RIGHT THE FIRST TIME

A PRACTICAL HANDBOOK ON HIGH SPEED PCB AND SYSTEM DESIGN

AUTHOR LEE W. RITCHEY

POWER SECTION BY JOHN ZASIO

EDITED BY KELLA J. KNACK

SPEEDING EDGE SUMMER 2003

Revised 9/06/03 and 0/15/08

Revised 6/2/03
Revised and edited 6/11/03
Revised and edited 8/14/03
Revised and edited 9/15/08

This book is dedicated to the memory of the late Dan Murphy. He paved the way for many of us.

Library of Congress Control Number: 2003109272

ISBN-0-9741936-0-7

ACKNOWLEDGEMENTS

As with all books, this one is the result of the work of many people. At the top of the list are the many students who have attended my “High Speed PCB and System Design” classes over the last ten years. There have been more than 5000 of you. You have challenged me with your questions, prompted me to go back and study topics I had long forgotten and pushed me for better answers. Even more, you have continually asked me when I was going to write the book that goes with the class. Without that continual prompting, this book might still be just a promise. To all of you, thanks for your encouragement and prompting.

The next group of people who played a big role in making this book a reality are all of the senior engineers who taught me when I was a young engineer. They took their time to explain difficult concepts to me; to prepare lectures and classes that I attended at conferences and to write articles and papers that have been in my reference library throughout my career and from which I drew upon for this book. As I went through my career, I vowed that if ever I had the opportunity, I would give back to my profession as it gave to me and all of the other young engineers that started out with me. More of the senior, seasoned, experienced engineers need to stay in the industry and share what they have learned with new engineers. If ever there was a time when this was needed, it is now because of the rapidly changing technologies with which we work.

I have worked with hundreds of very good engineers and designers. From each of them, I have learned things. Many of them have participated in experiments that were aimed at refining the rules I used to do design and to demonstrate concepts. Many of those experiments are in this book.

Early in my career, I switched from microwave and RF design to computer design. When this happened, I was privileged to work with two very good engineers at Amdahl Corporation, Dan Murphy and John Zasio.

With Dan, I later founded a design company known as Shared Resources. Dan was a great router developer and we used his skills to put together a design company and a PCB router that set the standards for how to route high speed PCBs. Many of the features in today’s PCB routers came from the work that Dan did. He made it possible for Shared Resources to design PCBs that could not be attempted with any of the then available design tools. Many of the design concepts in this book came from the work we did designing hundreds of high speed PCBs.

With John, I was able to work on many high-speed designs and to see how he performed analysis of some fairly complex problems. He has always had a design and analytical discipline that stands out among all the engineers I have known and with whom I have worked. I could always count on him to put together good analyses and to perform measurements with the rigor necessary to prove concepts in a conclusive way. He has been working at engineering since the early sixties and is still actively working on new and more complex designs. Whenever I needed a reliable ear to explore an idea, John has been there. For this book, he did the lab experiments that support the statements about bypass capacitors and he also wrote the section on that topic. We still find new things to explore.

In the fall of 1989, I was writing design articles for the early version of Printed Circuit Design magazine. A new editor came on the scene who did major damage to the articles I submitted. Some of them were sent back to me with so much red ink, it looked as though someone had been slain on them. This editor was Kella Knack who had come from Martin Marrietta where she had been a technical writer. I got used to her style and wrote many articles for the magazine. Time went by and Kella moved to the west coast to edit other magazines. We met, I found out she could talk techie talk, she had a sense of humor and liked to sail and backpack. One thing led to another and we became each other’s significant other. She has been my sounding board for this book. She is its editor and, in the end, she is the one who kept after me to get it finished. Without her, the book might have happened but it would not have the quality it does.

FORWARD TO RIGHT THE FIRST TIME BOOK

There were two engineers, a civil engineer, and an electrical engineer who worked together. A mechanical engineer came to visit their project. During a day off, these three colleagues decided to go fishing. They went to a local lake and set out in a small boat.

After a while, the civil engineer said he had to take care of a little personal business. He stepped out of the boat, walked across the lake to shore, did his business, walked back, got in the boat, sat down and continued fishing. A little while later, the electrical engineer said he had to take care of some personal business, got up, stepped out of the boat, walked to shore, took care of business, returned and continued to fish.

The mechanical engineer watched all this in amazement, scarcely believing that the other two had walked to shore and back without getting wet. Wanting to demonstrate that his skill set was as powerful as the other two, he announced that he, too, had to take care of some personal business. As he stepped out of the boat, he promptly sunk out of sight.

Seeing the mechanical engineer sink out of sight, the civil engineer turned to the electrical engineer and said, “Do you suppose we should tell him where the rocks are?”

Like getting to shore without getting wet, high speed design is about knowing where the rocks are. If you do, it’s pretty easy to get to shore without getting wet. If you don’t, you and your project are bound to sink out of sight. The problem is compounded by the enormous amount of fake rocks that are in print as rules of thumb and philosophical rules.

The intent of this book is to show the reader where the rocks are so that a high-speed design can be successfully realized in a straight forward and relatively easy manner. Along the way, the rules of thumb often presented as the “correct way” will be examined to see if they are valid or are the product of someone’s imagination.

The subject matter involved in high-speed design is quite large. So large, in fact, that one book does not and cannot adequately cover it all. There have been many books written at the theoretical level on this topic. Among these are “High Speed Signal Propagation” by Howard Johnson and Martin Graham (Prentice Hall, 2003), and “Introduction to Fields and Waves” by Holt, Wiley and Sons, 1963.

This book will not attempt to repeat the information in these books. Instead, where appropriate, I will refer the reader to the appropriate book to learn more. The reader is advised to obtain copies of these books, or their equal, as part of his or her technical library.

This book will focus on the practical business of turning theoretical concepts into finished PCBs that work right the first time. It is based on more than thirty years experience designing high-speed products ranging from microwave transponders to super computers to terabit routers to network interface cards. What has been learned designing hundreds of high speed PCBs and dozens of high performance systems will be shared with the reader. Not to leave out the other end of the high-speed spectrum, I have worked on elevator controllers, hand held computers, cell phones and PCs that have needed the same design techniques. High speed is high speed, no matter what the product.

As this book is being written, I am actively engaged in the design of next generation products. On the one end are networking products that have 4.8 GB/s and higher data paths in the backplane and at the other are the next round of handheld computers. The knowledge gained working on these products is shared as well.

Many of the illustrations in this book have been taken directly from the screens of oscilloscopes and spectrum analyzers. The intention is to show actual waveforms of real circuits, both failing and functioning properly. This has been done to make it easier for the reader to see how to set up instruments to make measurements and to see what real waveforms look like. The intent is to get as close to the work bench with real PCBs as possible. Some of these measurements were only days old when this book went to press. As a result, their formats may be varied and look like they have been hastily done. In some cases they have been, so that the latest information is available to the reader. I feel that accuracy and timeliness should win over finesse in presentation. For those who are bothered by the changing styles from illustration to illustration, I apologize. This is a practical handbook aimed at those who must get results immediately and is aimed at helping them along the way. Perhaps, someday it may be a textbook. When that happens, it won’t be as current as it should be and will be of less use to the reader. It will be more like a history book. I will try to make sure this never happens.

Note to the reader concerning current flow. When Ben Franklin did his experiments on current flow, he estimated that current flowed from positive to negative. Later it was demonstrated that current was electron flow, moving from negative to positive. The former is called conventional current. The power section of this book discusses current in this manner. The rest of the book refers to current flow as electron flow from negative to positive.

As we got deep into covering the important topics of high speed PCB and system design, it became readily apparent that there was too much information to cover in a single volume. Therefore, this book constitutes Volume 1 and covers all the fundamentals. Volume 2 will cover all the advanced topics and will be available sometime in 2004.

TABLE OF CONTENTS

Figures, Tables and Equations

CHAPTER 1: INTRODUCTION

Figure 1.1. Golden Gate Bridge

Figure 1.2. Boeing 777

What do the Golden Gate Bridge and the Boeing 777 pictured in Figures 1.1 and 1.2 have in common? Both of them were built with no prototypes. In other words, they were right the first time. How was this possible? Sufficient engineering analysis was done prior to construction to guarantee that each piece of each product would function properly. Why perform this level of engineering analysis? To do otherwise, the cost is too high.

Why not apply this approach to electrical engineering projects? Why is this approach not "standard practice" in electrical engineering?

Currently, there are two major camps representing the common methods for developing electronic products. One camp is the proponent of the "hardware prototyping" or the trial and error method while the other camp champions the analytical design or "virtual prototyping" method. In my experience, I liken the former methodology to pocketknife engineering where the design is continually whittled until it does its intended job. In this case, the "whittling" is done by building successive prototypes and correcting design problems found in previous versions. Certainly, early airplanes and bridges were built the same way. When the cost of failing got high enough, engineering skills were developed to improve the chance of being right the first time.

Electrical engineering has evolved in the same manner. Early approaches involved spinning several versions of a design in hardware, identifying the various problems and then creating another hardware iteration. Eventually, the cost of spinning new boards and silicon became cost prohibitive and the tools and skills needed to do virtual prototyping or front-end analysis were developed over time and are in use today in many places.

The question that might be asked is why isn't the second methodology the universal approach to designing electronic products? In my more than forty years experience observing the electrical engineering scene, I have witnessed the evolution of these two design methodologies. In the era before transistor-transistor logic (TTL) came into being, most electrical engineering was done using analytical methods. In fact, one of the fears of a young engineer coming out of college was that his skills weren't good enough to perform at this level. And, of course, they weren't. For this reason, most electrical engineering companies had apprentice programs for new engineers wherein these young engineers were mentored by senior engineers who had experience designing whole products. As time went by and a new engineer gained skill, he was given more complex tasks to do until his skill level reached a point where he could be made responsible for a whole project. This is the way most trades and engineering disciplines are operated. It is the way that I was introduced to electrical engineering.

Then, in 1969, TTL burst upon the scene and, suddenly, nearly anyone could get a working logic system by just connecting up gates and other logic elements. There wasn't much need for Ohm's Law, Maxwell's Equations or network analysis. The only places where real electrical engineering was needed were in supercomputers and RF/microwave products where things happened fast enough that trial and error didn't work. As a result, electrical engineering evolved into computer science and the gold went to those who were excellent at logic design. Engineers good at analytical design were shunted aside and, for those of us with this skill set, electrical engineering became computer science. As time passed, IC design took on the same characteristics, wherein gates and latches were hooked up with no concern for the electrical phenomena associated with it.

Rules of thumb, their origins, validity, etc.

Out of the trial and error design approach grew a large number of "rules of thumb". These rules of thumb were intended to guide engineers as they made early design choices. Early on, rules of thumb were approximations of actual detailed analysis. They allowed an engineer to get a quick sense of what would be needed to achieve a particular goal. Once the actual design started, these rules of thumb were replaced by detailed analysis that established precisely the correct values for components and other electrical devices.

As the trial and error design method became more widely used in the design of electronic products, these rules of thumb often became the only design rules used. Their origin and validity became blurred and they were often applied in cases where they were of little or no value. Even worse, engineers accepted them without any supporting proof that they added value or, at a minimum, didn't degrade performance. Because the speed of logic circuits was slow enough, it mattered little which rules were used. As time passed, more rules of thumb were added, including the no routing via rule, the no right angle bend rule, decoupling capacitor usage rules, etc. These rules were added without the necessary rigorous analysis to insure they fit or, for that matter, were even valid. Many of them added cost without adding benefit. But even worse, some rules of thumb were added--such as ground plane splitting--that actually degraded performance.

Throughout this book, these rules of thumb will be examined to see what they do and if they are of value. When rules of thumb are demonstrated to have value, the reader is cautioned to use them only for doing estimates. When actual design decisions are being made, the analytical methods presented in this book should always be used.

In teaching my classes, when I am asked about the validity of these rules of thumb, I often reply that the only place I know where the rules of thumb are always valid is in a butcher shop when the butcher puts his thumb on the scale to increase the weight of the meat and the cost of the sale. In this instance, the rule only benefits the butcher. The rest of us need to learn analytical methods for arriving at answers to design questions. It turns out that the analytical methods are quite direct and relatively easy to master.

Cost vs. Price

After observing the decision making process in electronic companies for many years, my CPA once told me that he thought electrical engineers were the smartest dumb people he knew. After I recovered from being offended by the remark, I asked him what he meant by this statement. He observed that electrical engineers designed all sorts of innovative, clever products, but often had no idea what things cost and often seemed not to care or think that cost containment was part of the job. These same engineers had the financial success or failure of the company they worked for in their hands and didn't seem to know it.

In the area of printed circuit boards (PCBs), this is especially true. A PCB is a major element of any system. It is the carrier of all of the components. Its quality is key to insuring a product meets its performance and reliability goals. The PCB also contains two key components of the system--all of the transmission lines and the interplane capacitance--that support the very fast edges. If the PCB is bad, it and all of the components on it are lost. If it is a prototype and it is bad, the PCB, the components and the time spent developing it are all lost.

More often than not, PCBs are treated as commodities and are purchased based on price alone. This strategy might work if the method of producing PCBs was so well defined that all producers made PCBs of equal quality. The reader is referred to the document describing PCB vendor selection in the back of this book to see the many instances where the PCB fabrication process can induce failures in a PCB. The price of a PCB is virtually never its cost. When deciding on a PCB supplier, the risk of introducing bad PCBs into the process must be included in the decision making process. When this has been done, it is rare that the low bidder will be selected. After all, in an industry of "equals", how does one become the low bidder?--either by pricing below cost or by taking shortcuts. John Ruskin observed:

> It is unwise to pay too much.
> But, it is worse to pay too little.
> When you pay too much, you lose a little money that is all.
> When you pay too little, you sometimes lose everything, because the thing you bought is not capable of doing the thing you bought it to do.
> The common law of business prohibits paying a little and getting a lot- it cannot be done.
> If you deal with the lowest bidder, it is well to add something for the risk you run and, if you do that, you will have enough money to pay for something better. John Ruskin 1819-1900

A second area within the development process where cost and price are often confused is in selecting design tools and choosing a design process. Often, companies will opt for the hardware prototyping design process thinking that it is cheaper than doing all of that tedious analytical work. I have often heard the statement, "We don't have time to do all that analysis. We've got to get the hardware out." Or, "We can't afford all the tools required to perform the analysis." A critical look at the cost of "doing it again" or the cost of shipping unreliable products quickly shows that this is not the case. This is yet another instance in the development process where my CPA's observations were correct--electrical engineers often don't know where the costs are.

Possible vs. Reasonable

The concept of possible versus reasonable is an important one to explore. Designs are often based on what is possible rather than what is reasonable. For example, a PCB fabricator may say that it is possible to drill and plate an 8 mil hole in a 100 mil thick PCB. However, it is not reasonable to do this. Why? The drilling and plating of such a hole requires that all of the PCB manufacturing processes be at their very best to achieve this objective. It may even be necessary for a process engineer to handle this type of PCB individually. When the question is put in the context of what hole size will result in millions of plated through holes that are reliable over time, the question will result in a very different answer.

The test of reasonableness is an important one. It needs to be asked every time that a tolerance or a methodology is being considered in order to insure that the rules being applied fit the situation at hand. This is especially true when it comes to PCB fabrication. In the process of winning orders, PCB salesmen often promise something that is not reasonable. Worse, they often encourage it. I don't think that this is done deliberately. Rather, it is likely done out of ignorance. There is an old saw, "What is the difference between a used car salesman and a computer salesman? Answer: "The used care salesman knows when he is lying to you!" If you substitute PCB salesman for computer salesman, this is often the case in the PCB industry. This may seem a little cynical, but it is my experience that this is true. Worse, very often a fabricator sees a new PCB design as a way to extend its capability into a new technology and will tackle a design that is likely to fail. This extending of capability is a noble goal, but not with someone's prototype. If the PCB turns out bad, the fabricator will often remake it at no charge, making it seem like the customer has lost nothing. **Nothing could be further from the truth**. The customer is out the elapsed time to remake the PCB as well as the cost of all the components and this will always be much more expensive than the cost of the PCBs. In this case, the reasonableness test applies to the selection of the supplier.

ON MAKING THINGS TOO SIMPLE

Everything should be as simple as possible,

And no simpler.

Albert Einstein

Visible vs. Significant

As part of composing a set of design rules using simulations or testing, it is important to view each result in the context of "will the phenomena being studied have a material effect on the design?" The object of the rule creation process is to include only rules that have a benefit that justifies their cost. As an example, a via can be used to route a signal from one layer of a PCB to another. It looks like a tiny parasitic capacitance added to the net. When such a net is tested with a measurement system, such as a TDR (time domain reflectometer), a tiny reflection will be seen as caused by the parasitic capacitance of the via. As will be demonstrated later in the book, in virtually all cases this reflection is of no consequence. The via is visible, but not significant. Placing a no via restriction on routing a PCB will make it more difficult to route without improving its performance.

There exist simulations that show tiny ripple voltages on power planes as current is drawn from them to perform switching operations. These are visible in the simulation, but they are rarely large enough to be of any consequence. Adding rules such as edge plating a PCB or adding rows of "ground" vias around the periphery of the PCB to control this adds cost without adding benefit.

Developing a Rule Set

Design rules come from a variety of places: Some rules are explained by saying "we have always done it this way;" some are described as "standard practice;" some are explained as having been developed by the senior engineering staff and must be followed without deviations; some are borrowed from other designs and others are taken from applications notes.

Still others are taken from fellow engineers. While others come from using analytical methods to develop design rules that are inappropriate for the design being undertaken.

Clearly, the giver of a design rule should be certain that the rule being given is valid, be prepared to demonstrate the derivation of the rule and show why it is valid. After all, this is the way the scientific method works and is the principle upon which electrical engineering is based. A theorem or hypothesis is deemed invalid until its validity is proven with analytical work, lab testing or both. Sadly, this is often not the case. The publishers of applications notes should do the same. Actually, publishers of application notes and guidelines should be held to an even higher standard. This is especially true since users of these application notes and guidelines expect the design rules contained within them to be valid and often bet not only their designs but also their entire companies on them. The unfortunate truth is that most applications notes are not prepared with this level of rigor.

Even worse, some engineers who give out rules of thumb and are asked to support their rules with scientific analysis or experiment are offended that the questioner doesn't trust them. Such rules should always be viewed with suspicion. The "we've always done it this way," design rule set is of equal concern. Many times, it has been said that if the only reason one has for doing something is this, it is a convincing reason to stop. There is a high likelihood a rule with this as its only basis for validity is doing harm rather than good.

Many projects start off with some design rules where their validity is uncertain or the point at which they begin to fail is unknown. To the extent that this is the case, allowances must be made for the uncertainty this creates and preparations made to encounter some failures as a result. A list of rules that are not demonstrated to be valid should be made and a program to investigate them put in place.

This book is intended to help electrical and electronic engineers, as well as PCB designers, achieve the objective of being right the first time when designing systems with PCBs in them. It draws on more than three decades of virtual prototyping of high speed PCBs used in all types of high-speed products, ranging from super computers to terabit routers, as well as low-cost PC cards. All of the steps in the process, from selecting tools to selecting suppliers, will be discussed in detail.

CHAPTER 2: THE ELECTRICAL ENGINEERING PROBLEM

It is useful at the start of a project to ask the question "What problem are we trying to solve?" When this is done with an electrical engineering project, the following is the result.

All electronic systems generate voltage waveforms that contain information. This information may be analog, digital or RF. Circuits take these voltage waveforms and perform some operation. This operation may be to open a garage door; to play music on a speaker; to calculate the value of π to five thousand places or to send a picture to grandma over the Internet. This, then, defines the problem. It is to create voltage waveforms of sufficient quality to convey the intended meaning, transport those voltages to circuits utilizing them and then use the waveforms to perform an operation, all the while making sure that other signals or noise do not corrupt the operation. At the same time, the intended signals or voltage waveforms must not corrupt other signals. This corruption might take the form of crosstalk or EMI (electromagnetic interference).

It is worth noting that the usable signal in all cases is a voltage waveform. This is true whether the signal is RF, microwave, analog or digital. It is also true no matter what type of driver is being used. It is true even if the signal source is described as operating in the "current mode." As will be seen later, it is not possible to send "voltage" waveforms from one place to another. In order to deliver a voltage waveform it is necessary to create and send energy in the form of an electromagnetic wave, either down a transmission line or through space. Understanding this and how electromagnetic waves behave is fundamental to successfully designing high-speed electronic devices.

TYPES OF HIGH SPEED PCBS

Over time, the world of high-speed electronics has been split into two general classes of PCBs: RF/microwave/analog and Digital. Table 2.1 lists these two classes and their key characteristics.

RF, MICROWAVE, ANALOG PCBs	DIGITAL-BASED PCB
Low circuit complexity	**Very high circuit complexity**
Precise matching of impedance	**Tolerant of impedance mismatches**
Minimizing signal losses essential	**Tolerant of lossy materials**
Small circuit element sizes	**Small circuit element sizes desirable**
Usually only 2 layers	**Many signal and power layers**
High feature accuracy need	**Moderate feature accuracy needed**
Low/uniform dielectric const	**Dielectric constant secondary**

Table 2.1. Characteristics of Two General Classes of High Speed PCBs

Industry has developed design tools, materials and manufacturing methods that are optimized for each major class of PCBs. The rate at which speeds of logic circuits have increased has moved many digital designs into the speed range normally considered to be RF or microwave. It is useful to compare these two classes to see how they differ.

Circuit complexity- RF/microwave/analog circuits tend to be relatively simple with few components. It is often possible to memorize their schematics. On the other hand, digital electronics tend to have very large schematics with high lead count components. It is rarely possible to memorize their schematics. Where this impacts design most is in choosing design tools. With the simpler circuits, the focus is on the ability to create complex geometric shapes on one or two layers. With digital designs, the databases needed to represent the design at various stages are very large. Keeping the databases in step with each other, as well as insuring that the databases are accurate, is a major concern. **For this reason, successful digital design requires that design tools facilitate database management and checking.**

Impedance matching- At each point along a transmission line where the impedance changes for any reason, some of the signal energy traveling on it is reflected back to the source. As a result, that signal energy is no longer available to create the voltage waveform that the receiver or user needs to function properly. In order to maximize performance in products such as radios, every effort is made to make transmission lines that are uniform in impedance. This is not especially difficult, as there tends to be only a few paths that must be managed.

In digital circuits, there are so many signal paths that function as transmission lines that it is not possible to achieve precision impedance matching on all of them. As a result, logic circuits are designed such that there is a tolerance for impedance mismatching. This tolerance does not allow for arbitrarily large amounts of mismatch. Therefore, it is necessary to invoke a level of impedance matching that is adequate for proper operation. One of the problems this creates in a design is to know how much tolerance for impedance mismatching a logic family has and insure that the design rules maintain this level of precision.

Signal losses- These losses occur due to absorption in the dielectric and resistive losses due to skin effect loss at high frequencies. Absorption in the dielectric, or dielectric loss, is managed by selecting insulating materials that have low loss tangents. With RF and microwave circuits this is usually handled by selecting a material with a Teflon® base. Unfortunately, Teflon®` is a dimensionally unstable material. It is also a non-stick material. As a result, it is difficult if not impossible to manufacture the large, high layer count PCBs required by digital circuits from it. The same part of the tolerance budget used for impedance mismatch will be used to compensate for dielectric losses.

Skin effect losses are associated with current flowing in conductors crowding into a thin layer near the surface at high frequencies. This loss gets larger as frequencies increase. It is compensated for by plating the outer surfaces of conductors with gold to increase conductivity or by increasing the width of traces to create more surface area. The first approach can only be used on surface layers of microwave/RF PCBs, while the second approach is used on logic PCBs.

Circuit element sizes- All components mounted in packages have their performance degraded due to unwanted package parasitics, such as lead inductance, lead capacitance and lead-to-lead crosstalk. The faster a component must operate, the more these unwanted parasitics degrade performance. Minimizing the size of the package in which the component is mounted minimizes these parasitics. RF and microwave components are often mounted directly on a PCB without a package. Logic circuits often have hundreds of leads. In order to make room for all of the contacts, the package cannot be small. Therefore, the design rules used with logic circuits must allow for these parasitics.

Layer count- Minimum layer count is always a desirable goal as it keeps the cost of a PCB down. Most of the time, it is possible to implement RF and microwave designs in two or three layers. High-speed digital designs cannot be done in two layers for several reasons. The most obvious is the need for layers to hold all of the signal wires or transmission lines. A not so obvious reason is the need for a robust power distribution system of low inductance planes and the capacitance that exists between them.

Feature accuracy- Figure 2.1 shows the surface layer of a microwave PCB. Odd shapes are visible on the left hand side of the PCB. These are band pass filters, power splitters, directional couplers, tuned circuits and other components. They are formed by etching precise shapes in the copper foil of the outer layer of the PCB. The frequency at which these functions are executed is directly dependent on how accurately they are etched. In a digital PCB, these functions will be performed by discreet components so this kind of feature accuracy is not required. What is needed in a digital PCB is accurate trace width and dielectric thickness control in order to insure that transmission line impedances are kept within allowable tolerances.

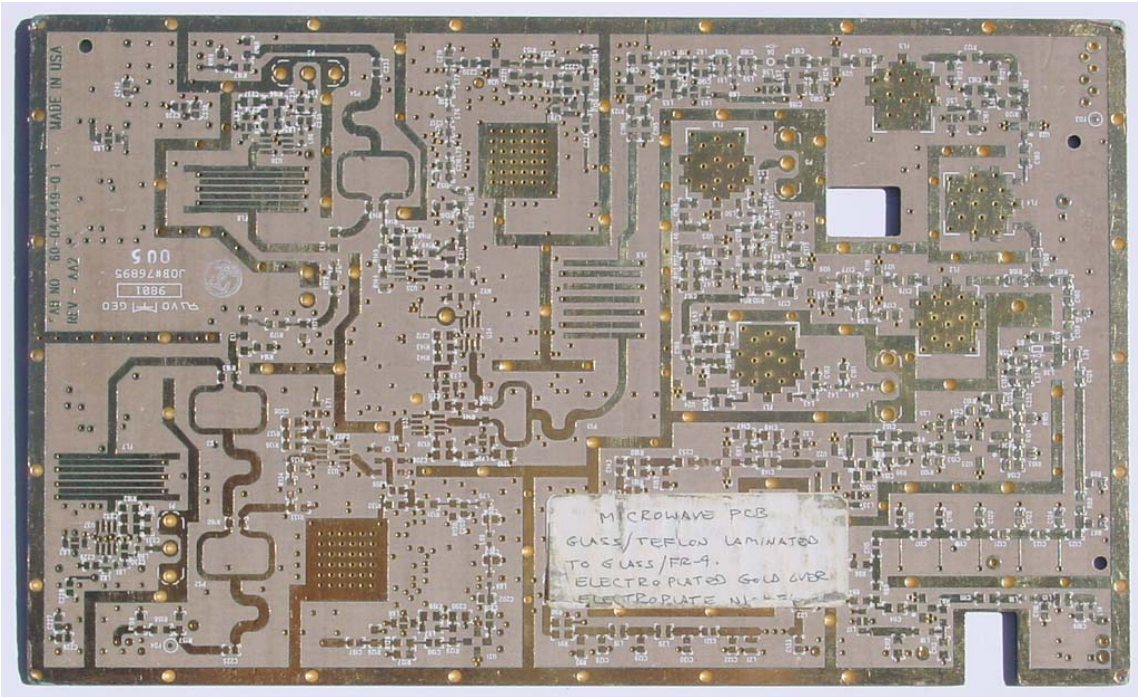

Figure 2.1. A Typical Microwave PCB Surface Layer

Low/Uniform dielectric constant- The dielectric material in the substrate supports the transmission lines in a PCB. This dielectric material increases the unwanted parasitic capacitance of every transmission line. In order to send fast signals over the transmission lines, it is necessary to charge up and discharge this parasitic capacitance. The larger the parasitic capacitance, the more energy it will take to perform a given operation. Therefore, it is desirable to minimize this parasitic.

One approach is to use low dielectric constant materials, such as Teflon®. This works for two layer RF/microwave PCBs. However, logic circuits need a dielectric that can withstand high temperatures, be dimensionally stable and laminates well. These requirements in turn dictate the choice of dielectric. The designer will need to compensate for the dielectric constant that results--a topic that will be discussed in later chapters.

CHAPTER 3: MAJOR ELEMENTS IN AN ELECTRONIC SYSTEM

As noted previously, an electronic system creates and uses voltage waveforms. This requires several different elements including: voltage waveform creators (drivers or sources); transmission lines or wires on which these signals or voltage waveforms are delivered; loads (circuits or receivers) and a power source or power system. They occur in order as follows:

Sources
Transmission lines
Loads
Power Sources

Table 3.1. Major Elements in an Electronic System

Sources or **Voltage waveform creators** vary greatly. Table 3.2 lists potential sources of voltage waveforms. All of these have characteristics such as output impedance, dynamic range, maximum slew rate (rise or fall time) or rate of change in voltage or frequency.

Antennas
Wave guides
Transducers
Logic drivers
Analog drivers
RF drivers
Transformers
The Sun

Table 3.2. Types of Signal Sources or Voltage Waveform Generators

Transmission lines are used to move a voltage waveform, actually an electromagnetic wave, from its source to its user and can take several forms. Table 3.3 lists many of them. All of these are transmission lines of one sort or another that guide the electromagnetic energy from the source to the user. While it is true that the objective is to deliver a voltage waveform, creating and sending energy in the form of an electromagnetic field is the only way to accomplish this.

Coaxial cables
Twisted Pairs
Twin leads
Ribbon cables
PCB traces
Waveguides
Space

Table 3.3. Methods of Moving A Voltage Waveform or Electromagnetic Field From Its Source to its User

Loads are the consumers of the voltage waveforms transmitted from the sources over the transmission line media as electromagnetic energy or waves. Table 3.4 lists many types of loads. These loads respond to the voltage waveforms presented at their inputs and cause some action, such as a logic operation, to take place. Nearly all loads absorb little, if any, of the energy contained in the electromagnetic field. As a result, this unabsorbed energy can reflect back along the transmission line causing corruption of the voltage waveform.

MOS transistor gates
TTL Emitters
ECL bases
Parallel terminators
Transformers

Table 3.4. Types of Loads

Most of these inputs look like tiny parasitic capacitors that take some charge from the electromagnetic field but do not absorb significant energy. The one exception is the terminating resistor. It does absorb the energy contained in the electromagnetic field and thereby prevents reflections.

Power sources supply the energy required by signal sources to create the electromagnetic fields involved in voltage waveforms. Table 3.5 lists several types of power sources.

Batteries Discrete capacitors Plane capacitors Voltage regulators Transformers

Table 3.5. Power Sources

Each of these power sources can be thought of as a reservoir of charge that will be drawn from by the driver circuits to perform some function. In logic circuits, the most common operation is to charge up the parasitic capacitance of the signal lines or transmission lines in order to change their logic state from a zero to a one. Each of these sources has some "output" impedance that is not zero. As a result of this non-zero impedance, when charge is drawn from the power source, its terminal voltage drops. This terminal voltage drop is usually called "ripple" when it occurs at a rapid rate and "droop" when it occurs slowly.

CHAPTER 4: ASSUMPTIONS OFTEN MADE ABOUT ELECTRONIC SYSTEMS

When a logic design is done, the elements are symbolic. That is, they are logical operations that are performed on a logic input signal with no delay or other phenomena, such as reflections, occurring. All that a designer need do is combine the right collection of gates, latches, memory, etc and the design is done. With slow logic circuits, it has been possible to take these logic designs and turn them into functional products. It has not been necessary to be concerned with electrical effects in order to succeed. Just connect up the parts and things always work. This design process is often called TTL engineering since the most common logic family that comprises these slow circuits is TTL.

Out of TTL engineering has grown a collection of assumptions about electronics. Among these are:

- There is no time delay between inputs and outputs.
- The voltage waveforms at both ends of a wire or connection are the same
- The power distribution system is ideal
- IC packages are ideal
- Driver outputs are ideal
- Signal inputs are ideal
- Connectors are ideal
- Wires have no delay
- Wires are ideal
- Noise is not a problem

As speeds of components have increased, the foregoing assumptions no longer apply. Circuits that used to work every time are flaky or don't work at all. From these failures have come trial and error fixes, rules of thumb, and the notion that some kind of black magic must be at work that defies explanation. The chapters that follow will demonstrate that all of the components are actually well behaved and follow a well-defined set of the laws of physics. They are indeed not ideal but are good enough to allow designs to work as long as their non-ideal effects are taken into account. **It might well be stated that high-speed design is the management of the non-ideal.**

CHAPTER 5: HOW DIFFERENT FROM IDEAL REAL SYSTEMS AND THEIR COMPONENTS ARE

As was hinted at in chapter 4, real parts are not ideal. Their departure from ideal is what makes circuits not behave as expected. The departures from ideal, outlined below, will be discussed in detail in later chapters. They include:

- Real devices have delays from input to outputs. These delays vary significantly from device to device, even among devices of the same type that meet manufacturer specifications.
- The voltage waveforms at the two ends of a wire or transmission line are normally different.
- There is time delay from one end of a wire or transmission line to the other.
- The power distribution system has many defects that can compromise performance.
- IC packages have unwanted parasitics, mainly inductances, that can degrade performance
- Driver outputs are not ideal. They have limits, such as finite, non-zero, rise and fall time and a non-zero output impedance.
- Signal inputs are not ideal. They have unwanted parasitic capacitance.
- Connectors are not ideal. Their impedance may not match the line into which they are connected. They have crosstalk. They have resonances.
- Wires are not ideal. They have parasitic inductance and parasitic capacitance as well as resistance.
- Noise can couple into a circuit from several sources.
- Capacitors are not ideal. They have parasitic resistance and parasitic inductance.

It is this departure from ideal that gives rise to the need to understand the physics involved in performing electronic operations. This understanding and applying the appropriate laws of physics is often called signal integrity engineering. It is my hope that by the time the reader has reached the end of this book, these topics will not seem mysterious or magic. Furthermore, I hope that the reader will understand that this is the stuff of all electrical engineering and not of some specialists.

With such non-ideal components, how is it possible to make high-speed circuits work properly? First, every connection needs to be treated as a transmission line and the rules that govern transmission lines need to be applied. Second, the circuits that are designed to work at high speeds have some tolerance for a non-ideal environment. The design rules used with these circuits must account for this degradation while keeping it within each circuit's tolerance.

CHAPTER 6: TRANSMISSION LINES

What is a transmission line? In its simplest form, a transmission line is any pair of conductors that is used to guide energy, in the form of an electromagnetic field, from one place to another. In a PCB, a transmission line is a trace and one or two planes. Table 6.1 lists examples of transmission lines with which the reader should be familiar.

Power lines are transmission lines
Waveguides are transmission lines
TV twin lead is a transmission line
Coaxial cable is a transmission line
Twisted pairs are transmission lines

Table 6.1. Examples of Transmission Lines

In all of these cases, except the power lines, the objective is to deliver a voltage from a source or output to an input rather than to deliver a bundle of electromagnetic energy. The power lines are the only ones used to deliver energy to users. The movement of energy in all of these transmission lines is managed with the same rules. This movement can be described and analyzed using Maxwell's Equations.

CHAPTER 7: WHAT'S MOVING ON A TRANSMISSION LINE?

In order to properly manage the behavior of signals on transmission lines, it is necessary to understand what is actually moving on those lines. Up to this point in this book, electromagnetic fields have been mentioned many times. In fact, electromagnetic fields are the entities that are in motion on a transmission line. Attempts are often made to describe what is happening with a signal by talking about current flow or current loops. This always falls short of accounting for signal behavior. **The reason is, current is a consequence or symptom of the presence of an electromagnetic field, not the reverse.**

> **Current flow is a consequence or symptom of the presence of an electromagnetic field, not the reverse.**

In order to understand the movement of energy on a transmission line, it is useful to use a mechanical model. A mechanical transmission line illustrates how energy, in the form of an acoustic wave, moves on the transmission line. Electromagnetic energy behaves in the same way. Figure 7.1 is an example of a mechanical transmission line. Most readers have seen such a device in action. In this case, the transmission line is made of steel and the energy traveling down it is in the form of an acoustic wave. The movement of acoustic energy in this structure can be described with the same equations used to describe the movement of electromagnetic energy in an electrical transmission line. The three variables in the mechanical transmission line are spring force, mass and resistance. Later, it will be seen that the equivalent three variables in an electrical transmission line are inductance, capacitance and resistance. (A question that is often asked is "Where is the spring in the mechanical transmission line?" It is in the elasticity of the steel that makes up the balls.)

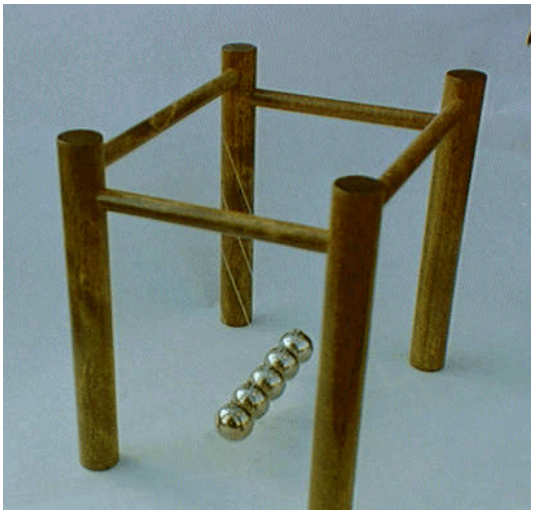

Figure 7.1. A Mechanical Transmission Line

When a ball at one end of the transmission line is swung up and away from the string as in Figure 7.2, it contains potential energy. Upon release, the potential energy changes to kinetic energy. When the ball collides with the end of the transmission line as in Figure 7.3, all of its kinetic energy is transferred into the transmission line. That energy travels at the speed of sound through the transmission line as an acoustic wave and is transferred into the last ball. The balls in the middle of the transmission line don't move, they simply compress and expand (spring force) as the acoustic wave passes through them. The last ball swings up as high as the original ball turning the kinetic energy back into potential energy as depicted in Figure 7.4.

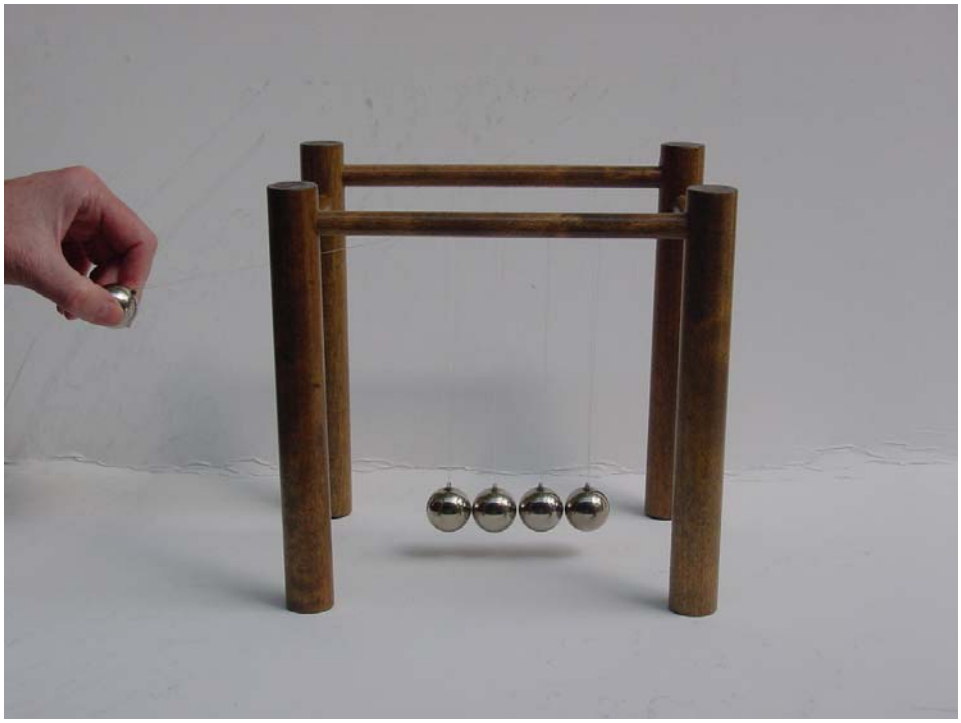

Figure 7.2. Potential Energy Stored in Elevated Ball

Figure 7.3. Kinetic Energy From Falling Ball Transferring into Transmission Line

Figure 7.4. Kinetic Energy Transferring to Ball at Far End of Transmission Line

Since there is nothing at the “receive end” of the transmission line to absorb the energy transferred into the last ball, it reverses direction, falls back, collides with the transmission line and transfers the energy back into the transmission line. In principle, the process repeats itself forever. In actuality, each time the energy traverses the transmission line, a small amount of it is absorbed in the line resistance, slowly dampening out the swing.

The fact that there was nothing to absorb the energy as it arrived at either end of the transmission line resulted in it being “reflected” back down the line. If this “reflection” is undesirable, something must be placed at the “load” end of the line to absorb the energy. Figure 7.5 shows such an absorber. In an electrical transmission line this absorber would be called a parallel termination.

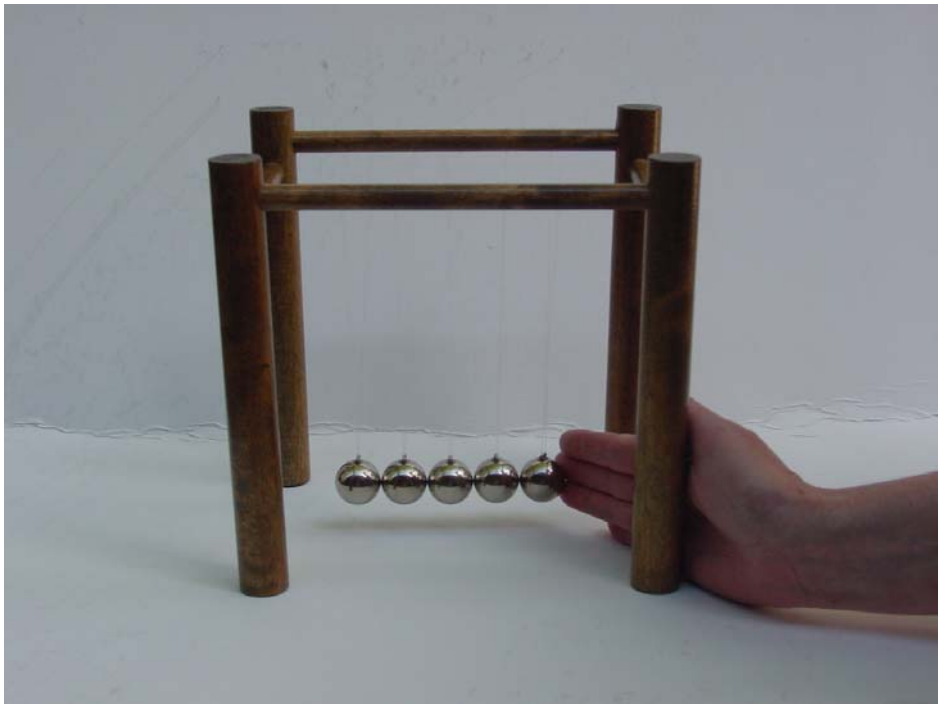

Figure 7.5. Absorbing the Acoustic Energy at the Load Using a "Parallel" Termination

Another possibility is to allow the energy to “reflect” back at the load end and travel back to the source. The energy could be absorbed there as depicted in Figure 7.6. In an electrical transmission line this type of absorber would be called a series termination.

Figure 7.6. Absorbing the Acoustic Energy at the Source with a "Series" Termination

An electromagnetic transmission line behaves in exactly the same manner as the mechanical transmission line. An electromagnetic wave is launched down the transmission line. If there is no absorber or terminator at the load end of the line, all of the energy reflects back toward the source or driver. If this reflection is undesirable, the energy must be absorbed at the load using a parallel termination. If it is not possible to use a parallel termination, the energy will reflect back down the transmission line and reflect off the source or driver end unless the source or driver is of such an impedance that it can absorb it. If, upon arrival back at the source, the electromagnetic field sees an impedance of the same magnitude of the transmission line, this impedance will absorb the energy and reflections will stop. This method of absorbing the energy is called a series termination. Allowing the energy to reflect off the open end of the transmission line and absorbing it at the driver or source is often referred to as reflected wave switching.

If there is something in the middle of the transmission line that causes a change in its impedance, such as a connector or a lumped load, some of the energy will reflect back to the source at this change and not be available to develop a signal voltage at the load. This degrades the signal. In order to keep the signal degradation within reasonable limits, the impedance of all parts of a transmission line must be held within limits determined by the tolerance of the using circuit. **This is known as controlling the impedance.**

The operations discussed above are fundamentally what happen with all transmission lines, whether they are used for RF, microwave, analog or digital. The basic concept is quite simple. The art in engineering transmission lines is in knowing the possible sources of degradation, knowing how much degradation is acceptable and then invoking design rules that keep degradation within limits that result in transmission lines which deliver a satisfactory signal or voltage waveform. There is no magic, black or white. However, successful transmission line design does require an understanding of the behavior of electromagnetic fields.

CHAPTER 8: BASICS OF ELECTROMAGNETIC FIELDS

Figure 8.1 shows the electromagnetic field surrounding a stripline transmission line. This is an end-on view of a trace traveling between two planes. Notice that there are two fields--the magnetic field and the electric field. Any time energy moves from one place to another on a transmission line or in space it does so as an electromagnetic field. It is important to remember that when energy is transmitted, both fields--electric and magnetic--are always present. Those two fields are always at right angles to each other. The voltage difference between the transmission line trace and the surrounding planes is a measure of the strength of the electric field. The magnitude of the current flowing in the transmission line is a measure of the strength of the magnetic field. The product of these two factors is a measure of the energy or power flowing on the transmission line in the electromagnetic field (P =I x V).

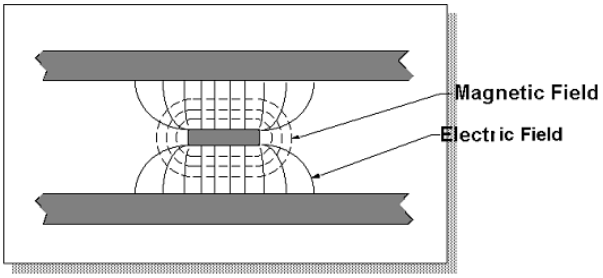

NOTE:
CURRENT FLOWS ONLY WHEN A MAGNETIC FIELD IS PRESENT.

A VOLTAGE DIFFERENCE BETWEEN THE LINE AND ITS SURROUNDINGS EXISTS ONLY WHEN AN ELECTRIC FIELD IS PRESENT.

Figure 8.1. An End-on View of a Stripline Transmission Line Showing the Electromagnetic Field Traveling on It

It is possible for one field to exist without the other such as when there is a capacitor charged to some voltage or there is a transformer with a steady current flowing in its primary circuit. In both cases, only one of the two fields is present. The capacitor has an electrostatic field between its plates. The transformer has a fixed magnetic field surrounding its primary windings. Energy is stored in each field, but it is not moving.

Electromagnetic fields are not only the province of high-speed signaling. They are present in all electronic circuits, no matter the speed at which they operate. This includes things such as flashlights, switching power supplies and AM radios. Electromagnetic fields are of concern in high-speed circuits **because they change rapidly**. These rapidly changing electromagnetic fields are affected in an adverse way by the parasitic inductance and capacitance that exist in the circuits through which they travel.

> Electromagnetic fields are present in all electronic circuits.
>
> Managing them carefully becomes important when they change rapidly.

Because most engineers are first introduced to current flow and voltage drop, it is not uncommon to incorrectly draw the conclusion that the current flow results in the creation of the electromagnetic field. This has given rise to many explanations of what takes place in a transmission line circuit in the context of current flow or current loops. This method of looking at what is happening falls short of accounting for what actually occurs. As a result, many effects that do occur, such as EMI, don't seem to behave as expected. If the analysis or explanation is based on the behavior of the electromagnetic fields that are present and an environment is created that properly deals with them, the current flow that results is automatically handled. The reverse is not necessarily true.

Consider a transformer. An electromagnetic field emanates from the primary winding. The magnetic lines of force in this field slice through the windings of the secondary inducing a voltage in the secondary winding, but no energy is flowing. If there is a path for energy to flow from the secondary, such as a load, then an electromagnetic field flows through the load, inducing a current flow in it.

Current flow is induced by electromagnetic fields.

Electromagnetic fields are not induced by current flow.

It is not uncommon to struggle with the concept of current flow occurring as a result of the presence of an electromagnetic field. To illustrate that the current flow exists because an electromagnetic field travels on or gets close to a conductor and induces a current flow in it, consider the signals listed in Table 8.1. All of these are examples of electromagnetic energy traveling from one place to another and there are no electrons in the path to create a current flow.

.

• **Radar Waves**	**Television Signals**	**Light waves**
• **AM Radio Waves**	**Pager Signals**	**EMI energy**
• **FM Radio Waves**	**RFI Energy**	**Microwave Signals**

Table 8.1. Examples of Electromagnetic Energy Moving from One Place to Another Without Electrons in the Path

Mastering the concept of electromagnetic energy moving from one place to another in the form of an EM (electromagnetic) field is crucial to understanding how high speed signaling takes place. It is also crucial to understanding how EMI happens.

Notice that EMI and RFI are on the same list with useful signals. This is no coincidence. EMI and RFI are made of the very same energy as those useful signals. That is the reason they are referred to as interference. The energy in EMI and RFI escapes from a product and interferes with other products. The methods of escape are the same as those used to create the useful signals. A source of RF energy is coupled to a radiating surface (antenna) in both cases. Containing EMI and RFI is accomplished by removing the RF energy source or by making sure that the RF energy cannot reach a radiating surface or antenna. The methods for controlling EM will be covered in Volume 2. It's not black magic nor is it mysterious.

CHAPTER 9: DIGITAL vs. RF/MICROWAVE vs. ANALOG

There is the notion that RF/microwave or analog and digital circuits are different and different rules are needed to successfully manage them. In some ways they are different, but in more ways, they are the same. All three rely on voltage waveforms to convey information. All three require an EM field to propagate the voltage waveform from the source or driver to the receiver. All three require the same design discipline to insure the signal being sent is not unduly distorted by reflections, loss and noise from outside sources.

The only difference between these signal types is in magnitude of the voltage waveform and tolerance to distortion and noise.

RF and microwave circuits tend to have very small amplitude signals at the receiver. So, the receiver environment needs to be well protected from noise sources that could distort these small signals. Usually, RF and microwave circuits also have a broad band of frequencies in the received signal. In order to prevent phase or amplitude distortion of this band of frequencies, the circuits that process them need to have uniform gain and wave velocity across the pass band.

Analog or video circuits convey information in the shape or amplitude of the voltage waveform being sent. As a result, it is important to design the circuits over which the signals travel such that distortion and external noise are held to a minimum.

Digital circuits convey information at two different voltage levels. In rare cases, four voltage levels are used. The drivers of digital circuits are designed to create signals with voltage levels that are larger than those required by the receivers. As a result, digital circuits can tolerate a significant amount of noise and distortion (loss) and still convey information successfully.

In all of these cases, transmission lines are used to guide the signal as an electromagnetic field from the source to the load. The same transmission line management principles apply in all cases. Once the skill of creating and managing transmission lines for one is mastered, it is a small leap to successful management of the others.

CHAPTER 10: TIME AND DISTANCE

As circuits operate increasingly faster, the time available to perform each operation becomes increasingly less. There comes a point where the time required for a signal to travel from one part of a design to another is a significant part of the total available time. Similarly, when the rise and fall time of a switching edge becomes short compared to the length of the transmission line over which it travels, effects such as crosstalk and reflections become important.

Developing a sense for time versus distance is a critical skill when creating design rules for fast circuits. This skill makes it easier to decide when a circuit requires transmission line management techniques and when it doesn't.

Electrical signals travel as electromagnetic energy. This EM energy travels at the speed of light, 186,000 miles per second or 300,000,000 meters per second in a vacuum. Doing a little math, this comes out to roughly one foot or 30 centimeters in a billionth of a second or a nanosecond (10^{-9} second). When this electromagnetic energy must travel in a dielectric or insulator, such as in a PCB, it is slowed down by the dielectric. In a PCB, the dielectric slows the signal down to roughly half the speed of light in a vacuum. As a result, a signal will travel roughly 6" or 15 centimeters in one nanosecond along a trace in a PCB. It is handy and reasonable to think of nanoseconds in a PCB as 6 inches. We will call on this relationship many times as the size versus speed discussion unfolds. It is useful to remember this relationship. .Equation 10.1 shows the relationship between velocity and the relative dielectric constant of the material used to construct the transmission line.

$$\sqrt{e_r} = \frac{C}{V}$$

Where e_r is the relative dielectric constant of the insulating material, C is the speed of light and V is the velocity in the dielectric. (Dimensional units, if the same scale, for C and V are unimportant)

Equation 10.1. Velocity vs. Material Relative Dielectric Constant

The rise and fall times of currently available logic devices range from about 200 picoseconds (200 x 10^{-12} seconds) to two nanoseconds. If these times are converted to travel time on a PCB trace, the corresponding lengths are 1.2 inches (3 centimeters) to 12 inches (30 centimeters). Such a time-to-distance conversion is a useful way to assess whether signal lines are "long" or "short". As will be illustrated in a later chapter, when a transmission line is long enough to be an appreciable portion of the "length" of a switching edge, transmission line problems such as reflections and crosstalk can become sources of failure and must be carefully managed. (It should be noted that newer ICs designed for circuits such as 10 GB/S signal paths have switching edges as fast as 50 picoseconds (50 x 10^{-12} seconds.) That is about 0.3 inches! At that rise time, everything in the path is long including the wires in the package from the die to the signal pin that connects to the PCB.

Figure 10.1 illustrates what is meant by rise and fall time for logic signals. The classic definition of rise time is the time required for a signal to travel form its 10% value to its 90% value. Fall time is the reverse of this. For very fast logic circuits, such as ECL and GaAs, the definition for rise time is 20% to 80%. The difference in measurement from the classic definition occurs because very fast switching edges tend to have rounded shoulders. In these very fast circuits, the 10%-90% measurement would yield rise and fall times that are slower than the actual edges.

These two different definitions point out a problem with defining how fast a circuit is using rise and fall time. A 1 nanosecond edge that switches 5 volts is much faster that a 1 nanosecond signal that switches 1 volt. A better, more accurate way to express the speed of an edge is with slew rate (volts per nanosecond for example.) This measurement more accurately communicates speed. In the end, what is of concern is the speed at which a signal travels through the forbidden zone between a logic zero and a logic one or the reverse.

There are integrated circuits that are being specified this way. Specifically, there are several FPGA manufacturers that have output drivers that are called "variable slew rate" outputs. These drivers allow the designer to select a rise or fall time that is just fast enough to perform the operation being specified (the advantage of this will become apparent later in this book when EMI and power subsystem design are addressed in detail). In essence, drivers with variable slew rate outputs allow a designer to create edges just fast enough to get the job done. This is far more preferable than being forced to use standard parts that have ultra fast edges particularly when the application doesn't require them.

For most designs, the only measurement of speed is rise and fall time. This measurement is often referred to as edge rate even though the two terms aren't exactly interchangeable. When edge rate is used in this book, it refers to rise or fall time. In a perfect world, all signal edge speeds would be specified as slew rate. Unfortunately, there is so much legacy information in analytical tools and instrumentation that relies on rise and fall time that we will continue to analyze circuits using this measure of speed.

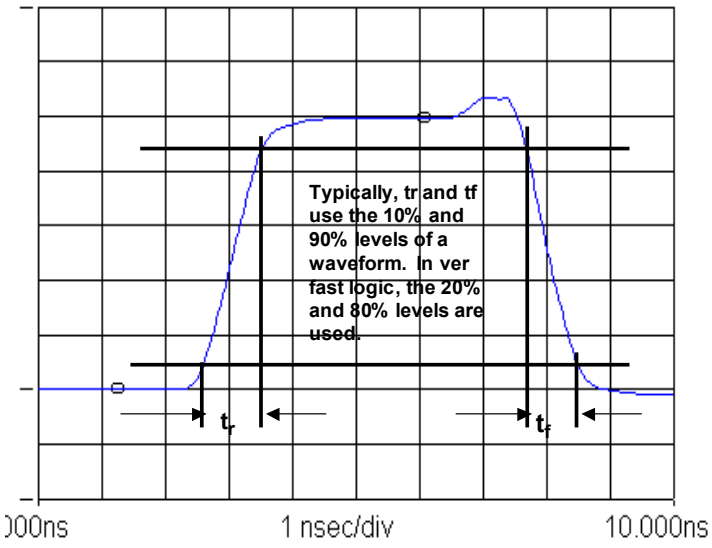

Figure 10.1. Rise and Fall Time

From the calculations above, it is easy to see that many of the signal paths in a PCB will be long enough to have potential reflection and crosstalk problems.

In a PCB a nanosecond is 6" or 15 cm.
An inch is 167 picoseconds
A centimeter is 67 picoseconds

When developing design rules, the relationship of the length of conductors in a design to the duration of the rise and fall time of each switching edge is crucial. To determine whether or not a design is high speed is accomplished by comparing these two factors to one another. The methodology for doing this comparison will be discussed later.

Table 10.1 lists several common logic families and their fastest rise times. These times have been converted to lengths in a PCB using the typical velocity of approximately 6" per nanosecond. The column titled "Transition Electrical Length" lists the electrical length of each of the rise times. The column titled "Critical Length" lists the approximate electrical length at which two transmission lines running in parallel will achieve worst-case backward crosstalk in a PCB.

The name "**Transition Electrical Length**" (TEL) was coined to provide a name for the length of a switching edge in a PCB or other transmission line. Sometimes this length is incorrectly labeled "critical length". Calculating TEL for a given rise time is a method for obtaining a first order estimate of whether or not a design is fast enough to require the use of transmission line management rules. If a design has lines as long as a TEL for that logic family, the designer can be assured that the design will require careful management of impedance in order to avoid reflection problems.

The term "**Critical Length**" was a label created by RF engineering decades ago to describe the length at which two transmission lines running in parallel achieve maximum backward coupling or crosstalk. In the case of RF engineering, this phenomenon can be used to take a fixed sample of the energy in a main line by deliberately coupling two lines together. In logic design, no such beneficial results come from tightly coupling lines to one another. Instead, tight coupling results in unwanted crosstalk. The critical length calculation provides a first order estimate of the lengths at which crosstalk is likely to be a problem.

LOGIC TYPE	TYPICAL EDGE SPEED (nSEC) RISE/FALL TIME	TRANSITION ELECTRICAL LENGTH IN FR-4 (INCHES)	TRANSITION ELECTRICAL LENGTH IN FR-4 (CM)	CRITICAL LENGTH (INCHES)	CRITICAL LENGTH (CM)
STANDARD TTL	5	30	76.20	15.00	38.10
FTTL	1.2	7.2	18.29	3.60	9.14
10 K ECL	2.5	15	38.10	7.50	19.05
BTL	0.5	3	7.62	1.50	3.81
LVDS	0.3	1.8	4.57	0.90	2.29
100K ECL	0.5	3	7.62	1.50	3.81
GaAs 106	0.3	1.8	4.57	0.90	2.29
GLT+ (PentPro)	0.3	1.8	4.57	0.90	2.29
0C-48	0.2	1.2	3.05	0.60	1.52
OC-192	0.05	0.3	0.76	0.15	0.38

Table 10.1. Several Logic Families, Their Rise Time and the Length in a PCB

Lines as long as or longer than the **Transition Electrical Length** (TEL) are likely to fail from reflections unless impedance control is used.

Lines as long as or longer than the **Critical Length** are likely to fail from crosstalk unless care is taken with trace-to-trace spacing.

CHAPTER 11: INDUCTANCE

Inductance is a property of a conductor that "impedes" the flow of electromagnetic energy along it. Inductance is referred to as a reactive component because its effect is only visible when an electromagnetic field traveling through it is changing. It "reacts" to this changing field by "impeding" its flow. The changing electromagnetic field causes a change in the rate at which current flows through an inductor resulting in a voltage drop across the inductance.

Equation 11.1 illustrates the relationship between the rate of change in current, the size of the inductance and the voltage drop that develops across it. (Remember, the current flow is changing because the electromagnetic field is changing.) Notice that there is no voltage drop across an inductance if the current flow through it or the electromagnetic field traveling on it is constant. This is the reason that parasitic inductances in circuits can be ignored at slow speeds.

$$V_L = L\frac{di}{dt}$$

Where V_L = the voltage, in volts, across the inductor, L = the inductor value in Henrys, di = the change in current in Amperes, dt = the change in time of the di in seconds

Equation 11.1. Voltage Drop Across an Inductance

Inductance is measured in Henrys. Named after Joseph Henry, an early researcher into electromagnetics, a Henry is the inductance for which the induced voltage in volts is numerically equal to the rate of change in current in amperes per second. All conductive structures-- wires, lead frames, resistors, capacitors standoffs, etc.--have some parasitic inductance. A nanohenry, (10^{-9} henry, abbreviated nH) is the relative size of the inductances that appear as parasitic elements in component leads and transmission lines. The area of electronic circuits where these parasitics have the most adverse effect is IC power leads, capacitors, planes or anywhere in the power path.

Table 11.1 lists the parasitic inductance of some common electrical components.

1" of 50 ohm transmission line-	10 nanohenrys
1 cm of 50 ohm transmission line-	4 nanohenrys
Power lead of 20 pin DIP-	13 nanohenrys
Power lead of a 208 pin QFP-	9 nanohenrys
1206 chip capacitor-	1.5 nanohenrys
0805 chip capacitor-	0.85 nanohenrys
0603 chip capacitor-	0.63 nanohenrys
0402 chip capacitor-	0.35 nanohenrys
Mounting for 1206 capacitor-	2 nanohenrys
100 mil tall via-	1 nanohenry
Power lead of a 249 pin BGA-	0.2 nanohenrys

Table 11.1. Parasitic Inductance of Some Typical Components

If a sine wave is impressed across an inductor, it will exhibit an impedance to current flow through it that can be calculated using Equation 11.2.

$$X_L = 2\pi fL$$

where X_L = inductive reactance in ohms, f = frequency in hertz, L = inductance in Henrys

Equation 11.2. Inductive Reactance Equation

Note that when the frequency, f, is zero (DC), the inductive reactance of an inductor of any size is zero. Said another way, at DC, any parasitic inductance that exists in the path of an electromagnetic field has no effect and can be ignored. That is precisely how the behavior of battery circuits and low frequency circuits is analyzed. The same is true of TTL circuits. Table 11.2 lists the reactance, in ohms, of the items in Table 11.1 at a variety of frequencies.

All real components with lengths greater than zero have parasitic inductance.

	Inductance (nH)	10 KHz	100 KHz	1 MHz	10 MHz	100 MHz	1 GHz	10 GHz
249 pin BGA	0.2	0.00001256	0.0001256	0.001256	0.01256	0.1256	1.256	12.56
0402 Chip Capacitor	0.35	0.00002198	0.0002198	0.002198	0.02198	0.2198	2.198	21.98
0603 Chip Capacitor	0.63	0.000039564	0.00039564	0.0039564	0.039564	0.39564	3.9564	39.564
0805 Chip Capacitor	0.85	0.00005338	0.0005338	0.005338	0.05338	0.5338	5.338	53.38
1206 Chip Capacitor	1.5	0.0000942	0.000942	0.00942	0.0942	0.942	9.42	94.2
1 cm 50 ohm line	4	0.0002512	0.002512	0.02512	0.2512	2.512	25.12	251.2
208 Pin QFP	9	0.0005652	0.005652	0.05652	0.5652	5.652	56.52	565.2
1 inch 50 ohm line	10	0.000628	0.00628	0.0628	0.628	6.28	62.8	628
20 pin DIP	13	0.0008164	0.008164	0.08164	0.8164	8.164	81.64	816.4

Table 11.2. Impedance, in Ohms, of the Parasitic Inductances of Table 11.1 Components at Several Frequencies

Notice that the inductive reactances of what seem to be very small lead inductances become quite large impedances at the high frequencies involved in microwave and very fast logic circuits. Imagine what happens when a component that is expected to operate at 2.5 GB/S is packaged in a 208 pin QFP. The current required to allow the part to operate would have to pass through the inductance shown in this table. As we will see later, parts in the wrong package (with power lead inductances that are too high) are a common cause of circuit malfunctions.

The above discussion has focused on the unwanted parasitic element of inductance. Inductance is used in beneficial ways as well. Energy can be stored in the magnetic field that surrounds an inductor. This is one of the two elements in a switching power supply in which energy is stored as part of converting one DC voltage to another. Inductors can also be used as part of filter circuits that selectively block certain frequencies or as loads of RF amplifiers to selectively amplify certain frequencies. When inductors are used this way, they are typically illustrated in circuit diagrams with the symbol in Figure 11.1.

Ideal Inductor

Figure 11.1. Typical Ideal Inductor Symbol

The symbol in Figure 11.1 suggests that an inductor only contains inductance. Ideal inductors do look this way. However, there are no ideal inductors. All real inductors also have parasitic resistance from the metal used to make them and parasitic capacitance between their two terminals. Figure 11.2 is a more realistic representation of an inductor. Later in the book, the effects of these parasitic elements-- resistance and capacitance--will be illustrated.

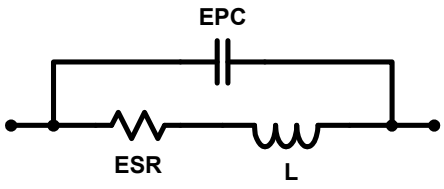

Real Inductor

ESR = equivalent series resistance, EPC = equivalent parallel capacitance

Figure 11.2. Typical Real Inductor

If one carefully examines the circuit in Figure 11.2, it can be seen as a parallel-tuned circuit. In fact, at some frequencies the inductor behaves this way, rather than as intended. When this happens, the circuit in which it is placed doesn't perform as expected. This circumstance is one of many that often leads engineers not expecting this type of behavior to suspect that some form of black magic is at work.

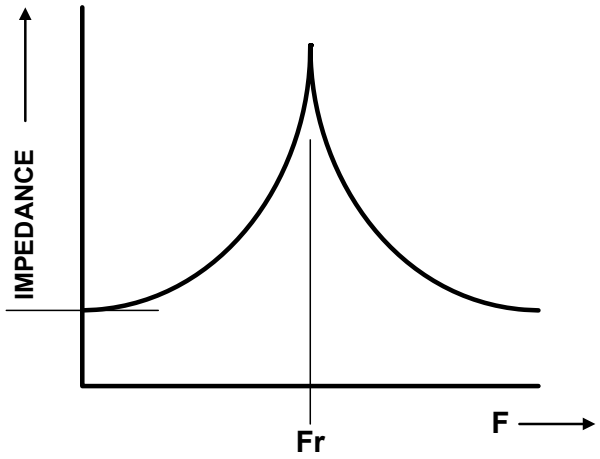

Where: Fr = Self Resonant Frequency of the Tuned Circuit, F = Frequency

Figure 11.3. Impedance of a Parallel-Tuned Circuit vs. Frequency

Figure 11.3 shows how the impedance across a parallel-tuned circuit (the inductor with its parasitic components) varies with frequency. At low frequencies, the inductor presents a very low impedance (the ESR or equivalent series resistance) at the left side of the chart. As the frequency increases, the inductor exhibits an ever larger impedance until it reaches the frequency where the reactance of the inductor exactly equals the reactance of the parasitic capacitance, or the self resonant frequency, Fr. At this frequency, the inductor's impedance is at its highest. Beyond Fr, the impedance drops off due to the parasitic capacitance.

Because of the parasitic components in the inductor, its value as an inductor producing high impedances is limited to a narrow range of frequencies around Fr.

CHAPTER 12: CAPACITANCE

Capacitance is the property of a collection of conductive structures that causes electrical charge to be stored on them when a voltage difference has been impressed between them. The basic unit of capacitance is a Farad (named after Michael Faraday, an early researcher into electromagnetic effects). A farad is the capacitance of a capacitor in which a charge of one coulomb produces a one-volt potential difference between its terminals. The capacitance that exists between transmission lines or wires and the surrounding structures (in PCBs, this is primarily the planes), must be charged up in order to change the voltage on the transmission line from a logic zero to a logic one. Energy in the form of an electromagnetic wave is used to do this. This "parasitic" capacitance impedes the changing of the signal voltage of a transmission line. Depositing charge onto signal lines to change their logic states or voltages is the primary power consumer in switching circuits.

Equation 12.1 illustrates the relationship between the charge stored on a capacitor and its terminal voltage.

$$V_c = \frac{1}{C}\int IdT$$

Where V_c = the voltage, in volts, across the capacitor, C = the capacitance value in farads, I = the current flowing into the capacitor in amperes, dT = the duration of the current flow in seconds.

Equation 12.1. Voltage Drop Across a Capacitor

Note that over time the current flow raises the voltage across the capacitor. The larger the capacitor, the more current flow over time or charge that will be required to raise its terminal voltage to a given value. Stated another way, if the parasitic capacitance of a circuit is large, it will take more power to perform a given logic operation. Indeed, the primary way to increase the speed of logic circuits for a given power consumption has been to reduce the parasitic capacitance of each circuit element. This has been made possible by increasingly smaller IC transistors and wires.

A picofarad, (10^{-12} farads, abbreviated pF) is the relative size of the parasitic elements that exist on real transmission lines. As already mentioned, the place where this has the most effect is on transmission lines that must be charged and discharged to change logic states. Table 12.1 lists the parasitic capacitance of some common electrical components.

Input to a memory IC-	5.0 pF
1" of 50 ohm transmission line-	3.5 pF
1 cm 50 ohm transmission line-	1.4 pF
Active Probe-	1.0 pF
100 mil tall, 13 mil dia. via-	0.3 pF

Table 12.1. Parasitic Capacitance of Some Typical Components

If a sine wave is impressed across a capacitor it will exhibit an impedance to the flow of current through it that can be calculated using Equation 12.2.

$$X_C = \frac{1}{2\pi fC}$$

Where Xc = capacitive reactance in ohms, f = frequency in hertz, C = capacitance in farads

Equation 12.2. Impedance of a Capacitor at a Given Frequency

Notice that when the frequency, f, is zero (DC), the capacitive reactance is infinite. Said another way, at DC or low frequencies, any parasitic capacitance that exists in the path of an electromagnetic field has no detectable effect and can be ignored. Table 12.2 lists the reactance, in ohms, of the items in Table 12.1 at varying frequencies.

	Capacitance (pF)	10 KHz	100 KHz	1 MHz	10 MHz	100 MHz	1 GHz	10 GHz
100 mil tall via	**0.3**	**53,078,556.26**	**5,307,855.63**	**530,785.56**	**53,078.56**	**5,307.86**	**530.79**	**53.08**
Active Probe	**1.0**	**15,923,569.45**	**1,592,356.94**	**159,235.69**	**15,923.56**	**1,592.35**	**159.23**	**15.92**
1 cm 50 ohm line	**1.4**	**11,373,976.34**	**1,137,397.63**	**113,739.76**	**11,373.98**	**1,137.40**	**113.74**	**11.37**
1 inch 50 ohm line	**3.5**	**4,549,590.54**	**454,959.05**	**45,495.91**	**4,549.59**	**454.96**	**45.50**	**4.55**
Memory IC input	**5**	**3,184,713.38**	**318,471.34**	**31,847.13**	**3,184.71**	**318.47**	**31.85**	**3.18**

Table 12.2. Impedance, in Ohms, of the Parasitic Capacitances in Table 12.1 Components at Several Frequencies

Notice that the impedance of all of these parasitic elements is extremely high at low frequencies. As a result, they have little or no effect on performance--so long as the frequencies involved are low. As frequencies approach a GHz, parasitic capacitance from structures such as plated through holes, oscilloscope probes and the input capacitance of each load can have a noticeable effect on signal quality. Later in this book, it will be shown that frequencies involved in very fast switching edges are in the GHz range. One can see that some of the parasitic capacitance elements that normally occur in signal paths have detectably low impedance and can have an adverse effect on signals.

As with inductance, capacitance is used in beneficial ways. Two of the ways capacitance can be used beneficially are as charge storage devices, usually called decoupling capacitors or bypass capacitors, and as coupling capacitors to pass an AC signal while blocking a DC signal. Another beneficial use of capacitance is as part of a switching power supply to store charge when converting one DC voltage to another. When capacitors are used in this manner, they are typically illustrated in circuit diagrams with the symbol shown in Figure 12.1.

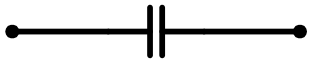

Ideal Capacitor

Figure 12.1. An Ideal Capacitor

The symbol in Figure 12.1 suggests that capacitors have no unwanted parasitic elements. In fact, ideal capacitors don't. However, there are no ideal capacitors. Real capacitors look like the circuit shown in Figure 12.2.

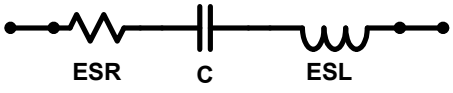

Real Capacitor

Where ESR is equivalent series resistance and ESL is equivalent series inductance

Figure 12.2. Real Capacitor

If one examines the circuit in Figure 12.2, it can be seen as a series resonant circuit. Such circuits behave quite differently than might be expected. Figure 12.3 shows how the impedance across the terminals of a real capacitor behaves with frequency. Instead of the impedance getting increasingly lower as the frequency increases, there is an impedance low when the reactance of the capacitor is exactly equal to the reactance of the inductor. At the frequency Fr (the self-resonant frequency of the network), the two reactive components cancel each other out and the impedance that is seen is the ESR (Rc) of the capacitor. Beyond that frequency, the impedance goes back up due to the reactance of the parasitic inductance and the capacitor begins to behave like an inductor rather than a capacitor. The adverse effect of this behavior will be discussed in detail in a later section of this book that addresses bypass capacitors and power subsystem design.

