures.

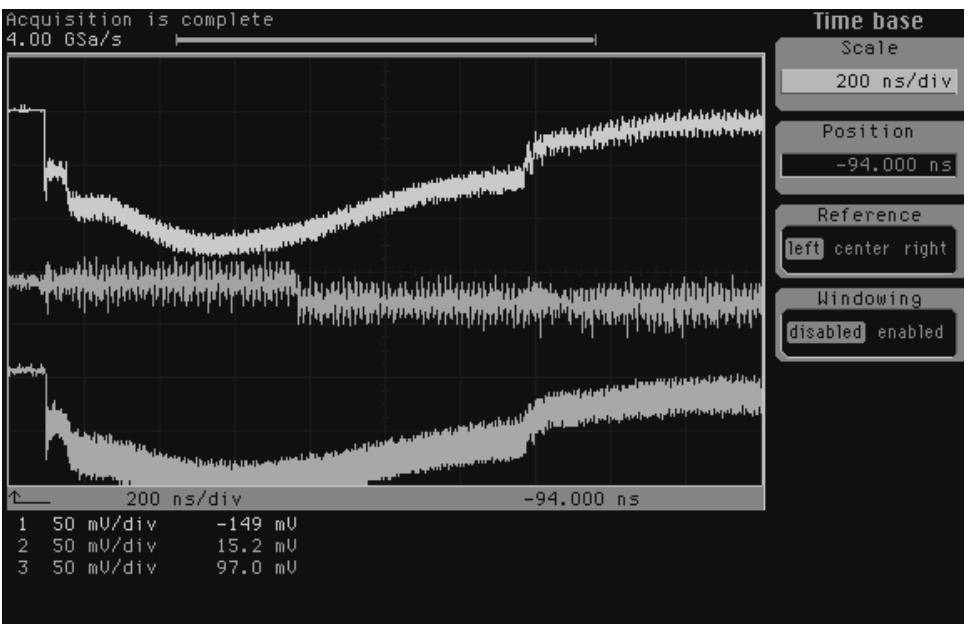

Figure 13-2 Example of the measured voltage between three different pairs of Vdd and Vss pads of a processor chip showing as much as a 125 mV drop. The initial step down is when the processor comes out of an idle state. The three traces are three different locations on the die. The precise shape is related to the microcode running on the processor.

13.2 The Root Cause

If the problem is a voltage drop or droop on the power supply rails on the pads of the chip, why not just use a "heftier" regulator, one that can supply a more rock stable voltage? Why not pay extra for a regulator with a 1% regulation or even 0.1% regulation? This way, the voltage from the regulator will be absolutely stable, no matter what, right?

What the chip cares about is the voltage on its pads. If there were no current flow in the PDN interconnects from the regulator pads to the chip pads, there would be no voltage drop in this path and the constant regulator voltage would appear as a constant rail voltage on the chip pads.

If there were a constant DC current draw by the chip, this DC current would cause a voltage drop in the PDN interconnects due to the series resistance of the interconnects. This is commonly referred to as the *IR drop*. As the current from the chip fluctuates, the voltage drop in the PDN would fluctuate and the voltage on the chip pads would fluctuate.

Now add not just the resistive impedance of the PDN, but also the complex impedance, including the inductive and capacitive qualities of the PDN interconnects. The impedance of the PDN, as seen by the pads on the chip, in general, is some impedance verses frequency, Z (f). This is diagrammed in Figure 13-3.

As fluctuating currents with some spectrum, I(f), pass through the complex impedance of the PDN, there will be a voltage drop in the PDN:

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

where:

V(f) = the voltage amplitude as a function of frequency

I(f) = the current spectrum drawn by the chip

Z(f) = the impedance profile of the PDN as seen by the chip pads.

