CHAPTER 1

Signal Integrity Is in Your Future

"There are two kinds of designers, those with signal-integrity problems and those that will have them."

on a white board at a large systems company

Ironically, this is an era when not only are clock frequencies increasing and signal integrity problems getting more severe, but the time design teams have available to solve these problems and design new products is getting shorter. Product design teams have one chance to get a product to market; the product must work successfully the first time. If identifying and eliminating signal integrity problems isn't an active priority as early in the product design cycle as possible, chances are the product will not work.

TIP As clock frequencies increase, identifying and solving signal-integrity problems becomes critical. The successful companies will be those that master signal-integrity problems and implement an efficient design process to eliminate these problems. It is by incorporating new design rules, new technologies, and new analysis tools that higher performance designs can be implemented and meet ever-shrinking schedules.

In high-speed products, the physical and mechanical design can affect signal integrity. An example of how a simple two-inch-long section of trace on a printed

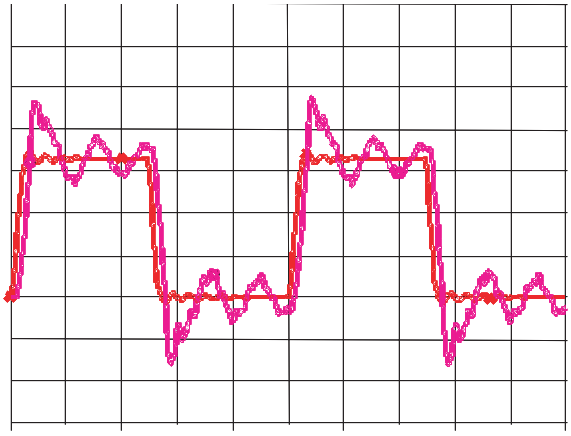

Figure 1-1 100-MHz clock waveform from a driver chip, when there is no connection (smooth plot) and when there is a two-inch length of PCB trace connected to the output (ringing). Scale is 1 v/div and 2 nsec/div, simulated wtih Mentor Graphics Hyperlynx.

circuit board (PCB) can affect the signal integrity from a typical driver is shown in Figure 1-1.

The design process is often a very intuitive and creative process. Feeding your engineering intuition about signal integrity is critically important to reaching an acceptable design as quickly as possible. All engineers who touch the product should have an understanding of how they influence the performance of the overall product. By understanding the fundamental principles of signal integrity at an intuitive and engineering level, every engineer involved in the design process can see the impact of their decisions on system performance. This book is about the fundamental principles needed to understand signal-integrity problems and their solutions. The engineering discipline required to deal with these problems is presented at an intuitive and a quantitative level.

1.1 What Is Signal Integrity?

In the good old days of 10-MHz clock frequencies, the chief design challenges in circuit boards or packages were how to route all the signals in a two-layer board and how to get packages that wouldn't crack during assembly. The electrical properties of the interconnects were not important because they didn't affect system performance. In this sense, we say that “the interconnects were transparent to the signals.”

A device would output a signal with a rise time of roughly 10 nsec and a clock frequency of 10 MHz, for example, and the circuits would work with the crudest of interconnects. Prototypes fabricated with wire-wrapped boards worked as well as final products with printed circuit boards and engineering change wires.

But clock frequencies have increased and rise times of signals have decreased. For most electronic products, signal-integrity effects begin to be important at clock frequencies above about 100 MHz or rise times shorter than about 1 nsec. This is sometimes called the *high-frequency* or *high-speed regime*. These terms refer to products and systems where the interconnects are no longer transparent to the signals and, if you are not careful, one or more signal integrity problems arise.

Signal integrity refers, in its broadest sense, to all the problems that arise in high-speed products due to the interconnects. It is about how the electrical properties of the interconnects, interacting with the digital signal's voltage and current waveforms, can affect performance.

All of these problems fall into one of the following categories:

1. Timing

2. Noise

3. Electromagnetic interference (EMI)

Timing is a complicated field of study by itself. In one cycle of a clock, a certain number of operations must happen. This short amount of time must be divided up and allocated, in a budget, to all the various operations. For example, some time is allocated for gate switching, for propagating the signal to the output gate, for waiting for the clock to get to the next gate, and for waiting for the gate to read the data at the input. Though the interconnects affect the timing budget, timing is not covered in this book. We refer any interested readers to a number of other books listed in the reference section for more information on this topic. Instead, this book concentrates on the effects of the interconnects on the other generic high-speed problem, too much noise.

We hear about a lot of signal-integrity noise problems, such as ringing, ground bounce, reflections, near-end cross talk, switching noise, non-monotonicity, power bounce, attenuation, capacitive loading. All of these relate to the electrical properties of the interconnects and how the electrical properties affect the waveforms of the digital signals.

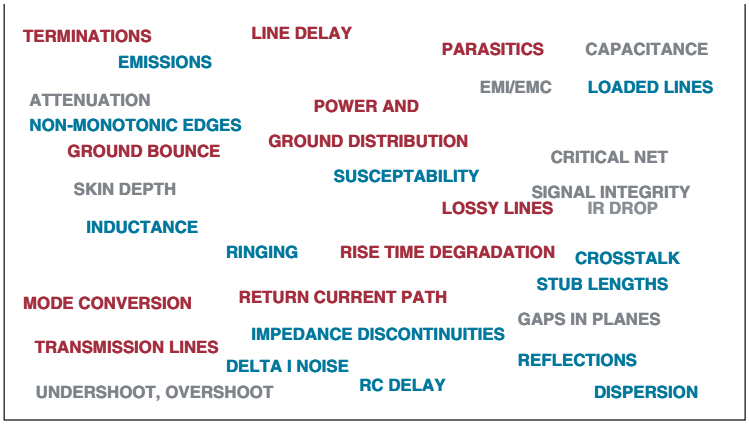

Figure 1-2 The list of signal-integrity effects seems like a random collection of terms, without any pattern.

It seems at first glance that there is an unlimited supply of new effects we have to take into account. This confusion is illustrated in Figure 1-2. Few digital-system designers or board designers are familiar with all of these terms as other than labels for the craters left in a previous product-design minefield. How are we to make sense of all these signal-integrity problems? Do we just keep a growing checklist and add to it periodically?

All the effects listed above, associated with signal-integrity noise problems, are related to one of the following four unique families of noise sources:

1. Signal quality of one net

2. Cross talk between two or more nets

3. Rail collapse in the power and ground distribution

4. Electromagnetic interference and radiation from the entire system

These four families are illustrated in Figure 1-3. Once we understand the origin of the noise associated with each of these four families of problems, the general solution for finding and fixing the problems in each family will become obvious. This is the power of being able to classify every signal-integrity noise problem into one of these four families.

These problems play a role in all interconnects, from the smallest on-chip wire to the cables connecting racks of boards and everywhere in-between. The

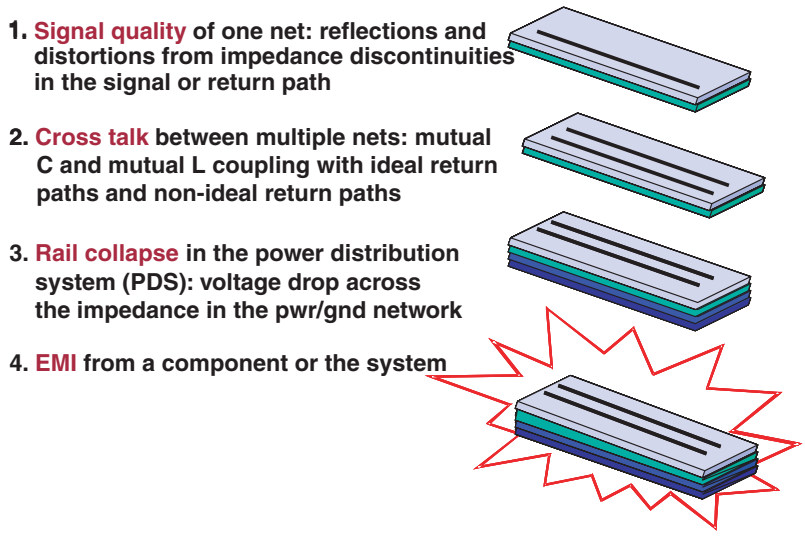

Figure 1-3 The four families of signal-integrity problems.

principles and effects are the same. The only differences in each physical structure are the specific geometrical feature size and the material properties.

1.2 Signal Quality on a Single Net

A net is made up of all the metal connected together in a system. For example, there may be a trace going from a clock chip's output pin to three other chips. Each piece of metal that connects all four of these pins is considered one net. In addition, the net includes not only the signal path but also the return path for the signal current. Signal quality on a single net depends as much on the physical features of the signal trace as on the return path. An example of two different nets on a circuit board is shown in Figure 1-4.

When the signal leaves the output driver, the voltage and the current, which make up the signal, see the interconnect as an electrical impedance. As the signal propagates down the net, it is constantly probing and asking, "what is the instantaneous impedance I see?" If the impedance the signal sees stays the same, the signal continues undistorted. If, however, the impedance changes, the signal will reflect from the change and continue through the rest of the interconnect distorted. If there are enough impedance changes, the distortions can cause false triggering.

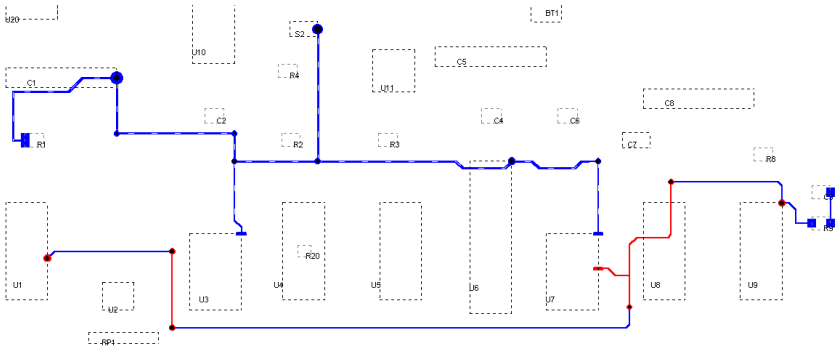

Figure 1-4 Example of two nets on a circuit board. All metal connected together is considered one net. Note: one net has a surface-mount resistor in series. Routed with Mentor Graphics Hyperlynx.

