1

INTRODUCTION TO POWER INTEGRITY

This book examines the design concepts of power delivery to modern microprocessors and other related high-speed silicon devices. Today this field is termed power integrity. This chapter provides the background information on what has driven the need for platform power integrity analysis in this relatively new field. The platform is essentially the computer board with its multiple silicon devices, in addition to the power sources, or *converters*, that power them. The subject of power conversion will be examined, in particular as it applies to areas relevant to power integrity engineering. For computer systems the power conversion is mainly in the DC–DC conversion area. The chapters that follow will discuss areas relevant to power integrity analysis—circuit and field theory, modeling, the power delivery network (PDN) and boundary analysis, and other system considerations—and end with an examination of system noise, loadline, and measurement techniques. The last chapter will introduce silicon power integrity, along with some advanced interrelated topics, because of the increasing interest now being given to silicon-level power and the problems associated with on-silicon and on-package power delivery.

In the present chapter, power integrity is defined in terms of the paths that make up the PDN. The paths and all their components comprise the PI (power integrity) domain. A historical review of the voltage and current changes over time (using the microprocessor as an example) is provided to show how silicon has been one of the driving forces behind the need for such fundamental power integrity analyses today. The concept of *first principles* is discussed, because utilizing known equations and

Fundamentals of Power Integrity for Computer Platforms and Systems, First Edition. J. Ted DiBene II.

circuit analysis helps one gain insight into complex problems prior to embarking on sophisticated modeling with numerical tools. A discussion of the limitations and boundaries in power distribution analysis follows covering the circuit limitations (noise sensitivity and silicon process technology) of many advanced devices today that can compromise the accuracy of results.

1.1 DEFINITION FOR POWER INTEGRITY

Power integrity as a field of study includes *power conversion*, *power distribution analysis*, *circuit analysis*, and often the *package/board/silicon system analysis*. But PI is not limited to these subjects. The PI engineer should also be versed in thermal and mechanical basics because some problems needing to be solved may include these and components of other disciplines that impact the system under study. A simple, but somewhat limiting definition is:

Power Integrity *The study of the efficacy of the power delivered from the source to the load within an electronic system.*

Today, power integrity engineers versed in other disciplines may need to consider in their analyses the *system's source*, *load*, and *path*. In the past power integrity engineers often excluded the source and load parts of a system. This is understandable because many power integrity problems focused only on the power distribution path. However, today, having knowledge of both the silicon load *and* the power source allows PI engineers to comprehend fully the complexities of the problems that they face. Conversely, many power conversion engineers are required to cross over into the power integrity domain in order to solve their domain's problems. It is therefore reasonable for engineers from both disciplines to move regularly into each other's domains in order to solve their problems satisfactorily.

As Figure 1.1 shows, the primary power source is the power converter. This is a type of DC–DC converter at the motherboard of the server, forming an inter computer platform. The power source to this converter is typically neglected in an analysis. The power converter includes a certain amount of *decoupling* for filtering and charge storage. In the middle of the figure is the PDN, or *power distribution*

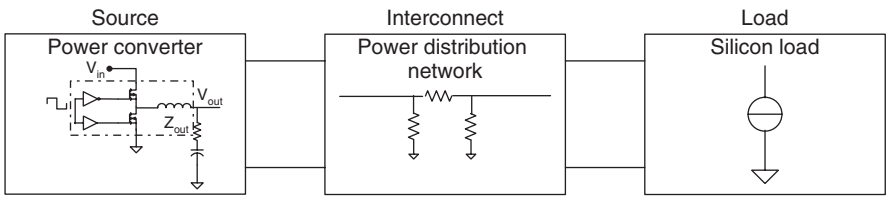

Figure 1.1 Power integrity domain and Scope of influence

network. The PDN typically comprises passive elements from the printed circuit board all the way to the level of the silicon. All of the decoupling and interconnections are included, from the output of the voltage regulator to the load. The printed circuit board and/or design package was where the PI engineer focused in the past. With the more complex recent systems, the PI engineer must often optimize the performance of the *silicon*-level passive and active components, so the circuit load must be considered as well. Note that Figure 1.1 gives a schematic representation of circuit load behavior. This is because modeling the actual behavior of load transistors under all possible conditions is virtually impossible. Nonetheless, knowledge of load behavior is required for PI engineers to do their jobs, and PI engineers must work closely with silicon development teams to gather the data necessary to perform their analyses.