This voltage drop in the PDN means that the constant voltage of the regulator is not seen by the chip, but is changed. In order to keep the voltage drop on the chip pads less than the voltage noise tolerance, usually referred to as the ripple, given the chip current fluctuations, the impedance of the PDN needs to be below some maximum allowable value. This is referred to as the target impedance:

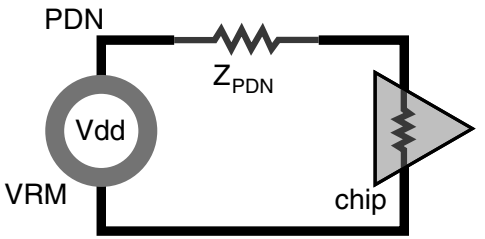

Figure 13-3 Diagram of the connections from the VRM, through the PDN to the chip pads and the voltage drop across the PDN interconnects due to the impedance of the PDN.

$$V_{ripple} > V_{PDN} = I(f) \times Z_{PDN}(f) \tag{13-2}$$

$$Z_{target}(f) = Z_{PDN}(f) < \frac{V_{ripple}}{I(f)} \tag{13-3}$$

where:

V_{ripple} = the voltage noise tolerance for the chip, in Volts

V_{PDN} = the voltage noise drop across the PDN interconnects, in Volts

I(f) = the current spectrum drawn by the chip, in Amps

Z_{PDN} (f) = the impedance profile of the PDN as seen by the chip pads, in Ohms

Z_{target} = the maximum allowable impedance of the PDN, in Ohms

As mentioned many times so far in this book, the most important step to solve a signal integrity problem is to identify the root cause of the problem. The root cause of rail collapse or voltage noise on the PDN conductors is that a voltage drop in the PDN interconnects is created by the chip's current flowing through the impedance of the PDN.

TIP If we want to keep the voltage stable across the pads of the chip, given the chip's current fluctuations, it means keeping the impedance of the PDN below a target value. This is the fundamental guiding principle in the design of the PDN.

13.3 The Most Important Design Guidelines for the PDN

The goal in designing the PDN interconnects is to bring the impedance below the target value from DC to high frequency. In general, this will be accomplished by following three important design principles. Though it may not always be possible to push these to the limit, it is always important to be aware of the directions to head, even if the real world constraints keep you from the ultimate path.

The three most important guidelines in designing the PDN are:

1. Use power and ground planes on adjacent layers, with as thin a dielectric as possible, and bring them as close to the surface of the board stack-up as possible.
2. Use as short and wide as possible surface trace between the decoupling capacitor pads and the vias to the buried power and ground plane cavity and place the capacitors where they will have the lowest loop inductance.
3. Use SPICE to help select the optimum number of capacitors and their values to bring the impedance profile below the target impedance.

Unfortunately, in the real world of practical product design, you may not always have the luxury of power and ground planes on adjacent layers or placed near the top of the board stack-up. There may be multiple voltage rails, and they may have odd and irregular shapes with many antipad clearance holes.

You may not be able to use as many capacitors as you think you need, nor place them in proximity to the devices they are decoupling. Even if you do the best you can, it will still be important to know if it is "good enough" before you build the product. The last place you want to find a problem is when you are making 100,000 units and finding that 1% of them are failing due to excessive ripple in the PDN.

The time to find this out is as close to the beginning of the design process as possible, and the only way to determine this is by using analysis tools that allow you to explore design space.

TIP It is essential to try to follow the three important design guidelines above, and at the same time, to use combinations of rules of thumb, approximations, and numerical simulation tools to predict the impedance profiles and voltage noise under typical and worst case conditions.

The most important principle to follow for cost-effective design is to add appropriate analysis as early in the design cycle as possible. This will reduce the surprises as the design progresses and result in a product with acceptable performance, at the lowest cost that works the first time.

13.4 Establishing the Target Impedance Is Hard

The first step in designing the PDN is to establish the target impedance. This must be done separately and independently for each voltage rail to all the chips on the board. Some designs may use as many as 10 different voltages. In each one, the target impedance may vary with frequency due to the specific current spectrum of the chip.

Suppose the current from the chip on one rail is a sine wave, with a peak to peak value of 1 A. The amplitude of the sine wave of current will be 0.5 A. This current from the chip is shown in Figure 13-4 in both the time domain and the frequency domain.