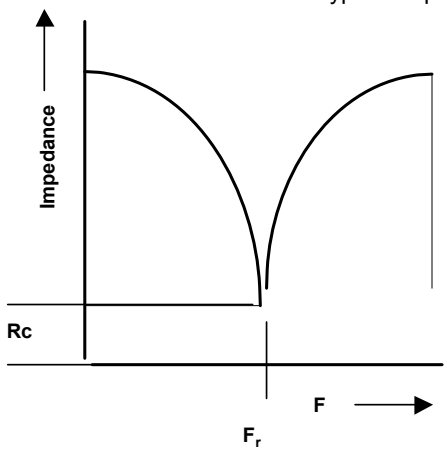

Where Rc = ESR of capacitor, Fr- resonant frequency of LC combination

Figure 12.3. Impedance of a Series Resonant Circuit (capacitor) as a Function of Frequency

CHAPTER 13: RESISTANCE

Resistance is a property of all real conductors. Like inductance, resistance "impedes" the flow of current in a conductor. However, unlike an inductor it is not frequency sensitive. It impedes current flow at all frequencies in the same amount. Also, unlike inductance, energy cannot be stored in a resistance. Instead, as current flows through a resistance, energy from the electromagnetic field is absorbed or dissipated and is lost in the form of heat. It is this property of conductors that is exploited in a beneficial way to create heaters. In the transmission lines and components of electronic circuits, the energy loss in the parasitic resistance is not beneficial. Rather, the signal amplitude is diminished by it.

Resistance is measured in ohms. An ohm (named after George Simon Ohm, a German mathematician who experimented with electricity in the early 19th century) is the quantity of resistance that develops a voltage drop of one volt across it when a current of one ampere is flowing through it. Equation 13.1 is the famous "Ohm's Law" used to calculate voltage drop.

$$E = IR$$

Where E = voltage in volts, I = current in amperes, R = resistance in ohms

Equation 13.1. Ohm's Law

Table 13.1 lists the parasitic resistance of some common electronic components.

Power lead of a 20 pin DIP-	5 milliohms
1" 50 ohm trace, 5 mils wide, 1.4 mils thick-	66 milliohms
15" x 18" 1 ounce power plane-	1.5 milliohms
, 50 watt light bulb-	200 ohms

Table 13.1. Parasitic Resistance of Some Typical Components

As with inductance and capacitance, resistance is used in beneficial ways, such as for terminations, voltage dividers, current-limiting devices and heaters. When resistance is used in this manner, it is usually represented with the symbol shown in Figure 13.1.

Ideal Resistor

Figure 13.1. An Ideal Resistor

The symbol in Figure 13.1 suggests that resistors have no unwanted parasitic elements. In fact, ideal resistors don't. As with the other two components, there are no ideal resistors. Real resistors look like that shown in Figure 13.2.

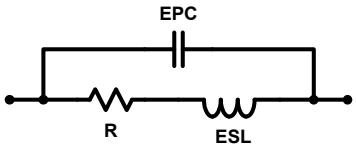

Where: EPC = equivalent parallel capacitance, ESL = equivalent series inductance, R = resistance

Figure 13.2. A Real Resistor

As with a real inductor, the parasitic components (in this case inductance and capacitance) result in a parallel-tuned circuit. At very high frequencies, the parasitic inductance has a noticeable effect on the behavior of the resistor. In logic circuits, parasitic inductance can be an issue when resistors are used as terminations for gigabit and higher transmission lines.

Skin Effect Loss

The discussion above asserts that parasitic resistance is not affected by frequency. This is not entirely true. At very high frequencies, the current flow in a conductor ceases to flow uniformly through the whole cross section of a conductor. Instead, it crowds near the surface. This phenomenon is called **skin effect loss**. When the current is forced to flow only in the surface layers of a conductor, it appears that the parasitic resistance of the conductor has increased. It seems that the conductor is hollow. In fact, it might as well be since the high frequency currents are only flowing in the outer layers of the conductor.

In logic circuits, calculating skin effect losses is a relatively complex piece of analysis. This is due to logic signals being made up of a wide bandwidth of frequencies each affected differently by skin effect loss. Successfully accounting for skin effect loss can only be done with analytical tools that take into account the geometry of the transmission line, its length and the types of signals being transmitted. Howard Johnson, etal, deals with this topic in great detail in the book, “High Speed Signal Propagation.”

CHAPTER 14: FUNDAMENTALS OF TRANSMISSION LINES

The previous three chapters have shown that real transmission lines have parasitic elements--resistance, capacitance and inductance--that effect how they behave. It is possible to represent a transmission line with a schematic like that shown in Figure 14.1. In fact, this is the usual way in which transmission lines are analyzed.

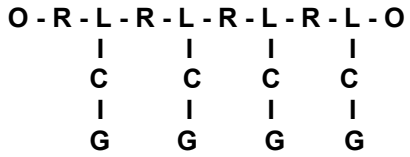

O = END OF LINE
G = GROUND/PARTNER PLANE
C = CAPACITANCE PER UNIT LENGTH
R = RESISTANCE PER UNIT LENGTH
L = INDUCTANCE PER UNIT LENGTH

$$X_L = 2\pi fL \qquad X_C = \frac{1}{2\pi fC}$$

Reactance of an Inductor — Reactance of a Capacitor

Model assumes ground is a plane of negligible inductance and resistance.

Figure 14.1. A Schematic Representation of a Transmission Line

This schematic representation of a transmission line is said to be “single ended”. This means that the transmission line travels over a perfect plane as represented by the symbol G. In fact, there are no perfect planes. All planes have parasitic inductance and resistance. However, these plane parasitics are so small compared to the parasitic inductance and resistance of the transmission line (as much as three orders of magnitude less) that this is a reasonable model.

The model is characterized on a per unit length basis. As long as the transmission line is uniform in cross section along its length so that the values of the parasitics are the same per unit length, this is an accurate model. Once an inch or a meter of the line has been analyzed, it is possible to state what the behavior of the entire transmission line will be, no matter how long it is.

Recalling Equation 11.2, the impedance of an inductance as a function of frequency and Equation 12.2, the impedance of a capacitance as a function of frequency, it can be seen that at DC or very low frequencies, the effect of the parasitic inductance and capacitance is nil. All that an electromagnetic wave seems to see that would impede the flow of energy down the transmission line is the parasitic resistance of the conductors. Indeed, when designing battery circuits and other power distribution circuits where the frequencies involved are low, this is the only parameter that is used. The parasitic inductance and capacitance are ignored.

When the frequency of the electromagnetic energy increases, the impedance of the parasitic inductance begins to increase, tending to impede the flow of energy down the transmission line. At the same time, the impedance of the parasitic capacitance begins to decrease, tending to “short” the energy to the plane. When the frequency of the electromagnetic energy approaches a kilohertz, these two reactive components work in such a way that they dominate the behavior of the transmission line. Together they present an impedance to the EM field that is much larger than the parasitic resistance. They create an “impedance” that is constant with frequency.

There are many ways to represent a transmission line in a schematic. Figure 14.2 shows some of them. In most of the symbols of Figure 14.2, there are one or more “ground” symbols denoting that one side of the transmission line is ground or tied to ground. This is a bookkeeping convenience when drawing transmission lines. In real circuits, the plane over which a transmission line travels may well be the Vdd plane. In the context of transmission lines, when the word ground is used, it means AC ground, not DC ground. The Vdd plane may well be AC ground. In fact, if the power system is properly designed, it will be AC ground.

From this point forward, when the ground symbol is used with a diagram containing transmission lines, it does not necessarily mean the DC ground in the system but, rather, the reference half of the transmission line. The concept of ground will be explored in more detail in the following chapter.

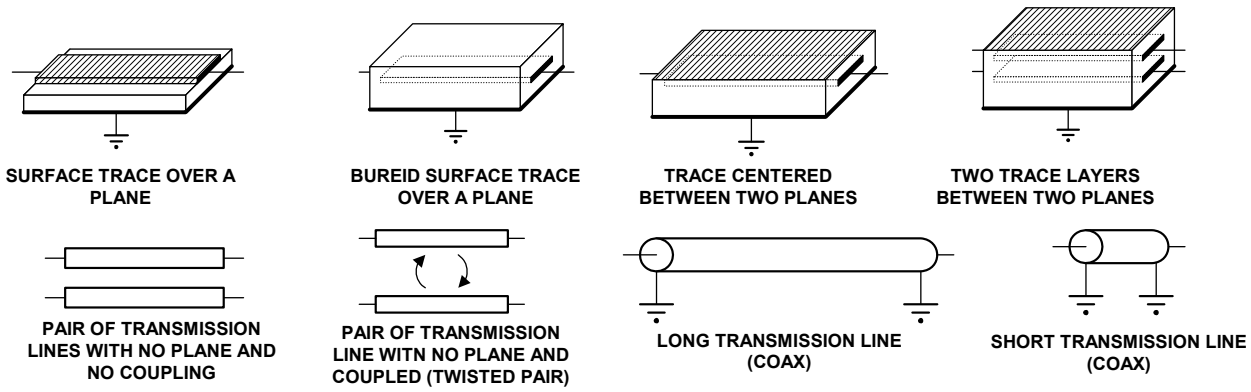

Figure 14.2. Ways to Draw Transmission Lines

The top row of drawings in Figure 14.2 depicts a way to draw transmission lines as they appear in PCBs. The left two are often referred to as **microstrip** transmission lines. Microstrip means that the transmission line is traveling over a plane or has only one plane as a partner. The line could be on the surface or buried. The right two drawings of the top row are referred to as **stripline** transmission lines. Stripline means that the transmission line travels between two planes or has two planes as partners. This transmission line could be centered between the two planes or placed off center.

The bottom row of Figure 14.2 shows another representation of transmission lines that is easier to draw than the top row. The left two drawings in the lower row symbolize a paired transmission line, such as a TV twin lead or unshielded twisted pair (UTP), and is the type of line over which most Internet and telephone signals are sent in buildings. If the two wires are close to each other, they will couple as shown in the second drawing.

The right two drawings in the lower row are a common way to symbolize a transmission line that has a partner such as a plane or the outer shield of a coaxial cable. The ground symbol represents the “AC” neutral portion of the transmission line system. As mentioned before, the appearance of the ground symbol in such a drawing does not mean that this is a connection to logic ground, chassis ground or RF ground.

CHAPTER 15: THE CONCEPT OF GROUND AND POWER PLANES

Before going further into transmission line theory, it is worth spending some time exploring the concept of ground. There is a great deal of confusion as to what ground is and what it can do with respect to transmission lines, EMI and high-speed signals. The term is used to describe a wide variety of things in electronics. For example, the green wire that comes out of the AC wall plugs is called ground, the chassis of a product is called ground, the reference terminal of an analog circuit is called ground, the negative terminal of a power supply is called ground and the plane half of a transmission line is called ground. This can be pretty confusing.

How can all of these be ground? Of course, they are not all ground. The term ground is being misused in all these cases. **By definition, ground is the one place in a system that has been identified as the point from which voltage measurements are made.** In order to avoid confusion, all of the other items listed in the above paragraph need other names.

CHASSIS GROUND

When the term "chassis ground" is used, it usually refers to the case of a product connected to the green wire coming from the AC outlet that powers the product. It is common to represent this connection with the symbol in Figure 15.1.

Figure 15.1. The Symbol Used to Represent "Chassis Ground"

The reason for the "green wire" is safety. The green wire pin on an AC connector is wired through the walls of the building back to a conductor that is literally driven into the ground or earth. The purpose of this connection to the "chassis" or case of a product is to protect the user from electrical shock in the event that the hot leads of the AC power connection should inadvertently make contact with the case. This is the only function of this connection. It plays no role in the containment of EMI. To suggest that it does or can presumes that this connection is a low impedance at the frequencies of concern for EMI (30 MHz to 1 GHz). Close examination of this path reveals that it contains thousands of nanohenrys of parasitic inductance. A quick calculation of the inductive reactance of such a parasitic inductance reveals that this is a virtual open circuit, even at the lowest frequencies of concern.

The term "chassis ground" is often used in discussions regarding EMI. Usually the reference is being made to the case surrounding a product. This case is expected to be a containment vessel for the EMI that has escaped from the circuits inside the product. Such a containment vessel is more properly called a **Faraday Cage**. A connection between the Faraday cage and "chassis ground" is not necessary for it to do its job as a containment vessel. If there is any doubt about this, consider a cell phone. It has a Faraday Cage surrounding the phone circuits, passes CISPRB EMI standards and has no connection to the green wire at all.

DIGITAL GROUND

Digital ground is one of the terminals of the power supply that provides power to the logic circuits in a product. In most cases this is the negative terminal of the power supply. It could be the positive terminal as with ECL logic. It is common practice to represent logic ground with the symbol in Figure 15.2.

Figure 15.2. The Symbol Used to Represent "Logic Ground"

One of the terminals of a logic power supply is named ground to provide a reference from which to measure logic levels of the logic signals. This is the only reason to declare one of the terminals "ground". It is common practice to make a connection between logic ground and the green wire ground of a product; however, it is not necessary to do so. There are cases where it is not allowed, such as a product that has an RS-232 interface.

When EMI is being considered, it is desirable to make a connection between logic ground and the Faraday cage, so that unshielded signal wires exiting the product don't vary with respect to the Faraday cage and result in possible emissions.

ANALOG GROUND

Analog ground is a power supply terminal for the analog circuits in a product and is used to reference analog signals. It can either be the positive or negative terminal. Usually, analog circuits are coupled to digital circuits as with an A/D

(analog to digital) converter. In this case, both circuits share the same ground. Analog ground is often represented with the symbol shown in Figure 15.3.

Figure 15.3. The Symbol Used to Represent Analog Ground

Sometimes there is confusion as to the need for a separate analog ground and digital ground when a mixed technology part, such as an A/D converter or read channel, is being used. This is not necessary, nor is it advisable. Once the reason for having two different terminals is understood, it can be seen that two different grounds are not needed. The handling of mixed analog and digital designs will be covered in Volume 2.

FARADAY CAGES

When the discussion of EMI containment takes place, it is common to hear the word chassis ground used. The chassis isn't the component that is containing EMI. The containment vessel is a metallic container surrounding the circuit that is emitting energy that could cause EMI. This containment vessel is a Faraday Cage. Often, parts of the chassis, such as the side panels of the card cage, comprise part of the Faraday Cage. Perhaps this is the reason that the word chassis is incorrectly used to describe the EMI containment vessel. In actuality, in a typical card cage oriented product, the Faraday Cage is comprised of the sides of the card cage, the face plates of the plug-in modules, the planes in the backplane, and some sort of EMI tight screens on the top and bottom. How well this six -sided structure contains EMI is heavily dependent upon how well the six surfaces are bonded to each other and how many signal wires breach this enclosure. This topic will be discussed in detail in Volume 2.

For convenience, it would be good to have a symbol to use when labeling points in a design as being connected to the Faraday Cage potential. There isn't such a symbol in common use. For purposes of discussing Faraday Cage potentials the symbol in Figure 15.4 will be used in this book. It might be good for anyone discussing EMI containment to adopt a similar symbol to allow differentiation between this potential and all of the other "grounds" in a system.

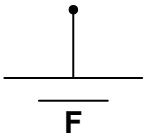

Figure 15.4. A Symbol for Representing Faraday Cage Connections

POWER PLANES

In this book, when discussing transmission line management, the planes of a PCB are referred to as power planes rather than ground and power planes. The context in which this name is used is as a partner for a transmission line. At the speeds involved in high-speed signals, all of the planes in a PCB are shorted together and can be used interchangeably as partners or "ground" planes for these signals. Why is this so? In the process of engineering the power subsystem, a low impedance voltage source is created. This is created through the combination of using bypass capacitors with different values and the plane capacitance that exists between the power and ground planes to create a very low impedance (typically, less than 20 milliohms) between the pairs of planes across the entire range of frequencies involved in the switching signals. This amounts to a short circuit when compared to the 50 ohms of a typical transmission line.

Because of the "short circuit" between power and ground planes, it is reasonable to consider them equivalent when routing high-speed signals. It is not necessary to use only the layers called "ground" for high-speed signals. It is true that the ripple on the Vcc or Vdd plane will couple onto signals that route over these planes. In order to make sure this ripple noise is within acceptable limits, appropriate power supply engineering is necessary. This engineering will be covered in a later chapter.

> For purposes of routing high speed signals, Vcc or Vdd and ground planes can be used interchangeably. It is not necessary to route high-speed signals only over the DC ground layers.
>
> Proper power supply engineering makes this possible.

CHAPTER 16: IMPEDANCE

Impedance is a major property of transmission lines. Impedance is the resistance that a transmission line presents to the flow of energy along the transmission line. It is composed of the three parasitic elements shown in Figure 14.1--resistance, capacitance and inductance. (Conductance has been omitted from this discussion because it has little effect on transmission line performance until signal rates exceed 1 GB/S). Looking at this figure, it can be seen that at DC or low frequencies, the parasitic resistance is the primary determining factor in transmission line impedance. When designs intended to handle DC voltages are created only parasitic resistance is used to calculate current flow. For these designs, both the parasitic inductance and capacitance can be ignored because their effect is not noticeable.

As the frequency gets higher than a few kilohertz, the reactance of the parasitic inductance tends to block or impede the flow of energy. At the same time, the parasitic capacitance tends to shunt the energy to "ground" or the plane. These two elements work together in such a way that a constant impedance is seen by the electromagnetic field at all frequencies. Equation 16.1 can be used to calculate the impedance of any transmission line. All that is required is a method for determining Ro, Co and Lo. (Note: Equation 16.1 doesn't exactly reflect how Ro behaves, but it is accurate enough to portray the idea.) Chapter 24 covers this in detail.

Where Zo is impedance in ohms, Lo is parasitic inductance in henrys per unit length, Co is parasitic capacitance in farads per unit length, Ro is parasitic resistance in ohms per unit length and Go is transconductance per unit length in mhos.

$$\mathbf{Z}_0 = \sqrt{\frac{\mathbf{R}_0 + \mathbf{j\omega L}_0}{\mathbf{G}_0 + \mathbf{j\omega C}_0}}$$

Equation 16.1. The "Impedance" Equation

Notice the equation states that impedance is not a function of the length of a transmission line. As long as the cross section remains the same along the length of the line, the impedance will be the same whether the line is one cm or one kilometer long. Also, notice that the impedance is not a function of frequency. The impedance will be the same whether the frequency of the EM energy is 1 MHz or 1 GHz. This holds true so long as none of the three variables changes with frequency.

Ro is the parasitic resistance of the conductors in the transmission line. For the dimensions of traces in a printed circuit board, this parasitic is expressed in milliohms per inch or centimeter. The impedance produced by Co and Lo will be in the range of 50 ohms for most PCB traces. Figure 16.1 is a plot of trace resistance versus trace width and thickness.

The impedance of a transmission line is not a function of its length.

As long as the cross section does not change, the impedance will be the same whether the line is 1 centimeter long or 1 kilometer long.

The impedance of a transmission line is not a function of frequency.

As long as the dielectric medium has a dielectric constant that is constant with frequency, the impedance will be constant, no matter what the frequency.

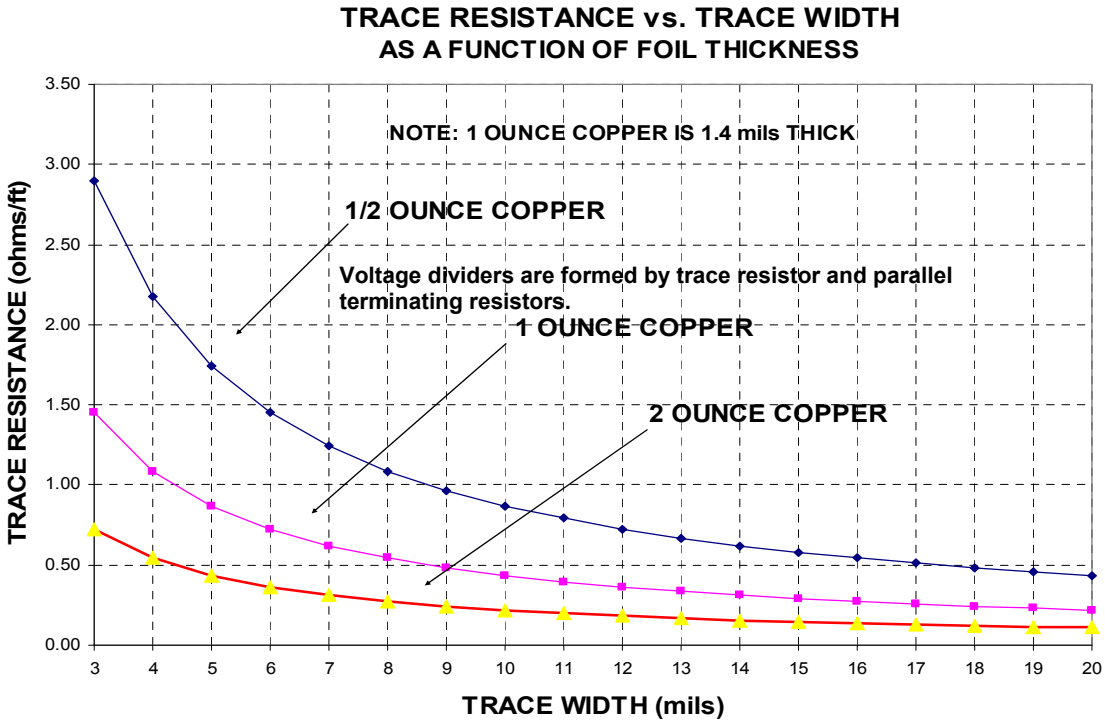

Figure 16.1. Trace Resistance vs. Trace Width and Trace Thickness

For most transmission line design work in PCBs, the resistive component is small enough that it can be neglected. When Ro is ignored, Equation 16.2 can be used to calculate impedance. Equation 16.2 is often called "the impedance equation," and it is one that should be committed to memory with the same rigor as Ohm's law.

$$Z_O = \sqrt{\frac{L_O}{C_O}}$$

Where Zo is impedance in ohms, Lo is parasitic inductance in henrys per unit length, Co is parasitic capacitance in farads per unit length

Equation 16.2. "The Simplified Impedance Equation"

However, it must be stated that the phenomena called skin effect loss can make the contribution from resistance significant at very high frequencies with long paths. How skin effect loss affects signals will be discussed and demonstrated in Volume 2.

With most of the dielectrics available for PCB fabrication, the dielectric constant changes with frequency. This will cause Co as well as the impedance to change. Figure 16.2 shows how the relative dielectric constant of several PCB laminates varies with frequency.

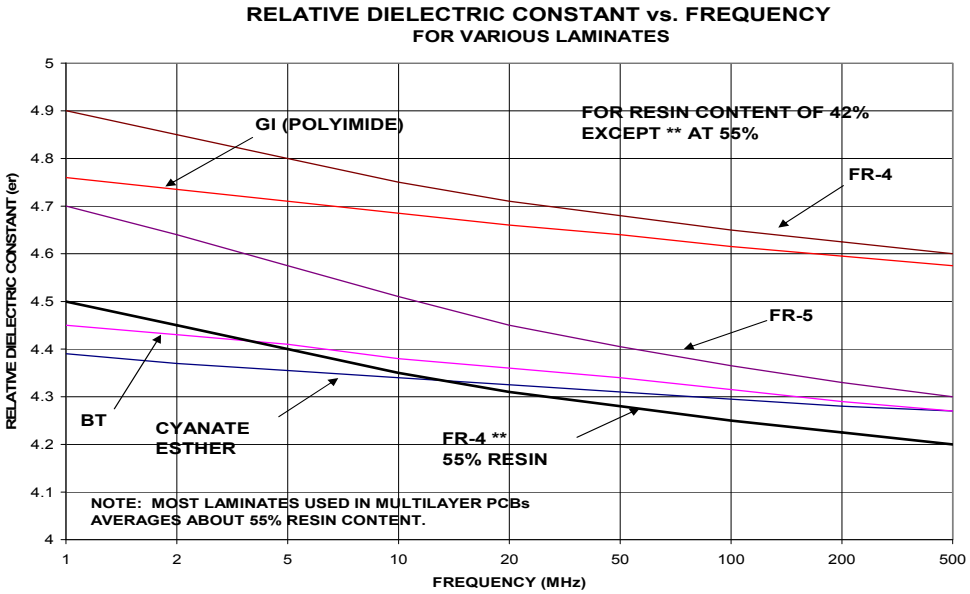

Figure 16.2. Relative Dielectric Constant vs. Frequency for Several PCB Laminates

Table 16.1 lists a variety of PCB laminate materials with their relative dielectric constants given for a resin content of 55% measured at the frequency of a 125 picosecond TDR (Time Domain Reflectometer) edge rate. The table also shows the velocity of travel for EM energy in each dielectric. Notice that all of the commonly used PCB laminates have approximately the same relative dielectric constant and wave velocity. To provide a comparison, the dielectric constants of air and water are also included. The reason that the e_r of the laminates commonly used to make multilayer PCBs--FR-4, Cyanate Ester (CE), Getek, Polyimide, and BT--are similar is that they all have glass as a reinforcement. Note that the e_r of glass is 6.

MATERIAL	e_r	VELOCITY (in/nSEC)	VELOCITY (pSEC/in)
AIR	1.0	11.76	84.9
PTFE/GLASS	2.2	7.95	125.8
ROGERS RO 2800	2.9	6.95	143.9
CE/GOREPLY	3.3	5.97	167.0
GETEK	3.9	6.21	161.0
CE/GLASS	3.7	6.12	163.0
SILICON DIOXIDE	3.9	5.97	167.0
BT/GLASS	4.0	5.88	170.0
POLYIMIDE/GLASS	4.1	5.82	172.0
FR-4/GLASS	4.1	5.82	172.0
GLASS CLOTH	6.0	4.70	212.0
ALUMINA	9.0	3.90	256.0
WATER	73.0	0.4	2200.0

Table 16.1. Some Typical PCB Laminate Materials with Relative Dielectric Constants (55% Resin)

CHAPTER 17: REFLECTIONS--WHAT CAUSES THEM, WHAT THEY DO TO A SIGNAL

As an electromagnetic wave propagates along a uniform cross section transmission line it does so with little or no loss. (The losses are derived from the parasitic resistance of the conductor and losses in the dielectric, both of which are small.) If the impedance of the transmission line changes for any reason, some of the energy in the EM field is reflected back toward the source. The EM field that continues on is diminished by the amount of the energy reflected. The amount of reflection and its polarity can be calculated using Equation 17.1. This equation predicts the percentage of the incident EM field that will be reflected back to the source based on the two impedances on each side of a change. The equation predicts the voltage amplitude of the reflection.

$$\% = 100\ \frac{Z_I - Z_O}{Z_I + Z_O}$$

Where Z_I = the load impedance or down stream impedance, Zo = the output or upstream impedance

Equation 17.1. The Reflection Equation

Notice that when the load impedance or the downstream impedance at a change is higher than the upstream impedance, the polarity of the reflection is positive. This reflected voltage adds to the incident voltage. At the extreme, if the downstream impedance is infinity or the transmission line ends in an open circuit, all of the incident energy is reflected. None of it is absorbed, much like the mechanical transmission line example depicted in Chapter 7. In this case, the entire voltage amplitude and all of the energy in the EM field is reflected back to the source and adds to the incident voltage level. The effect of this is to double the voltage at the open circuit end of the line.

> The logic voltage doubles at the end of a transmission line that is an open circuit. This happens no matter what the line impedance is.

At the other extreme, when the impedance at a change is lower than the line or upstream impedance, the polarity of the reflection is negative. The reflected voltage subtracts from the incident voltage. At the extreme, if the downstream impedance is a short circuit or zero ohms, all of the energy in the EM field reflects back toward the source. None of the energy is absorbed. The reflected voltage is inverted and subtracts from the incident voltage, resulting in zero volts.

> A short circuit at the end of a transmission line reflects all of the energy in an EM field back toward the source inverted.

Figure 17.1 depicts a 50-ohm transmission line driven by a logic signal. It is terminated in an impedance that has three different values. The waveforms that appear at the output of the driver are plotted in Figure 17.2. First, a zero to one logic transition takes place with a terminating resistor of 50 ohms that is the same value as the line impedance This results in a signal that has no reflection (the center trace). The termination absorbs all of the energy in the EM field.

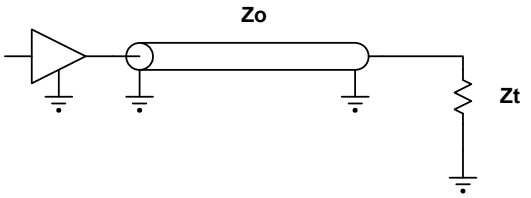

Figure 17.1. Driven Transmission Line Parallel Terminated

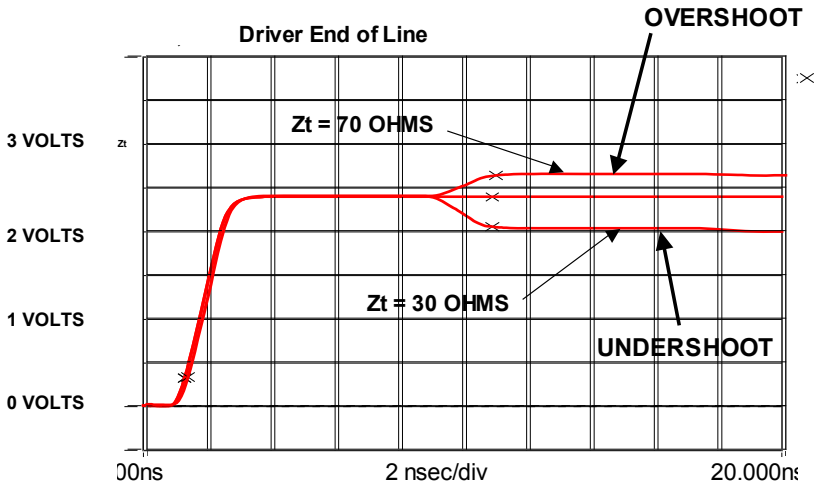

Figure 17.2. Driver Waveforms for a Parallel Terminated Transmission Line

When the value of the termination is set to 70 ohms (a downstream impedance higher than the upstream impedance), some of the energy in the EM field is reflected back to the source or driver. The polarity of the reflection is positive and adds to the incident voltage level. The name given to this "positive" or "adds to" reflection is **overshoot**. This kind of reflection does not degrade logic signal levels and is usually not harmful. It can be harmful if the sum of the incident voltage level and the reflected voltage exceeds the input voltage rating of a logic circuit. (This is the doubling that TTL engineers see at a load that mystifies them and results in the addition of AC terminations.)

When the value of the termination is set to 30 ohms (a downstream impedance that is lower than the upstream impedance), again, some of the energy in the EM field is reflected back to the source or driver. This time the polarity of the reflection is "negative" or "takes away" from the incident voltage level. The name given to this "negative" or "takes away" reflection is **undershoot**. For some reason, this kind of reflection is called "**ring back**" by some in our industry. The name ring back is not needed to describe this type of reflection. It is merely undershoot. This kind of reflection degrades the logic level and is always harmful. At any point in the design process where an engineer has any control over the way in which impedance mismatches occur, it is desirable to insure that downstream impedances are always higher than or equal to upstream impedances. One way to accomplish this is in choosing the values of termination resistors.

The example shown in Figure 17.1 is for a rising edge or a logic transition from a 0 to a 1. If the switching edge is falling or the logic transition is from a 1 to a 0, (the case that produces overshoot), with a termination resistor value higher than the line impedance, it results in a reflection that is negative going. This reflection adds to the incident logic level, just as it did in the previous case. Even though the reflection is negative going, it is still overshoot because it added to the incident logic signal. Similarly, when the termination resistor is smaller in value than the line impedance, the reflection takes away from the incident logic signal. Even though the reflection is positive going, it is still undershoot.

Reflections that add to the incident voltage level are **overshoot**, no matter if they are positive going or negative going.

Reflections that subtract from the incident voltage level are **undershoot**, no matter if they are positive going or negative going.

Now that the concept of reflections and what causes them has been explained, the behavior of a logic circuit can be illustrated and the resulting waveforms will make sense. Figure 17.2 is a typical 5V CMOS driver driving a 50-ohm transmission line that is 12 inches or approximately 2 nanoseconds long. There is a CMOS receiver at the end of the line that is a tiny parasitic capacitor. For these purposes, it looks like an open circuit.

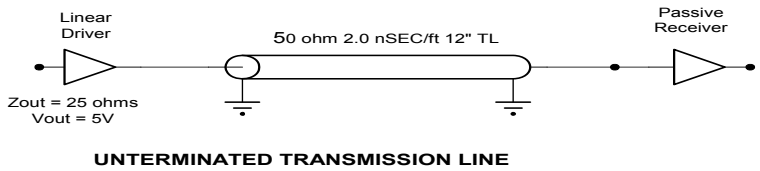

Figure 17.2. Typical Unterminated 5V CMOS Circuit

When the logic level in the circuit in Figure 17.2 switches from 0 to 1, the voltage waveforms shown in Figure 17.3 result. Notice that there are two waveforms separated in time by the electrical length of the transmission line. The first waveform in time is the one at the driver output. The second waveform in time is the one at the input to the load or receiver. Also, note that the two voltage waveforms are quite different. This is common to most transmission lines. The voltage waveform of interest and the one that must be engineered to meet the input conditions of the load, is the voltage waveform at the "load" end of the line. When conditions are right to create a good "input" waveform, the voltage waveform at the driver will likely seem to look bad. Last, remember that the voltage waveform is a manifestation of the magnitude of the electric field portion of the EM field. (The current flow is a manifestation of the magnitude of the magnetic field portion of the EM field.)

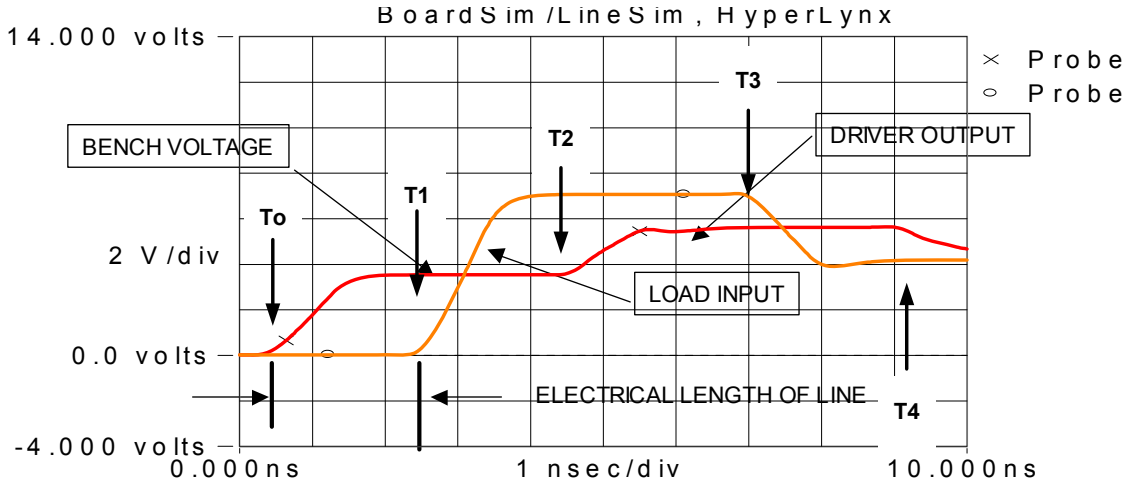

Figure 17.3. The Voltage Waveforms for an Unterminated 5V CMOS Circuit, 0 to 1 Transition

The following explanation will provide a walk through of the events that occur when the logic level switches from a 0 to a 1. The example being used is a 5V CMOS circuit. However, the same events would occur with any logic family. Understanding how a driver interacts with a transmission line and the loads along the way is vital to successful transmission line management.

At time, To, the logic circuit changes state internally from a 0 to a 1. At that moment, the equivalent circuit is shown in Figure 17.4.

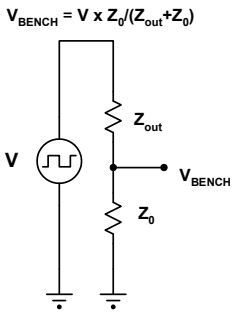

Figure 17.4. Equivalent Circuit of CMOS Circuit in Figure 30 at To

Notice that a voltage divider, made up of the output impedance of the driver and the line impedance, has formed. This divider causes the voltage level at the input to the transmission line to be less than the expected 5V CMOS logic 1. This level, in this case 3.3V, is the actual signal level that starts down the transmission line. It is created by the voltage divider made up of the 25-ohm output impedance of the driver and the 50-ohm impedance of the transmission line or 2/3 of 5 volts. The voltage level is a measure of the size of the EM field that is propagating down the line. The name for the

voltage level that starts down a transmission line is called the **"bench voltage"**. Also note that the load at the end of the transmission line is not visible to the driver. They are separated by 2 nanoseconds of time. Said another way, drivers can't see what the loads are at the other end of the transmission line. All the driver can see is the input to the transmission line.

> The driver cannot tell what is connected to the far end of a transmission line. It only sees the impedance of the transmission line itself. Once the reflected wave arrives at the driver, it is possible to calculate what the load was on the transmission line.

The EM field, whose voltage amplitude is 3.3V, travels down the transmission line and appears at the other end 2 nanoseconds later, T_1. The "load" is a tiny parasitic capacitor from the input to the logic circuit. This "parasitic" capacitor absorbs only a very small fraction of the energy in the EM field. For all practical purposes, the transmission line end is an open circuit. Because the transmission line is an open circuit, all of the energy in the EM field reflects back toward the source. The polarity of the reflection is positive and adds to the incident voltage, resulting in a "doubling". The amplitude of the voltage at the load at this moment is 2 times the bench voltage or 6.6V.

The energy in the EM field is now traveling back to the source. Two nanoseconds later, at T_2, it arrives back at the source and sees the equivalent circuit shown in Figure 17.5.

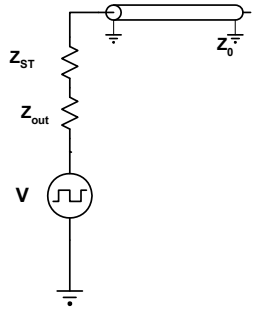

Figure 17.5. Equivalent Circuit Seen at T_3 By Reflected EM Field on Arrival Back at the Driver

If there is no series termination, Z_{ST}, as in this case, the impedance that is attached to the end of the transmission line is Zout of the driver or 25 ohms. The voltage source acts as a short circuit. Two things happen. First, the arriving signal is divided down by the ratio of the two impedances. Second, the remaining energy in the EM field is inverted by this high to low impedance change and reflected back down the transmission line toward the load end. Two nanoseconds later, at T_3, it arrives at the load end and appears inverted. As before, there is no load at this end of the transmission line, so all of the energy in the arriving EM field is reflected back toward the source, doubling as it does so.

In theory, this process repeats itself forever. In actuality, the signal becomes successively smaller on each round trip and finally dies out. The decrease in amplitude occurs because some of the energy is absorbed in the 25-ohm output impedance of the driver each time the EM field arrives back at the source. This diminishing waveform is often mistakenly called **"ringing"** because it often looks like a dampened sine wave. However, this is not ringing but is instead a series of reflections. Ringing is a phenomenon that occurs when a parallel resonant circuit made up of an inductor and a capacitor are excited by an impulse of energy, much like striking a bell.

Once the behavior of a transmission line driven by a logic state change is understood, it is possible to perform some diagnostic work by "reading" the waveforms that can be measured at the source and the load. For example, knowing that the line impedance is 50 ohms, the beginning voltage level is 5 volts and the bench voltage is 3.3 volts, it is possible to calculate the output impedance of the driver. Knowing the output impedance of the source, the beginning voltage level and the bench voltage it is possible to calculate the line impedance. This latter calculation is the one used to measure line impedance with all of the standard impedance measuring tools such as TDRs.

It is also possible to perform diagnostics on circuits that are not performing properly. Consider the clock circuit (Figure 17.6) and the resulting waveforms (Figure 17.7).

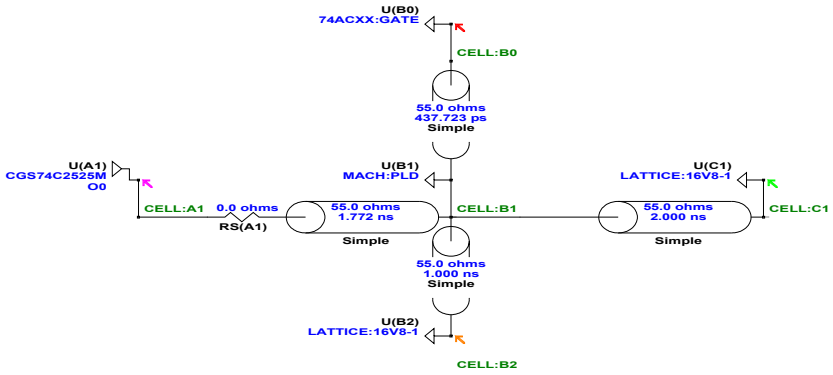

Figure 17.6. A Clock Circuit With Two Loads on Stubs

These waveforms are the four nodes in the circuit shown in Figure 17.5. This is a "clock" circuit with two of the loads on stubs off the main signal branch traveling from the source, A1, to the last load, C1. The two loads on stubs are B0 and B2.

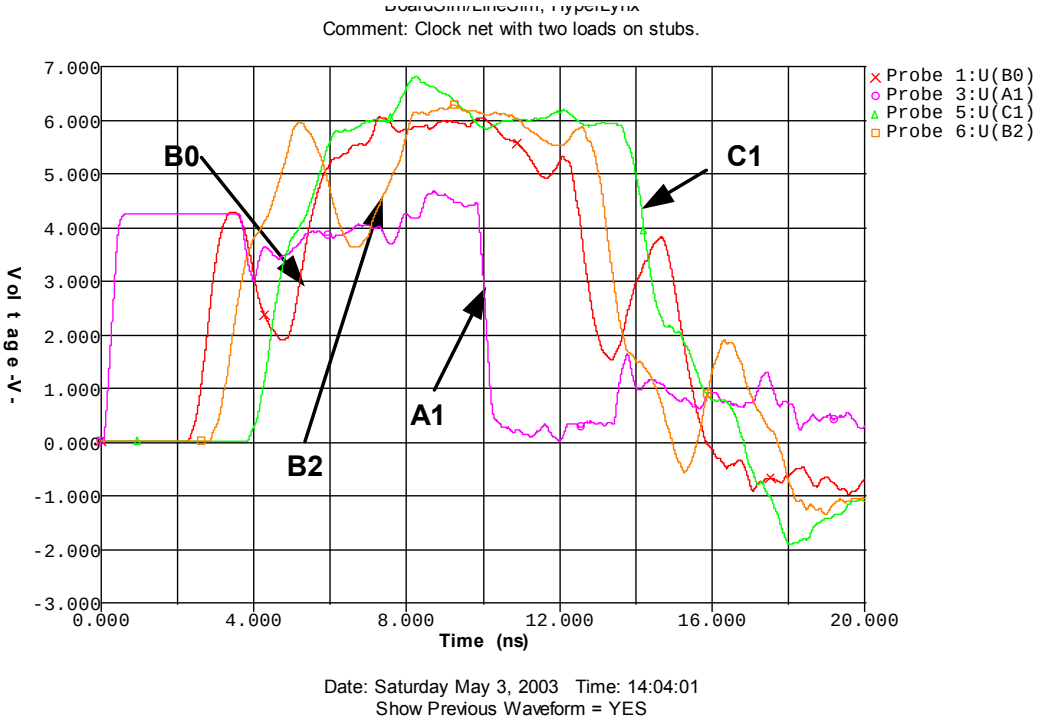

Figure 17.7. Waveforms of Clock circuit With Loads on Stubs

The waveforms occur in time, one after the other, based on where they occur along the transmission line. The first waveform is the driver A0. The second and third are B0 and B2 and the last one is C1. Looking at the A0 waveform first, it starts with a bench voltage that is larger than V/2 or 2.5 volts. It is approximately 4.2 volts. This is a clue that the output impedance of the driver is less than that of the transmission line it is driving, in other words, less than 50 ohms. It is possible to calculate the output impedance using Ohms law and the equivalent circuit in Figure 17.4. The value is approximately 6 ohms.

When this "bench voltage" arrives at the end of the line, an open circuit, it should double to 8.4 volts. The voltage level at C2 is much less than this. It is approximately 6 volts. This is because the PN diode junctions between the substrate and the input transistors that are normally reversed biased isolate the transistors from the substrate and each other. When an input voltage exceeds Vcc + 0.7V, these PN junctions begin to conduct and "clamp" the input voltage. This appears to be a desirable event. However, it is not. "Turning on" these parasitic diodes can have several undesirable side effects including:

- A parasitic 4 layer switch or SCR may exist between Vcc and Ground or Vdd and Vss that, when turned on, creates a short circuit that can cause a catastrophic failure.
- A parasitic path may exist between two adjacent inputs such that a quiet input is disturbed.

- Electrons may be injected into the substrate and disturb a logic state elsewhere in the IC.

Insuring that the combination of reflections, crosstalk, Vcc/ground bounce and other noise never cause one of these parasitic diodes to conduct, is a major goal of signal integrity engineering. It should be noted that these failures are not clock related. They can occur even if there is only a single switching edge per day. Whether or not such a failure occurs is a function of how fast a switching edge is compared to the length of the transmission line on which it is traveling, not its clock rate.

> Failures related to switching waveforms are a function of the edge rate of the signal, not its clock frequency. First and foremost, signal integrity engineering is based on rise and fall times.
> Clock frequencies are important only when a logic operation requires longer than a clock cycle to complete.
>
> **It's the edge rate that matters most.**

These failures can occur on any signal line. It doesn't matter whether it is a clock, data or a reset signal. Therefore, the concept of declaring some signals critical, such as clocks and data buses, and others "non-critical" is a sure way to encounter unwanted logic failures. What makes this kind of logic failure so troubling is that it can occur on any logic line at any time. It is likely to be a combination of overshoot, ground bounce, power supply ripple and crosstalk that only happens occasionally. Avoiding this kind of intermittent logic failure is the most compelling reason to perform rigorous signal integrity engineering prior to routing a PCB as well monitoring PCB layout to insure adherence to the SI design rules.

Continuing with the examination of the waveforms in Figure 17.7, it can be seen that the waveform at the input to C1 is a reasonable approximation of the desired square wave shape. Both B0 and B2 have reversals that can be taken as double clocking by the logic loads. This waveform reversal is a clue that one of two things has occurred. Either there are loads on stubs, as in this case, or there is a cluster of loads midway along a net. Either of these conditions is unacceptable. Figure 17.8 is the waveform set that results when the stubs are removed.

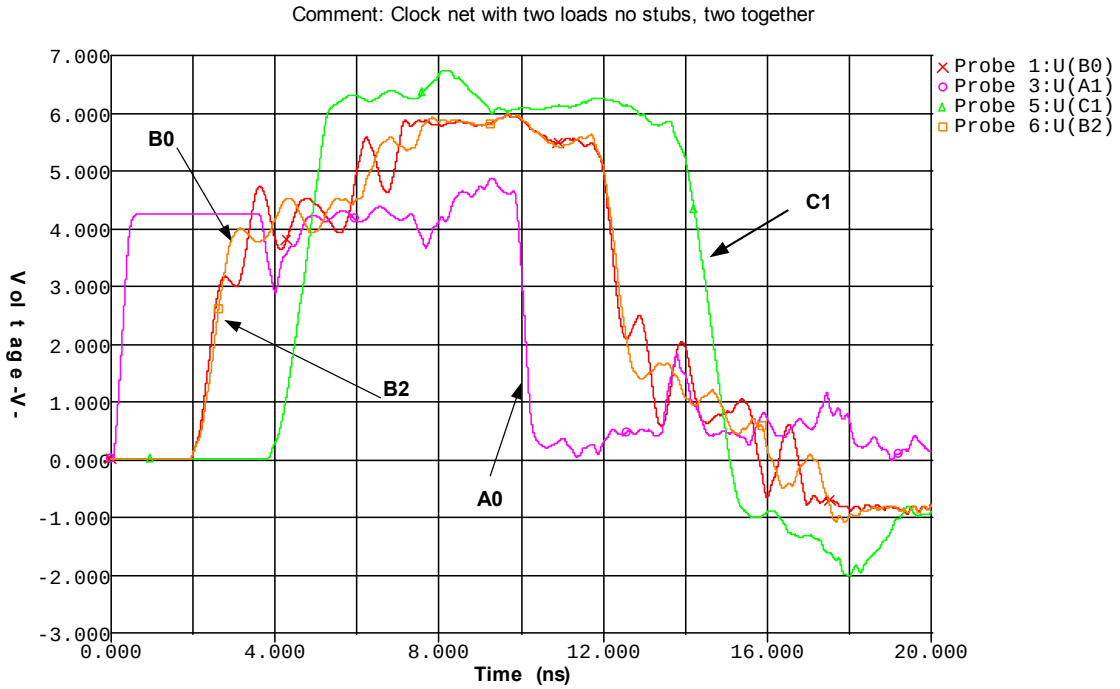

Figure 17.8. Waveforms of Clock Circuit with Stubs Removed

Notice that the waveform reversal is almost gone from B0 and B2 waveforms. Three loads being clustered near the middle of the transmission line cause the remaining waveform reversal. This reversal can only be removed by spreading the loads along the length of the transmission line.

The falling edge, the logic 1 to logic 0 transition, exhibits the same kinds of problems as the rising edges. When the bench voltage is more than V/2, there are waveform reversals and overshoot. With the falling edge, the overshoot drives the inputs below zero volts by more than a volt. This causes the same kind of failures as occurs when the voltage exceeds Vcc.

Before closing out this discussion, it is useful to see what happens when the parasitic diodes are removed from the inputs of the loads. Figure 17.9 is the same circuit as the one depicted in Figure 17.5 except equivalent input parasitic capacitances have been substituted for the actual gate inputs.

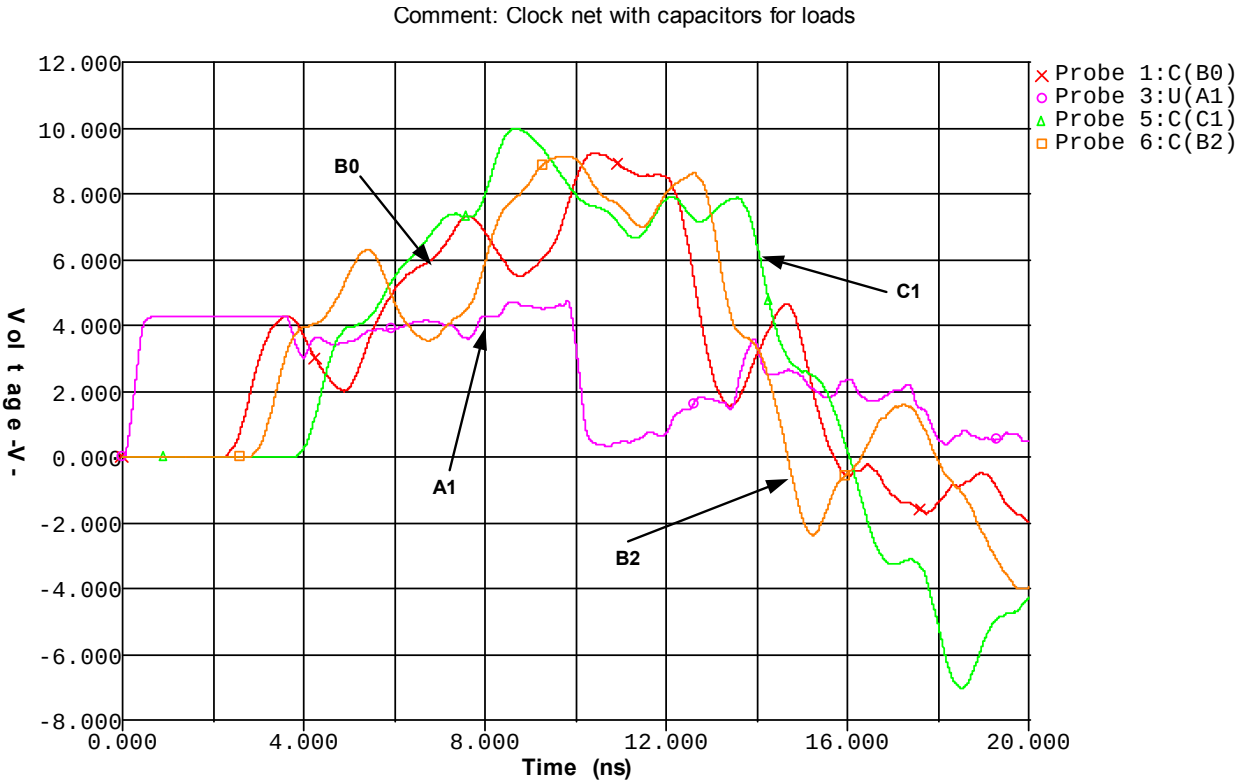

Figure 17.9. Clock With Two Loads on Stubs, Parasitic Input Capacitances Substituted for Gates

Notice that the vertical axis scale has been changed to 2 volts per division in order to fit all of the waveforms onto the screen. Without the parasitic input diodes clamping the waveform, reflections reach 10 volts positive and 7 volts negative. This illustrates the effect that the input diodes had as clamps. Said another way, these parasitic diodes have been forward biased with this much energy. It is this excess energy that couples to neighboring inputs, turns on parasitic switches and disturbs logic cells away from the offending input. In order to insure these reflections are not hidden by input parasitic diode structures, it is necessary to perform transient simulations with the equivalent input parasitic capacitances.

In the best case in Figure 17.9, if the objective is to create good clock edges for all three inputs, it is only achieved with C0 and then only by severely overdriving the inputs. B0 and B2 do not have smooth edges on which to clock data. Improving these two signals requires more attention to driver characteristics and termination rules than has been exercised in this example. These additional transmission line management techniques will be covered in later chapters of this book.

Waveform reversals on rising and falling edges are caused by stubs or clusters of loads in the middle of a transmission line.

In order to make full amplitude of reflections visible, loads should be simulated only with their parasitic input capacitance.

CHAPTER 18: WHAT IS MEANT BY SIGNAL INTEGRITY ENGINEERING

The term signal integrity management or engineering is often used to describe the steps involved in making sure that all signals arrive at their loads with enough adequate quality to insure proper logic operations. It might be useful to study just what "adequate quality" means. Figure 18.1 defines the job of signal integrity engineering. The lines surrounding this signal are often described as the envelope of acceptable values of voltage and timing for a waveform. It is this envelope that must be described to a board-level signal integrity analysis tool for each net in order for the tool to determine if all the signals are within specifications.

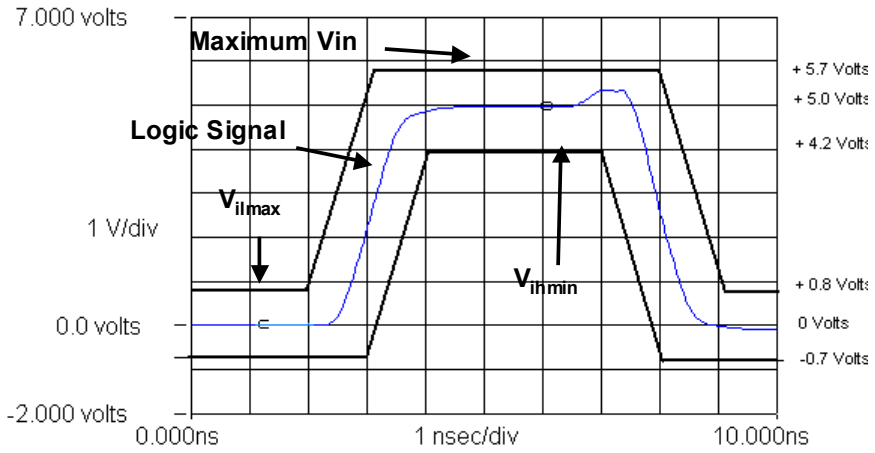

Figure 18.1. Envelope of Acceptable Values for 5-Volt CMOS Logic

The signal integrity goals have been met as long as the voltage waveform at the input to each device stays within these limits to assure proper logic operation. The waveform may have significant "noise" riding on it and it may transition at different times, but as long as it arrives within this envelope, the logic operation will be stable. Design rule generation starts with determining the limits shown in this diagram. For each logic family and each type of logic path, the time windows and the voltage windows must be established. Once this has been done, the time window can be devoted to those things that cause timing variations and the voltage window can be devoted to things that cause voltage variations.

The principal sources of timing variations are the time delays in the signal wires and the variations in propagation delay that are an inherent part of all semiconductor processing. The principal sources of voltage variations are reflections, crosstalk, Vcc and ground bounce and power supply variations (ripple). These are each handled by part of the signal integrity engineering process. This process is comprised of three parts--timing analysis and control, power supply engineering and voltage transient analysis and control.

> Design Rule creation begins with creation of the **envelope of acceptable values** and the rise and fall times for each logic family in a design. From this information, routing rules and power subsystem design is done to insure proper, reliable logic operation.

In almost all cases, the maximum allowable logic 1 is Vdd or Vcc plus the turn-on voltage of the parasitic diodes on the inputs. When input protection diodes have been added, these diodes are likely to be Shottky diodes that conduct at 0.3V. If no input protection diodes have been added, the parasitic diode will be a silicon PN junction that conducts at 0.7V. The minimum allowable logic 1 will be determined by the minimum logic 1 voltage level needed by an input to reliably detect that logic level. Any logic 1 that stays within the voltage window between these two levels is acceptable, even though it may have significant transients riding on it.

The minimum allowable logic 0 is Vss or 0 volts minus the same parasitic diode turn on voltages that apply to a logic 1. The maximum allowable logic 0 will be determined by the maximum voltage level at a logic input at which the input can

reliably detect a logic 0. Any logic 0 that stays within the voltage window between these two levels is acceptable, even if it may have significant transients riding on it.

From the discussion above, a successful PCB design process is more than the two-part schematic creation and PCB layout methodology used by "TTL" designers. In order to succeed, several new steps are needed. Among these are timing analysis, power subsystem design and transient analysis. In order to accurately complete any of these three operations, information from the proposed PCB layout must be fed back into the schematic and layout tools to determine if the SI goals will be met. This involves an iterative loop between all of the steps in the process. This can only be achieved if these operations are combined into a tightly linked process.

There is the notion that SI analysis can be done after the PCB is routed. It's possible, although it is very difficult to do, primarily because of the difficulty in collecting simulation models for every part on the PCB. If all of the models are obtained, the post-route SI analysis process is still flawed. It's rather like locking the barn door after the horse is gone. Without doing pre-route analysis by imposing routing rules on the routing process, the chances of achieving a routed PCB free of SI violations is very low. Usually, the number of SI violations is so high that it is not possible to correct them without starting over. If the PCB is routed again without the use of well thought out SI rules, the result will be the same.

Certainly, information obtained in the first post-route analysis can be used to refine the next route of the PCB. This will reduce the number of post-route failures. Using this process, it is possible to achieve a design that meets all of the SI goals. Along the way, all of the analytical work needed to arrive at a stable design must be done. However, if this analytical work is done before the first route is attempted and proper routing rules are imposed on it, the probability of a "right the first time" PCB becomes very high. This pre-route methodology is what this book is about. The process isn't as hard as it might seem. It's all about "knowing where the rocks are."

CHAPTER 19: WHEN IS A DESIGN HIGH SPEED?

It seems to be well understood that there is a speed below which a design is considered “low speed” and does not require transmission line management, as with TTL, and a speed above which this discipline is required, as with ECL. What is not well understood is when a design moves from the “low speed” class to the “high speed” class. A common way for deciding when a design makes this transition has been to look at the frequency of the clock. If the clock frequency seems low, a design has often been deemed to be low speed.

From the waveforms displayed in Chapter 17 (Figures 17.6, 17.7 and 17.8), it can be seen that the reflections at the open end of the transmission exceed the input voltage rating of the logic family being driven. In this case, it is 5 volt CMOS which can tolerate logic 1 voltages as high as +5.7 volts and logic 0 voltages as low as -0.7 volts. Clearly, this condition must be avoided. The excess voltage problem is caused by the fast edge, not the clock frequency or the rate at which edges occur. When the edges are fast enough compared to the length of the transmission lines this over voltage condition happens. At what point is this true? In order to develop a feel for when this condition begins to be a problem, the circuit in Figure 17.2, a 5-volt CMOS circuit will be used. This analysis will be done on the falling edge, the logic 1 to logic 0 transition. This measurement is used because the falling edges of most logic families are faster than the rising edges. Once the length of transmission line at which the voltage doubling violates the input voltage rating is established, this can be used as a measure of when a design becomes high speed for either logic transition.

Figure 19.1 is comprised of the same set of conditions as Figure 17.2 and shows the falling edge.

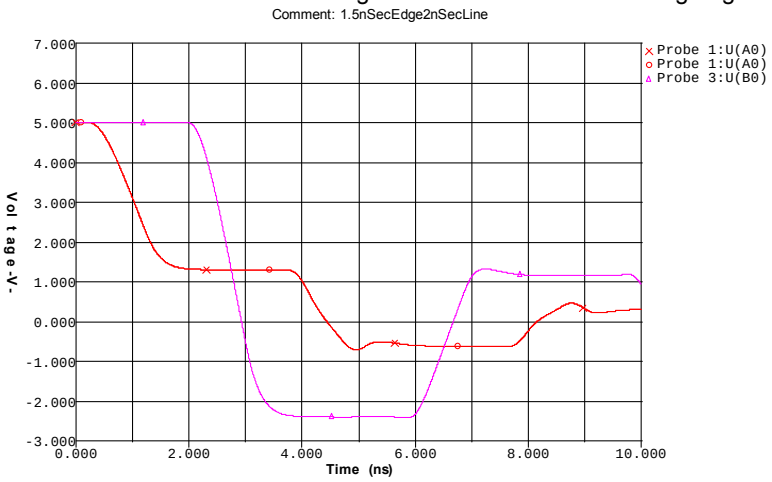

Figure 19.1. 5-Volt CMOS Circuit with 1.5 Nanosecond Edge on a 2 Nanosecond Long Transmission Line

It can be seen in Figure 19.1 that the overshoot has gone approximately 2.4 volts below ground, a condition that is unsatisfactory. At 12 inches, the transmission line is longer than one TEL. One TEL is 1.5 nanoseconds. In Figure 19.2 the line has been shortened to ½ the TEL.

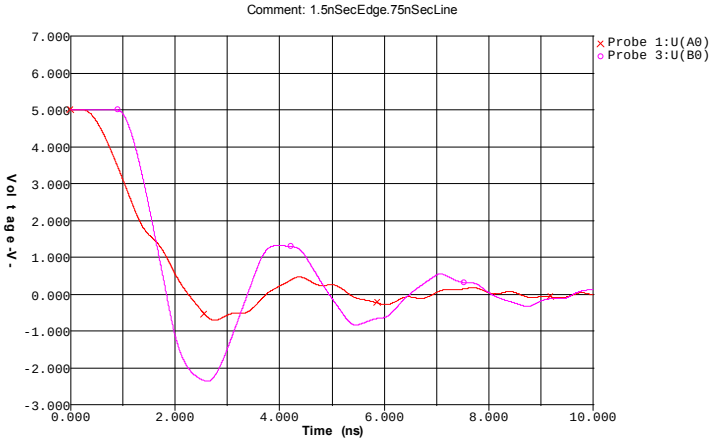

Figure 19.2. 5-Volt CMOS Circuit with 1.5 Nanosecond Edge on a 0.75 Nanosecond Long Transmission Line

The waveforms in Figure 19.2 are beginning to resemble a set of dampened sine waves. They are still the result of reflections. The line on which they are reflecting has gotten so short that there hasn't been time to see the "flat" portions that were visible in Figure 19.1. The line is still long enough that the voltage has had time to double at the load end of the line. Clearly, at this ratio of rise time to line length, the design is still "high speed." In Figure 19.3 the line has been shortened to ¼ of a TEL or 375 picoseconds.

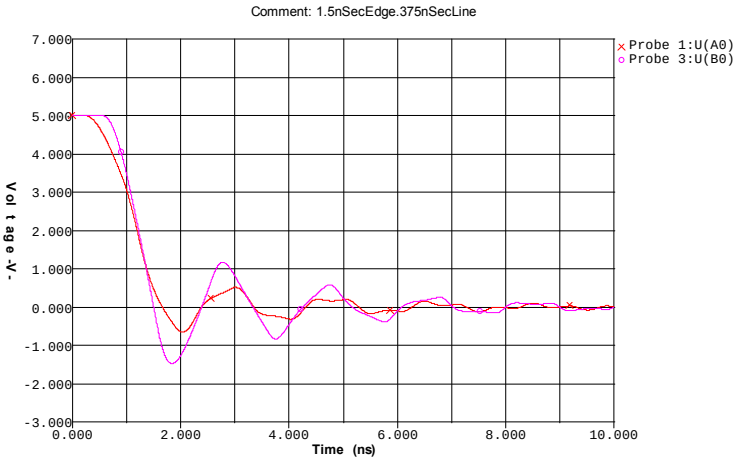

Figure 19.3. 5-Volt CMOS with 1.5 Nanosecond Edge on a 0.375 Nanosecond Long Transmission Line

Finally, at a length of ¼ TEL, the amplitude of the reflection at the load has been reduced from -2.4 volts to -1.4 volts. At last, the line is growing short enough that signal integrity concerns from reflections are diminishing. The energy to cause a failure ranges between -0.7 volts and -1.4 volts. For most CMOS circuits, this is not enough to induce failures. For ECL circuits this would not be true. **A line length less than ¼ TEL would be the boundary between low speed and high speed.**

Now that a boundary has been established between low and high speed, this can be turned into a length rule for the purposes of estimating when transmission line rules need to be applied to a line. Lines less than this length can be treated as "don't care." Lines longer than this length must be managed using signal integrity rules. The example has a TEL of 1.5 nanoseconds. ¼ of this is 375 picoseconds. In most printed circuit PCBs, signals travel at approximately 6 inches in one nanosecond. Therefore, for 1.5 nanosecond edges, lines longer than 2.25 inches will require impedance control and some form of reflection management.

Table 19.1 lists several common logic families with their rise times and TELs listed. Logic families in current use all have rise and fall times of 0.5 nanoseconds or less. Even logic parts that are intended for general use have such fast edges. This is attributable to the continually shrinking size of the transistors that comprise these logic parts. As a result of these ever faster edges, the length of a transmission line at which the line becomes "high speed" grows smaller and smaller. Table 19.1 shows these dimensions for rise and fall times ranging from 1 nanosecond to ¼ nanosecond.

Rise Time (nSEC)	TEL in Inches	TEL in cm	¼ TEL in Inches	¼ TEL in cm
1	6	15.2	1.5	3.81
0.5	3	7.62	0.75	1.90
0.25	1.5	3.81	0.375	0.95

Table 19.1. Rise Time vs. TEL and ¼ TEL

From the data in Table 19.1, it is clear that virtually every signal on most current PCBs has the potential to fail from overshoot. This ratio of TEL to rise and fall time is the primary determining factor for invoking impedance control.

> Every signal path longer than ¼ **TEL** has the potential to malfunction from overshoot. This is the first test to determine when a design becomes high speed.

CHAPTER 20: CONTROLLING REFLECTIONS BY USING TERMINATIONS

In Chapter 19, the transient behavior of switching signals was demonstrated. It was shown that the EM energy traveling down a transmission line will reflect back and forth along the line unless it is absorbed. The circuit in Figure 17.3 was fast enough that this reflected energy caused the voltage that appeared at the input to the load to be double the voltage level that started down the transmission line. In the process, the doubling exceeded the maximum allowable "1" voltage of +5.7 volts. Figure 20.1 shows both the rising and falling edges of this circuit. The falling edge also doubles and goes below ground by more than 2 volts, exceeding the limit of -0.7V. The two horizontal lines show the maximum allowable signal swing.

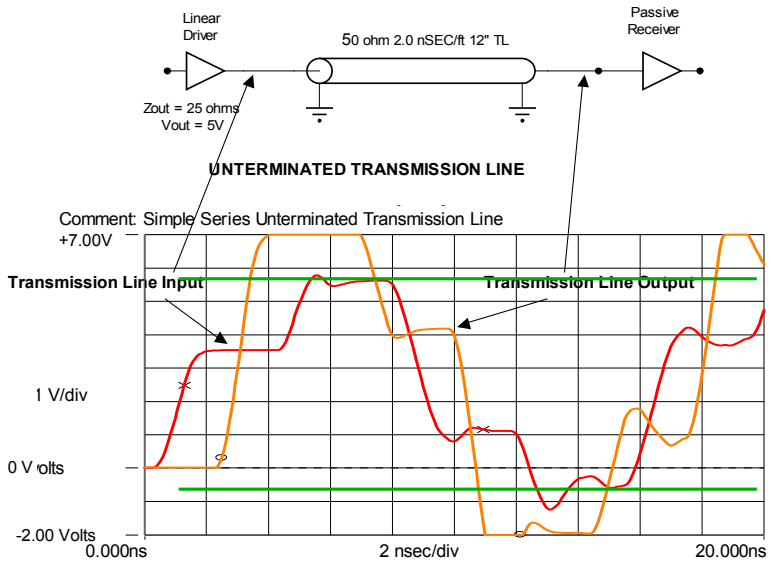

Figure 20.1. Rising and Falling Edges of 5V CMOS Circuit Showing Overshoot on Both Edges

The reason for the excessive voltages is that the size of the signal that started down the transmission line was too large. When it doubled, the resulting voltage was too large. The value of the voltage that started down the transmission line is determined by the voltage divider formed by the output impedance of the driver and the impedance of the transmission line as shown in Figure 17.4. If something is done to make these two impedances the same size, the 5 Volt starting signal will be divided in half and the signal that starts down the line will be +2.5 Volts. When this signal arrives at the open end of the line and doubles, it will be 5 volts--just what was needed. Figure 20.2 shows this. The divider ratio was adjusted by adding a 25-ohm resistor to the output of the driver.

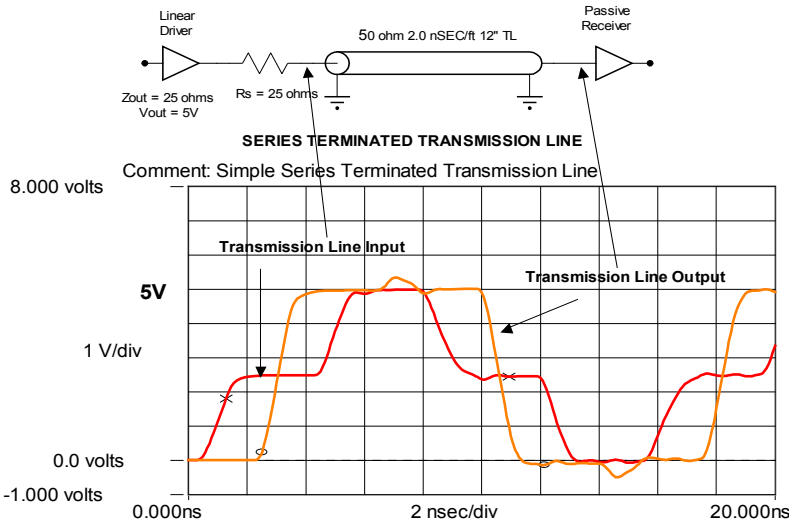

Figure 20.2. Rising and Falling Edges of 5V CMOS Circuit Showing Effect of a Series Termination

As before, the 2.5V signal traveled down the transmission line arriving at an open circuit. The open circuit did not absorb the energy in the EM field, so it all reflected back toward the source. On the outbound trip, the parasitic capacitance of the transmission line was charged to V/2 or +2.5 volts. On the return trip, the parasitic capacitance was charged the rest of the way up to +5 volts. When the EM field arrives back at the source it encounters the equivalent circuit shown in Figure 20.3.

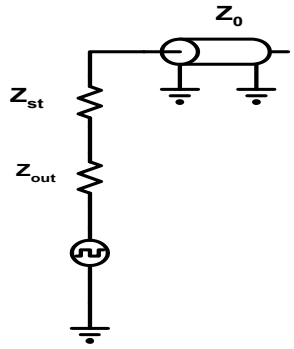

Figure 20.3. Equivalent Circuit at Source as Seen by Reflected Wave

Zout is 25 ohms, Zst (series terminator) is 25 ohms, totaling 50 ohms, and the voltage source is a short circuit. The transmission line impedance is 50 ohms. The effect is to provide the transmission line with a perfect termination. The 50-ohm termination absorbs all of the energy in the returning EM field so there is no reflection. The signal is stable at +5 Volts. When the signal switches from a logic 1 to a logic 0, the same events take place. The load is provided with the square wave that was intended and the input voltage rating of the part has not been violated.

The circuit in Figure 20.2 is said to have been "series terminated". The switching that results is often called **reflected wave switching**, because data becomes good all along the line only when the reflected wave passes by on its way back to the source. Notice that the voltage waveform at each end of the line is different. Only the load end of the line has valid logic levels at all times. Anywhere between the driver and the load, the voltage levels are half way between a 1 and a 0 for some period of time. This is an invalid logic state. Edge sensitive loads, such as clock inputs, could not be located anywhere except at the end of the line farthest from the driver.

The reflected wave switching shown in Figure 20.2 is the basis for the PCI bus. It is the lowest power method for doing high speed signaling. However, it has limitations when used with a bus such as PCI. The limitation is the time duration of the two bench voltage levels while the signal makes a round trip on the bus. No logic operations can be performed until this "dead time" elapses. This is the reason that the original 33 MHz PCI bus had a bandwidth limitation--it limited the amount of fast CPU performance that was available to a user. The original 33 MHz PCI bus was allowed to be 30 inches long. The round trip delay on such a bus is 10 nanoseconds. The total time in a clock period is only 30 nanoseconds. On each switching edge, 10 nanoseconds is consumed as dead time. This leaves only 10 nanoseconds for the two logic levels. Increasing the clock frequency does not reduce the dead time. It only reduces the data good time.

Now, we have 66 MHz and 100 MHz PCI bus systems. How is this possible? The 66 MHz PCI bus specification states that the maximum bus length cannot exceed 9 inches. The round trip delay of this bus length is 3 nanoseconds. Out of this 15-nanosecond clock period, only 6 nanoseconds is used as dead time, 9 nanoseconds remains for logic operations. Success! The 100 MHz PCI bus has a clock period of only 10 nanoseconds. For this to work, the bus length is restricted to 5 inches or a round trip delay of about 2 nanoseconds.

There is a pattern developing in the above discussion. In order to use series terminated logic in a bus-organized system, it is necessary to reduce the system size as the clock frequency increases. This minimizes the dead time. It is difficult to build meaningful systems of this kind at clock frequencies above 100 MHz. How do supercomputers with clock frequencies over a GHz work?