For the PI engineer, at the start of a study, there are many complementary components to consider. Often these are thermal and even mechanical issues that contribute to delivery problems. It is then up to the PI engineer to rework the filter structures in the PDN and, together with the board and silicon teams, to ensure that the power delivery path is performing efficiently. Such analysis requires knowledge of the board's layout, its components, the power source, the load characteristics, the design package, noise coupling to other planes, EMI (electromagnetic interference) issues, and other items that may contribute (potentially) to the results of PI modeling. The assumptions that go into the analysis are a critical part of a PI engineer's responsibilities and the next chapters will explore in detail these assumptions.

1.2 HISTORICAL PERSPECTIVE ON POWER INTEGRITY DRIVERS

The idea of analyzing the power distribution path is not novel. Engineers have been working with the concept of measuring voltages and currents on power lines since the 1920s [1]. However, the need for advanced power integrity techniques in the computer was not realized until recently. The transition from virtually *no* power integrity analysis needed to its being *required* on virtually every platform developed today has been more dramatic than many technologists could have realized. Though noise, EMI, and signal fidelity have traditionally been areas of focus for the system designer, the need for advanced power integrity analysis, relative to the advent of the microprocessor, is still a recent event. Many conferences today are dedicating significant blocks of time to the multitude of papers written on the subject in just the past few years. The main factors behind this trend are a culmination of system metrics: the need for more stringent voltage and current requirements, the increase in voltage rail proliferation (internal and external to the silicon), a dramatic increase in platform and silicon signal densities, and device and platform cost pressures, to name just a few. As evidenced by the previous issues, many of these developments are clearly platform dependent. The issues though vary across each platform type. For example, voltage proliferation and current and voltage requirements may be the main issue on a server platform, whereas cost will typically dominate in a desktop, laptop, or tablet today. Nonetheless, the problems are very similar between them

and the migration toward a deeper understanding of the state of the art is clearly needed today.

One very dramatic change that has occurred over the past two decades is in silicon power requirements. The relatively fast decline in source voltages and the sudden increase in currents delivered to silicon over this time has necessitated a stringent examination of the power delivery path—particularly for complex devices, such as microprocessors [3].[1] Figure 1.2 shows the voltage decline over time for microprocessor devices for the past 20 years. The changes in supply voltage have come about for a number of reasons. First, as complexity and transistor density has increased, so has the power required for the device. This has necessitated that the voltage be dropped to help reduce the overall power to the device. Second, as the silicon process geometries have shrunk, so has the requirement to reduce the voltages to the devices to prevent their damage. Thus manufacturing constraints and device physics have also driven the need to reduce the voltage to these devices as much as any factor. This has in turn driven the power converter industry to follow suit to supply suitable power to these complex devices [2]. The simple graph in Figure 1.2, however, may be a bit misleading. The voltages to these devices have *indeed* been reduced over the past decade. However, today, many high performance silicon devices and the *system on a chip* (SoC) devices are highly integrated, meaning they have other functional units in them besides the processing units. Thus this graph represents—for the silicon parts in more recent years—merely the voltage trends for the processing *elements*. For example, many processors have multiple processing *cores*, IO, memory, graphics, and power management units in them as well. The implication here is that most high-performance silicon chips now require *more* than one supply voltage to power these different units, and most of the supply rails for these units require *different* voltages.

The other important development to note is the *lower* core voltages in use today compared to just a mere 20 years ago. This change alone has led the rise toward power integrity analysis. Today, most core voltages are at or below one volt.

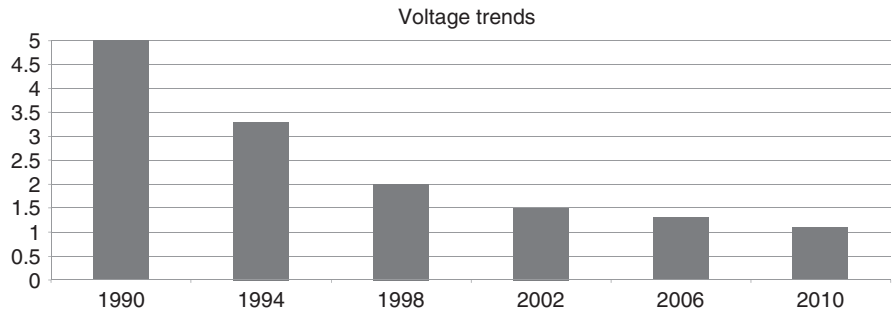

Figure 1.2 Microprocessor voltage trends over the past 20-year period

[1]The trend is not limited to microprocessors. Many other high-performance devices had also been operating at these voltage levels at this time.

