CHAPTER 1

Engineering the Power Delivery Network

1.1 What Is the Power Delivery Network (PDN) and Why Should I Care?

The power delivery network consists of all the interconnects in the power supply path from the voltage regulator modules (VRMs) to the circuits on the die. Generally, these include the power and ground planes in the boards, cables, connectors, and all the capacitors associated with the power supply. Figure 1.1 is an example of a typical computer board with multiple VRMs and paths delivering the power and ground to the pads of all the active devices.

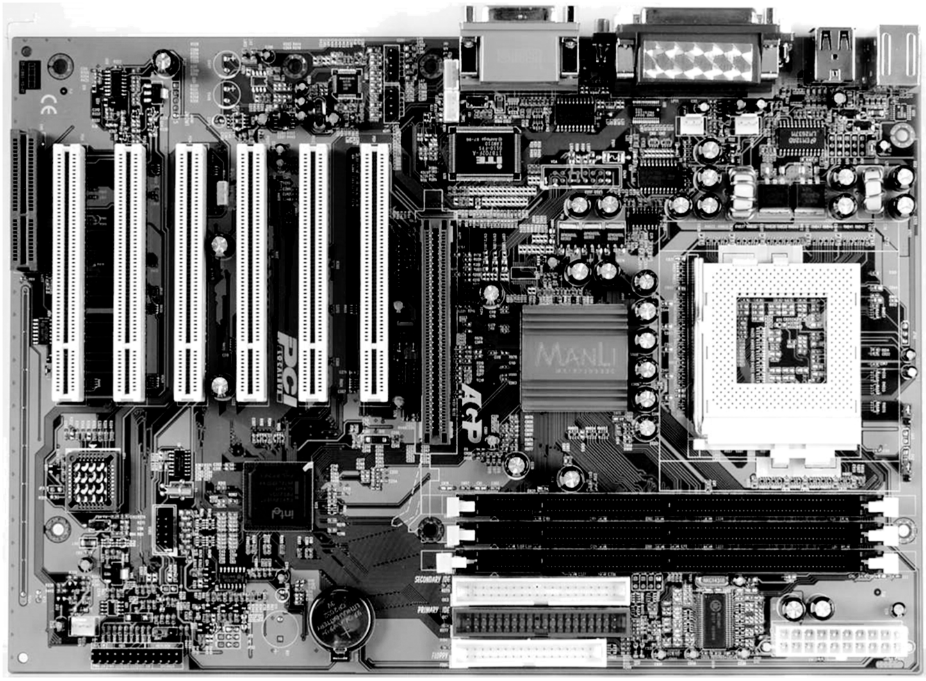

Figure 1.1 A typical computer motherboard with multiple VRMs and active devices. The PDN includes all the interconnects from the pads of the VRMs to the circuits on the die.

The purpose of the PDN is to

- Distribute low-noise DC voltage and power to the active devices doing all the work.
- Provide a low-noise return path for all the signals.
- Mitigate electromagnetic interference (EMI) problems without contributing to radiated emissions.

In this book, we focus on the first role of the PDN: to distribute a DC voltage and power to all the active devices requiring power and to keep the noise below an acceptable level. Unsuccessful noise control on the PDN will contribute to contraction of the eye of any signal. The amplitude of the eye in the vertical direction collapses from voltage noise. The time of the signal crossing a reference spreads out in the horizontal direction creating jitter and reduction of the eye opening. Internal core circuits might suffer setup and hold-time errors, leading to functional failures.

TIP The consequence of not correctly designing the PDN is increased bit error ratios from enhanced vertical noise and jitter on both I/O circuits and internal-to-the-chip circuits. Excessive horizontal noise in core circuits might lead to setup and hold-time violations.

Depending on the circuit of the switching gates, the PDN noise will add to the signal coming from the transmitter (TX). This can also appear as noise on the voltage reference at the receiver (RX). In both cases, the PDN noise will reduce the noise margin available from other sources.

Figure 1.2 shows an example of the measured voltage noise between the core power and ground (Vdd and Vss) rails on a microprocessor die at three different on-die locations and two different voltage rails. In this example, the voltage noise is 125 mV. In many circuits, a large fraction of this voltage noise will appear superimposed on the signal at the RX.