Suppose the EM energy is absorbed at the load end of the line by placing a termination there, as shown in Figure 20.4. The events in the first part of the operation are the same as all of the previous examples. The bench voltage of +3.3 Volts is launched down the transmission line as the signal. 2 nanoseconds later, the EM field arrives at the load end of the line. This time there is a 50-ohm termination (an absorber) that consumes the energy in the EM field. There is no reflection. Figure 20.5 shows this for both the rising and falling edge. We have the same waveform at all points along the transmission line. There don't appear to be any illegal logic states or reflections. Unfortunately, this is a 5-volt CMOS circuit. The minimum logic 1 for this logic family is +4.2 volts. The logic 1 in this example does not reach this level. Even though there are no reflections, this circuit will not work. Something must be done to raise the logic 1 level. The divider formed by the output impedance and the line impedance set the logic 1 level. One of these will need to change. It is difficult to change line impedances enough to remedy this problem, so the driver output impedance will need to be reduced. Figure 20.6 illustrates this. A new driver has been located with an output impedance of 5 ohms. This time the circuit has a 3.3-volt CMOS driver.

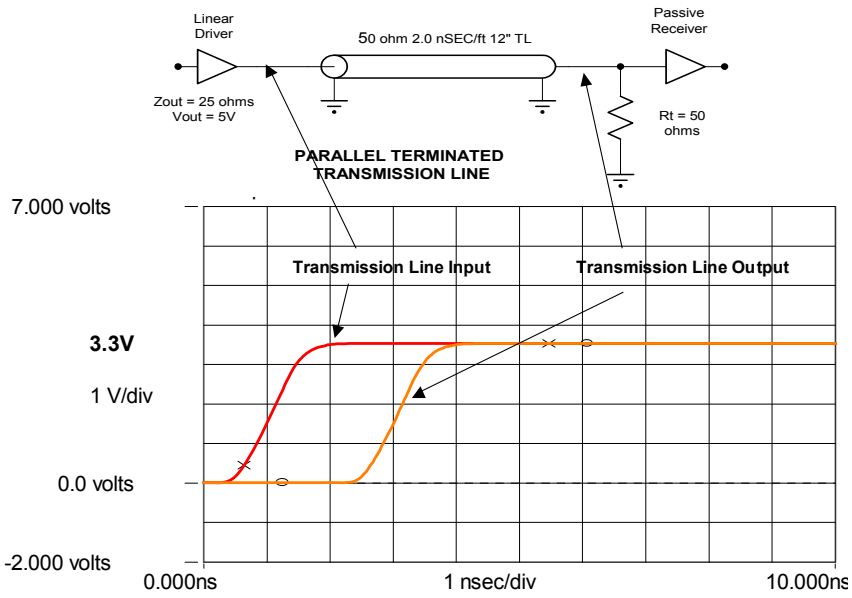

Figure 20.4. 5V CMOS Circuit With Parallel Termination

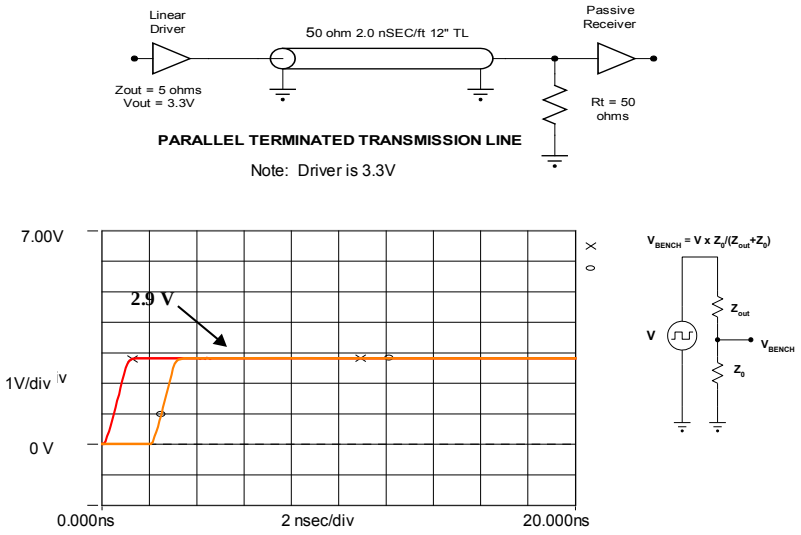

Shows Logic 0 to Logic 1 Transition

Figure 20.5. 5-Volt CMOS Circuit With Parallel Termination Showing both Rising and Falling Edges

Notice that in this example the bench voltage, which is the logic 1, is 10/11 of V or 3 volts. This is a proper logic 1 level for this circuit. All of the conditions have been met. There are no illegal logic states. A load can be placed anywhere along a transmission line with the assurance that it will always see a proper logic signal. This is called “parallel” termination. It is the terminating method used for all very high-speed logic paths. However, it should be noted that even this signaling protocol has drawbacks. In this case, it’s power consumption. With the 3.3 Volt signal swing, the power per signal line approaches 1/5 watt, which is far too high to allow for practical systems. For this reason, the signal swings of all logic families intended for parallel termination are small. For example, ECL signal swings are approximately 1 volt; GTL signal swings are 800 millivolts and LVDS signal swings are 400 millivolts.

This low level logic works very well at high speeds. However, due to the small signal swings, it does not have a very big noise margin. Noise management becomes a very important part of the design task with these logic families. This is especially important when they are used on a mixed logic system that contains 3.3 volt or 5 volt CMOS circuits.

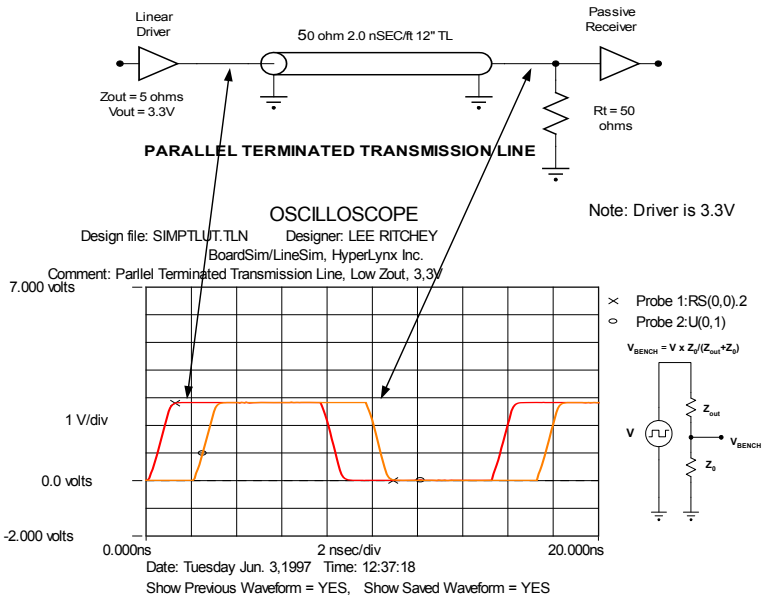

Figure 20.6. 3.3-Volt CMOS Circuit With Parallel Termination Showing both Rising and Falling Edges

> When using a parallel termination, the bench voltage is the logic 1 voltage.
>
> In order to create a logic 1 voltage large enough for proper operation, the output impedance of the driver must be much less than the line impedance.

What about AC termination?

AC termination is sometimes suggested as a way to control the voltage doubling at the open end of a transmission line. What is AC termination? Where is it used? AC termination attaches the parallel terminating resistor to the end of a net with a small capacitor. The objective is to provide the termination during the time edges are switching and disconnect it when the logic levels are “steady state”. This method was devised when TTL edges got fast enough that they exceeded ¼ TEL and caused excessive voltages at the inputs of gates. The result of attaching an AC termination to the end of a net is a rising or falling edge that has an RC time constant that effectively slows down the edge while it limits overshoot. If this edge degradation is acceptable, an AC termination might be a way to cope with fast edges.

Figure 20.7 is the same circuit from Figure 17.2 with an AC termination. Notice that when the value of the resistor and capacitor is chosen such that the overshoot does not exceed Vdd +0.7 Volts, the signal is beginning to look like a sine wave. The edges are no longer sharp. If the clock frequency is increased much beyond the 66 MHz in this example, not only is the waveform more like a sine wave than a square wave, it will no longer be capable of maintaining the required signal swing. This problem occurs when an AC termination is attempted with DRAM arrays. It is not a well-behaved methodology at high clock rates and should instead only be considered as a “band-aid” solution for a circuit that should have been designed with a true series or parallel termination.

What About Diode Terminations?

Using diodes at the receiver end of a transmission line as a substitute for a resistive termination is another example of a band-aid approach. Rather than designing transmission lines with proper terminations that prevent overshoot from becoming excessive, a pair of diodes is attached between the signal line and the two power rails and oriented such that when overshoot exceeds Vdd, one diode turns on as a clamp as depicted in Figure 20.8. When the overshoot attempts to go below Vss, the other diode turns on as a clamp. This does work. However, the diodes must be Shottky diodes in order to turn on quickly enough. The cost per line to solve problems this way is quite high.

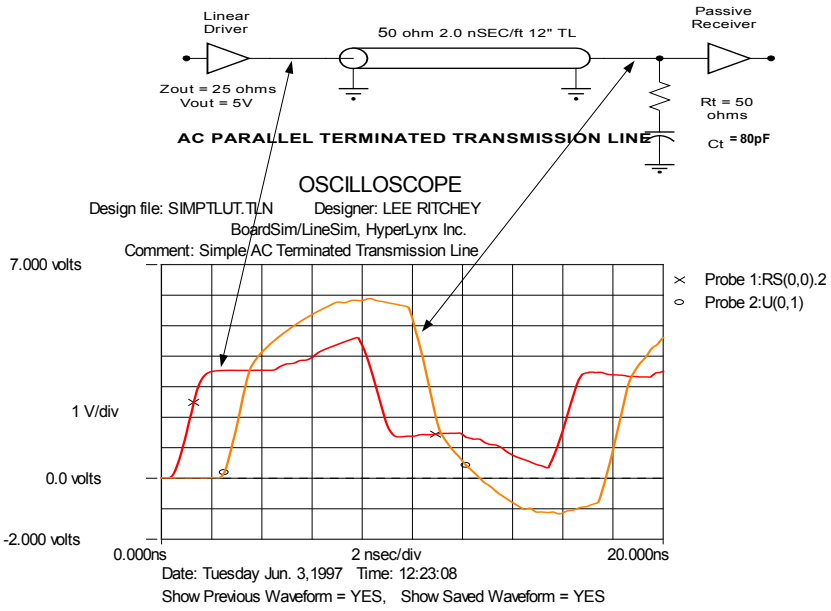

Figure 20.7. 5-Volt CMOS Circuit with AC Termination

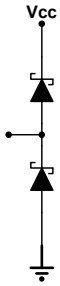

Figure 20.8. Diode Clamp Termination

What about Thevenin Terminations?

Parallel terminations have been shown in this section as being connected to ground. This is a symbolic ground. The actual parallel terminations always connect to a special terminator voltage, not ground or Vdd or Vee. In the case of ECL which operates between ground and -5.2 volts, the terminating resistors actually are connected to a special Vtt supply that is -2.0 volts. GTL terminations connect to +1.2Volts. Parallel terminations for 2.2 Volt CMOS connect to +1.1 volts.

When using these logic families, it is necessary to add a power supply and power plane to supply this terminator voltage. When there are only a few circuits that need parallel terminations, as is the case when PECL is used for an interface to a transceiver, this is a large expense for just a few lines.

There is another solution to this problem. It is possible to use a two-resistor network to emulate the terminator impedance and the terminator voltage. This is called a Thevenin equivalent. Figure 20.9 shows an example of a Thevenin equivalent terminator.

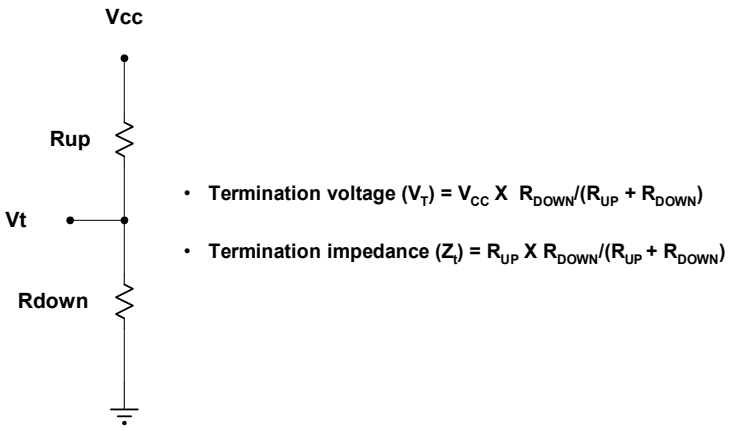

Figure 20.9. A Thevenin Terminating Network

To arrive at the values of the two resistors required to create the equivalent voltage and impedance, solve the two simultaneous equations in Figure 20.9.

Thevenin Networks as Pull-ups or Pull downs

A Thevenin network can be used to create a pull up to some voltage other than Vdd or a pull down to some voltage other than ground. An example of this is the resistor networks on VME bus backplanes. Figure 20.10 is an example of such a pull up. A TTL output has an unsymmetrical output. The impedance of the output when switching from a 1 to a 0 is much lower than when switching from a 0 to a 1. As a result of this lack of symmetry, the rise time can be too slow to make timing margins. Adding a pull up to +3V, the maximum logic 1 for TTL, provides more current to charge up the line. This results in an improved rising edge while only moderately degrading the falling edge.

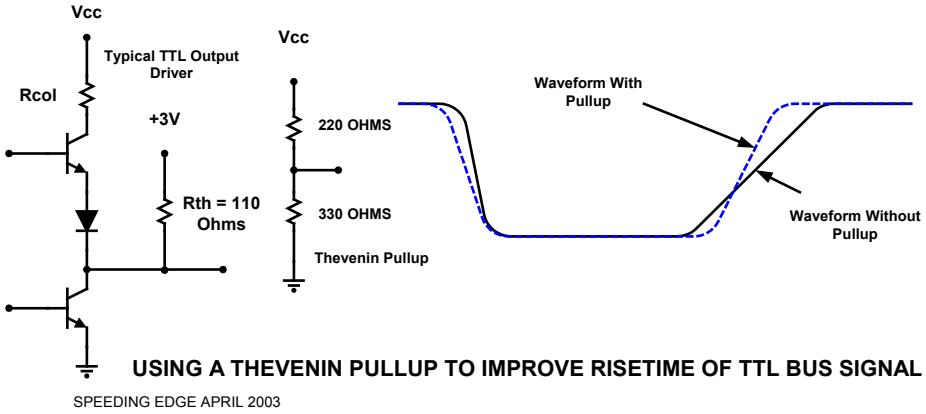

Figure 20.10. A Thevenin Network Used as a Pull up for a TTL Bus on a VME Backplane

Summing it all up, there are five ways to terminate a transmission line. Table 20.1 lists these terminators and their important characteristics.

TERMINATOR TYPES AND PROPERTIES					
TERMINATOR TYPE	ADDED PARTS	DELAY ADDED	POWER REQUIRED	PART VALUES	COMMENTS
SERIES	1	YES	LOW	Rs = Zo-Zout	GOOD DC NOISE MARGIN
PARALLEL	1	SMALL	HIGH	R = Zo	VERY HIGH POWER CONSUMPTION
THEVENIN	2	SMALL	HIGH	CALCULATE	HIGH POWER CONSUMPTION FOR MOS AND TTL DRIVERS
AC NETWORK	2	SMALL	MEDIUM	R = Zo,	DETERMINE VALUE OF C WITH SIMULATOR
DIODE PAIR	2	SMALL	LOW	N/A	NOT A TERMINATION, CLAMPS OVERSHOOT
PREPARED BY Ritch Tech FROM DATA SUPPLIED BY MOTOROLA, 1997					

Table 20.1. Terminator Types and Properties

Figure 20.11shows where each type of termination is located on a network.

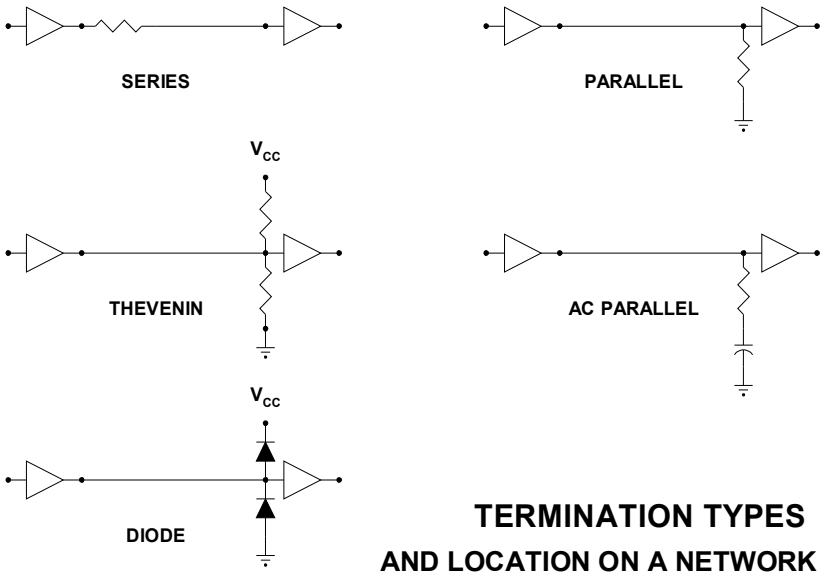

Figure 20.11. Terminator Types and Locations

While there are five types of terminations in the above charts and explanations, only three of them are necessary. These are the series termination, the parallel termination and the Thevenin equivalent parallel termination. All logic intended for high speed signaling is capable of being handled with one of these terminations. If a set of design rules seems to need AC terminators or diode terminators, it would be a good idea to go back through the decision making process that led to their use. It's more than likely that some error was made in devising the design rules.

What Happens When a Series and a Parallel Termination Are Used on the Same Net?

This question comes up in almost every one of my classes. For some reason, there is the perception that both a series and a parallel termination are needed on a network. Figure 20.12 is an ECL net that has a series termination at the output of the driver and a parallel termination at the load end. Note that the signal that arrives at the load never gets up to the -0.8 Volts that is needed for an ECL logic 1. This is because the series termination and the transmission line have divided down the output signal before it starts down the transmission line. Since there is a parallel termination at the load end, there isn't any way for this signal to double as a way to reach a proper logic 1. In this case, the "series" termination was

intended to serve as a current limiting resistor, which it is. Unfortunately, the transmission line also sees it as a series termination.

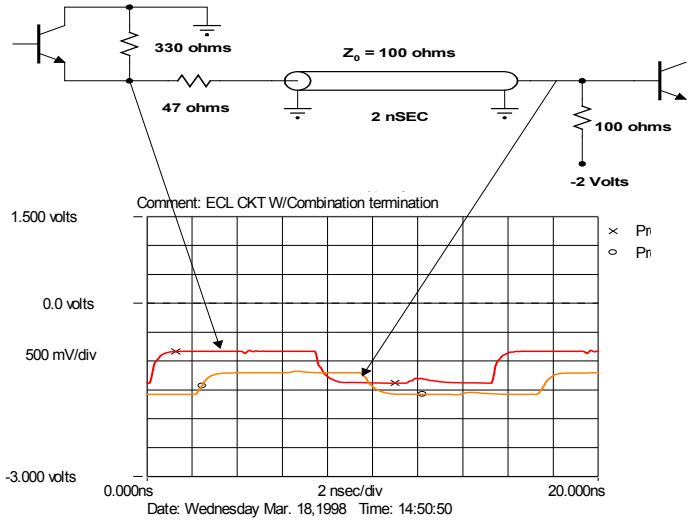

Signal is eroded. Don't use both series and parallel terminations on the same line.

Figure 20.12. An ECL Network with Both a Series and a Parallel Termination

As often happens, there are exceptions to rules. There are occasions when a termination is needed at both ends of a transmission line. Two that come to mind are a video driver that has an emitter follower for an output and OC-48 drivers.

In the case of the video driver, emitter followers have a tendency to oscillate. A common method for preventing this oscillation is to place a small resistor in series with the emitter as it drives the transmission line. When this is done, the signal that is launched into the transmission line is reduced in amplitude similar to the ECL signal shown in Figure 20.11. This problem is overcome by designing the video amplifier such that it creates a larger starting voltage.

In the case of 0C-48 drivers, small reflections are created by the imperfections, such as connectors, in the transmission line path. These small reflections arrive back at the driver that is usually a pseudo current source, meaning that it has a high output impedance. The energy in the small reflections is reflected off the driver's high impedance and travels back toward the load. Upon arrival at the load, these reflections add to the jitter. By adjusting the output impedance of the driver so that it exactly matches the line impedance, these small reflections are absorbed, improving jitter. The two ends of the transmission line are terminated--the driver end has a series termination and the load end has a parallel termination. In this case, the driver design must account for this change.

With commercially available logic components, the user has no way to adjust the driver characteristics to allow for the usage of both a series and a parallel termination

.

> A transmission line driven by commercially available logic parts is either series or parallel terminated, but never both.

CHAPTER 21: TERMINATOR TYPE, TERMINATOR PLACEMENT AND NET SEQUENCING

In Chapter 20, two methods were demonstrated for controlling reflections. These included the use of series terminations at the output of a driver and parallel terminations at the input to the last load or the load end of the transmission line. **These are the only two choices.** Deciding which method to use involves a trade off between power consumption, system bus speed and the type of logic selected for the task at hand.

SERIES TERMINATED TRANSMISSION LINES

The lowest power consumption method for high speed logic signaling involves placing series terminations at the output of each driver. This is the lowest power consumption method because energy is only consumed in the circuit when a logic line is switched from a logic 0 to a logic 1. Figure 21.1 shows the current and voltage waveforms for a transmission line that is series terminated (for example, Figure 20.2). The power required can be calculated by multiplying the current waveform by the voltage waveform at the input to the transmission line. One of these power pulses occurs each time that the logic level changes from a logic 0 to a logic 1. This is the reason that power consumption is a function of clock rate in most CMOS circuits. Later on, it will be demonstrated that the inability of the power subsystem to supply current at the frequencies contained in these current pulses is the major source of EMI in most products as well as a major source of "ripple" and Vcc/ground bounce that induce logic failures.

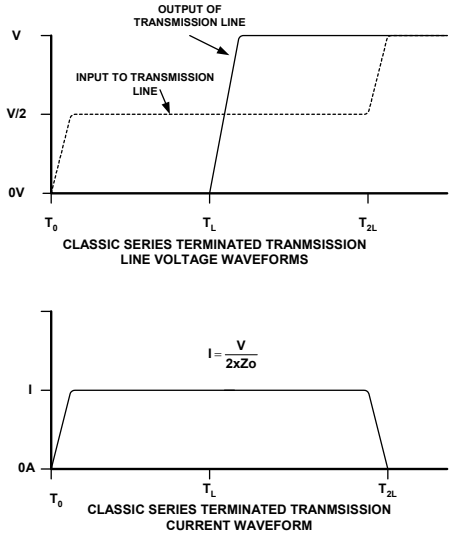

Figure 21.1. Voltage and Current Waveforms for a Series Terminated Transmission Line

Clearly, if the objective is to achieve the highest data rate with lowest power consumption, series terminated transmission lines are the right choice. Why isn't this the only method used if it is the lowest power consumption choice? As can be seen from the voltage waveform in Figure 20.2, only the last load has a square wave signal with no illegal logic states. The driven end of the line and all points along the net are held at a voltage level of V/2 for part of every switching cycle. As a result, systems that use series terminated lines for buses have a "down time" that lasts as long as the round trip delay on the longest line in the bus. When this "down time" is no longer acceptable, an alternative approach is needed. That approach is parallel termination. Before discussing this approach, a few more observations about series terminated lines are in order.

As can be seen from the discussion regarding matching a driver to a transmission line, the output impedance of the driver must be equal to or less than the impedance of the line being driven. This is one of the criteria that will be used to select drivers for this service.

> When selecting drivers for series terminated transmission lines, the output impedance of the driver must be equal to or less than the impedance of the transmission line.

If the output impedance of the driver is less than the impedance of the transmission line, a discrete series termination resistor will need to be added at the output of the driver such that the sum of the output impedance and the series termination resistor equals the TL (transmission line) impedance. Equation 21.1 can be used to calculate the value of the series termination resistor.

$$\mathbf{Z_{st} = Z_0 - Z_{out}}$$

Where Z_{st} = Series termination resistor, Z_0 = Transmission line impedance, Z_{out} = Output impedance of driver

Equation 21.1. Series Termination Value

The actual value of the series termination should be less than that shown in equation 21.1. This is to allow for the tolerance variation of the transmission line impedance that occurs in the normal manufacture of PCBs--usually ±10%. Why not match the transmission line exactly? Remember that a series termination has a dual purpose. The first is to launch a signal that is V/2 and the second is to absorb the reflected energy as it arrives back at the source. In order to insure at least a V/2 signal at all times, it is necessary to account for the PCB impedance tolerance. Launching a signal that is less than V/2 results in a voltage level at the load that is less than V that compromises the noise budget. Figure 21.2 illustrates the problem.

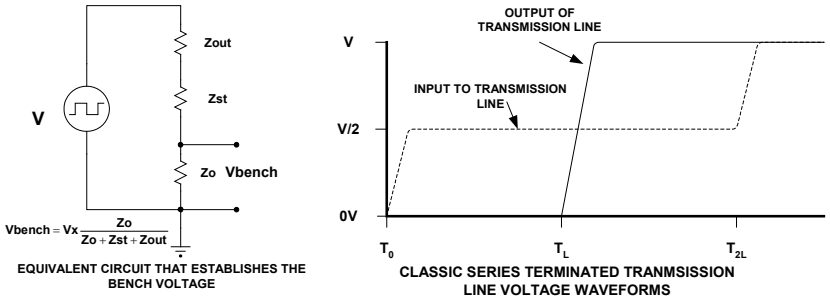

Figure 21.2. Voltage Waveforms on a Series Terminated Transmission Line

The worst-case tolerance occurs when the PCB impedance is on the low side of its tolerance range. For this condition, calculate the value of Z_{st} such that the sum of Z_{out} and Z_{st} equals the lowest value of Z_0 that will be encountered. In this way, the voltage level launched into the transmission line (the bench voltage) will never be less than V/2.

> The value of a series terminating resistor should be selected on the low side of the PCB tolerance range in order to insure that the voltage waveform launched into the transmission line is never less than V/2.

A fair question is how far can the series termination be from the output pin of the driver and still function as a series termination? Figure 21.3 shows the equivalent circuit that results when a series termination is added to a TL. Effectively, the resistor has been inserted between two segments of a transmission line.

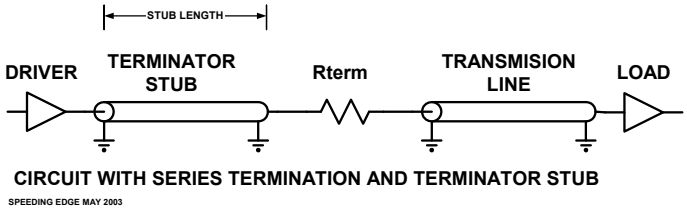

Figure 21.3. Equivalent Circuit of a TL With a Series Terminating Resistor

As the terminator stub increases in length, there comes a point where the resistor behaves as though it is between two transmission lines. Instead of achieving the neat V/2 waveform at the input to the TL, unacceptable waveform degradation occurs. In order to determine how long the "terminator stub" can be without incurring unacceptable waveform degradation, it is necessary to simulate the circuit in a spice modeler or an IBIS-based modeling tool. The length of the terminator stub is made successively longer and the effect it has on the waveform at the load is observed. At some length, the degradation will violate the input requirements of the load. This will establish the maximum length for the series terminator stub. This length is supplied to the PCB layout staff as a guide for placing series termination resistors. It is important to do this analysis at the fastest rise or fall time of the circuit being analyzed.

> The maximum allowable length of a series termination stub--the distance this resistor can be located away from the driver and still do its job--can only be determined with a good simulation tool.

PARALLEL-TERMINATED TRANSMISSION LINES

When the "dead time" associated with the bench voltage level of series terminated logic is unacceptable, a terminating technique that does not create this condition must be employed. Figure 17.1 depicts such a circuit. Figure 21.4 shows the voltage and current waveforms for a parallel-terminated transmission line. Notice that the current flows continuously at a logic 1. The power consumption of this type of transmission line is not clock frequency dependent. Current flows for half of every clock cycle, no matter what the clock frequency is. This is the disadvantage of parallel termination. The advantage is that there is no "dead time." At all times, the logic levels everywhere along the TL are valid, allowing buses to be built without being concerned about invalid logic states. This also allows clocks to be distributed to loads along the TL, not just at the ends of lines as is required with a series termination.

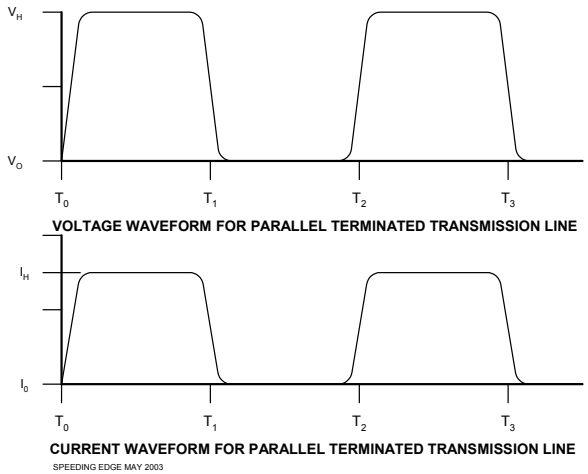

Figure 21.4. Voltage and Current Waveforms for a Parallel-Terminated Transmission Line

It is important to remember that the logic 1 level for a parallel-terminated TL is the bench voltage created by the voltage divider that is comprised of the output impedance of the driver and the TL impedance. In order for this level to be a valid logic 1, the output impedance of the driver must be much less than the TL impedance.

> When selecting drivers for parallel-terminated transmission lines, the output impedance must be much less than the TL impedance, preferably smaller than 10:1.

The value of the terminating resistor for a parallel-terminated TL should be the same value as the TL impedance. The actual value should be selected such that when the worst case impedance tolerance build ups occur, the reflection from

the termination is overshoot, which adds to the incident voltage level or creates no reflection at all. As with series termination, the major tolerance build up is the ±10% impedance of the PCB. In this instance, the worst-case impedance occurs when the PCB is on the high side of its impedance tolerance range. For this condition, if the parallel termination is this value, there will be no reflection. When the PCBs are at the low side of the impedance range, the mismatch at the termination will be low to high, resulting in overshoot that adds to the incident logic level. If the system impedance is 50 ohms, the parallel terminations should be 55 ohms.

> The value of the parallel termination resistor should be on the high side of the PCB impedance tolerance in order to insure that reflections at the TL terminator interface create a condition of overshoot but not undershoot.

Parallel terminations should be located after the last load. There are occasions, such as large buses that terminate on a BGA or PGA, where the problem illustrated in Figure 21.5 will occur. In the top circuit, the signal passes under the BGA in order to reach the signal pin. Then, it must travel back out to the terminator pin. The result is two wires under the part for every signal. This often results in a routing problem. The lower circuit illustrates stopping at the terminator first. Now, only one wire must be routed under the BGA for each net. This is often called terminator end swapping. Is this acceptable and, if so, when?

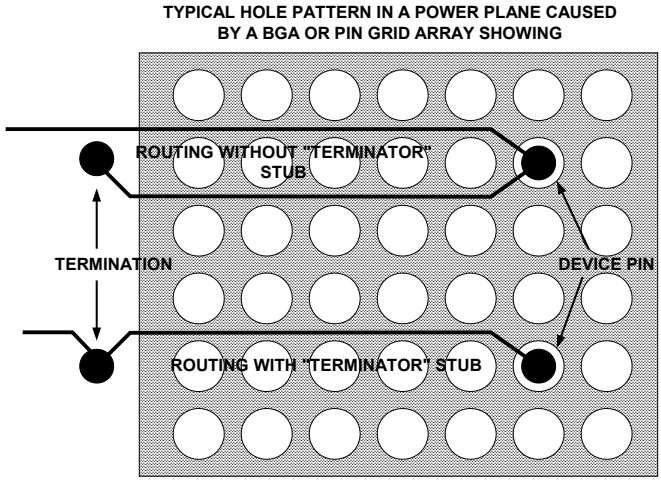

Figure 21.5. A Parallel Terminated Signal Terminating on A BGA

It can be seen from Figure 21.5 that terminator end swapping places the input, or load, at the end of a stub or short transmission line. This results in two things. First, from the termination on, the impedance is half of Z_0. Second, the load is at the open end of a transmission line. If this segment is long enough, voltage doubling will occur. In order to determine how long the "terminator" stub can be, it is necessary to simulate the circuit, slowly increase the length of this segment and observe the quality of the signal at the load. Once the length at which degradation becomes unacceptable is determined, this value can be given to the PCB layout staff for use in placing parallel terminating resistors.

TERMINATING VARIOUS NET TYPES

When using terminations to control reflections, an assortment of net types will be encountered. In all cases, it is desirable to avoid placing loads on side branches or stubs. (Stubs and what they do to signals will be explored in Chapter 22.) Therefore, the points in a parallel-terminated net need to be arranged in a serial string, starting with the driver and ending with the termination. In a series terminated net, the points in the net need to be arranged so that they start at the driver and pass through the loads in a serial string. Figure 21.6 shows the most common parallel-terminated nets.

The top net in Figure 21.6 is the default arrangement of the points in a parallel-terminated net. In some cases, the time delay traveling across a large array of loads in this manner is unacceptable. In this case, the second net topology, splitting the array and placing the driver in the middle, may be used. Memory arrays are a common place where this topology is used. When this topology is chosen, it is important to recognize that it is comprised of one net, but has two independent

transmission lines, each of which must be properly terminated. This results in the driver seeing a load of $Z_0/2$, a double terminated net.

D - L - L - L - L - T

A STANDARD SINGLE ENDED NET

T - L - L - L - L - D - L - L - L - L - T

A MEMORY ARRAY

T - D - L - L - L - L - L - L - D - T

A BUS DRIVEN AT BOTH ENDS

C - L - L - L - T

A CONNECTOR NET DRIVEN FROM OFF CARD

C - L - L - L - D

A CONNECTOR NET DRIVING OFF CARD

WHERE: T = TERMINATOR, L = LOAD, D = DRIVER, C = CONNECTOR

Figure 21.6. Typical Parallel Terminated Nets

In this case, the driver must support two transmission lines each of impedance Z_0. Recalling that the logic 1 of a parallel terminated net is the bench voltage created by Z_{out} and Z_0, the driver will need a much lower output impedance in order to insure a proper logic 1.

Some nets have more than one driver or bidirectional pins as is depicted in the third example in Figure 21.6. In this example, the signal may travel in either direction along the transmission line. As a result, it is necessary to place a parallel termination at each end of the transmission line. The drivers will have the same problem as the previous net.

Some nets contain connector pins. When this occurs, it is reasonable to assume that the rest of the net is on a different PCB. In this case, termination must be done with care to insure that the proper number of terminations is used. If the portion of the net being considered contains a connector pin and no driver, the connector pin is treated as though it is a driver and the termination is placed after the last load. This is shown in the third example in Figure 21.6. If the portion of the net being considered contains a driver, the connector pin is treated as though it were the terminator. This is shown in the bottom example of Figure 21.6.

Figure 21.7 shows the various types of series terminated nets. As can be seen, there are only two choices.

L - L - L - L - L - ST - D

MEMORY ARRAY

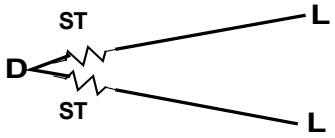

SERIES, TERMINATED SPLIT CLOCK LINE

WHERE: ST = SERIES TERMINATOR, D = DRIVER, L = LOAD

Figure 21.7. Series Terminated Nets

The default arrangement of pins in a series terminated net is shown at the top of Figure 21.7. The series termination is located close to the output pin of the driver. The net then passes through the loads in a serial fashion. It must be remembered that this is a series terminated net. A half amplitude signal will start down the transmission line, double at the open end, reflect back and bring the whole line to full amplitude. Loads between the driver and the last load or open end of the TL will set at an illegal logic state until the reflected wave passes by. Level sensitive inputs, such as clocks, cannot tolerate this condition. Such inputs will need to be placed only at the end of the TL.

If there is more than one such input and they cannot both be connected to the same end, then a “V” net shown in the bottom of Figure 21.7 will be required. Again, this has the effect of loading the driver with two transmission lines. Care will need to be exercised when doing this to insure the bench voltage launched down each TL is at least V/2. The familiar

“star” method of connecting up many clock inputs in this manner will not work, due to the fact that the driver will be overloaded. When more than two clock pins must be driven simultaneously, a fan-out clock buffer will be required.

When driving two branches with a series terminated driver as shown in the bottom of Figure 21.7, there will be a reflected wave returning it the driver from each branch. If the branches are not equal in length, these reflections will arrive at different times and produce interference signals that can degrade both signals. In order to avoid this problem, both branches must be the same electrical length.

Terminating backplane buses using parallel terminations is illustrated in Figure 21.8.

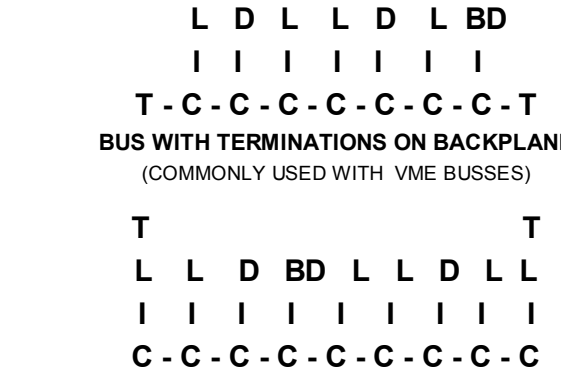

Figure 21.8. Parallel Termination of Backplane Buses

When any card can be plugged into any slot, as shown in the top of Figure 21.8, the drivers can be located anywhere along the length of the bus. As a result, signals can travel either direction on the bus. This forces the terminators to be located on the backplane at each end of the bus. The disadvantage to this approach is that small parts must be assembled onto the backplane. This is typical of a bus such as the VME bus. If the locations of cards is known at all times and there are always cards in the end slots, the termination resistors can be located on the end cards. The advantage is that no small components need to be soldered to the backplane.

Terminating backplane buses using series terminations is illustrated in Figure 21.9. Since the series terminations are always placed at the output of the drivers on the daughter cards, there are no terminations on the backplane. This seems quite simple. However, it must be remembered that there are half amplitude signals starting down each line that will double at the end and reflect back creating data good on the return trip. In order to avoid false triggering from the half amplitude signals, there must be a clocking mechanism in the backplane that supplies clock edges to each slot at the correct time.

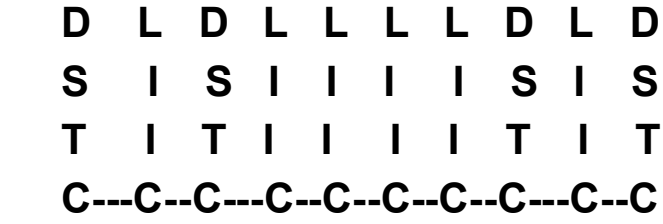

Figure 21.9. Series Termination of Backplane Buses

Buried Resistors

When many terminating resistors are needed, it may be difficult to fit them all onto the surface of the PCB. One solution is to form the resistors in the terminator voltage plane. Figure 21.10 illustrates this technique.

BURIED RESISTORS BUILT INTO ANTIPAD OPENING

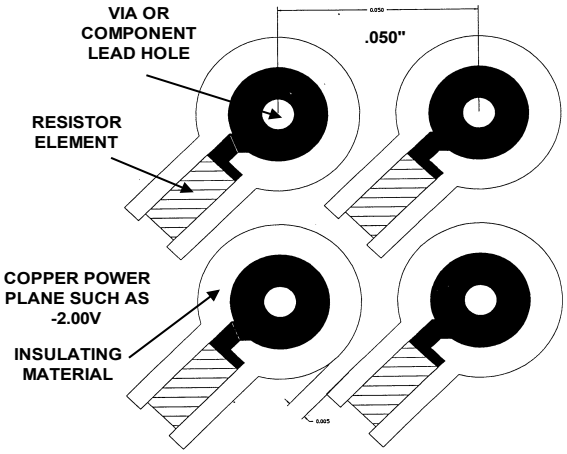

Mahler, Bruce, Ohmega Technologies, Inc. "Planar Resistor Technology for High Speed Multilayer PCBs." Electronic Packaging and Production. January 1986

Mahler, Bruce and Signer, Roy M., "Thin Film Resistor Technology." Printed Circuit Design. June 1986.

Figure 21.10. Buried Resistors formed in the Vtt Plane

The pattern shown in Figure 21.10 creates parallel terminations between component pins and the Vtt plane of an ECL design. Also shown are two articles that describe how to design and build these resistors. This technique works well for parallel terminations where one end of the resistor connects to a component pin and the other to the Vtt plane. It does not work well for series terminations where both ends must connect to nets in the PCB.

This solution has two drawbacks. The first is resistor accuracy. When the size of the resistor element is small, as in Figure 21.10, the tolerance approaches ±20%/. The second drawback is cost. Turning a power layer into a buried resistor layer in a PCB will add as much as $200 to the cost of the PCB panel from which the PCB is built. At this cost, buried resistors will always cost more than their discrete counterparts. Why use buried resistors, then? The original and only reason is when there is not enough space on the surface of the PCB to use discrete resistors.

At this point, some observations about termination accuracy are in order. It is possible to purchase discrete termination resistors, either as chip resistors or as R-pack, that have an accuracy of ±1% at no special price premium. The manufacturing process used to fabricate PCBs cannot sustain impedance accuracies tighter than ±10% day in and day out. This is due mainly to the limitation on trace width etching accuracy and laminate thickness tolerances. To specify a PCB impedance accuracy tighter than ±10% will incur premium charges. These premium charges will cover the cost of discarding PCBs that have an impedance outside the tolerance range of the manufacturing process. It is advisable to create a high-speed design rule set that can operate successfully with this ± 10% accuracy. In fact, when it is time to create a design rule set, the ±10% impedance tolerance of normal PCB production will be the starting point. All logic families designed for high-speed service can be made to work properly with this tolerance.

> The normal tolerance on impedance of the PCB manufacturing process is ±10%. High-speed design rules must allow for this tolerance.

What happens when it is not possible to terminate a transmission line?

In the preceding chapters, the effects of drivers and loads not matching the transmission lines has been discussed. And, a variety of reflections at impedance mismatches have been shown. The value of terminators as a method for controlling these reflections has also been shown. What happens when terminations are not used? The following list details some of the things that might occur:

- 100% positive reflection will occur at the open end of the TL, doubling the amplitude
- The reflected energy will travel back to the driver and reflect again, as much as 80%
- Loads near the driver will be false triggered by these reflections

As a result of the foregoing, it's important to remember the following guidelines:

- Level sensitive loads must be grouped at the open end of the TL
- If level sensitive loads cannot be grouped at the open end, they will need their own net
- Focusing on impedance control when there are no terminations is not useful
- If possible, slow down the edges to minimize reflections or
- Keep net lengths less than ¼ TEL

Sequencing or Scheduling

Once the impedance management strategy has been devised for each net on a PCB, the next step is to order the points in those nets in such a way that the EM fields that travel on a transmission line are well organized. The name for this step in the process is "sequencing" or "scheduling." This operation cannot be done until after the components have been placed on the PCB surface and the locations of all the points in each net are known. Some CAD systems contain routines that will perform the sequencing or scheduling automatically. These automatic tools contain some default sequencing rules. In order to ensure that the automatic sequencing rules are appropriate for a given design, it is advisable to check those rules prior to using them.

For designs that are done on a toolset that does not do automatic sequencing or scheduling, it is worth looking at ways to reduce the size of the manual sequencing operation. The two metrics that can be used to reduce the number of nets in a design and must be examined include:

- Nets with only two pins will automatically be sequenced correctly
- Nets that are shorter than a TEL will usually not need any special sequencing.

This reduces the manual sequencing job to nets of three or more points that are "long."

Sequencing Rules

The rules to be followed for sequencing high-speed nets to create well-designed transmission lines include:

1. Drivers/sources should always be at one end of the net
2. Parallel terminations should always be at the opposite end of the net from the driver
3. Series terminations should always be at the driver end of the net as close as possible to the driver (see Figure 21.3)
4. When multiple drivers or bidirectional drivers are used and parallel terminations are being used, two possibilities exist: If all the drives are clustered together, they are treated using the first two rules. If they are spread along the length of the net, it will be necessary to place a termination at each end of the net. (Deciding what constitutes "close together" can only be done by simulating the proposed layout and observing the quality of the waveform).
5. When multiple drivers or bidirectional drivers are used with parallel terminations, a termination resistor must be placed at each end of the net.
6. In all of the cases above, the loads are connected in a serial string to avoid stubs.
7. On unterminated nets, edge sensitive loads must be located at the end of the net farthest from the driver
8. For parallel-terminated technologies with nets that contain both a connector pin and a driver, treat the connector pin as though it is a parallel termination and rules 1 and 2 are followed.
9. For parallel-terminated technologies with nets that contain a connector pin and no driver, treat the connector pin as if it is a driver and place a parallel termination at the end of the net away from the connector pin.
10. For series-terminated technologies with nets that contain a connector pin and a driver, place the series termination at the output of the driver and treat the connector pin as the last load.
11. For series-terminated technologies with nets that contain a connector pin and no driver, simply connect the loads and connector pin in a serial string and use no terminations.
12. Stubs should be avoided at all times.
13. When stubs cannot be avoided, their length should be kept short enough that their effect on signal quality is acceptable. (Before good SI tools existed, the maximum length was estimated by calculating the highest frequency in the signal and insuring that the stub length was less than 1/8 wavelength at this frequency.)

Forced Sequencing

There are nets with topologies complex enough that it is not possible to sequence their points using automatic tools. In this case, the designer needs to be able to specify the order of connection and freeze it. This operation is called “forced sequencing”. Large memory arrays are examples of this kind of net. In order to determine an order that will result in proper signal quality, the possible ways in which the pins could be connected are analyzed using an SI tool. Once a topology and order of connection is determined that produces the proper quality signal, this order of connection is given to the PCB designer. Figure 21.11 illustrates a net where the sequence must be forced in order to insure that the pins are connected in the proper order.

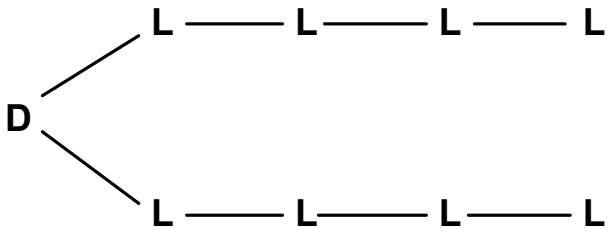

Figure 21.11. A Forced Sequenced Net

Schematic tools designed for high speed PCB layout often contain fields in the net list that allow the sequencing of a forced sequenced net to be locked down. In this way, the sequencing information is not lost when PCB placement or other operations are done. This “forced sequencing” methodology is tool specific, so a designer should consult the tool provider for instructions on how to use the forced sequencing feature. Auto sequencing is also tool specific, so the same advice applies. Designers who do not have a tool with these features as part of the tool set will need to be vigilant as layout proceeds to insure all nets, when routed, have the proper sequence.

CHAPTER 22: STUBS ON TRANSMISSION LINES?

In an earlier chapter, stubs along a net were mentioned as having an adverse effect on the quality of high-speed signals. As shown in Figure 17.7, stubs or branches off the main line of a signal path can cause waveform reversals. Figure 22.1 illustrates what a stub means.

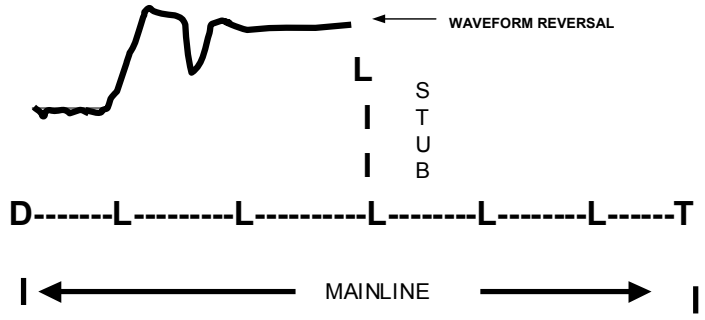

DEFINITION: THE MAINLINE IS THE LONGEST PATH FROM THE DRIVER TO THE FARTHEST LOAD. A STUB IS ANY BRANCH OFF THE MAINLINE.

Figure 22.1. A Net with a Stub or Branch off the Mainline

Figure 22.2 illustrates what happens when a stub whose length is ¼ wavelength at some frequency is excited by a sine wave at that frequency.

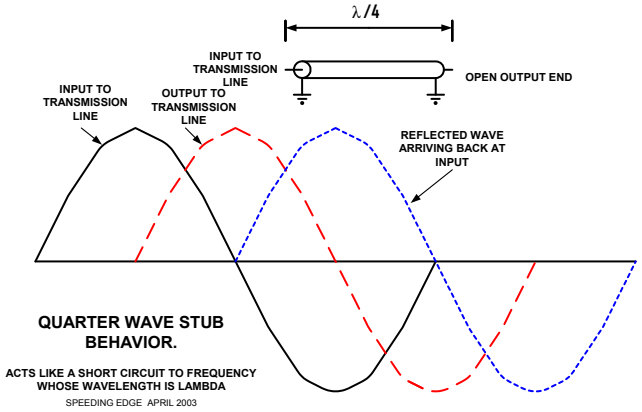

Figure 22.2. A Quarter Wave Stub Excited by a Sine Wave of Frequency at a Quarter Wavelength

The sine wave signal is launched into the transmission line. A quarter wavelength later (or 90° into the cycle), it arrives at the open end. There is no absorber or termination there, so all of the energy is reflected back to the source without being inverted. A quarter wavelength later (or 180° into the cycle), it arrives back at the source, exactly 180 degrees out of phase with the original signal, canceling it out. The quarter wave stub has "shorted out" the signal. This shorting out will occur at all frequencies at which the stub is an odd multiple of a quarter wave.

Digital data paths can see data traffic that ranges from DC to the highest clock rate for the data path, so the data path must not contain stubs that can short out any of those frequencies. In the example shown in Figure 17.7, the energy that was being "shorted out" was not the clock frequency. Instead, it was the frequencies that made up the switching edge. These are much higher than the clock frequency. In order to guarantee that stubs don't adversely affect the switching edges, it is necessary to determine the frequency content of all the switching edges that will travel on the TL. This analysis is usually done with a Fourier Transform, a mathematical operation that converts time domain signals to frequency domain signals and vice versa. Figure 22.3 is a graph showing the first harmonic of an edge of a given rise time and, from that frequency, the length of a quarter wave stub in FR-4 PCB material.

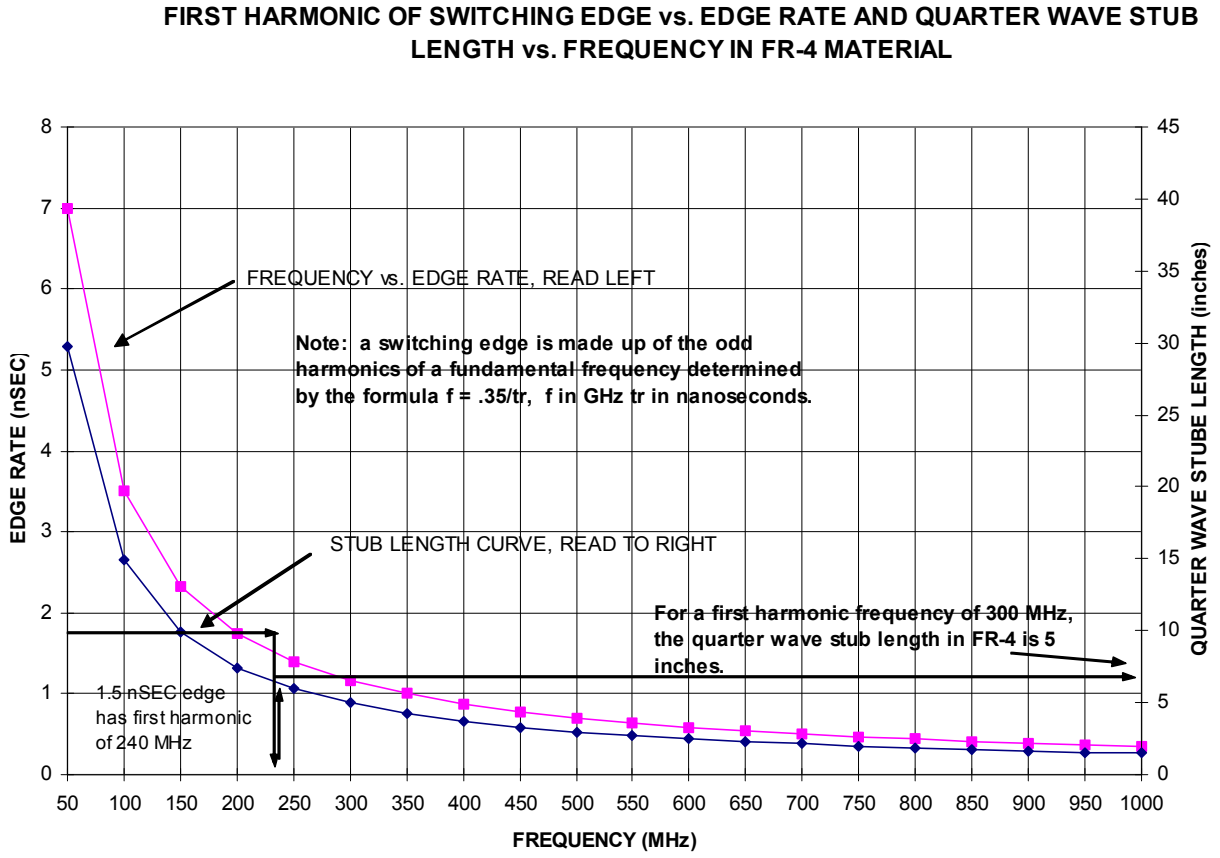

Figure 22.3. Edge Rate vs. First Harmonic and Quarter Wave Stub Length in FR-4

An approximation formula was used to calculate the first harmonic. The equation in the middle of Figure 22.3 was used for this calculation. For more precise calculations of frequency content of a switching edge, a Fourier Transform should be used. The approximation provides a sense for the frequencies involved in fast edges. The example shown is for a 1.5 nSEC edge. Notice that the first harmonic of this edge is about 240 MHz. A quarter wave stub at that frequency is approximately 6 inches and an eighth wave stub is about 3 inches. Thus, for circuits with 1.5 nSEC edges, stubs must be kept to less than 3 inches in length.

Most modern logic components have edges faster than 0.5 nSEC. If the analysis is done for that edge, the first harmonic is about 720 MHz and a quarter wave stub is 2 inches! Logic circuits with random data patterns and 0.5 nSEC edges can have frequencies anywhere from DC to well beyond 750 MHz. The band pass of the logic path must be flat over this entire range. At no frequency in this entire range can stubs be allowed that could “short out” the signal as it switches. That means data paths must be exceptionally broadband, more broadband than virtually any RF circuit!

As switching edges fall below 300 pSEC, the length of the component leads reaching from the PCB up into the IC package are often so long that they create stubs that adversely affect signal quality. That is the main reason GB/S and higher data paths are not multi-drop. The path up into a circuit in the middle of a net produces an excessively long stub. Such nets must be comprised of only two points in order to avoid stub-related failures.

From the brief introduction to the frequency content of switching waveforms in the above discussion it can be seen that digital signals have such high frequency components that EMI scans show frequencies over a very broadband. Most of these frequencies are a function of how long the transmission lines are and how fast the switching edges are. Very few of them are related to the clock frequency, as was the case when logic was much slower.

CHAPTER 23: PROPERTIES OF TRANSMISSION LINES THAT AFFECT IMPEDANCE

In Chapter 16, Equation 16.2 stated that the impedance of a transmission line is determined by the capacitance and inductance that is distributed along the length of the transmission line. For convenience, it is repeated here.

$$Z_O = \sqrt{\frac{L_O}{C_O}}$$

Equation 23.1. "The Impedance Equation"

From this equation it can be seen that changing either L_0 or $C_{0,}$ will change the impedance of a transmission line. In Chapter 17, it was demonstrated that signal reflections are caused by changes in impedance. For convenience, the reflection equation is repeated here.

$$\%= 100\,\frac{Z_I - Z_O}{Z_I + Z_O}$$

Equation 23.2. "The Reflection Equation"

From Equation 23.1, it is not obvious what variables will have an effect on impedance. Equation 23.3 is the classic surface microstrip equation. It illustrates the variables in a PCB that determine impedance. This equation is shown for illustration purposes only, so that those variables can be examined. Later, it will be shown that this equation, as well as all other equations used to calculate impedance, have a limited range over which they are valid. More accurate methods are available and will be discussed in detail in subsequent sections of this book.

e_r = RELATIVE DIELECTRIC CONSTANT
H = HEIGHT OF TRACE ABOVE PLANE
W = TRACE WIDTH
T = TRACE THICKNESS
Z_O = TRACE IMPEDANCE IN OHMS
ANY DIMENSION SYSTEM IS APPLICABLE
NOTE: VALID FOR 5<w<15 MILS

A more precise calculation can be obtained using a 2D field solver which the author recommends.

$$Z_0 = \frac{79}{\sqrt{e_r + 1.41}} \ln\left(\frac{5.98\,H}{0.8W + T}\right)$$

e_r value is that obtained from velocity measurements made with a TDR.

Equation 23.3. A Surface Microstrip Transmission Line Impedance Equation

Equation 23.3 shows that there are four variables that determine the impedance of a transmission line on a surface layer of a PCB. These include: height of the trace above the plane over which it travels; the width of the trace; the thickness of the trace and the relative dielectric constant of the insulating material used to support the trace. Once the variables that affect impedance are known, it is possible to determine what features in a PCB will have a meaningful affect on impedance. Table 23.1 lists these.

Changes in trace width in the same layer (usually caused by trace necking)
Changes in trace thickness
Changes in height above the plane
Stubs along the transmission line
Loads along the transmission line
Connector transitions
Poorly matched terminations
No terminations
Large power plane discontinuities
Changes in relative dielectric constant

Table 23.1. Sources of Impedance Change or Mismatch

Notice that right angle bends and vias are not on this list, even though there is a popular notion that these are sources of impedance mismatches. Neither of these is a significant source of impedance mismatch. This is why they are not on this list and, later in this book, it will be shown why this is the case.

The common characteristic of the items in Table 23.1 is that they can have a measurable effect on one or both of the variables in Equation 23.1, L_o and C_o. On further examination, it is possible to determine which variable is affected by each source. Table 23.2 shows which variable is affected.

Change in trace width- C_o
Change is trace thickness- C_o
Change in trace height- C_o
Stubs along the line- C_o
Loads along the line- C_o
Connectors- C_o
Large power plane discontinuities- C_o
Relative dielectric constant – C_o

Table 23.2. Variables in Impedance Equation Impacted by Sources in Table 23.1

Notice that, with the exception of terminator mismatches, all of the sources of impedance mismatch are caused by something that changed C_o. Within the limits of trace dimensions in PCBs, compared to changes in C_o, L_o is relatively constant. Knowing this helps in troubleshooting impedance problems or in designing controlled impedance signal paths.

The two main sources of impedance mismatch and reflections are:

Poorly matched terminations

And

Changes in C_o

Virtually all impedance changes along the length of a transmission line stem from changes in C_o. Understanding what can change TL capacitance and learning to manage those changes is the secret to good impedance control.

Table 23.3 shows the range of relative dielectric constant that occurs in the laminate known commonly as FR-4. Not only does the relative dielectric constant change with frequency as illustrated in Chapter 16, it also varies with the amount of glass and resin used to make the laminate. Notice that there are four ways to make a 4-mil thick piece of laminate, three ways to make a 5-mil thick piece of laminate and four ways to make a 6-mil piece of laminate. Also note that the ratio of glass to resin is different in each of these formulations as is the relative dielectric constant. If a PCB stackup is designed to use one of these formulations and the fabricator uses one of the others, the impedance will not come out as expected. This is the most common reason that changing fabricators results in PCBs with different characteristics. To avoid this kind of problem, it is necessary to specify, on the fabrication drawing, which laminate formulation is required in each opening in the stackup.

To avoid fabricator to fabricator variations in impedance of a PCB, it is imperative to specify the laminate to be used in each dielectric opening of the stackup.

Thickness	Construction	Resin Content	e_r @ 1 MHz	e_r @ 1 GHz
.002	1 x 106	69.0%	3.84	3.63
.003	1 x 1080	62.0%	4.00	3.80
.004	1 x 2113	54.4%	4.19	4.00
.004	1 x 106 + 1 x 1080	57.7%	4.11	3.91
.004	1 x 2116	43.0%	4.54	4.37
.005	1 x 106 + 1 x 2113	52.8%	4.24	4.05
005	1 x 2116	51.8%	4.26	4.08
.006	1 x 1080 + 1 x	52.2%	4.25	4.06
.006	1 x 106 + 1 x 2116	50.8%	4.29	4.11
.006	2 x 2113	43.5%	4.52	4.35
.007	2 x 2113	49.6%	4.33	4.14
.008	1 x 7628	44.4%	4.49	4.32
.010	2 x 2116	51.8%	4.26	4.08
.014	2 x 7628	38.8%	4.69	4.53

Table 23.3. Typical e_r values for Standard FR-4 Laminate

(Data courtesy of ParkNelco)

CHAPTER 24: METHODS FOR CALCULATING AND MEASURING IMPEDANCE

This chapter deals with the design and measurement of transmission lines in PCBs. There are four basic types of transmission lines in PCBs. These are shown in Figure 24.1.

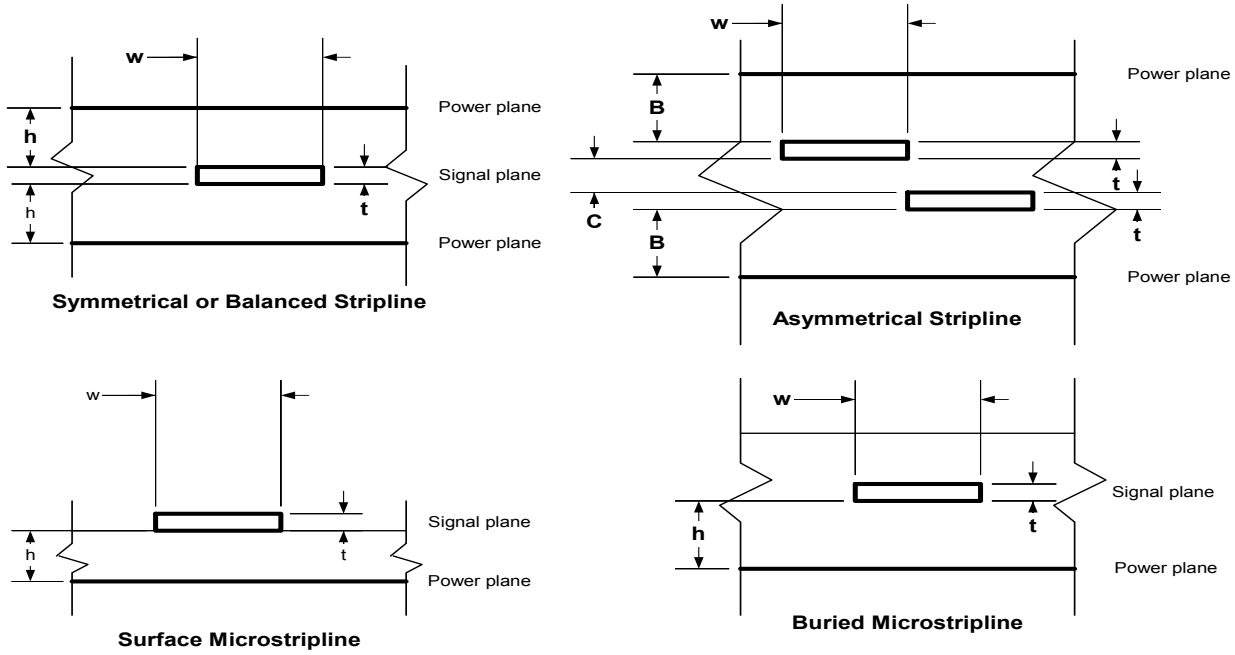

Figure 24.1. Four Major Types of PCB Transmission Lines

Two of these are “stripline” and two are “micro-stripline”. Stripline refers to a transmission line that has two planes as partners. Micro-stripline refers to a transmission line that has only one plane as a partner.

Symmetrical stripline is the most ideal of the four transmission lines shown here. However, in order to build PCBs with symmetrical striplines, it is necessary to add a plane every time a signal layer is added. As a result, the layer count goes up rapidly adding extra cost and extra thickness. Placing two signal layers between each pair of planes, asymmetric stripline, helps solve this problem. The quality of these transmission lines is more than good enough for any transmission lines that are encountered in logic designs. The only potential risk to signal integrity occurs if a transmission line in one layer runs directly over the top of another in the other layer. In this case, crosstalk between the two signals will be excessive, resulting in failures. This is best handled by orthogonal routing that involves routing signals in one layer in the X direction and in the other layer in the Y direction.

Micro-striplines can either be placed on the surface or buried in the dielectric. As will be seen in the chapter that addresses PCB fabrication processes, it is more difficult to control trace width and impedance on outer layers than on buried layers. As a result, of these two microstripline options, buried microstrip layers are the preferred choice for transmission lines.

Differentially coupled trace pairs (differential impedance) have been left out of this discussion because, in spite of popularly held beliefs, differential impedance and tightly coupled lines are not a beneficial method for handling differential signals. The reason for this will be explored in depth in Chapter 31.

PCB stackups are created by using combinations of the four types of transmission lines. The toolset that creates the stackups then uses Maxwell’s Equations to exactly calculate L_o and C_o. Once these two variables are known, Equation 16.2 is used to compute Z_o.

The advantage of the equations is that they are free and relatively easy to use. The disadvantage of them is that they all only provide partial solutions. The equations were created by taking actual measurements of real PCBs and then curve fitting the variables to produce the desired result. Within their “sweet spot” they produce reasonable results. Outside the

sweet spot, results may vary widely. Equation 23.3 is the most common equation used for surface microstrip transmission line impedance calculations. Equation 24.1 is one of the equations used to calculate the impedance of buried microstrip transmission lines. This equation was developed by Martin Marietta in the mid 1980s to provide a method for predicting the impedance of buried microstrip transmission lines. Notice that there is no dimension to the surface of the PCB. It was discovered that when the TL is submerged in the dielectric 5 mils or more, nearly all of the field lines are confined in the dielectric and the influence of the air is nil. This will be demonstrated in a later graph by observing the velocity of signal travel in a buried microstrip layer as compared to one on the surface or one in a strip line layer. There are two ways to calculate the impedance of each of the layers in a stackup. The first involves using the traditional equations that have evolved over time and the second involves using a 2D (two dimensional) field solver. A 2D field solver is a mathematical tool that allows the user to describe a transmission line of any geometry with any combination of dielectric constants, height of the line above the power plane, trace width and trace thickness.

Z_O = TRANSMISSION LINE IMPEDANCE (OHMS)

H = HEIGHT OF LINE ABOVE POWER PLANE

W = TRACE WIDTH

T = TRACE THICKNESS

e_r = RELATIVE DIELECTRIC CONSTANT

A more precise calculation can be obtained using a 2D field solver which the author recommends.

Valid for 5 < W < 15 mils, valid for any dimension system

Assumes at least 5 mils of dielectric lying on top of trace.

$$Z_0 = \left(43.037 \ln \frac{H}{W}\right) + 5.048\left(\frac{T}{W}\right) + \frac{106.76}{1.09\sqrt{e_r}}$$

Equation 24.1. An Equation for Calculating the Impedance of Buried Microstrip Transmission Lines

Equation 24.2 is an equation for calculating the impedance of centered or asymmetric stripline transmission lines. The origin of this equation is Digital Equipment Corporation and it was also derived in the mid 1980s.

- **Z_0 = TRANSMISSION LINE IMPEDANCE**
- **B = TRACE TO PLANE SPACING**
- **C = TRACE PLANE TO TRACE PLANE SPACING**
- **T = TRACE THICKNESS**
- **W = TRACE WIDTH**
- **e_r = relative dielectric constant of insulator**
- **FOR C = 0, equation applies to centered stripline**
- **Valid for 5 < W < 15 mils**

A more precise calculation can be obtained using a 2D field solver which the author recommends.

$$Z_0 = 80\left[\frac{1-\frac{B}{4(B+C+T)}}{\sqrt{e_r}}\right]\ln\left[\frac{1.9(2B+T)}{(0.8W+T)}\right]$$

Equation 24.2. An Equation for Calculating the Impedance of Stripline Transmission Lines

All if these equations are based on the dimensions in Figure 24.1.

One of the reasons that these equations are valid only over a limited range is because of the assumption illustrated in Figure 24.2. The capacitance that makes up Co has two components: the capacitance formed between the underside of the trace and the plane, Cpp, and the capacitance formed by the fringing fields, Cfringe. The assumption that is often

made in these equations is that the ratio of these two capacitances is constant. Clearly, this cannot be true. When the trace is very narrow, Cpp gets very small, while Cfringe remains constant. When the trace gets very wide, Cpp gets large while Cfringe remains constant. At both extremes, the accuracy of the equation is compromised.

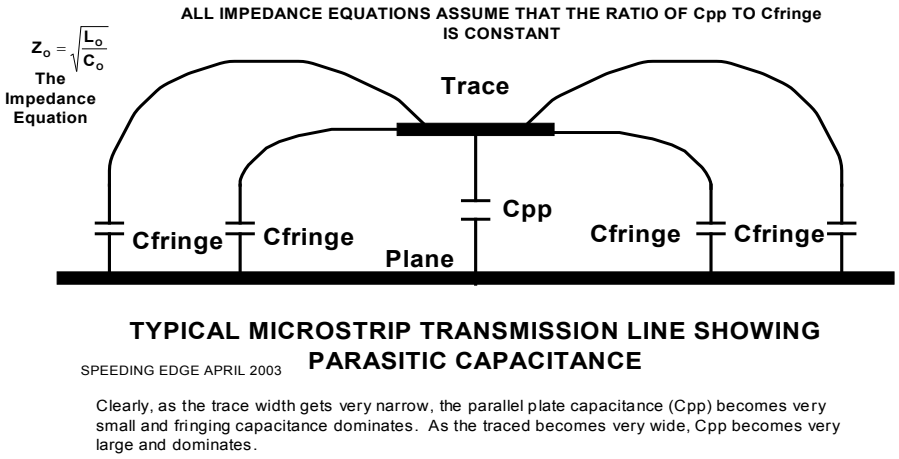

Figure 24.2. Fringing Capacitance vs. Parallel Plate Capacitance for a Microstrip Transmission Line

Comparison of Field Solver to Equations

Figure 24.3 is a plot of impedance versus trace width for the three most common types of transmission lines in PCBs--surface microstrip (SMS), buried microstrip (BMS) and centered stripline (CSL). All transmission lines are 1 ounce thick (1.4 mils) and are 5 mils above the nearest plane with a dielectric whose relative dielectric constant is 4. The only variable is trace width.

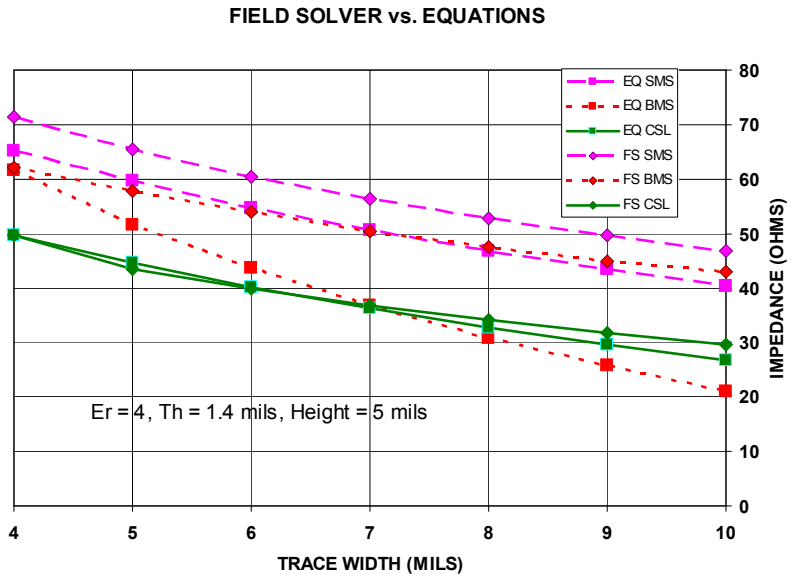

SMS = Surface microstrip, EMS = embedded microstrip, CSL = centered stripline

Figure 24.3. Comparison of Impedance Predicted by Equations vs. 2D Field Solver

The three curves with the square markers are the impedances predicted by the three impedance equations presented earlier. The three curves with the diamond markers are the impedances predicted by a 2D field solver. As can be seen for the centered stripline case, the 2D field solver and the equation results (Equation 24.2) are very close to each other. For the two microstrip cases, there is a significant difference.

Which is right? It has been shown in many papers that 2D field solvers predict impedance that agrees within the measurement accuracy of the test equipment used to measure impedance. It has also been shown that equations

commonly predict the wrong impedance. Because of this equation error problem, the users of equations have often resorted to iterative, trial-and-error methods for adjusting the equations to allow them to predict the correct impedance. Some PCB fabricators have gotten quite good at "fudging" the values of e_r used in their calculations in order to get good answers from the equations. The problem with this "fudging" is it requires considerable trial and error with new materials until they are "dialed in." Even if there is time to "dial in" a new material by manipulating the e_r values to arrive at accurate impedance calculations, it leaves the PCB engineer not knowing the true e_r of the material. This is needed in order to accurately calculate velocity for the purpose of predicting the time of flight along traces.

Even with a good impedance predicting tool, such as a 2D field solver, impedances are still often calculated wrong. This is because the e_r value given for a particular laminate may be incorrect. The most common published values for e_r are measured at 1 MHz. Figure 16.2 shows that the relative dielectric constant of all common PCB materials decreases with frequency. It also shows that the relative dielectric constant varies with the ratio of glass to resin in the material. Therefore, getting the impedance calculation correct requires using a good tool; knowing the glass to resin ratio of the laminate and knowing the frequency at which the transmission line will be used. Once all three of these factors have been taken into account, the results are predictable and repeatable.

It follows that once the material makeup of the PCB has been established, (i.e., the glass to resin ratio and laminate type), any deviation from this during fabrication may result in incorrect impedances. Therefore, the stackup on the fabrication drawing must list all of these parameters and they must not be changed without first revisiting the impedance calculations.