Any feature that changes the cross section or geometrical shape of the net will change the impedance the signal sees. We call any feature that changes the impedance a discontinuity. Every discontinuity will cause the signal to be distorted from its original pristine shape, to some extent. For example, some of the features that would change the impedance the signal sees include the following:

1. A line-width change
2. A layer change
3. A gap in return-path plane
4. A connector
5. A branch, tee, or stub
6. The end of a net

These impedance discontinuities can arise from the cross section, the topology of the routed traces, or the added components. The most common discontinuity is what happens at the end of a trace, which is usually either a high-impedance open or a low impedance at the output driver.

TIP The way to minimize the problems associated with impedance changes is to keep the impedance the signal sees constant throughout the net.

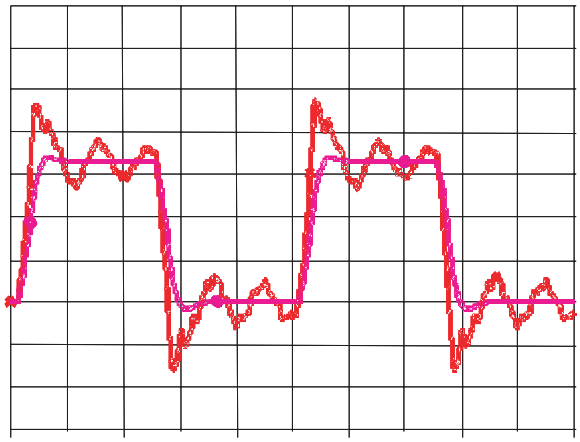

Figure 1-5 Ringing in an unterminated line and good signal quality in a source-series terminated interconnect line. The PCB trace is only two inches long in each case. Scale is 1 v/div and 2 nsec/div. Simulated with Mentor Graphics Hyperlynx.

This strategy is typically implemented by doing three things. First, use a board with constant, or "controlled," impedance traces. This usually means using uniform transmission lines. Second, use routing rules that allow the topology to maintain a constant impedance down the trace. Finally, use strategically placed resistors to manipulate the reflections and keep the received signals looking clean.

Figure 1-5 is an example of poor signal quality due to impedance changes in the same net and when a terminating resistor is used to manage the impedance changes. Often, what we think of as "ringing" is really due to reflections associated with impedance changes.

Even with perfect terminations, the precise board layout can drastically affect signal quality. When a trace branches into two paths, the impedance at the junction changes and some signal will reflect back to the source while some will continue down the branches in a reduced and distorted form. By rerouting the trace to be a daisy chain, the signal sees a constant impedance all down the path and the signal quality can be restored.

The impact on a signal from any discontinuity depends on the rise time of the signal. As the rise time gets shorter, the magnitude of the distortion will increase. This means that a discontinuity that was not a problem in a 33-MHz design may be a problem in a 100-MHz design. This is illustrated in Figure 1-6.

The higher the frequency and the shorter the rise time, the more important it is to keep the impedance the signal sees constant. One way of achieving this is by

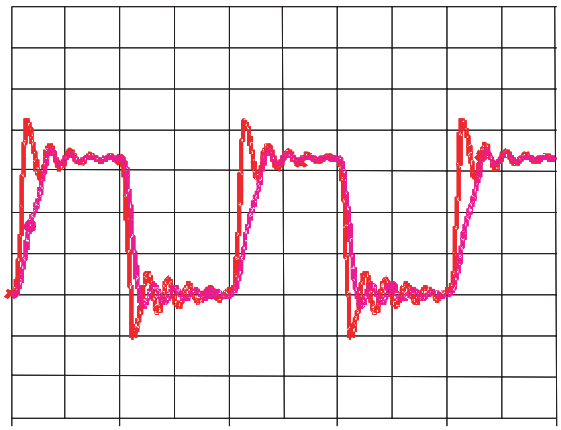

Figure 1-6 25-MHz clock waveforms with a PCB trace six inches long and unterminated. The slow rise time is a 3-nsec rise time. The ringing is from a rise time of 1 nsec. What may not have been a problem with one rise time can be a problem with a shorter rise time. Scale is 1 v/div and 5 nsec/div. Simulated with Mentor Graphics Hyperlynx.

using controlled impedance interconnects even in the packages, such as with multilayer ball grid arrays (BGAs). When the packages do not use controlled impedance, such as with lead frames, it's important to keep the leads short, such as by using chip-scale packages (CSPs).

1.3 Cross Talk

When one net carries a signal, some of this voltage and current can pass over to an adjacent quiet net, which is just sitting there, minding its own business. Even though the signal quality on the first net (the active net) is perfect, some of the signal can couple over and appear as unwanted noise on the second, quiet net.

TIP It is the capacitive and inductive coupling between two nets that provides a path for unwanted noise from one net to the other.

Cross talk occurs in two different environments: when the interconnects are uniform transmission lines, as in most traces in a circuit board, and when they are not uniform transmission lines, as in connectors and packages. In controlled-impedance transmission lines where the traces have a wide uniform return path, the relative amount of capacitive coupling and inductive coupling is comparable. In this case, these two effects combine in different ways at the near end of the

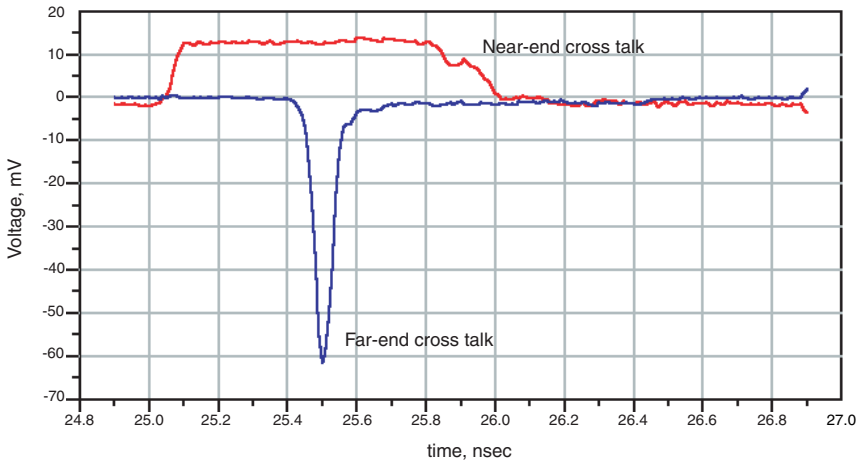

Figure 1-7 Measured voltage noise at the near end and the far end of a quiet trace when a 200-mV signal is injected in the active trace. Note the near-end noise is about 13% and the far-end noise is nearly 60% of the signal. Measurements performed with an Agilent DCA86100 with time domain reflectometer (TDR) plug-in and GigaTest Labs Probe Station.

quiet line and at the far end of the quiet line. An example of the measured near- and far-end cross talk between two nets in a circuit board is shown in Figure 1-7.

Having a uniform plane as the return path is the configuration of lowest cross talk. Anything that changes the return path from a uniform plane will increase the amount of coupled noise between two transmission lines. Usually when this happens, for example, when the signal goes through a connector and the return paths for more than one signal path are now shared by one of the pins rather than by a plane, the inductively coupled noise increases much more than the capacitively coupled noise.

In this regime, where inductively coupled noise dominates, we usually refer to the cross talk as switching noise, delta I noise, dI-dt noise, ground bounce, simultaneous switching noise (SSN), or simultaneous switching output (SSO) noise. As we shall see, this type of noise is generated by the coupled inductance, which is called mutual inductance. Ground bounce, reviewed later in this book, is really due to high mutual inductance between adjacent signal and return paths. Switching noise occurs mostly in connectors and packages. An example of SSO noise from the high mutual inductance between adjacent signal and return paths in a package is shown in Figure 1-8.

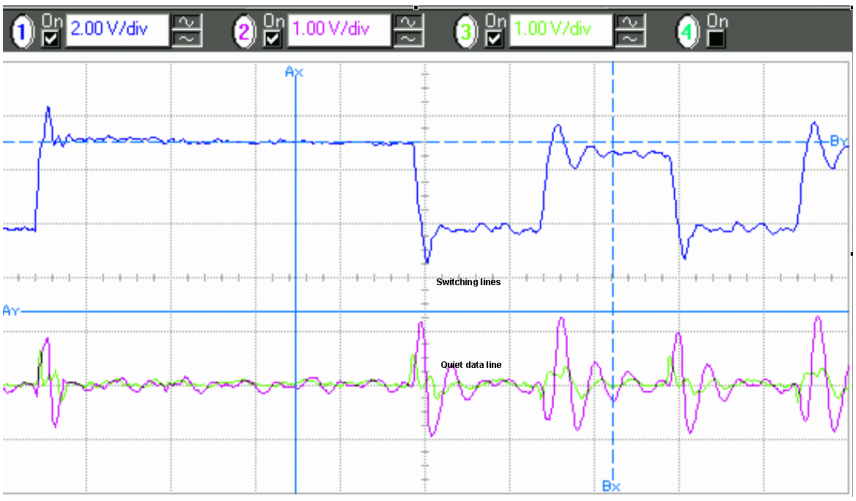

Figure 1-8 Top trace: measured voltage on active lines in a multiline buss. Bottom trace: measured noise on one quiet line showing the switching noise due to mutual inductance between the active and quiet nets in the package.

TIP SSO noise, where the coupled, or mutual inductance, dominates is becoming one of the most important issues in connector and package design. It will only get worse in next generation products. The solution lies in careful design of the geometry of paths so that the mutual inductance is minimized.