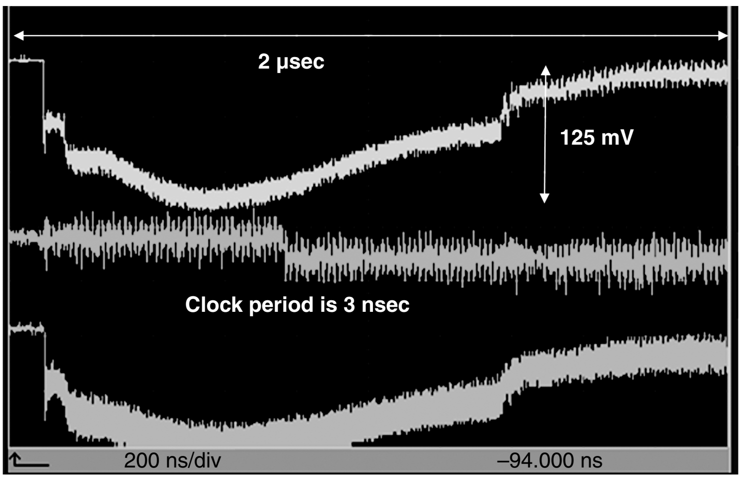

Figure 1.2 Example of the noise between the Vdd and Vss rails in a microprocessor running at 300 MHz clock, measured at three different locations. More than 125 mV of noise is present.

Even if this noise by itself is not enough to cause a bit failure, it will contribute to eye closure, and with the other noise sources might result in a failure.

Voltage noise on the power rails of the chips also affects timing. The propagation delay, the time from which an input voltage transition propagates through the sequence of gates contributing to an output voltage transition, depends on the instantaneous voltage level between the In CMOS technology, the higher the drain-to-source voltage, the larger the electric fields in the channels and the shorter the propagation delay. Likewise, the lower the Vdd to Vss voltage, the longer the propagation delay.

This means that voltage noise on the Vdd to Vss rails on die directly contributes to timing variations in the output signals called jitter. A higher voltage on the Vdd rail "pulls in" a clock edge, whereas a lower rail voltage "pushes out" a clock edge. Figure 1.3 is an example of the measured jitter induced on a high-end FPGA test chip from voltage noise on the PDN.