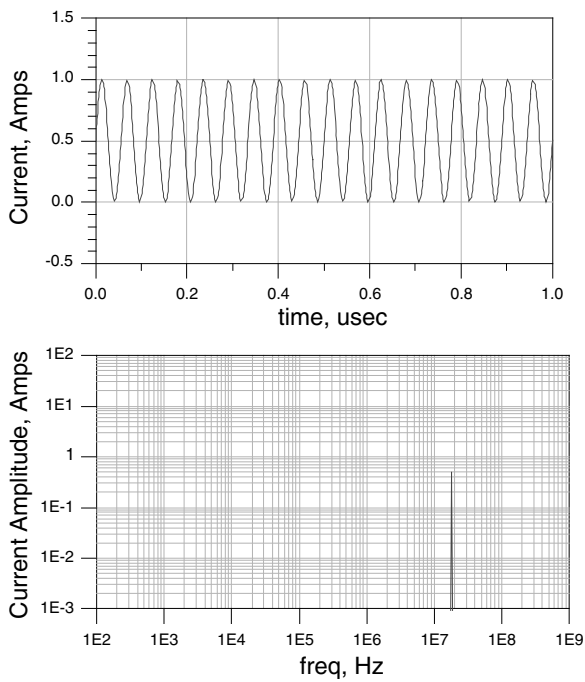

Figure 13-4 Example of the current waveform in the time domain (top) and the frequency domain (bottom) for a sine wave current draw.

When this frequency component of the current flows through the specific impedance profile of the PDN, a voltage noise will be generated in the PDN. An example of an impedance profile with the frequency component of the current, and the resulting time domain voltage noise across the chip pads, is shown in Figure 13-5.

When the sine wave current passes through an impedance that is too large, the voltage generated is above the ripple spec, which is typically +/– 5%, shown as the reference lines.

There is the potential for the current draw through a chip to be at almost any frequency from DC to above the clock frequency. This means that unless the precise current spectrum from the chip is well known, for all the possible microcode that could be running through it, we have to assume the peak current could be anywhere from DC to the bandwidth of the signals.

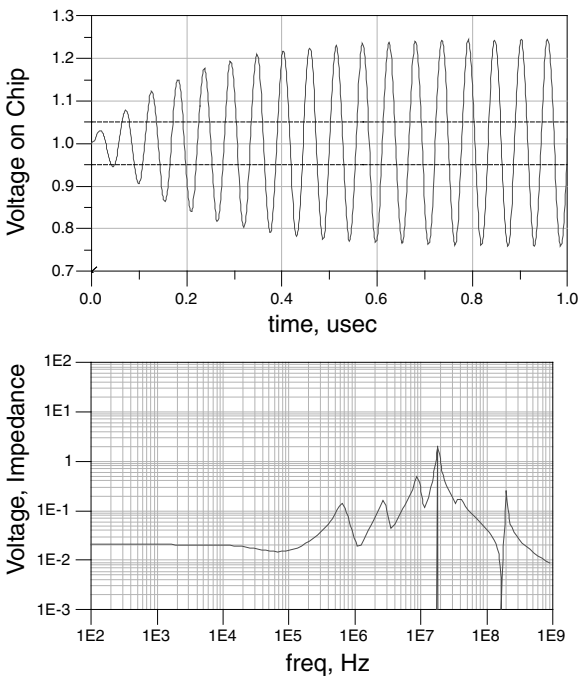

Figure 13-5 Voltage noise on the chip pads (top) as the sine wave current flows through the PDN impedance profile of simulated PDN (bottom). The spike in the PDN profile at about 20 MHz shows where the sine wave frequency component of the current is with respect to the PDN impedance peaks.

In a few rare cases, if it is known that the chip processing will have a high frequency fall off of the current draw above some frequency, it may be possible to put some constraints on the current spectrum. This should always be done, whenever possible.