This means that even the smallest noise voltage on the power bus may influence functional operation and cause data corruption. In 1990, the 5-V power rail to the processor (or the high-performance computing device) was typically specified at +/−5%, or +/−250 mv. Over time the tolerance for this power source has decreased to a voltage setpoint of only +/−1% today. Most complex devices require this tight regulation to ensure proper frequency and efficacy of operation. If a supply voltage were varied over the 500-mv range (+/−250 mv), for a 1V rail the voltage tolerance would be today +/−25%! Clearly, it is highly unlikely that the device would even function under such conditions. Thus, the noise margins (by necessity) have been required to shrink along with the voltage source to ensure the correct operation of the processor. This has made the job of the power integrity engineer that much more important.

An interesting change was also made in *current* over this same period; see Figure 1.3. The increase in current is graphed in a *log scale* to show the progression of growth. Notice that the *peak* currents[2] have increased by more than *two orders of magnitude* over just one decade.[3] So dynamic was this change that independent groups and even whole *companies* were created to help with the high current distribution problems that affected the motherboards and processing systems.[4] This change *alone* has necessitated the need for more expertise in power system and power integrity analysis.

Figure 1.4 gives a pictorial representation of this evolution. In 1990, the typical personal computer had noncolor screen (though color screens were coming into use at that time), a large ATX box that sat under the desk, and a large dot-matrix printer. Wireless communications in the office were either not available or

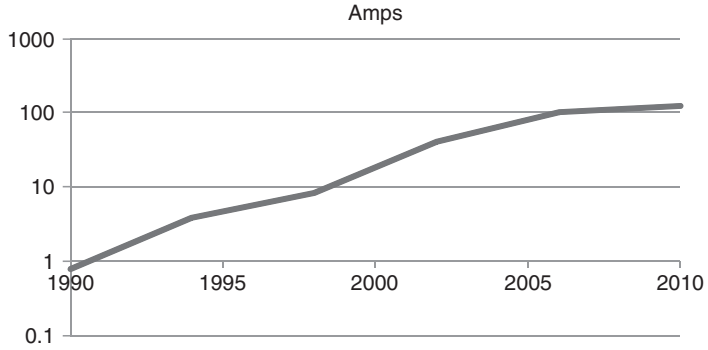

Figure 1.3 Peak current trends for high-performance silicon over the past 20-year period

[2]This trend does not represent the continuous, or steady-state current changes. These are typically lower than the peak currents.

[3]The current trend data are highly dependent on specific devices and are not representative of every device. The trends are also independent of manufacturer and considered generic.

[4]Many of the companies are focused on delivering CAD software for analysis purpose. A simple search on the internet on the subject of Power Integrity reveals a number of these.

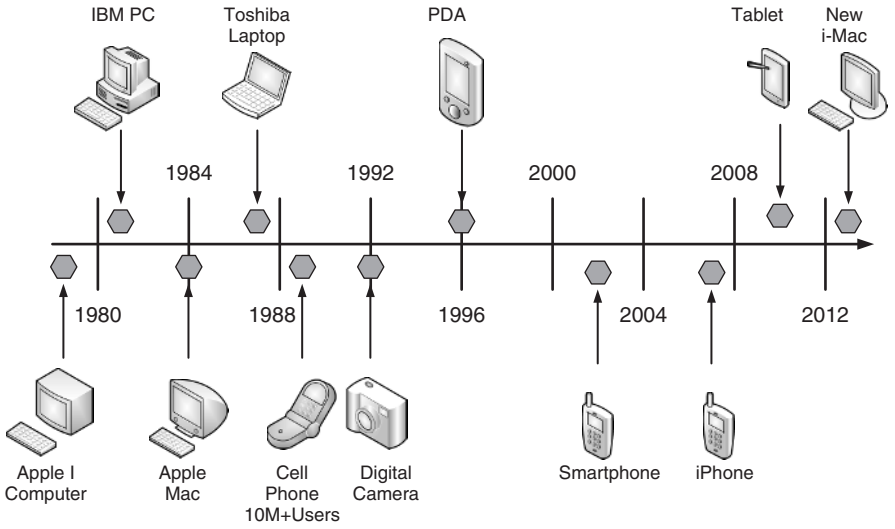

Figure 1.4 Evolution of modern computers

just starting. Today, many people use a tablet or smartphone to communicate in their personal and professional lives. The Internet has also revolutionized the way people use electronic systems.