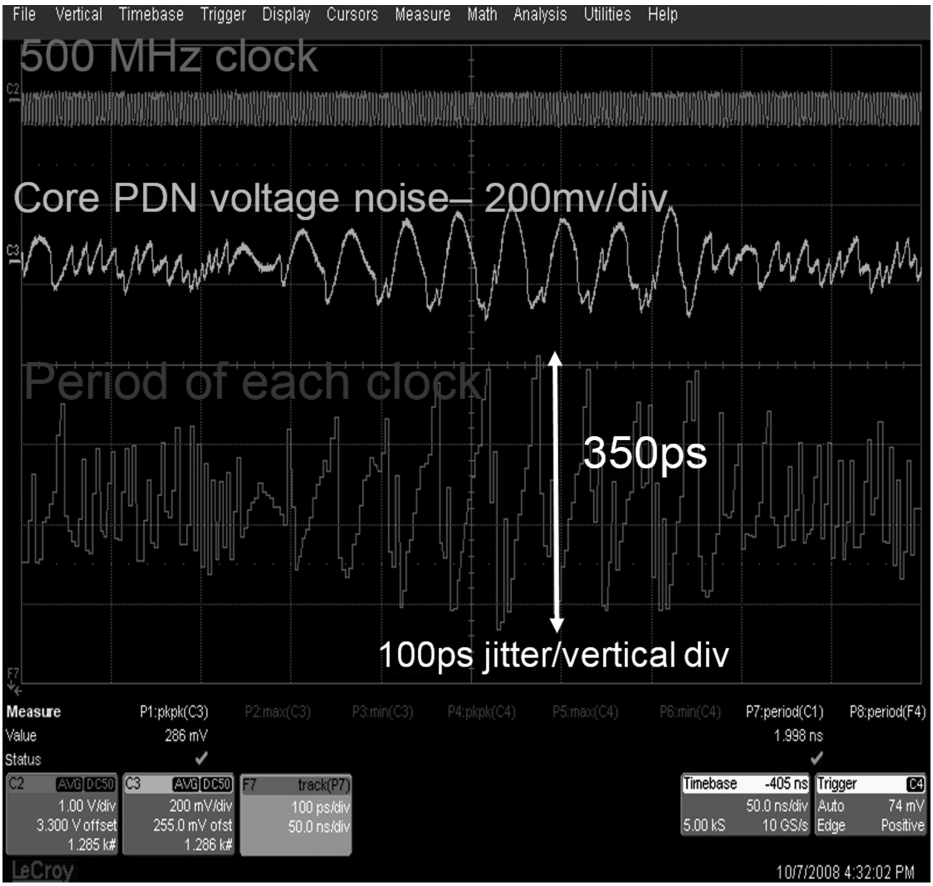

Figure 1.3 Measured jitter on a clock signal in the presence of Vdd to Vss voltage noise.

In this example, a clock distribution net shares the Vdd rail with a number of other gates. These gates were switching with a pseudo-random bit sequence (PRBS), drawing large currents from the PDN and generating large transient voltage noise. This voltage noise, as applied to the clock distribution network gates, caused timing variations in the clock signal. The period jitter measurement, the period of time from one clock edge to the next clock edge, appears as the period of each clock. This measurement demonstrates the direct correlation between the voltage noise on the die and the jitter on the clock.

In this example, the sensitivity of the jitter from PDN noise is about 1 ps of jitter per mV of voltage noise. A 100 mV peak-to-peak PDN noise would contribute to 100 ps peak-to-peak jitter. In a 2 GHz clocked system, the period is only 500 psec. The jitter from the PDN noise alone would consume the entire timing budget.

TIP In this example, the jitter sensitivity to PDN noise is about 1 ps/mV. This is a rough estimate of the sensitivity to expect in many devices.

1.2 Engineering the PDN

To meet both voltage noise and the timing budgets, the voltage noise on the PDN must be kept below some specified value. Depending on the system details, this voltage noise limit is roughly about ±5% of the supply voltage. In typical CMOS-based digital systems with single-ended signals, the total noise margin for the receiver is about 15% of the signal swing. Unless there is a compelling reason not to do so, we usually partition this budget equally between the three dominate sources of noise: reflection noise, crosstalk, and PDN noise. This is the origin of the typical specification being 5% PDN noise allowed.

In some applications, such as analog-to-digital converters (ADCs) or phase locked loops (PLLs), performance is very sensitive to voltage noise and the PDN noise must be kept below 1%. The voltage noise must be kept below the limits from DC all the way up to the bandwidth of the signals, which might be as high as 5 GHz to 10 GHz.

As with all signal integrity problems, the first step in eliminating them is to identify the root cause. At low frequency, the voltage noise across the PDN is usually due to the voltage noise from the VRM and so the first step in PDN design is selecting a VRM with low enough voltage noise under a suitable load current.

However, even with the world's most stable VRM, voltage noise still exists on the pads of the die. This arises from the voltage drop across the impedance of the entire PDN from transient power currents through the gates on the die. Between the pads of the VRM and the pads on the die are all the interconnects associated with the PDN. We refer to this entire network as the *PDN ecology*.

> **TIP** The PDN ecology is the entire series of interconnects from the pads on the die to the pads of the VRM. These all interact to create the impedance profile applied to the die and influence PDN noise.

As applied to the pads of the die, these interconnects contribute to an impedance profile. Figure 1.4 shows a typical example.

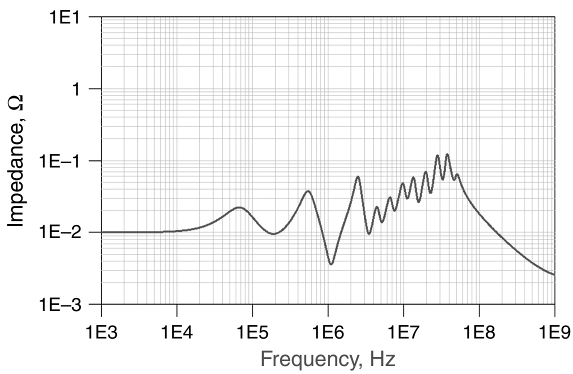

Figure 1.4 Example of an impedance profile of the entire PDN ecology, as applied to the pads of the die.