By understanding the nature of the capacitive and inductive coupling, it is possible to optimize the physical design of the adjacent signal traces to minimize the coupling. This usually can be as simple as spacing the traces farther apart. In addition, the use of lower dielectric constant material will decrease the cross talk for the same characteristic impedance lines. Some aspects of cross talk, especially switching noise, increase with the length of the interconnect and with decreasing rise time. Shorter rise-time signals will create more cross talk. Keeping interconnects short, such as by using CSPs and high-density interconnects (HDI), will help to minimize cross talk.

1.4 Rail-Collapse Noise

Noise is generated, and is a problem, in more than just the signal paths. It can also be a disaster in the power- and ground-distribution network which feeds

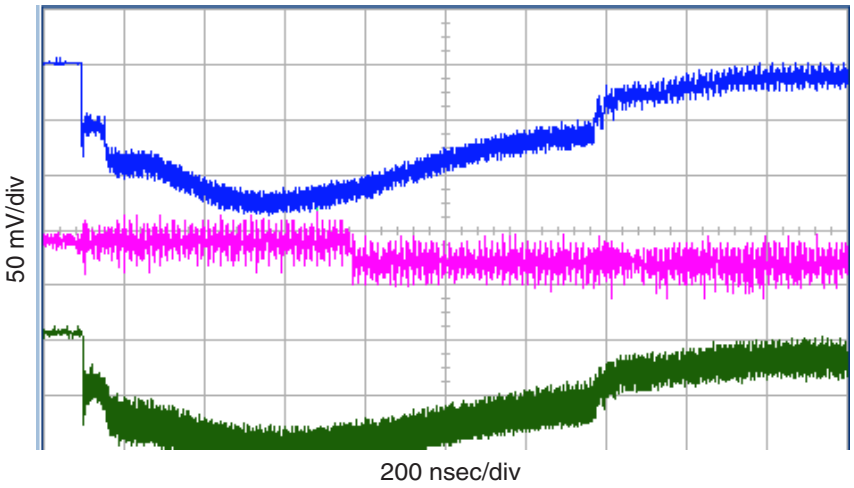

Figure 1-9 Measured Vcc voltage at three pin locations on the package of a microprocessor just coming from a "stop-clock" state. Nominal voltage is supposed to be 2.5 V, but due to voltage drops in the power distribution system, the delivered voltage collapses by almost 125 mV at some times.

each chip. When current through the power- and ground-path changes, as when a chip switches its outputs or core gates switch, there will be a voltage drop across the impedance of the power and ground paths. This voltage drop will mean less voltage gets to the chip, which will see a decrease, or collapse, of the voltage between the power and ground rails. One example of the change in the voltage across a microprocessor is shown in Figure 1-9.

In high-performance processors and some ASICs, the trend is for lower power-supply voltage, but higher power consumption. This is primarily due to more gates on a chip switching faster. In each cycle, a certain amount of energy is consumed. When the chip switches faster, the same energy is consumed, but more often, leading to higher-average power consumption.

These factors combine to mean higher currents will be switching in shorter amounts of time, and the amount of noise that can be tolerated will decrease. As the drive voltage decreases and the current level increases, any voltage drops associated with rail collapse become a bigger and bigger problem.

TIP The goal in designing the power and ground distribution is to try to minimize the impedance of the power-distribution system (PDS).

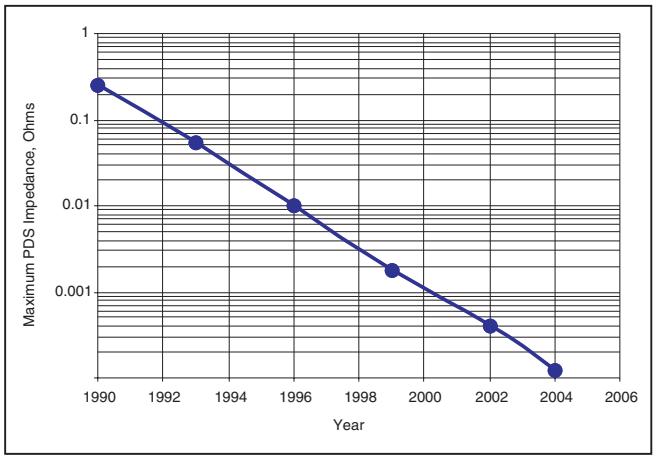

Figure 1-10 Trend in the maximum allowable impedance of the power distribution system for high-end processors. Source, Sun Microsystems.

In this way, even though there is current switching in the PDS, the voltage drop across a lower impedance may be kept to an acceptable level. The requirements for the impedance of the PDS have been evaluated for high-end processors by Sun Microsystems. Their estimate of the required impedance of the PDS is shown in Figure 1-10. Lower impedance in the PDS is increasingly important and harder to achieve.

If we understand how the physical design of the interconnects affects their impedance, we can optimize the design of the PDS for low impedance. As we shall see, designing a low-impedance PDS means including features such as the following:

1. Closely spaced adjacent planes for the power and ground distribution with as thin a dielectric between them as possible
2. Low-inductance decoupling capacitors
3. Multiple, very short power and ground leads in packages
4. On-chip decoupling capacitance

An innovative technology to help minimize rail collapse can be seen in the new ultrathin, high-dielectric constant laminates for use between power and ground layers. One example is C-Ply from 3M Corp. This material is 8 microns thick and has a dielectric constant of 20. Used as the power and ground layers in an other-

wise conventional board, the ultralow loop inductance and high distributed capacitance dramatically reduce the impedance of the power and ground distribution. One example of the rail-collapse noise on a small test board using conventional layers and a board with this new C-Ply is shown in Figure 1-11.

1.5 Electromagnetic Interference (EMI)

With board-level clock frequencies in the 100-MHz–500-MHz range, the first few harmonics are within the common communications bands of TV, FM radio, cell phone, and personal communications services (PCS). This means there is the very real possibility of electronic products interfering with communications unless their electromagnetic emissions are kept below acceptable levels. Unfortunately, without special design considerations, EMI will get worse at higher frequencies. The radiated far-field strength from common-mode currents will increase linearly with frequency and from differential-mode currents will increase with the square of the frequency. With clock frequencies increasing, the radiated emissions level will inevitably increase as well.

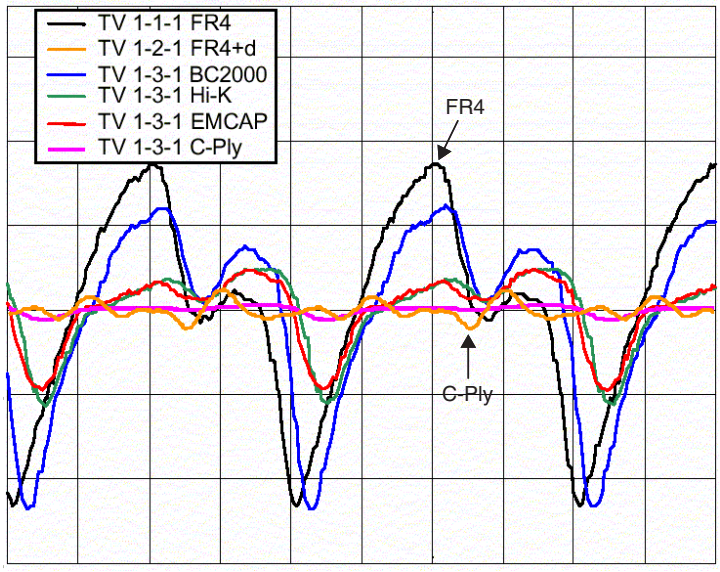

Figure 1-11 Measured rail voltage noise on a small digital board, with various methods of decoupling. The worst case is no decoupling capacitors on FR4. The best case is the 3M C-Ply material, showing virtually no voltage noise. Courtesy of National Center for Manufacturing Science. Scale is 0.5 v/div and 5 nsec/div.

It takes three things to have an EMI problem: a source of noise, a pathway to a radiator, and an antenna. Every source of signal-integrity problem mentioned above will be a source of EMI. What makes EMI so challenging is that even if the noise is low enough to meet the signal-integrity noise budget, it may still be large enough to cause serious radiated emissions.

TIP The two most common sources of EMI are (1) the conversion of some differential signal into a common signal, which eventually gets out on an external-twisted pair cable, and (2) ground bounce on a circuit board generating common currents on external single-ended shielded cables. Additional noise can come from internally generated radiation leaking out of the enclosure.

Most of the voltage sources that drive radiated emissions come from the power- and ground-distribution networks. Often, the same physical designs that contribute to low rail-collapse noise will also contribute to lower emissions.

Even with a voltage-noise source that can drive radiation, it is possible to isolate it by grouping the high-speed sections of a board away from where they might exit the product. Shielding the box will minimize the leakage of the noise to an antenna. Many of the ills of a poorly designed board can be fixed with a good shield.

A product with a great shield will still need to have cables connecting it to the outside world for communications, for peripherals, or for interfacing. Typically, the cables extend outside the shielded enclosure, act as the antennas, and radiate. The correct use of ferrites on all connected cables, especially a twisted pair, will dramatically decrease the efficiency of the cables as antennas. A close-up of the ferrite around a cable is shown in Figure 1-12.

The impedance associated with the I/O connectors, especially the impedance of the return-path connections, will dramatically affect the noise voltages that can drive radiating currents. The use of shielded cables with low-impedance connections will go a long way to minimize EMI problems.

Unfortunately, for the same physical system, increasing the clock frequency generally will also increase the radiated emissions level. This means that EMI problems will be harder to solve as clock frequencies increase.

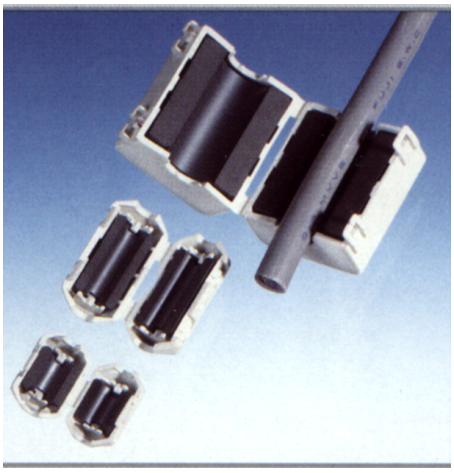

Figure 1-12 Ferrite choke around a cable, split apart. Ferrites are commonly used around cables to decrease common mode currents, a major source of radiated emissions. Courtesy of IM Intermark.