What is *not* shown in the Figure 1.4 is the evolution of the *silicon* over the same period of time. *Moore's Law* has continued to hold true over this time window as well, as the number of transistors has doubled in computing hardware every 18 months. No doubt, as transistor density increases, so will the need to power these transistors with high-quality power.

The merge of communications and computers that has occurred over the past 20 years has changed the way people live their daily lives. This change has also motivated the development of power integrity as a separate field of study. While, in many ways, the field of power integrity is very new, it is expected to grow rigorously in the coming years to meet advances in silicon and platform power delivery. And, as technology progresses, so will the demand for well-trained power integrity engineers.

1.3 FIRST PRINCIPLES ANALYSIS

As computers have evolved, so has the software that runs on computer systems. For the developer or engineer, the modeling requirements have become ever more sophisticated with more options and methods to analyze complex problems. CAD tools that utilize SPICE or numerical solvers with advanced models to solve for electromagnetic fields have grown substantially over the past two decades. In fact,

today, more engineers have become heavily reliant upon these tools to resolve PI problems.

Moreover, in today's fast-paced world of research and development, engineers are faced with solving complex circuit problems that require a strong background in other fields. That is to say, despite the fact that more sophisticated tools are now at their disposal, engineers cannot bypass this deeper learning and dive headlong into a simulation effort with the goal, inevitably, to deliver a solution that is satisfactory and meets the required deadline. The result is often error in the analytical process—or a less than optimal solution—and more lost time in the end to find a more adequate solution based on the first simulation. Those engineers coming out of school who lean too heavily upon the available tools and less upon the basics to solve difficult problems soon become frustrated.

However, the blame is not with the CAD companies and the ease in which their tools generate data. It is that professional engineers often lack preparation in analytical skills.

Not surprisingly, there is a good solution to this problem. In most cases, seemingly overly complex problems can be solved by understanding the *first principles* behind such problems before embarking on a detailed model and simulation effort. This does not mean that the final solution to the problem will be any simpler than the numerical modeling effort; it just means that the *estimates* and *boundaries* placed upon the problem (initially) will lead to a deeper understanding of the problem and thus a more efficient path to the solution.

1.3.1 Steps to Solve Power Distribution Problems

The *first principles analysis* is simply a way of using basic equations and physics (tools engineers have encounterd in their education) to describe a problem, and then using those tools to solve the problem. Usually the mathematical representations used to describe the problem are close *approximations* to the exact equations. For this reason, it is important to have the understanding that comes with the mathematics. Indeed, many complex problems can be solved quickly, using pen and paper, a simple spreadsheet, or a math program, and thus give the engineer quick insight into the boundaries of the solution. *First principles* will allow engineers to solve a related problem in order to gain insight into the solution of the problem of interest. While the results are typically an approximation to the actual solution, they are satisfactory for the initial analysis.

Invariably, the key for gaining a reasonable result to a complex problem is twofold: first, the assumptions must be within the bounds of physical behavior, and second, the limitations of the physical and mathematical models must be understood up front. The steps in solving a problem using the *first principles method* are straightforward and have been used by engineers in schools and in their professions for many years, though perhaps not known by this term. The steps outlined below are for solving a simple AC, or frequency-dependent, power integrity distribution problem. However, the same method could be used for any engineering problem.

1. Determine the result you expect to achieve, (Is this an impedance profile? Filter analysis? What is the output you expect to have?) Here, the objective is to capture the AC impedance across a given bandwidth.
2. Determine the relative positions of the components involved in the distribution and their dimensions relative to each other. Typically, this means sketching a plan and/or a cross-sectional view of the problem. For this problem, a cross-section of the various distribution networks may be drawn with placement of the components in the "plan view" of the networks.
3. Create a simple schematic of the problem. If the problem is too cumbersome or large, break it into subsections and analyze each subsection one at a time. Comfine them into larger sections and re-analyze. Here again, draw the network components (R, L, and C's).
4. Fill in the values for the schematic, starting with the known quantities (e.g., the capacitors).
5. Extract the interconnecting values from the basic equations and add these to the schematic to complete it. Here, one would use the tools of the later chapters of this book to extract these values, which are typically approximations to the exact values. Check to determine if each approximation is satisfactory for the desired estimation.
6. Determine, using simple hand calculations, a math program or, using SPICE, the solution, then examine the results.
7. Check the data against the simplifications. Typically, this means shorting or opening parts of the circuit to see if the results make sense. This does not always guarantee the assumptions or analysis is correct, but it is a good first check on the data.