What Frequency Should be Used to Calculate and Measure Impedance?

It has been established that impedance in a PCB varies with the frequency at which the measurement is made. This brings up the question, what frequency should be used in the design process? Most transmission lines in PCBs are intended for logic signals. Logic signals are not a single frequency. Rather, they are a collection of rising and falling edges and it is these rising and falling edges that can potentially become sources of reflections if impedance is not properly managed. How does one get from edge to frequency? It is possible to perform a mathematical operation on a switching edge of some rise time, tr, called a Fourier transform. This will convert the switching edge from the time domain to the frequency domain, yielding the frequencies of interest for the impedance calculation. These will be a spectrum of frequencies starting at what is often called the first harmonic of the switching edge along with the odd harmonics of this frequency. Most of the energy of interest will be in the first harmonic and it is this frequency that will be used to calculate impedance. Equation 24.3 yields a reasonable approximation of this first harmonic.

$$\mathbf{f} = 0.35 / \mathbf{tr}$$

Where f = frequency in GHz and tr = rise time in nanoseconds.

Equation 24.3. An Equation for Approximating the First Harmonic of a Switching Edge

From this equation, it can be seen that the first harmonic of a 350 pSec edge is about 1 GHz. This is very far away from the 1 MHz value used to characterize many laminates. The fastest edge that will travel on the transmission line should be used to calculate impedance.

> The fastest edge that will travel on a transmission line must be used to calculate impedance.

Why 50 ohms?

The question is often asked, why do we use 50 ohms for the impedance of PCBs? Why not use 60 or 75 ohms? The answer lies in the limitations of materials and processes and the layer requirements for PCBs. Figure 24.4 shows the impedance achieved with a 5-mil wide trace in the three types of transmission lines--surface microstrip, buried microstrip and stripline. The 5-mil trace width has been chosen as the narrowest trace width that is practical in volume manufacturing. The height above the nearest plane has been increased from 5 mils to 20 mils to see how high the impedance could be on each of these three layer types.

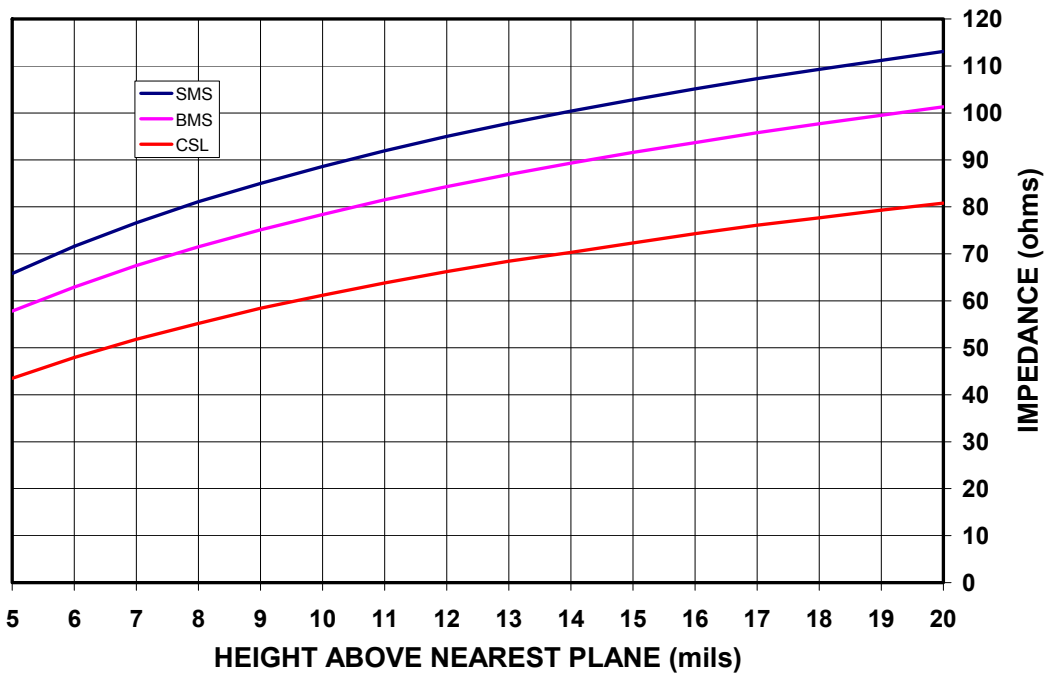

Even with the narrowest production trace width, stripline layers cannot achieve high impedances. In all cases, high impedances require very thick dielectrics, making PCBs excessively thick, as well, and subject to severe crosstalk.

Figure 24.4. Impedance vs. Height Above Nearest Plane, 5 mil Trace Width

If the only signal layers in a PCB are surface microstrip, it is possible to have impedances ranging from 65 ohms to 113 ohms. 50 ohms would be possible only if the height above the plane were reduced to less than 5 mils, possible, but on the low end of practical material thickness. As soon as more than two signal layers are needed, the additional layers will have to be buried microstrip or stripline. If buried microstrip is chosen, the possible impedances range from 58 to 101 ohms, so the highest impedance that could be achieved is still quite high. However, adding the buried microstrip layer would force the surface microstrip layer farther away from the plane and it would be difficult, if not impossible, to make the impedance of these two layer types the same.

If four or more routing layers of the same impedance are needed, all but two of them will have to be stripline layers. In this case, the possible impedances range from 42 ohms to about 80 ohms. So, it would be possible to create a PCB with signal layer impedances up to 80 ohms. However, to do this the dielectric thicknesses would have to be very large. This would result in a very thick PCB. As will be seen in a later chapter, the crosstalk would also be excessively high.

Bearing in mind the constraints of crosstalk and reasonable PCB thickness, dielectrics should be kept under 6 mils. With this constraint, it is not possible to achieve impedances higher than about 55 ohms on stripline layers. This will force the impedance on the buried microstrip layers to be the same.

As experience is gained in creating PCB stack ups that balance the needs of manufacturability, crosstalk and impedance against each other, it will become clear that 50 ohms is the best compromise. Other impedances, higher and lower than this, will be difficult to achieve. If a 38-ohm impedance, such as that required by Rambus™, is used, manufacturability will be compromised and if the 65+ ohms required by some protocols is used, crosstalk will be excessive.

Velocity vs. Location in PCB

For economic reasons, it is desirable to use as many signal layers in a PCB for “high speed” signals as possible. It is sometimes asserted that the only layers good enough for such signals are the stripline layers. If this is so, buried microstrip and surface microstrip signal layers are not usable. This means that additional stripline layers will be required with the result that the finished PCB will cost more. In order to use signal layers in a PCB interchangeably, three requirements must be met. These include: the impedance must be the same; the wave velocity must be the same and the

crosstalk must be the same. Figure 24.5 is a plot of signal velocity vs. relative dielectric constant for the usual range of materials used to build PCBs. The left axis is velocity in inches per nanosecond and the right axis is the reciprocal or picoseconds per inch. The two scales allow for easy conversion from relative dielectric constant to either unit of measure. Equation 10.1 was used to create these graphs.

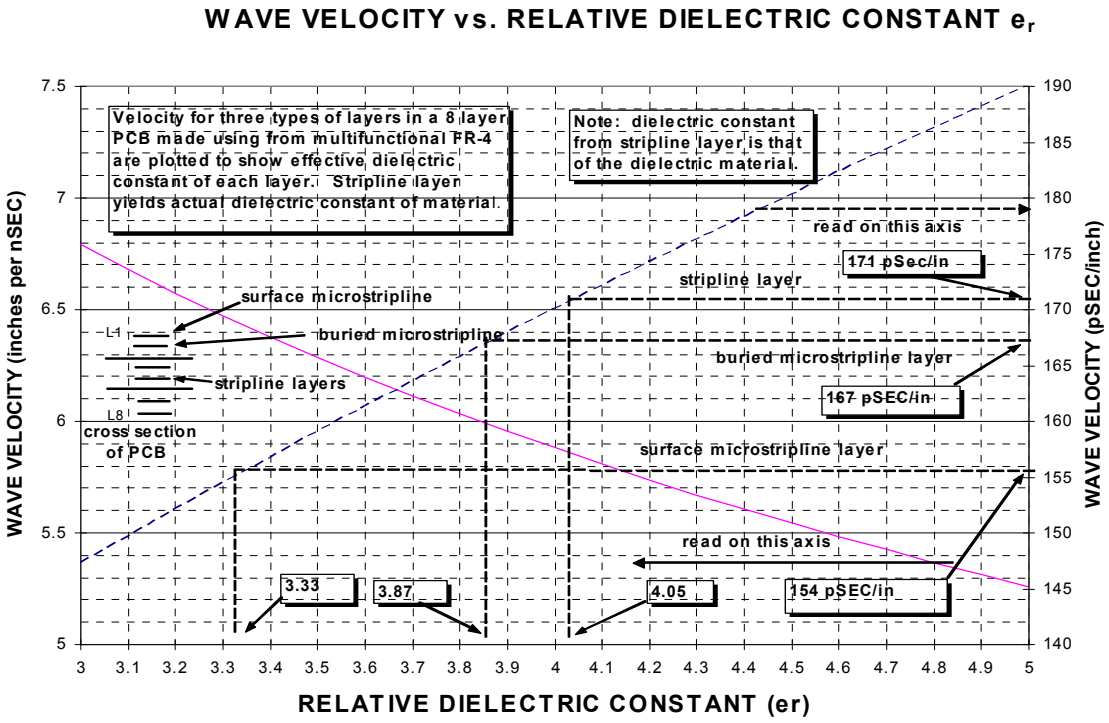

Figure 24.5. Velocity vs. Trace Layer Location in an Eight Layer PCB

On the left hand side of the graph is the stackup of an eight-layer PCB. For those not familiar with this notation, short bars represent signal layers and long bars represent planes. This PCB has two surface microstrip layers, two buried microstrip layers, two stripline layers and two planes. Part of the production test for this PCB involved impedance and velocity for each layer. The velocity of each layer type is plotted. The velocity of the surface microstrip layers was measured at 154 picoseconds per inch; the buried microstrip layers at 167 picoseconds per inch and the stripline layers at 171 picoseconds per inch. The effective dielectric constant for each of these three layer types is 3.33, 3.87 and 4.05, respectively. The velocity on the buried microstrip layer is virtually the same as the stripline layer. They are so close that it would be difficult to measure a difference. One of the goals of interchangeability has been met. These two layers are the same velocity. In order to use them interchangeably, they have to be of the same impedance, which can be done with careful selection of trace width. They also have to have the same level of crosstalk—the subject of which will be addressed in the chapter dealing with crosstalk. As a result of the foregoing, it is possible for buried microstrips to be used interchangeably with stripline layers for “high speed” signals.

Measuring Impedance

Figure 24.6 illustrates the most common method used to measure the impedance of PCB traces. An instrument known as a Time Domain Reflectometer (TDR) is used for this task. A TDR is really two instruments in one. There is a pulse generator that produces signals with very fast rising edges. These signals are sent out on the test cable to the trace under test. Attached to the output port of the TDR is a sampling oscilloscope capable of responding to very fast signal transitions. This sampling oscilloscope monitors the voltage at this point. The screen can display the signal that is sent out of the instrument and all of the reflections that occur from changes in impedance along the structure under test. Reflections from impedance changes in the impedance of a trace under test are measured. The size of the signal sent out is measured and the size and polarity of the signal that is reflected is measured. Knowing that the test cable and the test output are 50 ohms, it is possible to use Equation 17.1 to calculate the impedance of the trace under test.

Most TDR instruments have the Equation 17.1 built into the operating software. All the operator needs to do is place a cursor on the point of interest on the test trace and the instrument makes the measurements, does the calculations and displays the impedance. Some instruments, such at those from Polar Instruments, have additional features that provide a graphic on the screen that show the minimum and maximum allowable impedances. This creates a “go-no go” criterion that makes rapid production testing easy. Another feature of these tools is a data logging capability that produces test records as part of a test report.

Just as accurately calculating impedance requires a statement for the frequency at which the impedance is being calculated, accurately measuring impedance has the same potential for error in terms of the rise time of the test signal. For instance, if the impedance was designed for a 500-picosecond edge and the impedance is measured with a 40-picosecond edge, this will result in two different impedances. Table 24.1 shows how the impedance varies when the same traces are measured with different rise time edges from a TDR.

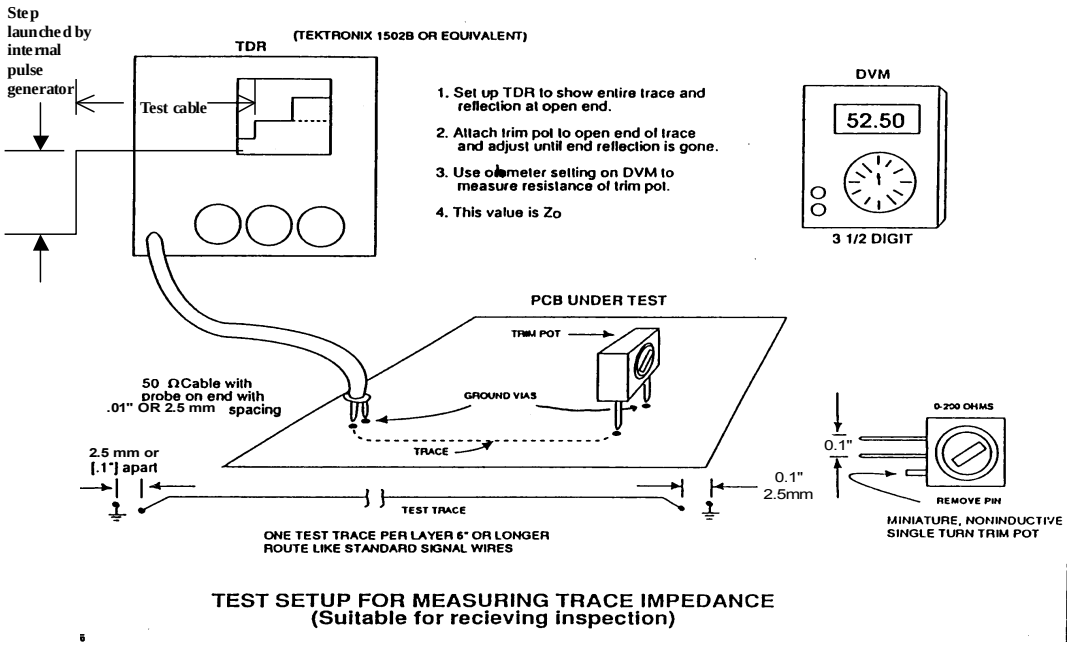

Figure 24.6. A Time Domain Reflectometer Setup for Measuring Impedance

	40 picosecond edge	125 picosecond edge	175 picosecond edge
L6	55.1 ohms	53.9 ohms	53.0 ohms
L8	57.4 ohms	56.5 ohms	52.6 ohms
L11	54.7 ohms	53.2 ohms	52.8 ohms

Table 24.1. Impedance vs. Rise Time of TDR Test Edge

Table 24.1 shows that measured impedance varies significantly as the rise time of the test edge is changed. Specifying the rise time of the edge used to test impedance is as important as all of the other parameters involved in calculating impedance.

It is important to agree on the rise time of the test edge as part of the test specification for a controlled impedance PCB. The rise time chosen must be within the range of rise times that can be produced by the test equipment. Since most of the PCB fabricators manufacturing controlled impedance PCBs have production test equipment with a 175 picosecond rise time, it is recommended that the rise time of the test edge be in the range of 125- to 175-picoseconds. This is not ideal, but it is a reasonable compromise to account for the differences existing between the equipment available to engineers and that which is available to manufacturers.

2D Field Solvers

As noted earlier in this book, it has been suggested that 2D field solvers provide the most accurate method for calculating impedance. A fair question is how can these toolsets be found. All of the SI toolsets that are part of the PCB design systems supplied by the major EDA vendors have 2D field solvers built into them. For those who don't have these PCB design toolsets, such as PCB fabricators, a number of companies sell stand-alone 2D field solvers. These vendors include:

Polar Instruments
Applied Simulation technology
Ansoft

All of these tools can perform the necessary calculations required to arrive at accurate impedances.

CHAPTER 25: RIGHT ANGLE BENDS AND VIAS POTENTIAL SOURCES OF REFLECTIONS AND OTHER PROBLEMS

For some reason, right angle bends and vias have been the source of much concern when it comes to routing PCBs. Each is thought to be a significant source of signal integrity problems including reflections and EMI. Much work has been done to dispel these concerns, yet there is still a group of engineers who won't allow their use. It is odd that the proponents of the "no via" or "no right angle bends" rules have no tests or analyses that demonstrate that their use is a problem, even when not allowing them significantly complicates the PCB layout process. In this chapter, tests will be made to show conclusively whether or not these two features should be of concern. In addressing this topic and to provide additional information, references will be made to other works on the same subject.

A Test PCB

Figure 25.1 depicts layer one of a 16-layer test PCB that contains more than one hundred different test structures intended to measure whether or not various "rules of thumb" are valid.

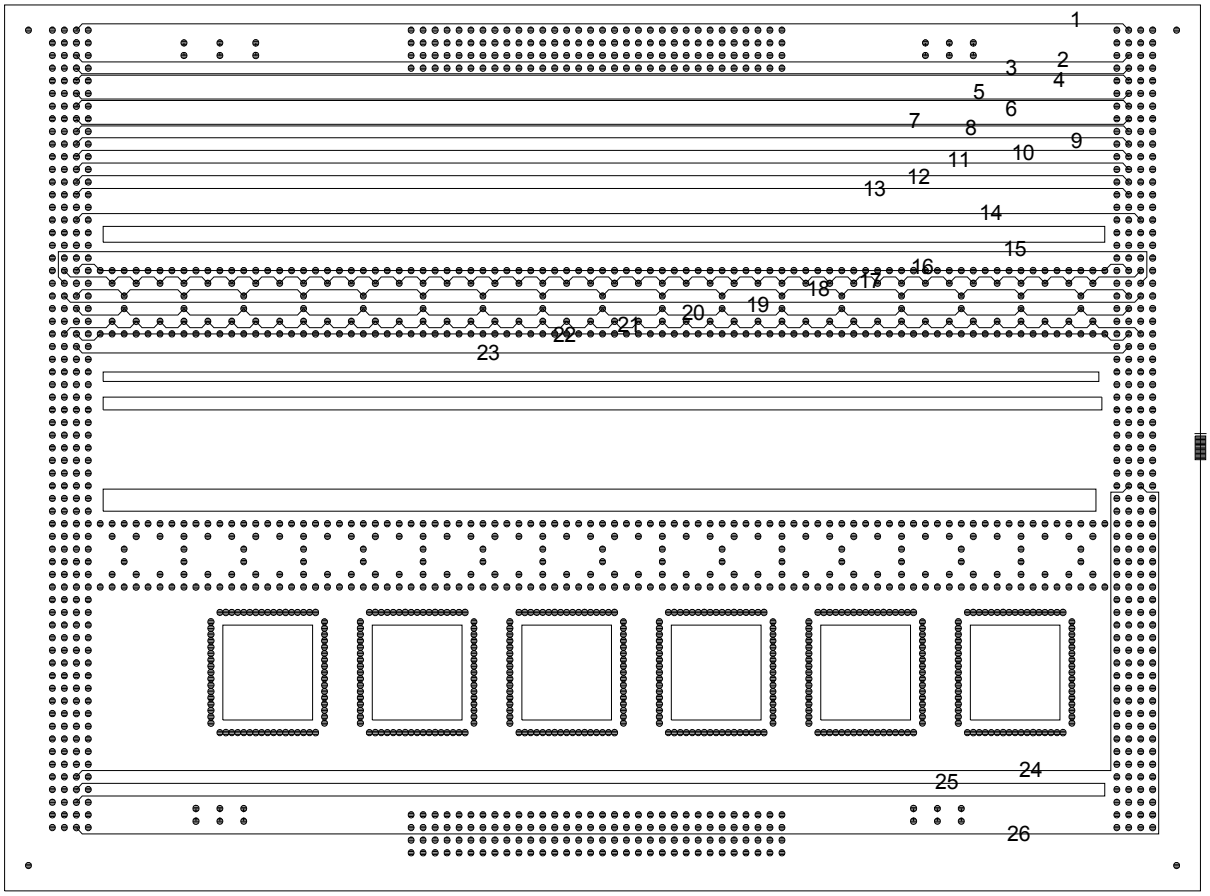

Figure 25.1. Layer 1 Artwork of Test PCB

Right Angle Bends

Trace 15 has a right angle bend near each end. Traces 24 and 26 have right angle bends in them in their lower right hand corners. To test these three traces, a TDR with a rising edge of 125 picoseconds is used to determine if these right angle bends cause any changes in impedance. If there is any meaningful degradation from these bends, there should be a detectable change in impedance that is visible on the TDR trace. A lowering of impedance would be depicted as a reflection that is "negative" going. A rise in impedance would be depicted as a reflection that is "positive" going.

Figure 25.2 contains the three TDR plots for the three test traces. The first dip in impedance is where the test cable makes contact with one end of the trace at a via. The second dip in impedance is the test via at the opposite end of the trace. The flat portion between these two dips is the trace under test.

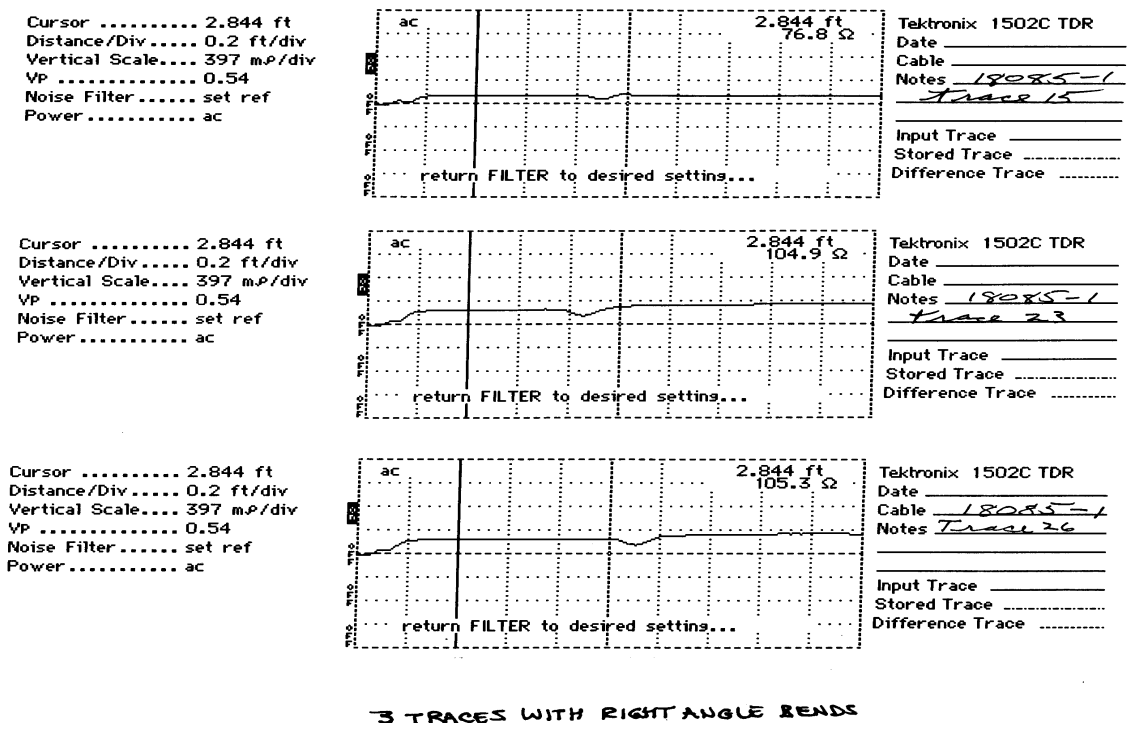

Figure 25.2. TDR Records of Three Traces with Right Angle Bends

From the three tested traces, it can be seen that the right angle bends caused no detectable change in impedance, even with an edge of 125 picoseconds. In a related test PCB, done by a group of engineers and tested at University of Missouri, Rolla, a right angle bend was tested with an 18 picosecond edge and the same results were achieved. In that same test, as reported in an article in Printed Circuit Design Magazine, titled "90° Corners, The Final Turn" by Douglas Brooks, (January 1998), right angle and acute angle bends were tested for reflections and as potential sources of EMI with negative results.

If one closely examines a right angle bend, it can be determined that for a very short distance the effective width of the trace is wider. However, for such a short distance there is no detectable change in either C_o or L_o, the two variables that have to be affected in order for a reflection to occur. So where did this rule come from? Over a long period of time, I have investigated this question, and I have determined that the rule arose as a result of one particular figure (Figure 7-17) appearing on one particular page (page 144), of the Motorola ECL Systems Handbook that was published in 1973. This figure shows two traces, side-by-side, one with two right angle bends and the other with those bends rounded. Below that picture are two oscilloscope traces. The trace for the right angle bend example has two small dips in the middle of it while the trace for the rounded corners has none. This suggests that right angle bends cause a problem. Backing up to page 140 of the same handbook, Figure 7-14 shows a trace with two right angle bends and no disturbance at all from either one of the bends. Seems like a conflict. I called the author to question this and was told that Figure 7-17 was flawed and should not have been published! So, for all this time, engineers have been preventing the use of right angle bends based on flawed data. (Note: The Motorola ECL Systems Handbook is available as a PDF file from On Semiconductor at www.onsemi.com.

Right angle bends do not cause signal integrity problems at any practical edge speed.

Right angle bends also do not cause EMI.

Right angle bends are not acid traps.

There is no good technical reason to prevent the use of right angle bends to route traces in a PCB.

When I cover this topic in my classes, sometimes students will ask “what about if you have a large number of right angle bends in the same trace?” Same answer. Other students raise the question “What about acid traps formed by right angle bends?” This hasn’t been a problem in the history of PCB manufacture. It is a puzzle where this rule came from, but it is also not valid.

Vias

Routing vias are another source of signal integrity concern. Often, their use is forbidden for fear that signals will be degraded in an undesirable way. It has been known for quite some time that a via or plated though hole used to change layers in a PCB looks like a very tiny capacitor. The question is how big is that parasitic capacitance and will it have an adverse affect on a switching edge?

The test PCB in Figure 25.1 has two arrays of test traces running horizontally across the PCB. In each array, the center trace had no vias along its length. Above and below this “unloaded” trace are traces with a via every 0.4 inches, a via every 0.2 inches and a via every 0.1 inches. The traces above the unloaded trace have vias drilled 13 mils in diameter. The traces below the unloaded trace have vias drilled 30 mils in diameter. The upper array has the traces on layer 1 and the lower array has the traces on layer 4.

The principle on which this test is based is that adding additional capacitance along the length of a transmission line will lower its effective impedance. If the impedance of each transmission line in the test is measured with a TDR, those with more vias per unit length should have a lower impedance. Once we get the impedance values for each trace, we can work backward, using Equation 16.2, to calculate how much additional capacitance per unit length was required to lower the impedance.

Figure 25.3 has TDR traces for the four test lines. The top trace is the trace with no vias along its length; the second trace has a via every 0.4 inches; the third trace has a via every 0.2 inches and the bottom trace has a via every 0.1 inch. As was expected, the impedance drops as more and more vias are added, the impedances being 104.6 ohms, 81.9 ohms, 65.8 ohms and 48.5 ohms, respectively.

With this data in hand, a straightforward set of calculations can be performed to arrive at the capacitance of a single via. Table 25.1 contains that data. It can be seen that the average capacitance of the 13-mil drilled vias is 0.25 pF and the 30 - mil drilled vias is 0.6 pF. (Note: it is the size of the drilled hole, not the finished hole diameter that determines the parasitic capacitance.)

Once the value of the parasitic capacitance of a via is known, it is possible to perform analysis to assess its affect on signal quality. This is done by using an SI tool to construct a model of a transmission line. Once the model has been built and the performance established without the routing via, it can be added and the effect seen.

Figure 25.4 is a view of the waveforms on a 12-inch long 50-ohm transmission line with a 0.3 pF routing via added in the middle of the net. This net has no load on it. There is a small reflection at the point where the via is located. It is negative going, showing that the via has lowered the impedance at that point by increasing C_o for a short distance.

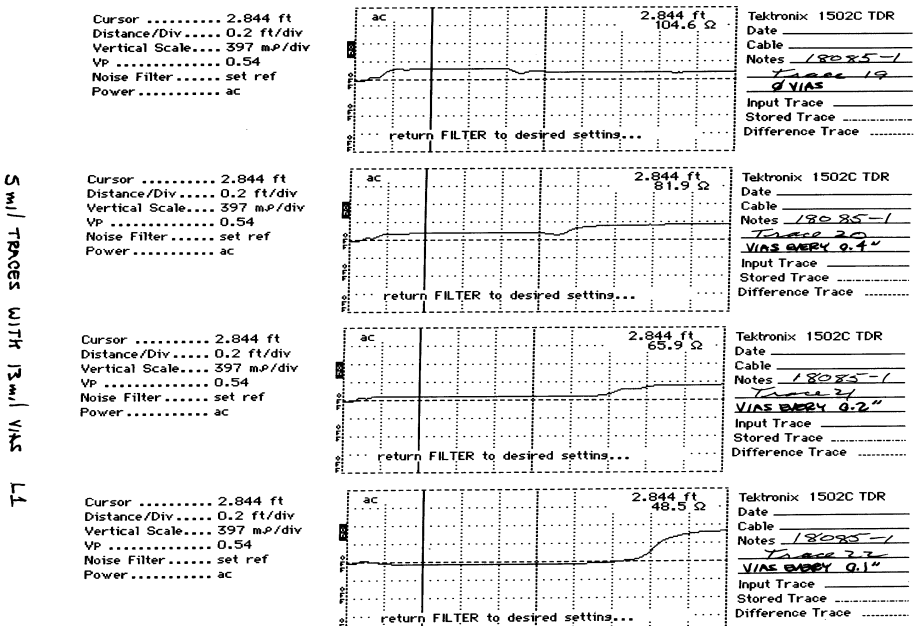

Note how adding more vias (capacitance) increases electrical length and reduces impedance.

Figure 25.3. TDR Screens for Four Test Traces with Vias Added

		CALCULATION OF VIA CAPACITANCE BASED ON 14 LAYER TEST PCB								
TRACE	Zo (ohms)	ZoxZo	Lo (nH)	Co(pF)	UNLOAD Co	C FROM VIA	VIAS/IN	C per VIA	VIA SIZE	AVE C
LAYER 1				pF	pF	pF		pF	mils	pF
15	32.5	1056	8	7.57	0.71	6.86	10	0.69	30	
16	49	2401	8	3.33	0.71	2.62	5	0.52	30	0.56
17	65.2	4251	8	1.88	0.71	1.17	2.5	0.47	30	
18	105.8	11194	8	0.71	0.71	0.00	0	0.00		
19	82.3	6773	8	1.18	0.71	0.47	2.5	0.19	13	
20	66.1	4369	8	1.83	0.71	1.12	5	0.22	13	0.22
21	49.1	2411	8	3.32	0.71	2.60	10	0.26	13	
LAYER 4										
48	26.7	713	8	11.22	2.44	8.79	10	0.88	30	
49	35.8	1282	8	6.24	2.44	3.81	5	0.76	30	0.76
50	44.7	1998	8	4.00	2.44	1.57	2.5	0.63	30	
51	57.3	3283	8	2.44	2.44	0.00	0	0.00		
52	50.3	2530	8	3.16	2.44	0.73	2.5	0.29	13	
53	44	1936	8	4.13	2.44	1.70	5	0.34	13	0.32
54	37.1	1376	8	5.81	2.44	3.38	10	0.34	13	
APRIL 1998		PREPARED BY LEE W. RITCHEY								

Vias, when used in traces, are capacitive, not inductive. The capacitance value of a via is small compared to the capacitance of a trace (3.5pF/inch for 50 ohms). In general, vias are not visible to signals with edge rates slower than .3 nSEC. NOTE: Test PCB as 100 mils thick.

Table 25.1. Calculation of Via Capacitance for 13-mil and 30-mil Drilled Holes in 100-mil Thick PCB

Figure 25.5 is the same net with a CMOS load added at the end farthest from the driver. Notice that the reflection caused by adding the parasitic input capacitance of a load is much larger than that caused by the via. Why isn't that troubling to

an engineer who is concerned about a routing via? Perhaps the “no via” rule or “because we’ve always done it this way” excuse were used as arbitrary guidelines rather than relying on critical analysis.

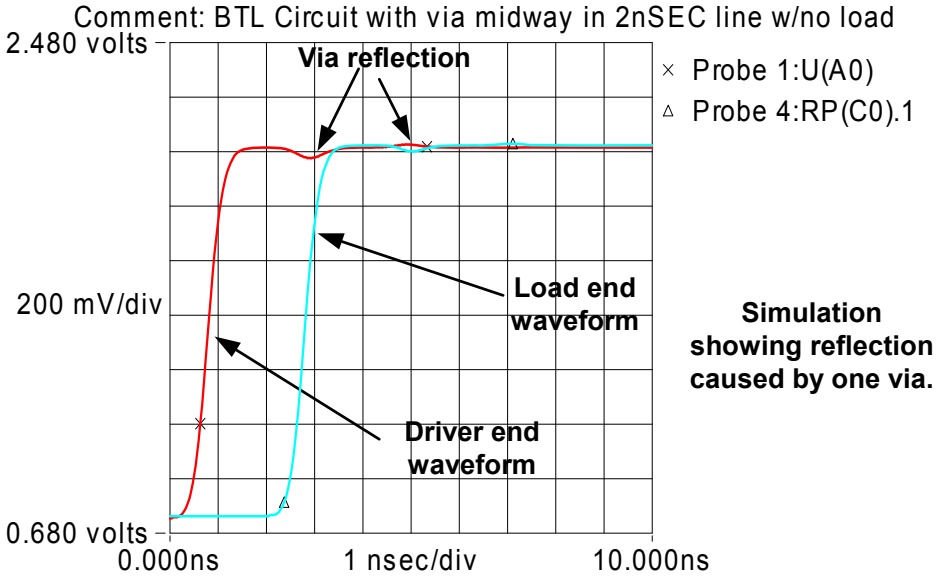

Figure 25.4. Signals on a 50-Ohm Transmission Line with 0.3 pF Routing Via Added in the Middle

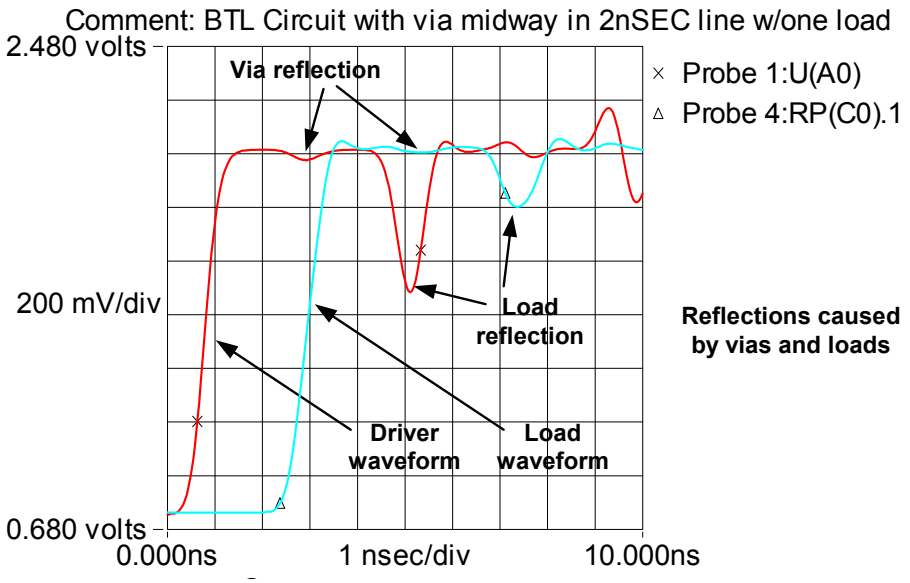

Figure 25.5. Signals on a 50-Ohm Transmission Line with 0.3 pF Routing Via and a Single Load at the End

As can be seen from the above tests and simulations, vias are not as insignificant as right angle bends are in terms of causing SI problems. They are, in fact, visible. However, the effect that routing vias (13-mils in diameter or less) have on signals is so small as to be unnoticeable. Constraining the routing of traces by prohibiting the use of routing vias introduces a complication without a benefit.

There is a case where using a routing via to change layers is more noticeable. When the two layers used to route signals on are not of the same impedance, there will be a reflection at the point where the trace changes layers. Figure 25.6 illustrates what happens when the two routing layers are at the opposite ends of the ±10% tolerance range normally found in the manufacture of PCBs.

In this case, the reflection is significant. One might be tempted to force the signals to remain on a single layer in order to avoid this type of reflection. This would only solve part of the impedance tolerance problem. Allowances still have to be made when the tolerance of a PCB arrives at one end of its ±10% tolerance range and the terminator is at the opposite end of its tolerance range. When the design rules allow for the normal impedance tolerances of drivers, transmission lines

and terminators, the amount of reflections seen in Figure 25.6 will be allowed for and there won't be a need to restrict via usage.

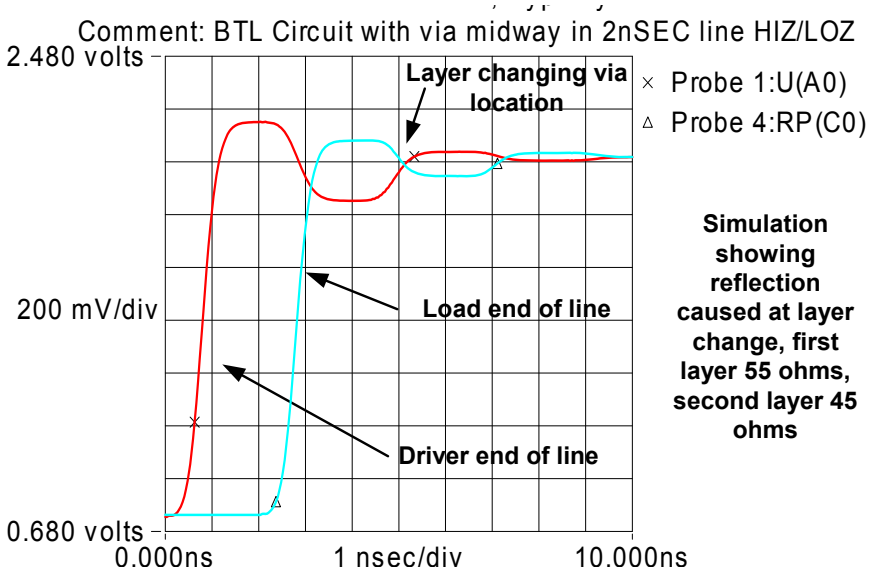

Figure 25.6. Signals on a 50-Ohm Transmission Line When a Routing Via is Used to Change Layers--Layers at Tolerance Limits, one at 55 Ohms, the other at 45 Ohms

The forgoing discussions deal with vias as parasitic capacitors, leaving the impression that vias are only capacitive. In fact, they are distributed networks and, as a result, also have inductance along their length. When a via is used to change layers in a signal trace, the effect of this inductance is least when the two layers are adjacent to each other. When using a via for routing, it has its biggest impact as a parasitic capacitor because the inductance is not in series with the signal. When the via is used to change from the top layer to the bottom layer of a PCB, the inductance is in the signal path. The effect, in this case, is to lessen the effect of the parasitic capacitance.

When a via is used to make a power connection to a plane layer, the parasitic capacitance is not visible. What is visible is the parasitic inductance. The effect that this inductance has on power connections will be examined in Chapter 35, in the discussion of bypass capacitors and their mountings.

> Vias used to route signals from one layer to another do not have a significant impact on signal quality, even with 200 pSEC edges.

The effect of large vias, such as those used in press fit connectors in thick backplanes, will be examined in Volume 2. At 2.4 GB/S and higher, the reflections from these vias or plated through holes cause jitter to become worse.

Looking back at Table 25.1, it can be seen that the average capacitance of a 13-mil drilled diameter via, 100 mils tall, is approximately 0.25 pF. The average capacitance of a 30-mil drilled diameter via, 100 mils tall, is approximately 0.6 pF. From a previous test PCB, it was determined that a 40-mil drilled diameter, 100 mils tall, is approximately 0.8 pF. If one calculates the outer surface area of the three cylinders formed by the plating in their holes, it will turn out that the parasitic capacitance is proportional to the area of this cylinder. From this observation it is possible to extrapolate the parasitic capacitance of vias that are taller or larger in diameter than those shown here. In actual practice, the measured parasitic capacitance of 30-mil diameter vias in 250-mil thick backplanes has scaled accurately from the 30-mil via in this test by multiplying it by 2.5

CHAPTER 26: TYPES OF DRIVERS OR SOURCES

Logic, analog and RF signals are created by drivers. These drivers launch energy in the form of an electromagnetic field down a transmission line. As mentioned earlier, the objective is to deliver a voltage waveform to the load or loads along or at the end of the transmission line. The voltage waveform is the electric field portion of the EM field.

There are three basic types of drivers. These are described as voltage sources, current sources and matched impedance drivers.

VOLTAGE SOURCES

Voltage sources can be represented with symbols such as that shown in Figure 26.1.

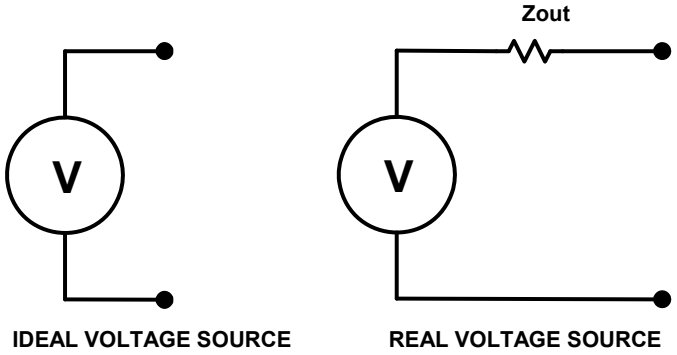

Figure 26.1. Voltage Sources

The symbol on the left side of Figure 26.1 is what is called an ideal voltage source. This means that no matter how much load current is drawn from the driver, the output voltage remains constant. Said another way, the load impedance can vary over a wide range and the output voltage of the source will remain the same. The voltage source is said to have an output impedance that is zero.

Unfortunately, real voltage sources are not perfect. They have output impedances that are not zero as is shown in the right side of Figure 26.1. When a load, such as a transmission line, is connected to such a source, the voltage that appears at the load is less than V. The current that flows into the load causes a voltage drop across Zout. As a result of this phenomenon, the signals that arrive at the loads will be less than V. When circuits are designed with real drivers, it is important to characterize drivers by determining their output impedances and allow for the signal degradation that results.

Many familiar logic families that are intended to be voltage sources have relatively high output impedances. In fact, they are so high that they cannot successfully drive transmission lines at high speeds. Most of the logic drivers that are rated in milliamps are in this category. To successfully drive a series terminated transmission line, a driver must have an output impedance that is the same as or less than the impedance of the transmission line. To successfully drive a parallel terminated transmission line, a driver must have an output impedance much less than the impedance of the transmission line.

Table 39.2 lists the output impedances of several logic families. Table 26.1 lists logic families that are designed to be voltage sources.

TTL- all types
ECL
LVCMOS
CMOS- all types
GaAs

Table 26.1. Logic Families Designed to Operate as Voltage Sources

CURRENT SOURCES

Current sources can be represented by the symbols in Figure 26.2.

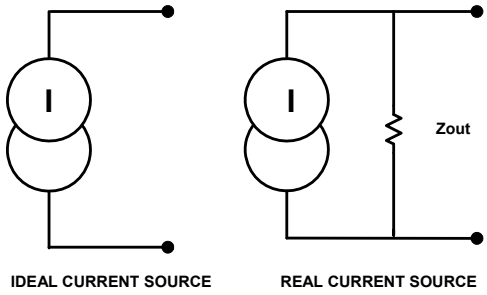

Figure 26.2. Current Sources

A current source has the property that no matter what load impedance is attached to it, the output current remains constant. Therefore, the output voltage is directly proportional to the load impedance. If the current source is 1 ampere and the load impedance is 1 ohm, the output voltage will be 1 volt. If the load impedance is 2 ohms, the output voltage will be 2 volts and so on. If the load impedance is infinite, the output voltage will be infinite. As a practical matter, the latter case does not exist. All real current sources have a limit to the voltage they can produce.

Like real voltage sources, real current sources are not perfect. Their output impedances are not infinite. The symbol on the right side of Figure 26.2 more accurately represents real current sources. The output impedance, Zout, is very high, but not infinite. The current, I, will divide between this impedance and the load impedance. The result is the voltage seen at the load is always less then I x Zload.

Table 26.2 lists logic families that are designed to operate as current sources.

LVDS
CML

Table 26.2. Logic Families Based on Current Sources at Drivers

The most common form of current source based logic currently in use is LVDS. The value of I is 4 milliamps. In normal use, each side of the differential pair in this logic protocol drives a 50-ohm transmission line that is parallel terminated in 50 ohms. As a result, 200 millivolts will develop across this impedance. One side of the differential pair is sourcing the 4 milliamps and the other is sinking the 4 milliamps. This results in one voltage of +200 millivolts and the other of -200 millivolts. The two together yield a total of a 400 millivolts difference voltage across the two terminating resistors. If the terminating resistors are 45 ohms each, the difference voltage will be 90 ohms times 4 milliamps or 360 millivolts. It follows that the voltage seen at the input to a load driven by a current source driver is proportional to the size of the termination resistor or load impedance.

Matched Impedance Drivers

Some drivers or sources have output impedances that are designed to match the impedance of the transmission line to which they are attached. The most common such source has an output impedance of 50 ohms. The purpose of matching these drivers or sources to their loads is to minimize reflections in the overall transmission line. Such a driver looks like the diagram on the right side of Figure 26.1. The value Zout is chosen to match the impedance of the transmission line that will be driven.

Until recently, matched impedance drivers or sources were seen only as the signal sources for instrumentation such as signal generators. However, as logic family speeds have increased, it has become necessary to design logic drivers that match the impedance of the transmission lines they are expected to drive. Series terminated transmission lines driven by CMOS drivers are in this group. This has become necessary due to the need to add a series-terminating resistor at the output of every driver. Adding a series-terminating resistor on the PCB for every net results in designs that are too crowded to permit successful layout.

Signal Generators
Advanced CMOS Drivers
High Performance LVDS

Table 26.3. Matched Impedance Sources or Drivers

CHAPTER 27: TYPES OF LOADS

Loads are the receivers of the voltage waveforms sent down transmission lines by drivers. The kinds of inputs to logic devices are listed in Table 27.1.

TTL--NPN transistor emitter
CMOS--Insulated gate of an FET
ECL--NPN transistor base
LVDS--Insulated gate of an FET

Table 27.1. Logic Device Inputs

All of the loads listed in table 271. draw little, if any, current from the transmission line to which they are attached. Compared to the impedance of the transmission line, their impedances are very high. All of them have parasitic elements that are an unavoidable consequence of the way ICs are designed.

One of these elements is the parasitic capacitance associated with the metal used to make the connection to the transistor as well as the capacitance of the transistor itself. This is the parasitic that is most visible to fast signals. This parasitic capacitance must be charged up and discharged in order to change logic states. It can range from 3 pF to as much as 10 pF.

A second parasitic is the PN diode junction that exists between the input transistor and the substrate. This junction is designed to be reverse biased during normal operation. This isolates the transistor from the rest of the substrate. When the input voltage is taken below ground by a reflection far enough to cause the junction to be forward biased, it will conduct current into the substrate. Logic failures and device damage can occur when this happens.

A third parasitic that can exist at the input to a load is a pair of input protection diodes that are intended to protect the input from over voltages. Figure 27.1 illustrates such an input circuit. These diodes conduct when the input voltage goes above Vdd or below ground by more than 0.3V. In some cases, there can be undesirable interaction between these diodes when they are conducting and other circuits are nearby. This unwanted interaction may cause logic failures.

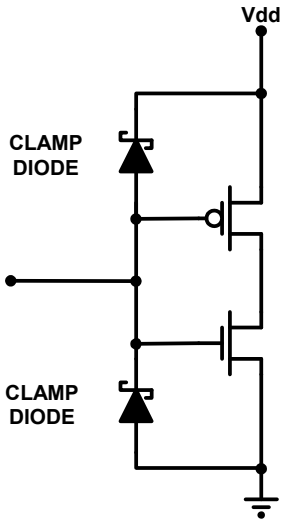

Figure 27.1. A CMOS Input With Input Protection Diodes

In normal operation, the inputs to logic circuits are very small parasitic capacitors that must be charge and discharged in order to change logic states.

CHAPTER 28: BUS PROTOCOLS

There are a number of common bus protocols used in the design of logic products. Most of them have standards that define how to interface to them and what performance can be expected from each. Without question, the most common bus protocol in the world is the PCI bus. This bus protocol is used in virtually all personal computers made with microprocessors from Intel and AMD. Because of the availability of many adapters and add-in specialty cards, as well as a worldwide understanding of the operating software that surrounds the PCI bus, it is used in many other products as well.

Figure 28.1 depicts the diagram for a data path on a PCI bus. It uses a series terminated logic family. It relies on reflected wave switching. The impedance of the transmission lines is nominally 55 ohms. Each driver is matched to the bus with a series termination in its output. There are no terminations on the bus itself. Remember, this is a "reflected wave" switching protocol as discussed in Chapter 20. This means data is only good along the whole bus after the reflected wave arrives back at the driver from which it was launched. Not shown in this diagram is a clock distribution network on the backplane that insures the clock edges line up with the data good periods. This protocol has a speed limitation, as discussed in Chapter 20, that is determined by the length of the bus and its round trip delay. When this limitation is reached, parallel-terminated buses must be used.

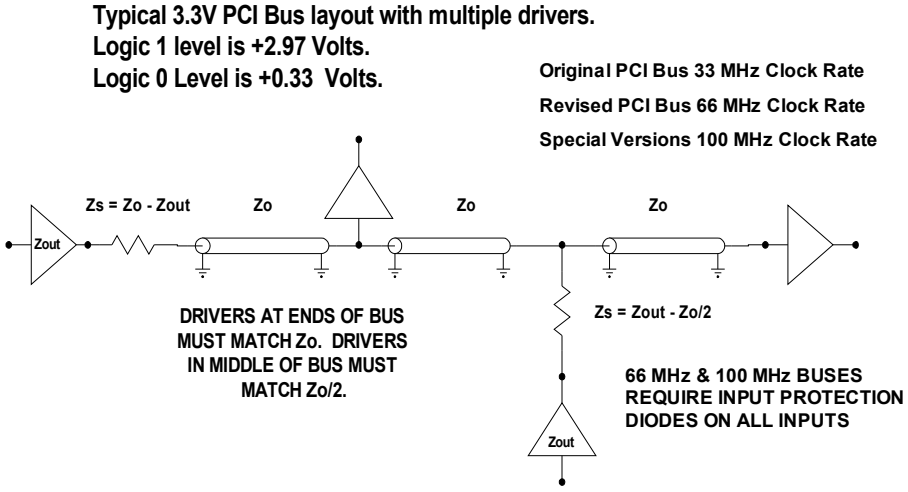

Figure 28.1. A PCI Bus

Figure 28.2 illustrates a bus using Series Stub Terminated (SST) logic. This is a variation of a series terminated logic bus. This bus protocol was devised to allow plugging in a varying number of memory SIMMs on a PC motherboard. The SIMMs have stubs that run up into the SIMM PCB for a distance that would cause malfunctions like those discussed in Chapter 22 to occur. The 22-ohm resistors on the SIMMs at the point where contact is made with the bus on the motherboard provide enough isolation of the SIMM stub from the bus so that signal quality is not adversely affected. This works well, so long as memory bus speed is not pressed much above 200 MB/S per line. At these higher data rates, the signal quality at the inputs to the components mounted on the SIMM becomes marginal and a true parallel-terminated bus is required.

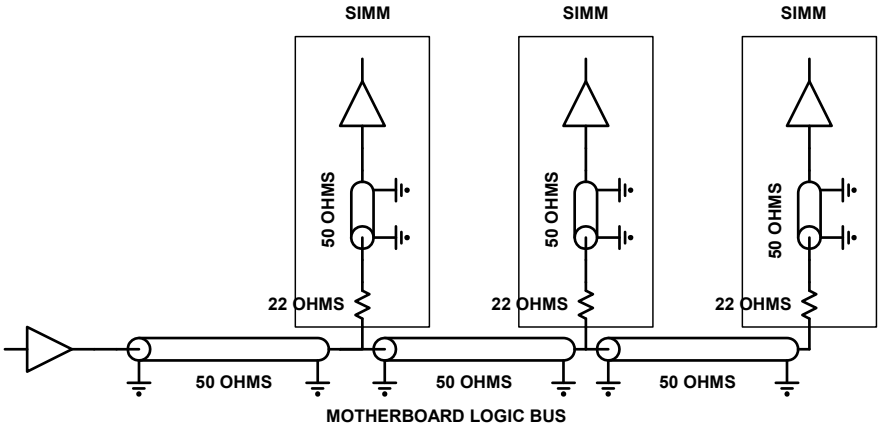

SERIES STUB TERMINATED LOGIC BUS

Figure 28.2. An SSTL Bus

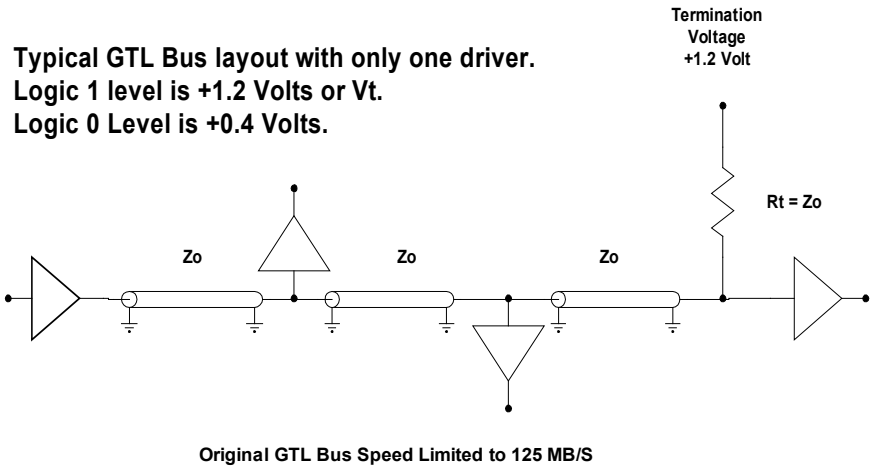

Figure 28.3. A GTL Bus

The diagram in Figure 28.3 is a GTL bus. The G stands for Bill Gunning of XEROX PARC laboratories in Palo Alto, California. During the mid 80s, large systems with daughter boards and backplanes were being designed for products such as workstations, massively parallel computers and servers. The logic in these machines was TTL and 5V CMOS. Neither of these two logic families is capable of driving a backplane bus. In order to solve this problem, level translators were used at bus interfaces with Emitter Coupled Logic (ECL) running on the backplane. The backplane bus was parallel terminated allowing data to be good on the outbound signal rather than having to wait for the reflected wave to arrive back at the driver. This protocol had two drawbacks. The first was the need for two power supply voltages (+5V and -5.2V) and the second involved a dead time related to the time delay through the level shifter at each end of a data path. This was much like the time lost in a series terminated bus with the round trip delay. Eventually, this dead time limited machine performance when this method was used to drive backplanes. (It was difficult to get backplane bus speeds much above 60 MHz.)

In the mid 80s, Bill Gunning and others decided to design input and output cells for CMOS circuits that emulated the ECL. The result was Gunning Transistor Logic (GTL). With GTL, no level shifting was needed; the drivers were capable of driving 28-ohm backplane buses at 125 MHz; the bus could be parallel terminated; there were no indeterminate logic levels and data was good on the outbound signal. This bus protocol found its way into PCs as the bus for the cache memory. Nearly all PCs contain a GTL bus used in this manner.

This solved the speed problem for a while. Soon, products needed to drive backplane buses at rates higher than 125 MB/S. This gave rise to the BTL bus protocol.

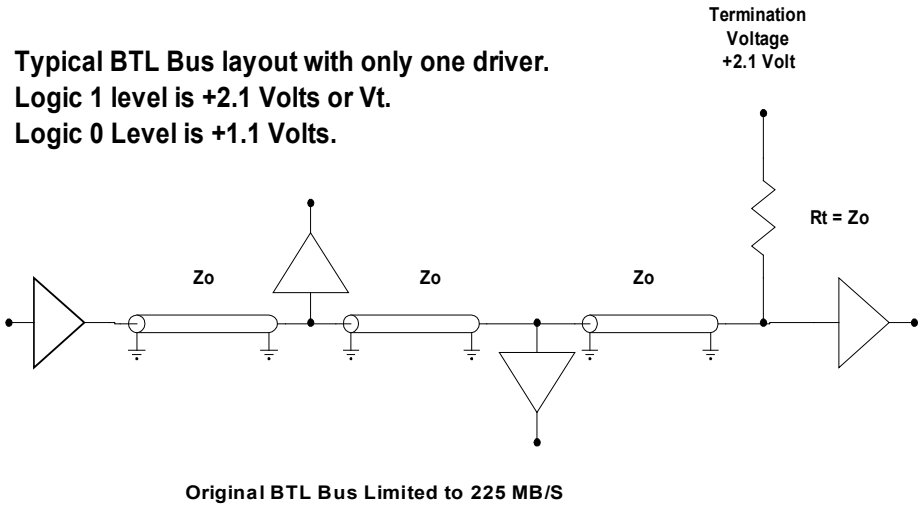

Figure 28.4. A BTL Bus

Figure 28.4 is a diagram for a BTL bus. In this case, the B stands for bipolar. The major difference between this bus and GTL is the driver. In GTL, the driver is an open drain NMOS transistor. The high resistance and slow speed of the NMOS transistor limited the overall speed of the GTL bus. BTL uses an open collector NPN transistor to solve these problems.

Products such as large routers and servers used the BTL bus protocol in their backplanes. When higher data rates were required than could be achieved with the maximum clock rate of the BTL bus, the problem was solved by creating increasingly wider buses. Eventually, bus widths reached 256 bits. Switching 256 single-ended lines such as this created very large simultaneous switching noise currents. These current transients began to limit performance due to phenomena such as Vcc and ground bounce and very large voltage transients on the Vdd supply. Doubling performance by going to 512-bit wide buses was physically impossible. Differential signaling has provided the path to higher bandwidth backplanes. This started out as LVDS used in laptop computers.

Typical LVDS Bus layout with a single driver.
Logic 1 level is +1.475 Volts.
Logic 0 Level is +0.925 Volts.
Vref = + 1.150 Volts

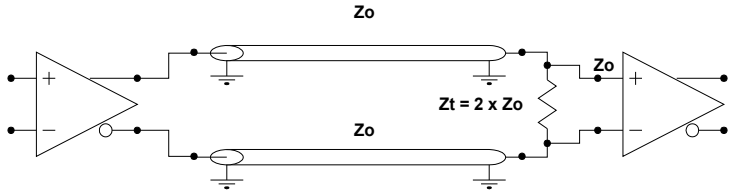

Lay out each line as a stand alone transmission line.
Match lengths within limits set by LVDS standard.
Center of terminating resistor goes to a virtual ground.
Therefore, each line is terminated in Zo.

Figure 28.5. An LVDS Bus Path

Figure 28.5 is a typical Low Voltage Differential Signaling Bus path (LVDS). This signaling protocol was devised by a consortium of laptop computer manufacturers and National Semiconductor to solve the problem of getting the high bandwidth graphics data from the laptop motherboard through the hinge and up to the display. This path was so small it was not possible to create a grounding system that would work properly for single ended switching. This was not the first time this problem appeared. In the early days of ECL-based computers, it was necessary to transmit low level ECL signals from one gray box to another in a computer room. There was no good way to keep the logic ground in two boxes close enough to each other to allow single ended ECL logic signaling. This gave rise to the original differential signaling protocol and was the underlying idea for LVDS. Both solutions address poor grounding in the form of ground offsets between the two ends of the data path.

The original LVDS protocol had a data rate of 75 MB/S on each pair of signal paths. This method of signaling provides very good immunity to ground offsets and noise and can be clocked at extremely high data rates. As a result of this, it has been adopted for a wide variety of other product families including:

Fiber channel for disc drives
Infiniband as a replacement for the PCI bus
PCI Express, another replacement for the PCI bus
2.4 GB/S Ethernet links
4.8 GB/S Ethernet links
9.6 GB/S Ethernet links
SONET
SCSI Buses
Universal Serial Bus (USB)
Hyper Transport

All of these signaling protocols use some variation of LVDS.

Chapter 31 will provide a detailed discussion of how the LVDS signaling protocol deals with ground offsets. It will also address the methodologies for managing the two members of each differential pair.

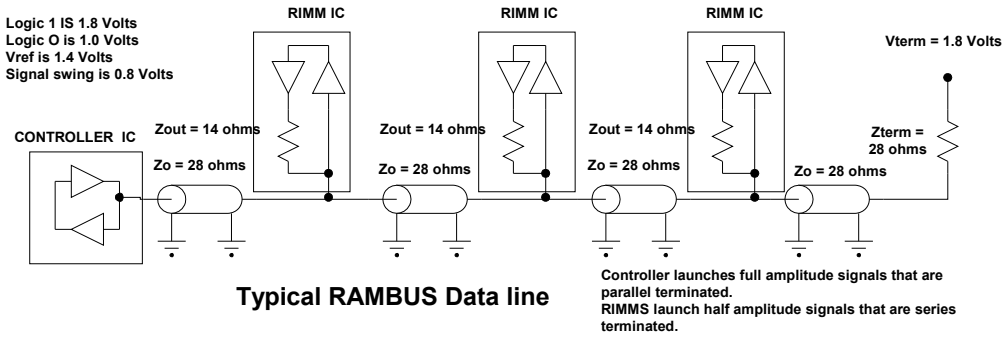

Figure 28.6. A Typical Rambus Data Line

From the beginning of computer design, main memory has always been slower than the CPU, often by an order of magnitude. This disparity in the speed of the two main components in a computer has limited the speed with which processors can operate on data that must go into and out of main memory. Fast cache memory has been located inside the CPU or right next to it to help solve this problem. Sooner or later, even with this solution and clever operating software, this disparity in speed always slows down the CPU. The Rambus® memory architecture was intended to solve this problem. By using specially designed memory ICs and a specially designed memory controller, it is possible to build a main memory with a data rate of 800 MB/S per data line. Figure 28.6 is an illustration of a typical Rambus® data line. When data is sent out from the controller IC to the RIMMs, the driver functions as a parallel-terminated transmission line. When a RIMM sends data back to the controller, the path is series terminated.

This memory organization has proved difficult to implement in actual designs. Some of the reasons are the need for very tight timing between all of the elements in the chain and the fact that the bus needs an impedance of 30 or so ohms, a very difficult impedance to fabricate and keep within tolerance.

CHAPTER 29: CROSSTALK OR COUPLING

The term crosstalk is used interchangeably with coupling to describe the unwanted interaction of the electromagnetic field traveling on one transmission line with a neighboring transmission line. The neighboring transmission line could be in the same signal layer or in an adjacent layer. This unwanted interaction takes the form of energy coupled into the neighbor as a disruptive noise signal.

In Chapter 8, the EM field was shown to have two components--the electric field and the magnetic field. Each of these two components can couple energy from a driven line into a quiet line. When the electric field dominates it is because the stray capacitive coupling between the two transmission lines is larger than the stray magnetic coupling. This is often called capacitive coupling. This form of coupling occurs when traces are run one over the top of the other in the adjacent signal layers of a PCB. The magnetic field dominates when the magnetic coupling between the two transmission lines is larger than the capacitive coupling. This is often called magnetic coupling. This form of coupling occurs when traces run side by side in the same signal layer of a PCB.

In both cases cited above, both types of coupling are always at work. Each is discussed as though they were independent. The reason for focusing on one type or the other is to draw attention to the main source of crosstalk failure and focus efforts on making sure that it is within the tolerances of the logic family being used, knowing that if the primary source of crosstalk is under control, the lesser form will be as well.

Capacitive Crosstalk

As mentioned earlier, crosstalk from the stray capacitive coupling between two TLs dominates when two traces run one over the top of the other in adjacent signal layers. This type of crosstalk looks as though a very small capacitor has been connected between the two signals. The faster the edges, the larger the coupled signal is. For the very fast edges of modern electronics, a stray capacitance of only a few picofarads results in enough coupling to exceed the noise budget of any logic family. Many hundreds of hours have been spent trying to develop an algorithm that will accurately predict how long two signals can run one over the top of the other without creating excess coupling. None of those efforts has resulted in a reliable predictive algorithm. Instead, controlling coupling between signals in adjacent layers is accomplished by imposing routing rules that force signals in adjacent layers to route at right angles to each other. In this way, there is no detectable coupling, either capacitive or inductive.

PCB layout tools intended for high speed PCB design have controls built into them that allow a designer to constrain routing to be X in one layer and Y in the adjacent layer. This is the mark of a well-designed PCB layout tool. However, most of these routing algorithms will occasionally violate the “X only” or “Y only” rule As a result, post-route analysis needs to be done to insure that this type of routing violation has not occurred.

Inductive Crosstalk

Inductive crosstalk occurs because the stray magnetic coupling is large enough to couple significant energy from one TL to its neighbor. This occurs when traces run side by side in the same layer. Stray capacitive coupling exists, as well, but it is much smaller due to the fact that the two traces are edge on with very little surface area in common.

Inductive coupling is unavoidable. It is necessary to route signals side by side in the same layer in order to fit all the wires into the PCB. Because of this, it is necessary to understand clearly how this form of crosstalk operates. Coupling due to transmission lines running in parallel takes two forms: forward crosstalk and backward crosstalk. Figure 29.1 illustrates these two types of crosstalk.

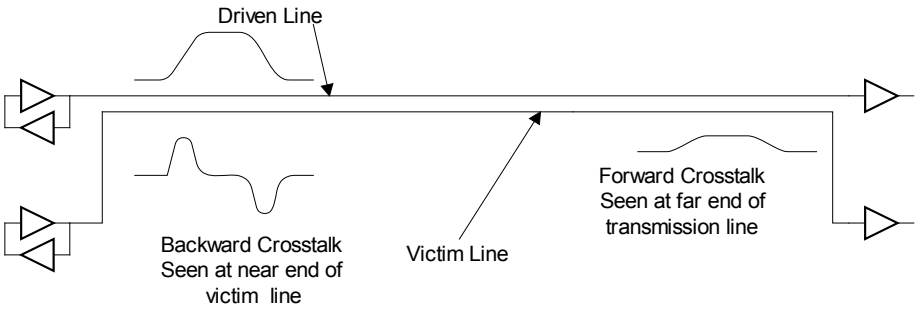

Two types of coupling resulting from wires travelling side by side for an extended distance. Often called parallel crosstalk.

This type of crosstalk is dominated by the inductive component of the coupled signal rather than the capacitive component.

Backward crosstalk will grow rapidly with parallel length and reach a saturation value after which further parallelism will not cause an increase. Foreward crosstalk will grow very slowly until its amplitude equals the inducing signal. It's shape will mimic the inducing signal.

Figure 29.1. Two Coupled Transmission Lines Showing Forward and Backward Crosstalk

The terms forward and backward crosstalk were developed by RF engineers to describe the behavior of closely coupled transmission lines. The crosstalk signal observed in a victim line depends on the point on the victim line where it is measured as well as how the four ends of the two lines are terminated. For a complete treatment of all of the variations of crosstalk that are possible, the paper on crosstalk written by John de Falco of IBM that is cited in the bibliography of this book should be read.

By definition, backward crosstalk is the noise signal observed on the end of the victim line nearest the driver in the driven line. Forward crosstalk is the noise signal observed on the end of the victim line farthest from the driver of the driven line.

Figure 29.2 illustrates how these two forms of crosstalk vary as the length that two transmission lines run parallel increases.

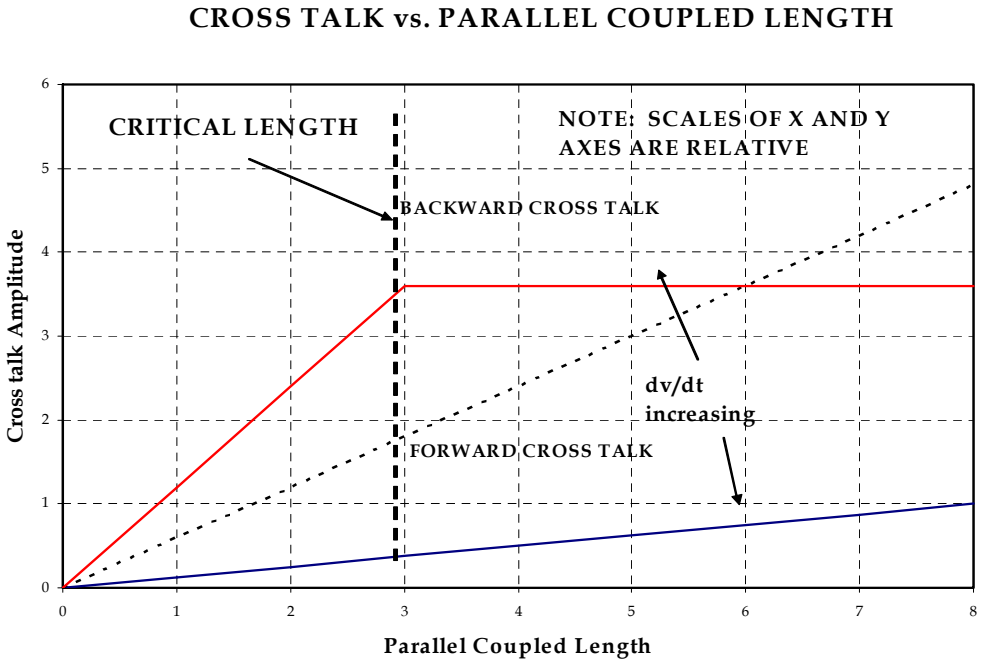

Figure 29.2. Forward and Backward Crosstalk as a Function of Parallel Coupled Length

The graph in Figure 29.2 is a qualitative curve meaning that it portrays how the two forms of crosstalk vary with length of parallelism. Notice that forward crosstalk grows somewhat gradually, but continues to increase as long as the two TLs run parallel. If the two TLs run parallel long enough, the forward crosstalk signal will be the same size in both lines. For this to occur, the lines would have to run parallel much longer than is possible in PCBs. This type of crosstalk is at work in the cables phone companies use. It is the most common reason that other phone conversations can be heard faintly in the background. Phone cable manufactures resort to all sorts of twisting and interweaving to reduce this type of crosstalk. Fiber optics have made most of this kind of crosstalk go away.

Notice that backward crosstalk grows much more rapidly with coupled length than does forward crosstalk. Also, notice that at some point backward crosstalk reaches a maximum or saturation point. Continuing to run parallel after this point does not result in more backward crosstalk. This saturation point is called the "Critical Length." (Note: often, critical length is incorrectly used to describe the length of a transmission line beyond which it is necessary to control impedance.) In RF applications, the fact that backward crosstalk reaches this saturation point is exploited to build directional couplers that allow a sample of the signal in the driven line to be taken at the backward end of the victim line. This is a useful application of backward crosstalk. In logic circuits there is no such useful application. Backward crosstalk always results in unwanted noise that can cause logic malfunctions.

If one is to control backward crosstalk by restricting the length that two transmission lines run in parallel, the implication is that the two traces will need to be separated from each other at some length less than this critical length so that crosstalk is some value less than the saturation value. Critical length is a function of the rise time of a switching edge or the frequency of signal traveling on the driven line. Figure 29.3 illustrates how the critical length varies with the rise time of the signal being sent.

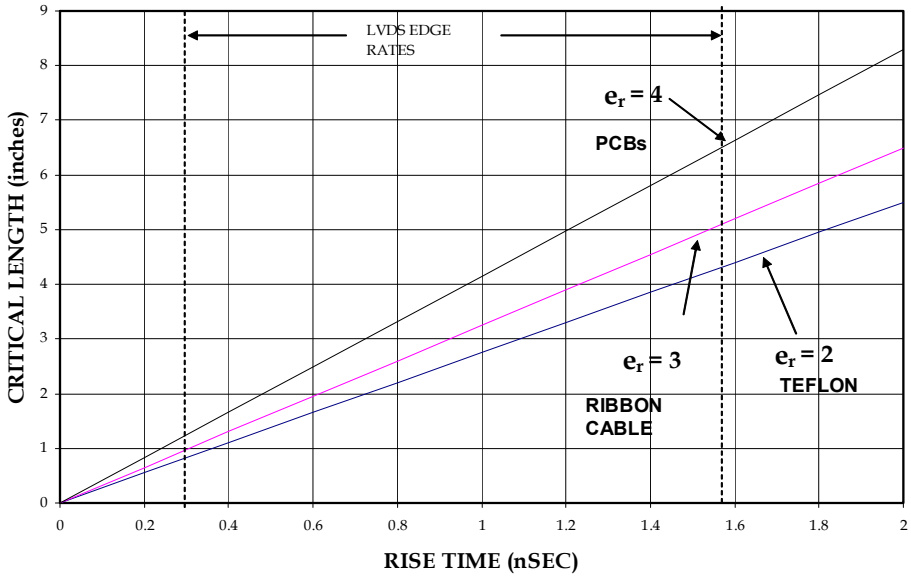

Figure 29.3. Critical Length vs. Rise Time and Relative Dielectric Constant

If one were designing a PCB with switching edges of 1.5 nSEC, the critical length would be approximately 6 inches. Controlling backward crosstalk by limiting the length of parallel run would require imposing a parallel run restriction of some distance less than 6 inches, such as 3 inches. It might be possible to route PCBs with this restriction, but it might also be difficult. The fastest rise and fall time of most modern logic devices is less than 0.5 nSEC. At this rise time, critical length is a little over two inches! Controlling crosstalk by restricting the length of parallel run would require running parallel less than 1 inch. With such a restriction, it would be virtually impossible to route any PCB. This has been the operating condition for super computers for the last 20 years. How were the PCBs in such designs routed successfully?

Controlling backward crosstalk by limiting length that two traces run parallel does not work.

Knowing that it is necessary to route signals in parallel longer than the critical length, the only method that is available for controlling backward crosstalk is to impose geometries on the traces that keep the crosstalk within design limits, no matter how long they run in parallel. Once two traces have been run side by side longer than the critical length, crosstalk is a function of edge-to-edge separation of the traces and height above the nearest plane. Trace width is not a variable nor is relative dielectric constant of the insulating material. Figure 29.4 is a plot of crosstalk vs. edge-to-edge separation and height above the nearest plane for an asymmetric stripline coupled pair.

