PRENTICE HALL

A Signal Integrity Engineer's Companion

REAL-TIME TEST AND MEASUREMENT AND DESIGN SIMULATION

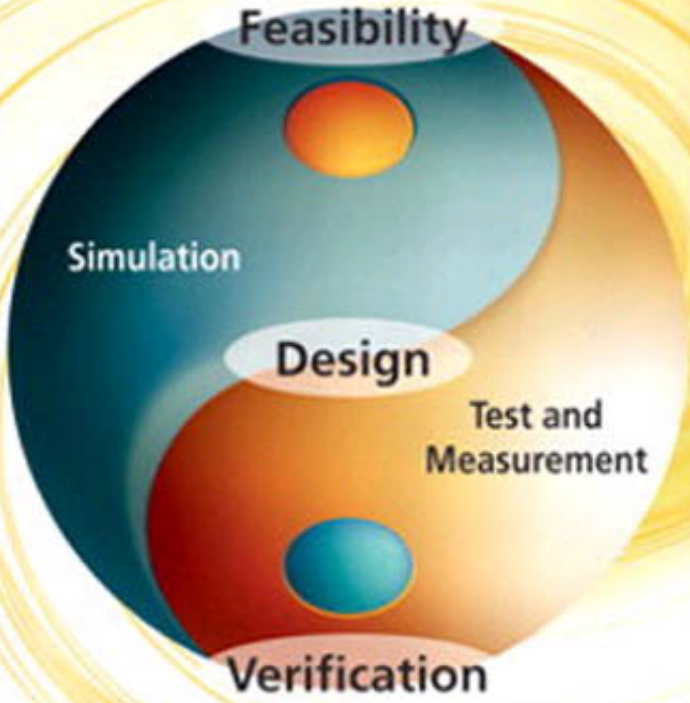

Geoff Lawday ▪ David Ireland ▪ Greg Edlund

Foreword by Chris Edwards, Editor, *IET Electronics Systems and Software* magazine

Prentice Hall Modern Semiconductor Design Series

Prentice Hall Signal Integrity Library

Visit us on the Web: www.prenhallprofessional.com

Library of Congress Cataloging-in-Publication Data:

Lawday, Geoff, 1946-
A signal integrity engineer's companion : real-time test and measurement and design simulation / Geoff Lawday, David Ireland, and Greg Edlund.
p. cm.
ISBN-10: 0-13-186006-2 (pbk. : alk. paper)
ISBN-13: 978-0-13-186006-3
1. Electronic apparatus and appliances—Testing. 2. Electronic apparatus and appliances—Design and construction. 3. Signal processing—Simulation methods. 4. Switching circuits—Reliability. 5. Oscillators, Electric—Testing. I. Ireland, David, 1957- II. Title.
TK7870.23.L39 2008
621.382'2—dc22
2008011981

ISBN-13: 978-0-131-86006-3
ISBN-10: 0-131-86006-2

Text printed in the United States on recycled paper at Courier in Westford, Massachusetts.
First printing June 2008

Credits

Editor-in-Chief
Mark Taub

Proofreader
Linda K. Seifert

Acquisitions Editor
Bernard Goodwin

Technical Reviewers
Sam Shaw, James Peterson, Andy Martwick, Mike Resso

Marketing Manager
Curt Johnson

Publishing Coordinator
Michelle Housley

Managing Editor
Kristy Hart

Multimedia Developer
Dan Scherf

Project Editors
Lori Lyons, Andy Beaster

Cover Designer
Louisa Adair

Copy Editors
Gayle Johnson, Geneil Breeze

Composition
Bronkella Publishing

Indexer
Brad Herriman

Graphics
Laura Robbins

Foreword

It is easy to be misled by the rhetoric of the day. Even the world of electronic engineering is not immune. Words and phrases such as “digitalization” and “digital convergence” carry the subtext that you only have to worry about ones and zeroes; that analog is being shown the door by a growing band of electronics designers in their quest to render more into binary logic. But it’s an illusion. Far from being squeezed out, the word analog is seeping into all areas of electronic designs in increasingly subtle and potentially damaging ways.

Matters are made even worse by the ease with which digital control can be used to massage and reshape the form of signals. Pre-emphasis is almost trivial to implement in the circuits that precede a driver. Although this processing can improve the ability of a receiver to decode the signal, it can have deleterious effects on other receivers in the vicinity. Worse than that, the interference can depend heavily on the data being transmitted. From that, it is not hard to see how intermittent, apparently random “Heisenbugs” can pop up during operation and promptly disappear the moment you try to add instrumentation to work out what is going wrong.

Signal integrity has been a problem for many years but the issues were often isolated to small parts of a system design. Today, all the trends point to signal corruption getting worse and worse. Switching speeds are going up and the voltages provided on supply rails are going down.

But the trends are not all technical. As you can read in the introduction that follows, some of the biggest problems can result from commercial decisions: the pressure to reduce manufacturing costs are pushing designers to consider cheaper components, packages, and substrates; and designers have less time to get the job done.

In some markets, such as cellular handsets, companies want to be able to produce variants very quickly. They might be on the shelf for only six to nine months. Marketing may not know more than six months out whether the design will be a flyer or a dud: All they can do is extrapolate current trends and hope. The closer they define a product to when it is meant to go on sale, the better their chance of getting it right. But that is no help to an engineering team trying to make the phone work.

The core chips may have been designed for a different phone. But a new screen or keyboard, or a switch from a candy-bar to a flip-phone design means that the board layout has to change. And with a small rearrangement of the components on that board, you can suddenly find that the design decisions made by the chipmakers are at odds with the requirements of the new layout. What was meant to be a quick-turnaround project to create a simple variant of an existing phone for a new market suddenly looks a lot less tractable.

In this book, the authors take you through the methods available to digital designers to ensure that they are not vulnerable to the little tricks that the analog properties world can play on them. It is a comprehensive treatment that shows how working in the virtual and real worlds provide a combined methodology for avoiding signal-integrity problems. It is tempting to think of signal integrity as a subject dominated by “black magic” techniques. But there is plenty of science to help the time-starved engineer ensure that a high-speed, low-voltage bus will work in the final system.

This book demonstrates how modeling and behavioral simulation let the engineer make sensible decisions early on in the project. It covers tricky subjects such as the modeling of transmission lines—a skill that will prove vital in the coming years.

Equally important is the ability to work out where things are going wrong in the prototype, and to track down the source of the problem. Chapters on probing, oscilloscope use, and time-domain reflectometry provide practical advice on the best way to look beyond the ones and zeroes the logic is meant to see into the electromagnetic soup that the real world is made from.

It is not just about the wired world either: the last chapter concentrates on the wireless world and the challenges raised by new software-defined radio architectures.

This is a book that I am sure will be an essential addition to every electronic engineering lab as more people find they have to grapple with the analog infrastructure that underlies every software-driven digital system. In such a digital world, a book like this has never been so important.

Chris Edwards, Editor, *The IET Electronics Systems and Software* magazine
London, UK

Chris Edwards reports on electronics, IT, and technology matters. He has more than 15 years of journalism experience as an editor and writer. He is currently the freelance editor of *Electronics Systems and Software*, published by the Institution of Engineering and Technology (IET), and a regular contributor to the magazines *Engineering and Technology* (formerly *IEE Review*), *Information Professional*, as well as the new tech magazine for teenagers, *Flipside*.

Preface

WHAT THIS BOOK IS ABOUT

We live in the high-speed digital age where embedded system developers in many cases must apply new design methodologies and perform complex signal integrity tests and measurements. This book guides the reader through the full life cycle of embedded system design from specification and simulation to test and measurement. A significant feature of the book is the explanations of the thorny issues of signal integrity engineering that are so often the cause of delay in a product's development—time to market is the signal integrity engineer's Achilles' heel. By considering the whole life cycle from simulation through to test and validation you can see the new interplay of simulation and real-time test, which drives the new design methodologies. A case in point is the design and implementation of a new high-speed serial bus where the simulation and real-time test are inextricably linked, especially where designs incorporate device driver pre-emphasis or receiver equalization.

The celebrated teacher and philosopher Archimedes said, "What we must learn to do is to learn by doing," and this book endeavors to do just that by presenting practical applications so that you can learn from the authors' experiences

in embedded system design, simulation, test, and measurement. What's more, you are encouraged to consider and adopt good practices in signal integrity engineering throughout the entire life cycle of the design.

Of particular importance today is the need to meet regulatory compliance and interoperability requirements. Consequently this book treats the demands of compliance testing and interoperability design as a foremost topic. Alongside today's compliance testing is the migration to high-speed serial buses, where both topics lend themselves to some prime examples of good practice in signal integrity engineering. Therefore, the design and testing of high-speed serial buses with their associated low-voltage differential signaling are pivotal themes throughout several chapters of this book. Nevertheless, the fundamentals of signal integrity engineering underpin the understanding of compliance testing and interoperability design, and you are encouraged to refresh or learn the basics of signal acquisition, test, measurement, and device simulation, which for the most part is provided in this book.

Work as a team has become routine for engineers designing, developing, and testing digital systems. The members of the team build up companionships and regularly look to each other for advice, especially in today's complex high-speed digital world. Today, the task of developing or testing a digital system is complex, and even the design of what appears to be a simple circuit can be problematic. For example, an engineer can design a simple digital circuit that meets all the requirements, and then a year down the line a chip is changed because the new chip is cheaper and has the same functionality. What the engineer didn't know is that the new chip is now made with a smaller device size that reduces its cost and its switching time, leading to faster edge rates and signal integrity problems! This leads to the well-documented quote "There are two kinds of design engineers, those that have signal integrity problems, and those that will." As we have seen if a design is sensitive to edge rates, the component specification must make edge rate a formal product parameter since it is just not possible to anticipate the evolution of a silicon fabrication process. This book aims to be a companion to the engineer and part of the engineer's team by providing an understanding of design specification, simulation, test, and measurement along with some significant advice on maintaining signal integrity throughout the life cycle of a design.

Although we can take the big view, there are significant little problems that the engineer needs to know but, as they say, "is afraid to ask." For example, if a designer suspects ground bounce, or more accurately transients in the signal reference, and wants to measure the ground line with an oscilloscope, where is the probe ground connected? Actually, it is normally connected to a solid logic zero. Moreover, why is it that a state-of-the-art oscilloscope can give a 50% measurement error when measuring ground bounce? Well, the bandwidth of an oscilloscope is typically specified at the –3 dB point, and a voltage measured at the limit

of the specified oscilloscope bandwidth will be shown as half the real voltage. Also it is important to measure the voltage but think in terms of current—since the current spike on a ground rail generates crosstalk and spurious switching. And we could go on; there are a myriad of signal integrity challenges from intermittent setup and hold violations, resulting in problems ranging from metastability to electromagnetically induced crosstalk. The ability to foresee signal integrity problems and how to avoid them is fundamental to this book.