This *first principles* method will work for most problems, such as time-domain analyses, with minimal modifications. It is also possible to combine it with numerical results obtained from data of other programs, published papers, and so forth. In any case, some simplifications will need to be made to link data to the simple model obtained by these from first-principles steps.

These seven steps are essentially a variant on the *scientific method*. So the process is clearly not new. These steps will always cut down the number of simulations and modeling iterations used because they provide the engineer with a strong foundation on where to start in the analysis process and give some insights that can improve on the early assumptions. Later in this book, some further tools that support this procedure will be discussed, along with some illustrative examples. It is, however, assumed that the reader has basic knowledge of the equations used in solving the type of power integrity problems illustrated previously. The rest of this section sets the *foundation* for the analyses used throughout this text.

It should be noted that it is not the intention here to replace detailed numerical modeling with the *first principles* method. The main purpose is to show how a strong analytical foundation can help the power integrity engineer solve problems more efficiently using these additional skill sets.

1.3.2 Limitations in the Analytical and Numerical Process

Numerical modeling is used extensively in power integrity engineering, and sophisticated CAD tools, such as SPICE (circuit simulator), Maxwell (field solver), and math tools, are an essential part of the engineers toolbox. However, recently engineers have been depending too much on these tools without understanding their limitations. Many 3D modeling tools are based on either the *finite element method* or the *finite difference/volume* method. Additionally, excellent advancements have been seen in the *PEEC* (partial equivalent element circuit) method in the past decade. All of these methods have their advantages and disadvantages—which are usually centered around accuracy, convergence, time, and cost. The main problem that they all share is in the assumptions that are made by the engineer doing the modeling. Moreover, this is less of a problem with the tool than it is with the user's ability to correctly utilize the tool. It is beyond the scope of this text to re-visit all the boundary conditions for these methodologies, the reader is encouraged to look over Chapter 3, where some of these methods are discussed. Here, some basic warnings are provided on both utilizing the *first principles method* and on reporting data using these tools.

Currently, there are limitations to what the power integrity engineer can analyze. Much of this is due to the modeling constraints—such as in the software and accuracy of the assumptions—that go into the analyses. As systems have become more sophisticated, so has the need to have better and more accurate sets of tools. When it comes to modeling a problem, the PI engineer is mostly interested in the physical representation of the circuit structure and thus the issue becomes how to represent accurately the complex electromagnetic behavior with simple passive elements. For most problems, a simple circuit approximation to the actual is adequate. For example, extraction of parasitic resistance, conductance, inductance, and capacitance of a circuit may sometimes be done by the *first principles method*, at least at the beginning of the analysis. However, in more complex cases, advanced modeling tools are required to adequately extract the values. This is where the tools, and the assumptions that go into the tools, become important.

Prior to performing a detailed modeling analysis—which may take weeks or even months—it is good practice to assess the boundaries of the accuracy of the effort—and if it is *sufficient* for the given objective. Too often, engineers start with the software and then determine that the assumptions they were working are inaccurate. Figure 1.5 illustrates where (typically) the accuracy coverage of the problem is and at which point in the process flow the accuracy limits will be available. Usually, a *first principles analysis* (point A in figure 1.5) will yield a 50% accurate solution; that is, if the exact solution is 2×, the expected accuracy will be *at worst* between 1× and 3×. This should be the goal for the PI engineer when starting out on a problem, such as the analysis of a power distribution network, to ensure that the accuracy of the results is well understood. In addition, it is surprising how often even a *first principles analysis* is more than sufficient for many engineering challenges. Moreover, if a *first principles analysis* is performed, the engineer should check the accuracy of the equations used to make sure that they are within these