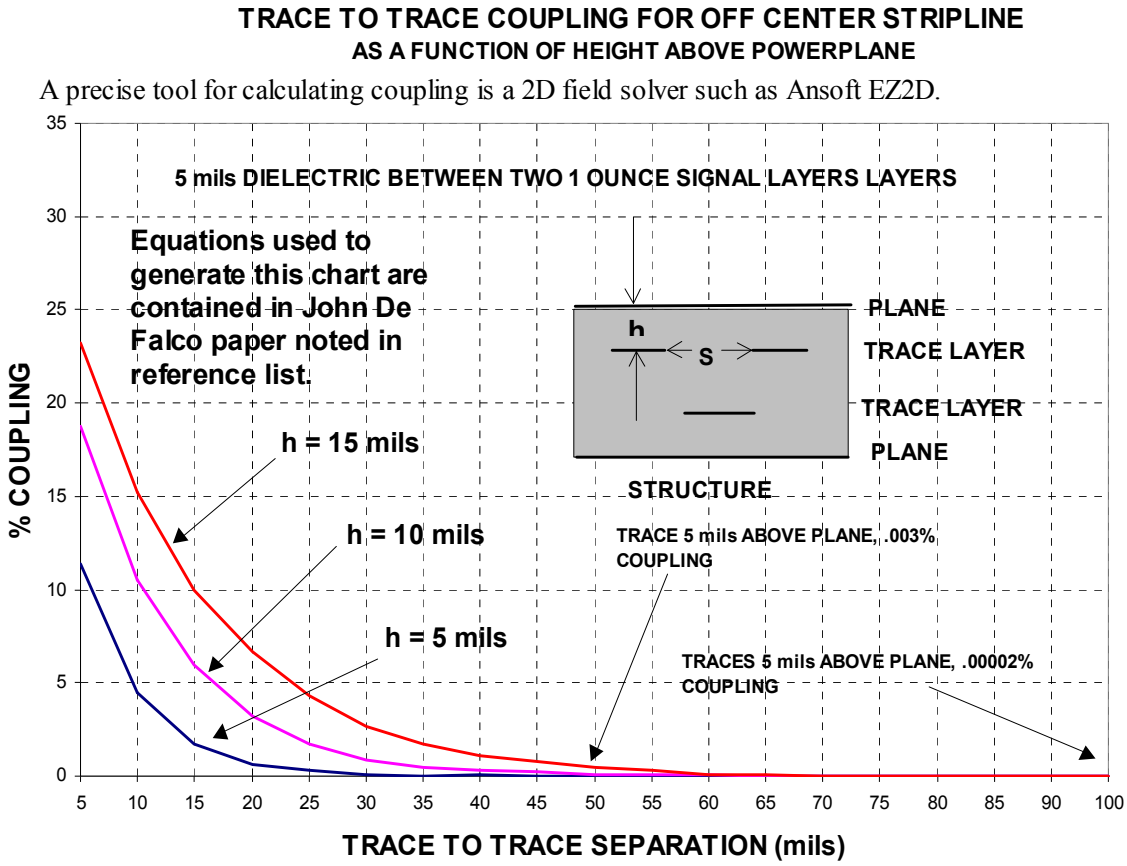

Figure 29.4. Crosstalk vs. Height Above the Nearest Plane and Edge to Edge Separation for Asymmetric Stripline

This chart shows the amplitude of the coupled noise in the victim line as a percentage of the signal in the driven line. To use these curves to create design rules, the allowable noise amplitude in the victim line is first established. The maximum signal swing in the driven line is calculated. From these two numbers, the allowable percentage of the driven signal can be calculated. Various combinations of edge-to-edge separation and height above the nearest plane can be tried until trade offs are made between PCB thickness, impedance and routability that result in the most manufacturable PCB that also meets electrical requirements. Notice that when this step has been done, routing rules consist of trace widths and trace spacing. If this spacing is followed there is no need to do board level simulations to insure specifications have been met. Routing rules and checking rules are all done by performing geometry checks to insure the spacing rules have been followed.

Several things are worth noting. First, the slope of the curves as traces are routed closer and closer to each other is quite steep. **The trend in PCB design and fabrication to route PCBs with increasingly finer traces and smaller spaces carries with it a significant risk of failing from crosstalk alone**. Second, as traces are separated, crosstalk drops rapidly. At separations of 50 mils, there is virtually no interaction between traces on the same layer. This is a direct indication that the energy in the EM field of a transmission line stays very close to the TL on which it is traveling. It does not "stray around" and disturb other signals that are far away. Further, it does not wander around and leak out the sides of a PCB or backplane requiring edge plating to contain EMI. Third, as the height above the plane is reduced, the crosstalk also diminishes. Stackups should be designed with the smallest distance between trace layers and the nearest plane that results in a practical, manufacturable PCB.

Figure 29.4 is saturated crosstalk for asymmetric striplines. Figure 29.5 is a similar set of curves for surface or buried micro-stripline traces. As can be seen by comparing Figure 29.4 with Figure 29.5, for the same height above the plane

and the same edge-to-edge separation, crosstalk is higher in microstrip configurations. This is because there is only one partner plane for the microstrip configuration. In Chapter 24, a case was made for using buried microstrip transmission lines interchangeably with stripline layers for high-speed signals. In order to achieve this, three conditions must be met: these two entities must have the same velocity, same crosstalk and same impedance. The velocity condition is met by embedding microstrip lines in the dielectric. If one compares the two 5-mil height above the plane curves in the two figures, it can be seen that crosstalk is nearly the same. With minor adjustments of either height or spacing in the micro-stripline layer, crosstalk will be the same. All that remains is to adjust the trace widths in the two layer types to arrive at the same impedance and the goal has been met.

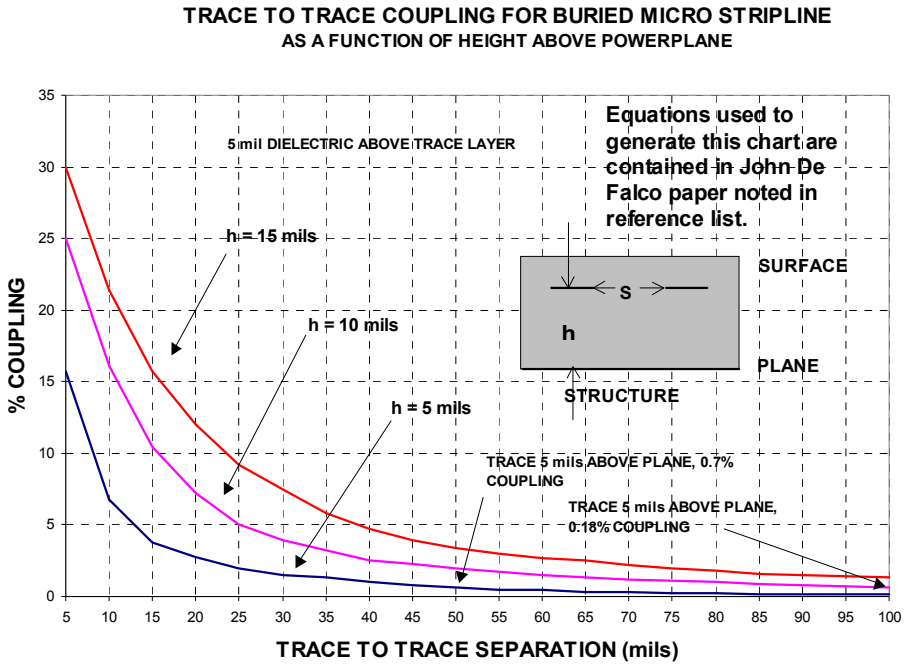

Figure 29.5. Crosstalk vs. Height Above the Nearest Plane and Edge to Edge Separation for Surface or Buried Micro-stripline

> Buried micro-stripline transmission lines can be used interchangeably with stripline layers for high-speed signals. They can be made electrically identical.

Figure 29.6 is an example of a typical, six-layer PCB. On the left, is the PCB as designed before the need for good transmission lines. Below the stackup on the left are listed the impedance and crosstalk for the two types of signal layers when routing is done with 5-mil lines and 5-mil spaces. The impedance and crosstalk are both unacceptable for a high-speed design. The stackup on the right-hand side of Figure 29.6 shows how rearranging the layers improves both impedance control and crosstalk.

The example on the left in Figure 29.6 is a common stackup for legacy PCB designs, done before logic rise and fall times were fast enough to warrant concern about either crosstalk or reflections. This is representative of the “TTL” era of logic design. As technology has evolved, the slow TTL parts have been replaced with less costly equivalents. The rise and fall times of these lower cost replacements are fast enough that both reflections and crosstalk are very important. Using these new, improved parts on legacy PCB designs, such as that illustrated on the left side of Figure 29.6, often results in failures from both crosstalk and reflections. Since purchasing more of the older, slower parts to continue building these legacy designs is not an option, the only alternative is to redesign the whole PCB or change the stackup as shown on the right hand side of the figure. Often, this restacking of the PCB will extend the life of such a design. Notice that the impedance is now the same in all signal layers, eliminating the reflection problem while reducing crosstalk by more than half.

As with most things in high-speed design, there is significant interaction between the design rules that are involved in creating stable designs. I liken the rule creation process to having a long carnival balloon with a lump in it. Squeezing down on the lump to make it go away often results in the lump popping up somewhere else on the balloon. The same kind of interaction is at work in Figure 29.6. Restacking the layers to improve crosstalk and reflections, separated the power planes from each other, reducing the inter-plane capacitance needed to support the fast switching edges that are

characteristic of the new logic parts. In addition, the inductance of the power planes is increased, further degrading the quality of the power subsystem. In the chapters of this book that deal with designing the power subsystem, it will be demonstrated that a large, high quality, low inductance plane capacitor is essential for supplying the switching currents associated with driving the transmission lines. Something will have to be done to the stackup to add back the Interplane capacitance. The solution will be described in the power system design section of this book.

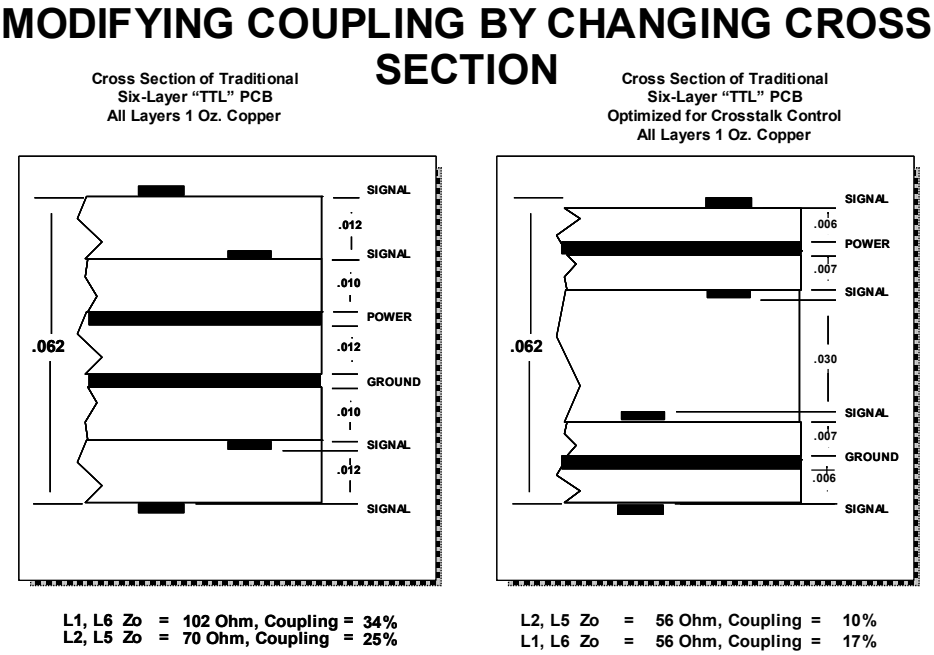

Figure 29.6. A Typical 6 Layer PCB Stackup Without SI Control and With SI Control

Guard Traces

Guard traces or ground traces are often represented as a way to control crosstalk. Is this really true? Imagine that a design is being routed in asymmetrical stripline with 5-mil lines and 5-mil spaces with the height above the nearest plane of 5 mils. Under these conditions, from Figure 29.5, it can be seen that the crosstalk is 12%, a number that is likely to cause logic failures in some systems. A decision is made to insert a guard trace. In order to make room for the "guard" trace, the spacing must be increased to make room for another 5-mil trace and another 5-mil gap. Now, the separation between victim and driven line is 15 mils. The crosstalk at this spacing is 2%, or a 6:1 reduction. Looks like the guard trace did its job. In actuality, it was not the guard trace the caused the change in crosstalk, it was the increased spacing.

The trace that was added was not "ground" even if its ends were connected to the ground plane nor is it capable of blocking the EM field. **Wires of any sort, no matter what their ends are connected to, cannot block EM fields. If this were true, transformers and motors would not work.** The trace is a distributed LC network that will resonate at some frequency. If the geometry of the trace happens to be right, it can resonate at a frequency of interest in the design creating a band pass filter, increasing coupling rather than reducing it. Figure 29.7 illustrates two ways to create band pass filters by arranging trace segments side by side.

Notice that it is possible to create a band pass filter by attaching both ends of each trace segment to the underlying plane, as is done with a "guard" trace, by attaching alternating ends to the underlying plane or by leaving both ends unconnected. The microwave PCB in Figure 2.1 has three band pass filters built by connecting alternating ends to the underlying plane. Notice how wide the spacing is between segments.

In the mid 1980s, Martin Marietta developed a high-speed design handbook for the US military titled, "Design Guideline for Very High Speed Integrated Circuits". This document is listed in the bibliography of this book. As part of the research that supported those guidelines, test PCBs were built that contained traces running parallel that had "guard" traces between them and identically spaced structures without guard traces. The conclusion was:

"Increasing the conductor spacing around a critical signal line is recommended rather than grounding an interposed conductor. Grounding an interposed conductor does have a small effect on crosstalk, but does not eliminate it. The

grounded conductor is not a shield. Increasing the spacing is recommended because of its simplicity. Grounding an interposed conductor reduces via sites, possibly creating interconnect problems."

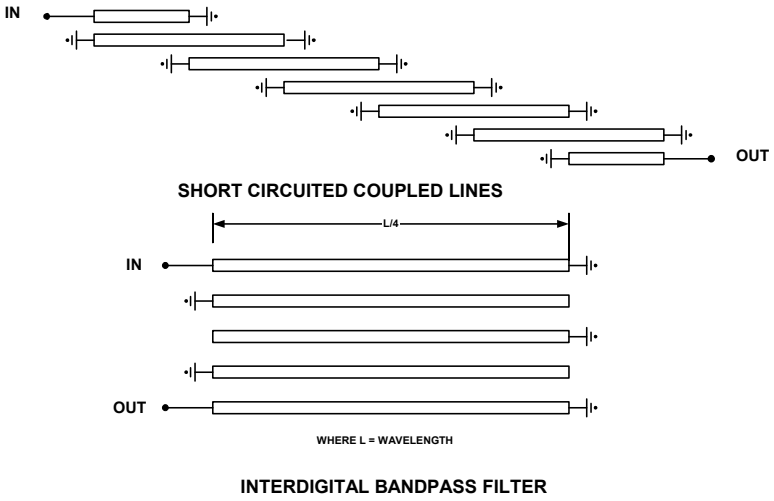

Side by side routing of traces can create band pass filters such as these

Figure 29.7. Methods for Creating Band Pass Filters Using Parallel Traces

Guard traces do not stop crosstalk. They have the potential for making it worse by creating an unwanted band pass filter. **Do not use them.**

Figure 29.8 is a sample of the artwork in a backplane designed in the late 1980s for a supercomputer of the time.

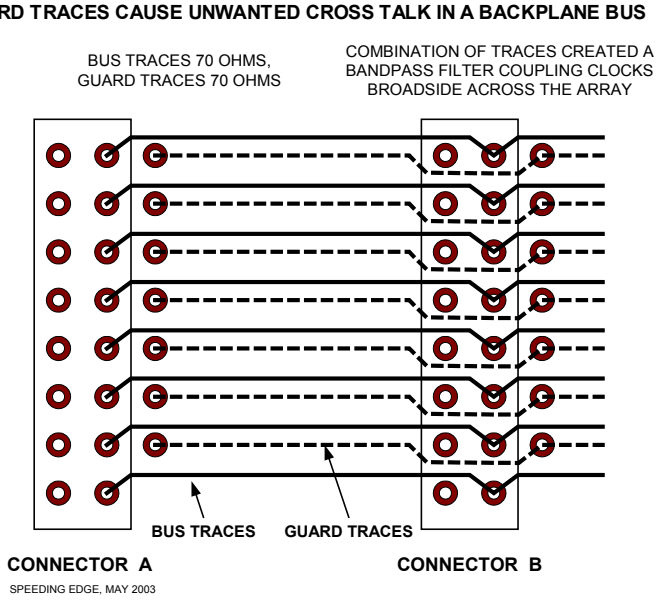

Figure 29.8. Artwork for a Failed Supercomputer Backplane of the Late 1980s

The engineers on this project were concerned that the impedance of the backplane transmission lines would overload the drivers. To reduce the possibility of this happening, the impedance of the backplane traces was set at 70 ohms. In order to achieve this high impedance on a stripline layer, it was necessary to increase the height above the nearest plane to 20

mils. From Figure 29.4 it can be seen that crosstalk gets larger as the distance from the nearest plane increases. To compensate for this, the designers added “guard traces” in between each bus line as shown in Figure 29.8. The length of these segments happened to be such that they resonated at the clock frequency of this computer. The result was unwanted coupling between signals running across the backplane of such magnitude that the computer was unstable. I was called in to help solve the problem. The only fix was to start over with proper spacing of the lines in the backplane. Starting over is rarely possible, so this design became one of the many Silicon Valley tombstones caused by incomplete understanding of transmission line behavior.

Figure 29.9 depicts crosstalk in one of the backplanes designed for the sonar system in the submarines that were featured in the movie, “Hunt for Red October.” My design company, Shared Resources, did the layout work for those backplanes. Again, the engineers on the project were concerned about the low impedance of the backplane transmission lines overloading the drivers and designed the highest impedance possible, 68 ohms. To do this, it was necessary to increase the spacing above the nearest plane to 13 mils. The crosstalk was twice what the system could tolerate and the product was “flaky”. I have these waveforms because of a weeklong investigation into the source of the flaky behavior. The problem was remedied by redesigning the backplane for a 55-ohm impedance. The result was a 2:1 reduction in crosstalk.

Actual coupling in a 34 layer back plane.

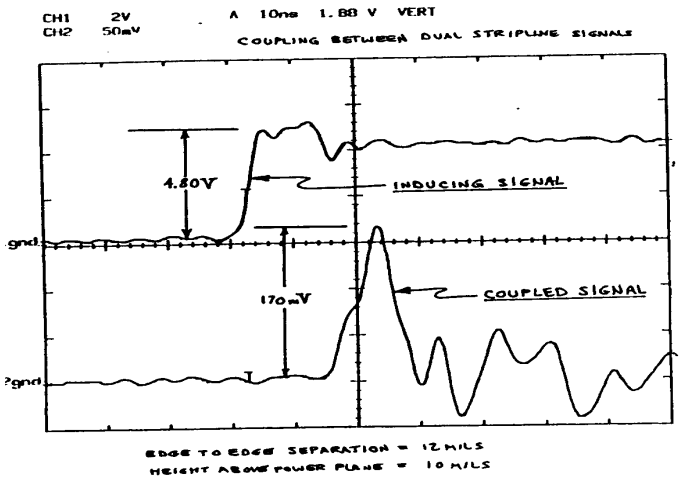

Figure 29.9. Crosstalk in a Submarine Sonar Backplane Due to Signal Layers Too Far from Planes

The purpose of this discussion is to illustrate the risk involved in striving for impedances higher than 50 ohms. Unavoidably, crosstalk will increase. Figure 29.10 illustrates how impedance and crosstalk increase as the distance to the nearest plane increases in the backplane design show in Figure 29.9. What can be seen is that crosstalk increases at a much higher rate than does the impedance. I know of no logic family intended for high speed signaling that is incapable of driving 50-ohm transmission lines. Designs should not be exposed to the extra crosstalk risk associated with higher impedance transmission lines.

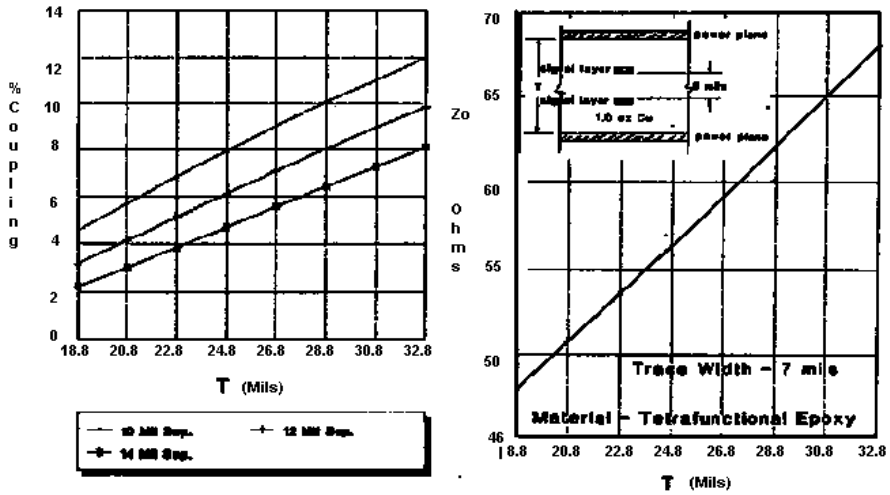

COUPLING AND IMPEDANCE vs THICKNESS

Note: Coupling increases at a much faster rate than does impedance as height above the plane is increased.

Figure 29.10. Impedance and Crosstalk vs. Height Above Nearest Plane in Submarine Backplane

All logic families intended for high-speed design are capable of driving 50 ohm transmission lines.

Determining How Much Crosstalk Can be Tolerated

A reasonable question to ask is how much crosstalk can be tolerated. In the chapter on design rule creation based on noise margin analysis, the method for establishing the allowable crosstalk will be covered. Tradeoffs will be made between the sources of noise,one of which is crosstalk.

CHAPTER 30: SINGLE ENDED SIGNALING

The term single ended signaling refers to a data path made up of a driver and a transmission line traveling over a plane or between a pair of planes and one or more inputs of loads. It is the least expensive method of sending logic signals from one place to another, since it requires only one wire and one signal pin for each data path with all data paths sharing the same "ground" plane.

Figure 30.1 is a typical single ended signal path. This circuit is from the BTL logic family, has multiple loads and is parallel terminated. Other variations are nets with single loads and with series terminations. Other logic families that have single ended logic paths include TTL, LVCMOS, ECL, GTL, LVTTL and SSTL.

In all of these cases, the signal is referenced at both ends as well as at the intermediate loads to the plane over which it travels. In most cases, this plane is logic ground but it does not have to be. In any case, all of the logic levels are referenced to the logic ground terminal of the power subsystem. Any noise that is injected into the ground path that connects the elements of the signal path erodes the logic levels that appear at the logic inputs. The most common types of noise are DC and AC voltage drops that occur in the ground structure because currents flow in this path and Vcc and Ground bounce that result as the current is required to charge up and discharge the logic lines as logic levels switch from 0 to 1 and pass back through the power leads of the IC.

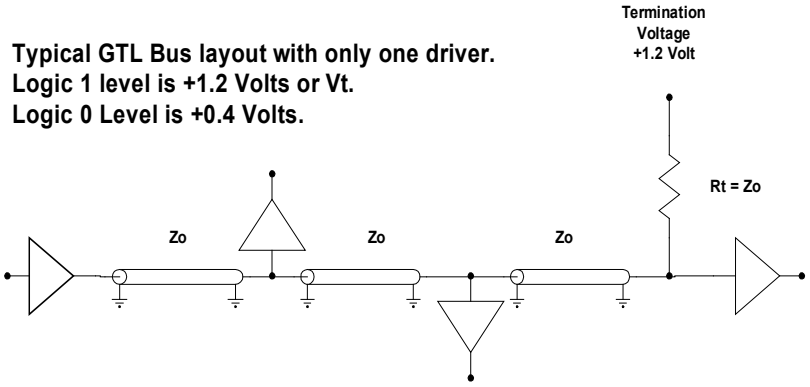

Figure 30.1. A Typical Single Ended Signal Path

In most systems, the current required to charge and discharge the parasitic capacitance of the transmission line and the inputs of the logic devices is the major source of "ripple" on Vcc as well as of both Vcc and Ground bounce. It is this current transient that limits the practical width that can be used for data buses. Figure 21.1 illustrates the current flow that takes place in a series terminated transmission line when the logic line switches from a logic 0 to a logic 1. For convenience, it is repeated here as Figure 30.2. The current waveforms for a parallel-terminated transmission line are shown in Figure 21.4.

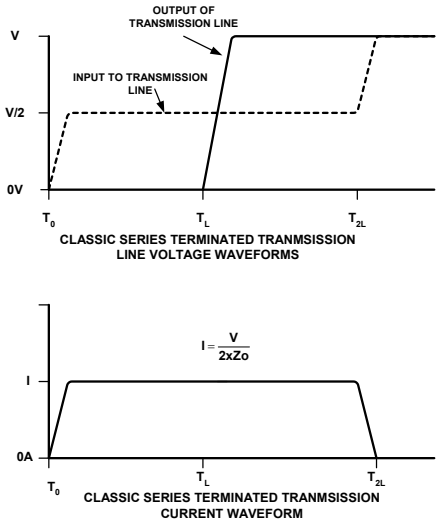

Figure 30.2. Current Flow From Vcc When a Series Terminated Transmission Line Switches From 0 to 1

The current, I, can be calculated for any logic family. For example, if the logic family is 3.3V CMOS, the peak current is approximately 33 mA per line. If a data bus is created using this form of logic, the peak current required by the bus when all bits transition from 0 to 1 simultaneously is 33 mA times the number of bits in the bus. Table 30.1 shows the peak current for several logic families at various bus widths.

PEAK CURRENT DRIVING 50 OHM LINES					
	16 BITS	32 BITS	64 BITS	128 BITS	256 BITS
1.8V	0.288	0.576	1.152	2.304	4.608
2.5V	0.400	0.800	1.600	3.200	6.400
3.3V	0.526	1.052	2.104	4.208	8.416
5.0V	0.800	1.600	3.200	6.400	12.800

Table 30.1. Peak Current Required By Various Data Bus Widths

The currents shown in table 30.1 must be supplied by the power system. When the logic lines switch from logic 0 to logic 1, this current must pass through the inductance of the Vcc leads of the IC package. When the logic lines switch from logic 1 to logic 0, the parasitic capacitance of the transmission lines must be discharged through the inductance of the ground leads of the IC package. This produces two unwanted types of noise. When the lines switch from 0 to 1, the current required to charge up the transmission lines must be drawn from the capacitance of the power subsystem. If this capacitance is not large enough and capable of responding at the speed of the switching edges, ripple will be present on Vcc. The power leads of the IC package always contain unwanted inductance. These current transients flow through this inductance producing Vcc and ground bounce, both of which can cause logic failures. Examples of ripple on Vcc from these current transients are shown in the section of this book that addresses power subsystems. Examples of Vcc and ground bounce are shown in the section that addresses IC packages.

> The worst-case switching current transient that the power subsystem must be designed for is when all of the bits of the largest data bus switch from 0 to 1.

As the speeds of logic circuits have increased, the two noise sources cited in the previous paragraph become so large that it is not possible to build packages that have inductances which are low enough and power subsystems that have sufficient high quality capacitance to contain this type of noise. When this happens, it is necessary to change from single ended switching to differential signaling.

> The current transients required to charge up and discharge logic lines as they change state are the primary sources of ripple on Vdd and the source of Vcc and ground bounce in IC package leads. They limit the practical width that can be used for single ended data buses.
> These current transients are a major source of logic failure.

CHAPTER 31: DIFFERENTIAL SIGNALING

Differential signaling is a method of moving data from a source to a load that has much greater immunity to noise than does single-ended signaling. It also creates less noise as it switches. It has been used since the beginning of large scale computing to move logic signals from one big box to another. With these advantages, it seems like the best choice for all logic paths. Why hasn't it been the signaling protocol of choice all along? The reason, of course, is that each data path requires two wires, two connector pins, two drivers and two receivers. This extra cost wasn't needed as long as data rates remained relatively low and both ends of the signal path were on the same ground plane. As single-ended data paths became very wide buses and the rising and falling edges of the logic signals became very fast, the switching noise discussed in Chapter 30 made it difficult to meet noise targets and differential signaling moved back into favor. The discussion of LVDS in chapter 28 is an example of this.

The features of differential signaling that warrant its use are often misunderstood. As a result of this, design rules have crept in that are either not beneficial or may be harmful. In order to clear up these misunderstandings, it would be useful to revisit the definition and benefits of differential signaling.

In its most basic form, differential signaling consists of two equal and oppositely changing signals that terminate on a receiver that is capable of detecting the voltage difference between those two signals and producing a single-ended logic 1 or 0 based upon which of the two signals is higher in voltage. The circuits that are used to receive these signals (such as ECL, the original differential signaling protocol) are capable of dealing with significant ground offsets. This is the original reason for using differential signaling--to allow successful logic signaling in an environment with ground offsets that are too large to allow single-ended logic circuits to properly operate. It is the main reason that differential signaling is in use again today. The LVDS contained in a laptop is a good example of this. The entire Ethernet network system is another.

> The primary benefit of differential signaling is its ability to cope with substantial ground offsets between the two ends of a data path.

ECL
LVDS
Biphase TTL Clock Trees
Ethernet Links
RS-422

Table 31.1. Types of Differential Signaling Circuits.

Table 31.1 lists several types of differential signaling protocols. All of them were created to cope with ground offsets between the two ends of the data path. As will be shown later, all of them take two equal and oppositely changing signals, detect when they cross and then generate a single-ended logic state change.

Differential signaling has been credited with other benefits that are inaccurate. One of the most common characteristics ascribed to differential signaling is that side-by-side routing of the pair of traces in a PCB provides common mode noise rejection. As illustrated in Figure 31.1, this is a false assumption.

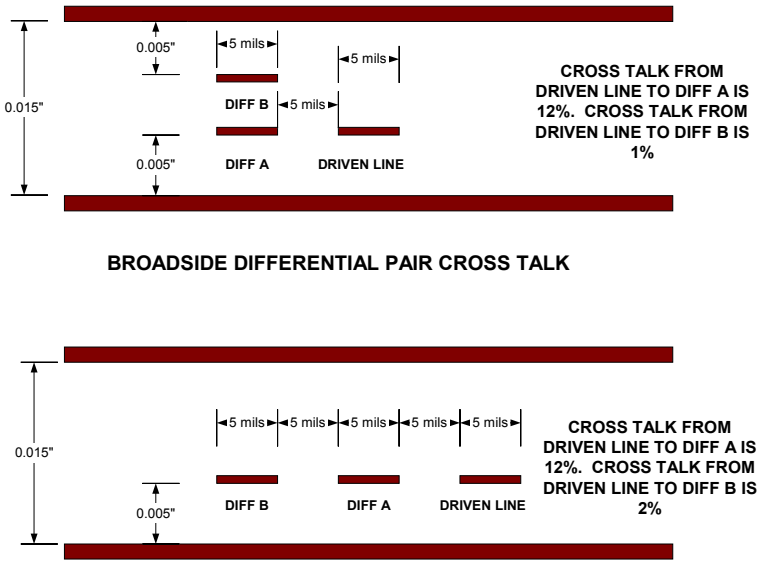

Figure 31.1. Side-by-Side Routing of a Differential Pair with a Noisy Line Routed Next to It

It is important to understand that the concept of side-by-side routing cannot guarantee common mode noise rejection in a PCB. Failure to understand this or to depend on it as a way to avoid noise from crosstalk, is one design methodology that results in a flaky system.

Why do people think that side-by-side routing of traces creates common mode noise coupling or noise rejection? It probably stems from the fact that two wires side-by-side in space have this characteristic. This works in free space but not next to a plane because common mode coupling requires the EM field that intercepts each wire be the same strength so that the same size noise signal is induced into each wire. In free space, both wires do experience the same size EM field. In a PCB, this is not possible due to the interaction of the field with the plane. The crosstalk curves in Figures 29.4 and 29.5 can be revisited to see how the strength of the EM field diminishes as it gets farther away from a signal line.

> Side-by-side routing of a differential pair in a PCB does not produce common mode noise coupling or noise rejection from nearby signals. **It produces differential mode coupling.**

Figure 31.2 is the classic ECL differential driver receiver pair that began the use of differential signaling over long, low quality connections. Differential signaling is being used in this circuit because the ground connection between Box A and Box B is not clean enough to permit single-ended signaling. Let's examine how this circuit functions. In Box A, the logic signal A is created as two copies, one the inverse of the other, both very tightly timed to each other. The objective is to deliver each of these two signals to the appropriate transistor base in Box B with minimum distortion. How is this done? This is a simple problem. Each signal needs to travel on a transmission line that is properly terminated. There is nothing special about this operation. It is done the same way as with other logic signals. Each transmission line is routed in such a way that its impedance is uniform and it is kept away from noise sources, such as other transmission lines, and is terminated in its characteristic impedance to Vtt. (It should be noted that the other member of the pair is also a noise source.)

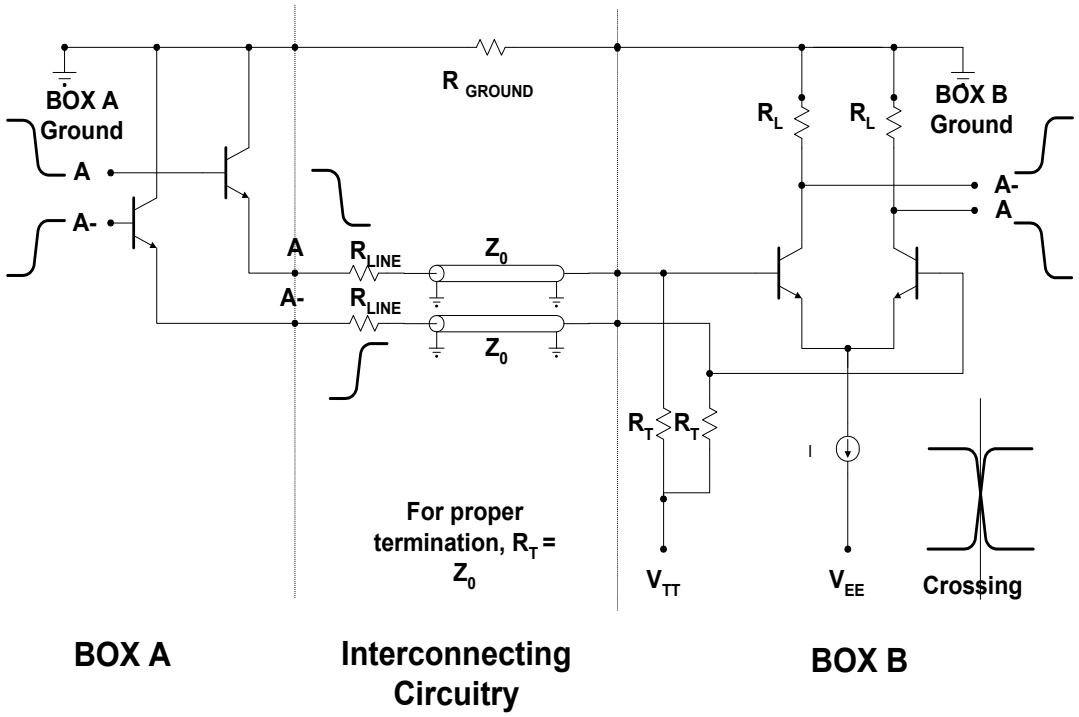

Figure 31.2. An ECL Differential Signaling Path

The real functionality of this circuit appears in Box B. The two signals terminate on a bipolar transistor pair that is joined at their emitters creating an emitter-coupled pair. This circuit is often incorrectly called a differential amplifier. This is a misnomer. This circuit is actually a current switch. The current, I, is steered completely up one side or the other of the current switch. It is never divided between the two sides as happens with a differential amplifier. If this condition occurred, the logic states would be uncertain.

So, the objective is to apply a voltage difference between the two bases that will insure complete current steering to one side or the other of the current switch. For ECL, the voltage difference required to do this is on the order of 15 millivolts, a very small difference. Since the starting signal in box A is approximately 1000 millivolts, significant attenuation can occur while a valid logic signal is still being delivered.

The real reason for using a differential signaling protocol such as this is to allow for ground offsets between the two ends of the signal path. How is this accomplished? Notice that the current switch is riding on top of a current source. A current source has an “infinite” output impedance. This means that no matter what the load impedance is, the value of the current it generates does not change. As a result, the emitters of the current switch can move positive and negative and it can still do its job. The current switch can ride up and down with the circuits in Box A. The two collectors of the current switch transistors are also current sources. This also allows them to move up and down. The collectors can move positive until they approach a voltage value below ground that is I x R_L below zero volts. The combination of these two “current” sources allows the current switch (emitter-coupled pair) to move positive and negative with the circuits in Box A over a broad range of voltages referred to as the “compliance window”. For ECL, this compliance window is approximately 1.5 volts. This is how the ground offset problem is addressed.

Functionally, what is the circuit in Box B really doing? It is detecting the moment at which the two voltage waveforms, A and A-, cross each other as they change logic states. At that moment, the current, I, is switched to the other side of the current switch and the logic levels at the two collectors change states. The crossing is converted into a logic state change. So, the circuit is actually a **crossing detector**. The secret to its proper operation is to preserve the crossing. That is the primary design consideration. How is this accomplished? It is accomplished by making sure that the two edges arrive at the same time. This is achieved by making sure the two signal paths are the same electrical length.

How close to the same length do these signal paths need to be? Figure 31.3 depicts a routed pair and where the waveforms cross. If the lengths are unequal, the switching edges will be skewed with respect to each other like the two waveforms at the bottom of Figure 31.3. The current switch will still detect a crossing and create a logic state change. However, due to the fact that the crossing occurs where the slope of the waveforms is somewhat flat, the precision with which the actual time of crossing is detected from one data transition to the next will be lessened. This will result in jitter as the time of logic state change will vary from one bit to the next. In order to reduce jitter to a minimum, it is necessary to

insure that the two edges cross in the “straight” part of their switching edges. This is shown in the upper pair of waveforms in Figure 31.3.

Calculating the degree of length mismatch that can be tolerated is a matter of measuring the length of these rising and falling edges (they should be the same in a differential circuit) and converting this time into length by multiplying it by how fast the signal travels on the transmission lines.

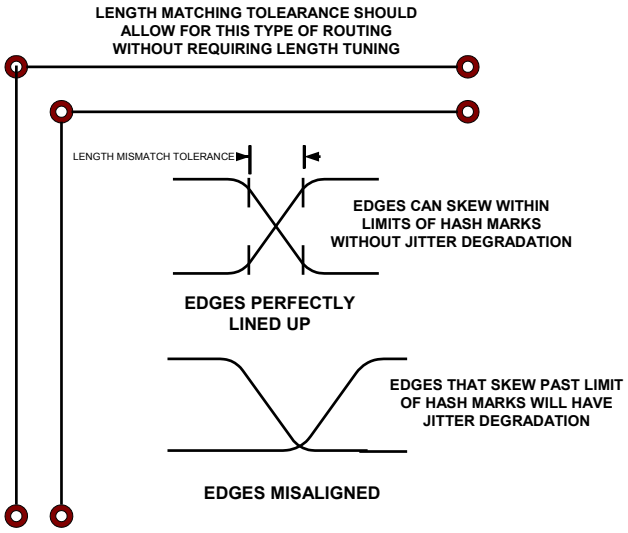

Figure 31.3. Differential Switching Edges and Routing Choices

To perform this calculation, the rise or fall time at the receiver should be used, not the rise or fall time at the driver. When this calculation is complete, it will be seen that at 2.4 GB/S, a length mismatch of as much as 0.300 inches is tolerable. This amount of mismatch will permit the kind of routing shown in Figure 31.3 without requiring that extra length tuning or matching steps be done. This is of significant benefit to the person doing the PCB layout as it saves precious layout and development time.

The allowable timing mismatch for the LVDS used in a PC laptop is 400 picoseconds. **If this time is converted to length, the result is 2.4 inches!** The data rate for this protocol is 75 MHz on each pair or a bit period of 13.3 nanoseconds. With this much margin, it is almost impossible to lay out the two members of an LVDS differential pair such that the length matching requirement is violated.

> The accuracy of length matching required by a differential pair is determined by calculating or measuring the fastest rise and fall time of the two signals as they arrive at the receiver. This time is converted to length using the signal velocity in the transmission line.

Notice that as the logic state changes in the differential circuit in Figure 31.2, the current that flows from the power supply remains steady. All that happens is the path it takes through the circuits in both Box A and Box B change. This is one of the more significant advantages of differential signaling over single ended signaling and the reason that backplane data paths are being changed from single-ended parallel data buses to differential serial links. The magnitude of current flowing through the power leads of the IC package remains constant. As a result, the Vcc and Ground Bounce or simultaneous switching noise (SSN) that develops across the inductance in the power leads is zero. This is yet another advantage of differential signaling and a major reason why it is finding its way into new memory and backplane architectures.

Differential signaling steers the current it draws from V_{CC} or V_{DD} from one side of the circuit to the other. The current drawn for the power subsystem remains constant. This reduces Simultaneous Switching Noise, SSN, to nearly zero. This is a major advantage of differential signaling.

Some other observations can be made based on the circuit in Figure 31.2. An important observation is that the two signals traveling from Box A to Box B are completely independent of each other. Neither is aware the other exists. Neither needs to be next to the other to do its job. In fact, there is no beneficial interaction between the two. Therefore, side-by-side routing does not provide any special benefits. It is often said that tight coupling, meaning routing the two traces very close to each other, has a benefit. It is not clear where this idea started, but it is incorrect. When these two signals are routed very close to each other, they do interact, but in an unfavorable way.

When one signal is routed close to the other, several undesirable things happen. The first is the impedance of each line is lowered by the presence of the other. As a result of this, the terminating resistor placed at its end will need to be a lower value and the resulting voltage across it will be smaller. If this is not possible, each trace will need to be made narrower in order to restore the original impedance. As each trace gets narrower, the skin effect loss increases. Another undesirable side effect of routing pairs tightly together is that it induces crosstalk from one line to the other. This crosstalk is not beneficial. It actually degrades the switching edges resulting in a poorer signal at the receiver.

Differential pairs do not need to be routed side by side for proper operation. In fact, if they are routed too close to each other, destructive coupling occurs that causes the edges to be slowed down.

If differential pairs are routed side by side, it is imperative that they be kept far enough apart that they don't destructively interact.

The foregoing discussion illustrates the concept of differential signaling. However, ECL differential signaling is not the most commonly used technology in current use. In fact, ECL or PECL as a logic family is virtually gone from the design scene. The only places where the ECL or PECL logic family appears is at the interfaces to older designed transceivers. In this application, some form of LVDS, such as differential signaling, is replacing it.

The most common form of differential signaling in current use is LVDS or one of its derivatives. Figure 31.4 depicts a typical LVDS differential signal path.

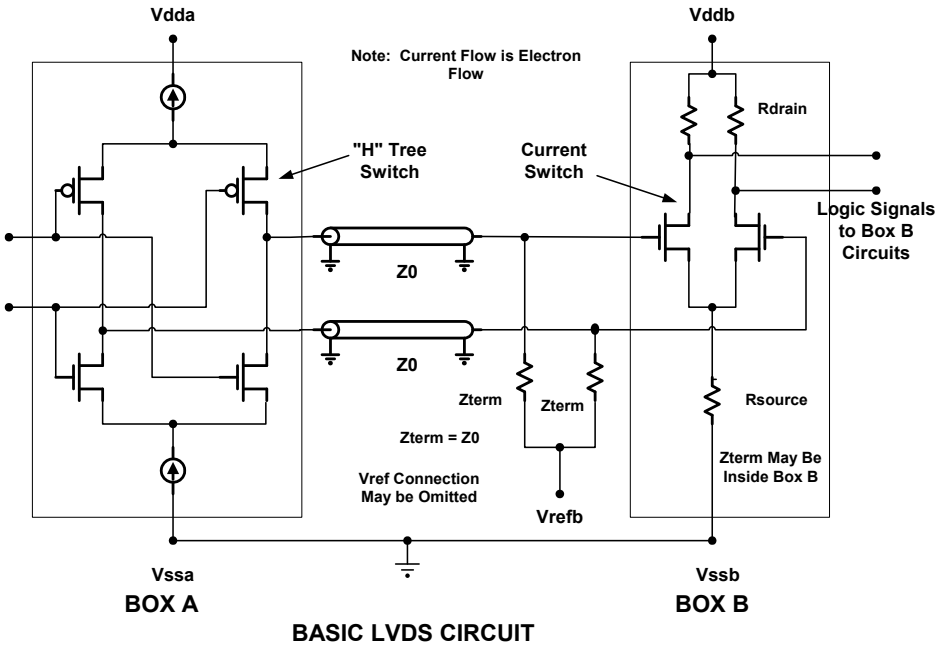

Figure 31.4. An LVDS Differential Signal Path

The major difference between an LVDS differential signaling circuit and an ECL differential signaling circuit is depicted in Box A. A constant current is passed through an "H" tree switch made up of four transistors. These transistors are switched on in pairs allowing the current to flow out of Box A through one transmission line and into the V_{ref} node of Box B through the termination resistor. This current develops a voltage drop across the termination resistor. A second current flows out of Vref through the second terminating resistor, it's transmitted back into the H switch and through the second current source. This develops a voltage drop across the second terminating resistor that is equal to and opposite the first voltage developed across the other terminating resistor. This is illustrated in Figure 31.5.

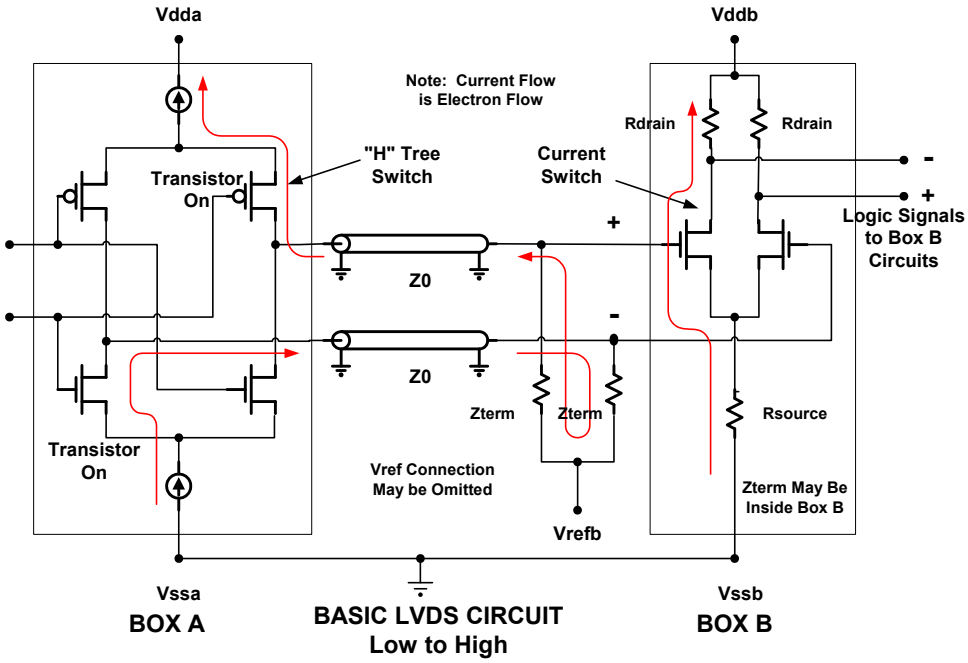

Figure 31.5. An LVDS Differential Signaling Circuit Showing Current Flow for One Logic State

Notice that the magnitude of the difference voltage that appears at the input to the current switch in Box B is determined by the magnitude of the current generated in Box A and the size of the terminating resistors in Box B. For standard LVDS circuits, the magnitude of the current is 4 milliamps. Each transmission line is 50 ohms. Each transmission line is

terminated in 50 ohms to V_{ref}. 4 milliamps multiplied by 50 ohms results in 200 millivolts across each resistor or a total difference voltage of 400 millivolts.

There is no net current flowing into or out of the V_{ref} terminal in Box B. As a result, this connection can be and usually is omitted. This results in two terminating resistors being placed in series across the ends of the two transmission lines. When this is done, it is common to use a single resistor of value two times Z_0. A common error is to call this a differential termination and then seek to create a differential impedance between the two transmission lines. This is not necessary or of any value. The two transmission lines should each be 50 ohms with as little interaction as possible occurring between them.

Figure 31.6 shows the current flow after a logic state change.

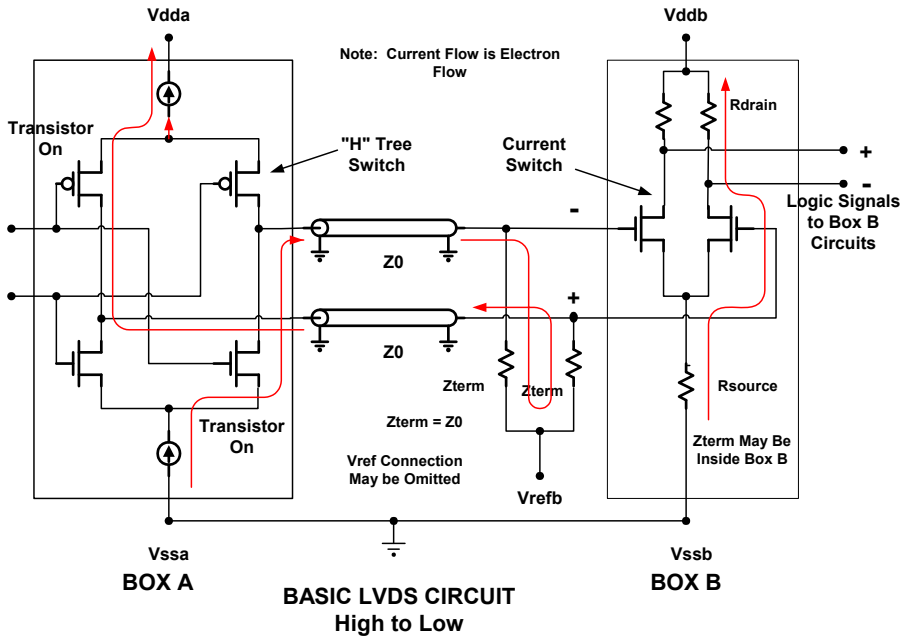

Figure 31.6. An LVDS Differential Signaling Circuit Showing Current Flow for The Second Logic State

As can be seen from the above three figures, the circuit in Box B is a current switch. The objective is the same for this circuit as for the ECL current switch--to completely steer the current up one side or the other. The voltage difference to do this with CMOS circuits is on the order of 200 millivolts. For LVDS, the starting voltage at the driver is 400 millivolts. Half of the signal can be lost due to reflections, crosstalk and loss in the transmission lines and the circuit will still operate correctly.

Compensating For Large Losses

Variations of the LVDS circuit discussed above are used in all of the current high-speed serial links such as PCI Express, Hyper Transport, Fiber Channel, InfiniBand and the differential serial links used for 2.4 GB/S data paths in protocols such as XAUI. All of these protocols operate at very high data rates. At these data rates, losses in the dielectric material that makes up the PCB, skin effect loss in the traces themselves, and losses as the signals pass through connectors can be significant. These losses are often dealt with by using some form of pre-emphasis. Pre-emphasis is accomplished by adding circuits to the H switch in Box A that increase the amplitude of the signal that is sent down the transmission line.

Both skin effect loss and dielectric loss are frequency sensitive. The higher the frequency, the more loss a signal will experience. For logic signals this loss is data-pattern sensitive. The more often the data bits switch, the more loss there will be because the frequency components in a logic signal are higher the faster the data pattern changes. The less often the data bits switch, the smaller the loss will be. Because the loss is sensitive to data pattern, it is necessary for the circuit in Box A to monitor the last data bit sent and add pre-emphasis when the next bit is different from the last bit and omit pre-emphasis when the next bit is the same as the last bit.

Losses at high frequencies and how to deal with them will be discussed in Volume 2.

Managing The Transition From A Differential Pair Cable To A PCB

It is common to send differential signals from box to box using shielded or unshielded twisted pair cable. The impedance of such a transmission line is measured between the two wires in the differential pair. The question of how to deal with the transition from the cable to the PCB arises. Usually, the differential impedance of such twisted pairs is on the order of 100 ohms. In order to transition the signal from the cable wires to the PCB traces without introducing an impedance discontinuity, it is necessary for the cable to see the same 100 ohms between the two wires. This is accomplished by creating two independent, single-ended transmission lines or traces on the PCB that are each Zdiff/2. There is no need to be concerned about the differential impedance between the two traces. Each of these two transmission lines needs to be terminated by Zdiff/2 to the specified termination voltage of the logic being used.

If the two traces on the PCB are routed close to each other, each will lower the impedance of the other. If this is done, it will be necessary to narrow the width of each trace to arrive back at the 50 ohm single ended impedance needed by the cable.

If the twisted pair is shielded, the shield is part of the Faraday cage. Its purpose is to prevent an EMI problem should there be a common mode signal riding on both members of the twisted pair. Since the shield is an extension of the Faraday cage, it should be connected to the Faraday cage at the point where the cable exits the cage. If there is no Faraday cage, the shield should be connected to logic ground.

CHAPTER 32: THE POWER SUBSYSTEM

By John Zasio

The power subsystem is required to deliver adequate voltage to the load circuits for all operating conditions and with all manufacturing tolerances. An ideal power subsystem, as depicted in Figure 32.1, would be a fixed voltage source with zero output impedance at all frequencies. The ideal power subsystem would deliver a constant voltage to the load independent of load current changes.

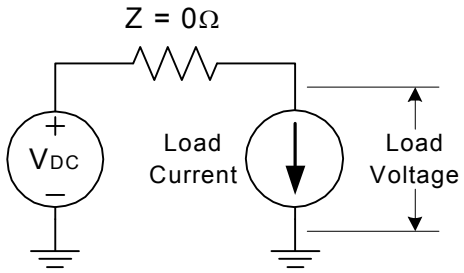

Figure 32.1. Ideal Power Subsystem

The electronic products we build today consist primarily of CMOS integrated circuits mounted on epoxy glass PC boards and are interconnected with transmission lines. Over time, the CMOS technology continues to decrease in feature size, resulting in a decrease in power per logic gate; an increase in gate count per chip; an increase in clock frequency and a lower power supply voltage. With lower power supply voltages, higher power per chip, and higher clock frequencies, the power distribution system becomes more difficult to design.

Typical Products

A typical product today contains microprocessors (uP), ASICs, FPGAs, SRAM, DRAM and high speed interconnects.

Current CMOS technology is built with a minimum feature size of 0.13 um. Logic chips (uP, ASICs and FPGAs) require a core voltage of 1.2V to 1.5V, depending on the design. I/O voltages tend to be different and higher than the core.

Microprocessors have power dissipation ranging from a few watts up to 50 watts and core clock frequencies greater than 2.0 GHz. They have I/O speeds running up to 900 MHz.

High-end ASICs tend to be somewhat larger. A large 0.13um CMOS ASIC can have more than several hundred million transistors. This will include several hundred thousand DFFs (D-Flip-Flops) running off the same clock tree.

SRAMs and DRAMs tend to have higher voltages. Current DRAM technology is DDR-1 (Double Data Rate, Version 1) that operates at 2.5V and has data rates of 400 Mb/s (Mega bits per second) per pin. DDR-2 will be here by the end of 2003, operate at 1.8V and have data rates of up to 800 Mb/s.

In today's power subsystems, there is usually a small amount of glue logic and a little analog logic that will use 3.3V as well as a very small amount of 5V circuits. The result is a several hundred-watt PC board with four power supply voltages. The low voltage supplies must supply high currents of up to 50A per board.

Older products tended to have the power supplies built into the chassis and delivered power to several PC boards. This has become impractical due to the high-current, low-voltage requirements of the cards. High-end products today use a bulk power supply to convert line voltage to 48 Vdc that is distributed to the PC boards. Small DC/DC converters are used on the PC boards to convert the 48 Vdc into low-voltage with high current.

Design Issues

1. CMOS logic circuits consume power only when they switch states. This causes the power supply current to change in both magnitude and frequency depending on the activity of the chip. The current can go to zero in the test or reset state. Typically, once out of reset, the clocks are on and continue to run at a fixed frequency. However, the logic can go from idle to full activity at any rate from extremely low frequency up to the clock rate. A typical uP or ASIC will consume 25% of the maximum power in the clock tree. The remaining 75% of the power is variable.

2. ASIC clock trees tend to clock a very large number of DFFs at exactly the same time. This produces a very large current spike once per clock cycle. The result is a load frequency spectrum that is much higher than the clock frequency.
3. ASIC I/O power (Vddq) is usually a different voltage than the core and can go to near zero in the idle state. There is usually one clock pin for many data pins so the clock power may be only a few percentage of the total.
4. DC/DC converters have very low output impedance but only at low frequencies. At typical 1.8V 40A converter will have less than a 1 mΩ impedance from DC up to about 2 KHz. Above 2 KHz, the output impedance of the converter increases directly proportional to the frequency of the load current.
5. Decoupling capacitors can be used to provide a low impedance to load changes at higher frequencies than the DC/DC converter can track, but all capacitors have a small amount of series inductance that is proportional to their physical size. Even the small ceramic capacitors have enough inductance to limit their use to frequencies of no more than a few hundred MHz.
6. PC board capacitance between the voltage and ground planes is required to decouple at frequencies above 200 MHz.
7. The DC resistance of the PC board power plane becomes significant with a low voltage DC/DC converter delivering up to 40A from a single pin.

Requirements

Most ASICs and microprocessors require a power distribution system to deliver a voltage with a tolerance of +/- 5%. This 5% includes all conditions of line voltage, load current, temperature and manufacturing tolerances of the components being used. A practical approach to this problem is to break up the 5% into several categories and design specifically for each category.

Category	Description	Assignment
1.	Power Module (including drift with age)	+/- 2%
2.	DC Resistive Distribution Drop	+/- 1%
3.	AC Ripple (peak-to-peak)	+/- 2%

Total		+/- 5%

Example

For example, let's examine a PC board with ASICs that run on a core supply of 1.5V and have a maximum load current of 20A. The clock trees consume 25% of the power and 75% is activity dependant. Table 32.1 depicts the power distribution for the PCB.

Category	Tolerance	Voltage	Current	Required Impedance
Power Module	+/- 2%	+/- 30 mV		
DC Drop	+/- 1%	+/- 15 mV	20 Adc	0.75 mΩ
AC Ripple p-p	+/- 2%	+/- 30 mV	15 Aac	2.0 mΩ

Table 32.1. Power Distribution Example Requirements for a PCB

The required impedance is calculated using Equation 32.1. as follows:

$$Z = (\text{Voltage Change}) / (\text{Current Change})$$

$$\text{Ie: } Z = 15\text{mV} / 20\text{A} = 0.75\ \text{m}\Omega$$

Equation 32.1. Impedance Calculation for Power Distribution System

The load current can change at any frequency from DC up to the clock frequency and it will have harmonics well above the clock frequency. Therefore, the required impedance of the distribution system must also cover this frequency range.

Power Distribution Circuit

The components available for the power distribution circuit are not ideal. They all have frequency dependent impedance. Therefore, a variety of components must be used to cover the frequency range of interest. Figure 32.2 depicts the circuit diagram for a typical power distribution system.

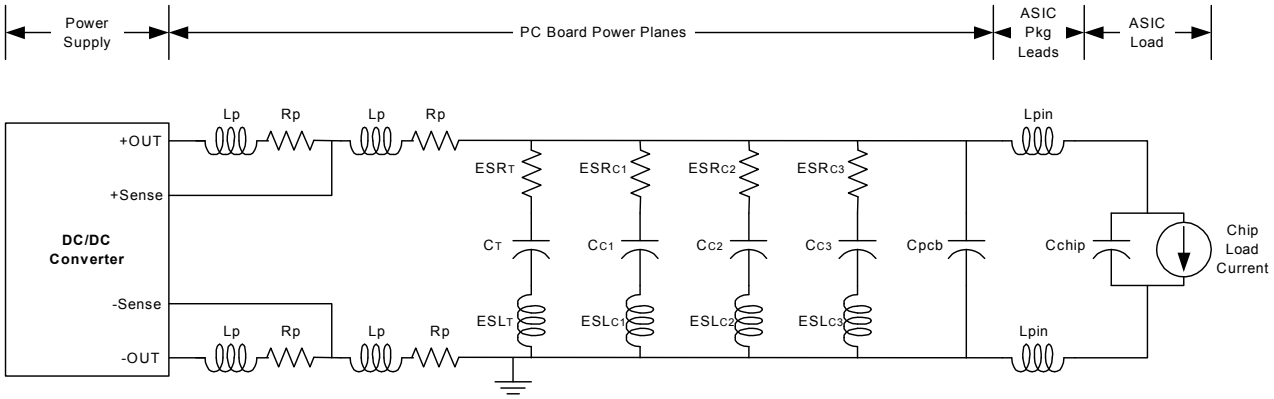

Figure 32.2. Circuit Diagram For A Typical Power Distribution Circuit

The PCB power planes are a critical part of the power distribution system. They provide the low DC resistance, the high frequency capacitance, and the low inductance connection to the discrete capacitors. The power plane resistance (Rp) and inductance (Lp) are not exactly as shown in the above circuit schematic. Instead, they are distributed between all the capacitors and the loads.

The DC/DC converter provides the power and the low-frequency response to load changes. Decoupling capacitors are used to provide a low AC impedance between the low-frequency point where the DC/DC converter cuts off to the high-frequency point where the PCB planes become effective.

There are usually several types of decoupling capacitors and large quantities of each type are required. Capacitors with large values (and large physical size) cover the lower frequencies, while capacitors with small values (and small physical size) cover the higher frequencies.

There is an inductive connection to the integrated circuit (IC) chip due to the package leads. Inside the chip there is a significant amount of capacitance formed by the SRAM cells and the circuits that are not switching. A typical IC chip may have 20 nF of capacitance. This is crucial because it is used to absorb the very high frequency components of the clock current spike.

DC/DC Converter Output Characteristics

Figure 32.3 depicts an oscilloscope waveform showing the output voltage response of a small on-board DC/DC converter with a load current change of 10 Amps. The load current is a 1.0 KHz square wave with a rise time of 50 ns. The time scale is 100 microseconds per division. The amplitude of the voltage change spike is -105 mV. 100 microseconds after the current change, the power supply voltage has recovered to a -26 mV drop. Part of the drop is due to the DC resistance of the power distribution planes. This example shows that the DC/DC converter recovers from a load change in about 200 microseconds. This is equivalent to a frequency response of about 2.5 KHz. At higher frequencies, good quality decoupling capacitors are needed.

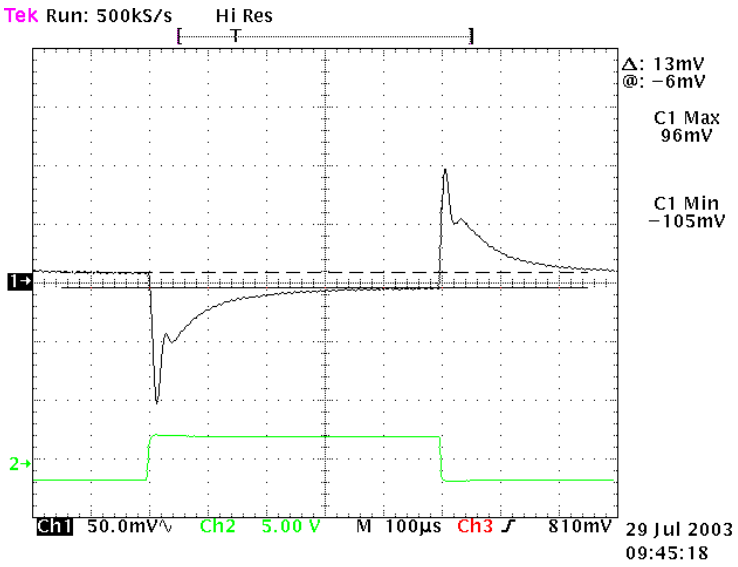

Figure 32.3. DC/DC Converter Output Response to a 10A Load Current Change

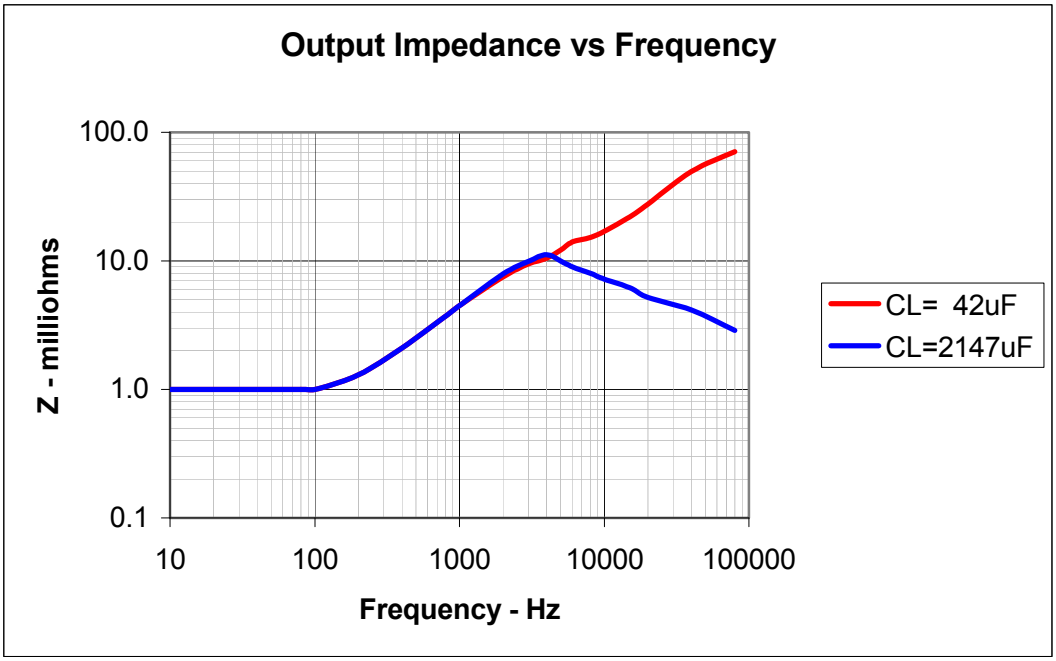

Figure 32.4. DC/DC Converter Frequency Response

Figure 32.4 is a plot of the output impedance vs frequency for the same DC/DC converter. The converter is a Quarter Brick 1.5V unit rated for a 30A load. The test was performed with a DC load of 10A and an AC Sine Wave load of 200mA. The measured output impedance of the supply mounted on the board is 1.0 milliohms. The supply itself has an impedance of about 0.1 milliohms measured at the sense point with a slow changing DC load. The 1.0 milliohms in the above plot is the DC resistance of the power planes in the PC board.

The Red curve shows the response with capacitive load (CL) of 42uF of high frequency decoupling capacitors. The Blue curve shows the response with a quantity of six T520 330uF capacitors (Total Measured CL=2147uF). It is necessary to add more large capacitors to the board if the required impedance must be kept low.

PC Board Stackup

A high performance PCB typically has 20 or more conductive layers. Half of these layers is used for signal wires (controlled impedance transmission lines) while the other half is for power distribution. Voltage and ground planes are

placed adjacent to each other in order to create sufficient high frequency parallel plate capacitance. These power planes also act as return current paths for the signal wires.

Figure 32.5 depicts a cross section view of a high performance PCB. The long horizontal bars are power and ground planes. The short horizontal bars are signal wires.

Figure 32.5 depicts a 22-layer PCB with 10 signal layers, 10 power layers and two outer layers for surface mount component pads. Dielectric thickness can be as low as 3 mils with conventional processing. Copper thickness is 1 ounce or ½ ounce copper. The adjacent voltage/ground planes act as a capacitor with no lead inductance. This produces the capacitance necessary for frequencies above 200 MHz.

The adjacent voltage/ground planes also act as a very low impedance transmission line to connect from the IC to the discrete decoupling capacitors.

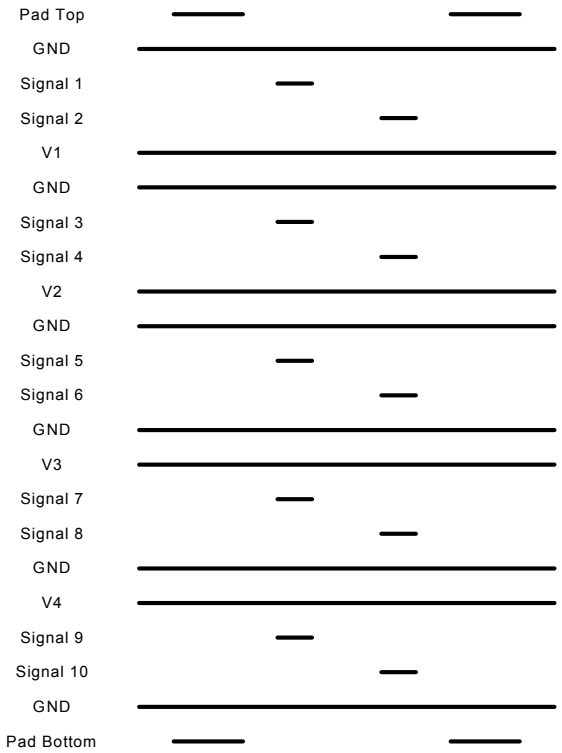

Figure 32.5. Typical High Performance Board Stackup

CHAPTER 33: POWER DISTRIBUTION DC DROP

By John Zasio

In a typical PCB, there are many ground planes to return the DC current to the power supply but only one plane to distribute the current for each voltage. However, these ground planes must carry the return current for all the voltages. There is a desire to get almost zero DC voltage drop across these planes but with practical limits on copper plane thickness, this is not possible. Therefore, designs must be executed to get an acceptable drop with the copper available.

Required Information

Based upon today's copper plane thicknesses, Equations 33.1 and 33.2 can be used to determine the rectangular plane and circular ring plane resistance respectively.

Copper Plane Characteristics:

	1/2 oz	1.0 oz
Plane Thickness	0.55 mils	1.15 mils
Sheet Resistance (R_S)	1.26 mΩ/sq	0.59 mΩ/sq

$$\mathbf{R = R_S \times L / W}$$

Where: R_S is the resistance of a square section of the copper plane measured from opposite ends of the square; L= length and W = width.

Equation 33.1. Rectangular Plane Resistance

And, the circular ring resistance can be calculated using Equation 33.2.

$$R = \frac{Rs}{2\pi} * \ln \frac{R_2}{R_1}$$

Where: R_1 = Inner Radius, R_2 = Outer Radius

Equation 33.2. Circular Ring Plane Resistance

DC Drop Example 1: Long Narrow PCB With Power Feed At One End

In order to understand the concept of DC drop, a long, narrow PCB with uniform loads and the power supply connected to one end of the board will be used as is depicted in Figure 33.1. The 1.5V, 20A example from Chapter 32 is used. This PCB has a length of 20 inches, a width of 5 inches and ½ ounce copper planes. This PCB is also assumed to have an ideal voltage source of 1.5 volts.

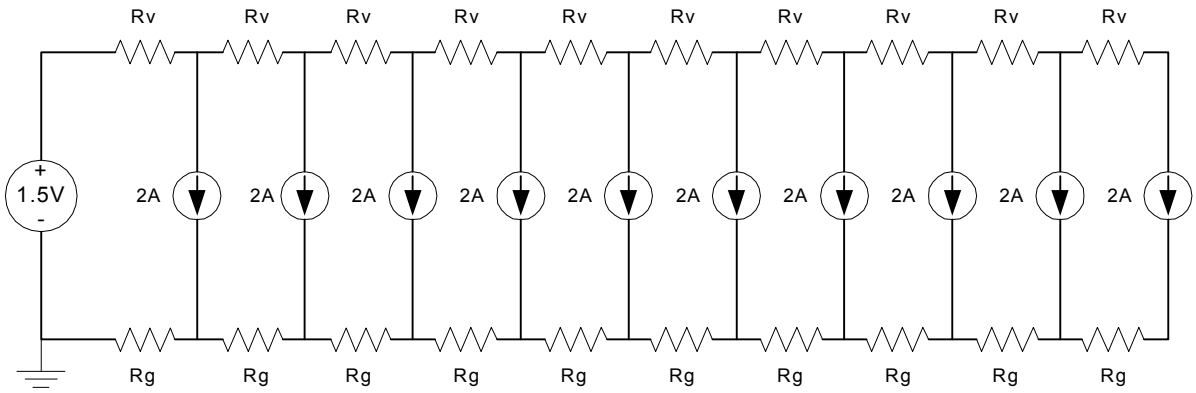

Figure 33.1. Power Distribution Drop – Example 1

In Figure 33.1, it is assumed that all of the power supplies on the PCB are not drawing the maximum load and there is the equivalent of two ground planes for the return current. Each resistor in the above circuit is for a 2-inch length by 5-inch width section of the PCB. Rv is resistance of the voltage plane and Rg is resistance of the ground plane.

Therefore: Rv = 1.26 mΩ/sq x 2" / 5" = 0.504 mΩ
Rg = 0.5 x 1.26 mΩ/sq x 2" / 5" = 0.252 mΩ

All of the current (20A) flows through the first pair of resistors closest to the voltage source. The second pair of resistors has 18A and each resistor down the line has the current decreased by 2A. This produces an exponential curve for the voltage drop along the planes as depicted in Figure 33.2.