1.6 Two Important Signal Integrity Generalizations

Two important generalizations should be clear from looking at the four signal-integrity problems above.

First, each of the four families of problems gets worse as rise times decrease. All the signal-integrity problems above scale with how fast the current changes or with how fast the voltage changes. This is often referred to as dI/dt or dV/dt. Shorter rise times mean higher dI/dt and dV/dt.

It is unavoidable that as rise times decrease, the noise problems will increase and be more difficult to solve. And, as a consequence of the general trends in the industry, the rise times found in *all* electronic products will continually decrease. This means that what might not have caused a problem in one design may be a killer problem in the next design with the next generation chip sets operating with a shorter rise time. This is why it is often said, "There are two kinds of engineers, those that have signal integrity problems and those that will."

The second important generalization is that effective solutions to signal-integrity problems are based heavily on understanding the impedance of interconnects. If we have a clear intuitive sense of impedance and can relate the physical design of the interconnects with their impedance, many signal-integrity problems can be eliminated during the design process.

This is why an entire chapter in this book is devoted to understanding impedance from an intuitive and engineering perspective and why much of the rest of this book is about how the physical design of interconnects affects the impedance seen by signals and the PDS.

1.7 Trends in Electronic Products

A few trips to the local computer store over the last 10 years will have given anyone a good sense of the incredible treadmill of progress in computer performance. One measure of performance is the clock frequency of a processor chip. The trend for Intel processor chips, as illustrated in Figure 1-13, shows a doubling in clock frequency about every two years.

This trend toward ever-higher clock frequency is driven by the same force that drives the semiconductor revolution—photolithography. As the gate channel length of transistors is able to be manufactured at smaller size, the switching speed of the transistor increases. The electrons and holes have a shorter distance to travel and can transit the gate, effecting transitions, in a shorter time when the channel length is shorter.

When we refer to a technology generation as 0.18 microns or 0.13 microns, we are really referring to the smallest channel length that can be manufactured. The shorter switching time for smaller channel-length transistors has two important consequences for signal integrity.

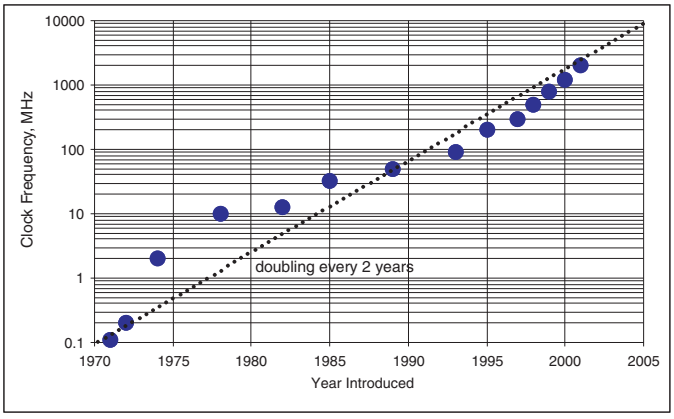

Figure 1-13 Historical trend in the clock frequency of Intel processors based on year of introduction. The trend is a doubling in clock frequency every two years. Source is Intel Corp.

The minimum time required for one clock cycle is limited by all the operations that need to be performed in one cycle. Usually, there are three main factors that contribute to this minimum time: the intrinsic time for all the gates that need to switch in series, the time for the signals to propagate through the system to all the gates that need to switch, and the settling time needed for the signals at the inputs to be read by the gates.

In single-chip microprocessor-based systems, such as in personal computers, the dominant factor influencing the minimum cycle time is the switching speed of the transistors. If the switching time can be reduced, the minimum total time required for one cycle can be reduced. This is the primary reason clock frequencies have increased as feature size has been reduced.

TIP It is inevitable that as transistor feature size continues to be reduced, rise times will continue to decrease and clock frequencies will continue to increase.

The projections from the 2001 Semiconductor Industry Association, (SIA) International Technology Roadmap for Semiconductors (ITRS) for future on-chip clock frequencies, based on projected feature size reductions, compared with the Intel processor trend are shown in Figure 1-14. This shows the projected trend for clock frequency increasing at a slightly diminishing but still growing rate for the next 15 years as well.

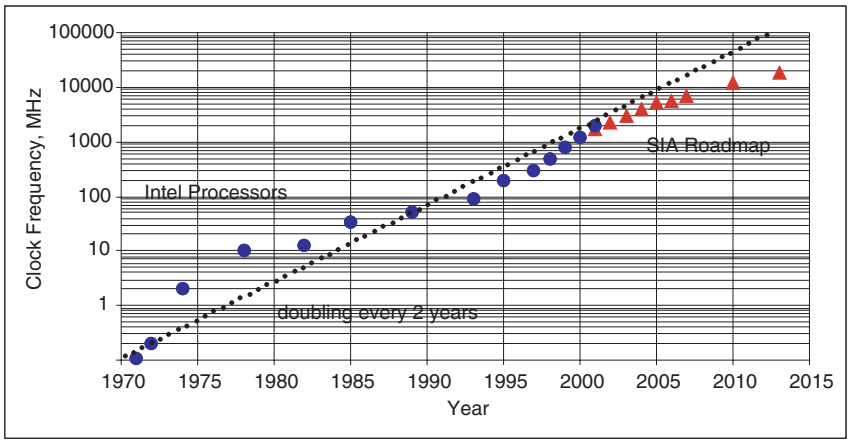

Figure 1-14 Historical trend in the clock frequency of Intel processors based on year of introduction. The trend is a doubling in clock frequency every two years. Also included is the Semiconductor Industry Association roadmap expectations. Source is Intel Corp. and SIA.

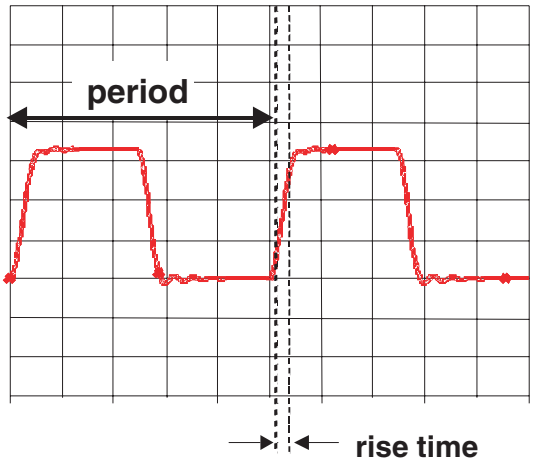

Figure 1-15 The 10–90 rise time for a typical clock waveform is roughly 10% of the period. Scale is 1 v/div and 2 nsec/div, simulated with Mentor Graphics Hyperlynx.

As the clock frequency increases, the rise time of the signals must also decrease. Each gate that reads either the data lines or the clock lines needs enough time with the signal being either in the high state or the low state, to read it correctly.

This means only a short time is left for the signal to be in transition. We usually measure the transition time, either the rise time or the fall time, as the time it takes to go from 10% of the final state to 90% the final state. This is called the 10–90 rise time. Some definitions use the 20% to 80% transition points and this rise time is referred to as the 20–80 rise time. An example of a typical clock waveform and the time allocated for the transition are shown in Figure 1-15. In most high-speed digital systems, the time allocated to the rise time is about 10% of the clock cycle time, or the clock period. Based on this generalization, the rise time is roughly related to the clock frequency by:

$$RT = \frac{1}{10 \times F_{clock}} \tag{1-1}$$

where:

RT= the rise time, in nsec

F_{clock} = the clock frequency, in GHz

For example, when the clock frequency is 1 GHz, the rise time of the associated signals is about 0.1 nsec, or 100 psec. When the clock frequency is 100 MHz, the rise time is roughly 1 nsec. This relationship is shown in Figure 1-16.

TIP The treadmill-like advance of ever-increasing clock frequency means an ever-decreasing rise time and signal-integrity problems that are harder to solve.

Even if the clock frequency of a product is low, there is still the danger of shorter rise times as a direct consequence of the chip technology. A chip-fabrication factory, usually called a "fab," will try to standardize all their wafers on one process, to increase the overall yield. The smaller the chip size, the more chips can be fit on a wafer and the lower the cost per chip. Even if the chip is going to be used in a slow-speed product, it may be fabricated on the same line as a leading-edge ASIC, with the same small-feature size.

Ironically, the lowest cost chips will always have ever-shorter rise times, even if they don't need it for the specific application. This has an unintended, scary consequence. If you have designed a chip set into your product and the rise time is 2 nsec, for example, with a 50-MHz clock, there may be no signal-integrity

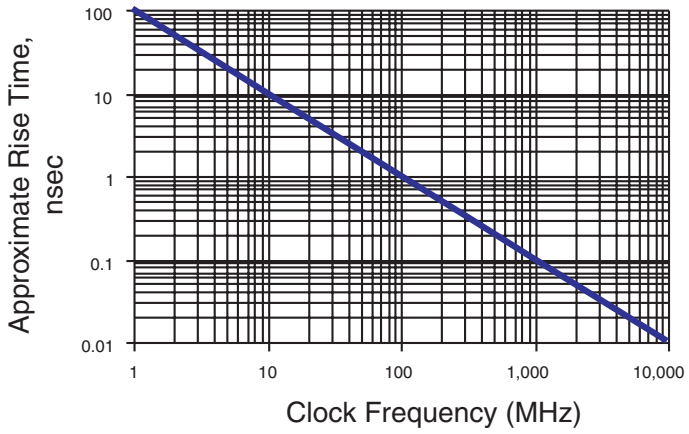

Figure 1-16 Rise time decreases as the clock frequency increases. Signal-integrity problems usually arise at rise times less than 1 nsec or at clock frequencies greater than 100 MHz.

problems. When your chip supplier upgrades their fab line with a finer-feature process, they may provide you with lower-cost chips. You may think you are getting a good deal. However, these lower cost chips may now have a rise time of 1 nsec. This shorter rise time may cause reflection noise, excess cross talk, rail collapse, and they may fail the Federal Communications Commission (FCC) EMI certification tests. The clock frequency of your product hasn't changed, but, unknown to you, the rise time of the chips supplied to you has decreased with the newer, finer-feature manufacturing process.