While the current draw from a chip is rarely a pure sine wave, there are always sine wave frequency components in the current. The precise spectrum of the current amplitudes will interact with the impedance profile of the PDN completely independently of each other, but the resulting voltage waves will add together. Sometimes they can add and still meet the ripple spec, while at other times, they can add and exceed the ripple spec, depending on the precise overlap of current peaks and impedance peaks.

Figure 13-6 is an example of a 1 A peak to peak square wave current draw for two slightly different modulation frequencies. The square wave current will have sine wave frequency components at odd multiples of the square wave frequency. Above roughly the fifth harmonic, depending on the rise time, the amplitude of the sine wave harmonics will drop off much faster than 1/f. As the modulation frequency changes, the frequency distribution of the harmonics shifts and interacts differently with the PDN impedance profile.

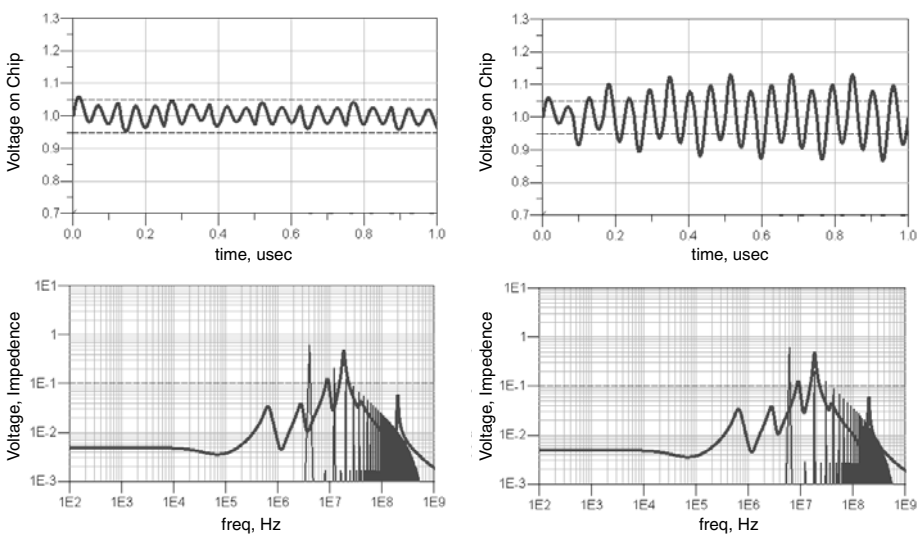

Figure 13-6 The resulting ripple noise when the same 1 A current draw has a slightly different modulation frequency, where one of the harmonic components overlaps an impedance peak of the PDN impedance profile.

A very slight frequency shift in the modulation of the current can mean the difference between acceptable performance and a failure. Unfortunately, the engineer designing the PDN has very little control of the current draw spectrum of the chip. It is whatever the chip is going to do, depending on its operation.

This means that unless there is precise information about the specific, worst case spectrum of the current draw of the chip, a conservative design must assume a worst case current that could occur at any frequency from DC to the bandwidth of the clock, which is a few times the clock frequency.

In practice, it is not the peak current but the maximum transient current that interacts with the higher frequencies of the PDN. If there is a steady state DC current draw from the chip, the sense lines of the VRM can usually compensate to keep the rail voltage close to the specified voltage value. It's when the current changes from the DC value, either increasing or decreasing, at frequencies above the response frequency of the VRM, that the current will interact with the PDN impedance.

The maximum impedance for the PDN, the target impedance, is established based on the highest impedance that will create a voltage drop still below the acceptable ripple spec. This is given by:

$$Z_{PDN} \times I_{transient} = V_{noise} < V_{dd} \times \text{ripple}\% \tag{13-4}$$

or

$$Z_{target} < \frac{V_{dd} \times \text{ripple}\%}{I_{transient}} \tag{13-5}$$

where:

V_{dd} = the supply voltage for a specific rail

$I_{transient}$ = is the worst case transient current

Z_{PDN} = the impedance of the PDN at some frequency

Z_{target} = the target impedance, the maximum allowable impedance of the PDN

V_{noise} = the worst case noise on the PDN

ripple% = the ripple allowed, assumed to be +/– 5% in this example

TIP The optimum PDN impedance should be below the target impedance, but not too far below the target impedance.