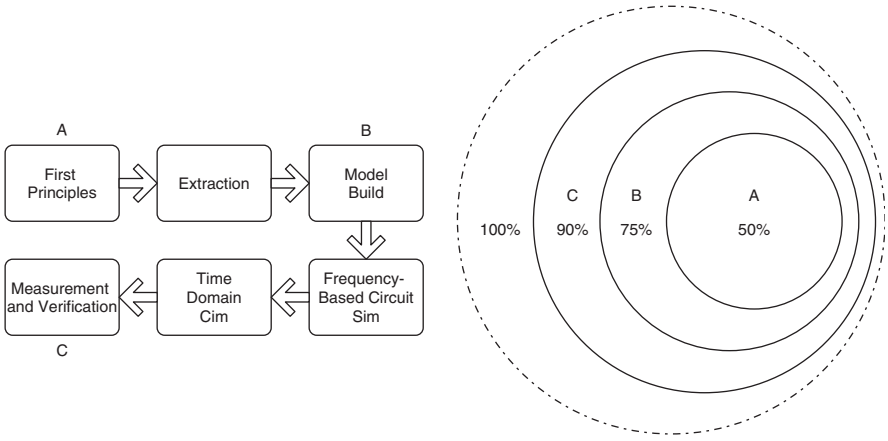

Figure 1.5 Coverage and accuracy chart for power integrity problem solution space

bounds. If the equation is, say, to *estimate* the inductance of an element in the path geometrically, this should be stated up front and the accumulation of the errors statistically should be noted at this point with the other components as well. As one progresses into the model building phase (B), the expected accuracy will typically improve—thus an error of 25% or better would be expected here. After this point in the analysis, it is usually good to assess whether or not more information is needed to make a decision on a particular project. If more accuracy is needed, the engineer will continue to the measurement and verification phase (C) where there is also a *tolerance*, but where an accuracy of at least 10% is sufficient. Next the data are validated to ensure that the load is accurate by using dynamic load tools and other instrumentation. Engineers may jump from one step to the next, sometimes bypassing a step along the way—provided that the data are sufficient from the previous step in the analysis.

Unfortunately, in practice, engineers too often report data based on inaccurate assumptions and analyses. For example, suppose that a power integrity engineer is asked to report on the droop deviation for a time-domain transient analysis in a PDN circuit. It is not unusual for engineering teams to have been part of such a review and to have reported the data in the following manner: the engineer completes the analysis, and a meeting is then called; then the engineer compiles the data in a format (e.g., a slide presentation) in order to share the data with the team. The final result is reported using a number based on SPICE modeling and the 3D extraction of the parasitics:

$$\Delta V = 7.849326 \text{ mV} \tag{1.1}$$

The first question one might ask is: What is wrong with this number? As will be seen shortly, this type of reporting is completely incorrect, but far too often this is how the data are presented. To begin with, based on the sphere in Figure 1.5, dimensionally, this much accuracy is not possible. Cleary one would expect the tolerance to be at

best 25%. The engineer should have performed a simple dimensional analysis of the problem, to begin with. This way, the value has a *range* of accuracy—even if the simulation was perfect! The new number, simply based on the known tolerance of the elements (not even taking into account the other potential errors), would be

$$\Delta V \cong 8 \mp 2 \text{ mV} \tag{1.2}$$

Even if there were no errors in the simulation, the assumptions were perfect, the conditions were perfect, and the modeling was perfect, this would be the engineer's starting point. One might presume that the team would start to ask a number of questions if the data were reported in this manner. Unfortunately, it is common for the team to review the data and *not* question the accuracy, since it came from the tool, of course, and isn't the tool always correct! Many engineers do not even question the data until much later on in the process—often as late as when the lab measurement data comes back—at which time the data are predominately found to be very *different* from the simulations. At this point, gross errors in the results can cost significant money, time, and resources to recover. Invariably, this becomes problematic for the data owner, and those that may require the information. Moreover, reporting the data in this manner lacks certain clarifications, though it may look perfectly good to many on the first review. Missing in the analysis are some very fundamental facts that would lead a reviewer to question the result. With the present model, in a through review one would expect a number of questions to follow;

- How were the individual elements modeled? What were the assumptions in the model for the location of the ground reference for each parasitic?
- What was the tolerance for each parasitic? For example, if it was a plane inductance, what tolerance for the dielectric thickness was used, and what temperature used (for mechanical expansion)? What connections were assumed and what were the boundary conditions of the initial setup in the model?
- What was the mesh size and convergence criteria set to for the CAD program? How many iterations did it take? Did the engineer perform an accuracy assumption prior to running the simulation to determine what accuracy was expected just from the simulation?
- In the SPICE setup, what was the time-domain step size used? Did this person break down the PDN to a small enough lumped element so that step sizes did not matter?
- Was this analysis checked per an existing analysis so that a known baseline from either a measurement or a known simulation run would verify the results?