TIP As all fabs migrate to lower-cost, finer-feature processes, all fabricated chips will have shorter rise times, and signal integrity problems have the potential to arise in all products, even those with clock frequencies below 50 MHz.

It's not just microprocessor-based digital products that are increasing in clock frequency and decreasing in rise times. The data rate and clock frequencies used in high-speed telecommunications products are blowing past the clock frequencies of microprocessor digital products.

One of the most common specifications for defining the speed of a high-speed serial link is the optical carrier, or OC, spec. This is actually a data rate, with OC-1 corresponding to about 50 Mbits/sec. OC-48, with a data rate of 2.5 Gbits/sec, is in high-volume deployment and widely implemented. OC-192, at 10 Gbits/sec, is just ramping up. In the near future, OC-768, at 40 Gbits/sec, will be in wide-scale deployment.

The OC designation is a specification for data rate, not for clock frequency. Depending on how a bit is encoded in the data stream, the actual clock frequency of a system can be higher or lower than the data rate.

For example, if there is one bit per clock cycle, then the actual clock frequency will be the data rate. If a data bit is encoded on both the leading and trailing edge of the clock, then a 2.4-Gbits/sec data rate can be obtained with a 1.25-GHz clock. The more bits in the data stream used for error correction and overhead, the lower the data rate, even though the clock frequency is constant. To keep signal-integrity problems at a minimum, the lowest clock frequency and longest edge rate should be used. This has inspired a growing trend to encode four to eight data bits per clock cycle, using multilevel signaling and multiple signal lines in parallel.

Given these caveats, as a general rule of thumb, the clock frequency of a high-speed serial link can be roughly approximated by the data rate. The trend is

definitely toward higher data rates. In a few years, clock frequencies of 40 GHz will be common in telecommunications systems. The entire clock period of a 40-GHz clock signal is only 25 psec. This means that the rise time must be significantly less than 10 psec. This is a short rise time and will require extremely careful design practices.

High-speed serial links are not just limited to telecommunications applications. All high-speed digital busses will soon migrate to clock frequencies near or above 1 GHz. Many high-speed serial link busses have been proposed for on-board data flow for general high-speed digital products. Among these are the following:

- RapidIO with a clock frequency about 1 GHz
- Hypertransport with a clock frequency about 800 MHz
- 3GIO with a clock frequency about 1.25 GHz
- Serial ATA with a clock frequency up to 3 GHz
- SCID with a clock frequency about 1.6 GHz
- Infiniband with a clock frequency about 1.25 GHz
- FibreChannel 4025 with a clock frequency about 2.125
- Gigabit Ethernet with a clock frequency about 625 MHz

1.8 The Need for a New Design Methodology

We've painted a scary picture of the future. The situation analysis is as follows:

- Signal integrity problems can prevent the correct operation of a high-speed digital product.
- These problems are a direct consequence of shorter rise time and higher clock frequencies.
- Rise times will continue their inevitable march toward shorter values and clock frequencies will continue to increase.
- Even if we limit the clock frequency, using the lowest cost chips means even low-speed systems will have chips with very short rise times.
- We have less time in the product-design cycle to get the product to market—it must work the first time.

What are we to do? How are we to efficiently design high-speed products in this new era? In the good old days of 10-MHz clock systems, when the interconnects were transparent, we did not have to worry about signal-integrity effects. We could get away with designing the product for functionality and ignore signal integrity. In today's products, ignoring signal integrity invites schedule slips, higher development costs, and the possibility of never being able to build a functional product.

It will always be more profitable to pay extra to design a product right the first time than to try to fix it later. In the product life cycle, often, the first six months in the market are the most profitable. If your product is late, a significant share of the life-cycle profits may be lost. Time really is money.

TIP A new product-design methodology is required that ensures signal-integrity problems are identified and eliminated from the product as early in the design cycle as possible. To meet ever-shorter design cycle times the product must meet performance specifications the first time.

1.9 A New Product Design Methodology

There are five key ingredients to this new methodology:

1. Understand the origin of signal-integrity problems and the general guidelines to minimize these problems.
2. Translate the general guidelines into specific design rules for each specific custom product.
3. Predict performance early in the design cycle by creating electrical circuit models for each component, critical net, and the entire system and by performing local and system-level simulation.
4. Optimize the performance of the design for cost, schedule, and risk by modeling and simulating at every step of the design cycle, especially at the beginning.
5. Use characterization measurements throughout the design cycle to reduce the risk and increase confidence to the quality of the predictions.

In addition to "understanding," the other key processes to this new design methodology are modeling, simulation, and characterization.

Simulation is about predicting the performance of the system before building the hardware. It used to be, only those nets in a system that were sensitive to signal-integrity effects were simulated. These nets are called "critical nets." Typically, they were clock lines and maybe a few high-speed buss lines. In 100-MHz clock frequency products, maybe only 5%–10% of the nets were critical nets. In products with clock frequencies at 200 MHz and higher, more than 50% of the nets may be critical, and the entire system needs to be included in the simulation.

TIP In all high-speed products today, system-level simulations must be performed to accurately predict signal-integrity effects.

In order to predict the electrical performance, which is typically the actual voltage and current waveforms at various nodes, we need to translate the physical design into an electrical description. This can be accomplished in one of two paths. The physical design can be converted into an equivalent circuit model and then a circuit simulator can be used to predict the voltages and currents at any node.

Alternatively, an electromagnetic simulator can be used to simulate the electric and magnetic fields everywhere in space, based on the physical design. From the electric and magnetic fields, a behavioral model of the interconnects can be generated which can then be used in a circuit simulator, or the electric and magnetic fields can be converted into voltage and currents to show performance.

1.10 Simulations

There are three types of electrical simulation tools that predict the analog effects of the interconnects on signal behavior:

1. Electromagnetic (EM) simulators, which solve Maxwell's Equations and simulate the electric and magnetic fields at various locations in the time or frequency domains.
2. Circuit simulators, which solve the differential equations corresponding to various circuit elements and include Kirchhoff's current and voltage relationships to predict the voltages and currents at various circuit nodes, in the time or frequency domains.
3. Behavioral simulators, which use models based on tables and transmission lines and other passive-element models based on transfer functions which

quickly predict the voltages and currents at various nodes, typically in the time domain.

Blame signal integrity on Maxwell's Equations. These four equations describe how conductors and dielectrics interact with electric and magnetic fields. After all, signals are nothing more than propagating electric and magnetic fields. When the electric and magnetic fields themselves are simulated, the interconnects and all passive components must be translated into conductors and dielectrics, with their associated geometries and material properties.

A device driver is converted into an incident electromagnetic wave, and Maxwell's Equations are used to predict how this wave interacts with the conductors and dielectrics. The material geometries and properties define the boundary conditions in which Maxwell's Equations are solved.

Though Figure 1-17 actually shows Maxwell's Equations, it is never necessary for any practicing engineer to solve them by hand. They are shown here only for reference and to demonstrate that there really is a set of just a few simple equations that describe absolutely everything there is to know about electromagnetic fields. How the incident electromagnetic field interacts with the geometry and materials, as calculated from these equations, can be displayed at every point in space. These fields can be simulated either in the time domain or the frequency domain.

Time Domain	*Frequency Domain*
$\nabla \bullet \varepsilon E = \frac{\rho}{\varepsilon_0}$	$\nabla \bullet \varepsilon E = \frac{\rho}{\varepsilon_0}$
$\nabla \bullet B = 0$	$\nabla \bullet \mu H = 0$
$\nabla x\, E + \frac{\partial B}{\partial t} = 0$	$\nabla x\, E + j\omega\mu H = 0$
$\nabla x\, B - \frac{\mu\varepsilon}{c^2}\frac{\partial E}{\partial t} = \mu_0$	$\nabla x\, H - j\omega\varepsilon E = J$

Figure 1-17 Maxwell's Equations in the time and frequency domains. These equations describe how the electric and magnetic fields interact with materials through time and space. They are provided here just for reference.

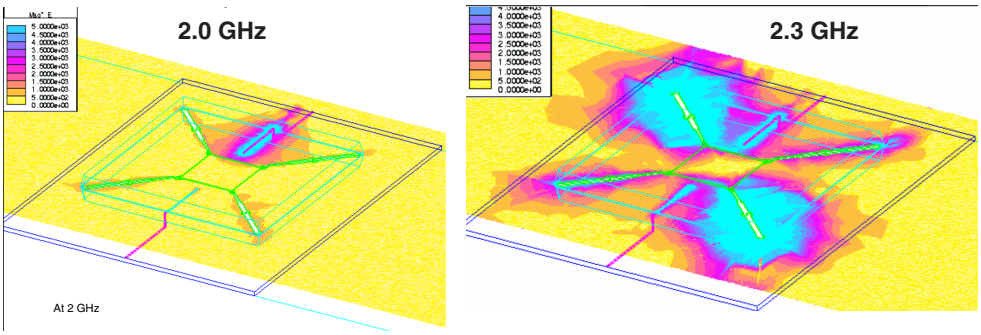

Figure 1-18 Example of an electromagnetic simulation of the electric fields in a 208-pin PQFP excited at 2.0 GHz (left) and at 2.3 GHz (right) showing a resonance. Simulation done with Ansoft's High Frequency Structure Simulator (HFSS).

An example of the simulated electric-field intensity inside a 208-pin plastic quad flat pack (PQFP) for an incident voltage sine wave on one pin at 2.0 GHz and 2.3 GHz is shown in Figure 1-18. The different shadings show the field intensity. This simulation illustrates that if a signal has frequency components at 2.3 GHz, it will cause large field distributions inside the package. These are called resonances and can be disastrous for a product. These resonances will cause signal-quality degradation, enhanced cross talk, and enhanced EMI. Resonances will always limit the highest bandwidth for which the part can be used.