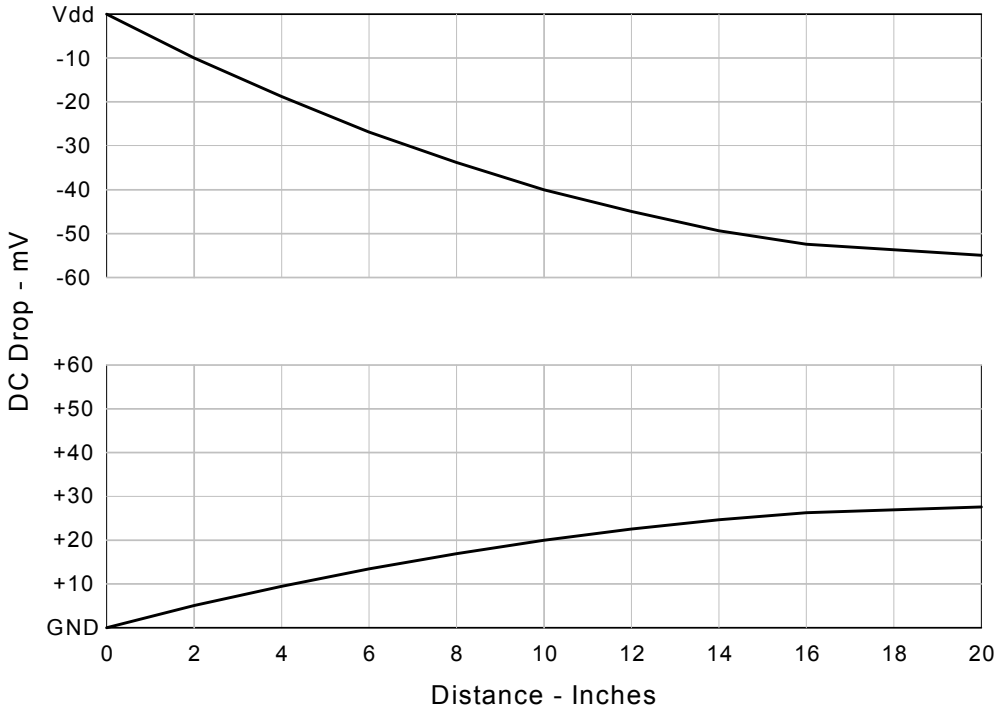

Figure 33.2. Example 1 -- DC Voltage Drop vs. Distance from Source

In Figure 33.2, the power plane voltage drop at the far end of the PCB is 55 mV and the voltage rise on the ground plane is 25 mV. This is far from the target of 15 mV. The original target was +/- 15 mV but, in this example, the drop is always minus.

DC Drop Example 2: Long Narrow PCB with Remote Sense

Modifying the circuit depicted in Figure 33.1 to include a real power supply with remote sensing, yields the second

example of power distribution drop as depicted in Figure 33.3.

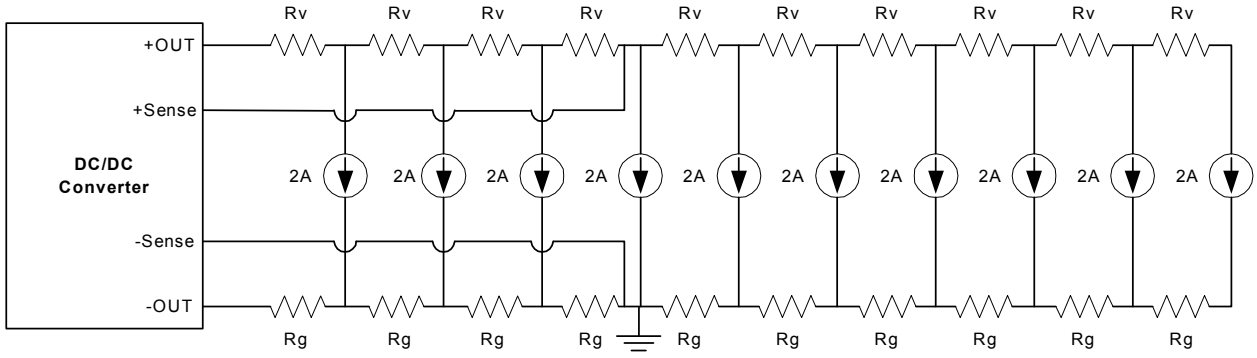

Figure 33.3. Power Distribution Drop – Example 2

In Figure 33.3, the sense terminals of the converter feed a regulation circuit that adjust the voltage on the output terminals so that the specified voltage exists at the point where the sense terminals touch the power planes. This increases the voltage on the +OUT terminal above VDD and lowers the voltage on the -OUT terminal to a value below GND. This results in the exponential curve for the voltage drop as shown in Figure 33.4.

The voltage on the VDD plane is now +23 mV at the first load and -23 mV at the last load. On the GND plane, the voltage is -11 mV at the first load and +11 mV at the last load. This is more than a 2:1 improvement because there is additional voltage drop in the wires from the converter to the first load.

The DC voltage drop is more than the desired total but this example is a very long thin board with the power source at one end. A real board would have an aspect ratio that is more square and would not have the power source at the end.

The sense wires feed a high impedance input to an amplifier inside the DC/DC converter and this input is very sensitive to electrical noise**. Therefore, the sense wires should be routed as a differential pair to reduce the effect of noise coupling.**

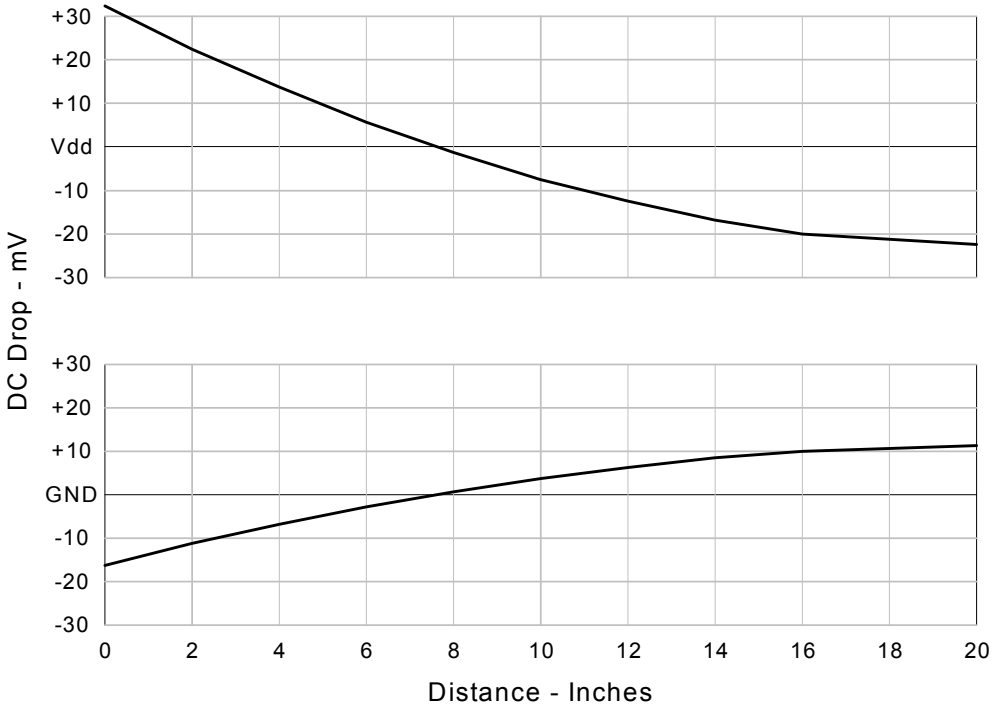

Figure 33.4. Example 2 – DC Drop vs. Distance with use of Remote Sense

DC Drop Example 3: Rectangular PCB with Power Source on Board

A real PCB is usually a rectangle with less than a 2:1 aspect ratio and it has the DC/DC converter placed away from the edge of the board as depicted in Figure 33.5.

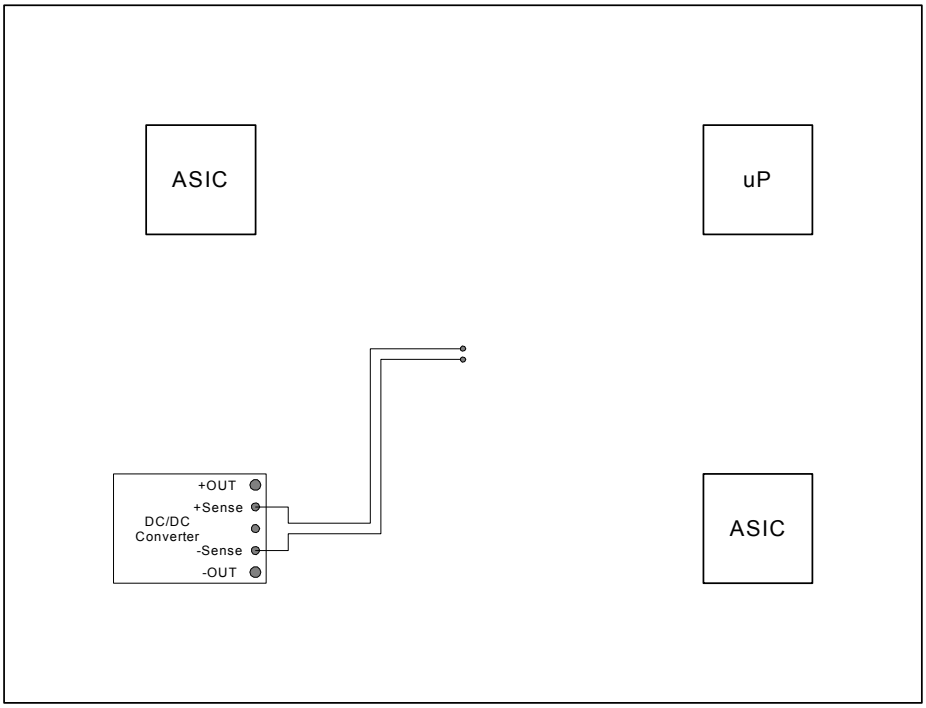

Figure 33.5. Example 3 – Rectangular PCB With Power Source On Board

In Figure 33.5, the nominal voltage will appear at the center of the PCB at the contact point of the sense wires to the power planes. The current will spread in a radial pattern from the power pins of the DC/DC converter. There may be a significant voltage drop from the power pins of the converter to the sense point but a low voltage drop from the sense point to the loads.

High Current Connections to PC Boards

An industry standard ¼ brick DC/DC converter is approximately 2.80" x 1.45" in size and has through hole power output pins with a diameter of 62 mils. A reasonable size via for this pin will have a diameter of 75 mils. In the example cited here, there is a maximum DC load of 20A. This is a very high current to drive radially from a 75 mil via into a ½ oz copper plane. The spreading resistance and resulting voltage drop will be significant.

Using Equation 33.2, we can estimate the spreading resistance from the pin and the resulting voltage drop. Assume a circular ring resistor with an ID of 75 mils and an OD of 1.5" (20x increase in radius).

Where R is the spreading resistance; Rs is the sheet resistance of the copper plane; R_1 is the inner radius of the circle and R_2 is the outer radius of the circle.

$$R = \frac{Rs}{2\pi} * \ln\frac{R_2}{R_1} \qquad R = \frac{1.26m\Omega}{2\pi} * \ln\frac{1500}{75} \qquad \text{R = 0.60 m}\Omega$$

Equation 33.3. Spreading Resistance from the Power Pin

With a 20A current, this will produce a voltage drop of 12 mV and a power dissipation of 0.24 Watts. It is not desirable to waste a quarter of a Watt if it can be easily avoided.

The power pins of the DC/DC converter are also used to conduct heat from the converter to the PCB power planes. The small hole from the power pin also creates a high thermal resistance to the power plane.

A solution to both problems is to make a large surface pad on both sides of the PCB surrounding the power pin via hole. Multiple small via holes should be drilled through these pads to make vias connecting to the power plane. The surface pads on a PCB are thicker than the internal planes because they are plated up when the vias are plated. A typical surface pad thickness is 2.0 mils. With pads on both sides of the PCB this gives almost an 8x lower spreading resistance than the internal voltage plane.

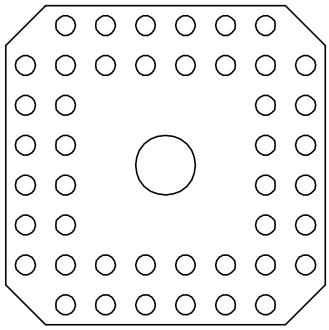

Figure 33.6. Power Supply Pin Surface Pad Layout

The pad in Figure 33.6 is 400 mils square and will make a connection to the internal voltage plane similar to a 400 mil diameter via. This will reduce the resistance to a 1.5" circle from 0.60 mΩ to 0.27 mΩ. The power dissipation is reduced to 106 mW but, more importantly, the maximum current density in the internal ½ oz copper plane is reduced by the ratio of the effective via size which is 400/75 or 5.3x.

Note: The via contacts to the power plane should be solid without any thermal ties. To facilitate soldering and unsoldering the power supply module, the 75 mil holes for the power pins should have two thermal ties no wider than 10 mils.

CHAPTER 34: DECOUPLING CAPACITORS

By John Zasio

Decoupling capacitors are used to provide a low impedance source of energy to high frequency changes in load current. There are many types and sizes of capacitors. In general, units with large capacitance are also large in physical size. The physical size, package type, material and mounting method are, in most cases, more important than the capacitance value. All capacitors have parasitic elements that determine their effectiveness.

Capacitor Types

Figure 34.1 depicts the many types of capacitors that are suitable for use as decoupling capacitors. Those capacitors, and the properties of each, are as follows:

- Aluminum Electrolytic
  - Large Capacitance, Large Physical Size, High Voltage, Low ESR
  - Capacitance degrades with life due to evaporation of electrolyte
- Low ESR Tantalum
  - 1 uF to 1000 uF, Medium Physical Size, Low Voltage, Low ESR
  - Failure mode is a short and a fire.
- Tantalum Organic
  - 1 uF to 1000 uF, Medium Physical Size, Low Voltage, Low ESR
  - Self-healing failure mode.
- Ceramic
  - Small Capacitance, Small Physical Size, High and Low Voltage, Very Low ESR
  - Lowest cost, Highest reliability
- Arrays
  - Ceramic capacitors with an array of contacts for very low ESL
  - High cost

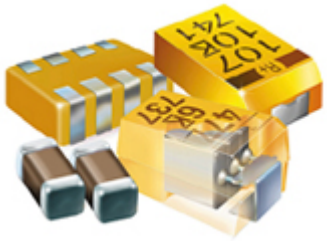

Figure 34.1. Ceramic, Tantalum, and Array Capacitors

Aluminum electrolytic capacitors are of use in high voltage applications such as on the 48V inputs to the DC/DC converters. Physically, they are too large for low voltage applications.

Tantalum oxide capacitors have a much higher capacitance per unit volume for low voltage applications than aluminum electrolytic capacitors. They must be operated at a 50% derating of the rated voltage to obtain a low failure rate. The failure mechanism is a catastrophic short that results in the capacitor causing it to catch fire and burn itself off the card.

A few years ago, Kemet introduced Tantalum Organic capacitors. These have the same physical size as tantalum oxide capacitors but they have very low ESR and a self-healing mechanism for failures.

Ceramic capacitors are the smallest and the lowest cost. They are available in many physical sizes, voltage ratings and capacitance values. There are also several types of dielectric material that are used to make these units. A high quality ceramic capacitor can be made from a dielectric, such as X7R, that is very stable with changes in voltage and has a very low ESR. As will be shown later, very low ESR presents resonance problems for decoupling applications. The Y5V dielectric has a higher dielectric constant that results in higher capacitance per unit volume and a higher ESR that is more suitable for decoupling applications. However, this dielectric is unstable with both temperature and voltage and is a poor choice for decoupling.

Ceramic capacitor arrays are chip capacitors with multiple contacts to lower the ESL. The cost for these units is at least 10 times that of an 0603 standard ceramic chip capacitor. It is usually a better choice to use two 0603 capacitors instead of one array to get the same effective ESL at a much lower cost. This is due to the fact that the mounting inductance will dominate over the low inductance of the array package (see Figure 34.10).

Capacitor Equivalent Circuit

Figure 34.2 depicts a decoupling capacitor equivalent circuit.

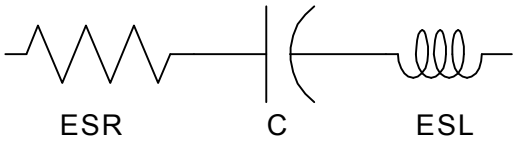

Figure 34.2. Decoupling Capacitor Equivalent Circuit

In Figure 34.2, the equivalent circuit elements are as follows:

Element	Impedance
ESR (Equivalent Series Resistance)	**Rs**
Capacitance	**XC = 1/2πfC**
ESL (Equivalent Series Inductance)	**XL = 2πfL**

The effective impedance of the decoupling capacitor equivalent circuit can be determined using Equation 34.1

$$Z = \sqrt{R_S{}^2 + (X_L - X_C)^2}$$

Equation 34.1. Effective Impedance Equation for Decoupling Capacitor Equivalent Circuit

Series Resonance

Series resonance occurs at the frequency where X_C (capacitive reactance) is equal to X_L (inductive reactance). At this frequency, the impedance of the capacitor becomes equal to the ESR. At all frequencies above series resonance, the capacitor becomes an inductor.

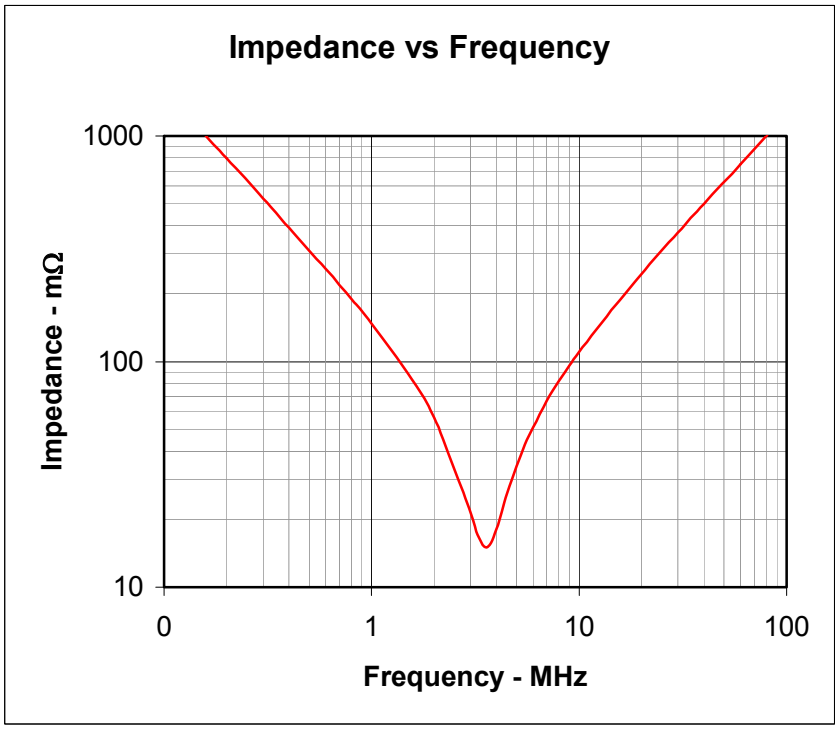

Figure 34.3. 0603 Ceramic Capacitor Impedance vs. Frequency (C=1uF, ESR=15mΩ, ESL=2nH)

Figure 34.3 is a plot of the impedance vs frequency for an 0603 ceramic 1uF capacitor with a typical PCB footprint. Series resonance occurs at 3.5 MHz. Above 3.5 MHz, the impedance of the capacitor is determined by the ESL.

Parallel Resonance

All capacitors mounted on a PCB will interact with the capacitance of the PCB power planes to form a high impedance spike (impedance hole) at a frequency somewhat higher than the series resonance point. Figure 34.4 is the equivalent circuit of a surface mount ceramic capacitor mounted on a small PC board.

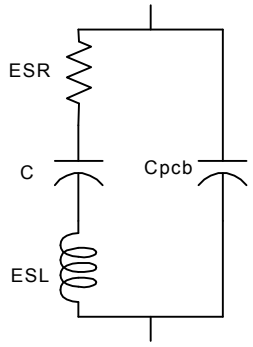

Figure 34.4. Equivalent Circuit of Ceramic Chip Capacitor on a PC Board

At high frequencies, the capacitance "C" is effectively a short circuit. Parallel resonance occurs at the frequency where X_L (for the ESL of the chip capacitor) is equal to X_C of the PCB plane capacitance. The PCB capacitor has an extremely low ESR. At the parallel resonance frequency, the magnitude of the impedance of the circuit is indirectly proportional to the ESR of the ceramic capacitor and can be calculated using Equation 34.2.

$$Z = \frac{ESL}{ESR \times Cpcb}$$

Equation 34.2. Impedance at Parallel Resonance Frequency

If ESR is extremely low, the impedance becomes extremely high. This typically occurs at a frequency of several hundred MHz and is a severe problem if this also happens to be the clock frequency or some harmonics of the data pattern.

Therefore, it is desirable to have ESR in a range sufficiently low to be an effective capacitor at low frequencies and sufficiently high to control the impedance spike at the parallel resonance frequency.

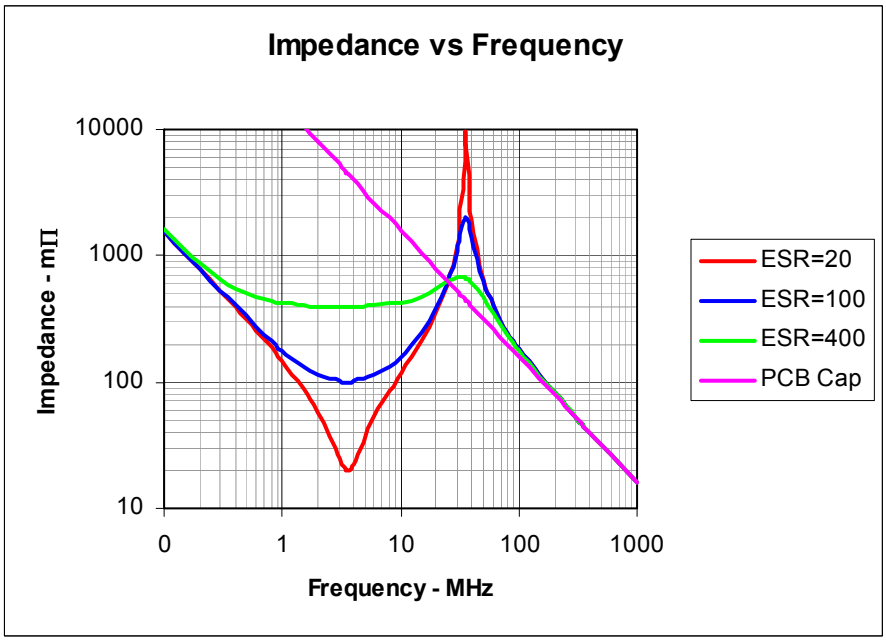

Figure 34.5. Impedance at Parallel Resonance vs. ESR (1uF 0603 Ceramic, 10nF PCB)

Figure 34.5 shows the series and parallel resonance impedance as a function of the ESR for the ceramic capacitor. The fourth plot in the graph is the impedance of the PCB capacitor. It should be noted that the parallel resonance impedance is always higher than the PCB capacitor impedance by itself. With an ESR of 20mΩ, the impedance spike is 9.5Ω and occurs at 35.6 MHz.

Data Sheets vs. Parameters of Interest

Capacitor manufacturers typically specify the capacitance and ESR value measured at 1 KHz. Typically, the ESL is not specified and is significantly affected by the footprint used to mount the capacitor on the PCB.

The capacitance, ESR and ESL are not constant at all frequencies. A capacitor is made up of many small elements. At low frequencies, all of the elements contribute to the measured value of the capacitor. At higher frequencies, the capacitive elements nearest the terminals can have an impedance that is low compared to the resistive connections to elements farther from the terminals. This causes non-linear values for C, ESR and ESL vs. frequency.

The best method for determining the characteristics of the individual capacitors, as well as the total decoupling system, is to measure them.

Capacitor Characterization Method

Capacitance meters are designed to operate at a low frequency where X_C is large compared to ESR and ESL. The characteristics of a decoupling capacitor are most interesting near the resonance frequency where the impedance may be only a few milliohms. Therefore, it is necessary to use a low impedance, high frequency measurement technique.

When measuring a resistor of a very low value, a four-point probe method is used. This is done because the resistance of the wires from the meter to the resistor is generally large compared to the resistor value. One pair of wires is used to drive a known current through the resistor and the second pair of wires is used to measure the voltage drop across the resistor. A similar approach is used to measure a low impedance value at high frequencies. A sweep generator can be used to drive a known amount of current into the load through one coax cable. A second coax cable connected to a spectrum analyzer can be used to measure the resulting voltage. Figure 34.6 depicts the capacitance measurement diagrams.

Circuit Diagram

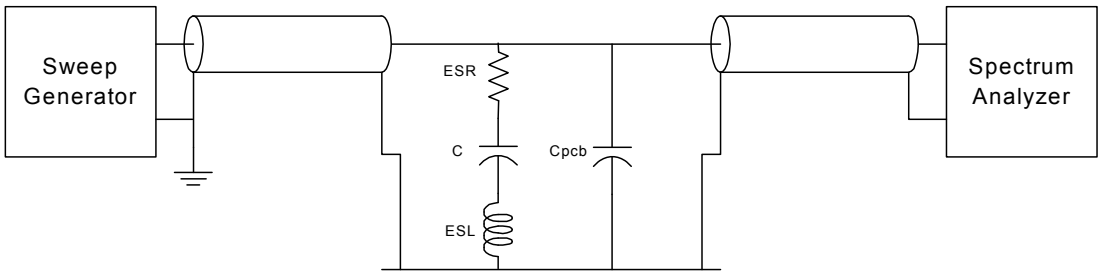

Physical Diagram

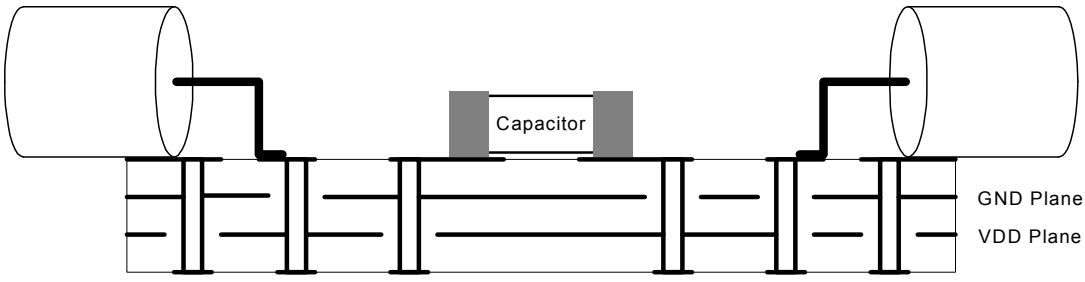

Figure 34.6. Capacitance Measurement Diagrams

A Hewlett Packard model 8594EM Spectrum Analyzer was used for all the measurements in this chapter. This instrument has a frequency range from 9 KHz to 2.9 GHz and it has a built in sweep generator that covers the same frequency range. The instrument has a built in 50Ω source terminator for the sweep generator. The receiver is also terminated in 50Ω.

Figure 34.7 is an equivalent circuit of the measurement setup. The sweep generator is set to 0 dBm output power setting (zero dB above one mW). All Spectrum Analyzer readings are in dB below the input power setting.

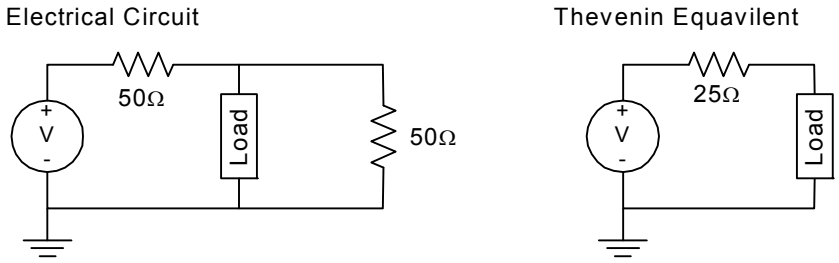

Figure 34.7. Capacitor Measurement Equivalent Circuit

Equation 34.3 provides the method for calculating the measured impedance vs. the measured output voltage. For a load with an impedance (Z_L) very small compared to 25Ω:

$$I = V_S / 25$$

$$V_L = Z_L \times I$$

$$Z_L = 25 \times (V_L / V_S)$$

Equation 34.3. Calculation of measured impedance vs. measured output voltage

The Spectrum Analyzer measurement in Equation 34.3 has the ratio of (V_L / V_S) expressed in dB (decibels).

$$dB = 20 \times \text{Log} (V_L / V_S)$$

Equation 34.4. Definition of dB for Voltage Ratios

Therefore, the impedance (Z_L) can be calculated using Equation 34.5.

$$Z_L = 25 \times 10^{(dB / 20)}$$

Equation 34.5. Impedance vs. Spectrum Analyzer Measurements in dB

For example, a reading of -40 dB is equal to a voltage ratio of 0.01, which gives a load impedance of 0.25Ω. Figure 34.8 depicts the spectrum analyzer screen capture of a PCB impedance vs frequency. The upper curve is the PCB plane capacitor by itself and the lower curve is the combination of the PCB plane capacitor and one 0.1uF capacitor.

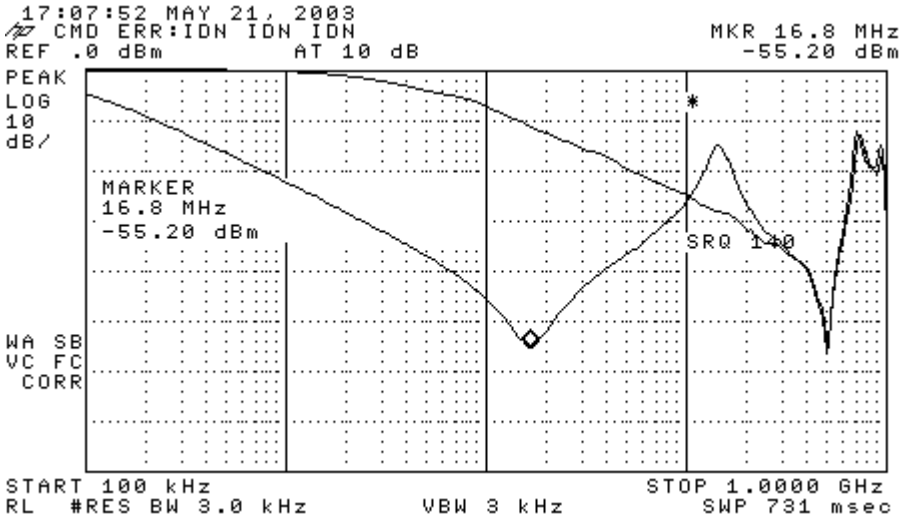

Figure 34.8. Spectrum Analyzer Screen Capture for AVX 0.1uF 0603 Y5V Cap on Small Test Card

Figure 34.9 depicts the same information expressed in ohms. This type of plot is obtained by capturing the Spectrum Analyzer data as a "Frequency-Amplitude" table, loading it into an Excel Spreadsheet, converting it to impedance and then generating the plot from Excel.

The capacitance for this unit measured on an LCR meter at 1 KHz is 90 nF (0.09 uF). The capacitance calculated from the impedance measurement at 20 MHz is 81.8 nF, a drop of 9.1%.

The series resonance occurs at 15.8 MHz and the ESR for this capacitor is 45mΩ. Parallel resonance occurs at 145 MHz and at this frequency the PCB capacitance is 1.11nF.

The ESL at the parallel resonance frequency is 1.09 nH (1.09 x 10^{-9} Henries). This includes the inductance of the capacitor as well as the inductance of the vias to the power planes.

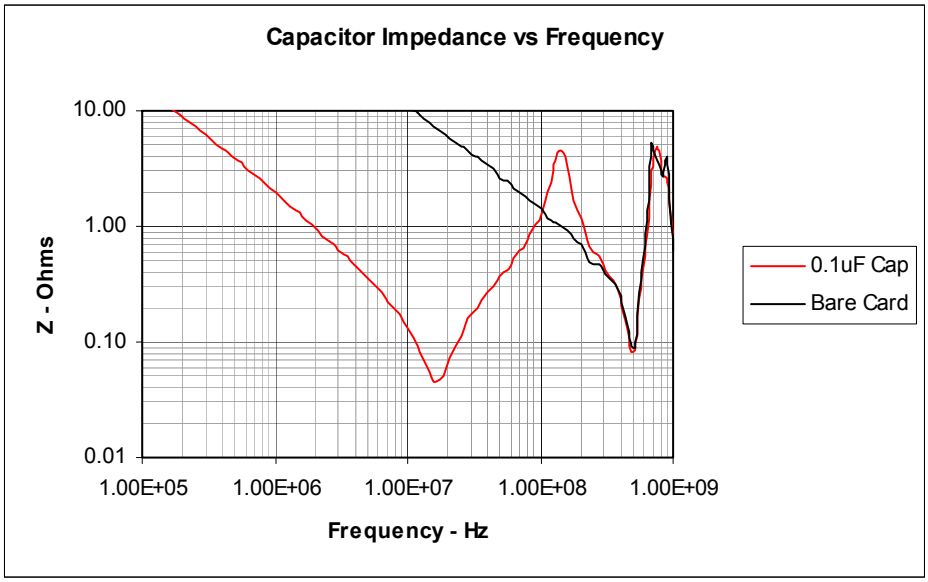

Figure 34.9. Impedance vs. Frequency for AVX 0603 0.1uF Capacitor

The black curve is the impedance of the plane capacitor by itself. The red curve is the combined impedance of the plane capacitor and the 0.1 uF capacitor.

Capacitor Measured Characteristics

Table 34.1 presents the measured characteristics of frequently used decoupling capacitors.

Type	Vendor	Vendor Part Number	Rated		Measured				
			Cap	Volts	Cap @ 1KHz	Cap @ 1MHz	ESR	ESL	F_{RES}
Tantalum Org.	Kemet	T520D337M006AS	330 uF	6.3	343uF	na	23mΩ	3.1nH	250 KHz
0603 Ceramic	AVX	0603ZG105ZAT2A	1.0 uF	10	1.13uF	650nF	20mΩ	0.7nH	6 MHz
0603 Ceramic	AVX	0603ZG105ZAT2A	0.1 uF	10	94.5nF	81nF	40mΩ	0.6nH	18 MHz
0603 Ceramic	AVX	0603ZG105ZAT2A	0.01 uF	10	9.5nF	9nF	60mΩ	0.5nH	50 MHz
0612 Ceramic	AVX	0612YC104MAT	0.1 uF	16	95.0 nF	92 nF	14mΩ	0.38nH	23 MHz
Array Ceramic	AVX	W3L1YC104MAT	0.1 uF	16	97.2nF	82nF	27mΩ	0.25nH	35 MHz

Table 34.1. Measured Characteristics for Decoupling Capacitors

The last column (F_{RES}) in the above table is the series resonance frequency of the capacitor. This is the frequency above which the capacitor becomes inductive and starts to be ineffective at providing low impedance. The above data was measured on a test board with the power planes very close to the side of the board where the capacitors are mounted. The Via-Length from the surface capacitor pad to the voltage plane is 11 mils. A discussion of the via length effect is contained in a following chapter.

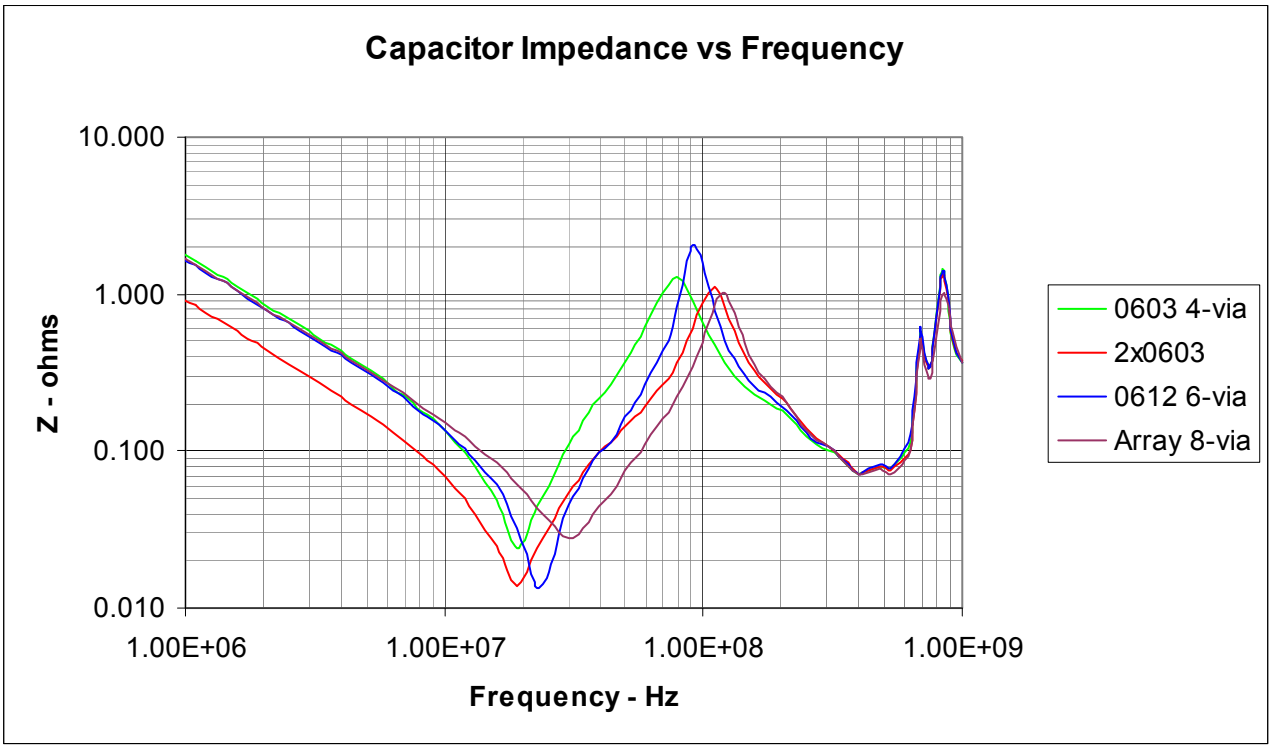

Figure 34.10. Capacitor Impedance Comparison

Figure 34.10 shows the measured frequency response of three different capacitor types. All capacitors are nominally 100nF. The test board has the power plane near the center of the board so the vias are a little longer (30 mils) and the effective series inductance (ESL) is larger than in the previous table.

Type	Footprint	Cap	ESR	ESL	Cost
		C - nF	ESR-mΩ	ESL - nH	Cents
0603	4-via	84	25	0.840	1
2x0603	4-via	168	14	0.410	2
0612	6-via	92	14	0.520	15
Array	8-via	90	27	0.315	40

Table 34.2. Capacitor Characteristics from Above Plot

The 0603 ceramic capacitor is a very common device made by many manufacturers. It's cost is approximately one cent when purchased in volume.

The 0612 ceramic capacitor has the metal contacts on the wide side rather than the narrow ends. This creates a lower ESL than a 1206 capacitor of the same physical size. This is a much lower volume part and is available only from a few manufacturers. This type of capacitor is used on IBM and Motorola PowerPC ceramic packages to provide high frequency decoupling on the package. Because it is a specialty item the cost is much higher.

The AVX IDC Array capacitor is the same physical size as the 0612 capacitor but instead on one contact on each side it has four and they alternate between VDD and GND connections. This produces a much lower inductance but at a much higher cost. Typical cost for this part is 40 cents each.

The Red curve in the above plot shows two 0603 capacitors in parallel on the PC board. This produces an impedance with a high frequency parallel resonance very close to that of the Array capacitor and at a higher frequency than the 0612 capacitor at a significantly lower cost.

Power Plane Capacitance

Even the 0603 10nF capacitors have series resonance at 50 MHz. At a few hundred MHz they are no longer of value in providing a low impedance path to the load. Above this frequency, it is necessary to use the capacitance of the power planes to provide decoupling. Equation 34.6 can be used to determine the capacitance of a parallel plate capacitor.

Where: E = 0.2249×10^{-12} F/in

Er = 3.8 to 4.2 for FR406

L = Length (in)

W = Width (in)

T = Dielectric Thickness (in)

Example: T = 3 mil, Er = 4, C= 300pF/in^2

Parallel Plate Capacitors: $$C = E \times Er \times \frac{L \times W}{T}$$

Equation 34.6. Capacitance of a Parallel Plate Capacitor

Using the stackup shown in Figure 32.4, there is a dielectric thickness from VDD to GND of 3 mils on one side and 9 mils through the signal planes to the second GND plane. This second GND plane contributes an additional 100 pF/in^2 for a total of 400 pF/in^2.

There are many vias in a PCB, especially under a BGA package. Each via requires an antipad (clearance area) in the power plane. This reduces the effective capacitance to approximately 85% of the above values.

A 12" x 18" board will have a total of 86 nF of high frequency decoupling capacitance. At 100 MHz, this provides an impedance of 4.6 mΩ.

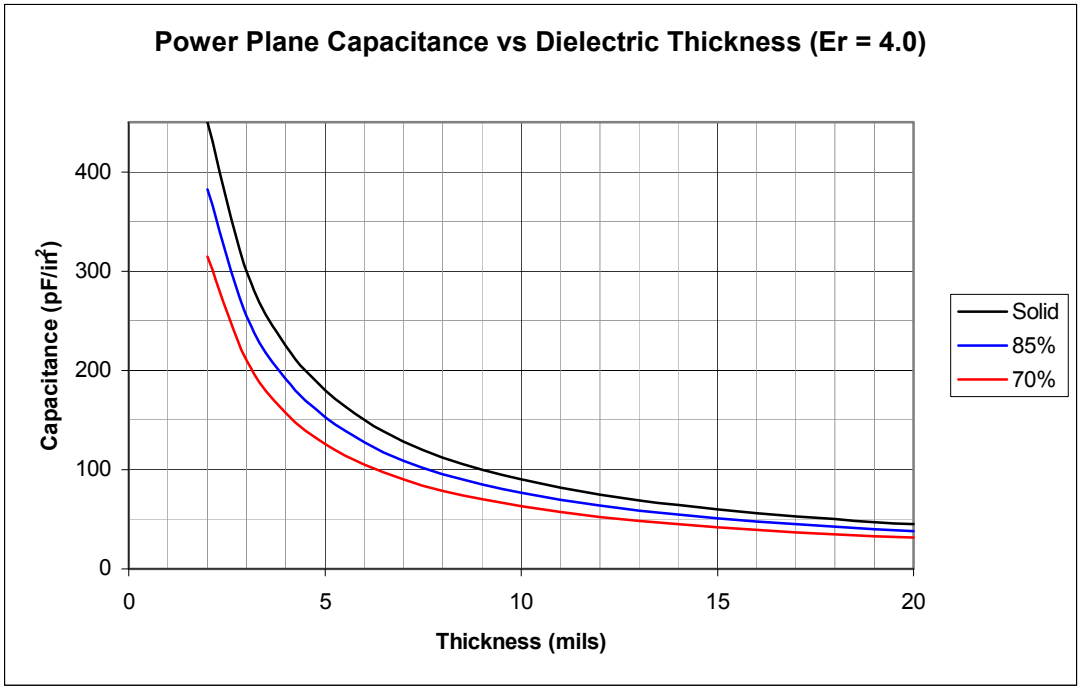

Figure 34.11. Power Plane Capacitance vs. Dielectric Thickness

Power planes have antipad openings or plane clearances for vias. Under a BGA this may reduce the effective area to 70%. Averaged over the whole board the effective area is typically 85%.

CHAPTER 35: POWER SUBSYSTEM INDUCTANCE

By John Zasio

Inductance is the most significant problem in designing a high frequency power distribution system. Inductance impedes the flow of current at high frequencies and, as a result, makes decoupling capacitors of little use at frequencies substantially above their series resonance. There are four major contributors to power distribution inductance.

1. Self inductance of the decoupling capacitor
2. Inductance of the capacitor footprint
3. Inductance of the vias from the capacitor footprint to the PCB power planes
4. Inductance of the power planes

To address these issue, components can be chosen and a design style used that can lower the inductances.

Physical Design Contributors to Inductance

The entire path including the ground return path determines the inductance of an electrical circuit. Smaller length conductors and close spacing between the conductors decreases the inductance. The inductance of a power via is a function of the distance to the return ground via. Inductance increases with effective via length. Figure 35.1 depicts those contributors that can add to inductance. In addition, Figure 35.1 shows the inductive path, highlighted in red, from the IC to the decoupling capacitor.

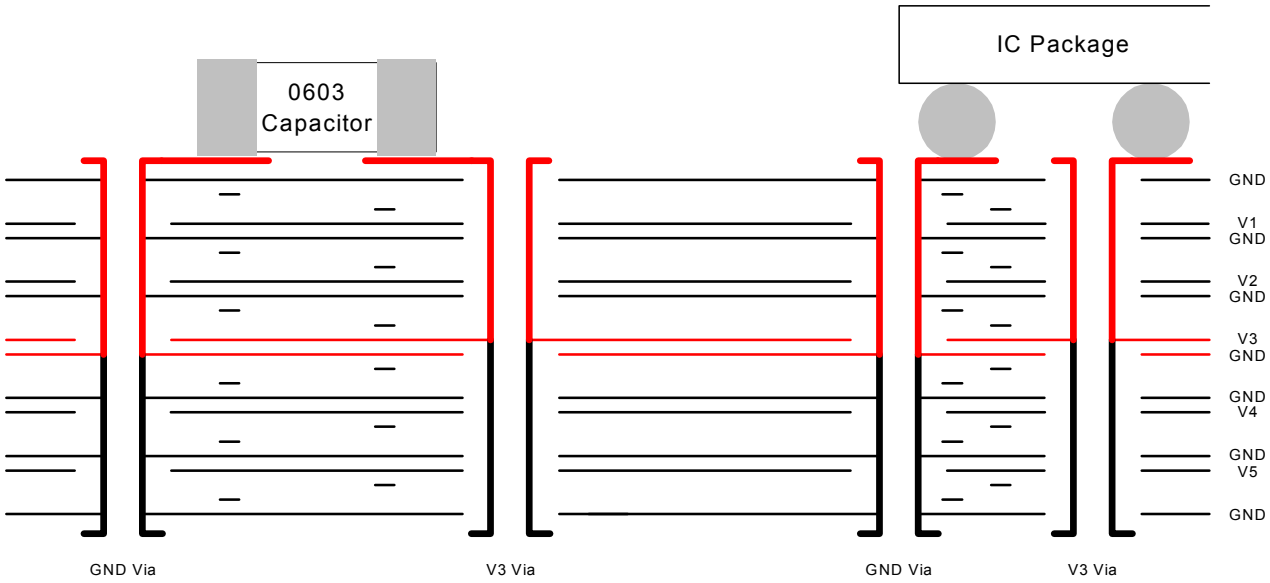

Figure 35.1. PCB Cross Section Showing Contributors to Inductance

The effective capacitor inductance is a function of the cross sectional area of the loop formed by the capacitor, the capacitor footprint and the via length to the power planes. The power plane pair between the capacitor and the IC adds to the effective inductance.

The IC power via length, spacing between VDD and GND vias and the package pin length also contribute to this inductive path. High performance IC packages have a large number of VDD and GND vias to form a parallel path to the power planes in order to decrease the effective inductance.

Capacitor Footprints

The ESL (Effective Series Inductance) of a capacitor is proportional to the physical size of the capacitor. Large electrolytic or tantalum capacitors have much higher inductance than small ceramic capacitors. A ceramic 0603 size capacitor has a lower ESL than a 1206 size capacitor.

The size and via placement of the capacitor footprint on the PCB has a significant effect on the ESL of a given capacitor. The examples of capacitor footprints presented in Figure 35.2 shows how the shape of the capacitor and via placement can change the ESL of a given capacitor by more than a factor of two.

All of the ESL values shown in Figure 35.2 are measured values with the capacitors connected to the V1/GND plane pair. Longer vias will increase the ESL.

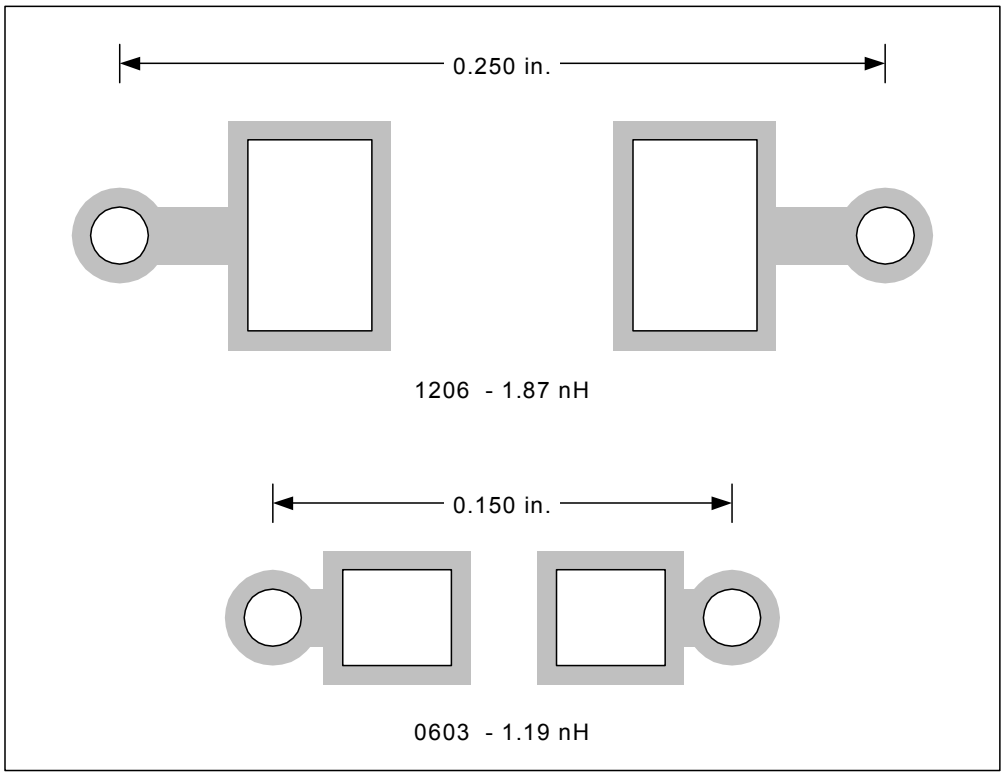

Figure 35.2. Typical 1206 and 0603 Ceramic Capacitor Footprints

As can be seen in Figure 35.2, the 0603 capacitor has 40% lower ESL than the 1206 capacitor when mounted with a typical footprint.

Figure 35.3 depicts how changing the position of the vias as well as adding vias can decrease the effective inductor area under the capacitor and substantially reduce the ESL.

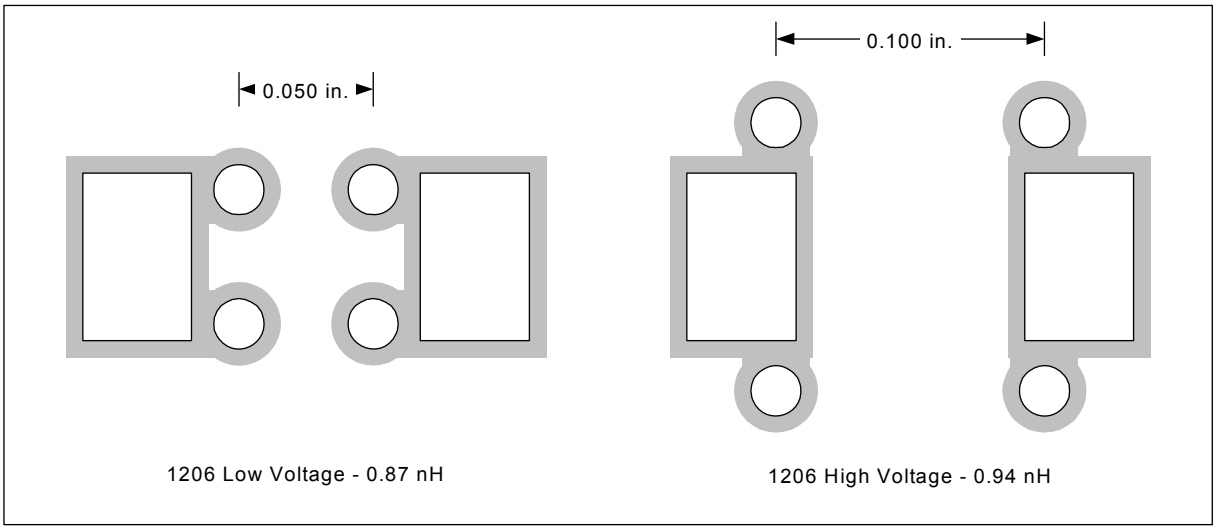

Figure 35.3. Four Via Footprints for 1206 Size Capacitors

As shown in Figure 35.3, there is sufficient space between the pads of the 1206 footprint to place the vias between the pads and to use four vias. With four vias placed between the pads, the ESL is reduced by 53% from the two via design. This results in half the required capacitors to obtain the same effective high frequency impedance or move the series resonant frequency higher.

High voltage applications such as the 48V inputs to DC/DC converters, such as those depicted in Figure 35.4, require a larger space between surface patterns and prohibit the vias from being placed between the pads. With the vias placed adjacent to the edge of the pads, the ESL is still significantly reduced from the two via design.

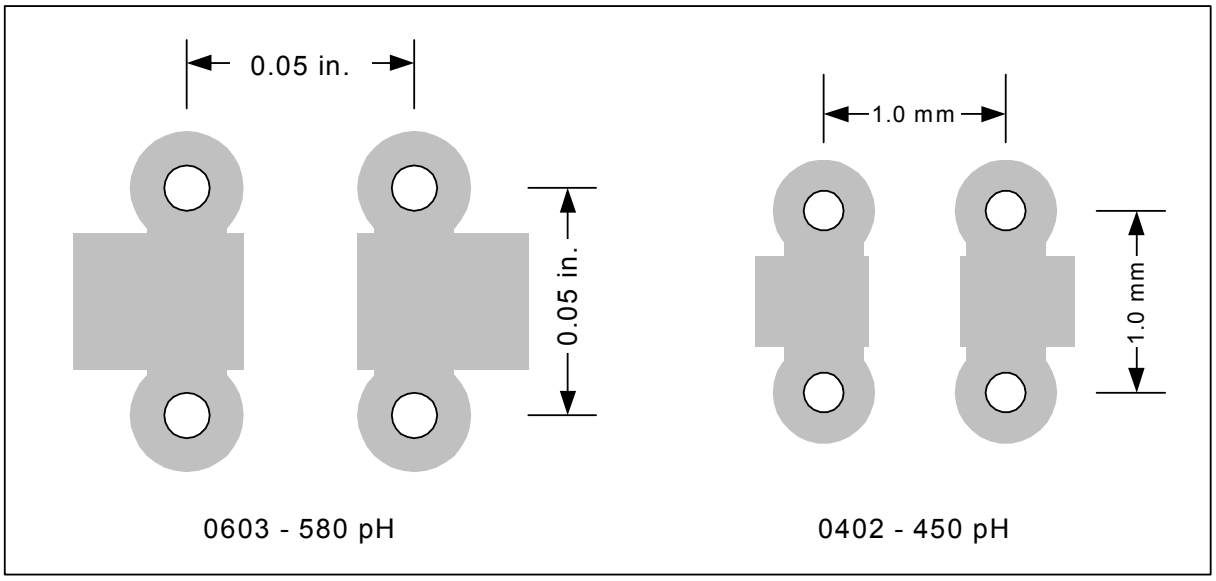

Figure 35.4. Four Via Footprints for 0603 and 0402 size capacitors

Capacitors with a case size of 0603 and smaller do not have sufficient space between the pads for the vias so they must be placed adjacent to the pad edges. This four via design has about 50% of the ESL of the two via design.

In Figure 35.4, the choice of the 0.05 inch spacing on the 0603 footprint is done so that it can be placed near the edge of a 50 mil pitch BGA and allow the same wire routing channels. Care must be taken with the design of the solder mask or the vias need to be plugged in order to guarantee that the solder will not drain off into the via holes during the reflow process.

The 0402 footprint via pattern shown in Figure 35.4 can fit the pitch of a 1.0 mm BGA**. However, this size capacitor is much more difficult to assemble on a large PCB and is generally not worth the trouble for the small decrease in ESL that is obtained.**

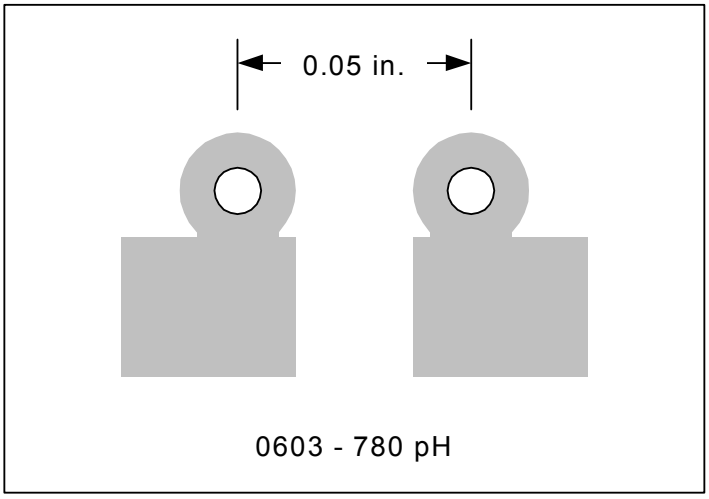

Figure 35.5. 2-Via Footprint for an 0603 capacitor with vias on the side

Some designers do not like the use of 4-via footprints because of the cost of drilling additional vias and the difficulty of controlling the tight geometries of the solder mask to prevent solder from running into the holes during the reflow process. If two vias are used they are much more effective if placed on the side of the pads rather than on the ends. As shown in Figure 35.5 two vias on the side produce a 200 pH larger ESL than the 4-Via example but substantially less ESL than two vias on the ends.

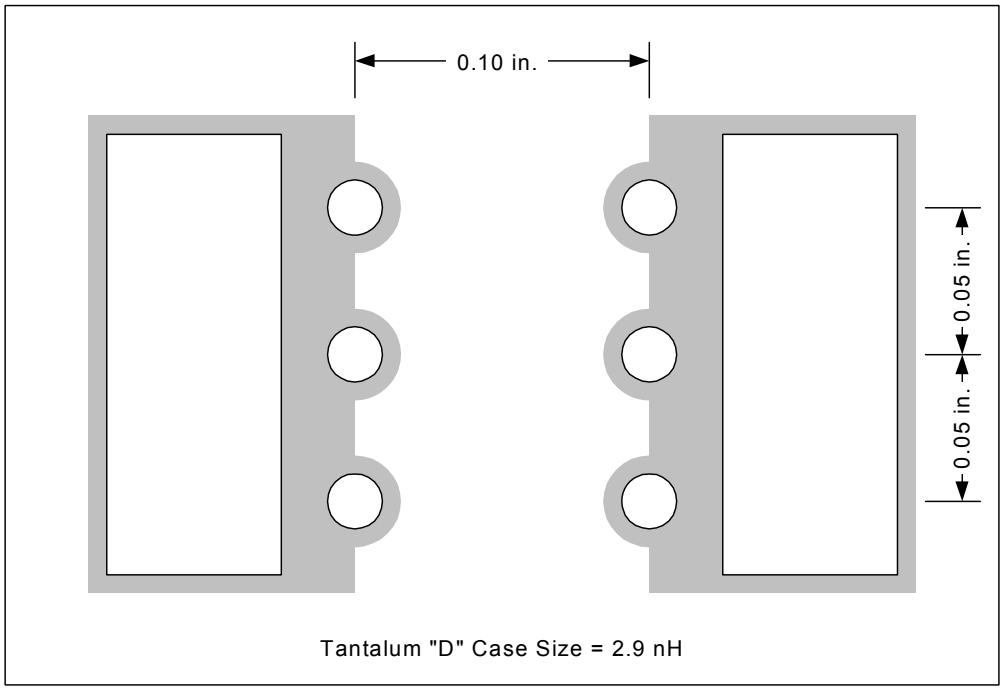

Figure 35.6. Six Via Footprint for a Tantalum Capacitor with a D-size Case

There are two reasons for multiple vias on large tantalum capacitor footprints--ESL and ESR. Tantalum capacitors are available with ESR as low as 15 mΩ. With these very low ESR capacitors, as shown in Figure 35.6, a six via pattern can have half the ESL and half the ESR of a two via footprint.

Capacitor Via Length to Power Planes

Referring to Table 35.1, the ESL of a capacitor is a function of via length to the power plane pair. The ESL is substantially less for a capacitor connected to V1/GND plane pair compared to the V5/GND plane pair.

Tests were performed with identical capacitors mounted side by side with identical footprints but with the capacitors connected to different PWR/GND plane pairs. The tests were done on multiple PCB stackups (all with thin dielectric between PWR/GND plane pairs) and with at least two different footprint designs. Equation 35.1 is a very close approximation to the measured ESL for an 0603 ceramic capacitor.

Where: C_{ESL} = ESL of Capacitor and footprint (pH)

Lv = Via Length (in mils) to PWR/GND

L_F = Length factor (pH/mil)

N_V = Number of Vias

T = Dielectric Thickness (cm)

Example: 0603 Ceramic Capacitor

C_{ESL} = 460pH, L_V = 13.5 mils, L_F = 35.5 pH/mil, N_V = 4

ESL= 460pH + (13.5 mils x 35.5 pH/mil/4 vias) = 580pH

$$ESL = C_{ESL} + (L_V \times L_F / N_V)$$

Equation 35.1. ESL vs Via Length and Via Count

The Length Factor (L_F) of 35.5 pH/mil works for all 0603 size capacitors with either two or four vias. The AVX IDC array capacitor has a slightly higher value of 44.4 pH/mil. Table 35.1 shows the ESL for several capacitors as a function of the via length. These values are calculated from the above equation and are within a few pH of the values measured on a 100 mil thick PC board with the stackup shown in Figure 32.5.

Vendor	AVX	AVX	Johanson	AVX
Case	0603	0603	0603	Array-IDC
Cap.	0.1uF	0.1uF	0.1uF	0.1uF
Vias	2	4	4	8
C_{ESL} - pH	950	460	580	151
L_F - pH/mil	35.5	35.5	35.5	44.4
V_L	ESL	ESL	ESL	ESL
mils	nH	nH	nH	nH
13.5	1.19	0.58	0.70	0.23
28.8	1.46	0.72	0.84	0.31
62.5	2.06	1.01	1.13	0.50
77.5	2.33	1.15	1.27	0.58

Table 35.1. Examples of Capacitor ESL vs Via Length

The data in the Table 35.1 shows that the ESL can double when the capacitor is connected to a PWR/GND plane pair on the opposite side of the PC board.

It should be noted that the measured ESL of a Johanson 0603 capacitor is 120 pH higher than the AVX capacitor. This is assumed to be due to a slightly thicker outer dielectric layer on the Johanson capacitor. This increased inductance was observed on several batches of both 0.1uF and 0.01uF 0603 capacitors.

Inductance of a Power Plane Pair

A pair of conductors that has a uniform cross section form a transmission line with a characteristic impedance (Z_0). If the wire is a signal trace with a width of about 3.75 mils and has a dielectric thickness of 3 mils, it will have a Z_0 of 50Ω. If the signal wire width is increased, the Z_0 of the line will decrease. A signal wire with a width of 1.0 cm (0.4 in) and the same dielectric thickness of 3 mils will have a Z_0 of 1.07Ω.

This value can be calculated from Equations 35.2 through 35.5.

Where: $E = 0.2249 \times 10^{-12}$ F/in

Er = 3.8 to 4.2 for FR406

A = Area (in^2)

T = Dielectric Thickness (in)

Example: Power plane between two GND planes as in Figure 35.1

Er = 4, T1 = 3 mils, T2 = 9 mils

C = 400 pF/in^2 = 62 pF/cm^2

$$C = E * Er * \frac{A}{T}$$

Equation 35.2. Parallel Plate Capacitance

Assume a transmission line with this 1.0 cm width as an example.

Example: $$Zo = \frac{\sqrt{4}}{3} * \frac{10^2}{62(pF)} = 1.07\Omega$$

$$Zo = \frac{\sqrt{Er}}{3} * \frac{10^2}{C(pF)}$$

Equation 35.3. Transmission Line Impedance vs Capacitance

A transmission line consists of uniform segments of capacitance and inductance. The impedance of this line can be calculated from the following equation.

$$Zo = \sqrt{\frac{L}{C}}$$

Equation 35.4. Transmission Line Impedance vs. L and C

Converting this equation to solve for inductance given we know the value of Z_0 and C gives the following.

Example: $L = 1.07^2$ x 62 pF/cm
$L = 71.6$ pH/cm
Or: Lsq = 71.6 pH/sq

$$L = Zo^2 * C$$

Equation 35.5. Transmission Line Inductance

This inductance for a one cm long and one cm wide transmission line can be thought of as inductance per square since the equations will produce the same value so long as the width and length are the same. This then follows the same form of calculations that are used for resistance based on the sheet resistance of ohms per square.

Let us consider an integrated circuit (IC) mounted on a PCB and a ring of 0603 decoupling capacitors that surrounds the package as shown in Figure 35.7. There are a total of 48 capacitors shown but these may be divided equally among the 4 power planes in the IC package. Therefore, there are only 12 capacitors for each plane.

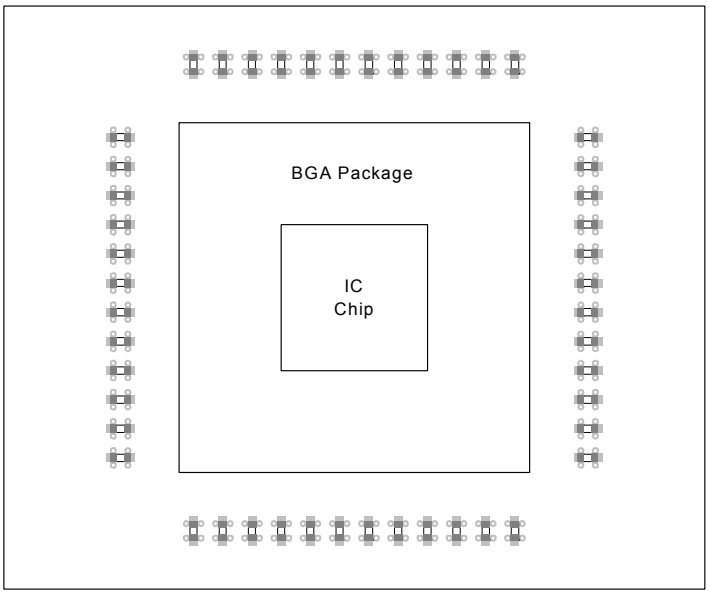

Figure 35.7. IC BGA Package and a Ring of 0603 Capacitors

In Figure 35.7, the BGA package is 1.5" square, the IC Chip is 0.8" square and the capacitors are places ¼ " from the edge of the package. The inductance of the power plane can be approximated by the following equation for the inductance of two parallel plates from an inner radius R1 to an outer radius R2.

Where: Lsq = Inductance per square
R1 = Inner Radius
R2 = Outer Radius

Example: R1 = 0.4" Radius from center to edge of IC chip
R2 = 1.0" Radius from center of IC to capacitor ring
Lsq = 71.6 pH/sq

$$L = \frac{Lsq}{2\pi} * \ln\frac{R2}{R1}$$

$$L = \frac{71.6\,pH}{2\pi} * \ln\frac{1.0}{0.4}$$

L = 10.4 pH

Equation 35.6. Circular Parallel Plate Inductance

At 580 pH per capacitor, the 12 capacitors in parallel would have an effective inductance of 48 pH if the capacitors are on the same side of the PC board as the BGA and are connected to the V1/GND plane pair. The addition of 10.4 pH for the inductance of the power planes is small compared to the inductance of the capacitors.

If the capacitors are on the backside of the PCB, the via length would be longer and the ESL of each capacitor would increase to 1.15 nH. The 12 capacitors in parallel would have an inductance of 95 pH that is almost double.

The message is that the horizontal inductance of the power plane pair is much lower than the ESL of the capacitors.

Capacitance Test Card

The test card used to evaluate capacitor placement and PCB resonance is shown in Figure 35.8. It is an epoxy glass (FR406) PCB that is 10.3" x 4.0" in size. The board has eight power planes, six signal layers and two outer pad layers.

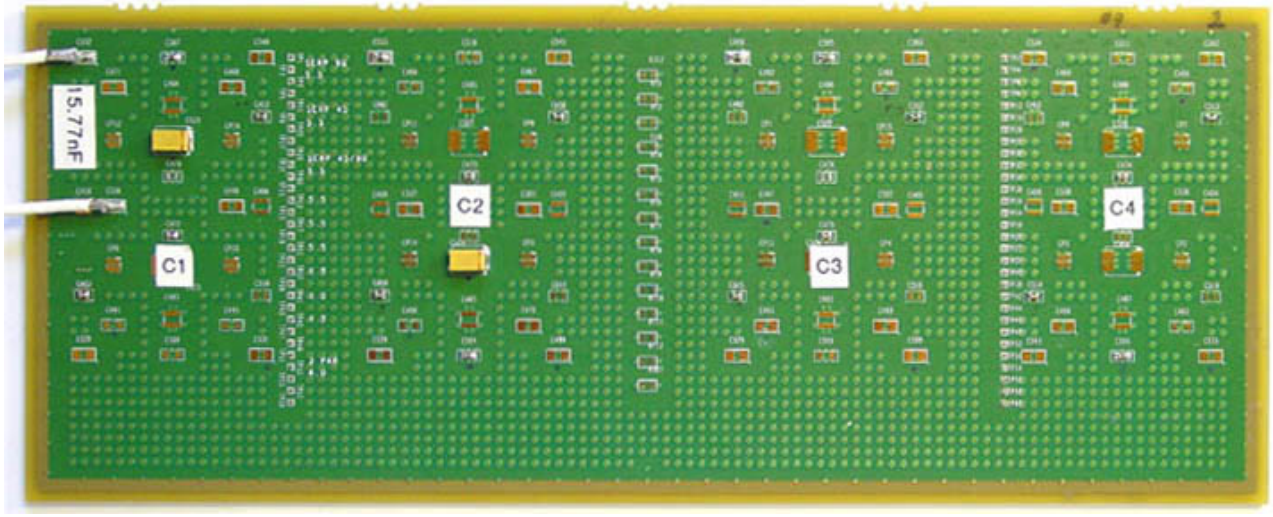

Figure 35.8. Photo of Capacitor Test Card

The Voltage/Ground plane pair used for these measurements has 15.77nF of capacitance. Two RG188 coax leads used to make the measurements are soldered to the left end of the board.

Capacitor Placement

Because of the low inductance of the power plane pair, the placement of the capacitors is not as critical as the consideration of via length and using the power plane close to the IC package side of the PC board.

To show the effect of capacitor placement, a single 0.1 uF 0603 capacitor was measured on four different locations on the PCB, as depicted in Figure 35.9. The two coax measurement leads were on one of the narrow ends of the board. The capacitor locations were at approximately 2.5" intervals down the long dimension of the board. Figure 35.10 shows a magnified view of the impedance vs. capacitor placement.

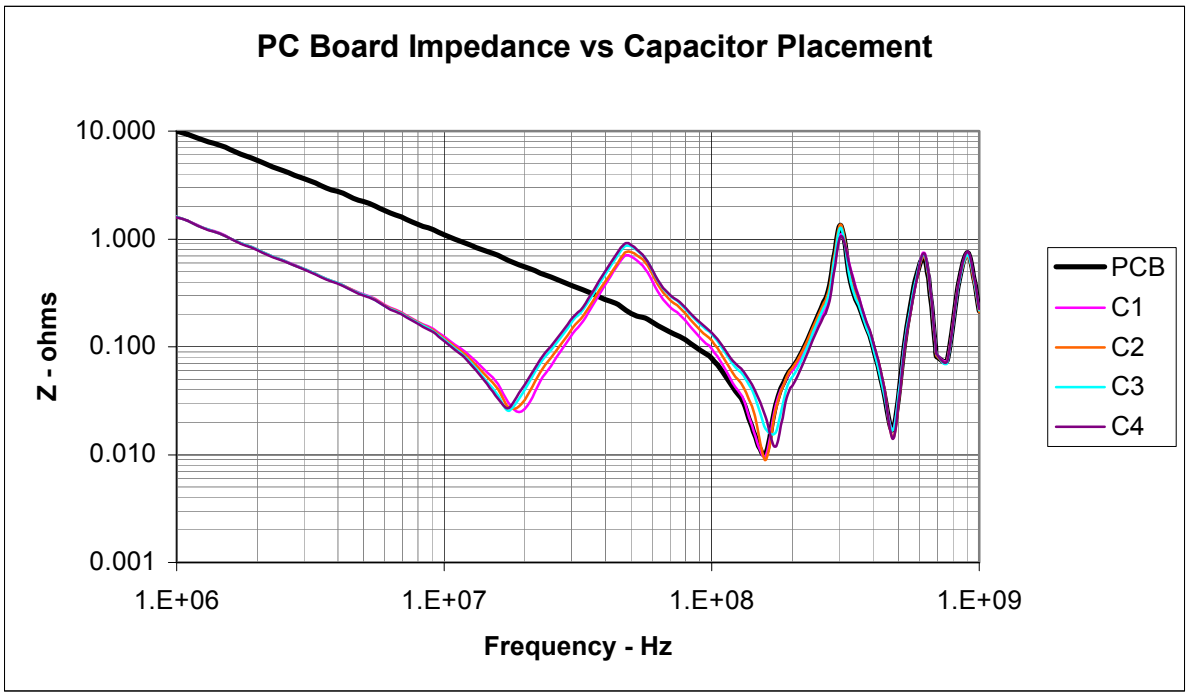

Figure 35.9. Impedance vs. Capacitor Placement

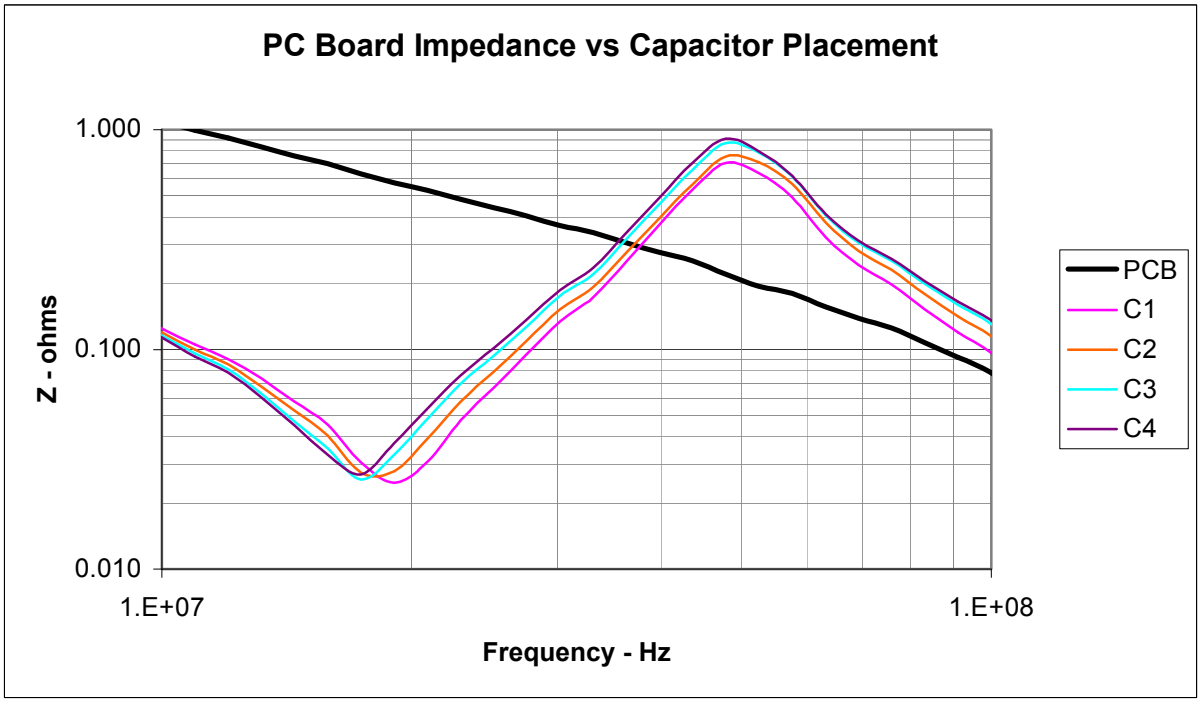

Figure 35.10. Magnified view of Impedance vs Capacitor Placement

Notice that the variation of the placement of the capacitor has only an 8% change in the series resonance frequency and almost no change in the parallel resonance frequency. The effects of the copper plane resistance seem to be larger than the inductive effects.