A feature of this book is the blend of source material. Whereas a theoretical text on signal integrity is built on scientific laws and notable hypotheses, this book has sourced its applications from the authors' professional experiences, published papers, and the work of associates. This book is a blend of source material coherently assembled and expanded to provide an understanding of modern signal integrity applications. Practical issues concern us most in this book. Each chapter focuses on a day-to-day activity of the signal integrity engineer, giving advice and illustrations from the industry. Practicality forms the central theme throughout the book.

The twenty-first century is a digital era of media convergence where mobile telephony, computing, and digital broadcasting merge, and consumers expect the media to be transparent to the technology. Put simply, the sports fan expects to view an event, in real time, on his or her mobile phone, laptop, and personal music player or digital TV at a reasonable cost and with absolute reliability. We could have taken any number of other examples from industry, medicine, or the military. Today, the digital designer or maintenance engineer is expected to be accomplished at signal integrity engineering, which is at the heart of the provision of systems that will make tomorrow's innovative technologies happen. This book reflects this trend.

THE INTENDED AUDIENCE

Signal integrity engineering is a young and evolving science where few who proclaim to know it all. Writing this book has been a journey of discovery for the authors, and we have every reason to believe it will be a rewarding journey for you. We have no doubt that some topics discussed in this book will provoke debate as there is much to be standardized in this branch of engineering. However, indisputable principles are presented in this book that underlie signal integrity engineering. These principles give the emergent engineer a basis on which to build the knowledge and understanding necessary for good signal integrity engineering. Therefore, this book is recommended reading for the student signal integrity engineer and the practicing engineer whereby the authors

present a wealth of applications that illustrate good practice and show the development, test, and validation of modern digital systems.

While the book is naturally partitioned into chapters of diverse topics a common thread runs throughout the book. Each chapter provides a guide for the reader by presenting the necessary prerequisites of a topic before detailing complex design or test applications. Consequently the experienced engineer can approach a topic by stepping through the beginning of a chapter and concentrating on the detail in the applications and advanced topics. No signal integrity book claims to be all encompassing, and this book is no exception. You may need to consolidate your understanding of the theory of signal integrity engineering via in-depth theoretical texts in the Prentice Hall SI series. Nevertheless, much of this book is self-contained in terms of addressing a wide audience in signal integrity engineering.

HOW THIS BOOK IS ORGANIZED

To guide you through the full life cycle of embedded system design, including specification, simulation, test, and measurement, the book is structured, where possible, chronologically to follow the development cycle. However to encompass the diverse aspects of signal integrity engineering and to provide a coherent thread as you read, chapter order is a compromise of product life cycle flow and a natural grouping of signal integrity engineering topics. Therefore both compliance and serial bus simulation are found toward the latter part of the book, whereas earlier topics are prerequisites for these more advanced subjects.

Chapter 1: Introduction: An Engineer's Companion

Chapter 1 takes you into the world of device and circuit simulation, which is a major phase in the successful development of a modern digital product. A thought-provoking example is given whereby a designer is under the intense pressure of time to market where it's easy to overlook a true understanding of operating margins—will a network continue to function reliably over the range of manufacturing and operating conditions it will encounter during the useful life of the product? What are the expected primary failure mechanisms, and how do they interact with one another? Some of these complex simulation questions will be considered in the body of the introduction, but more to the point, these discussions lead the way to the in-depth chapters that examine these concerns.

Following simulation, the introduction describes in detail a number of principal innovations in signal integrity engineering test and measurement. For example, to overcome some of the traditional signal integrity engineering problems,

device manufacturers currently use novel integrated signal processing functions within device drivers and receivers to apply signal pre-emphasis and equalization. However, incorrectly applied pre-emphasis generates unwanted overshoot, crosstalk, and noise. This is one illustration of the complexity in solving today's signal integrity problems. This chapter throws light on such issues and presents a pathway for the solution of such problems by describing the basics of eye diagrams, which form the basis of many of today's automated compliance tests and signal analysis measurements.

Chapter 2: Chip-to-Chip Timing and Simulation

This chapter 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 handful of CMOS IO circuits get used time and again, and Chapter 2 examines their pertinent electrical characteristics. The chapter 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. This chapter and others make repeated references to the accuracy and quality of the models we use in signal integrity simulations.

Chapter 3: Signal Path Analysis as an Aid to Signal Integrity

This chapter describes signal path analysis based on intuitive time-domain reflectometry (TDR) techniques. TDR measurement theory and its application are described in detail, given that TDR is fundamental to the understanding of signal integrity effects, such as impedance mismatches and circuit board (PCB) issues. In particular, it provides an ideal vehicle for illustrating some principal signal integrity challenges and their solutions. Another facet of this chapter is the introduction to Vector Network Analysis (VNA), which is an important frequency domain measurement methodology used, among other things, to accurately characterize high-speed signal paths. A particular feature of this chapter is the introduction to the design, development, and test of Low Voltage Differential Signaling (LVDS) signal paths. However an understanding of basic transmission line theory underpins good practice in signal path design.

Today, with high-speed digital signal transmission, even the shortest passive PCB trace can exhibit transmission line effects. Transmission line theory encompasses electromagnetic field concepts and generally attracts complex mathematical analysis. However, using TDR, leads intuitively to a basic understanding of transmission line theory, even though some of the basic concepts require a few simple calculations.

Chapter 4: DDR2 Case Study

The DDR2 case study tackles the million-dollar question for a common source-synchronous bus: Will the interface operate with positive timing margin over the lifetime of the product without incurring the high costs associated with excessive conservatism? This approach involves picking the interface apart piece by piece—understanding how many mV of crosstalk a DIMM connector generates and how many ps of eye closure go along with it. Under the pressure of a project schedule, it is often tempting to gather a set of models, construct a simulation, and be done with the exercise. This chapter challenges the reader to take a deeper look.

Chapter 5: Real-Time Measurements: Probing

Central to the book and this chapter is the challenge of data acquisition whereby the problematic issues of the ideal nonintrusive probe are examined. How is signal fidelity achieved in a modern signal integrity measurement, and how are unwanted measurement artifacts avoided? Analog signal measurement is carefully investigated to demonstrate the importance of correctly connecting to a system under test because the analog features of a high-speed digital signal determine signal integrity. The chapter provides practical advice on probing and how to avoid probe problems. Special attention is given to the probing of today's fast-edge signals and Low Voltage Differential Signaling (LVDS).

Chapter 6: Testing and Debugging: Oscilloscopes and Logic Analyzers

This chapter reviews modern signal integrity testing from both an analog and digital viewpoint, since at the high frequencies encountered in today's designs the two are inextricably linked. The emphasis is on frequencies over 1 GHz, where the measurement tools of choice are the digital oscilloscope and the logic analyzer. The chapter provides practical examples to show how detailed observation of both analog and digital waveforms, side by side, can provide the data necessary for the understanding of the most challenging signal integrity timing budget issues, such as setup and hold violations, data skew, and metastable states. Real-world illustrations show these problems and how they can be detected and debugged.

Chapter 7: Replicating Real-World Signals with Signal Sources

A core theme throughout signal integrity engineering is the behavioral analysis of a digital circuit in terms of its analog properties and notably how the behavior of a digital circuit is determined. Sections of this book are devoted to the methods used to simulate models where computer-generated outputs show signals and data in a variety of formats in response to an array of simulated inputs. However, the fundamental method used to determine the real-time characteristics and operation of an actual circuit, or prototype, is to externally control and observe the circuit or device under test. Most of this book is about signal acquisition and measurement, and this chapter provides the balance; it is about the other half of the story—the signal source—which is used to control a circuit. Put simply, this chapter is about the externally provided signal that is used as a real-time stimulus for electronic measurements. This chapter describes and demystifies the complex issues of modern signal sources and shows how they can be used to stress a digital circuit to expose signal integrity faults.

Chapter 8: Signal Analysis and Compliance

The proliferation of digital systems, such as the new high-speed buses, has created numerous interoperability and compliance standards. This chapter explains how real-time test and measurement is the cornerstone of compliance testing. It describes the various standards, with particular emphasis on high-speed serial buses, and shows why, at frequencies of more than 2.5 GHz, real-time test and measurement is the only way to achieve compliance. The practical use of logic analyzers and oscilloscopes for compliance validation is examined, whereby techniques such as automated eye diagrams and statistical analysis are discussed in detail. This chapter is essential reading for the signal integrity engineer who needs to understand the challenges of meeting regulatory compliance tests.

Chapter 9: PCI Express Case Study

High-Speed Serial (HSS) is the last of the three major paradigms for transferring data between chips, and PCI Express is an excellent example of a high-speed serial interface. The PCI Special Interest Group published a set of guidelines that will keep designs out of trouble. However, situations often arise that force a designer to depart from the well-traveled path, either by breaking one of the guidelines or by trading one off for another. In these cases, it is helpful to acquire an understanding of how much each interconnect component contributes to the jitter budget and how many picoseconds remain after accounting for all relevant effects. Starting with a set of models and design rules, Chapter 9 examines the characteristics of each component in the time and frequency domains and then combines them one at a time to arrive at a total jitter budget.

Chapter 10: The Wireless Signal

The success of cellular technology and wireless data networks has caused the cost of basic radio frequency (RF) components to plummet. This has enabled manufacturers outside the traditional military and communications markets to embed relatively complex RF devices into all sorts of commodity products. RF transmitters have become so pervasive that they can be found in any number of embedded systems. Therefore this book introduces RF test and measurement for completeness in the understanding of signal integrity engineering. Moreover, given the challenge of characterizing the behavior of today's high-speed logic devices, this chapter provides an understanding of how radio frequency parameters such as jitter are measured. Although this chapter provides a detailed decision of the real-time spectrum analyzer (RTSA), the chapter also offers a detailed introduction to the Swept Spectrum Analyzer (SA) and the Vector Signal Analyzer (VSA).

THE WEBSITE THAT ACCOMPANIES THIS BOOK, WWW.INFORMIT.COM/TITLE/0131860062

Color Pictures and Illustrations

Illustrations and pictures are used throughout a number of chapters in this book to allow the reader to become involved with instrumentation applications. In particular the chapters describing test and measurement use logic analyzer, oscilloscope, and spectrum analyzer displays to quantify and simplify what would otherwise be difficult and wordy descriptions. However, in keeping with most other books, the illustrations are understandable in monochrome, but some of the detail or features of a picture can be lost. Therefore, downloading the figure files from the website, www.informit.com/title/0131860062, allows the reader to view the descriptive color images and relate them to the accompanying text contained in the following three chapters:

- Chapter 6, "Testing and Debugging: Oscilloscopes and Logic Analyzers"
- Chapter 8, "Signal Analysis and Compliance"
- Chapter 10, "The Wireless Signal"

Simulation Models

Chapters 2, 4, and 9 demonstrate the allocation of picoseconds using case studies of three interface paradigms: common clock, source synchronous, and high-speed serial.

The model kit for Chapter 2 includes SPICE transistor models for an ancient 3.3 V 0.5 um CMOS process, IO circuits, and some simple networks.