And the list could go on. These are just some of the questions that could come up in an actual review process. And *here* is where it becomes critical for the power integrity engineer—and for engineers in general—to *ensure* that the proper

preparatory work is done on the problem. It may be worthwhile to put together a *checklist* to prevent errors in delivering such data. A few basic steps would be useful in the modeling and analytical process are:

Step 1. Start with an existing analysis that is similar enough (e.g., a known PDN) and look at the results of this droop event given the constraints. Most other engineers have gone down a similar path and have results that yield an excellent starting point—even so much that their initial models may be used.

Step 2. Review all of the parasitics—putting these in a table and looking at the tolerances of each separately is a good idea. The tolerances may be used for the SPICE analyses later on.

Step 3. Perform a *first principles* analysis on the PDN to make sure that, at least from a lumped element perspective, the data makes sense. Make sure that it is well understood that the accuracy of this first pass is a crude, though a reasonable, approximation to the final result. For example, if the droop result from this analysis is 100 mVpp, this would indicate that the 8 mVpp initial simulation may be suspicious.

Step 4. Set up the 3D modeling tool correctly, and once again, look for similar extraction results on already measured data if possible. Too often the setting up of the 3D modeling assumptions can result in gross errors that strongly influence the final results.

Step 5. Run the simulations based on the tolerances of every device and result in the system. This means that the engineer should run at least two boundary cases in addition to the nominal case. This will give the expected range of results.

These basic steps comprise a general approach to solving just about any problem. However, it is common to find that engineers can miss some basic steps in reporting out the data—and likely some errors in the assumptions as well. In this example a little more digging may reveal that the actual component values and tolerances, behind the result, that went into the engineers assumptions, may be quite different. This is not uncommon. In fact, after applying the basics to the earlier results—which were derived from an *actual* example—the new, re-run analysis resulted in the following data:

$$\Delta V \cong 42 \overline{+} 13 \text{ mV} \tag{1.3}$$

Note here that the following missed steps occurred in the first reporting: the number of dimensions in the initial case, and actual result reported, were misleading and led to the wrong result. Moreover, the assumptions that went into the modeling effort were incorrect. This is why the process of generating data must be reviewed carefully to get a timely and accurate solution space.

As with all analyses, proper data representation and assumptions are not the only problems the engineer must face. For power integrity, in particular, an understanding of the environment affecting the power path is also important. The *boundary conditions* for the PI engineer must be considered prior to embarking on a

complex and detailed analysis. The boundaries are typically focused around the *bandwidth* of the solution space. The bandwidth for solving power integrity problems is highly application and issue dependent, but the range for most problems is between DC and 1 GHz. This is a very wide range. Moreover, many noise-related PI problems have harmonics, and thus extending into this range and beyond must be considered—particularly at the silicon level. One could argue that within the silicon, power integrity frequencies extend well into the radio frequency regime. However, most problems fall into the noted bandwidth.

In summary, the critical action should be for the PI engineer to assess the proper boundaries of the problem prior to the simulation phase. This usually involves performing some basic computations, usually by hand or in a spreadsheet. What this does is set the level of complexity for the next phase of analysis and significantly speed up the process. That is, if the accuracy of the expected result is known *a priori*, then setting up the program to sample at a certain rate and through a given bandwidth will typically reduce the processing time. Understanding the boundaries of the problem will nearly always result in a more efficient and timely solution and will allow the engineer to be a more effective developer in the end. Chapter 4 provides some detailed examples that illustrate this process more thoroughly.

1.4 SCOPE OF THE TEXT

This chapter has provided a discussion of some of the key metrics in computer platforms that have driven power integrity and how the state of the art has evolved into an in-depth independent subject. The subject of boundary conditions was introduced and the limitations in power integrity associated with current modeling tools. However, the frequency bounds of power integrity will increase as we move forward into the next coming decades, since this is a dynamic environment.