For PCBs with well-designed plane capacitors, the location of the discrete capacitors makes little, if any, difference. Placement of these capacitors should be done such that routing and assembly are made as easy as possible.

Power Plane Resonance

The power plane pair in a PCB has parallel plate capacitance and inductance. This uniform capacitance and inductance makes the planes look like a very low impedance transmission line. An open-ended transmission line will resonate at frequencies where its delay is evenly divided into the wavelength of that frequency. For a quarter and three quarters of the wavelength, the input to the resonant transmission looks like a short. For a half and multiples of a half wavelength, the input looks like an open. The power planes have sheet resistance and, at high frequencies, added resistance due to skin effect loss. So, in reality, it is not totally a short or open but it does have large impedance variations.

The Voltage/Ground plane pair in the 10.3" x 4.0" PCB used for these measurements forms a transmission line with an impedance of 0.100. The board is fabricated with FR406 that has a relative dielectric constant (Er) of 4.0. The propagation velocity of a signal in this low impedance transmission line can be calculated by using Equation 35.7.

Where: Vp = Velocity of Propagation in inches/ns

c = Speed of light in air = 11.8 inches/ns

Er = 3.8 to 4.2 for FR406

For Er = 4 Vp = 6 inches/ns

$$Vp = \frac{c}{\sqrt{Er}}$$

Equation 35.7. Signal Velocity of Propagation

The distance from the measurement coax attachment point to the far end of the PCB is 9.7 inches. At a velocity of 6 inches/ns, a signal will travel from the measurement point to the opposite end of the board in 1.62 ns.

The "wavelength" (λ) of a signal is the physical distance the signal must travel in a transmission line to complete one cycle. A 1.0 GHz signal completes a cycle (period) in 1.0 ns. Therefore, in this PCB, the wavelength is 6 inches.

At 1 GHz λ = 6 in/ns x 1.0ns = 6 inches

At 500 MHz λ = 6 in/ns x 2.0ns = 12 inches

Table 35.2 shows the calculated and measured resonance frequencies for this test PCB.

Length	Vp	Delay	
in.	in/ns	ns	
9.70	6	1.62	
Resonance	Calculated	Measured	
Point	MHz	MHz	Z (Ω)
1/4 λ	154.6	154	0.008
1/2 λ	309.3	303	1.290
3/4 λ	463.9	460	0.016
1 λ	618.6	606	0.840

Table 35.2. PCB Resonance Frequency

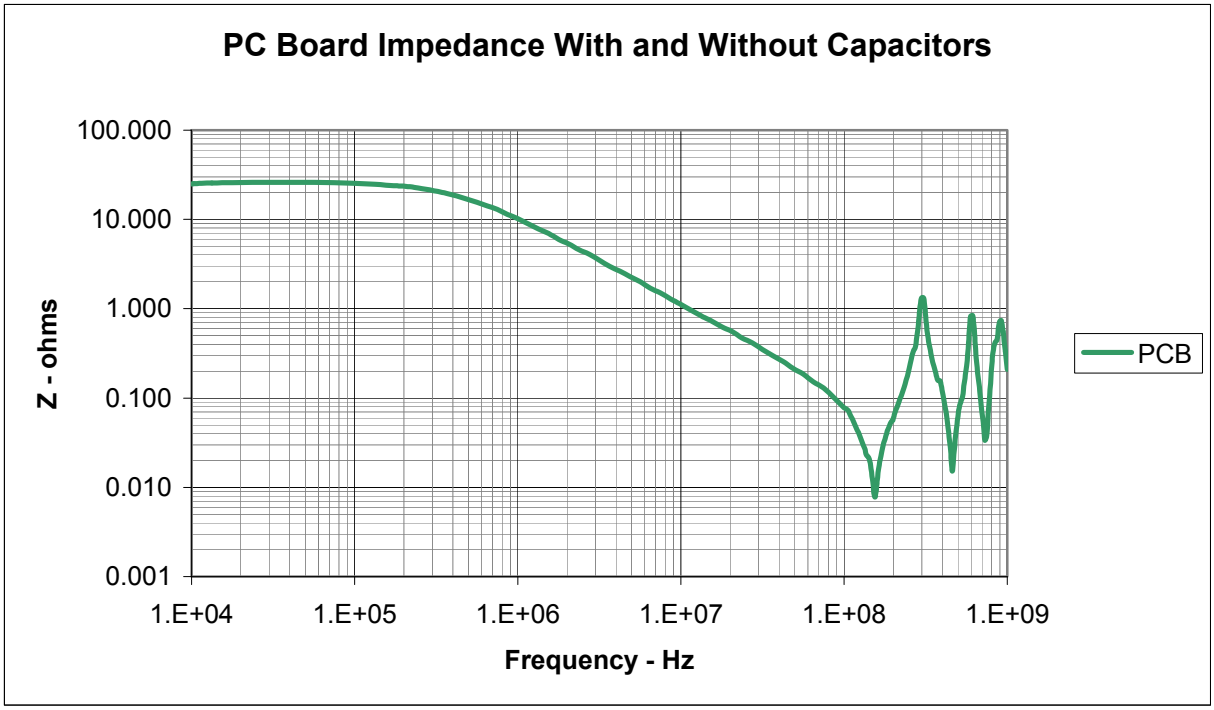

Figure 35.11. Test PCB Resonance Without Decoupling Capacitors

Figure 35.11 shows the ¼ λ low impedance resonance of 8 mΩ at a frequency of 154 MHz. The ½ λ high impedance (Impedance Hole) resonance of 1.29Ω occurs at 303 MHz. This impedance hole can be a significant problem if it occurs at the clock frequency or a harmonic of the clock frequency.

Adding decoupling capacitors to the board will reduce the impedance to acceptable levels and alter the high frequency impedance holes as shown in Figure 35.12.

In Figure 35.12, the red trace shows the impedance with a quantity of 1 x 330uF T520, 1 x 1.0uF 0603, 2 x 0.1uF 0603, and 6 x 0.01 uF 0603 capacitors. The blue trace is the same PCB with twice the number of capacitors. The low impedance from 10 KHz to 3 MHz is due to the 330uF tantalum capacitor. The 1.0uF ceramic capacitors take over from 3 MHz to 10 MHz. The 0.1uF ceramic capacitors are effective from 10 MHz to 30 MHz and the 0.01 uF capacitors from 30 MHz to 200 MHz.

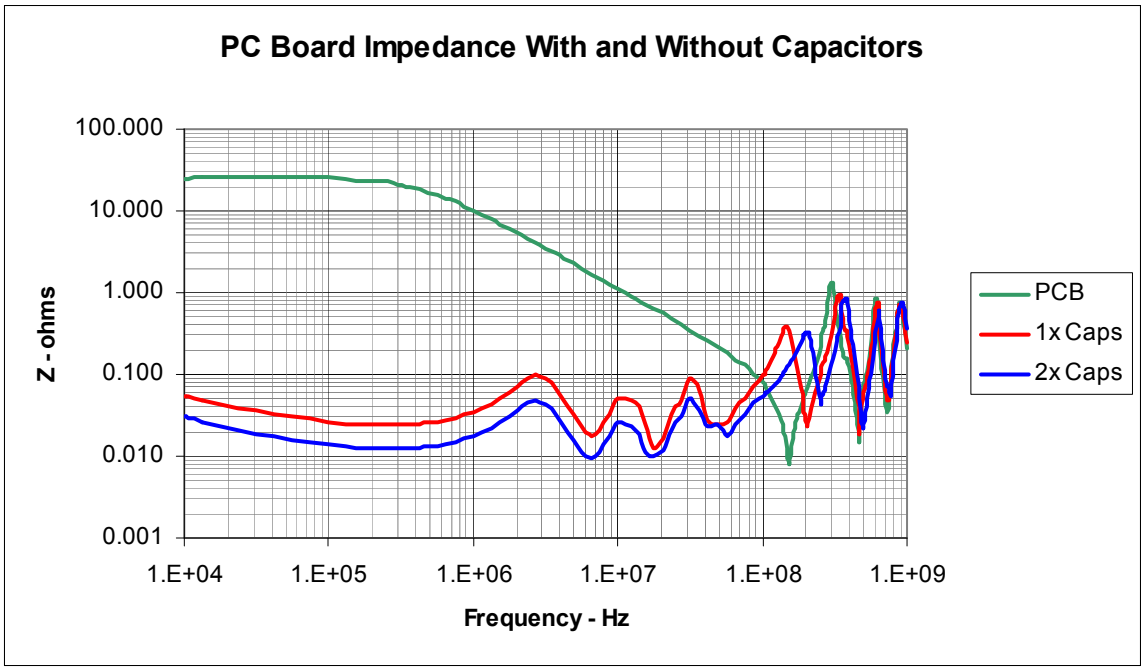

Figure 35.12. PCB Impedance With and Without Capacitors

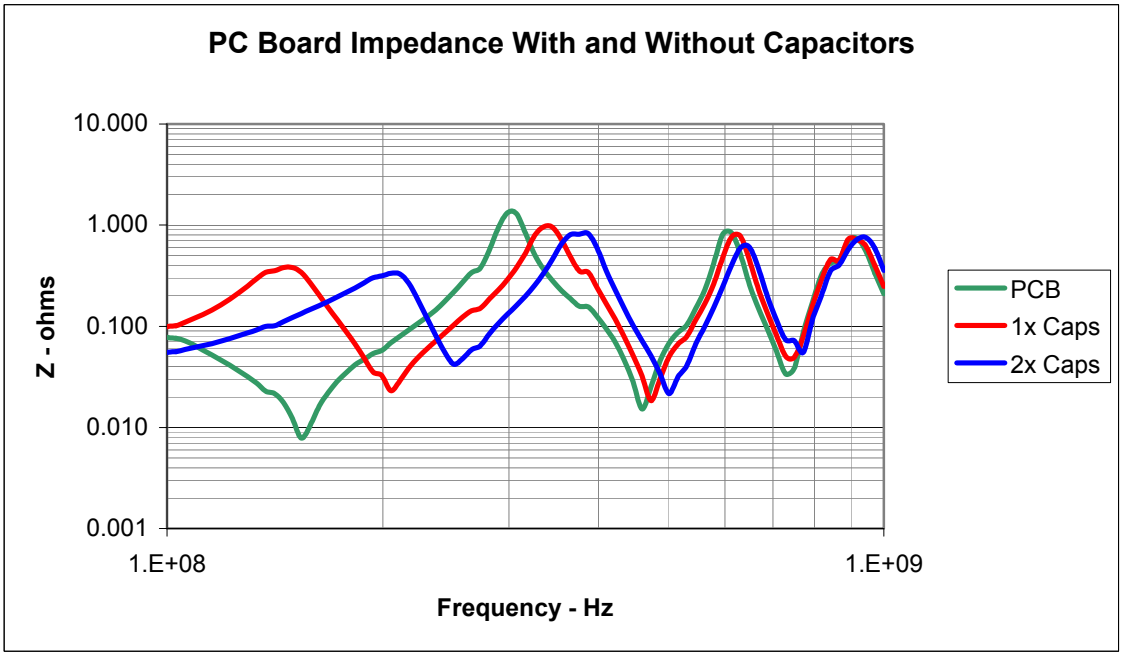

Figure 35.13. Magnified View of High Frequency Region

Figure 35.13 presents a magnified view of the PCB depicted in Figure 35.12. It should be noted that the with one set of capacitors (Red Trace), the first resonant point above 100 MHz is a positive spike (Impedance Hole) of 380 mΩ at 147 MHz due to the parallel resonance of the 0.01 uF capacitors and the PCB parallel plate capacitance. With 2x the number of capacitors (Blue Trace), this resonance point moves to 205 MHz and the impedance value drops to 330 mΩ. This is due to the fact that the total inductance of the 0.01 uF capacitors is halved

The first few PCB resonance points are moved to higher frequencies and lower impedance levels but they are not eliminated. PCB resonance points above 300 MHz are affected very little by the addition of capacitors.

CHAPTER 36: POWER DISSIPATION ESTIMATE

By John Zasio

In order to design the power distribution system, it is necessary to estimate the power consumption of all the loads on the PCB. The IC vendor provides some of the information but not in the required form. The data from the vendor normally is typical power dissipation and worst-case power dissipation for the entire chip. What is required is the load current on each of the power sources into the chip and the frequency at which this load changes.

Microprocessor Core

A typical microprocessor (uP) today is fabricated with a 130nm CMOS process. It may have a core clock frequency of 1GHz and a 72-bit I/O bus running at 133 MHz. The vendor's data sheet specifies a core voltage (Vdd) of 1.3 volts and the I/O voltage (Vddq) of 2.5V. Maximum power dissipation is specified to be 7 Watts. This includes the I/O and core.

If the total power for the microprocessor is specified as 7 Watts maximum and the I/O power is 1.0 Watts, then the core will dissipate about 6 Watts. Typically, a large IC running at a high clock frequency will dissipate1/4 to 1/3 of this power in the clock tree that is running all the time (except for JTAG test or Reset). The remainder of the core power is activity dependent and can change at any frequency.

The core of a large IC typically has 20 nF of internal capacitance due to idle circuits and memory cells. In addition, there may be a number of very high frequency decoupling capacitors mounted on the chip package. The clock tree on a large IC can have a current spike of several hundred amps with a width of a few hundred picoseconds. Most of this spike will be absorbed by the internal capacitors but there will be some significant ripple current at the clock frequency that must be decoupled on the PC board.

Core Maximum DC Current Idc-max = 6 Watts / 1.3 V = 4.62 Amps

Core Minimum DC Current Idc-min = 25% x Idc-max = 1.15 Amp

This change in load from 1.5 Amps to 4.62 Amps can occur at any frequency, depending on activity.

Microprocessor I/O Bus

A typical I/O bus will have a clock, 72 data pins and 40 address and control pins. All nets are capacitive loads consisting of the uP, two controller chips, the PCB vias and the PCB traces. Table 36.1 presents the load estimate for the microprocessor I/O.

Unit	Feature	Qty	Each	Total
			pF	pF
Processor	Pin	1	5.5	5.5
Load Chips	Pin	2	6.0	12
PCB	Via	6	1.3	7.8
PCB	Wire - in	7	3.4	23.8

			Total =	49.1

Table 36.1. Microprocessor I/O Load Estimate

The output circuit of the processor is a CMOS driver consisting of a P-channel MOSFET pull up to Vddq and an N-ch MOSFET pull down to GND as shown in Figure 36.1.

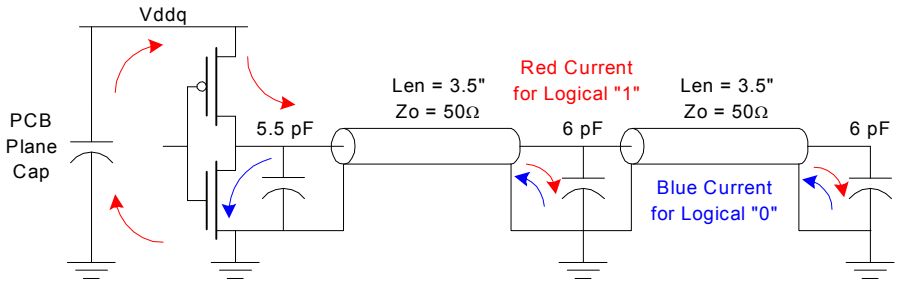

Figure 36.1. CMOS I/O Driver Circuit

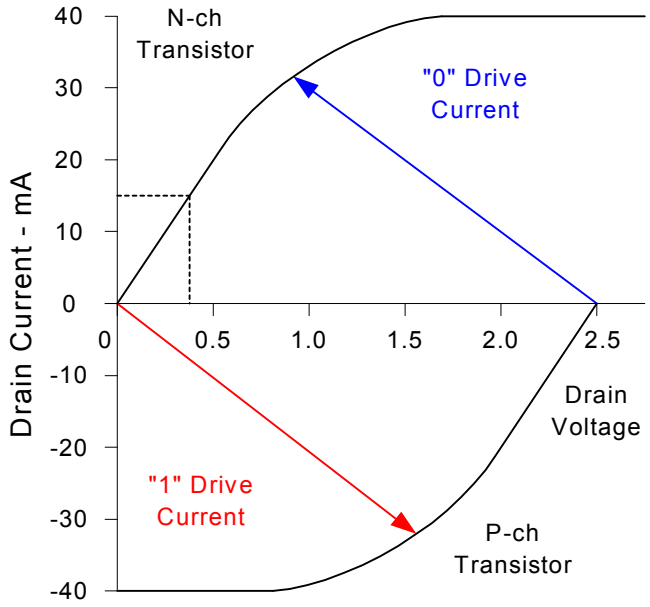

Figure 36.2. Output Transistor Drive Characteristic

Typically, a datasheet will specify the drive current at a TTL down level of 0.4V (as depicted by the dotted line in Figure 36.2.). This may be called a 15 mA driver. However, at a drain voltage of 2.5V, this transistor can produce 40 mA. This current can turn on in a few hundred picoseconds. It is the magnitude of this current, its rise time and its duration that is of interest in designing the power distribution system.

The drive characteristic of the transistor and the impedance of the transmission line connecting the driver to the load determine the amplitude of the Idd current. In Figure 36.2, the Red vectors indicate the current for a rising output signal and the Blue vectors represent the current for a falling output signal. It should be noted that current from the Vddq power supply occurs only on a rising transition. On a falling transition, the N-ch transistor discharges the capacitive load to ground.

The amplitude of the P-ch drive current is determined by the transistor drive characteristic and the initial output voltage step into the 50Ω transmission line. In Figure 36.2, this is represented by the Red vector and is a 50Ω load line. The load line shows that the transistor can pull the transmission line up to 1.6V at the time it turns on. The current spike amplitude can be calculated using Equation 36.1

Idd = dV / Zo
Idd = 1.6V / 50Ω = 32 mA

Equation 36.1. Current Spike Amplitude Equation

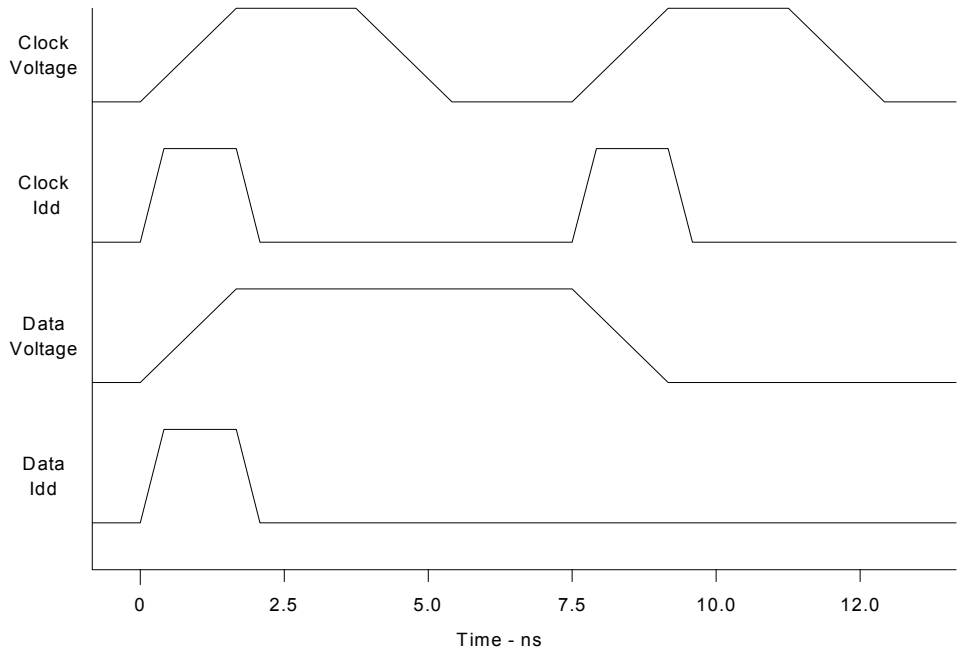

Figure 36.3. I/O Driver Waveforms

Figure 36.3 shows a diagram of the I/O driver waveforms for a Clock and Data pin. The clock voltage turns on and then off in each clock period but the data voltage waveform makes a transition only on the rising edge of the clock. The P-ch transistor turns on and draws current from the Vddq supply when the output voltage transitions from low to high. The N-ch transistor turns on and discharges the capacitive load to GND when the output voltage transitions from high to low. It is these current pulses from the Vddq supply that determine the power distribution requirements.

Clock Power: $P = C \times V^2 \times F$

Data Maximum Power $P = ½ \times C \times V^2 \times F$

Where: P = Power in Watts

C = Capacitive Load in Farads

V = Vddq in Volts

F = Clock Frequency in Hz

Clock Power: $P = 50pF \times 2.5^2 \times 133\ MHz = 41.6\ mW$

Data Maximum Power $P = ½ \times 50pF \times 2.5^2 \times 133\ MHz = 20.8\ mW$

Vddq =	2.5	Volts							
Freq. =	133	MHz							
P-ch Id =	32	mA							
			Maximum Power			Heavy use Power		Typical use Power	
I/O Type	Qty	Load	Idd Peak	Power Each	Power Total	Transiston Probability	Power Total	Duty Cycle	Power Total
		pF	mA	mW	mW		mW		mW
Clock	1	50	32	41.6	41.6	100%	41.6	100%	41.6
Address	40	50	1280	20.8	831.3	25%	207.8	50%	103.9
Data	72	50	2304	20.8	1496.3	50%	748.1	50%	374.1
			----------		----------		----------		----------
Total =			3616		2369		998		520

Table 36.2. Microprocessor I/O Power Estimate

Table 36.2 presents the estimate of the maximum power dissipated in the microprocessor I/O drivers. At idle, this drops to the power for the clock driver, which is 41.6 mW. Depending on the activity, the power can change from the minimum to maximum at any frequency. Therefore, the power distribution system must be able to maintain adequate low impedance throughout the entire frequency range from DC to several harmonics above the clock frequency.

I/O Peak Current **Ip = (72 + 40 + 1) x 32 mA = 3.62 Amps**

Equation 36.2. I/O Peak Current Equation

I/O Heavy Use DC Current Idc = 0.998 Watts / 2.5 V = 0.40 Amps

Equation 36.3. I/O Heavy Use DC Current Equation

The "Peak Current" occurs only when all the signals transition from low to high. This can occur but happens infrequently. However, the power distribution system must be designed to sustain all load changes that can occur. The peak current lasts only long enough to charge the capacitive loads. It gets averaged out over a cycle. The peak current can be calculated using Equation 36.2.

The "Maximum Power" and DC current can occur only when a very specific diagnostic test program is run to toggle every bit on every clock cycle. It is not possible to have these worst-case programs running on all the system buses at the same time.

The "Heavy Use Power" is the maximum power that can be sustained for short periods of time when the system is in normal operation and is used to size the power supplies. In order to determine this power, the designer must determine how the circuit is used and assign a transition probability to each type of signal. The clock will switch 100% of the time. The data bus ships binary data that has a 50% probability of two consecutive data bits having the same value. The address bus also ships binary data but typically operates in a small address range for a given program so all the bits are not likely to change. The Heavy Use Power can be calculated using Equation 36.3.

A microprocessor with a Cache has a high probability of getting data from the Cache and not going to I/O or main memory. The typical use is never more than 50%.

These numbers can be used for system power dissipation but since the data bus is bi-directional, a further refinement is needed to estimate the component power for thermal estimates. On read operations, only the microprocessor I/O data drivers are off and therefore dissipating zero power.

The numbers from Table 36.2 are used for the following purposes:

Parameter	Used For
Peak Current	IC Package ground bounce estimates
Maximum Power	Very High Frequency Decoupling
Heavy Use Power	Low frequency decoupling and Power Supply design
Typical Use Power	System thermal analysis.

If the "Maximum Power" is used to calculate the power supply size or the thermal characteristics, the system can be over designed by a factor of 3.

Series Terminated HSTL Bus

Series terminated drivers work well with point-to-point signals that have one driver, one load and a round trip transmission delay that is less than the clock period. A series terminated driver is designed to drive a half amplitude signal into a transmission line. The light load at the end of the line looks like an open circuit and produces a 100% reflection of the driving signal. This results in a full amplitude clean signal at the load and a lower amplitude current spike at the driver compared to the previous example. The reflection from the load travels back to the driver and is terminated by the driving impedance.

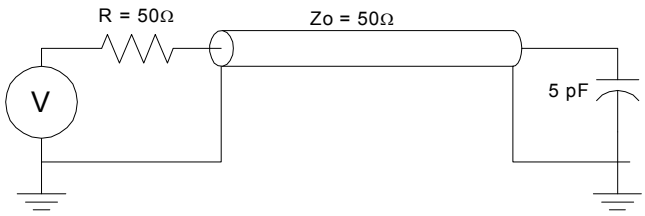

Figure 36.4. Equivalent Circuit for a Series Terminated Driver

The ideal series terminated driver would have a low output impedance voltage source and a resistor in series with the input of the transmission line. The resistor value should be the same as the impedance of the line. An equivalent circuit for a series terminated driver is depicted in Figure 36.4.

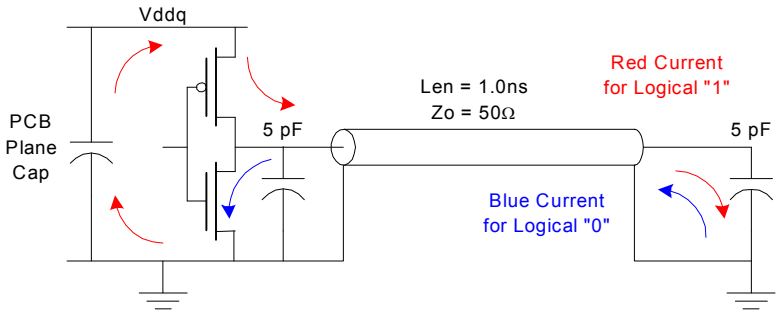

Figure 36.5. HSTL Series Terminated Driver and Load

There is no need for a separate series resistor if the driving voltage source impedance can be controlled to be equivalent to the transmission line impedance. The drive transistor size can be adjusted so that it produces a half amplitude transition into the transmission line. The impedance of the on drive transistor acts as the termination. Figure 36.5 shows a half-series transmission line with a series terminated driver and load.

Figure 36.6 shows the drive characteristic of the N-ch and P-ch transistor shown in Figure 36.5. The vectors represent dynamic load lines and the current transition. For a “0” to “1” transition, the P-ch transistor drives 18 mA (red vector) into the transmission line to produce a 0.9 volt step. For a “1” to “0” transition, the N-ch transistor draws 18 mA (blue vector) from the transmission line to produce a –0.9 volt step.

Current flows from the power supply only on positive transitions and this current lasts for a time equal to the round trip time of the transmission line. Once the reflection reaches the driver, the current ceases to flow. On negative transitions there is no power supply current. The N-ch transistor pulls the transmission half way from Vddq to ground. This ground current also lasts the round trip time of the transmission line.

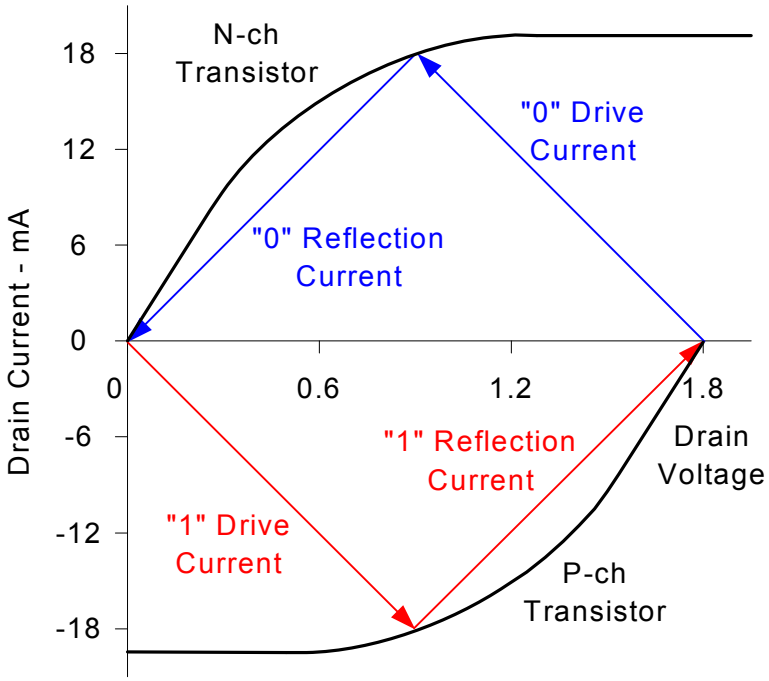

Figure 36.6. HSTL I/O Drive Characteristics

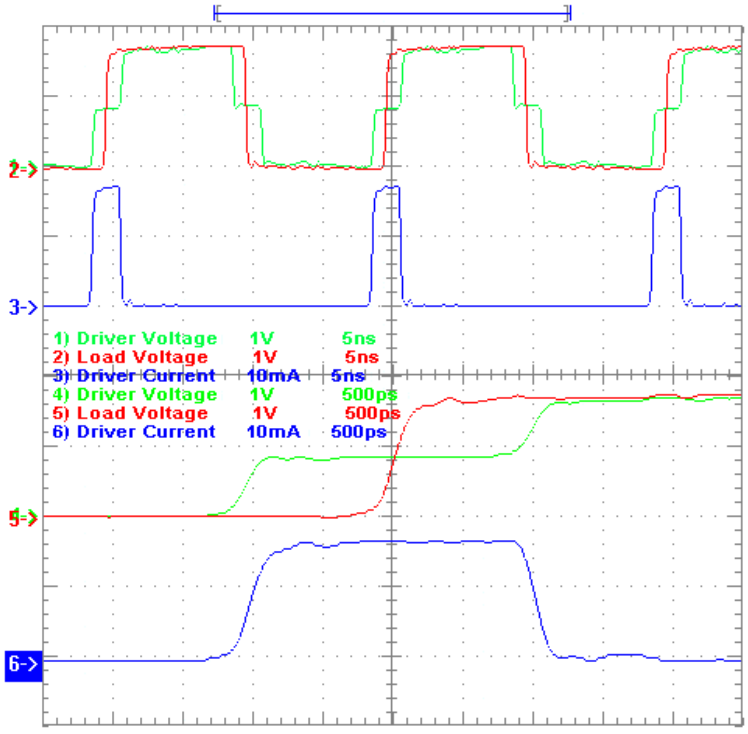

Figure 36.7. HSTL 100 Mb/s Waveforms

Figure 36.7 shows the voltage and current waveforms for the HSTL data signals clocked at 100 MHz (Data rate of 100 Megabits per second). The top half of the graph has a time scale of 5ns per division. The bottom half is expanded to show the same waveforms at 500ps per division. The Green waveform is the voltage at the driving end of the line and the Red waveform is the voltage at the receiver. The Blue waveform show the P-ch driver current pulses that are 18 mA in magnitude and last for a little over 2.0 ns or the round trip delay of the transmission line. The slight increase over 2.0 ns is due to the 5 pF capacitive load at the driver and receiver.

The above waveforms were measured on a Tektronix 11801C Sampling Oscilloscope with a 20 GHz bandwidth. The voltage and current waveforms shown in Figure 36.7 have about a 300ps rise time. This produces energy at very high frequencies. Figure 36.8 depicts the same current pulse waveform measured on the 50Ω input to a HP8594EM Spectrum Analyzer. This shows that most of the energy has a frequency spectrum far above the 100 MHz clock rate. Almost all of this energy must be supplied by the PCB power plane capacitance. The discrete decoupling capacitors are not effective above about 200 MHz.

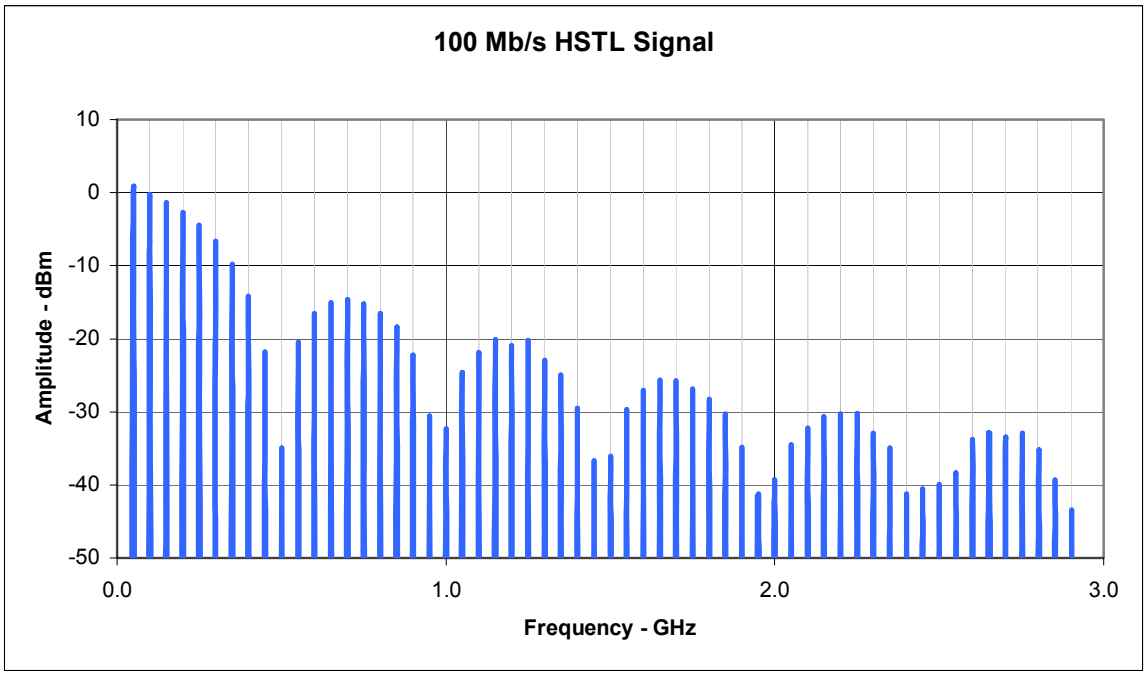

Figure 36.8. Spectrum Analysis of HSTL Waveform

Figure 36.8 depicts a single HSTL series terminated driver running at a clock rate of 100 MHz. Typically, HSTL buses are run at far higher clock rates and they have many signals in parallel on the bus.

Table 36.3 shows the bus capacitive load estimate for a 64-bit SPI-4.1HSTL bus running at 200 MHz. The current pulses will be the same amplitude and width as for the HSTL bus running at 100 MHz but they occur twice as often.

Unit	Feature	Qty	Each	Total
			pF	pF
Driver	Pin	1	5.0	5
Load	Pin	1	5.0	5
PCB	Via	4	1.3	5.2
PCB	Wire - in	6	3.4	20.4

			Total =	35.6

Table 36.3. HSTL Bus Capacitive Load Estimate

Vddq =	1.8	Volts								
Freq. =	200	MHz								
P-ch Id =	18	mA @ 0.9V								
			Maximum Power			Heavy use Power			Typical use Power	
I/O Type	Qty	Load	Idd Peak	Power Each	Power Total	Transiston Probability	Duty Cycle	Power Total	Duty Cycle	Power Total
		pF	mA	mW	mW			mW		mW
Clock	1	36	18	23.3	23.3	100%	100%	23.3	100%	23.3
Control	18	36	324	11.7	210.0	50%	39%	40.9	23%	24.1
Data	68	36	1224	11.7	793.2	50%	78%	309.3	47%	186.4
			----------		----------			----------		----------
Total =			1566		1026			374		234

Table 36.4. 64-bit HSTL Bus Power Estimate

The SPI-4.1 interface is used on a SONET OC192 Framer that runs at 9.953 Gbps. With a 64-bit wide bus, this translates into a maximum frequency of 155.52 MHz. The bus runs with a 200 MHz clock to allow for a little slack time for filling and emptying the FIFOs. Therefore, the maximum sustainable data rate is 155.52/200 or 78%. Typical use of a SONET link is 60% of full bandwidth. Table 36.4 shows the breakdown of the bus power estimate for the 64-bit HSTL bus.

It should be noted that with the reduced voltage of 1.8V and the series terminated drivers, both the peak current and power dissipation are much less for this bus compared to the example depicted in Table 36.5. This is true even though the bus in Table 36.4 is running at a much higher frequency.

SPI-4.2 LVDS Parallel Terminated Bus

The SPI-4.2 bus is used on a newer class of SONET OC192 framers. It uses LVDS (Low Voltage Differential Signals) that allow operation at 800 MHz with half the total pins used in the SPI-4.1 version. LVDS uses parallel 100Ω termination resistors that are usually built into the ASIC receiver cell. A parallel termination consumes DC power all the time but allows for much higher frequency operation than the series terminated approach.

The main advantage of terminated differential signals is that all of the transients on the two signal wires are of equal and opposite polarity. At high frequencies they cancel out. This applies not only to the power subsystem but also the SSO noise on the IC package and the ground return currents in the PCB.

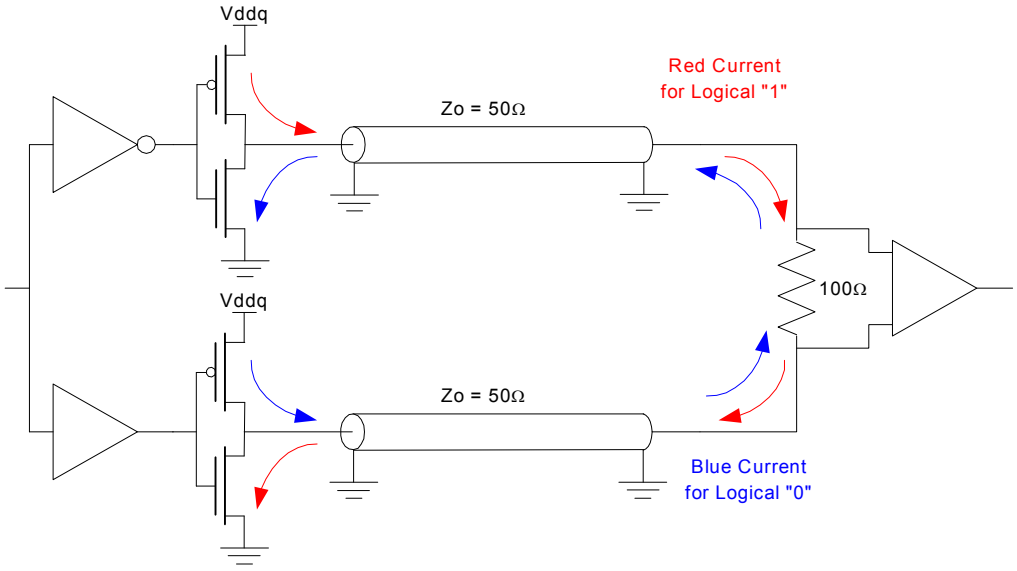

Figure 36.9. LVDS I/O Circuit Diagram

Figure 36.9 is the circuit diagram of an LVDS driver/receiver pair. The signal amplitude is typically 400 mV at the driver. This results in a current of 4 mA from the drivers. For a high signal level (“1”), the current (red arrows) flows from Vddq through the P-channel transistor at the top of the diagram through the terminating resistor and to ground through the N-ch transistor at the bottom of the diagram. For a low signal level (“0”), the current (blue arrows) flows from Vddq through the

bottom P-channel transistor; in the opposite direction through the terminating resistor and to ground through the top N-channel transistor.

The power supply current is the same 4 mA per signal for a “1” or a “0” level. There is a small amount of additional power dissipated in the driver circuit and receiver circuit but it is small compared to the DC load.

The amount of capacitance from the package, vias and transmission lines does not affect power dissipation. It only affects the delay of the path.

The power supply voltage, DC current and number of signals determines the power dissipation. For the SPI-4.2 interface there are 16 data signals, 4 control signals and a clock.

Power = Ns x Vddq x Is

Power = 21 x 1.8V x 4 mA = 151 mW

Equation 36.4. Power Dissipation Equation

The only special consideration is for thermal calculations. Some of the power is dissipated in the driver and some in the terminating resistor in the receiver chip. In most cases, the SPI-4.2 bus has an identical transmit and receive side going in opposite directions between the same two chips so this is also not an issue.

The power dissipation is a constant 151 mW independent of the data rate or line length. This is less power and half the pins of the previous bus but it does require higher speed ASIC technology. Equation 36.4 can be used to determine the system power dissipation.

The biggest advantage of this LVDS circuit is that the power dissipation is constant. **Very little decoupling is needed**.

DDR-SDRAM Memory Interface

Synchronous Dynamic Random Access Memory (SDRAM) has been in use for several years. The current form of this memory uses a Double Data Rate (DDR) on the bi-directional data bus between the memory DIMMs and the controller. The Command/Address bus operates in one direction with all signals driven by the controller and received by the memory chips. Clock rates for DDR-SDRAM range from 100 MHz to 200 MHz. This implies data rates of 200 Mb/s to 400 Mb/s.

Registered (buffered) DIMMs are typically used on high performance systems with four memory DIMMs. The Register/buffer chip of the DIMM powers the clocks and Command/Address signals that must drive all chips on the DIMM. The data pins are not buffered because they connect to one or two chips on the DIMM. The DDR-SDRAM subsystem block diagram is shown in Figure 36.10.

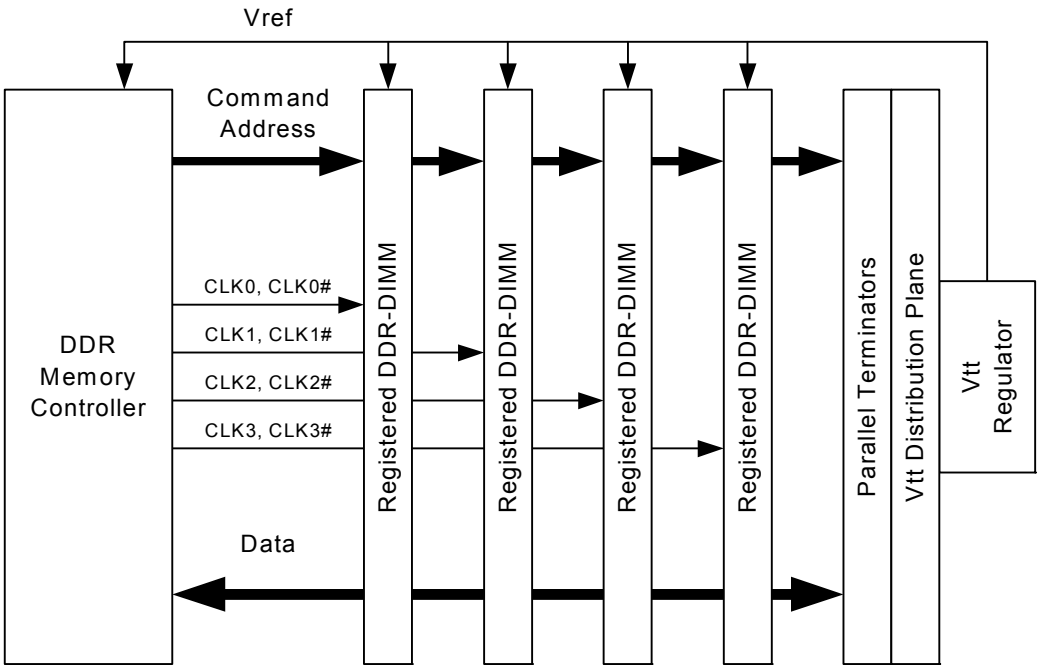

Figure 36.10. DDR-SDRAM Subsystem Block Diagram

DDR-SDRAMs operate with a 2.5V power supply. Clock signals are differential and terminated with a 120Ω resistor on each DIMM. All other signals are parallel terminated with approximately a 25Ω resistor to a 1.25V (1/2 Vdd) power supply. Figure 36.11 shows the DDR-SDRAM circuit diagram.

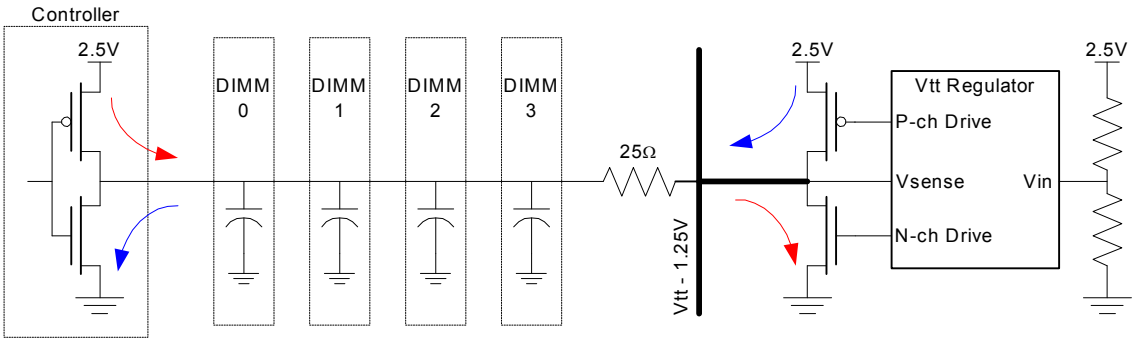

Figure 36.11. DDR-SDRAM Circuit Diagram

The power supply must be capable of sourcing or sinking load current. When the controller drives the signal to a "1", current flows from the 2.5V supply, through the driver P-ch transistor (red arrows), through the 25Ω termination resistor and to GND via the Vtt power supply N-ch power transistor. For a "0" signal level, the driver in the controller pulls the signal low. For the "0" signal level, current flows (blue arrows) from the 2.5V supply through the Vtt power supply P-ch power MOSFET, in the reverse direction through the 25Ω termination resistor and to GND through the controller N-ch transistor.

In active mode, the same DC current flows from the 2.5V supply for both a "0" or a "1" level. The data bus drivers are bi-directional and go into a high impedance state during the turnaround from write to read operations. There is at least one dead cycle where all drivers on the data bus are in the high impedance state. The Command/Address drivers operate in one direction so they do not turn off in active mode. However, many designs have tri-state drivers on the Command/Address lines to save power when the memory is inactive.

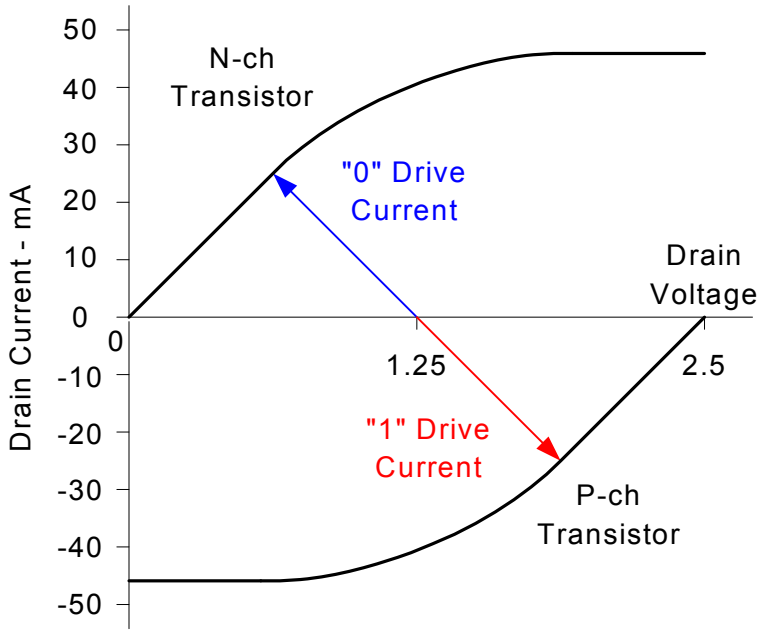

Figure 36.12. DDR-SDRAM Controller Drive Characteristics

Figure 36.12 shows the output driver DC characteristics on the controller. The drive current is approximately 25 mA for either a "0" or "1" level.

Signal	Type	Qty	Idd Each	Idd total	Itt Total
			mA		
Clocks	Diff	4	25	100	
Controll/Address	Single	24	25	600	600
Data - ECC - DQS	bi-di	80	25	2000	2000
				-------	-------
			Total =	2700	2600

Table 36.5. DDR-SDRAM Interface Power Supply Current

During active operation the bi-directional current of 2 Amps can turn on and off at any frequency depending on program activity. Switching from standby to active will turn on and off all of the current except for the clocks. Table 36.5 shows the characteristics and properties of the DDR-SDRAM interface power supply current.

Signal	Qty	Idd	Heavy Use			Typical Use		Standby	
		Circuit mA	Total mA	Duty Cycle	mW	Duty Cycle	mW	Duty Cycle	mW
Clocks	4	25	100	100%	250	100%	250	100%	250
Controll/Address	24	25	600	100%	1500	100%	1500	0%	0
Data - ECC - DQS	80	25	2000	100%	5000	50%	2500	0%	0
			-------		-------		-------		-------
		Total =	2700		6750		4250		250

Table 36.6. DDR-SDRAM Interface Power Dissipation

This design is good for applications that have low power utilization or in portable applications that conserve energy by going into the standby mode. The disadvantage of this approach is that it requires a small power supply plane in order to get the high frequency PCB plane capacitance. Table 36.6 presents the power dissipation breakdown for the DDR-SDRAM interface.

An alternate approach is to use parallel termination resistors at the end of each signal wire and eliminate the 1.25V power supply. This approach simplifies the design but dissipates power all the time in the terminators even if all the drivers are in the high impedance state. This alternate design approach is depicted in Figure 36.13.

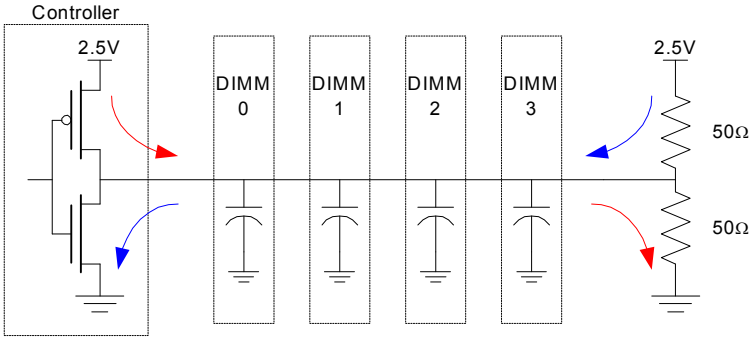

Figure 36.13. DDR Alternate Design Using Two Resistors

Total current per active signal increases from 25 mA, shown in Table 36.5, to 37.5 mA with the parallel termination as depicted in Table 36.6. In the standby case, the current increases from zero to 18 mA.

		Active	Standby									
Signal	Qty	Idd	Idd	Heavy Use			Typical Use			Standby		
		Circuit mA	Circuit mA	Duty Cycle	Total mA	mW	Duty Cycle	Total mA	mW	Duty Cycle	Total mA	mW
Clocks	4	25.0	25.0	100%	100	250	100%	100	250	100%	100	250
Controll/Address	24	37.5	18.0	100%	900	2250	100%	900	2250	0%	432	1080
Data - ECC - DQS	80	37.5	18.0	100%	3000	7500	50%	2220	5550	0%	1440	3600
					------	------		------	------		------	------
			Total =		4000	10000		3220	8050		1972	4930

Table 36.7. DDR-SDRAM Interface Power Dissipation with Parallel Terminators

PCB Total Power Estimate

A spreadsheet, as depicted in Table 36.8, works well to sum up the total card power. The estimates provided in Table 36.7 along with the component data sheets can be used to get the power for each component.

The example is for a small computer motherboard with a graphic controller chip.

	Heavy Use Power			Typical Use Power			Minimum Use Power		
Component	2.5 V	1.8 V	1.3V	2.5 V	1.8 V	1.3V	2.5 V	1.8 V	1.3V
			uP Core			uP Core			uP Core
	Watts	Watts	Watts	Watts	Watts	Watts	Watts	Watts	Watts
Microporcessor		1.0	6.0		0.6	4.0		0.1	2.0
uP Controller	6.7			4.3			0.8		
DDR-SDRAM	5.0			2.5			1.0		
Graphic Processor	0.6	6.0		0.5	5.0		0.4	4.0	
Graphic Memory Chip	1.5			1.0			0.4		
Misc	1.0	1.0		1.0	1.0		1.0	1.0	
	----------	----------	----------	----------	----------	----------	----------	----------	----------
Subtotal - Watts	14.8	8.0	6.0	9.3	6.6	4.0	3.6	5.1	2.0
- Current	5.9	4.4	4.3	3.7	3.7	2.9	1.4	2.8	1.4
Power Supply Efficiency	88%	86%	75%	88%	86%	75%	88%	86%	75%
Power Supply Loss	2.0	1.3	2.0	1.3	1.1	1.3	0.5	0.8	0.7
	----------	----------	----------	----------	----------	----------	----------	----------	----------
Power per Supply	16.8	9.3	8.0	10.6	7.7	5.3	4.1	5.9	2.7
Total Input Power			34.1			23.6			12.7

Table 36.8. PCB Total Power Estimate

CHAPTER 37: EXAMPLE POWER SUBSYSTEM DESIGN

By John Zasio

Power Supply Requirements

In Chapter 36, the requirements were presented for a total power supply tolerance of +/- 5%. Table 37.1 shows how these requirements apply to the card depicted in Table 36.8.

Voltage	Supply	DC Drop			AC Ripple			
	Regulation	Tolerance	Maximum	Required		Voltage	Maximum	Required
	2.0%	1.0%	Load	Impedance	Tolerance	Drop	Ripple	Impedance
	mV	mV	Amps	mΩ		mV	Amps	mΩ
2.50	50.0	25.0	5.9	4.2	2.0%	50.0	4.5	11.1
1.80	36.0	18.0	4.4	4.1	2.0%	36.0	1.6	22.5
1.30	26.0	13.0	4.3	3.0	2.0%	26.0	2.9	9.0
1.25	0.0	12.5	2.7	4.6	4.0%	50.0	2.6	19.2

Table 37.1. Power Distribution Requirements

The “Maximum Ripple” is the difference between the heavy load current and the minimum load current from the PCB total power estimate. The required impedance is calculated from the following equation based on the allowable voltage change and the current.

$$Z = dV / dI$$

Where: Z = Required Impedance

dV = Voltage Change

dI = Current Change

Equation 37.1 Power Distribution System Impedance Requirements Equation

In order to design a power distribution system to meet these impedance requirements, it is necessary to select a sufficient quantity and type of decoupling capacitors to provide the required impedance over the entire frequency range of interest.

For frequencies above 150 MHz, the discrete capacitors are no longer effective. **At these high frequencies, the PCB power plane capacitance is required to provide the low impedance**.

PCB Stackup

The PCB depicted in Figure 37.1 needs two main power planes and two smaller planes for the uP-Core and the DDR Vtt supplies. A 10-layer board was chosen for this example.

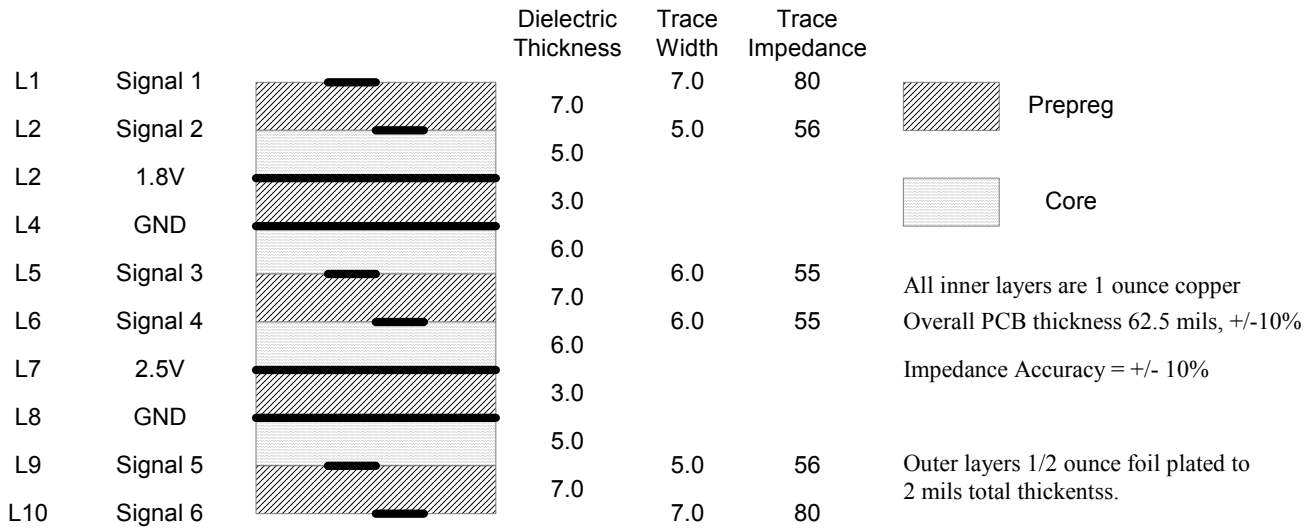

Figure 37.1. 10-Layer PCB Stackup

2.5V Distribution System

The process for designing the distribution system is as follows:

1. Estimate the power consumed by the load. In this example it is 14.8 Watts or 5.9 Amps.
2. Estimate the maximum change in load current. 4.5 Amps
3. Determine the peak.AC ripple voltage 2.5V x 2% = 50 mV
4. Calculate the required impedance. Z = dV / dI = 50mV / 4.5A = 11.1 mΩ
5. Design the PCB for sufficient power plane cap. Needed for frequencies above 150 MHz
6. Select a sufficient quantity and type of capacitors to obtain the required impedance over the entire frequency range.

The 2.5V supply is more typical of power distribution systems. For this design, we are assuming the PCB is 8” x 10” in size and the dielectric layers are 3 mils thick. This produces 32 nF of power plane capacitance. The listing of decoupling capacitors for the 2.5V distribution system is presented in Table 37.2.

		Each			Total		
Device	Qty	Cap	ESR	ESL	Ceff	ESR	ESL
DC/DC Converter	1		1.00E-04	7.00E-08		1.00E-04	7.00E-08
Tantalum 330 uF	3	3.30E-04	2.30E-02	3.00E-09	9.90E-04	7.67E-03	1.00E-09
0603 Ceramic 1.0 uF	3	1.00E-06	2.00E-02	7.00E-10	3.00E-06	6.67E-03	2.33E-10
0603 Ceramic 0.1 uF	6	1.00E-07	4.00E-02	6.00E-10	6.00E-07	6.67E-03	1.00E-10
0603 Ceramic 0.01 uF	12	1.00E-08	6.00E-02	5.00E-10	1.20E-07	5.00E-03	4.17E-11
PC Board 8" x 10"	1	3.20E-08			3.20E-08		

Table 37.2. Decoupling Capacitors For 2.5V Supply

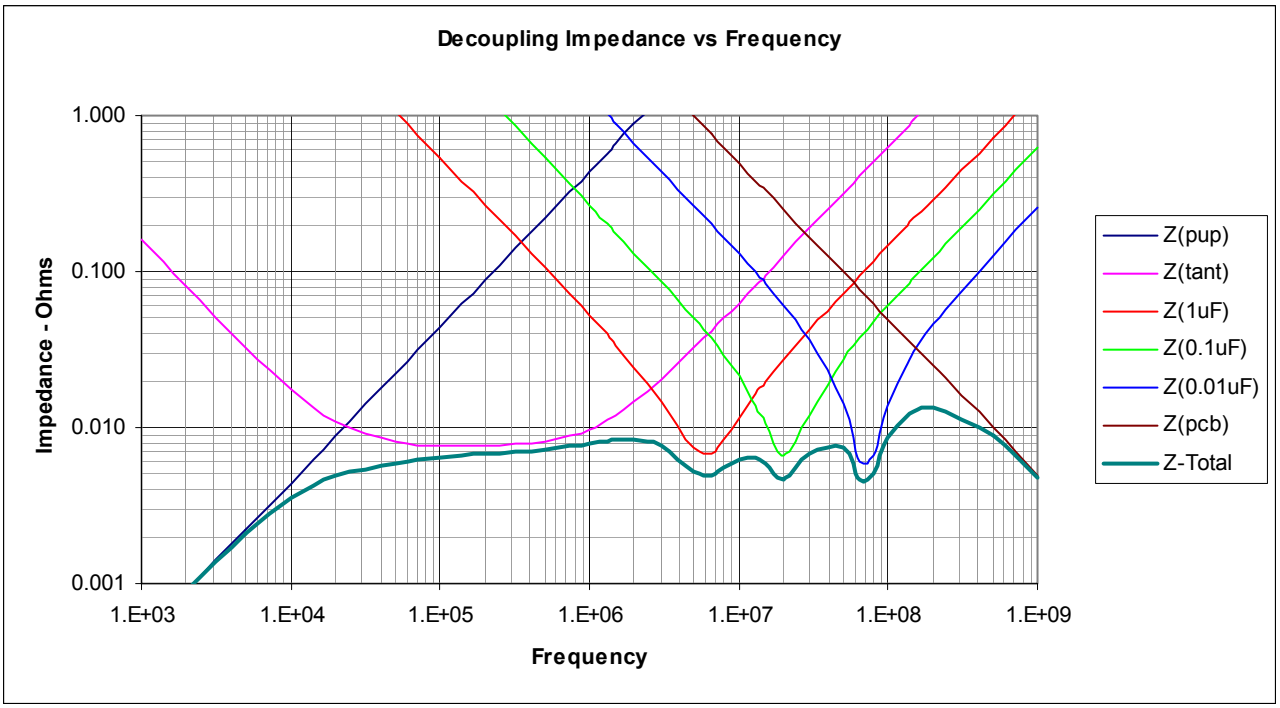

Figure 37.2. Impedance vs Frequency for 2.5V Supply

The impedance plot in Figure 37.2 shows the impedance of each capacitor type and the total impedance (Z-Total) of all the capacitors in parallel. This number of capacitors gives a total impedance below the 11 mΩ target for the entire frequency range. This plot was done in an Excel spreadsheet and does not accurately predict the PCB plane resonance at very high frequencies.

It should be noted that the quantity of decoupling capacitors used is larger for the lower value capacitors. This is because the lower value capacitors have a higher ESR that allows the use of more capacitors in parallel to achieve a low inductance without also having a very low ESR that may cause high frequency resonance.

1.8V Distribution System

The 1.8V supply has the lowest load and the smallest change in load current. Therefore, the required power supply impedance is significantly higher at 22 mΩ. The list of decoupling capacitors for the 1.8V distribution system is presented in Table 37.3.

		Each			Total		
Device	Qty	Cap	ESR	ESL	Ceff	ESR	ESL
DC/DC Converter	1		1.00E-04	7.00E-08		1.00E-04	7.00E-08
Tantalum 330 uF	2	3.30E-04	2.30E-02	3.00E-09	6.60E-04	1.15E-02	1.50E-09
0603 Ceramic 1.0 uF	2	1.00E-06	2.00E-02	7.00E-10	2.00E-06	1.00E-02	3.50E-10
0603 Ceramic 0.1 uF	4	1.00E-07	4.00E-02	6.00E-10	4.00E-07	1.00E-02	1.50E-10
0603 Ceramic 0.01 uF	8	1.00E-08	6.00E-02	5.00E-10	8.00E-08	7.50E-03	6.25E-11
PC Board 8" x 10"	1	3.20E-08			3.20E-08		

Table 37.3. Decoupling Capacitors For 1.8V Supply

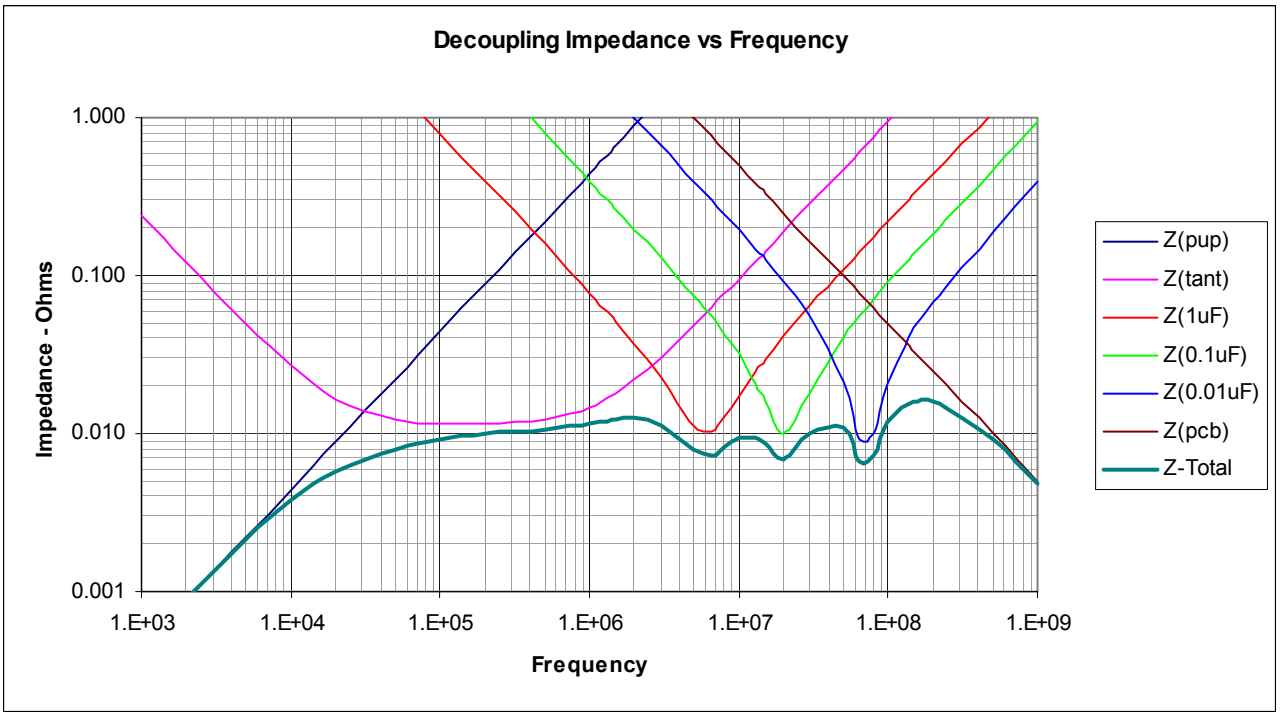

Figure 37.3. Impedance vs Frequency for 1.8V Supply

This plane requires a smaller quantity of decoupling capacitors than the 2.5V plane in order to achieve the target impedance that is shown in Figure 37.3.

DDR-SDRAM Vtt

The DDR-SDRAM Vtt power supply is a special case since it is driven from the 2.5V plane and it must source or sink 2.7 Amps. Because the power source is always from the 2.5V supply and the Vtt supply reverses current, it is best to split the decoupling capacitors with half going to GND and half to the 2.5V plane. If all the decoupling capacitors were between Vtt and GND then a large load transient between the 2.5V and Vtt supply would see the 2.5V capacitors in series with the Vtt capacitors as the return current path. This would result in a larger impedance.

Another major difference in this type of supply is that this is a local power supply for one memory subsystem. Typically, there is an insufficient number of planes in the PCB to dedicate an entire plane for this voltage. The DDR-DIMM array can be wired with only two signal layers. Signal planes can be filled in under the DIMM array in order to provide sufficient power plane capacitance as shown in Figure 37.4. The result is a much lower amount of capacitance on this plane than on a full board plane such as 2.5V. Figure 37.4 shows the PCB stackup using plane fill for the Vtt power plane.

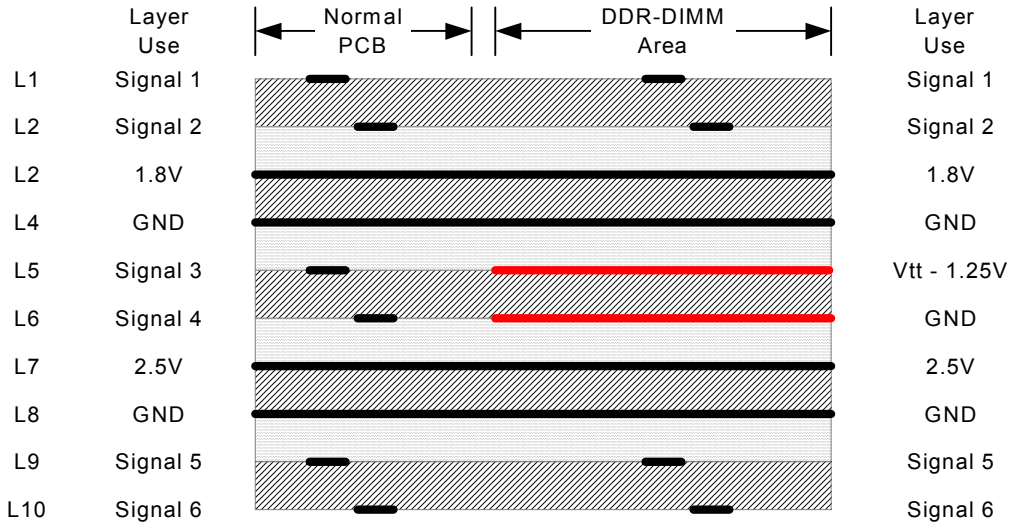

Figure 37.4. PCB Stackup showing Signal Plane Fill (Red) for the Vtt Power Plane

The absolute voltage value for this supply is not critical. What is critical is that it be approximately ½ the 2.5V supply voltage and have a reasonably low AC ripple due to changes in load current. The operational amplifier in the small power supply will keep a relatively constant DC voltage. Therefore, the requirement for this supply is +/- 1% DC Drop and +/- 4% for AC ripple. This makes the design a little more practical considering the small area of the power plane.

Memory DIMMs have high frequency 0402 decoupling capacitors and the memory chips themselves provide a significant amount of very high frequency capacitance. The DIMMs are 5.25" x 1.2" in area. Each DIMM has a power plane pair that contributes to the high frequency capacitance.

DDR-SDRAM memory chips are mounted on 184-pin DIMMs that are 5.25" in length. The socket takes up a footprint that is 6" x 0.6" in area. Four DIMMs plus the power supply and end of the controller chip typically take up an area that is 4" x 7". The plane fill for the Vtt layer shown in the stackup in Figure 37.4 has a GND plane on both sides. The dielectric thickness on the top side is 6 mils and on the bottom side it is 7 mils. This will produce a capacitance of 222 pF/in^2. The 5" x 7" area will have a total capacitance of 7.78 nF. The additional capacitance from the DDR-DIMM power planes will double this amount of capacitance. Figure 37.5 shows the decoupling capacitor connections for the Vtt plane.

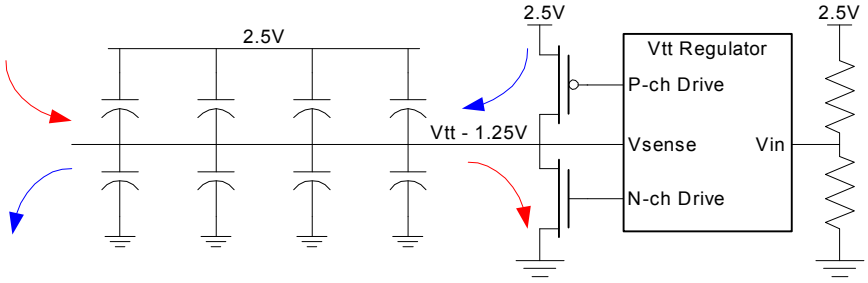

Figure 37.5. Vtt Decoupling Capacitor Connections

		Each			Total		
Device	Qty	Cap	ESR	ESL	Ceff	ESR	ESL
DC/DC Converter	1		1.00E-04	7.00E-08		1.00E-04	7.00E-08
Tantalum 330 uF	2	3.30E-04	2.30E-02	3.00E-09	6.60E-04	1.15E-02	1.50E-09
0603 Ceramic 1.0 uF	2	1.00E-06	2.00E-02	7.00E-10	2.00E-06	1.00E-02	3.50E-10
0603 Ceramic 0.1 uF	4	1.00E-07	4.00E-02	6.00E-10	4.00E-07	1.00E-02	1.50E-10
0603 Ceramic 0.01 uF	12	1.00E-08	6.00E-02	5.00E-10	1.20E-07	5.00E-03	4.17E-11
PC Board & DIMMS	1	3.20E-08			3.20E-08		

Table 37.4. Decoupling Capacitors for the Vtt Supply

Table 37.4 shows the listing of decoupling capacitors used for the Vtt supply. It should be noted that the only difference between this and the 2.5V supply is the PCB capacitance. This would result in a high impedance parallel resonance at about 300 MHz if it were not for the capacitance of the DIMMs. Figure 37.6 shows the amount of impedance vs. frequency for the 1.25V Vtt supply.