Chapter 4 features behavioral simulation of standard DDR2 IO circuits, lossy transmission lines, vias, and a DIMM connector. Although the chapter discusses DIMM connector crosstalk, the model kit only contains single-line models because the DIMM connector model is proprietary.

The PCI Express case study in Chapter 9 demands more accurate interconnect models: coupled s-parameter representation of a ball-grid array, the corresponding via field, and edge connector. The simulations utilize a simple 100 ohm de-emphasized driver model found in the Agilent ADS library. This model is not included in the kit.

All models are suitable for simulation in SPICE or a behavioral simulator.

1

Introduction: An Engineer's Companion

An engineer's companion is like any other companion: It's a fellow traveler and colleague who offers advice and support, sharing experiences with a friend, in what would otherwise be a solitary journey. Our journey will take us through the endeavors of embedded system design, simulation, prototype development, and test. The dangers are signal reflections, attenuation, crosstalk, unwanted ground currents, timing errors, electromagnetic radiation, and a host of other signal integrity (SI) issues. SI engineering is a relatively new branch of electronics engineering. For the most part, it relates to the analog factors that affect both the performance and reliability of modern high-speed digital signals and systems. In general, integrity has to do with truthfulness. When applied to digital electronics, such as communication and computer systems, it is specifically about signal accuracy and system reliability.

Although it is written for SI engineering professionals, this book is intended to support new engineers and students who have an interest in designing, simulating, and developing modern high-performance digital and embedded systems. Along with the central theme of how to think about SI engineering, this book includes practical guidance on how to achieve and interpret a simulation or real-time test and measurement. In many jobs within the SI industry, a technician,

engineer, or designer could be anyone who thinks about the reliability or operation of a modern embedded system. This book aims to address the concerns and uncertainties faced by these people. Because this book was written with a wide audience in mind, each topic is presented with a prerequisite theoretical or practical preamble that supports novice engineers and students and that can often be omitted by practicing engineers. This chapter keeps with the format of this book, in that it follows the development of an embedded system, from its simulation to prototyping and real-time test.

We live in the digital age, in which providers of telephony, computing, and broadcast systems are busily facilitating media convergence. Music, video, and information systems must be transparent to the newly integrated communication and computing systems. Consequently, consumers expect their modern high-performance telephony and computing systems to interactively communicate the latest news and entertainment while providing e-business transactions. Driving forces such as media convergence are challenging digital designers to work in an era of reasonably priced, high-performance, highly dependable digital systems that typically are portable and generally are required to have worldwide compliance. Today the issues of interoperability are paramount where modern systems and components from disparate manufacturers are often required to work together seamlessly. Moreover it's widely documented that communication and computing systems double in computational throughput every eighteen months. This implies that any absolute frequency or data rate quoted in any SI book will be out of date by the time the book is published—and this book is no exception. Nevertheless, it is anticipated that, similar to most of the books in this SI engineering series, the fundamental practical examples, guidance, and underlying theoretical concepts will remain relevant for many years to come. Today's cutting-edge digital interfaces will become the bread and butter of tomorrow's digital systems.

1.1 LIFE CYCLE: THE MOTIVATION TO DEVELOP A SIMULATION STRATEGY

Most signal integrity engineers would agree that the primary motivation for simulating chip-to-chip networks is to maximize the probability that those networks will function flawlessly on first power-up. Another compelling motivation is easy to overlook under the intense pressure of time to market: understanding operating margins (see Figure 1-1). It is tempting to stitch together IO circuit and interconnect models, run the simulations, check the results, and be done with the exercise. This may prove that the network will function under a given set of conditions, but will it continue to function reliably over the range of manufacturing and operating conditions it will encounter during the product's useful life? What are the expected primary failure mechanisms, and how do they interact with one another?

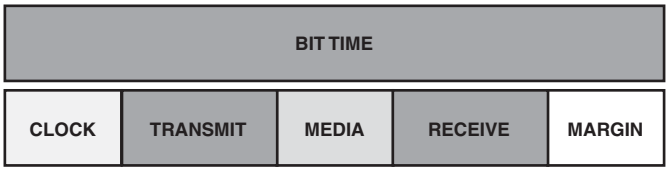

Figure 1-1 Operating margin.

These questions are indeed of primary importance, and it may be possible to answer them given unlimited resources and time. Unfortunately, most signal integrity engineers operate under somewhat different conditions. A contemporary PC board design may have upwards of a thousand nets that belong to two or three dozen different buses. The power spectrum will likely have significant content above 5 GHz. Supplying power and cooling to high-performance processors can place challenging constraints on layout and routing. On top of these technical challenges, the customer may require that the product be ready for manufacturing in a time period that severely stresses the team's ability to carry out the level of analysis required to ensure reliable operation. This is the irony of the business: the relevant physical effects become ever more difficult and expensive to analyze, while the market relentlessly exerts downward pressure on cost and schedule. These two freight trains are running full speed toward each other on the same track.

Given these technical and business challenges, is it still possible to achieve the goal of reliable operation of a system filled with dozens of digital IO buses over the product's lifetime? At times it may appear that the solution to this difficult problem is the empty set. In the heat of battle, the level of complexity can be so overwhelming that it seems impossible to satisfy all the constraints simultaneously. Nevertheless, it is the opinion of the authors that it is possible to successfully manage the signal integrity of a complex contemporary design if the lead engineers keep two important principles in mind. First, the signal integrity engineer must be involved at the very beginning of the design cycle, when the team is making critical architectural decisions and selecting component technology. Second, the team must develop a comprehensive simulation and measurement strategy that applies the appropriate level of analysis to each bus in the system.

1.1.1 The Benefits of Early Teamwork

It must be tempting for those who make weighty architectural decisions during the earliest stages of a new product to avoid addressing implementation details. After all, it is easier to define a product while cruising along at the top of the troposphere. The product must offer innovative, distinguishing features. It must also

offer Earth-shattering performance at a lower cost than the competition. Most important, it must be ready for manufacturing before the infamous marketing window closes. If this is a consumer product, this means having the product shrink-wrapped and on the shelves for Christmas shoppers. These are all difficult goals to achieve, even without having to worry about how to manufacture the product and make it reliable.

As any experienced engineer will attest, usually a price must be paid for making architectural decisions without input from those whose job it is to implement the architecture. Probably the worst possible scenario is a product that is marginally functional, because the marginal part does not become apparent until production is in full swing. This is even more treacherous than a product that overruns its budget, misses its milestones, and never makes it to market. Assembly lines come down. Companies—more than one—lose large quantities of money each day. There may be recalls. There will certainly be redesigns under intense pressure. At the end of the whole experience lies a painful loss of reputation. This is certainly a scenario that no Vice President of Technology or Chief Financial Officer would choose to put in motion if the choice were made clear. One thing is crystal clear: the company greatly enhances it chances of success by building a cross-disciplinary team in the project's early architectural phases. A minimum team would include a board layout designer, firmware programmer, and engineers from the disciplines of logic, mechanical, thermal, power, manufacturing, electromagnetic compliance, and signal integrity. A large company may have separate engineers to represent each of these disciplines, while in most other companies one person wears several hats. In either case, it is important that each discipline be represented on the team.

As an illustration of the importance of early teamwork, consider the following fictional scenario that is close enough to reality to be disturbing. Project X took off like a rocket from the very beginning. Conceived in a boardroom behind closed doors, it was already well-defined before anyone in development engineering heard of it. The senior engineering staff voiced their opinions about the unrealistic schedule, especially during a time when other high-profile projects were consuming most of the company's resources. However, commitments to customers had already been made, a marketing plan was in the works, finances were allocated, and the wheels were in motion—the first domino in a sequence of related events.

The second domino was in place before PC board placement and routing began. The Vice President of Engineering had defined a budget that was consistent with the price point that Marketing deemed competitive. This budget called for a four-layer PC board: signal-ground-power-signal. Upon seeing the form factor for the first time, the PC board designer expressed concern about routability in certain areas he perceived to be bottlenecks. The lead signal integrity engineer

expressed concern about high forward crosstalk due to microstrip transmission lines and high edge rates in the PCI Express (PCIe) signals. However, neither the board designer nor the signal integrity engineer could prove that a solution did not exist, so the design progressed as planned.

An exceptionally aggressive schedule became domino number three. To meet the schedule, the board design shop had two of their top designers work back-to-back twelve-hour shifts for three weeks, with one day off on the weekends. While it is true that an auto-router can save many hours of human labor, there is no substitute for an experienced designer when channels are fully allocated. The company's design process called for routing constraints to be in place before routing could begin. Again, schedule won the day. Routing began before the signal integrity team could assign constraints to the 800 nets on the board that required attention out of a total of 1,000 nets.

In a remarkable feat of skill and sweat, the design team generated Gerber files on schedule. Much to their credit, they reserved one day for a complete design review, with mandatory attendance by everyone on the team. They even invited a few seasoned veterans who had since gone on to jobs in other parts of the company. One of these veterans spotted the problem: a PCIe differential pair routed next to a Gigabit Media Independent Interface (GMII) clock signal. The edge rate of the 2.5 gigabits per second (Gbps) PCIe signal was clearly defined in the specification as 10 V/ns. The edge rate of the 125 MHz GMII clock was only 1 V/ns, making it vulnerable to crosstalk from aggressors with higher edge rates, but this information was unavailable in the component datasheet. The signal integrity engineer took an action item to acquire an IBIS datasheet for both the PHY chip and the IO controller. She was successful for the PHY chip, but the vendor for the IO controller required the company to sign a nondisclosure agreement before releasing the IBIS datasheet. This legal process took much longer than the one day allocated for reviewing the design, and the team sent the Gerber files to the PC board manufacturer. Domino number four.

The first set of boards came back from assembly and performed admirably on the benchtop. The software team did their job of loading new code, debugging, recompiling, and loading again. The hardware team did their job of measuring thermal characteristics, calculating power draw, dumping registers, and capturing traces for bus transactions on the logic analyzer. Knowing that there was exposure to crosstalk problems, the signal integrity engineer spent a lot of time probing the board, looking for them. She found some quiet line noise on the order of 200 to 300 mV, but the nets had enough noise and timing margin to tolerate it. After several weeks of intense work, the development team gave management the green light for production, and each team member spent a few days contributing to the final report before moving on to the next project.

Unbeknownst to the heroes of our story, the fifth and final domino was starting to fall at a semiconductor manufacturing plant on the other side of the planet. Although the IO controller chip had been in production for two years and the process had remained stable for most of this time, a recent drop in yield prompted the process engineers to make some tweaks that ultimately resulted in slightly lower edge rates. The IBIS datasheet for the IO controller contained edge rate information, but the component datasheet did not. The manufacturer did not feel a need to notify its customers, because all specified parameters remained within their limits.