Chapter 2 reviews power conversion devices used in computer system platforms. The discussion is necessarily brief, since there are many excellent texts on the subject for the student. The focus is therefore on power conversion as it relates to power integrity, rather than design. The review covers the basics of buck–level power conversion and how linear regulators operate. The discussion includes how to compute power losses and filter operation, how power switches work along with inductor and driver, some controller basics, as well as how current and voltage sensing operate. This introduction provides the baseline for later chapters where loadline and other subjects are discussed. Because the power conversion industry is also making advances in the area of platform power conversion some additional advanced subjects are briefly reviewed, namely coupled inductors along with multi-phase and tapped inductor topologies. The chapter ends with a discussion of platform-level power conversion, which will give the reader a basic understanding of the *source* block of the power integrity domain.

Chapter 3 reviews electromagnetic and circuit basics. Circuit extraction examples are introduced for various key geometries as a basis for the models

used in later chapters. The simple first principle analyses will be valuable to the PI engineer. Some additional areas covered are power plane extraction, which is a type of problem the PI engineer may often encounter, and basic modeling methodologies commonly in use in software programs. To give any one of these subjects justice requires multiple texts, so the reader is referred to references at the end of the chapter for additional reading.

Chapter 4 introduces system PDN analysis. This is the middle domain region of the power path for the PI engineer. Within the power distribution network (PDN) are frequency-dependent impedance network distributions, and these are introduced with an emphasis on using first principles to analyze their structures. Some additional tools are presented, mainly concerning the printed circuit board and package structures, since a frequency domain analysis will change depending on which portion of the network is the focus. To supplement the analysis, connector interconnects will be discussed with an emphasis on socket type connectors, and in particular, contact resistance and how to treat this in the overall PDN. The chapter ends with a detailed discussion of decoupling fundamentals. This is an area where PI engineers lend the most expertise to the system.

Chapter 5 moves into the area of modeling the source and load. The emphasis is on the load, since much of the source was covered in Chapter 2. The focus is on the basic methods for modeling dynamic sources. Mastering this chapter's content will enable the engineer to generate a sufficient representation of the full path for time-dependent modeling, centered on the voltage bus droop analysis. This is likely the most important metric in power integrity analysis. Both closed loop and open loop methods are discussed and make use of modeling the buck regulator in SPICE. After the droop analysis distribution path losses are revisited for the entire link path in a discussion of bounding the overall power droop budget. Last, an *impedance boundary analysis* is presented that includes a frequency dependent representation to allow for a more complete mathematical view of the voltage droop function.

Up to this point in the book, the main elements of the power integrity domain have been covered. The focus then shifts to system considerations in Chapter 6. The important subject of the *power loadline* is introduced. The power loadline involves both the distribution network and the power converter. Next, signal noise and power noise on the power planes with the emphasis of coupling and how it is generated are covered. The next areas discussed are noise and signal representations in the PCB, at the package and silicon levels, with the focus on the different techniques that can reduce noise coupling in planes and improve power integrity distribution paths. Chapter 6 closes with a discussion of different PI measurement techniques.

Chapter 7 provides an introduction to silicon-and package-level power distribution and design futures. The discussion includes an analysis of the power delivery path inside of a silicon device and of the layout and routing issues related to the PDN. Relative frequency dependencies are discussed along with relative dimensions in devices such as advanced processors and SoC devices.

REFERENCES

1. Carroll, J. S., Peterson, T. F., Stray, R. Power measurements at high voltages and low power factors. *Trans. American Institute of Electrical Engineers*, 1924.
2. Stanford, E. Microprocessor voltage regulators and power supply trends and device requirements. *Proc. Int. Symp. Power Semiconductor Devices & IC's*, 2004.
3. Kwangok, J. A power-constrained MPU roadmap for the international Technology Roadmap for semiconductors (ITRS). IEEE SoC Design Conf. (ISOCC), 2009.
4. Singh, D., Rajgopal, S., Mehta, G., Patel, R. Technology Roadmap for Semiconductors. *IEEE Journals and Magazines*, 2001.
5. Tiwari, V et al. *Reducing Power in High-Performance Microprocessors*. Santa Clara, CA: Design Automation Conf; 1999.
6. Sullivan, D. M. *Electromagnetic Simulation Using the FDTD Method*. IEEE Press, 2000.
7. Reddy, J. N. *An Introduction to the Finite Element Method*, 2nd ed., McGraw-Hill, 1993.