If the PDN impedance is kept below the target impedance at every frequency, the worst case voltage noise generated across it as the worst case, maximum transient current flows through it will be less than the ripple spec. If the PDN impedance is much below the target impedance, it means the PDN was overdesigned and costs more than it needs to.

TIP Whenever possible, the peak transient current should be used in estimating the target impedance. When the peak transient current is not available, it can be roughly estimated from the maximum current draw or from the power consumption of the chip.

While the worst case transient current is what is important, rarely is this provided in a spec sheet. Rather, it is the worst case peak current per rail that is on spec sheets. After all, this is important to estimate how large a voltage regulator module is needed. It must be capable of supplying the maximum current draw.

The peak current could be mostly DC with a 10% transient current, or it could be a very low quiescent current with most of the peak current being transient that would only last for a few microseconds. Without special knowledge of the behavior of the chips in each application, the conservative design has to plan for the worst case.

What fraction of the maximum current is transient? Obviously, it depends on the function of the chip. It could vary from 10% to 90% depending on the application. As a rough, general rule of thumb, the transient current can be estimated as being half of the maximum current:

$$I_{transient} \sim \tfrac{1}{2} \times I_{max} \tag{13-6}$$

where:

$I_{transient}$ = the worst case transient current from the chip

I_{max} = the maximum total current from the chip

Alternatively, the worst case power dissipation of the chip will always be provided in a chip's specs, since this is critical information when designing the thermal management approach for the package. It is not usually separated by voltage rail, so some assumptions would have to be made on the power consumption of each rail.

However, given the power consumption per voltage rail, the peak current draw of a chip can be estimated from:

$$I_{peak} = \frac{P_{max}}{V_{dd}} \tag{13-7}$$

And from this, the target impedance can be evaluated as:

$$Z_{target}(f) < \frac{V_{dd} \times ripple\%}{I_{transient}} = 2\frac{V_{dd} \times ripple\% \times V_{dd}}{P_{max}} \tag{13-8}$$

where:

I_{peak} = the worst case peak current in Amps

P_{max} = the worst case power dissipation in watts

V_{dd} = the voltage rail in volts

ripple% = the ripple spec in %

2 = comes from the transient current being ½ the peak current

For example, if the ripple spec is 5%, the target impedance is:

$$Z_{target}(f) = 0.1 \times \frac{V_{dd}^2}{P_{max}} \tag{13-9}$$

For example, in the case of a 1 volt rail and 1 watt power dissipation device, the target impedance would be about 0.1 Ohm:

$$Z_{target}(f) < 0.1 \times \frac{1^2}{1} = 0.1\Omega \quad (13\text{-}10)$$

Some chip vendors, especially FPGA vendors, also provide calculation tools that allow simple estimates of the current draw of specific voltage rails depending on the gate utilization. These can be used to estimate the target impedance specs of the rails. An example of the results of using one such analysis for the Altera Stratix II GX FPGA is shown in Figure 13-7.

Finally, the current requirements of the I/O voltage rails, typically referred to as either the Vcc or Vddq rails, can be estimated based on the number of gates that are switching.

If each output gate drives a transmission line with some characteristic impedance, then the load it sees, if only for a round-trip time, is the same as the characteristic impedance of the line it drives.