Shortly after these new components hit the assembly floor, the boards began experiencing high levels of fallout in the form of intermittent failures that were associated with traffic on the PCI Express and Gigabit Ethernet buses. When a few of the boards made it to customer installations, the crisis was officially under way. Management told the team to drop what they were doing and focus on resolving the crisis. One sleepless night, the neurons in the brain of the signal integrity engineer rewired themselves to remind her of the comment during the design review about crosstalk between PCI Express data and the GMII clock. The next day she decided to probe these signals, using a logic analyzer to trigger the oscilloscope when both were active at the same time. This did not happen often, because the two buses were asynchronous to each other. After three days of persistent testing she captured a waveform that showed a slope reversal—double clocking—right in the threshold region of the GMII clock receiver on the PHY chip and coincident with a transaction on the PCI Express bus.

Not many pleasant alternatives presented themselves to the team. It was not possible to slow down a PCI Express signal and still expect the bus to function. While an FET switch would sharpen the edge of the GMII clock signal, the IO controller was in a BGA package, and there was no way to install the FET switch between the pin and the net. In the end, it was decided to stop production and rush a new six-layer PC board through design.

During the lull between releasing the new design and shipping the new boards, upper management held an all-day process review. The veteran engineer from the original design review retraced the trail of circumstances that led to the failure. First and most important, the signal integrity team did not have a representative at the table when the architecture was being defined. He pointed out that a cross-disciplinary design team is the cornerstone of a healthy, solid design process. Second, there was no cost-performance analysis of the PC board stackup. This is admittedly one of most difficult things any engineer has to do. There was no excuse for the third contribution to the failure; schedule should never trump assigning design constraints unless the risk of nonfunctional hardware is deemed acceptable. Item number four was an unavailable model. This is understandable, because high-quality models are hard to come by, but it is possible to

have the necessary models when they are needed by planning ahead. Finally, if a design is sensitive to edge rate, the component specification should call out edge rate as a parameter, because it is not possible to anticipate the evolution of a silicon fabrication process.

The recommendations that came from the design process review were to establish a cross-disciplinary team for all new designs and to develop a comprehensive process for applying the appropriate level of analysis to each net in a PC board or system design.

Most of the time, circumstances do not present themselves to us in such an obvious, logical progression. Only careful retrospection reveals the sequence of events that led to a particular conclusion. To some degree, the engineer's job is to play the role of the prophet who can foresee these circumstances and avoid them without becoming the Jeremiah to whom nobody pays much attention.

1.1.2 Defining the Boundaries of Simulation Space

Avoiding situations like the one just described requires a sound understanding of the physics involved. You also need the insight to know what level of analysis is required for a given net or bus. Although it is certainly possible to acquire models for a thousand nets and simulate every one of them before releasing a design to manufacturing, the company that practices such a philosophy may not be in business very long. It would appear that one of the more critical tasks of designing any digital IO interface is establishing the boundaries of simulation space. These are the criteria you use to decide whether a net needs to be simulated or whether some other method of analysis is more appropriate. Make no mistake: simulation is expensive and should be used only when there are strong economic and technical motivations for doing so. Once this critical question of whether to simulate is answered, you can go about the tasks of actually running the simulations and interpreting their results.

An excellent way to begin the decision-making process is to compile a comprehensive list of all nets in the design and some relevant information associated with each bus. This begs for a definition of a bus, which is a word that is frequently used but seldom defined. For the purposes of this discussion, a bus is defined as a collection of data and control signals that have a common functional purpose and are synchronized to the same clock or strobe signal.

One example of a bus is the traditional 33 MHz PCI bus. It is composed of 32 address-data signals and a set of control signals, each synchronized by a common clock signal that originates from a clock source chip. DDR memory is a source-synchronous bus in which the transmitting chip sends a clock or strobe signal along with the data. The source synchronous bus facilitates faster data rates by eliminating skew between multiple copies of the same clock and transmitter

launch time from the timing budget (see Chapter 4, "DDR2 Case Study"). PCI Express is another example of a bus, although it is not necessary for each chip on a given PCI Express bus to share the same reference clock. PCI Express uses a clock-data recovery circuit. This means the receiving chip uses the same low-frequency reference clock as the transmitting chip. It boosts the clock to the data rate and infers the optimum sample point from the incoming data stream.

The analysis decision matrix shown in Table 1-1 allows the signal integrity engineer to view all the relevant electrical parameters of each bus at the same time. The engineer also can decide what level of analysis is necessary to ensure that each bus functions reliably over the product's lifetime. The simplest case might be a bus that a trusted colleague has analyzed in the past and that others have used successfully time and again. In this case no simulation is required—*provided* that all of the bus's electrical parameters are identical to the ones that were analyzed in the past. The next, more complicated case is the bus for which a designers' guide or specification exists. If a third party analyzed the bus and published a set of rules that, when followed, guarantee sufficient operating margins, simulation is not necessary. The job of the signal engineer defaults to describing design constraints to the CAD system and checking that they are met. Of course, the reliability of the source must be beyond reproach! Some buses may not require simulation but do require rudimentary hand calculation, such as the value of termination resistors, stub length as a function of rise time, or RC time constant of a heavily loaded reset net. Finally, if a bus passes through each of these three filters, it is time to assemble the models and fire up the simulator. The closer this process occurs to the beginning of the project, the higher the likelihood of success. The bus parameter spreadsheet should include the following items:

Table 1-1 Analysis Decision Matrix

Parameter	I2C	PCI-X	DDR2	PCIe	Units
Engineer					
Net count					
Data rate					Gbps
IO power supply voltage					V
IO circuit technology					
Input setup time					ps
Input hold time					ps
Input minimum edge rate					V/ns
Input high threshold					V
Input low threshold					V

Parameter	I2C	PCI-X	DDR2	PCIe	Units
Output rise time					ps
Output fall time					ps
Output maximum edge rate					V/ns
Output impedance					ohm
Output high level					V
Output low level					V
Pin capacitance					pF
System clock skew and jitter					ps
Net characteristic impedance					ohm
Termination					ohm
Maximum net length					in.
Number of loads					

It's helpful if the signal integrity engineer and the person who draws the schematics can agree on a naming convention that involves adding a prefix to the net name of each net in a bus. This will facilitate tracking coverage of all nets in a design. Someone who is good with programming can write a simple script that sorts and counts the nets from the "all nets" file, the best source of which is either the schematic entry or layout CAD tool. The goal is to make a decision about each net in each bus in the system: what level of analysis does it require? You can then keep track of the number of nets simulated and constrained and have an up-to-date measure of how close a project is to completion. Totaling the net count column will give you an excellent rough estimate of the work involved at the beginning of the project. Keeping track of who is analyzing each bus prevents unpleasant revelations toward the end of a project, such as "I thought so-and-so was working on that bus!"

1.2 PROTOTYPING: INTERCONNECTING HIGH-SPEED DIGITAL SIGNALS

Traditionally the professional engineer strives to design high-performance digital and embedded systems within tight time-to-market constraints, cost limitations, and quality demands while managing manufacturability requirements. Alongside the traditional concerns, the key challenge facing today's designer is the task of maintaining signal integrity in a modern high-performance digital system. A contemporary example of the changing landscape of embedded system design is the support given to the SI engineer by device manufacturers that integrate some ingenious circuitry within their devices to minimize or circumvent a number of SI

issues. A particular case in point is the interplay between the SI engineer and modern programmable logic device manufacturers. The inclusion of programmable pre-emphasis, deskewing, edge rate control, and equalization provides a range of solutions to a number of key SI concerns.

As system speeds increase and timing budgets decrease, there is less time for switching between logic levels. Consequently, digital signal edges become faster, which results in the need for rigorous design requirements if signal integrity is to be maintained. High-performance digital systems are prone to two fundamental sources of signal degradation:

- Digital degradation that is timing-related, such as synchronization and setup and hold violations, which often generate metastability or race conditions that typically produce erratic system behavior
- Analog degradation, such as indeterminate signal amplitudes, power supply and ground variations, glitches, signal overshoot, crosstalk, and unwanted noise, that generates a diversity of system malfunctions

Both of these phenomena typically have their origins in printed circuit board (PCB) design or signal termination, but there are a myriad of other causes. Not surprisingly, there is a high degree of interaction and interdependence among digital and analog signal integrity requirements. The analog aspects of digital system design tend to cause the most concern. High-speed signals transmitted along PCB tracks tend to suffer from high-frequency attenuation, which makes it difficult for a receiver to interpret the information. The effect is similar to a low-pass filter, which decreases a signal's high-frequency content.

The main causes of high frequency attenuation are PCB dielectric loss, which is a capacitive effect, and the skin effect, which limits the signal to the surface of a conductor. As the data rate increases or the edge rate becomes faster, the signal frequency increases, and the dielectric loss becomes the dominant factor in high-frequency attenuation. The effect of dielectric loss is proportional to frequency, whereas the skin effect is proportional to the square root of the frequency. The skin effect describes how high-frequency currents tend to travel on the surface of a conductor, rather than on the whole cross section of the conductor. This is caused by the conductor's self-inductance, which increases the inductive reactance with frequency, forcing the current to travel on the surface of the material. This reduces the effective conductive area of a PCB trace, increasing the trace's impedance, which causes the signal to be attenuated.

While other PCB anomalies such as poor termination can cause crosstalk and reflections, the problem of high-frequency signal loss is aggravated as signal frequencies increase and PCB tracts lengthen. For example, Figure 1-2 shows the

particularly demanding case of a 40-inch backplane where high-frequency signals are transmitted through a PCB stripline that is constructed with FR4-type PCB material.

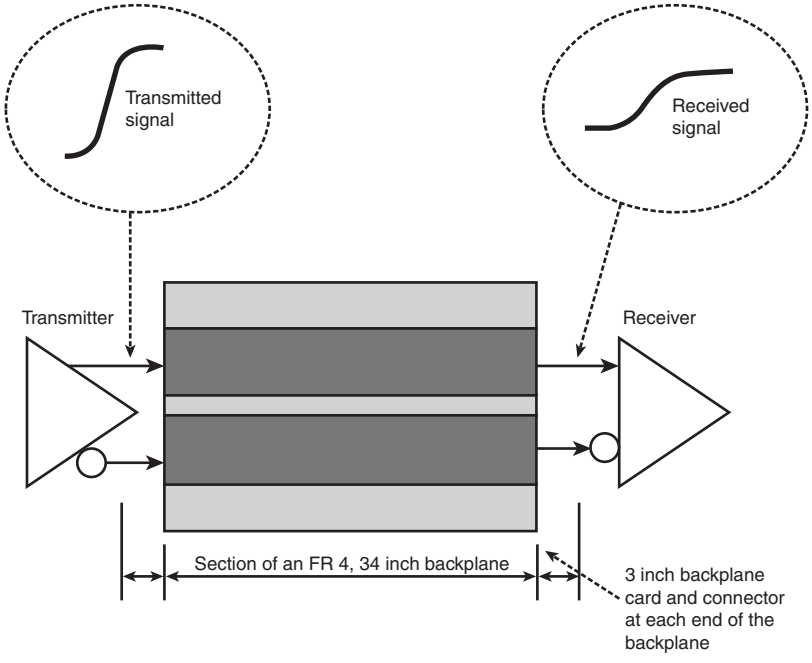

Figure 1-2 The particularly demanding case of a 40-inch backplane where the signals are transmitted through a PCB stripline that is constructed with FR4-type PCB material.

