Preface

My first assignment as a signal integrity engineer was to determine the bond wire signal-to-return ratio for a group of ASICs that would go into a supercomputer IO subsystem. This seemed like a daunting task at the time (it still does) because the answer to the question seemed to depend on so many other things: receiver input thresholds, crosstalk from several sources, timing between various noise events, and clock arrival at the receiving flip-flop, to name just a few. I asked the company librarian to help me with a literature search and found that most of the good papers on the topic of margins came from a small group of engineers who worked for IBM Corp. and Digital Equipment Corp. Little did I know at the time that I would have the good fortune to work with some of these engineers later in my career. Now it's my turn to spice up this well-established body of knowledge with some contemporary examples flavored by my own unique experiences. I hope this book provides insights that will help you make the tough decisions you face in your daily work.

We can talk about intersymbol interference and simultaneous switching noise until the cows come home, but the million-dollar question has always been, "How much intersymbol interference will an interface tolerate before the system crashes?" (Insert your own favorite effect in place of intersymbol interference.)

Early in my career, I looked for a neat formula that would answer this question for me, but I gradually came to realize that the answers were as unique as each new design, each engineer's mind, and even a company's corporate DNA. Perhaps the ambiguity explains why folks are reluctant to discuss the topic.

Nevertheless, that is exactly what I intend to do. Rather than attempt to give you a 250-page recipe for how to design your next digital interface, which my employer would not be happy about anyhow, I have tried to uncover what I see as the fundamental sources of failure and convey my own philosophy for preventing these failures from occurring. My challenge to you, the reader, is to use this book to develop your own unique approach to solving these core problems. No doubt as you're reading, you will think to yourself, "I could improve on that." After all, we *are* engineers. When you find mistakes (and I know you will), I hope you take them in the spirit of the book, which is communication between practicing engineers.

Chapter 1, "Engineering Reliable Digital Interfaces," is a bit of a lark. I cooked up a story to illustrate the multi-faceted nature of design for signal integrity. You will probably recognize a not-so-thinly-veiled element of realism in this introduction to the book. I wrap up this short chapter with a list of information that I collect during the early stages of a design to help me make decisions on the analysis required for a new card, board, or system design.

Chapter 2, "Chip-to-Chip Timing," covers the circuits used to store information in a CMOS state machine and how they fail. A set of SPICE simulations and spreadsheet budgets introduces the common-clock architecture, the first of three paradigms for transferring digital signals between chips. Even though the source-synchronous and high-speed serial paradigms are more prevalent in contemporary systems, the common-clock architecture is not dead yet. A solid approach for timing common-clock transfers is a useful thing to have in the toolbox.

IO circuits play a pivotal role in signal integrity engineering, yet we seldom get to lay our eyes on a schematic for one of them. A few CMOS IO circuits get used time and again, and Chapter 3, "Inside IO Circuits," examines their pertinent electrical characteristics. It also covers the assumptions we make when using behavioral models for these circuits. Studying these circuits provides a basis for understanding the more esoteric circuits. In this chapter and throughout the book, I make repeated references to the accuracy and quality of the models we use in signal integrity simulations.

Modeling forms the core of this book—both component modeling and interface modeling. It is essential to have a solid grip on the assumptions behind the models before stating that you trust the results of your simulations. Toward that end, Chapters 4, "Modeling 3D Discontinuities," and 5, "Practical 3D Examples," discuss the topic of 3D electromagnetic modeling in a manner that I hope will be digestible to the newcomer. Be advised that you may need to read them more than

once because field solvers are simply not as familiar to us as SPICE or IBIS. Both chapters are suitable for someone just getting started in 3D electromagnetic modeling and for those who want to ask their model providers the tough questions.

Chapters 6, "DDR2 Case Study," and 7, "PCI Express Case Study," tackle that million-dollar question: Will an interface operate with positive timing margin over the lifetime of the product without incurring the high costs associated with excessive conservatism? Here is where engineering can begin to lapse into religious fervor, mostly for lack of data. I have never been a big fan of throwing all the models into a pot, stirring it around, and taking a taste. My approach involves picking apart each interface piece by piece—understanding how many mV of crosstalk a connector generates and how many ps of eye closure go along with it. We owe it to our colleagues to show them our numbers and hope they will reciprocate. It's the only way to weed out the natural tendency toward covering one's own backside. Only when the numbers are on the table can you have a decent, meaningful argument.

For those new to signal integrity or simply in need of a refresher, Appendix A, "A Short CMOS and SPICE Primer," touches on some basic CMOS circuits and how we model them in SPICE.

When you're sitting in a class on electromagnetic theory, it's easy to get tangled up in vector calculus and lose sight of the physics. Appendix B, "A Stroll Through 3D Fields," hangs a high-level conceptual framework for field behavior. Think of it as a 5,000-foot flight over the scenery that attempts to place the physics in the context of engineering.

The people who will benefit the most from this book are new signal integrity engineers, those in transition from another discipline, and university students. Occasionally, I even run across the seasoned signal integrity engineer who has not had much exposure to chip-to-chip timing for whatever reason—well, more than occasionally. Because this is not a text on signal theory, I assume you have access to the other fine books on the market and can refer to them for concepts I do not develop here.

I hope you find this purchase worth your investment. Most of all, I hope it generates new thoughts that allow you to more effectively approach the problems you encounter in your day-to-day work.

Good hunting!