Any transient currents through this impedance profile generates voltage noise on the pads of the chip, independent of the VRM stability.

For example, Figure 1.5 shows the transient current spectrum drawn by the core power rail for a device when executing a specific microcode. Superimposed on the current spectrum is the impedance profile through which this current flows. The combination of the current amplitude and impedance at each frequency generates a voltage noise spectrum. This noise spectrum, when viewed in the time domain, results in a transient voltage noise.

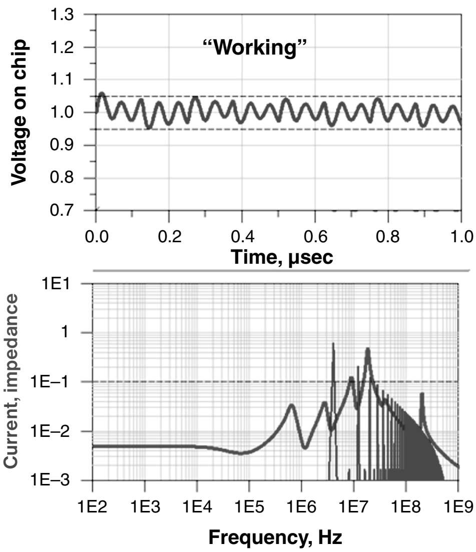

Figure 1.5 ***Left Side:*** PDN impedance profile and transient current spectrum result in acceptable voltage noise. ***Right Side:*** Slight change in current spectrum gives unacceptable voltage noise.

The left side of Figure 1.5 shows the transient current spectrum, PDN impedance profile, and resulting voltage noise on the power rail. This combination of current spectral peaks and impedance peaks results in acceptable noise. On the right is the same impedance profile, but with slightly different microcode algorithm driving the same gates at a slightly different frequency. A current spectral peak ended up overlapping a larger impedance peak and generating a rail voltage noise above the acceptable limit.

The actual voltage noise generated by the transient current through the impedance profile depends on the overlap of the current frequency components and the peaks in the impedance profile. If the voltage noise is below a specified level, PDN induced errors will not occur. If the microcode changes resulting current amplitude peaks and frequency component changes, their overlap with impedance peaks might create more voltage noise and product failure.

TIP The noise on the PDN depends as much on the impedance profile applied to the die as the spectrum of the transient current through the die. Microcode details and gate utilization have a strong impact on the PDN noise generated.

1.3 "Working" or "Robust" PDN Design

The variability in performance due to the specific microcode driving the switching of on-die gates makes testing a product for adequate PDN design difficult. A product might work just fine at boot up, or when running a specific software test suite if the combination of current spectral peaks and impedance peaks results in less than the specified transient noise. The product design may "pass" this test and be stamped as "working."

However, if another software suite were to run that drives more gates and causes them to switch at a different dominant loop frequency, which coincidently overlaps a peak in the PDN impedance profile, larger instantaneous voltage drops might result and the same product could fail.

Although having the product boot up, run a test suite and apparently work is encouraging, it does not guarantee "robust" operation. Products often "work" in evaluation but have field failures when driven by a broad range of customer software.

A robust PDN design means that any software code may run and generate the maximum transient current at any arbitrary frequency with any time domain signature. The resulting worst-case voltage generated by this current through the impedance profile is always less than an amount that would cause a failure.

The combination of the worst-case transient current and the voltage noise specification work together to set a limit for the maximum allowable PDN impedance such that the voltage noise will never exceed the specification.