Some effects can only be simulated with an electromagnetic simulator. Usually, an EM simulator is needed when interconnects are very non-uniform and electrically long (such as traces over gaps in the return path), when electromagnetic-coupling effects dominate (such as resonances in packages and connectors), or when it's necessary to simulate EMI effects.

Though all the physics can be taken into account with Maxwell's Equations, with the best of today's hardware and software tools, it is not practical to simulate the electromagnetic effects of other than the simplest structures. An additional limitation is that many of the current tools require a skilled user with experience in electromagnetic theory.

An alternative simulation tool, which is easier and quicker to use, is a circuit simulator. This simulation tool represents signals as voltages and currents. The various conductors and dielectrics are translated into their fundamental electrical circuit elements of resistances, capacitances, inductances, and their coupling.

Circuit theory is no less correct than the electromagnetic approach. It's just that some problems in signal integrity are more easily understood and their solu-

tions are more easily identified by using the circuit description over the EM theory description. There are some limitations to what can be simulated by a circuit simulator. Electromagnetic effects such as EMI, resonances, and nonuniform wave propagation cannot be taken into account by a circuit simulator. However, effects such as near-field cross talk, transmission line propagation, and switching noise can be accurately accounted for. An example of a circuit and the resulting simulated waveforms are shown in Figure 1-19.

A circuit diagram that contains combinations of fundamental circuit elements is called a schematic. If you can draw the schematic, a circuit simulator will be able to calculate the voltages and currents at every node.

The most popular circuit simulator is generically called SPICE (short for simulation program with integrated circuit emphasis). The first version was created at UC Berkeley in the early 1970s as a tool to predict the performance of transistors based on their geometry and material properties. It is fundamentally a circuit simulator. If you input a schematic, in a specialized format, the tool will solve the differential equations each circuit element represents, then calculate the voltages and currents either in the time domain, called a transient simulation, or in the frequency domain, called an AC simulation. There are over 30 commercially available versions of SPICE, with some free student/demo versions available for download from the Web.

Behavioral simulators use tables and specialized transfer functions to simulate voltages and currents. Their chief advantage over circuit simulators is in their computation speed. Many of the behavioral simulators use proprietary simulation engines and are optimized for particular types of circuits, such as lossless, coupled transmission lines.

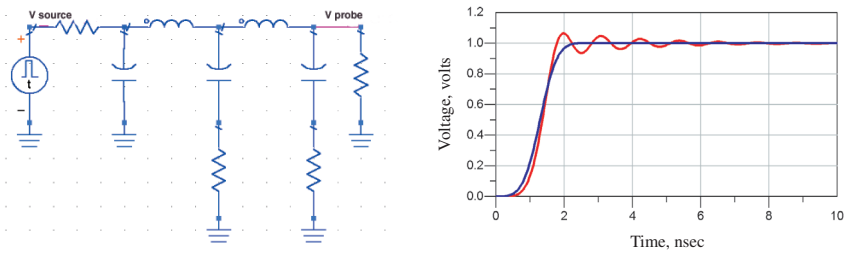

Figure 1-19 Example of a circuit model for the probe tip of a typical scope probe, approximately 5 cm long, and the resulting circuit simulation response of a clean, 1-nsec rise-time signal. The ringing is due to the excessive inductance of the probe tip.

1.11 Modeling and Models

Modeling refers to creating an electrical representation of a device or component that a simulator can interpret and use to predict voltage and current waveforms. The models for active devices, such as transistors and output drivers, are radically different from the models of passive devices, such as all interconnects and discrete components. For active devices, a model is typically either a SPICE-compatible model or an input/output buffer interface spec (IBIS) compatible model.

A SPICE model of an active device will use either combinations of ideal sources and passive elements or specialized transistor models based on the geometry of the transistors. This allows easy scaling of the transistor's behavior if the process technology changes. A SPICE model contains information about the specific features and process technology of a driver. For this reason, most vendors are reluctant to give out SPICE models of their chips since they contain such valuable information.

IBIS is a format that defines the response of input or output drivers in terms of their V-I and V-t characteristics. A behavioral simulator will take the V-I and V-t curves of the active devices and simulate how these curves might change as they are affected by the transmission lines and lumped resistor (R), inductor (L), and capacitor (C) elements which represent the interconnects. The primary advantage of an IBIS model for an active device is that an IC vendor can provide an IBIS model of its device drivers without revealing any proprietary information about the geometry of the transistors.

It is much easier to obtain an IBIS model than a SPICE model from an IC supplier. For system-level simulation, where 1,000 nets and 100 ICs may be simulated at the same time, IBIS models and behavioral simulators are typically used because they are more available and typically run faster than SPICE simulators.

The most important limitation to the accuracy of any simulation, SPICE or behavioral based, is the quality of the models. While it is possible to get an IBIS model of a driver that compares precisely with a SPICE model and matches the actual measurement of the device perfectly, it is difficult to routinely get good models of every device.

TIP In general, as end users, we must continually insist that our vendors supply us with some kind of verification of the quality of the models they provide for their components.

Another problem with device models is that a model that applies to one generation of chips will not match the next generation. With the next die shrink, which happens every six to nine months, the channel lengths are shorter, the rise times are shorter, and the V-I curves and transient response of the drivers change. An old model of a new part will give low estimates of the signal-integrity effects. As users, we must always insist the vendor provide current, accurate, and verified models of all drivers they supply.

TIP Though the intrinsic accuracy of all SPICE or behavioral simulators is typically very good, the quality of the simulations are only as good as the quality of the models that are simulated. The expression "garbage in, garbage out (GIGO)" was invented to describe circuit simulations.

For this reason, it is critically important to verify the accuracy of the models of all device drivers, interconnects, and passive components. Only in this way can we have confidence in the results from the simulations. Though the models for the active devices are vitally important, this book deals with models for the passive devices and interconnects. Other references listed in the bibliography discuss active-device models.

Obtaining models for all the elements that make up the system is critically important. The only way to simulate signal-integrity effects is to include models of the interconnects and passive components, such as the board-level transmission lines, package models, connector models, decoupling capacitors, and terminating resistors.

Of course, the circuit model can only use elements that the simulator understands. For most behavioral simulators, this means interconnects are described by resistors, capacitors, inductors, and transmission lines. For SPICE simulators, interconnects and passive components can be described by resistors, capacitors, inductors, mutual inductors, and transmission lines. In some SPICE and behavioral simulators, new ideal circuit elements have been introduced that include ideal coupled transmission lines and ideal lossy transmission lines as basic ideal circuit elements.

An example of a physical component, two surface-mount terminating resistors, and their equivalent electrical circuit model is shown in Figure 1-20. This model includes their inductive coupling, which gives rise to switching noise. Every electrical quality of their behavior can be described by their circuit model. This circuit model, or schematic, can accurately predict any measurable effect.

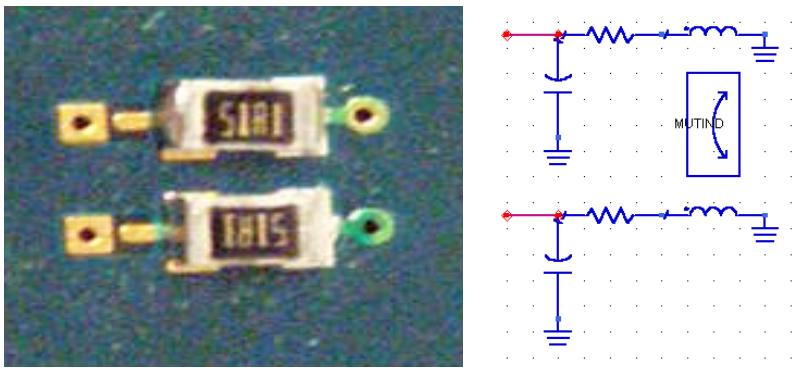

Figure 1-20 Two surface-mount 0805 resistors and their equivalent circuit model. This model has been verified up to 5 GHz.

There are two basic approaches to creating accurate circuit models of interconnects: by calculation and by measurements. We usually refer to creating a model from calculation as analysis and creating a model from a measurement as characterization.

1.12 Creating Circuit Models from Calculation

So much of life is a constant balancing act between the value received and the cost (in time and money). The analysis of all signal integrity design problems, as well as in most other fields, is no exception. We are constantly balancing the quality of an answer, as measured by its accuracy, for example, with how long it will take and how much it will cost to get the answer.

T I P The goal in virtually all product-development programs in today's globally competitive market, is to get to an acceptable design that meets the performance spec, while staying within the time, cost, and risk budget.

This is a tough challenge. The engineer involved in signal integrity and interconnect design can benefit from skill and versatility in selecting the best technology and establishing the optimum design rules as early in the design cycle as possible.

The most important tool in any engineer's toolbox is the flexibility to evaluate trade-offs quickly. These are really trade-offs between how choices of geometry, material properties, and design rules will affect system performance.

TIP The earlier in the design cycle the right trade-offs can be established, the shorter the development time and the lower the development costs.

To aid in the trade-off analysis, there are three levels of approximation used to predict performance or electrical properties. These three approaches are as follows:

1. Rules of thumb
2. Analytical approximations
3. Numerical simulations

Each approach represents a different balance of closeness to reality (i.e., accuracy, time, and effort required to get the answer). This is illustrated in Figure 1-21. Of course, these approaches are not substitutes for actual measurements. However, the correct application of the right analysis technique can sometimes shorten the design-cycle time to 10% of its original value, compared to relying on the build it/test it/redesign it approach.

Rules of thumb are simple relationships that are easy to remember and help feed your intuition. An example of a rule of thumb is, the self-inductance of a short wire is about 25 nH/inch. A wire bond 0.1 inches long, therefore, would have a self-inductance of about 2.5 nH.