Figure 1-3 shows the dominant high-frequency dielectric loss effect in a 40-inch PCB stripline. The loss is caused by the capacitance formed by the trace and ground plane with a dielectric constructed with PCB-type FR4 materials. All PCB laminate materials have a specific dielectric constant, which will affect the impedance of the trace, especially at high frequencies, where the trace behaves as a transmission line. The value of a PCB dielectric constant is determined by comparing the capacitive effect of the PCB material to the capacitive effect of a conductive pair in a vacuum, where the vacuum has a dielectric constant of 1. In Figure 1-3, the FR4 material has a dielectric constant of about 4 to 4.7. A lower dielectric constant can allow a PCB to support a longer transmission line before the high-frequency losses become significant, but this is a simplification. Determining dielectric loss is a complex topic. Many materials are used as PCB laminates, and many have better propagation characteristics than FR4. However, the

high-performance PCB becomes too expensive for large-volume, low-cost applications. Although the extensive length of the backplane used in this example has exaggerated the loss effects, Figure 1-3 clearly shows how the dielectric loss is directly proportional to an increase in signal frequency. The skin effect is somewhat constant at –10 dB throughout the frequency range 5 GHz to 10 GHz. Lossless transmission lines are a bit of a misnomer because they consider only impedance and timing. Attenuation is considered in a second-level approximation of the line.

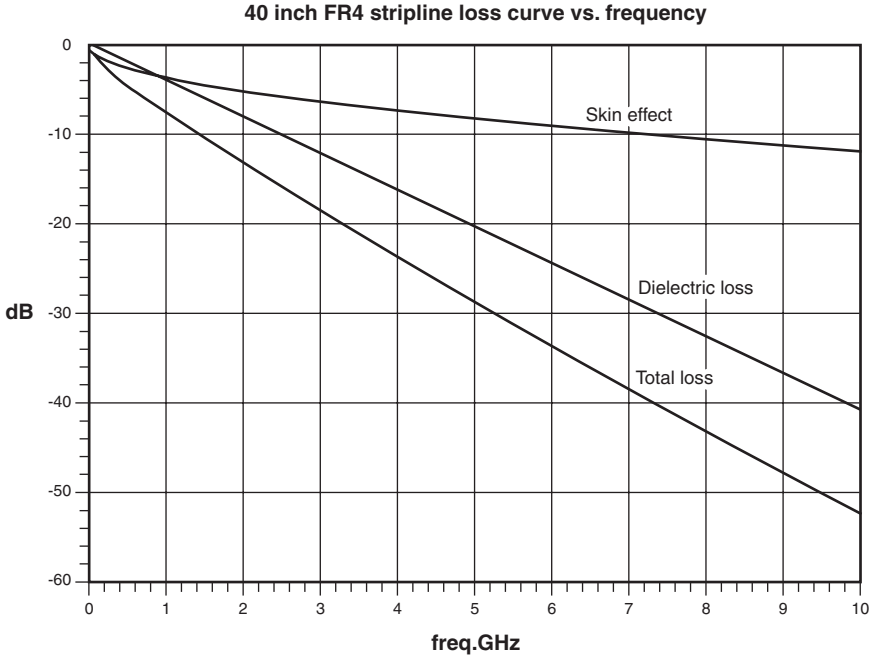

Figure 1-3 The dominant high-frequency dielectric loss effect in a 40-inch PCB stripline.

Both dielectric loss and skin effect can cause problems with intersymbol interference (ISI). The attenuation effectively prevents the signal from reaching its full amplitude within the required duration or its symbol time. As a result, the signal symbol, such as logic 1 or logic 0, spreads into the following symbol, mixing the symbols. The effect is pattern-dependent and is known as pattern-dependent jitter (PDJ) or data-dependent jitter (DDJ). If a string of data remains at the same level, such as a string of logic 0s, the energy in the signal has time to

reach its peak, allowing the data to be transmitted and received correctly. However, for a quick switching signal, with alternating logic 1s and 0s, full signal strength is not reached within each symbol period. This causes the symbols to merge and the system to malfunction.

To maintain signal integrity in a modern digital system, differential signals and integrated signal processing functions within device drivers and receivers are becoming more common. Nonetheless, differential signals demand that designers pay special attention to PCB layout. Poorly designed differential traces and terminations can cause many of the signal integrity problems associated with conventional single-ended systems. Also, an intimate knowledge of signal pre-emphasis and equalization is necessary, because incorrectly applied pre-emphasis in effect generates unwanted overshoot, crosstalk, and noise.

1.2.1 The Effects of Increasing the Drive Signal

The simple solution to overcoming signal loss, or attenuation, is to increase the signal strength to overcome the attenuation. Unfortunately, increasing the signal strengths does not solve the problem of selective loss of high frequencies, or high-frequency roll-off. Also, the PDJ would deteriorate, because each symbol would be unable to achieve its full strength within its allotted time slot. Also, as a result of the increased signal level, each signal symbol will probably spread even further into the next symbol. Increasing the signal strength also affects the noise in the system, because noise increases proportionally with the increase in signal strength. What's more, the overall power consumption of the logic driver, or transceiver, also increases as the driver buffer increases the amount of current flowing into the PCB trace. Consequently, a simple increase in signal strength is not a solution to either the dielectric or skin-effect losses. In addition, increasing the signal strength may in fact make the situation worse.

1.3 PRE-EMPHASIS

Pre-emphasis is a way to boost only the signal's high-frequency components, while leaving the low-frequency components in their original state. Pre-emphasis operates by boosting the high-frequency energy every time a transition in the data occurs. The data edges contain the signal's high-frequency content. The signal edges deteriorate with the loss of the high-frequency signal components. A simple pre-emphasis circuit can be constructed from a two-tap finite impulse response (FIR) filter. The circuitry works by comparing the previously transmitted data bit to the current data bit, where the circuit block Z^{-1} provides the delay for a single data bit. If the two bits—the delayed bit and the current bit—are the same level,

the current bit is transmitted at the normal level. If the two bits are different, the current bit is transmitted at a higher magnitude. Figure 1-4 shows the FIR filter block diagram and associated waveforms.

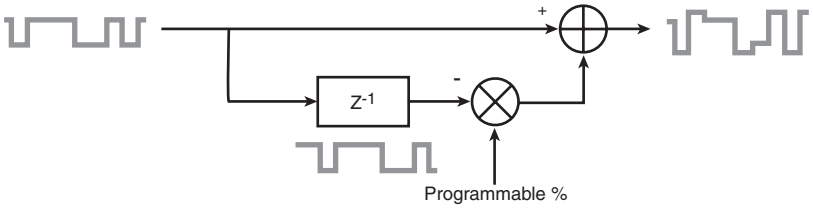

Figure 1-4 Block schematic of a FIR filter pre-emphasis circuit where the percentage of pre-emphasis is programmable.

The pre-emphasis circuit is primarily designed to overcome frequency-dependent attenuation.

1.3.1 Pre-emphasis Measurement

There are many different methods of measuring pre-emphasis. Although it is not important to follow a particular measurement method, it is important for the engineer to understand a particular definition when modeling his or her system. For example, Figure 1-5 shows the pre-emphasis measurement method used in an Altera programmable logic device. The waveform is part of a differential signal.

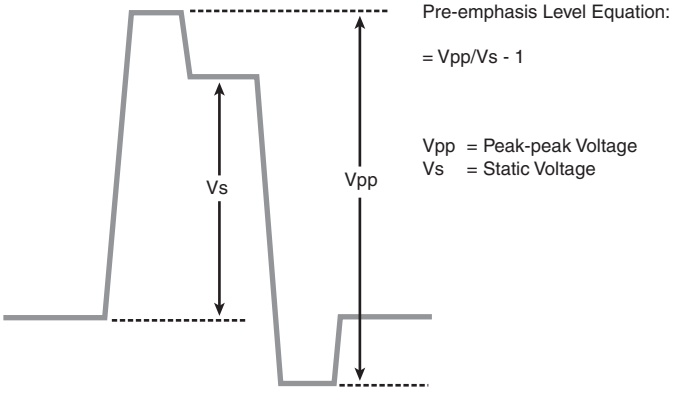

Figure 1-5 Waveform definition of pre-emphasis.

The pre-emphasis circuitry within a modern programmable device such as the Altera Stratix has an architecture that can be dynamically programmed to five

different levels of pre-emphasis. The exact value of pre-emphasis cannot be predetermined, because each device requires a percentage of pre-emphasis that is dependent on the output signal strength and transmission path characteristics. Table 1-2 shows five possible programmable pre-emphasis levels for a differential drive signal (V_{OD}) of 800 mV. In this case the internal termination is 100 ohms. The amount of pre-emphasis changes according to the transmission path parameters.

Table 1-2 Typical Percentage Pre-emphasis Levels for a Programmable Logic Device with an 800 mV Drive Signal (V_{OD})

Programmable Setting	Typical Pre-emphasis Level
1	11%
2	36%
3	65%
4	100%
5	140%

1.3.2 Receiver Equalization

An alternative to pre-emphasis is receiver equalization, which provides functionality in the receiver to help overcome the high-frequency signal losses of the transmission medium. Receiver equalization acts as a high-pass filter and amplifier to the data as it enters the receiver. In effect, equalization distorts the received data, correcting the distortion of the signal resulting from the high-frequency losses. This allows the receiver to rebuild the signal and interpret it successfully. External receiver equalization can be implemented with external filter networks. However, these filters require extra components with added PCB tracks and PCB stubs that require careful design if signal integrity problems are to be avoided. Furthermore, a fixed filter circuit is difficult to adapt for differing loss.

Some modern digital devices include equalization within the receiver. In a number of programmable logic devices, the equalization function is dynamically controllable. The equalization setting typically depends on the application and environment. For example, the receiver equalization would be up to 9 dB of gain for a 40-inch FR4 backplane. Moreover, data dispersion can be overcome when an equalizer is designed to cut off unwanted frequency components that spread symbols. The equalizer brings the symbols back into shape and time, thereby minimizing or eliminating PJD.