This maximum allowable PDN impedance with guaranteed performance is referred to as the *target impedance* in PDN design, and we derive it with [1]

$$Z_{target} = \frac{\Delta V_{noise}}{I_{max-transient}} \tag{1.1}$$

where

Z_{target} = the maximum allowable PDN impedance at any frequency

ΔV_{noise} = the maximum specified voltage rail noise to meet performance requirements

$I_{max\text{-}transient}$ = the worst-case transient current under any possible operation

For example, if the noise spec is set as ±50 mV and the worst-case transient current is 1 A, the target impedance is

$$Z_{\text{target}} = \frac{\Delta V_{\text{noise}}}{I_{\text{max-transient}}} = \frac{0.05\text{V}}{1\text{A}} = 50\ \text{m}\Omega \tag{1.2}$$

If either ΔV_{noise} or $I_{\text{max-transient}}$ is a function of frequency, then Z_{target} is a function of frequency.

In principle, the combination of the entire spectral distribution of currents and the entire impedance profile is what creates the worst-case peak voltage noise. Unfortunately, this can only be determined with a transient simulation including the details of the transient current waveform and the impedance profile of the entire PDN. In practice, the target impedance is a useful approximation as a figure of merit to help focus the design of the PDN on a good starting place.

TIP The target impedance is a useful figure of merit for the PDN. It is a good approximation of a design goal for a robust PDN design. The final evaluation of robust PDN design would come from a transient simulation of the entire PDN and the transient current waveforms.

A fully robust PDN is defined by this target impedance. If the impedance of the entire PDN ecology, as applied to the pads of the die, is below the target impedance at all frequencies, the maximum worst-case rail collapse noise generated by the transient current flowing through the PDN impedance will not exceed the noise spec except in a very rare rogue wave situation. Figure 1.6 shows an example of the impedance profile below the target impedance of 50 mΩ at all frequencies and an example of the resulting rail voltage noise with a high current load.

TIP The target impedance is the most important metric when evaluating PDN performance. The farther the PDN impedance is above the target impedance, the greater the risk of a failure.

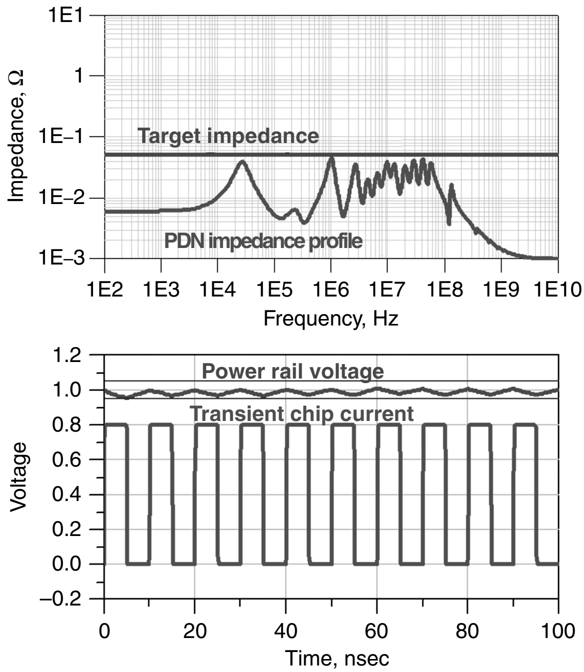

Figure 1.6 ***Top:*** The impedance profile of the PDN ecology engineered to be below the target impedance from DC up to a very high bandwidth. ***Bottom:*** The resulting Vdd rail noise under large transient current load showing the noise is always below the 5% spec limit. The square wave trace is the transient current as driven by a clock. It is plotted on a relative scale.

In practice, the maximum, worst-case transient current through the die will not be flat at all frequencies. The maximum current amplitude generally drops off at the high-frequency end, related to how quickly the maximum number of switching gates can be turned on. The precise details depend on the chip architecture, the number of bits in the pipeline, and the nature of the microcode. The effective rise time could be from the rise time of the clock edge to 100 clock cycles.

For example, if the clock frequency is 2 GHz, with a 0.5 ns clock period, and the maximum number of switching gates requires 20 cycles to build up, the shortest rise time for the turn on of the worst-case transient current would be 0.5 ns × 20 cycles = 10 ns. The amplitude of the maximum transient current frequency components will begin to roll off above about 0.35/10 ns = 35 MHz.

Above 35 MHz, the worst-case transient current spectrum would drop off at −20 dB/decade and the resulting target impedance would increase with frequency. The target impedance, in this example, assuming a 50 mV rail voltage noise spec and worst-case current amplitude of 1 A, is shown in Figure 1.7.