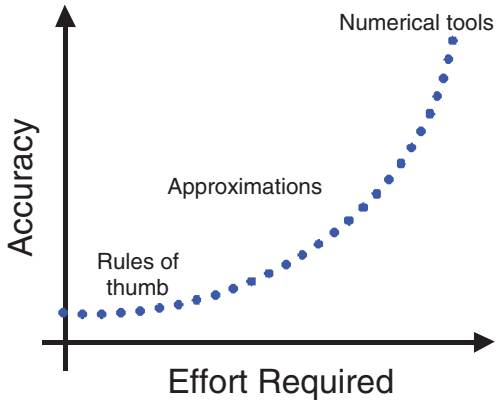

Figure 1-21 The balance between accuracy returned and effort required for the three levels of analysis. Each tool has its appropriate time and place.

A rule of thumb should be the first step in any analysis problem, if only to provide a sanity-check number that every answer would be compared to. When you are brainstorming design/technology/cost trade-offs or looking for rough estimates, comparisons, or plausibility arguments, using rules of thumb can accelerate your progress better than tenfold. After all, the closer to the beginning of the product-development cycle the right design and technology decisions can be made, the more time and money will be saved in the project.

Rules of thumb are not meant to be accurate, they are meant to give an answer quickly. They should never be used when signing off on a design. They should be used to calibrate your intuition and provide guidance to help make high-level trade-offs. Appendix A contains a summary of many of the important rules of thumb used in signal integrity.

Analytical approximations are equations or formulas. For example, an approximation for the loop self-inductance of a circular loop of wire is:

$$L_{self} = 32 \times R \times \ln\left(\frac{4R}{D}\right) nH \tag{1-2}$$

where:

L_{self} = the self-inductance, in nH

R = the radius of the loop, in inches

D = the diameter of the wire, in inches

For example, a round loop, 1/2 inch in radius or 1 inch in diameter, made from 10-mil-thick wire, would have a loop inductance of about 85 nH. Put your index finger and thumb together in a circle. If they were made of 20-gauge copper wire, the loop inductance would be about 85 nH.

Approximations have value in that they can be implemented in a spread sheet and they can answer what-if questions quickly. They identify the important first-order terms and how they are related. The approximation above illustrates that the inductance scales a little faster than the radius. The larger the radius of the loop, the larger the inductance of the loop. Also, the larger the wire thickness, the

smaller the loop inductance, but only slightly, as the loop inductance varies inversely with the natural log of the thickness, a weak function.

TIP It is important to note that with very few exceptions, every equation you see being used in signal-integrity analysis is either a definition or an approximation.

A definition establishes an exact relationship between two or more terms. For example, the relationship between clock frequency and clock period, F = 1/T, is a definition. The relationship between voltage, current, and impedance, Z = V/I, is a definition. However, just because a formula allows evaluation on a calculator to five decimal places doesn't mean it is accurate to five decimal places.

You should always worry about the accuracy of an approximation, which may vary from 1% to 50% or more. You can't tell the accuracy of an approximation by looking at its complexity or its popularity.

How good are approximations? If you don't know the answer to this question for your specific case, you may not want to base a design sign-off—where a variance of more than 5% might mean the part won't work in the application—on an approximation with unknown quality. The first question you should always ask of every approximation is: how accurate is it?

One way of verifying the accuracy of an approximation is to build well-characterized test vehicles and perform measurements that can be compared to the calculation. Figure 1-22 shows the very good agreement between the approximation for loop inductance offered above and the measured values of loop inductance based on building loops and measuring them with an impedance analyzer. The agreement is seen to be better than 2%.

Approximations are extremely important when exploring design space or performing a tolerance analysis. They are wonderful when balancing trade-offs. However, when being off will cost a significant amount of time, money, or resources, you should never rely on approximations whose accuracy is uncertain.

There is a more accurate method for calculating the parameter values of the electrical circuit elements of interconnects from the geometry and material properties. It is based on the numerical calculation of Maxwell's Equations. These tools are called field solvers, in that they solve for the electric and magnetic fields based on Maxwell's Equations, using as boundary conditions the distribution of conductors and dielectrics. A type of field solver that converts the calculated fields into the actual parameter values of the equivalent circuit-model elements, such as the R, L, or C values, is called a parasitic extraction tool.

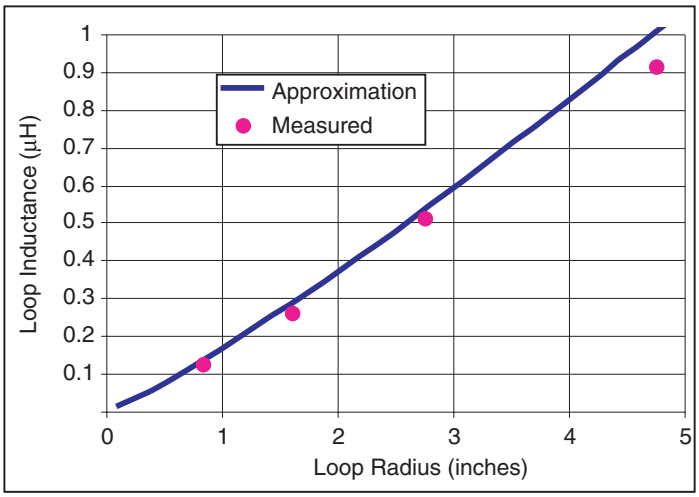

Figure 1-22 Comparison of the measured and calculated loop inductance of various circular wire loops as measured with an impedance analyzer. The accuracy of the approximation is seen to be about 2%.

When the geometry of the interconnect is uniform down its length, it can be described by its cross section and a 2D field solver can be used to extract its transmission line properties. An example of a typical cross section of a microstrip transmission line and the simulated electric field lines and equipotentials is shown in Figure 1-23. For this structure, the extracted parameter values were Z_0 = 50.3 Ohms and TD = 142 psec.

When the cross section is nonuniform, such as in a connector or IC package, a 3D field solver is needed for the most accurate results.

Before relying on any numerical-simulation tool, it is always important to have its results verified for some test cases similar to the end application for which you will be using it. Every user should insist on vendor verification that the tool is suitably accurate for your typical application. In this way, you gain confidence in the quality of your results. The accuracy of some field solvers has been verified to better than 1%. Obviously, not all field solvers are always this accurate.

When accuracy is important, as, for example, in any design sign-off, a numerical-simulation tool, such as a parasitic extraction tool, should be used. It may take longer to construct a model using a numerical-simulation tool than it would using a rule of thumb or even using an analytic approximation. More effort in time and in expertise is required. But what they offer is greater accuracy and a

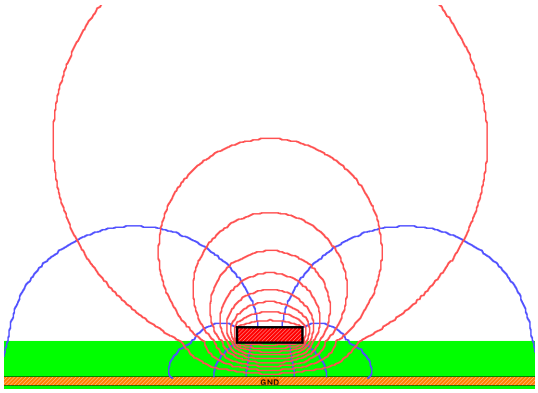

Figure 1-23 The results of a 2D field solver used to calculate the electric fields in a microstrip transmission line. The parasitic extraction tool used was Hyperlynx from Mentor Graphics. The accuracy of this tool has been independently verified to be better than 2%.

higher confidence the as manufactured part will match the predicted performance. As new commercially available numerical-simulation tools advance, market pressures will always drive them to be easier to use.

Using a combination of these three analysis techniques, the trade-offs between the costs in time, money, and risk can be balanced with a very good prediction of the possible performance gain.

1.13 Three Types of Measurements

T I P Though calculations play the critical role of offering a prediction of performance before the product is built, measurements play the critical role of risk reduction. The ultimate test of any calculation result is a measurement.

When it comes to measuring passive interconnects, as distinct from active devices, the measuring instrument must create a precision reference signal, apply it to the device under test, and measure the response. Ultimately, this response is related to the impedance of the device. In active devices, which create their own signals, the measurement instrument can be passive, merely measuring the created voltages or currents. There are three primary instruments used to perform measurements on passive components:

1. Impedance analyzer

2. Vector-network analyzer (VNA)

3. Time-domain reflectometer (TDR)

An impedance analyzer, typically a four-terminal instrument, operates in the frequency domain. One pair of terminals is used to generate a constant-current sine wave through the device under test (DUT). The second pair of terminals is used to measure the sine-wave voltage across the DUT.

The ratio of the measured voltage to the measured current is the impedance. The frequency is typically stepped up from the low 100-Hz range to about 40 MHz. The magnitude and phase of the impedance at each frequency point is measured based on the definition of impedance.

The vector-network analyzer also operates in the frequency domain. Each terminal or port emits a sine-wave voltage at frequencies that can range in the low kHz to over 50 GHz. At each frequency, the incident-voltage amplitude and phase, as well as the reflected amplitude and phase, are measured.

The reflected signal will depend on the incident signal and the impedance change in going from the VNA to the DUT. The output impedance of a VNA is typically 50 Ohms. By measuring the reflected signal, the impedance of the device under test can be determined at each frequency point. The relationship between the reflected signal and the impedance of the DUT is:

$$\frac{V_{reflected}}{V_{incident}} = \frac{Z_{DUT} - 50\Omega}{Z_{DUT} + 50\Omega} \tag{1-3}$$

where:

$V_{reflected}$ = the amplitude and phase of the reflected sine-wave voltage

$V_{incident}$ = the amplitude and phase of the incident sine-wave voltage

Z_{DUT} = the impedance of the device under test

50 Ω = impedance of the VNA

The ratio of the reflected to the incident voltage, at each frequency, is often referred to as one of the scattering, or S, parameters, signified as S11. Measuring S11 and knowing the source impedance is 50 Ohms allow us to extract the impedance of the device under test at any frequency. An example of the measured impedance of a short-length transmission line is shown in Figure 1-24.