1.3.3 Maintaining Signal Integrity in Legacy Systems

The expansion of high-speed interfaces has led to some dilemmas because designers generally need to use legacy systems to support existing components or interfaces and reduce the cost of replacing an entire infrastructure. For example, this means that backplanes designed to operate at 1 Gbps are now required to run at 2.5 Gbps and faster to support existing and new components. In some circumstances this may be possible with the combined use of pre-emphasis and equalization. Equalization can compensate for many of the issues of the legacy backplane, such as narrow PCB tracking, which typically suffers from increased signal attenuation as transmission frequencies increase. However, equalization is relatively new as an integral part of logic receivers, so it is possible that legacy cards plugged into a system will not include equalization. This means that higher levels of pre-emphasis are required to ensure reliable communications. It is not uncommon for pre-emphasis levels in excess of 100% to be used in legacy applications. Normally each system and application require a unique setting for pre-emphasis, equalization, and drive strength. It is therefore important to model the entire data communication path using accurate model descriptions for the PCB interconnect and transceiver interfaces to ensure that the entire system is matched from within the driver to the internal receiver circuitry. In conclusion, it is clear that skin effect and dielectric loss can cause significant attenuation to the high-frequency content of signals. For example, this can impact the success of communicating high-speed data via a conventional FR4 PCB. The use of pre-emphasis and equalization can help the signal integrity of a transmission path, provided that you carefully select the parameters.

1.3.4 Simultaneous Switching Outputs

Although high-speed data rates correlate with high-frequency signals, it is the signal edge timing that has the most detrimental effect on signal integrity. This is particularly true as systems migrate to dense, highly integrated, high-speed switching systems where typically hundreds of pins are switching with edge rise and fall times that generally are faster than 500 picoseconds. A consequence of a large number of quickly switching device connections is the unstable power supply voltage, where greater power demanded in short time periods causes transient disturbances. Traditionally designers have decoupled power supplies to minimize transient charges and stabilize power sources. Nevertheless, new high-speed, high-density systems require very careful design to minimize power supply transients, or the result is a phenomenon called simultaneous switching noise (SSN).

Moreover, as digital circuitry increases in speed and output switching times decrease, higher transient currents occur within the device output circuits as the effective output load capacitors discharge. This sudden flow of current exits the device through internal inductances to a PCB ground plane. This causes a transient voltage to develop, which is a voltage difference between the device output and the board ground. This is known colloquially as ground bounce, but in actuality an unwanted signal return current causes transient variations in the ground voltage. The signal return currents or bounce effect can cause an output low ground signal to be seen as a high-output signal by other devices in the system. You can reduce unwanted signal return currents by following a number of classic design rules. Nonetheless, a number of programmable device manufacturers now provide pin slew rate control, which allows the designer to slow an edge rate and therefore reduce ground current transients or ground bounce effects. Additionally, most modern devices include multiple power and ground pins. This allows the designer to locate a high-speed input or output pin close to a ground pin to reduce the effects of simultaneous switching outputs (SSOs).

The challenges of high-speed design require some additional effort to ensure signal integrity. This can be achieved by following some simple analog design rules and by using careful PCB layout techniques. Nonetheless, contemporary integrated circuit manufacturers are providing many features to compensate for PCB anomalies and support high-speed design. Programmable slew rate control and programmable on-chip termination technology are helping make the designers' work somewhat easier. However, programmable pre-emphasis, equilisation, and slow rate control only work in stable systems, and they are not a substitute for good design practice.

1.4 THE NEED FOR REAL-TIME TEST AND MEASUREMENT

As data rates increase, it is ever more difficult to detect and debug noise and signal aberrations in a prototype or production model. A rigorous regime of signal integrity measurements can provide the means for the engineer to trace sources of noise or glitches and provide the wherewithal for him or her to eradicate the root causes of signal aberrations. Apart for the requirements of test and debug, there is often a need for a product to operate in a global market, and many designers are increasingly concerned with compliance measurements. Complying with industry standards ensures interoperability among system elements, where discrete system components from various manufacturers successfully interconnect and communicate. Compliance measurements usually entail a series of prescribed acquisition and analysis steps, which are carried out on a completed product. However, successful compliance design and testing often depend on eliminating signal

integrity problems in the early phases of a design. Ideally the designer or test engineer locates signal integrity errors during the initial simulation of a product's development, as mentioned previously. However, a primary aim of the SI engineer must be to understand the role played by electronic bench instrumentation and portable test equipment in the pursuit of signal integrity. Today, the SI engineer needs to understand the latest methodologies that are used to achieve signal integrity. They are founded on the interrelationships between device simulation, circuit simulation, and real-time test and measurement. Apart from the fundamental theoretical concepts, or abstract mathematical models, the engineer must understand the practical issues that lie at the heart of signal integrity engineering. Remember that simulation is only a model, and a model is an imperfect replica of a real component or system. In many cases there is no substitute for a prototype and a real-time test phase, which at the very least will allow real-world data to be fed into the simulation.

One of the most demanding aspects of modern digital system design and debug is the successful measurement of the analog content of high-speed digital signals, where complex multilayer boards, the high density of interconnects, and highly integrated systems have made successful probing an exact science. Ideally the measurement system bandwidth, including that of the probe, should be at least three times the frequency of the signals to be observed. Above all, the edges of a digital signal must be considered when making a measurement. For instance, a standard 5 Gbps data rate serial bus signal requires at least a 15 GHz measurement bandwidth and a true differential probe that has a rise time of less than 35 picoseconds. Otherwise, the probe would at best fail to show any analog aberrations and at worse would simply stop the bus. Ultra-low loading and a diversity of attachment methods are needed to ensure fast, positive connection with minimal effect on signals.

Experienced engineers know that signal integrity is the result of constant vigilance during the design, simulation, and real-time test processes. It's all too easy for signal integrity problems to get compounded as a design evolves. An aberration that goes unnoticed in the early stages of a design can cause erratic behavior in a product, which can entail many hours of demanding test and debug to correct. An experienced engineer plans ahead and pays special attention to signal integrity issues in the early stages of a design. In particular, the engineer decides during the specification phase what signals need to be probed and how access is to be provided to signals of interest.

1.4.1 Timing Budgets and the Analog View

A signal integrity problem may first seem to be a misplaced digital pulse or timing error. However, the cause of the problem in a high-speed digital system can often be related to the analog characteristics of a digital signal. In many cases a digital problem can be fairly easy to pinpoint when an errant digital signal is successfully probed and the analog representation of the flawed digital signal is

exposed. Analog characteristics can become digital faults when low-amplitude signals turn into incorrect logic states, or when slow rise times cause pulses to shift in time. An innovative test and measurement technique is to use a logic analyzer and oscilloscope in unison. The test technique is to show a digital pulse stream along with a simultaneous analog view of the same pulses, where both waveforms are shown on a single display, as shown in Figure 1-6. This technique is becoming a standard debug methodology, and the debug method is frequently the first step in tracking down an SI problem. Figure 1-6 is a real-time view of a digital data stream exhibiting a timing error. The analog view clearly shows the cause of the timing error as an amplitude aberration on the trailing edge of the digital pulse.

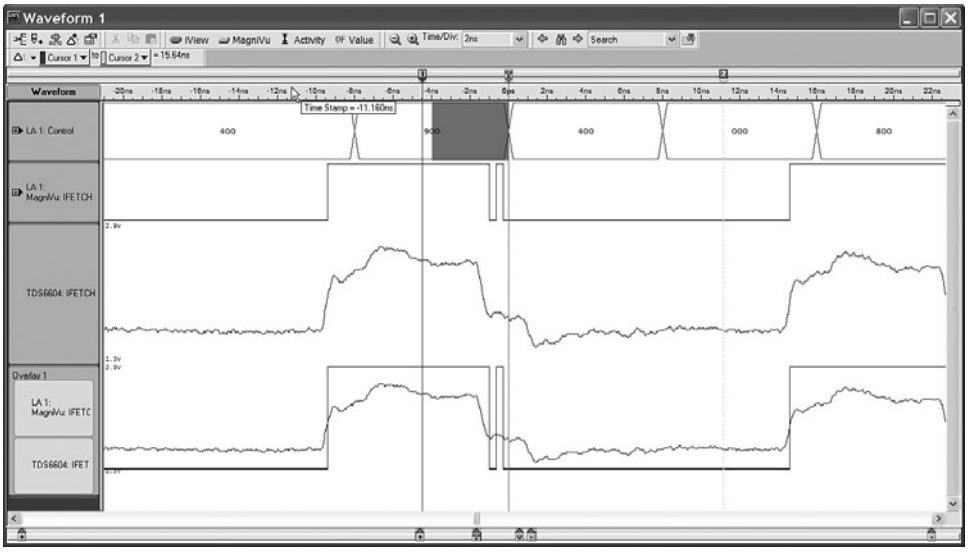

Figure 1-6 Digital pulse streams and a simultaneous analog view of the same pulses, where both waveforms are shown on a single display.

In any discussion of signal integrity, signal transitions deserve special attention. For example, the timing diagram shown in Figure 1-7 shows two digital inputs feeding an ordinary AND gate. The gray trace for Input A shows the correct pulse. Superimposed on it is an analog view, a distorted signal shown as a black trace, of the actual signal. Due to its slow rise time, the actual signal does not cross the required threshold value until much later than it should. Consequently, the output pulse from the AND gate is significantly narrower than it should be; the correct pulse width is shown in gray. The integrity of the signal on Input A is very poor, with serious consequences for the timing of digital elements elsewhere in the system. This type of SI problem typically causes timing errors in subsequent logic steps. Such dilemmas often require careful analysis for successful debugging.

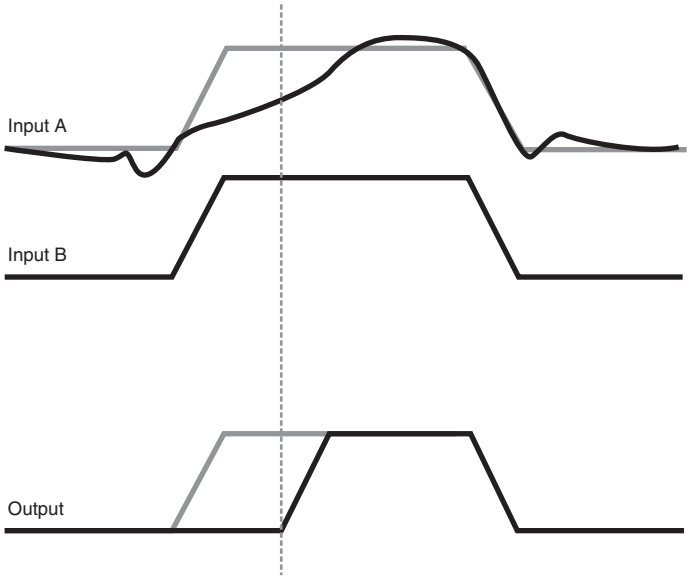

Figure 1-7 The gray traces have the required signal characteristics, and the black traces are the actual analog signals; this simple example is an AND gate. Clearly Input A has a slow rise time and crosses the logic threshold too late, which results in an output timing error.

Suppose the output went on to become part of a memory address. The short pulse might cause the memory to see logic 0 where logic 1 should exist and therefore select a different memory location than the one that is expected. The content of that location, of course, is inappropriate for the transaction at hand. The end result is an invalid transaction, raising the all-too-familiar question of whether the bug is inherently a hardware or software failure.