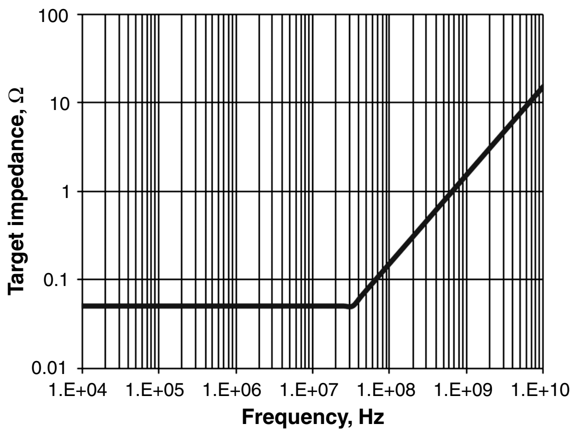

Figure 1.7 Target impedance when the transient current turns on in 20 clock cycles to a maximum of 1 A.

The consequence of this behavior is that the target impedance spec is relaxed at higher frequency. Estimating where this knee frequency begins is often difficult unless we know the details of the transient current and worst-case microcode.

This analysis points out that, in practice, accurately calculating the transient currents and the precise requirements for the target impedance of the PDN is extremely difficult. One must always apply engineering judgment in translating the information available into the requirements for a cost-effective design.

The process to engineer the PDN is to

- Establish a best guess for the target impedance based on what is known about the functioning and applications of the chips.
- Make engineering decisions to try to meet this impedance profile where possible.
- Balance the trade-offs between the cost of implementing the PDN impedance compared to the target impedance, and the risk of a field failure.

A rough measure of the risk of a failure of circuits to run at rated performance is the ratio of the actual PDN impedance to the target impedance, termed the *PDN ratio*:

$$\text{PDN ratio} = \frac{\text{Actual PDN Impedance}}{\text{Target Impedance}} \tag{1.3}$$

A ratio of less than 1 indicates low risk of a PDN-related failure. As this ratio increases, the risk increases as well. From practical experience, a ratio of 2 might still offer an acceptable risk, but a ratio of 10 will almost surely result in unacceptable risk. Even though many microcodes run at rated performance, some are likely to stimulate the PDN resonance and generate product stability issues.

Generally, achieving a lower impedance PDN, and consequently a lower risk ratio, costs more either due to more components required, tighter assembly design rules impacting yield, more layers in the board or package, increased area for die capacitance, or the use of more expensive materials. The balance between cost and risk is often a question of how much risk you are comfortable with. By paying more for added design margin, you can always "buy insurance" and reduce the risk. This is the fundamental trade-off in PDN design.

TIP An important metric of risk in PDN design is the PDN ratio, which is the ratio of the peak impedance to the target impedance. A PDN ratio of 2 or lower is a low risk whereas a PDN ratio of 10 or more is a high risk.

In consumer applications, often strongly cost driven, engineering for a higher risk ratio with a lower cost design might be a better balance. However, in avionic systems, for example, paying extra for a risk ratio less than 1 might be the cost-effective solution. Different applications have a different balance between cost and risk ratio.

1.4 Sculpting the PDN Impedance Profile

The goal in PDN design is to engineer an acceptable impedance profile from DC to the highest frequency component of any power rail currents. All the elements of the PDN should be engineered together to sculpt the impedance profile of the entire ecology. Although many elements interact, assigning some features of the PDN impedance profile to specific features in the PDN design is possible.

Figure 1.8 shows a simplified schematic of the entire PDN ecology. This includes the on-die capacitance, the possibility of on-package capacitors, the package lead inductance, the circuit board vias, the power and ground planes in the circuit board, decoupling capacitor, bulk capacitors, and VRM.