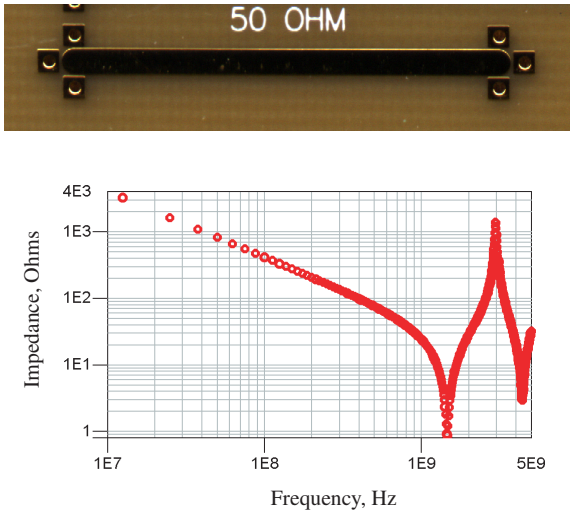

Figure 1-24 Measured impedance of a one-inch length of transmission line. The network analyzer measured the reflected sine-wave signal between the front of the line and a via to the plane below the trace. This reflected signal was converted into the magnitude of the impedance. The phase of the impedance was measured, but is not displayed here. The frequency range is from 12 MHz to 5 GHz. Measured with a GigaTest Labs Probe Station.

The time-domain reflectometer (TDR) is similar to the VNA but operates in the time domain. It emits a fast rise-time step signal, typically 35 psec to 150 psec, and measures the reflected transient amplitude. Again, using the reflected voltage, the impedance of the DUT can be extracted. In the time domain, the impedance measured represents the instantaneous impedance of the DUT. For an interconnect that is electrically long, such as a transmission line, the TDR can map the impedance profile. Figure 1-25 is an example of a TDR profile from a four-inch-long transmission line with a small gap in the return plane showing a higher impedance where the gap occurs.

TIP Though the same impedance of a DUT can be displayed in the frequency domain or the time domain, it is really a different impedance in each case. When displayed in the frequency domain, the impedance is the total, integrated impedance of the entire DUT at each frequency. When displayed in the time domain, it is the instantaneous impedance at each spatially distinct point on the DUT that is displayed.

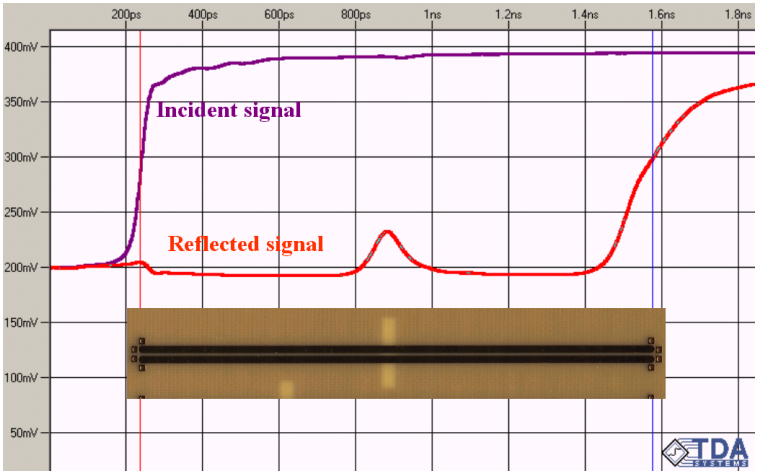

Figure 1-25 Measured TDR profile of a four-inch-long uniform transmission line with gap in the return path near the middle. The far end of the line is open. Measured with an Agilent 86100 DCA TDR and recorded with TDA Systems IConnect software, using a GigaTest Labs Probe Station.

1.14 The Role of Measurements

If it is possible to calculate the expected electrical performance of a component or system, why bother with a measurement? Why not just rely on the modeling and simulation tool? Measurements can only be performed on real devices. Isn't the goal to avoid the build-it/measure-it/redesign-it iteration loop that takes so much time?

Measurements such as the ones illustrated above play four critically important basic roles, at various stages in the product life cycle, all related to risk reduction and establishing higher confidence in the accuracy of the simulations. Measurements allow designers to do the following:

1. Verify the accuracy of the design/modeling/simulation process before significant resources are expended using an unverified process.
2. Create a model for a component at any stage of the design cycle as parts are made available or acquired from a vendor.
3. Emulate system performance of a component as a quick way of determining expected performance without building a model, at any stage of the design cycle as parts are made available or acquired from a vendor.
4. Debug a functional part or system, at any stage of the design cycle as parts are made available or acquired from a vendor.

An example from Delphi Electronics illustrates the incredible power of the combination of a design/modeling/simulation process that has been verified using measurements. One of the products Delphi Electronics makes is a custom flexible connector connecting two circuit boards for high-speed signals. An example is shown in Figure 1-26. They are used in servers, computers, and switching systems. The electrical performance of this connector is critical to the correct function of the system.

A customer comes to them with a set of performance specifications and Delphi is to deliver a part that meets these specs. The old way of designing the product was to make a best guess, manufacture the part, bring it back to the lab and perform measurements on it, compare to the specs, and redesign it. This is the old, build-it/test-it/redesign-it approach. In this old process, one iteration took almost nine weeks because of the long CAD and manufacturing cycle. Sometimes, the first design did not meet the customer specs and a second cycle was required. This meant an 18-week development cycle.

To shorten this design-cycle time, Delphi implemented a 2D modeling tool that allowed predicting the electrical properties of the connector based on the geometry and material properties. Through a few cycles of experimenting, using measurements with a TDR and VNA as the verification process, Delphi was able to fine-tune its modeling process to ensure accurate predictability of its product, before the part was sent to manufacturing. An example of the final agreement between what its modeling tool is able to predict and what was actually measured for a connector, compared with the customer's original specification, is shown in Figure 1-27.

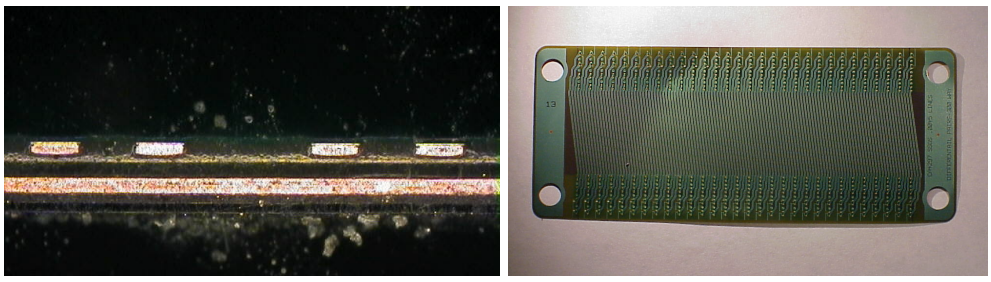

Figure 1-26 Gold Dot Connector from Delphi Automotive. Left: cross section of the two-metal-layer flex substrate. Right: top view of the hundreds of conductors connecting one board to the other, each with controlled electrical performance. Photo courtesy of Laurie Taira-Griffin, Delphi Electronics.

Parameter	Simulation	Measured	Goal
Single-Ended Impedance	52.1 Ohms	53 Ohms	50 +/-10% Ohms
Differential Impedance	95.2 Ohms	98 Ohms	100 +/- 10% Ohms
Attenuation (5GHz)	<.44 dB/inch	<.44 dB/inch	<.5 dB/inch
Propagation Delay	152 ps/inch	158 ps/inch	170 ps/inch
Single-Ended NEXT	<4.5%	<4.5%	<5%
Differential NEXT	<.3%	<.3%	<.5%
Data Rate	>5 Gbps	>5 Gbps	5 Gbps

Figure 1-27 Summary of the predicted and measured electrical specifications of the connector compared with the requirements for a particular connector. After the modeling/simulation process was optimized, the ability to predict performance was excellent.

Once the modeling/simulation process was in place and Delphi had confidence in its ability to accurately predict the final performance of the manufactured connectors, Delphi was able to reduce the design-cycle time to less than four hours. Nine weeks to four hours is a reduction of more than 100x. Measurements provided the critical verification step for this process.

1.15 The Bottom Line

1. Signal integrity problems relate to how the physical interconnects screw up pristine signals coming from the integrated circuits.
2. There are four general families of signal-integrity problems: signal quality on one net, cross talk between adjacent nets, rail collapse, and EMI.
3. Each of these problems gets worse and harder to solve as rise times decrease or clock frequencies increase.
4. It is inevitable that as transistors get smaller, their rise times will get shorter and signal integrity will be a greater problem.

5. To find, fix, and prevent signal-integrity problems, it is essential to be able to convert the physical design into its equivalent electrical circuit model and to use this model to simulate waveforms so as to predict performance before the product is built.
6. Three levels of analysis can be used for calculating electrical effects—rules of thumb, approximations, and numerical tools. These can be applied to modeling and simulation.
7. Three general instruments can be used to measure electrical properties of passives and interconnects: an impedance analyzer, a network analyzer, and a time-domain reflectometer.
8. These instruments play the important role of reducing the risk and increasing the confidence level in the accuracy of the modeling and simulation process.
9. Understanding the four signal-integrity problems leads to the most important ways to design them out of the product. Figure 1-28 summarizes the general solutions for the four families of signal integrity problems.

The rest of this book discusses the fundamental principles required to understand these problems and the specific techniques to minimize them in your product design.

Noise Category	*Design Principle*
Signal Quality	Signals should see the same impedance through all interconnects
Cross Talk	Keep spacing of traces greater than a minimum value, minimize mutual inductance with non-ideal returns
Rail Collapse	Minimize the impedance of the power/ground path and the delta I
EMI	Minimize bandwidth, minimize ground impedance, and shield

Figure 1-28 Summary of the four families of signal-integrity problems and the general design guidelines to minimize these problems. Even if these guidelines are followed, it is still essential to model and simulate the system to evaluate whether the design will meet the performance requirements.