Slow signal transition edges can lead to intermittent system faults even if they are not causing repeatable errors. Timing budgets in high-speed digital systems allow very little time for signal rise and fall transitions. Setup and hold times have scientifically decreased in recent years. Modern electronic memory systems are a typical example of setup and hold times in the low hundreds of picoseconds. Slowly changing edges can leave too little margin in the timing budget for data transactions to be valid and stable, as implied in Figure 1-8. The relationships shown in Figure 1-8 are exaggerated to emphasize the concept. These two simple examples show some of the potential unwanted effects resulting from too slow an edge transition. The majority of SI problems in high-speed digital systems commonly are related to the effects of fast switching edges and their associated high-frequency signal content.

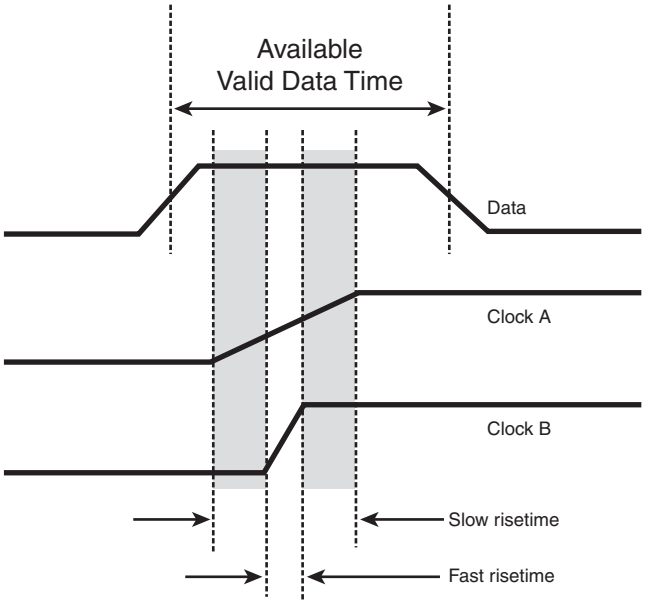

Figure 1-8 Although not to scale and exaggerated, Clock A is seen to clock the data with a slow rise time, leaving the system with greater susceptibility to noise, with the possibility of double clocking and metastability. The faster Clock B edge results in improved reliability, and the timing budget is improved.

One of the major timing concerns in SI engineering is the setup and hold timing values that are specified for clocked digital devices. Setup and hold timing are at the heart of application-specific integrated circuit (ASIC) functional verification measurements. Setup time is defined as how long the data must be in a stable and valid state before the clock edge occurs. Hold time is how long that data state must remain stable and valid after the clock edge. In the high-speed digital devices used for computing and communications, both setup and hold timing values may be as low as a few hundred picoseconds.

Transients, edge aberrations, glitches, and other intermittencies can cause setup and hold violations. Figure 1-9 is a typical setup and hold timing diagram. In this example, the data envelope is narrower than the clock. This emphasizes the fact that with today's high-speed logic, transition times and setup and hold values can be very brief, even when the cyclical rate or signal frequency is relatively slow.

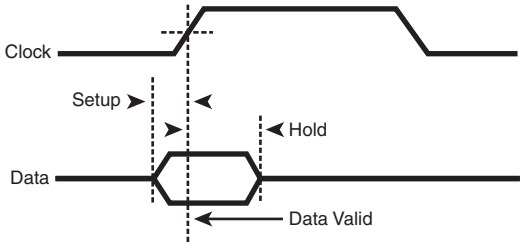

Figure 1-9 A typical setup and hold timing diagram.

There are three common approaches to evaluating a device's setup and hold performance. A number of other timing parameters can be verified using either a low-speed or high-speed functional test:

- Low-speed tests are a fairly coarse functional verification procedure, but they are often adequate. In some cases it is not necessary to take quantitative measurements of the actual setup and hold values that are specified on a device's data sheet. If a device can tolerate a broad clock placement range, the timing test may be as simple as running a low-speed functional data pattern, adjusting the position of the clock edge relative to the data, and observing the results on an oscilloscope. The oscilloscope trigger normally is set to a 50% level, which shows timing information before and after the clocking edge. A device tends to become metastable as it exceeds its setup and hold timing limitations, even at low speed. Metastability is an unpredictable state in which a device output may switch to either a logic 1 or logic 0 without any apparent regard for the logical input conditions. Similarly, excessive signal jitter may appear on the output when setup and hold tolerances are violated.
- High-speed tests typically require bursts of high-speed data, where the burst forms a functional test that exercises the device at rates approximating its intended operational frequency or higher. A signal source is used to deliver a block of data to the device under test (DUT) at data rates that are much higher than the basic low-speed functional test. However, this test process is still one of empirically finding a range of setup and hold values and specifying the system clock placement accordingly. Using a data generator with a repetitive data pattern, the recurring skew problems associated with high-speed setup and hold violations often can be isolated and rectified.
- Device self-test is an option on some contemporary source-synchronous and high-speed serial (HSS) components. The transmitter sends a known pattern to the receiver, which slides the location of the clock or strobe in time until data errors occur. The difference between the nominal clock or strobe location and the onset of data errors indicates how much margin exists in that interface.

1.4.2 Eye Diagrams

Eye diagrams have become one of the cornerstones of SI test and compliance measurements. The eye diagram is a regulatory measurement for validation and compliance testing of various industry-standard digitally transmitted signals. An eye diagram is a display that typically is viewed on an oscilloscope. It can efficiently reveal amplitude and timing errors in high-speed digital signals. The eye diagram shown in Figure 1-10 is built by overlaying digital signal waveform traces from successive logic signal cycle periods, or unit intervals. Waveform A in Figure 1-10 is a digital signal with exaggerated timing errors. Waveform B is the clock signal recovered from the signal shown as waveform A. Eye diagram A is generated by dividing waveform A into individual cycle periods and overlaying each cycle of waveform A. A number of measurements can be made on an eye diagram to quantify signal quality. One such measurement is the eye opening that relates to the signal timing, which is shown in Figure 1-10 as a gray area. Eye diagram B is produced from the recovered clock signal and shows the ideal eye opening, which indicates improved signal timing. Eye diagrams display serial data with respect to a clock that normally is recovered from the data signal using either hardware or software tools. In the eye diagram shown in Figure 1-11, the clock is recovered by a hardware-based reference or "golden" phase locked loop (PLL). The diagram displays all possible transition edges, both positive-going and negative-going, and both data states in a single window. The result is an image that somewhat resembles an eye, as shown in Figure 1-11.

In an ideal world, each new trace would line up perfectly on top of those that came before it. Also, the eye diagram would be composed of narrow lines representing the superimposed logic 1s and logic 0s. In the real world, signal integrity factors such as noise and jitter cause the composite trace to blur as it accumulates the logic 1s and logic 0s. The gray regions in Figure 1-11 have special significance; they are the violation zones used as mask boundaries during compliance testing. A compliance mask typically is produced by the instrument manufacturer in association with a standards body. In this case the gray polygon in the center defines the area in which the eye is widest. This encompasses the range of safe decision points for extracting the data content, the binary state, from a logic signal. The upper and lower gray bars define the signal's amplitude limits. If a signal peak penetrates the upper bar, for instance, it is considered a "mask hit" that will cause the compliance test to fail, although some standards may tolerate a small number of mask hits. More commonly, noise, distortion, transients, or jitter cause the trace lines to thicken. The eye opening shrinks, touching the inner gray polygon. This too is a compliance failure, because it reveals an intrusion into the area reserved for evaluating the logic state of the data bit. The compelling advantage of an eye diagram is that it enables a quick visual assessment of signal quality.

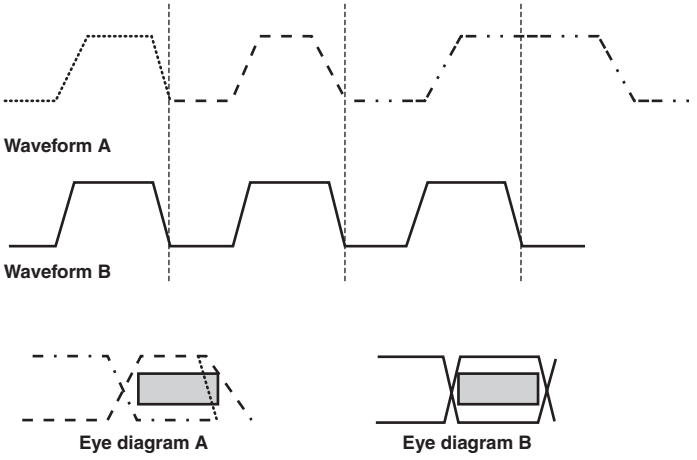

Figure 1-10 Eye diagram formation.

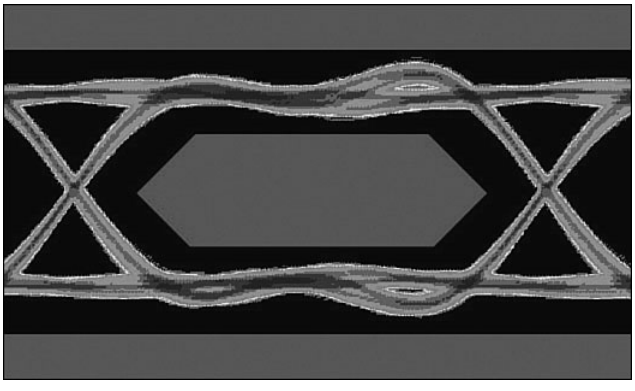

Figure 1-11 An oscilloscope display showing a hardware-generated eye diagram. The gray areas are the mask violation zones.

1.4.2.1 Simulated Eye Diagrams

Although the ideal eye diagram measurement is carried out in real time, it is sometimes impossible to extract. For example, a programmable logic device receiver equalization circuit typically is embedded within the device. However, it

is possible to simulate the results using SPICE modeling. The upper part of Figure 1-12 shows a simulated eye diagram into the programmable logic device receiver after traversing a 40-inch backplane, without pre-emphasis. The lower part of Figure 1-12 shows the same signal after it has passed through receiver equalization.