If n gates could switch simultaneously, the transient current draw could be:

$$I_{transient} = n\frac{V_{cc}}{Z_0} \quad (13\text{-}11)$$

And the target impedance of the VCC rail would be:

$$Z_{target}(f) < 0.05 \times \frac{1}{n}Z_0 \quad (13\text{-}12)$$

Power rail	Voltage (v)	ripple%	Max current (A)	Transient current amplitude (A)	Z_{target} (Ohms)
VCCT/R	1.2	2.5%	1.2	0.6	0.05
VCCH	1.5	2%	0.17	0.085	0.35
3.3 v Analog	3.3	3%	0.274	0.137	0.72
VCCP	1.2	2%	1.03	0.51	0.047

Figure 13-7 Example of a calculation of the target impedance of different voltage rails based on gate utilization of an Altera FPGA.

where:

$I_{transient}$ = the worst case transient current

n = the number of I/O that could switch simultaneously

V_{cc} = the voltage rail

Z_0 = the characteristic impedance of the transmission lines

For example, if the lines are all 50 Ohms and there are 32 bits switching simultaneously, the target impedance for the Vcc rail would be:

$$Z_{target}(f) < 0.05 \times \frac{1}{32} 50 = 0.08\Omega \quad (13\text{-}13)$$

Even with the peak current and the target impedance established, the current could be fluctuating at almost any frequency, due to the specific microcode or application running. This means that unless there is information to the contrary, it must be assumed this is the target impedance, flat from DC to very high frequency.

TIP The goal in designing the PDN is to keep the impedance of the PDN interconnects below this target value over a very wide bandwidth. A PDN above the target impedance may result in excessive ripple. A PDN impedance much below the target impedance may be overdesigned and more expensive than it needs to be.

When the impedance profile is kept below the target value, the worst case voltage rail noise will be less than the ripple spec. An example of a successful impedance profile is shown in Figure 13-8.

However, if there is a peak in the impedance profile that exceeds the target spec, and if the worst case current happens to fall on top of this peak impedance, there is the chance the ripple spec may be exceeded. This is shown in Figure 13-9.

Peak impedances in the PDN impedance profile are an important design feature to watch out for. Many aspects of PDN impedance design, especially the selection of capacitor values, are driven by the desire to reduce the peak impedances in the PDN.

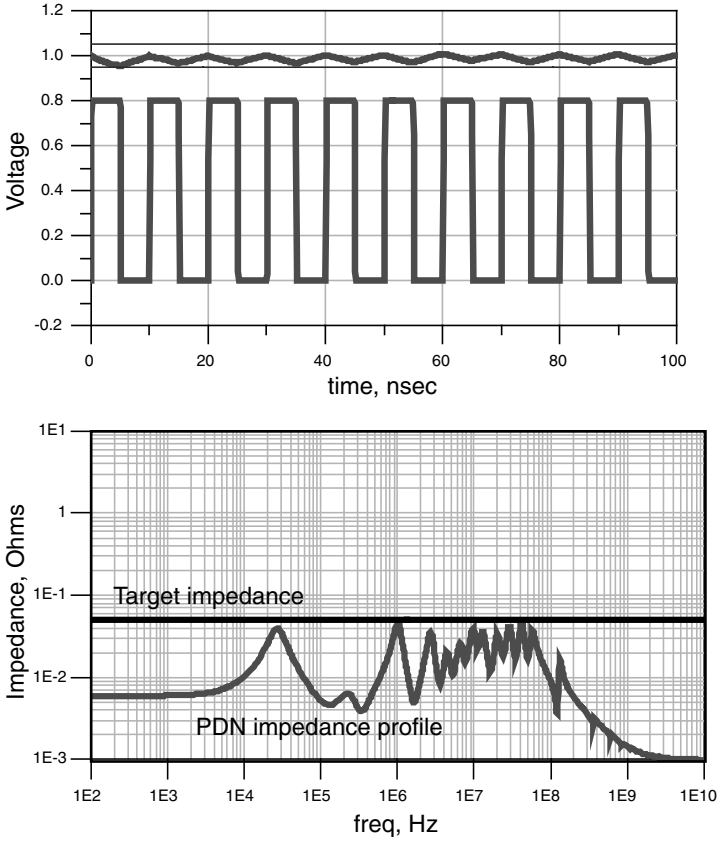

Figure 13-8 When the impedance profile (bottom) is below the target impedance, the worst case voltage noise (top) is below the ripple spec. The square wave is the current draw by the chip, while the flat curve is the voltage on the supply rail.