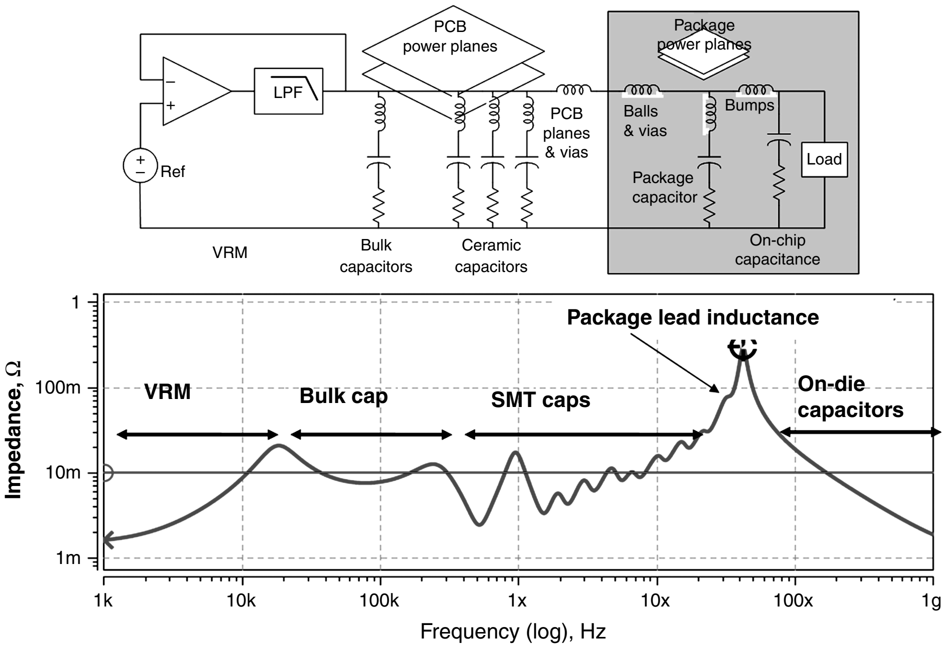

Figure 1.8 ***Top:*** Simplified schematic of the PDN ecology showing the major elements. ***Bottom:*** Resulting impedance profile identifying how specific design features contribute to specific impedance features. On the horizontal scale “x” is MHz.

Isolating functions of some PDN elements enables us to optimize parts of the PDN independent of the others, as long as we always pay attention to the interfaces where the impedance of one element interacts with the impedance of another. This is why so much of PDN design is about the interfaces between the parts.

In the journey ahead, we explore each of these elements that make up the PDN and how they interact to result in a robust and cost-effective PDN design. Ultimately, the power integrity engineer is responsible for finding an acceptable balance between cost, risk, performance, and schedule. The more we know about the details of the specific PDN elements, the more quickly we can reach an acceptable solution.

1.5 The Bottom Line

1. The PDN consists of all the interconnects from the pads on the die to the VRM and all of the components in between.
2. The purpose of the PDN is to provide a clean, low-noise voltage and ground supply to the devices and a low impedance return path for signals, and to mitigate EMC problems.
3. The typical noise spec on the PDN of 5% tolerance is based on an allocation of 1/3 the noise budget to each of the main sources of noise: reflection noise, cross talk, and PDN.
4. Voltage noise on the PDN is a result of transient power currents passing through the impedance of the PDN. The amount of noise is due to the combination of the impedance profile and the transient current spectrum.
5. Noise on the PDN can contribute to jitter. A typical value of the sensitivity is 1 psec/mV of noise. This number varies depending on the chip design and device technology node.
6. The impedance profile, as applied to the chip pads, is the most important metric for the quality and performance of the PDN. This is from DC to the highest frequency components of the switching signals.
7. The target impedance is a measure of the maximum impedance, which will keep the worst-case voltage noise below the acceptable spec.
8. The PDN ratio is the ratio of the actual PDN peak impedance to the target impedance. It is a good metric of risk. A PDN ratio greater than 10 is a high-risk design.
9. Sculpting the impedance profile requires optimizing both the individual elements of the PDN and their interactions. The entire PDN ecology must be optimized to reduce the peak values.
10. If you care about PDN design, this book is for you.

REFERENCE

[1] L. D. Smith, R. E. Anderson, D. W. Forehand, T. J. Pelc, and T. Roy, "Power distribution system design methodology and capacitor selection for modern CMOS technology," *IEEE Trans. Adv. Packag.*, vol. 22, no. 3, pp. 284–291, 1999.