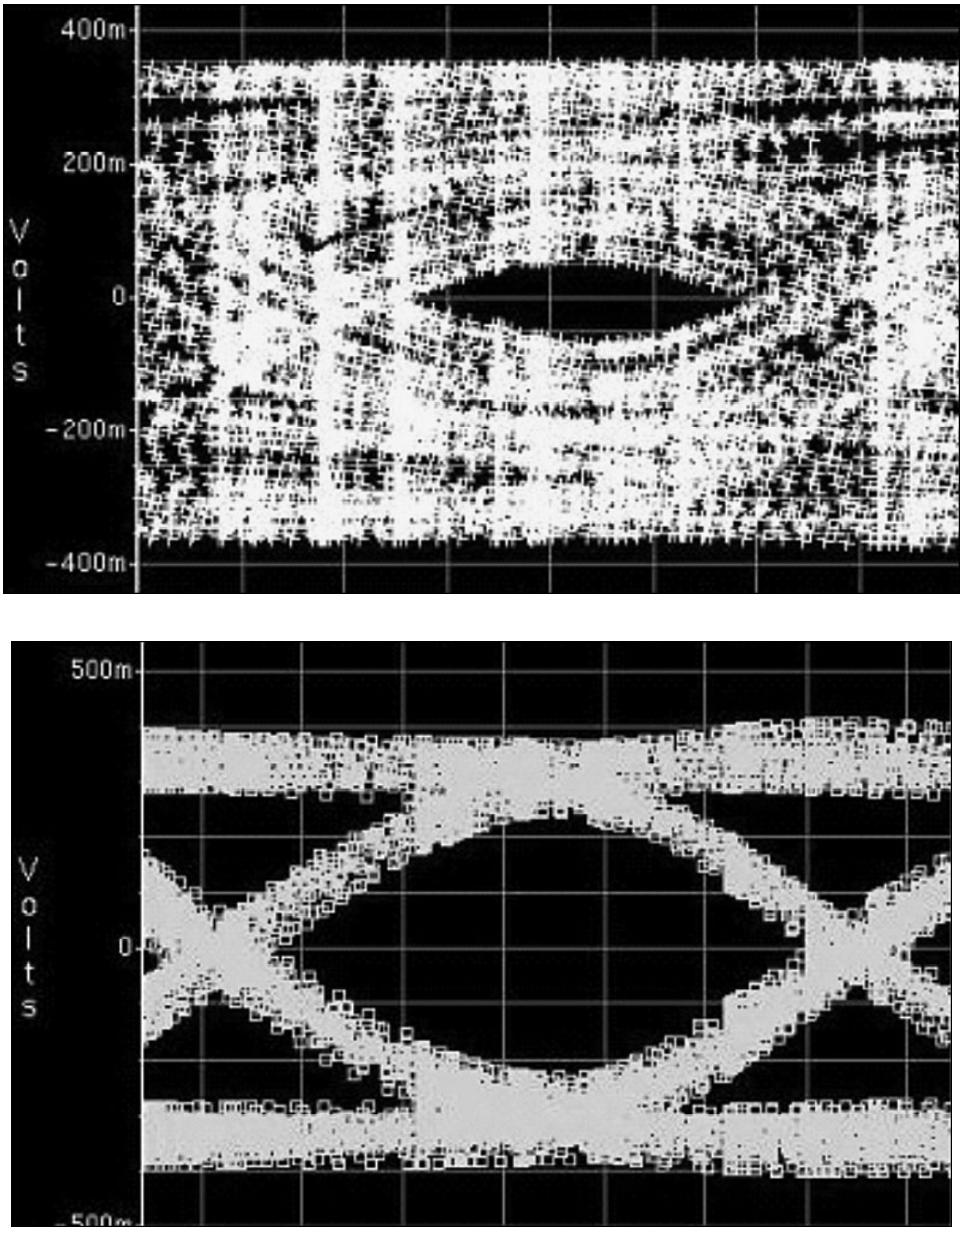

Figure 1-12 SPICE simulated eye diagrams.

The eye diagrams shown in Figure 1-12 are a clear example of how an equalization circuit significantly improves the quality of a received signal, allowing reliable detection of logic states. Another example, Figure 1-13, shows the effects of too much pre-emphasis. You must select the optimum settings for both pre-emphasis and equalization. Overcompensation can cause additional issues within the system. It adds extra jitter, which closes the eye, making it impossible for the receiver to interpret the information. Figure 1-13 clearly shows the effect of adding too much pre-emphasis and equalization.

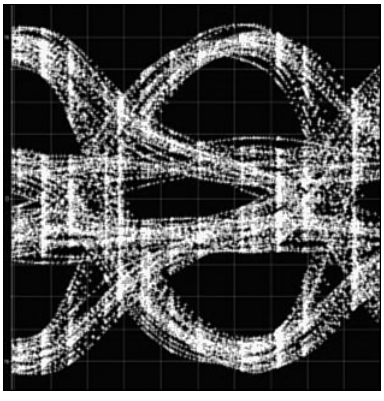

Figure 1-13 An eye diagram showing the exaggerated effects of too much pre-emphasis.

1.4.2.2 Real-Time Eye Diagrams

Real-time eye diagram debug methodologies often provide a shortcut that lets you quickly detect and correct SI problems. For example, some modern high-performance logic analyzers are combined with an oscilloscope and host troubleshooting tools that bring analog eye diagram analysis to the logic analyzer screen. The eye diagram is a real-time visualization tool that typically allows the designer to observe the data valid window, and general signal integrity, on clocked buses. This test methodology is a required compliance testing tool for many of today's buses, particularly the high-speed serial buses, but any signal line can be viewed as an eye diagram. Moreover, the logic analyzer eye diagram analysis can show wide parallel bus performance. It integrates hundreds of eye diagrams into one view that encompasses the leading and trailing edges of both positive-going and negative-going pulses that compose the bus signals. Figure 1-14 is an eye diagram where the contents of twelve address bus signals are superimposed. The benefit of observing bus lines simultaneously with an eye diagram is that it presents all possible logic transitions in a single view and allows fast assessment of the bus.

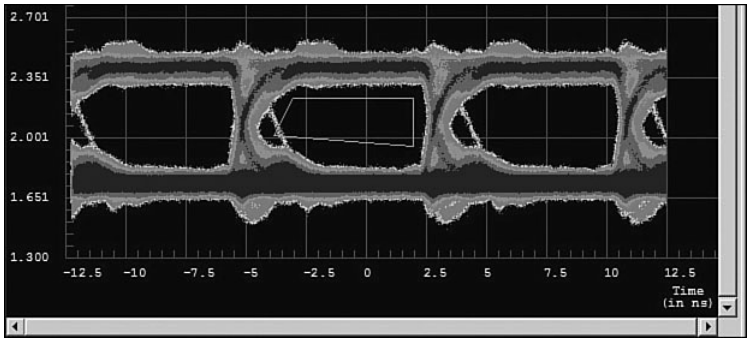

Figure 1-14 An eye diagram that simultaneously shows twelve bus signals.

An eye diagram can reveal analog problems in high-speed digital signals, such as slow rise times, transients, and incorrect logic levels. Figure 1-14 shows the performance of 12 parallel bus lines. The error encroaching into the mask is caused by an incorrect rise time in one of the bus signals.

The eye diagram shown in Figure 1-14 reveals an anomaly in the signals, which typically is shown in a distinctive color; the color indicates a relatively infrequent transition. In this example at least one of the signals has an edge that is outside the normal range. The mask feature built into the instrument helps locate the specific signal causing the problem. By drawing the mask in a particular way, such that the offending edge penetrates the mask area, the relevant signal can be isolated, highlighted, and brought to the front layer of the image. The result is shown in Figure 1-15, in which the flawed signal has been brought to the front of the display and highlighted in white. In this example the instrument identifies the aberrant edge and indicates a problem on the A3 (0) address bus signal. The origin of this particular problem is actually crosstalk, where the edge change is being induced by signals on an adjacent PCB trace.

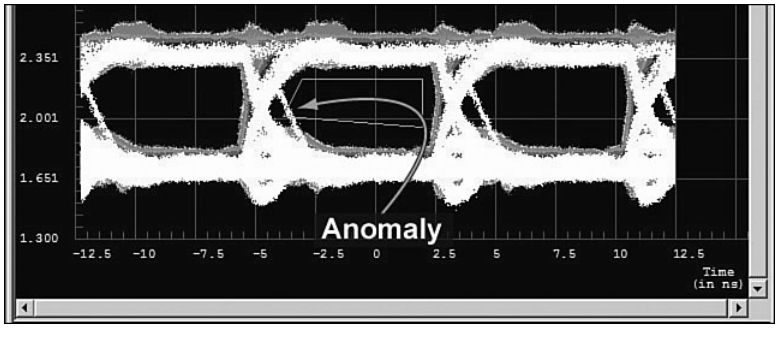

Figure 1-15 The flawed signal has been brought to the front of the display.

The real-time functional verification and troubleshooting phase of the design in this example has described how a common SI problem is detected and resolved. The logic analyzer is often the first line of defense when testing digital functionality. However, digital problems can stem from analog signal issues, including edge degradation due to improper termination or crosstalk, as demonstrated here. By teaming the logic analyzer with an oscilloscope and evaluating time-correlated digital and analog signals on the same screen, you can see problems affecting either domain using an eye diagram. Today, low-cost mixed signal instruments with real-time eye diagram capabilities are used to test and debug high-performance digital systems. Specialist oscilloscopes are used for the more demanding tasks of compliance and interoperability tests.

CONCLUSION

The design process has multiple steps, each with its own particular set of SI challenges in which the customer requirement is translated into specifications, simulation, and real-time measurement needs. This chapter has illustrated some of the more important points of SI engineering to promote the notion that today's SI engineers need innovative thinking if they are to keep pace with the digital bandwidth race. Increasing a system's operating rate is not simply a matter of designing a faster clock. As frequency increases, simulation issues become complex, and PCB traces on a circuit board become significantly more complicated. As frequencies increase, the trace begins to act like a capacitor. At the highest frequencies, trace inductance plays a larger role. All of these characteristics can adversely affect signal integrity. At today's clock frequencies, which are in the hundreds of megahertz and above, every design detail is important. Catching problems early and minimizing rework enables a new product to reach its market on time. History has shown that in the highly competitive embedded systems marketplace, the majority of profits go to the product that is first to market and that delivers premium performance.

This chapter has defined signal integrity problems as phenomena that can compromise a signal's ability to convey binary information. In real digital devices, these binary signals were shown to have analog attributes that result from the complex interactions of many circuit elements. These range from driver outputs to signal path transmission, terminations, and digital receivers. It is critical that today's designers follow a simulation, test, and debug strategy that is appropriate for the wide range of performance levels found in modern digital systems. The primary aim of this book is to advise you on and encourage the use of a range of best practices in the simulation, real-time test, and measurement aspects of SI

engineering. However, we also hope that you will discover how to design digitally and think in analog. Moreover, the majority of simulation, test, and measurement tools show voltage and time. In a number of SI issues, the engineer must think in terms of signal currents, ground voltage, and electromagnetic effects, along with trace distance. Today the SI engineer has an ever greater dependence on automated real-time measurements and built-in signal analysis tools, such as eye diagram generation software. And what is most exciting is the interplay of simulation and real-time test. Device manufacturers are working with instrument providers to implant simulation models in benchtop instruments, allowing the real-time measurement and analysis of a partially populated system. In terms of digital design, simulation, test, and measurement, we live in exciting times. But as with any picture, the observers can see in the picture only what their experiences in life have equipped them with. Put another way, this book cannot teach experience. It can only point the way to good practice and, like a good companion, give sound advice.

Index

Symbols

A

B

C

D

E

F

G

H

I

J–K

L

M

N

O

P

Q–R

S

T

U–V

W–Z