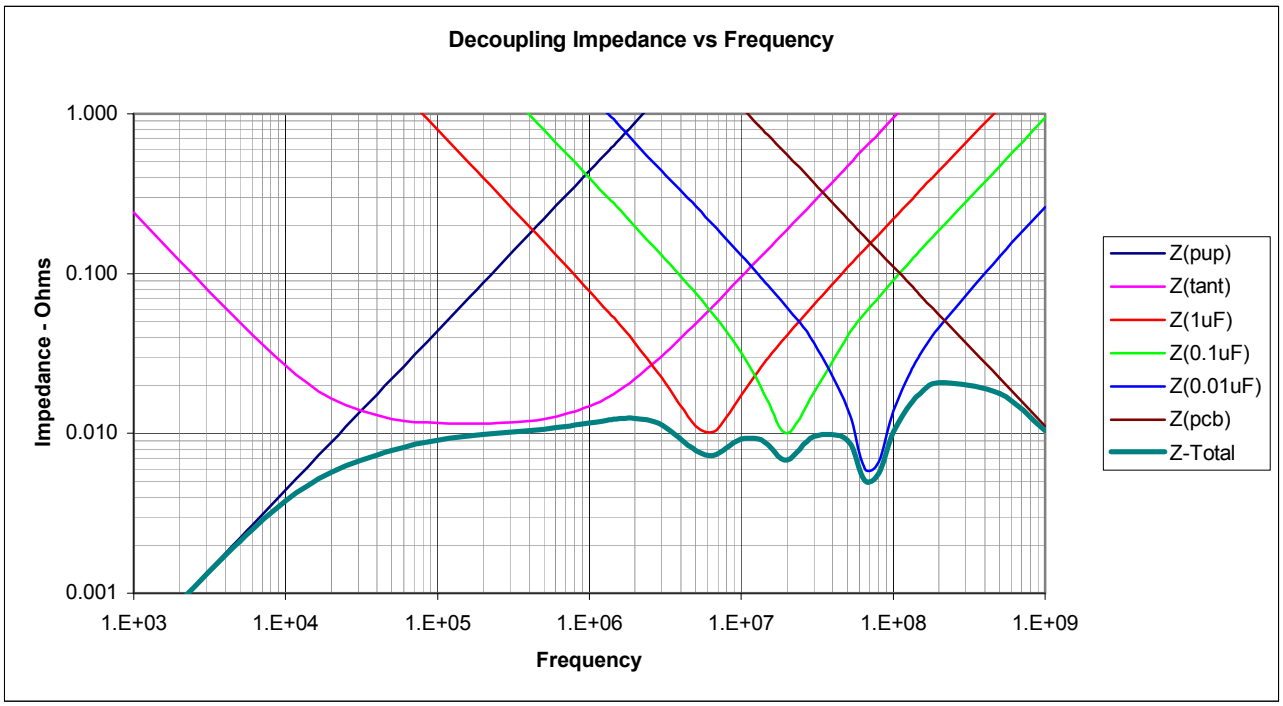

Figure 37.6. Impedance vs Frequency for the Vtt (1.25V) supply

1.0 GHz Microprocessor Core

Microprocessors are usually made from the most advanced technology and tend to have the lowest core voltage requirement of any chip on the card. The low core voltage and large variation in load current results in a low power distribution impedance requirement. The required impedance (Table 37.1) for this design is 9 mΩ.

Usually, there are several iterations of a particular microprocessor with the same footprint but with different core voltage requirements. Therefore, there is an advantage in providing a local power supply for each microprocessor chip. The local power supply voltage can easily be adjusted for a processor technology upgrade and there is typically an insufficient number of PCB power planes.

The power planes for the microprocessor can be “Plane Fill” on signal layers similar to that used for the DDR-SDRAM memory. The current generation of high-speed microprocessors tends to have a smaller number of I/O pins than the previous generation because the cache is inside the processor chip. This reduces the number of signal pins and allows the chip signal routing to consume only two layers directly under the chip. This allows space on the PCB signal layers for the “Plane Fill”. The area available for plane fill under the processor chip will be smaller than under the memory. An area that is 4” x 4” is reasonable to obtain for this plane. This will produce a capacitance of 3.5 nF.

Microprocessor cores have a substantial amount of very high frequency decoupling capacitance on the chip itself. A typical value for this core capacitance is 25 nF. Many of the microprocessors also have small ceramic capacitors mounted on the top of the ceramic package to aide the decoupling.

		Each			Total		
Device	Qty	Cap	ESR	ESL	Ceff	ESR	ESL
DC/DC Converter	1		1.00E-04	7.00E-08		1.00E-04	7.00E-08
Tantalum 330 uF	3	3.30E-04	2.30E-02	3.00E-09	9.90E-04	7.67E-03	1.00E-09
0603 Ceramic 1.0 uF	3	1.00E-06	2.00E-02	7.00E-10	3.00E-06	6.67E-03	2.33E-10
0603 Ceramic 0.1 uF	6	1.00E-07	4.00E-02	6.00E-10	6.00E-07	6.67E-03	1.00E-10
0603 Ceramic 0.01 uF	12	1.00E-08	6.00E-02	5.00E-10	1.20E-07	5.00E-03	4.17E-11
PPC Chip	1	2.50E-08			2.86E-08		
PC Board	1	3.60E-09					

Table 37.5. PPC-Core Decoupling Capacitors

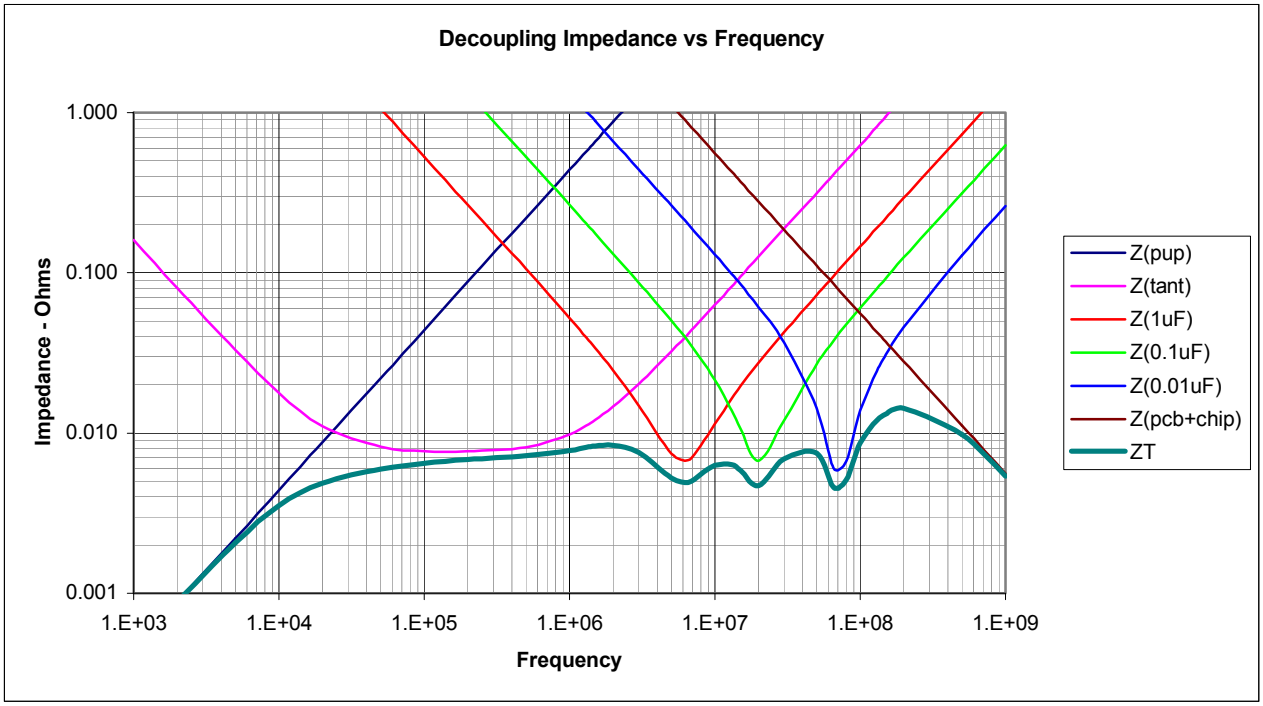

Figure 37.7. Impedance vs Frequency for PPC-Core

Determining How Much Plane Capacitance is Required to Support Fast Edges

The preceding chapters have made it very clear that the plane capacitance of a PCB is what supports the fast edges and signals in a logic design. A fair question is "How much is needed?" The answer depends upon what kinds of loads are placed on each power supply voltage. For nearly all logic designs, the loading that places the worst-case demand on plane capacitance is comprised of single ended data buses with all of the bits switching from a logic 0 to a logic 1 simultaneously. The current waveform shown in Figure 36.7 is typical of the load represented by this kind of logic.

The diagram on the left side of Figure 38.1 illustrates what is occurring. In order to switch a logic line from a 0 to a 1, it is necessary to charge up all of the parasitic capacitance on that line. This charge is drawn from the plane capacitor. It is reasonable to say that this operation is a charge transfer from one capacitor to another. As the charge is removed from the plane capacitor, its terminal voltage drops, resulting in a "ripple" voltage on that supply rail. This is the primary source of ripple. It is also a major source of EMI when a wire that exits the product is attached to VDD. This attachment happens most often when the wire has a logic signal of 1. Therefore, minimizing this ripple voltage has a beneficial effect on both EMI and logic stability.

The most precise way to determine the size of plane capacitor that will keep the ripple voltage within specification is to build the proposed plane structure in a 3D modeling tool such as Sigrity 2000, place a trapezoidal current waveform of the length and rise and fall time and proper current magnitude and observe the voltage ripple that results. If the ripple voltage is too large, the plane capacitor size is increased. (Place capacitance can be increased by reducing the dielectric thickness or by adding more planes until the desired capacitance has been reached.)

For most projects, access to such a modeling tool is not available. An alternative approach is to recognize that a charge transfer operation is taking place. Removing charge from one capacitor and placing it on the other causes the terminal voltage of the one capacitor to go down and the other to go up. Equation 37.2 shows the relationship between the size of the two capacitors and the amount the plane capacitor terminal voltage will change as this charge transfer takes place.

This equation is an approximation of the problem to be solved. It has been used many times to calculate the needed plane capacitance. In the next chapter, the ripple shown in Figures 38.3 and 38.4 is the result of using this equation. Figure 38.3 depicts a small PCB with a limited area allowing only the minimum plane capacitance. Figure 38.3 shows a very large PCB with far more plane capacitance than needed. The ripple that results in these two cases shows the affect of different size plane capacitors.

$$\frac{\Delta V}{V} \cong \frac{Cswl}{Cplane}$$

Where: ΔV is allowable ripple, V is power supply voltage, Cswl is the sum of the parasitic capacitance of all the lines that can switch simultaneously, and C plane is the plane capacitance needed.

Equation 37.2 Ripple Voltage vs. Plane Capacitance to Load Capacitance Ratio

The value of Cswl is calculated by summing up the parasitic capacitance of the transmission lines and the input capacitance of the inputs attached to those transmission lines. For most designs, the transmission lines will be 50 ohms. At this impedance, the parasitic capacitance is approximately 3.5 pF per inch. More precise values can be obtained using a 2D field solver to calculate Co.

A typical design might be a 32-bit bus, six inches long with a single load at the end of 2 pF. One member of the bus would have 6 x 3.5 pF + 2 pF or 23 pF. The total bus would have 32 x 23 pF or 736 pF. If 50 mV or ripple is allowed and the supply voltage is 2.5 volts, the plane capacitor would need to be 50 x 736 pF or 36 nF.

This is a conservative method for calculating the size of plane capacitor needed, but correlates well with the result shown in Figure 38.3. Clearly, if the plane capacitor is lower than that calculated, the ripple voltage will be higher. The example which follows, of a PCMCIA-based Ethernet adapter, illustrates this by measuring EMI before and after adding plane capacitance.

Figure 37.8 shows the ripple voltage associated with a 256-bit wide bus 3 inches long with a single load at the end of each line. This voltage waveform is exactly the same shape as the current pulse. This design should have had a plane capacitor or 1000 pF per square inch, but an error was made in the stackup resulting in only 100 pF per square inch. The symptom was the design was stable at low data rates but became “flaky” at high data rates. The 232 mV ripple spike happened when all 256 bits simultaneously switched from 0 to 1. The problem was solved by redesigning the stackup to provide the proper plane capacitance.

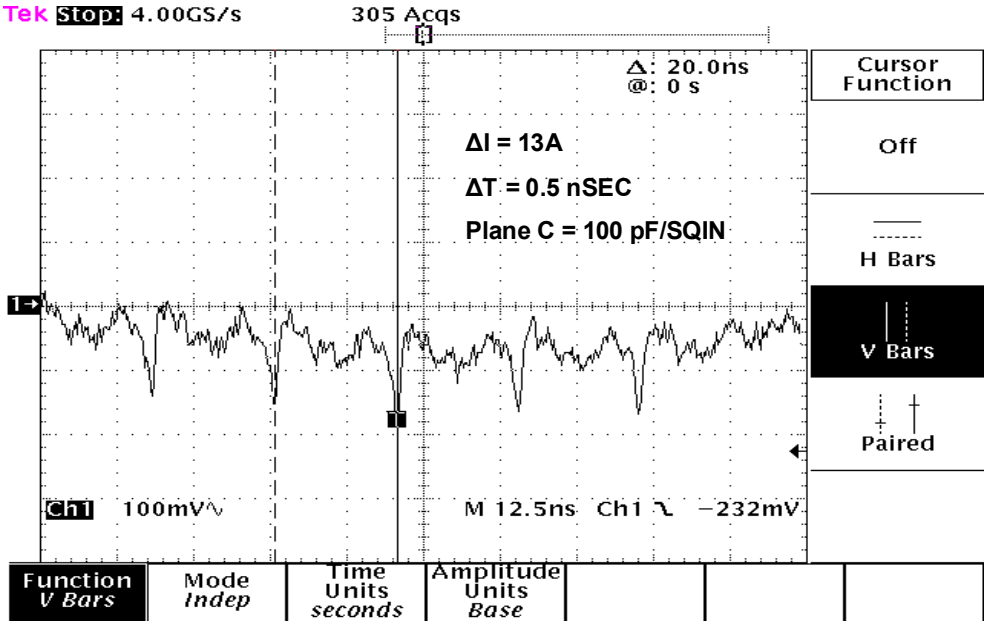

Figure 37.8. Voltage Spikes Associated With a 256-Bit Data Bus Switching 0 to 1

Signal Plane Fills as a Way to Increase Plane Capacitance

An obvious way to increase plane capacitance is by adding pairs of closely spaced planes to a stackup. This is always going to increase the PCB thickness and its cost. Often, there is significant unused space in signal layers that may be flooded with copper and this copper can then be used to create plane capacitance. This is a common way to add plane capacitance to very small PCBs such as PCMCIA and Flash cards. Figure 37.9 shows two PCB stackups that both have plane capacitance from power plane pairs and instructions on how to fill the unused space in signal layers to add capacitance.

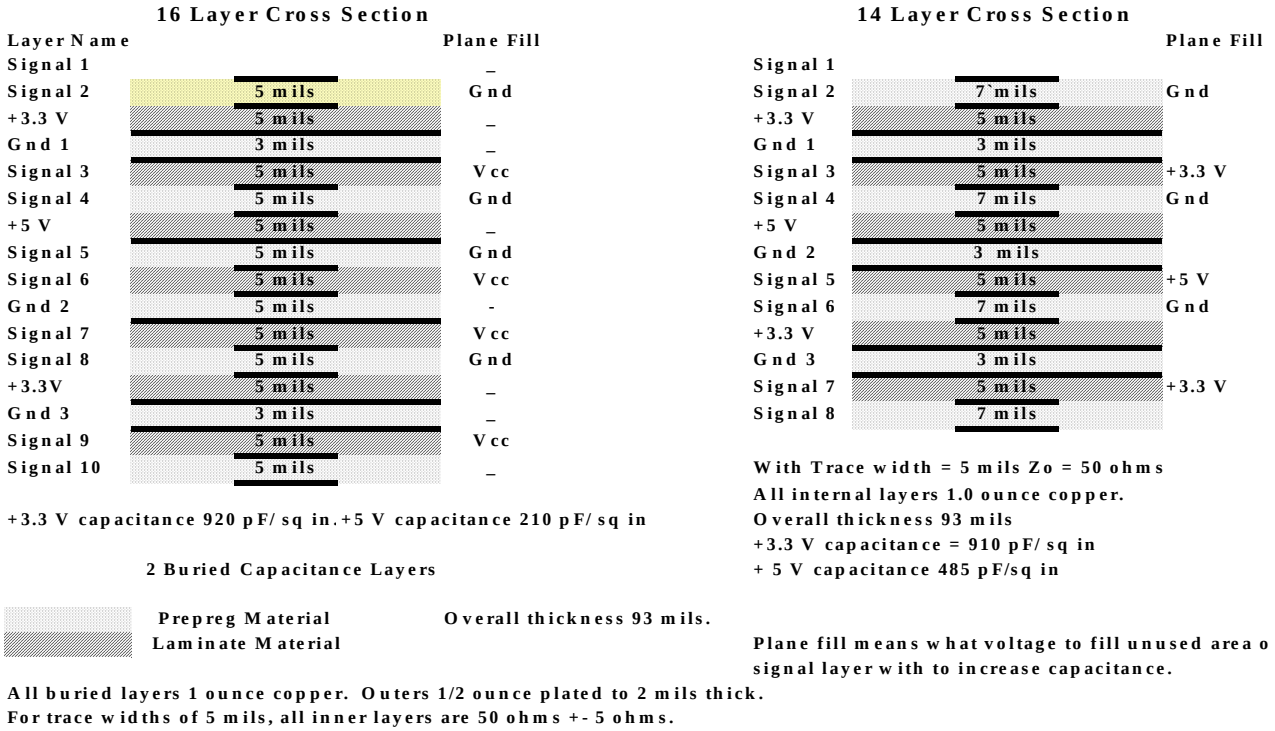

Figure 37.9. Two Stackups With Plane Capacitance and Signal Plane Fill

Very Small PC Board Power Subsystem Design

A consumer product such as a PCMCIA card has very severe size and cost requirements. These cards are typically six layer PC boards with only two power planes. Two small size power planes separated by two signal layers do not provide sufficient high frequency power plane capacitance. The small surface mounted 0402 capacitors do not help at frequencies above 200 MHz. The stackup for a six layer PCMCIA card is shown in Figure 37.10.

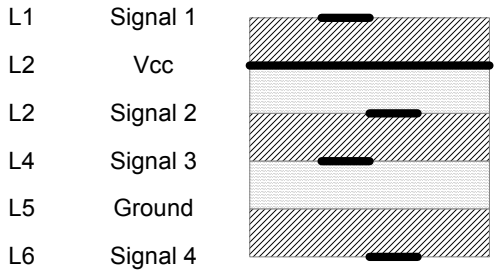

Figure 37.10. PCMCIA Card 6-Layer Stackup

An approach that can provide sufficient power plane capacitance is to use “Plane Fill” on all signal layers in areas not used for signal routing. Care must be taken to not leave isolated areas of plane fill. These areas must be connected to the Vcc or Ground plane with vias. The stackup for the PCMCIA card using plane fill is depicted in Figure 37.11. Figure 37.12 shows the same PCMCIA card with the signal layers filled with power planes. Figure 37.11 shows the emission test results for a six-layer PCMCIA card with and without signal layer plane fills. The PCMCIA card depicted contains a 100 BT Ethernet adapter.

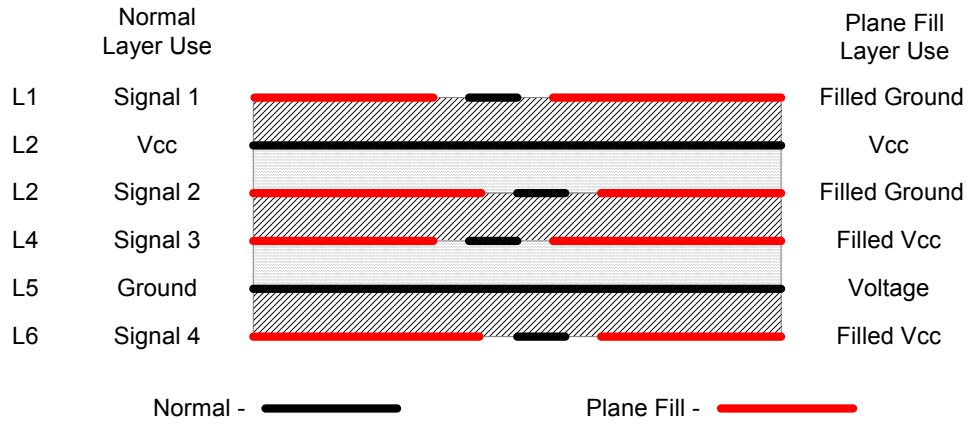

Figure 37.11. Stackup of PCMCIA Card With Plane Fill

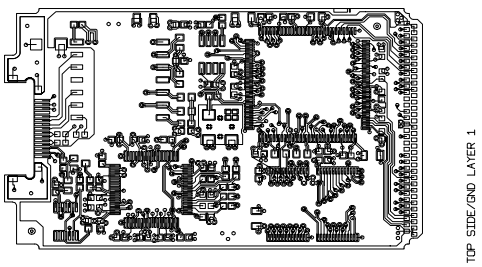

LAYER 1, SIGNAL FILLED WITH GROUND

LAYER 2, Vcc LAYER

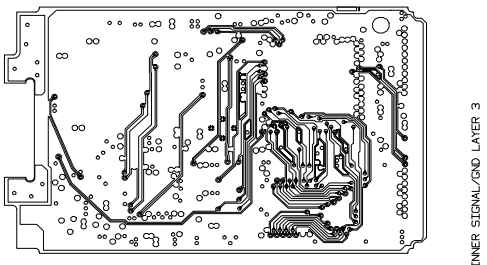

LAYER 3, SIGNAL FILLED WITH GROUND

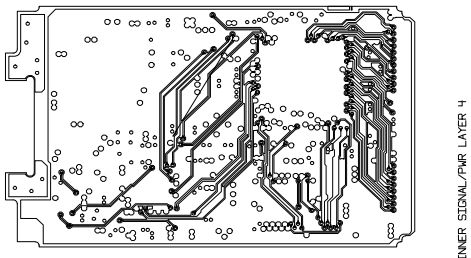

LAYER 4, SIGNAL FILLED WITH Vcc

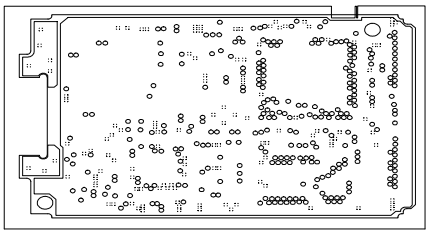

LAYER 5, GROUND LAYER

LAYER 6, SIGNAL FILLED WITH Vcc

Power plane capacitance without fill, 500pF. With fill 4100pF.

Figure 37.12. 6-Layer PCMCIA PCB Showing Signal Layers Filled with Power Planes

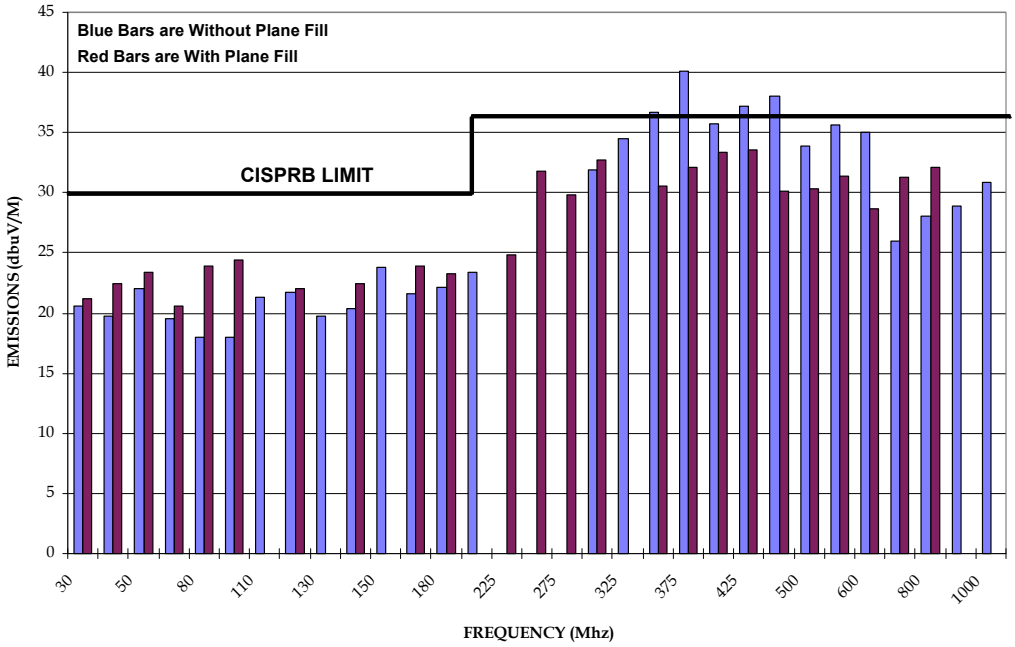

Figure 37.13. Emissions Test Results With And Without Signal Plane Fills

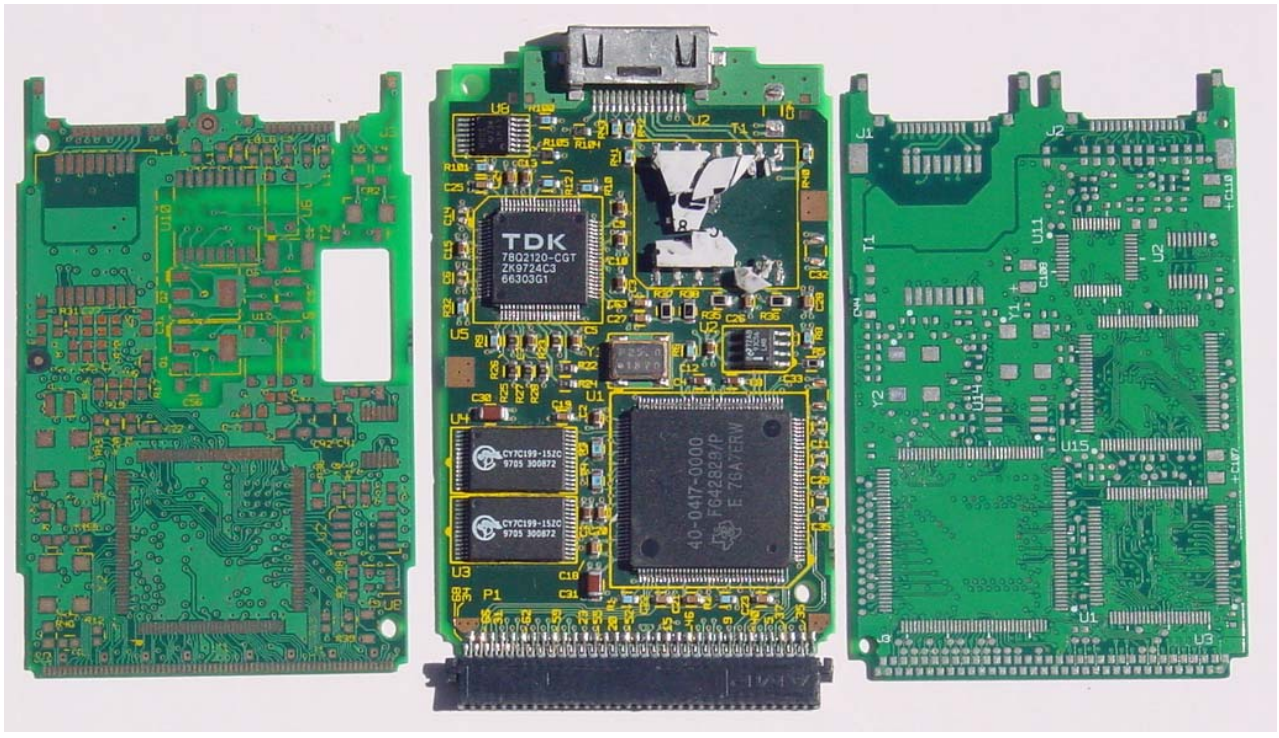

Figure 37.14. The Actual PCBs from Figures 37.10 and 37.11 Showing Plane Fill

Steps in The Power System Design Process

Now that the needs of the power subsystem have been identified, it is possible to outline the steps involved in completing the design. The steps that should take place are as follows:

1. Do a Fourier transform of time domain signals to determine frequency content.
2. Count the number of edges that can switch simultaneously. This will allow calculation of maximum delta I needed to support switching from 0 to 1.
3. Analyse Cpp vs Cswl to decide size of plane capacitor needed to support worst case switching while remaining within ripple budget.
4. Calculate size of plane capacitor, Cpp.
5. Calculate size and quantity of discrete capacitors needed to support the “low frequency,” (below 150 MHz) demand.
6. Model combination of discrete capacitors and plane capacitance to insure there are no impedance holes
7. Build actual PCB and sweep Z vs. F and measure ripple under worst-case conditions.

Verifying the Power Subsystem Has Proper Bypassing

Once the power subsystem design has been completed, it is imperative that it be checked to verify that the objectives have been met. Those objectives include ensuring that the impedance is low at all of the frequencies of interest and that the worst-case ripple is within specifications. This is accomplished with two tests—the sweeping impedance vs. frequency and actual worst-case load tests.

Sweeping impedance vs, frequency:

This test was described in Chapter 34. It uses a spectrum analyzer to create a plot of impedance vs. frequency. Two tests need to be made. The first involves sweeping the impedance vs. frequency of the bare PCB. As part of this test, the plane capacitance of each supply voltage is measured using an ordinary capacitance meter to verify that the target plane capacitance has been met. Next, the PCB is loaded only with the decoupling capacitors and the test is repeated. It is advisable to order two prototype PCBs that will be used only for the purposes of this test.

Worst-Case load testing:

This set of tests is performed on the assembled PCB. A code set is written that causes the data buses to switch from 0 to 1 and the ripple is measured on Vcc or Vdd. In addition, a code set is written that causes any processors to go from standby to active to standby and the ripple is measured on Vcc or Vdd. Figure 37.8 shows worst case ripple caused by a 256 bit wide bus switching from 0 to 1.

Figure 37.15 shows the test results obtained by sweeping impedance vs. frequency for a bare PCB on the left and the PCB with all bypass capacitors loaded on the right.

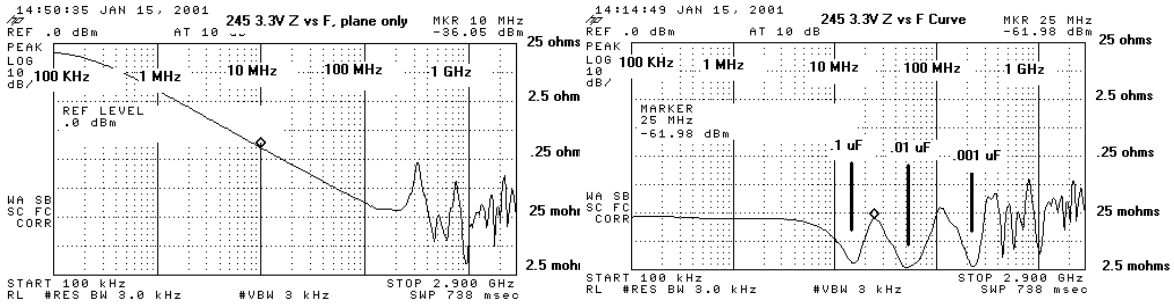

Impedance vs. Frequency, Power Place Capacitance only

Impedance vs. Frequency, Power Plane Capacitance & Decoupling Capacitors

8” x 8” PCB Capacitance information:

20 15 uF, 10 .1 uF, 50 .01 uF, 60 .001 uF,

50 470 pF, 50 100 pF, plane capacitor 44 nF

Note: 100 pF and 470 pF Capacitors have little effect

Figure 37.15 Power Supply Impedance vs. Frequency for an OC-48 Line Card

CHAPTER 38: IC PACKAGES--Vcc AND GROUND BOUNCE OR SSN

This chapter covers the effect that package parasitic inductance has on the performance of high-speed logic circuits. Specifically, the package parasitic involved in this case is the unwanted inductance in the power paths into and out of an IC package. As logic speeds have increased and data and address buses have become increasingly wider, the noise spikes created by the current transients involved in these switching events have become a major source of failure.

Vcc and ground bounce are the shift in the Vcc and ground rails on the IC die with respect to their respective levels on the planes of the PCB. This type of unwanted transient is most often the result of single-ended logic drivers charging and discharging transmission lines. Figure 38.1 illustrates the way in which Vcc and ground bounce are created.

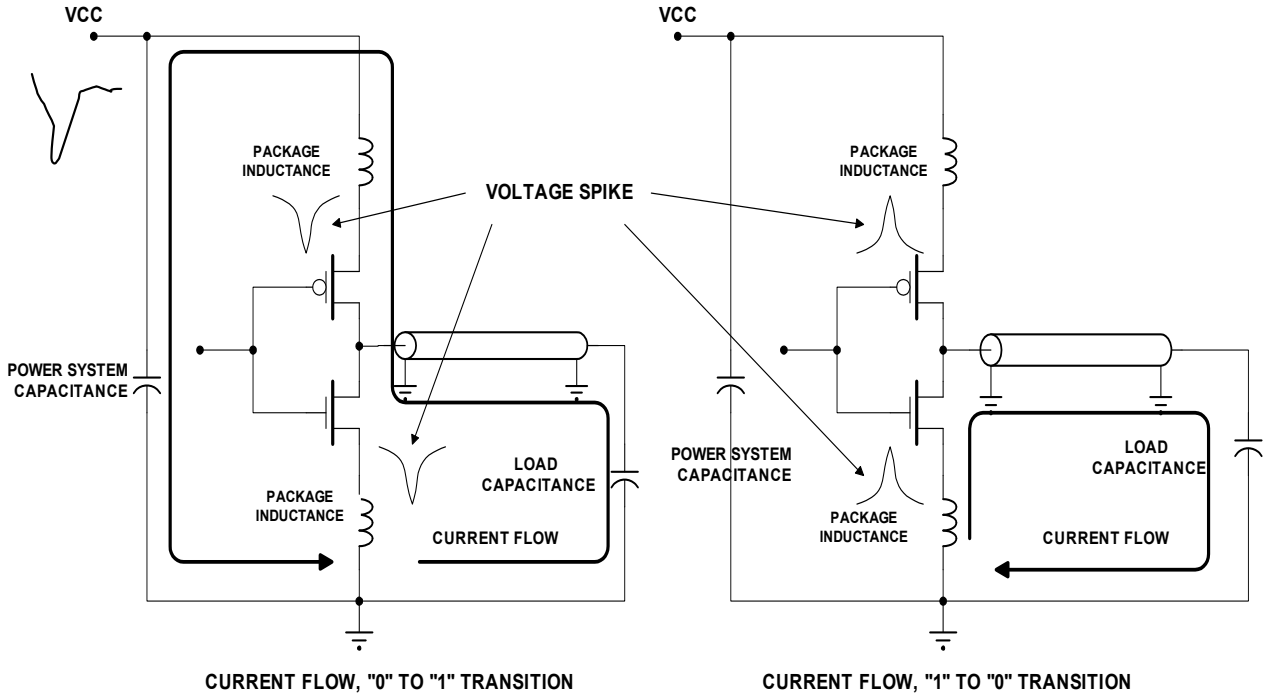

Voltage spikes are developed across the package inductances and are seen as Vcc and Ground bounce. Quiet outputs will move with the same waveform as the Vcc or Ground bounce. Note that both power terminals of the semiconductor die move in either case.

Figure 38.1. A Typical Single Ended Transmission Line Showing Vcc and Ground Bounce

The left side of Figure 38.1 shows the current path for the current required to charge up the parasitic capacitance of the transmission line and the parasitic capacitance of the loads as the transmission line is switched from a logic 0 to a 1. The right hand side of Figure 38.1 shows the current path associated with discharging the parasitic capacitance as the logic line switches from logic 1 to a 0. These transient currents are the primary source of simultaneous switching noise (SSN). The inductance shown in the diagram includes the inductance of the vias that connect the IC power leads to the power planes.

Notice that the Vcc terminal of the IC die is driven negative with respect to Vcc on the PCB power plane during a logic 0 to 1 transition. All of the terminals of the IC are driven negative at the same time (this is Vcc bounce). The effect is that this voltage spike appears on all quiet outputs and inputs. If the spike is large enough, it can cause a logic failure. During a transition from logic 1 to 0, the ground rail of the IC is driven positive with respect to ground on the PCB power plane (this is ground bounce). This spike appears on all lines as well and can cause logic failures.

Equation 38.1 can be used to calculate the magnitude of the voltage transient that results from a change in logic states.

$$V_L = L\frac{di}{dt}$$

Where: V_L = the voltage drop across the inductor, L is the inductance of the inductor in Henrys, di = the magnitude of the change in current in amps, dt = the time required to make the current change.

Equation 38.1. An Equation for Calculating the Voltage Drop Across an Inductor

Notice that voltage drops occur across inductors only when the current through them is changing or, more precisely, when the electromagnetic field traveling through them is changing.

Table 38.1 lists the lead inductances of some typical IC packages. The reason for the wide spread in inductance is the fact that the leads in most packages are of widely different lengths.

14 pin DIP	**3.2 - 10.2 nH**	**14 pin SOIC**	**2.6 - 3.6 nH**
20 pin DIP	**3.4 - 13.7 nH**	**20 pin SOIC**	**4.9 - 8.5 nH**
40 pin DIP	**4.4 - 21.7 nH**		
40 pin TAB	**1.2 - 2.5 nH**	**208 pin QFP**	**5.31 - 8.74 nH**
44 pin QFP	**6.07 - 7.06 nH**	**100 pin QFP**	**6.69 - 7.96 nH**
20 pin PLCC	**3.5 - 6.3 nH**	**119 pin PBGA**	**.15 - 5.7 nH**
28 pin PLCC	**3.7 - 7.8 nH**	**249 pin PBGA**	**.13 - 5.1 nH**
44 pin PLCC	**4.3 - 6.1 nH**	**624 pin CBGA**	**.5 - 4.75 nH**
68 pin PLCC	**5.3 - 8.9 nH**	**456 pin PBGA**	**.2 - 5.8 nH**

Table 38.1. Typical Lead Inductances of a Variety of IC Packages

To get some sense for the magnitude of Vcc and ground bounce that can occur in a common IC package, a simple calculation can be done. As an example, the 20 pin DIP (dual in line package) will be used. The power pins on this package are on the corners and have an inductance of 13.7 nanoHenries per power pin. The delta I in this case is when the logic state changes from 0 to 1 is 50 mA for a single output and the delta time is 2 nSEC for a 5V HCMOS part. Using these values in equation 38.1 results in a voltage spike of 342 mV. Imagine what happens when all 8 bits of a bus are changing simultaneously from 0 to 1. The voltage spike is 2.74 Volts. Next, the rise time is changed to 1 nSEC. The spike is 5.48 Volts. It is switching transients such as this that drove the change from DIP to PLCC packages. Companies that attempted to increase logic speeds while remaining in the tried and true packages found that their products did not work properly. It should be noted that this problem is package related and cannot be fixed by actions taken on the PCB.

> Vcc and ground bounce (SSN) are caused by excessive inductance in the power paths of the IC packages. No actions on the PCB can be taken to fix this problem. Changing to an IC package with lower power lead inductance is necessary.

The above discussion also applies to the QFP packages that are in common use in the industry. The inductances are not as high as for DIP packages. However, the width of data buses is much wider than 8 bits. Failures from SSN manifest themselves as infrequent failures. The reason is the worst-case noise spike occurs when all of the data bits change simultaneously from one logic state to the other. This happens only once in 2^N times, where N is the number of data bits in the bus. **Many of the "flaky" systems currently in manufacture or design are failing from this cause.** Failure to understand that this phenomenon is at work leads to designs that are never stable. Because of this potential failure mechanism, it is imperative that each part being considered as a driver of a single-ended logic bus be checked out prior to being used in order to insure SSN does not cause failures. Figure 38.2 illustrates a test setup used to measure worst-case SSN.

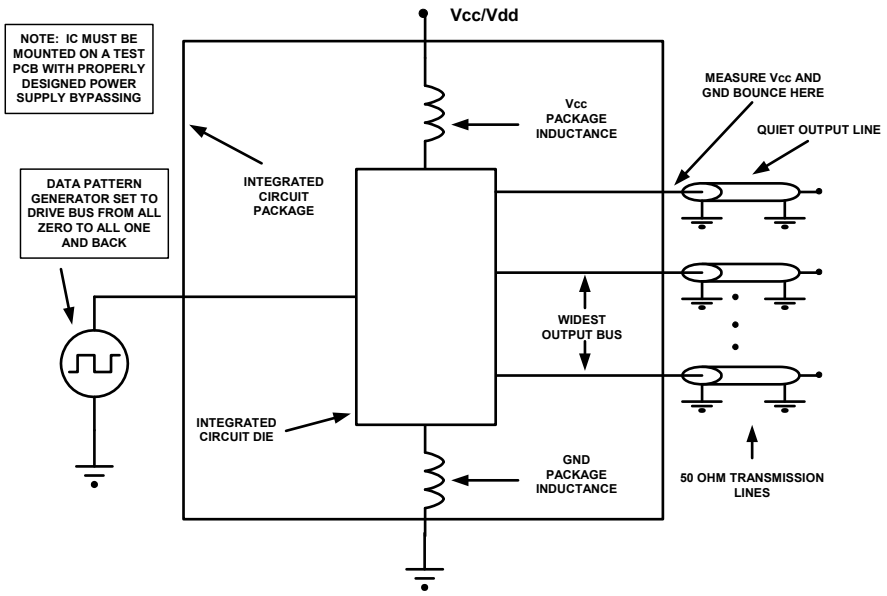

Figure 38.2. A Test Setup For Measuring Worst-Case Vcc and Ground Bounce

In order to measure worst-case Vcc and ground bounce, it is necessary to load all of the outputs of the widest bus with a realistic set of loads that can switch simultaneously. It is common to test IC outputs by attaching a large capacitor, for example 60 pF, to each output and observe the rise and fall time. This is not a realistic load. It "overloads" the output and results in a rise or fall time that is slower than will actually occur when the output drives a transmission line (see Figure 48.1). In addition, the current spike needed to charge up or discharge the capacitor is much larger than will occur in actual use. A realistic load is a 50 ohm transmission line as this is what the part will usually be expected to drive.

In Figure 38.2, all of the outputs are "loaded" with 50 ohm transmission lines. The measurement probe is connected to an output that is powered from the same internal power and ground bus as the data lines that will be switched. The IC is mounted on a PCB with a power plane structure that is capable of supplying the charging currents needed to switch all of the data lines from 0 to 1 simultaneously without drooping.

To observe Vcc bounce, the IC is driven with a signal pattern that will cause all of the data lines to switch from 0 to 1 simultaneously. While this is happening, the Vcc bounce associated with this event will appear on the quiet line with little or no attenuation. The current waveform that results was shown in Figure 21.2. The Vcc bounce waveform results from the rising edge of this current waveform.

To observe ground bounce, the IC is driven with a signal pattern that will cause all of the data lines to switch from 1 to 0 simultaneously. While this is happening, the ground bounce associated with this event will appear on the quiet line with no attenuation. The current waveform for this event is the same shape as that shown in Figure 21.2, but instead of flowing out of the power supply, it is discharged into ground.

The waveforms in Figure 38.3 are measured in this manner on an actual IC. In this case, the data bus being switched is 64 bits wide and the Vcc is 2.5 Volts. The top trace is a combination of Vcc and ground bounce. The excursions above Vcc are ground bounce and the excursions below Vcc are Vcc bounce. This set of waveforms was produced by toggling all of the data lines from 0 to 1 and then 1 to 0 many times. Notice that the excursions are about the same amplitude, 500 mV. From this it can be inferred that the inductance in both the Vcc and ground paths are about the same size. It is possible to get a rough idea how big this inductance is. This is accomplished by using equation 38.1

The peak current per output is 2.5V/100 ohms or 25 mA. The total current is 64 x 25 mA or 1.6A. The resulting voltage is 0.5 V. The delta time is 2 nSEC. The inductance is calculated to be approximately .625 nH. This inductance is a combination of the inductance of the IC package leads and the vias reaching down into the PCB to access the power and ground planes.

From earlier power supply discussions, it has been established that the current needed to perform this switching function is drawn from the plane capacitance built into the PCB. If this capacitance is not large enough, there will be ripple on Vcc every time the data lines switch from 0 to 1. This is shown in the upper left hand corner of Figure 38.1 as a "v" shaped dip.

The lower waveform in Figure 38.3 is the ripple on Vcc that corresponds to these switching events. In this example, the plane capacitor that was designed into the PCB to support this event was 24 nF. The resulting ripple is approximately 150 mV. For 2.5V logic, this is near the limit of what is acceptable. The only way to reduce this ripple is to redesign the PCB stackup to add more plane capacitance. Discrete capacitors have too much inductance to solve this problem.

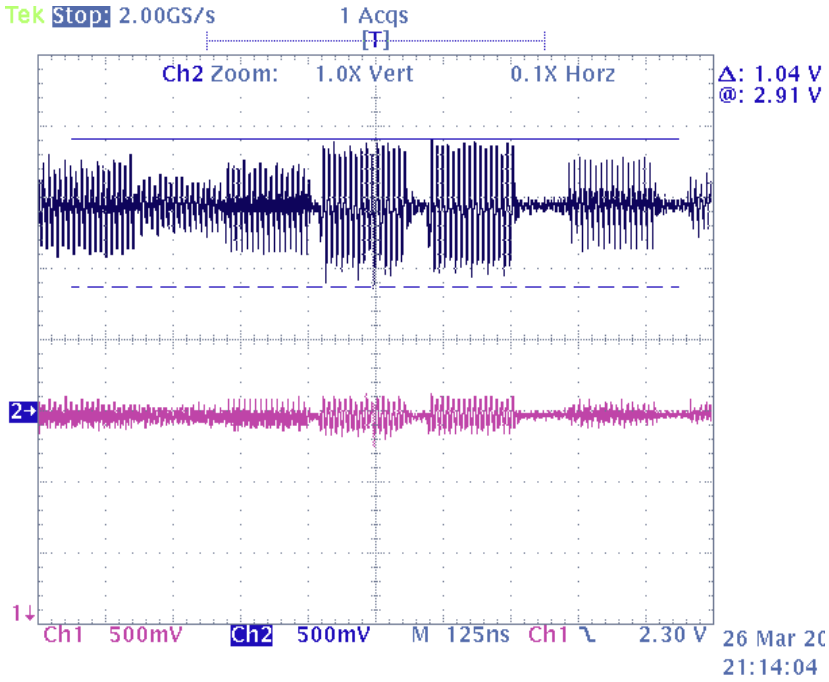

Figure 38.3. Actual Vcc and Ground Bounce on a 64 Bit Data Bus with Vcc Noise Shown

In addition to designing the PCB stackup to produce enough plane capacitance to supply the switching transients, it is important to minimize the inductance of the vias required to reach the planes. **Therefore, the plane pair that supports the largest data buses should be the first plane pair below the components**.

The waveform in Figure 38.3 is the power on reset line for this particular design. Each time all members of this data bus switched from 0 to 1, the system went through a power on reset cycle. This would happen during a memory access. The result was a system that could not be shipped and could not be fixed by taking any action on the PCB. The only remedy was to redesign the IC package. Such designs are often called Silicon Valley Tombstones.

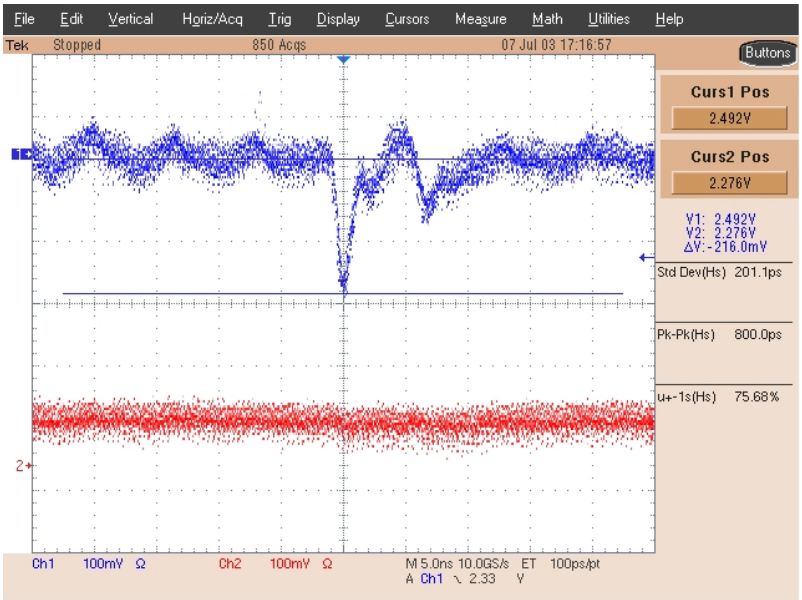

Figure 38.4. Actual Vcc Bounce on a 80 Bit Data Bus with Vcc Noise Shown

Figure 38.4 is the Vcc bounce waveform from another IC with an 80-bit data bus switching from 0 to 1 simultaneously. Also shown is the noise on Vcc. It should be noted that the Vcc noise related to this switching event is so small that it is difficult to see. This is because the plane capacitor that supports this switching event is 140 nF or seven times that in Figure 38.3.

The Vcc bounce in Figure 38.4 is from a 2.5V DDR data bus. The effective inductance that caused this noise can be calculated using Equation 38.1. In this case, the delta I per line is again 25 mA. The total delta I is 2A. The delta V is 216 mV. The delta T is 1.15 nSEC. From this, the equivalent L is 0.497 nH. Again, this is a combination of the package inductance and the inductance of the vias that connect to the power planes.

The example in Figure 38.4 has a relatively slow rise time of 1.16 nSEC. The component in this example is capable of producing rise times of less than 0.5 nSEC. Should one of these parts on the fast side of the performance spectrum be mounted in the same circuit, the Vcc bounce would be twice that shown, or 532 mV. This would be far in excess of the noise tolerance of this circuit. Redesigning the package to reduce the parasitic inductance is the surest way to fix this problem.

If redesigning the IC package is not a choice is there any other possible solution? In some cases there is. The source of the noise problem is many outputs switching simultaneously. Sometimes it is possible to stagger the clocking of the outputs such that only a portion can switch at the same time. This is often called a multiphase clock. Yet another choice might be to distribute the outputs on several Vcc and ground rails. This is possible with some of the FPGAs on the market.

One of the parameters in this equation is the rise time of the signal edge. It might be possible to slow down the edges.

Of course, the best solution is to select components that have packages with very low inductances in the power and ground rails. Table 38.2 shows the amount of Vcc or ground bounce that a variety of 2.5V wide data buses will create for 0.5 nSEC edges as a function of the total inductance in the power path.

INDUCTANCE	16 BITS	32 BITS	64 BITS	128 BITS	256 BITS
0.1	80	160	320	640	1280
0.2	160	320	640	1280	2560
0.3	240	480	960	1920	3840
0.4	320	640	1280	2560	5120
0.5	400	800	1600	3200	6400
PEAK I (A)	0.4	0.8	1.6	3.2	6.4
Vcc or ground bounce in mV. Inductance in nH. RISE TIME = 0.5 nSEC, 2.5 V LOGIC					

Table 38.2. Vcc and Ground Bounce vs. Package Inductance, 2.5V CMOS, 0.5 nSEC Edge

As edges get faster and data buses become wider, the design of the IC package must be done with great care if failures from Vcc and ground bounce are to be avoided.

> Every IC that is intended to drive a wide data bus must be checked to insure the package inductance is low enough to produce acceptably low Vcc and ground bounce.
>
> It is unwise to include a part in a design for which the vendor is not able to demonstrate acceptable Vcc and ground bounce numbers.

Vcc AND GROUND BOUNCE HAVE BECOME THE PRIMARY SOURCE OF INTERMITTENT FAILURES IN MOST NEW DESIGNS. THESE FAILURES ARE TRACEABLE TO POOR IC PACKAGE DESIGN.

How to Determine Vcc and Ground Bounce for New ICs Prior to Their Being Packaged

Often, it is necessary to choose a part while it is still in development. There isn't a real part to measure. It is possible to model the I/O and package in a SPICE modeling software package and calculate the anticipated Vcc and Ground bounce.

This is done by obtaining a SPICE model of the output driver, combining it with the predicted inductance of the package power and ground paths, driving transmission lines of the appropriate impedance and switching the outputs at the fastest rise and fall times expected in the final circuit.

Examples of Poorly Designed BGA Packages

In general, BGA packages have better Vcc and ground inductances than do other package types, such as QFP, PLCC and DIP. However, this is not always so. Some BGA packages are designed such that there are leads or wires from the power and ground contacts out to the balls that contact the PCB. When a new component in a BGA package is being considered, it is imperative that the package layout be examined to insure both Vcc and ground paths have low and equal inductances. Figure 38.5 is a photograph of two two-layer BGA packages with this problem. In both cases the die is mounted in the cavity over contacts that pass directly through the package to balls on the backside of the BGA. These are the ground contacts and are very low inductance. In both cases, the Vcc leads are traces that travel from the edge of the die to the edge of the BGA package as traces.

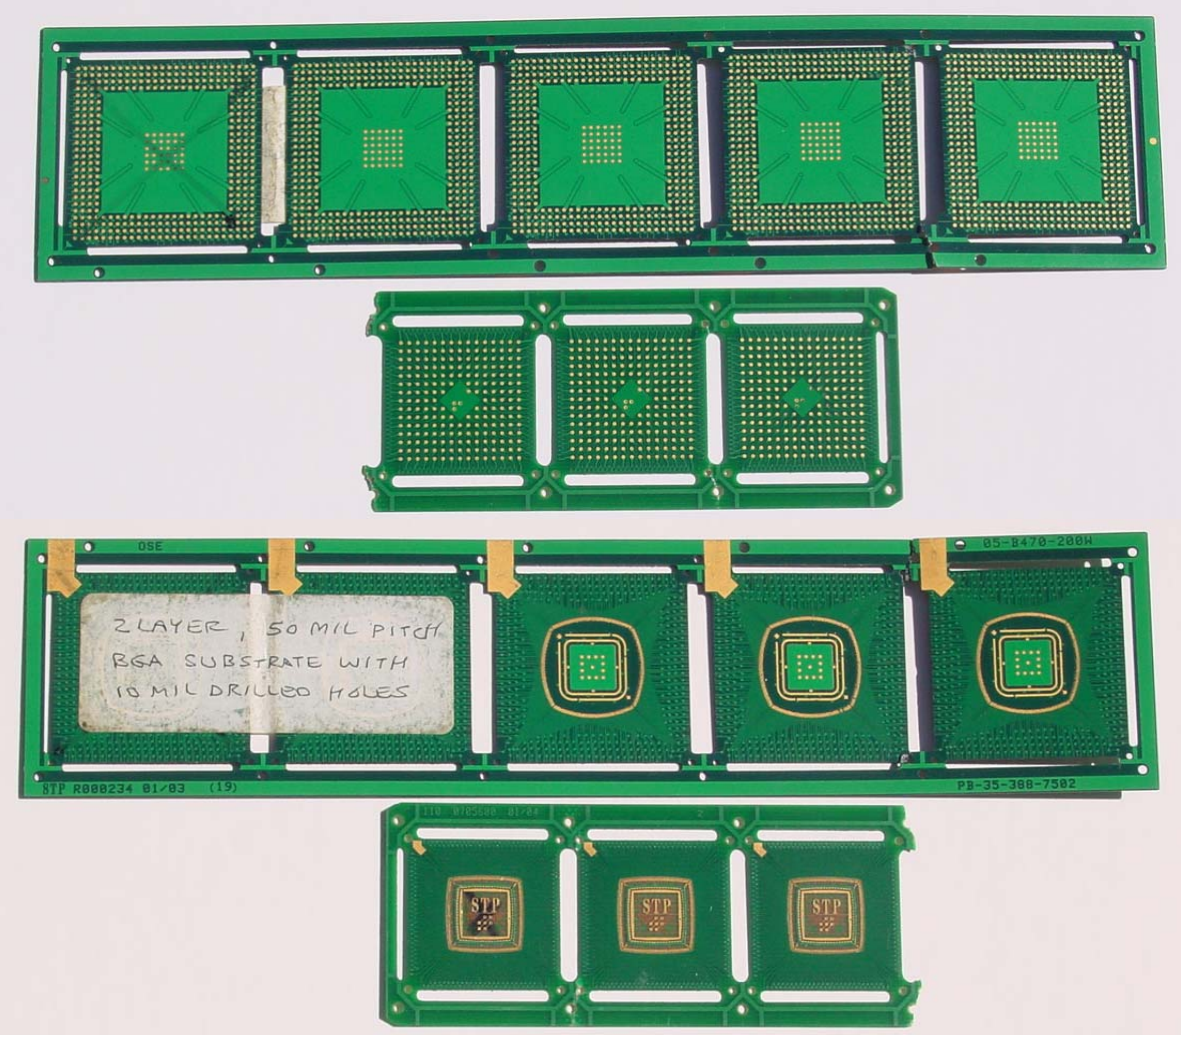

Figure 38.5. Two BGA Packages That Have Excessively High Inductance in the Vcc Leads

> With CMOS logic, it is imperative that the parasitic inductance in both the Vcc and the ground paths of an IC package be both small and equal.

Determining What Is an Acceptable Amount of Vcc and Ground Bounce

The amount of Vcc and ground bounce that can be tolerated is calculated by analyzing the noise from other sources and comparing it to the noise margin for the logic family being used. A portion of the total noise margin can be consumed by Vcc and ground bounce. The methodology used to perform this analysis is presented in the following chapters.

CHAPTER 39: NOISE MARGINS

Noise margins are a part of the characteristics of every logic family. They are the difference between the smallest signal that is created by a driver and the smallest signal required by an input or receiver to successfully detect the proper logic level. This difference is designed into each logic family to allow for noise that will degrade the signal as it travels from driver to receiver. Figure 39.1 illustrates the input and output characteristics of an ECL logic family.

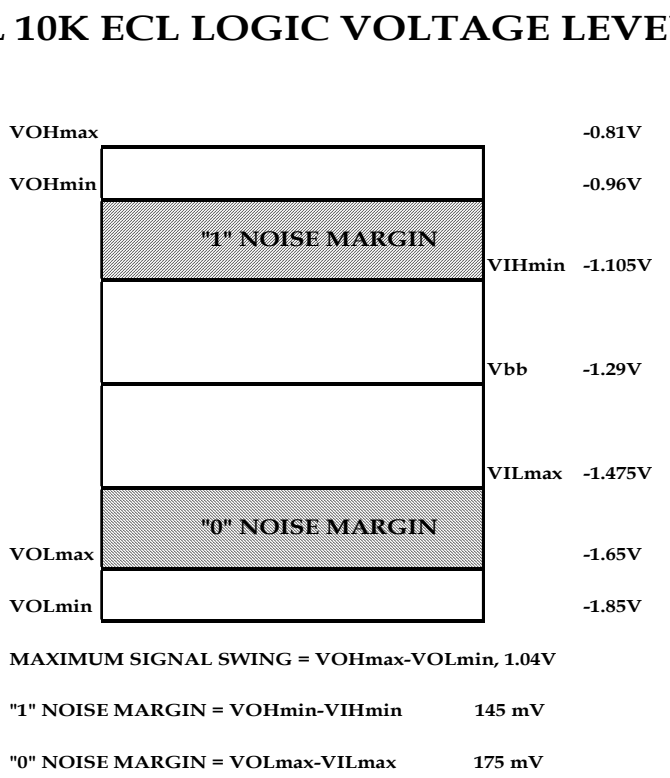

Figure 39.1. Noise Margin Diagram for an ECL Logic Family

The left hand side of this diagram shows the characteristic of the drivers in this logic family. The right hand side shows the characteristics of the receivers. This data is obtained from the manufacturer's data sheet and is the starting point for creating design rules.

The meaning of this data is:

VOH_{MAX}- This is the highest logic 1 voltage that will ever be generated by the driver. It occurs when the best part that comes off the IC manufacturing line and meets the manufacturer's specifications is used under ideal conditions.

VOL_{MIN}- This is the lowest logic 0 voltage that will ever be generated by the driver. It occurs when the best part that comes off the IC manufacturing line and meets manufacturer's specifications is used under ideal conditions.

The difference between VOH_{MAX} and VOL_{MIN} is the largest signal that will ever be generated by the driver. This is the signal that will create worst-case crosstalk and worst-case reflections. This signal amplitude will be used to calculate the noise generated in each of these two cases.

VOH_{MIN}- - This is the lowest logic 1 voltage level that will ever be generated by an output driver. It occurs when the worst part that comes off the IC manufacturing line and meets manufacturer's specifications is used under worst-case rated load conditions.

VOL_{MAX}- This is the highest logic 0 voltage level that will ever be generated by an output driver. It occurs when the worst part that comes off the IC manufacturing line and meets manufacturer's specifications is used under worst-case rated load conditions.

The difference between VOL_{MIN} and VOL_{MAX} is the smallest signal that will ever be generated by an output driver. The design objective when creating design rules is to select a set of design rules that insures that the erosion of this signal does not diminish it below the levels required by the inputs.

VIH_{MIN}- This is the smallest logic 1 at the input to a receiver which will guarantee that a logic 1 will always be detected.

VIL_{MAX}- This is the largest logic 0 at the input to a receiver which will guarantee that a logic 0 will always be detected.

Any voltage between VIH_{MIN} and VIL_{MAX} can result in an indeterminate logic state.

This set of voltage values defines the signal integrity engineering problem. The process of generating design rules involves calculating all of the sources of noise voltages that will result from that set of design rules and guaranteeing that when they are all summed together, which may occur on rare occasions, they do not exceed the noise margin of each logic family in the design. (This analysis must be performed for each unique logic family used in a design.)

> Design rule creation involves analyzing the amount of noise that each rule may generate and guaranteeing that when all the noises are summed together, they do not exceed the noise budget of each logic family being used.

Noise margin analysis and design rule generation is an iterative process that involves examining a proposed rule set to see how much noise each rule will generate. These noise sources are summed up and compared to the noise margin of each logic family. If the noise margins have been exceeded, the rules are examined to determine which rules can be most economically changed to reduce noise until the noise budget has been met. This process will be illustrated in Chapter 40.

Table 39.1 lists the ten possible sources of noise that can be generated by a design. As will be illustrated in Chapter 40 some logic families are not affected by all of these sources.

Reflections	**Cross Talk**
Ground Bounce	**Ground IR Drop**
Ground Offsets	**Reference Accuracy**
Thermal Offsets	**Terminator Noise**
Trace IR Drop	**Power Supply Variations**

Note: IR drop is voltage drop. I x R = V

Table 39.1. Potential Noise Sources in High Speed Logic Systems

Why sum up all of the noise sources?

The reason that noise margin analysis involves summing up all of the possible noise sources rather than taking some average or RMS value is that, on rare occasions, noise voltages will coincide forming a large noise spike that causes a logic failure. It is this kind of coincident noise that produces most of the systems that are described as “flaky.” For example, when a Vcc or ground bounce voltage spike coincides with a reflection or a crosstalk spike, a failure may occur. Remembering that the worst case Vcc or ground bounce spike is rare in itself, occurring only once in about 2^N transactions, an occurrence of this kind along with a worst-case reflection or crosstalk may be quite rare, but imminently possible. It is this potential for occasional coincident noise voltage spikes that is the compelling reason for doing noise margin analysis as a method for insuring that a design rule set is robust.

> The potential that noise voltages from multiple sources may coincide is the compelling reason that noise margin analysis of a set of proposed design rules must be done with care and adjustments made until the noise budget balances.

Table 39.2 lists the characteristics of several commonly used logic families. With the exception of the timing information, this table provides all of the information needed to create a complete set of design rules. The timing information must be obtained on a component-by-component basis and supplied to the timing analysis tool set.

This data was taken from manufacturers’ data sheets and is as accurate as those data sheets. The reader is cautioned to use this data as a guide only and to obtain actual data sheets from each manufacturer in order to create the noise margin

data for each new design. Figure 39.2 is an illustration of the input and output characteristics of several common logic families. It is courtesy of Texas Instruments.

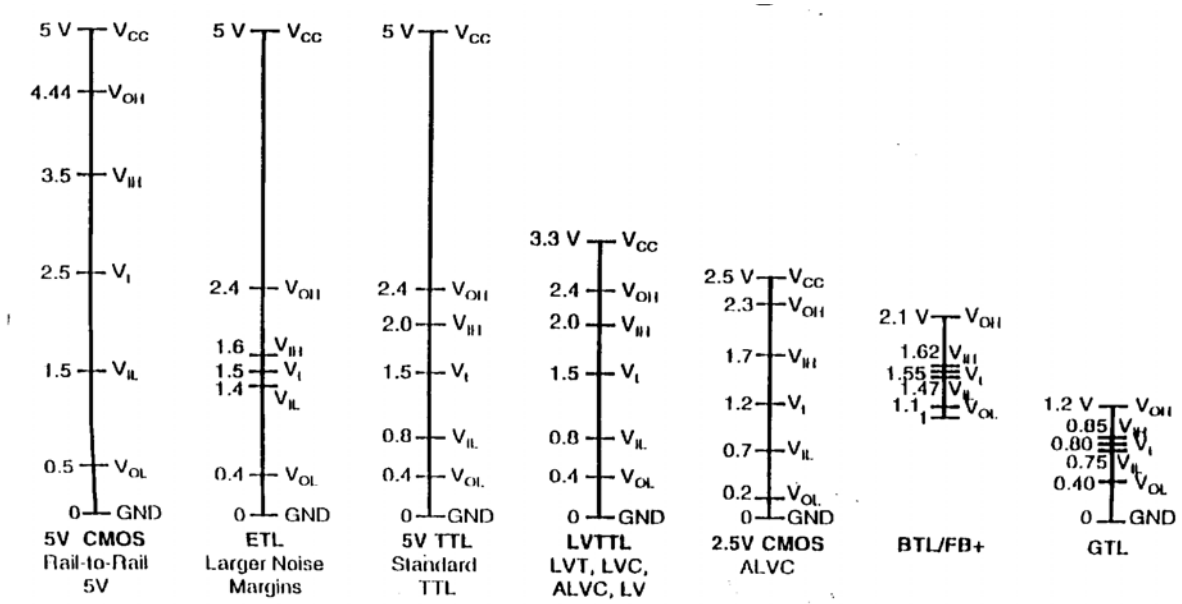

Figure 39.2. Input and Output Characteristics of Several Logic Families

Shown for each logic family is the minimum signal launched by a driver specified as V_{OH} and V_{OL} as well as the minimum input signal levels required by inputs for proper operation specified as V_{IH} and V_{IL}. The difference between these two is the noise margin available to be "spent" on noise sources. The closer these two levels are to each other, as with TTL, the less tolerant of noise a logic family is. The farther apart these two levels are, as with GTL or BTL, the more tolerant of noise a logic family is. All things being equal, choosing a logic family with a large spread will make the design rule creation process easier.

LOGIC FAMILY	Power Supply Voltage	LOGIC TYPE	Vohmax	Volmin	Maximum Output Signal	Vohmin	Volmax	Minimum Output Signal	Vihmin	Vilmax	Noise Margin	Termination Voltage	Reference Voltage	Zout	Fastest Rise Time	Output Current	Maximum Clock Frequency
TTL	+5V	BIPOLAR	3.3V	0.2V	3.1V	+2.4V	+0.4V	2.0V	+2.0V	+0.8V	400 mV	N/A	+1.4V	N/A	5 Nsec	N/A	20 MHz
LVTTL	+3.3V	BIPOLAR	+3.3V	+0.2V	3.1V	+2.4V	+0.4V	2.0V	+2.0V	+0.8V	400 mV	N/A	+1.4V	90/36 ohms	N/A	N/A	
HSTL-I	+3.3V	CMOS				+2.9V	+0.4V	2.5V	+1.75V	+1.55V	1.15V	N/A	+1.65V	20 ohms	1.5 nSec	8 mA	
HSTL-II	+3.3V	CMOS				+2.9V	+0.4V	2.5V	+1.75V	+1.55V	1.15V	N/A	+1.65V	11 ohms	1.3 nSec	16 mA	
HSTL-III	+3.3V	CMOS				+2.9V	+0.4V	2.5V	+1.75V	+1.55V	1.15V	N/A	+1.65V	9 ohms	1.0 nSec	24 mA	
HSTL-IV	+3.3V	CMOS				+2.9V	+0.4V	2.5V	+1.75V	+1.55V	1.15V	N/A	+1.65V	7.7 ohms	0.6 nSec	48 mA	
10K ECL @25C	-5.2V	BIPOLAR	-0.7V	-1.99V	1.299V	-0.96V	-1.65V	0.45V	-1.105V	-1.475V	175 mV	-2.0V	-1.29V	4 ohms	0.3 nSec	N/A	
PECL	+5V	BIPOLAR	4.19V	3.05V	1.14V	4.02V	3.37V	0.65V	3.87V	3.52V	150 mV	+3.0V	3.62V	4 ohms	0.3 nSec	N/A	
LVPECL	+3.3V	CMOS	+2.28V	+1.06V	1.22V	+1.92V	+1.43V	0.49V	+1.82V	+1.52V	90 mV	+1.3V	+1.67V	4 ohms	0.2 nSec	N/A	
GTL	+3.3V	CMOS	Vt	+0.2V	1.08V	Vt-.01V	+0.4V	0.05V	+0.65V	+0.55V	150 mV	+1.2V	+0.6V	16 ohms	0.2 nSec	N/A	125 MHz
GTLP	+3.3V	CMOS	Vt	+0.2V	1.08V	Vt-.01V	+0.6V	0.50V	+0.70V	+0.50V	100 mV	+1.2V	+0.6V	16 ohms	0.2 nSec	N/A	175 MHz
CMOS	+5.0V	CMOS	5.0V	0V	5.0V	4.4V	0.5V	3.9V	2.0V	0.8V	300mV	N/A	N/A	N/A	2.0 nSec	N/A	
CMOS	+3.3V	CMOS	3.0 V	0.2V	2.9V	2.0V	0.5V	1.5V	2.0V	0.8V	0 mV	N/A	N/A	N/A	0.5 nSec	N/A	
CMOS	+2.5V	CMOS	2.3V	0.2V	2.1V	1.8V	0.5V	1.3V	1.7V	0.7V	100nV	N/A	N/A	N/A	0.2 nSec	N/A	
CMOS	+1.8V	CMOS	1.8V	0V	1.8V	1.6V	0.2V	1.4V	1.17V	0.63V	430 mV	N/A	N/A	N/A	0.2 nSec	N/A	600 MHz
PCI-33	+3.3V	CMOS	+3.1V	+0.20V	2.09V	+2.97V	+0.33V	2.64V	+1.65V	+0.99V	660 mV	N/A	N/A	22 ohms		N/A	33 MHz
PCI-33	+5.0V	CMOS				+2.4V	+0.55V	1.85V	+2.0V	+0.8V	250 mV	N/A	N/A	22 ohms		N/A	33 MHz
PCI-66	+3.3V	CMOS	+3.1V	+0.20V	2.09V	+2.97V	+0.33V	2.64V	+1.65V	+0.99V	660 mV	N/A	N/A	22 ohms		N/A	66 MHz
PCI-66	+5.0V	CMOS				+2.4V	+0.55V	1.85V	+2.0V	+0.8V	250 mV	N/A	N/A	22 ohms		N/A	66 MHz
SSTL_2 class 1	+2.5V	CMOS	+2.3V	+0.2V	2.1V	+1.8V	+0.68V	1.14V	+1.60V	+0.90V	220 mV	+1.25V	+1.25V	96 ohms		7.6 mA	N/A
SSTL_2 class 2	+2.5V	CMOS	+2.3V	+0.2V	2.1V	+2.01V	+0.49V	1.52V	+1.60V	+0.90V	400 mV	+1.25V	+1.25V	32 ohms		15.2 mA	N/A
LVDS	+3.3V	CMOS	+1.6V	+0.9V	0.7V	+1.5V	+1.0V	0.50V	+1.35V	+1.15V	150 mV	+1.25V	+1.25V	High		4 mA	N/A
PCIExpress		CMOS DIFF	+1.200 V	0 V	1200 mV	+0.90 V	+0.10 V	800 mV	562 mV	438 mV	338 mV	500 mV	500 mV	N/A	60 pSEC	N/A	2.5 GB/S
RS-232	+-15V	BIPOLAR	+15V	-15V	30V	+5V	-5V	10V	+3V	-3V	2000 mV	0V	0V	N/A	0V/mS	N/A	20 KB/S
RS-422	+-15V	BIPOLAR	+6V	-6V	12V	+2V	-2V	4V	+200 mV	-200 mV	1800 mv	0V	0V	N/A	N/A	N/A	10 MB/S

Note: The above information was taken from selected manufacturer's data sheets. Before doing actual design calculations, the actual data sheet for the logic in use should be rechecked.

N/A means not available or not applicable.

Table 39.2. Characteristics of Commonly Available Logic Families

CHAPTER 40: DESIGN RULE CREATION USING NOISE MARGIN ANALYSIS

Creating a set of design rules for a high-speed logic system has three parts. These are: analyzing timing to insure that the design functions properly for all variations in component propagation delays and edge speeds; designing the power subsystem such that the voltage stability and ripple specifications of each logic family are met and creating a set of

transmission line rules that insure every signal arrives at its load or loads with the proper quality to guarantee stable operation under all conditions.

> Design rule creation has three parts. These are:
> - Timing Analysis
> - Power Subsystem Design
> - Transmission Line Rules

The design rule development process is not complex or difficult to do. Further, it can and should be done prior to launching the schematic design process. In this way, logic families that are not capable of functioning in the proposed design are discovered in time to replace them with more appropriate logic.

Timing Analysis

When designing the PCB, path length delays on the routed signal traces must be fed into the timing analysis tools to insure timing goals are met. This is first done at the pre-route stage after all of the components have been placed in such a manner that they satisfy thermal requirements, manufacturability goals and routability of the PCB itself. Since the PCB has not been routed at this point in the process, the path lengths are estimated by calculating the Manhattan distance between the points in each net. That information is then supplied to the timing analysis tool. (Manhattan distance is that length of wire or trace needed to connect any two points when traces must be routed only using X and Y routes. In dense PCBs, this X-Y discipline is required in order to fit the most wire in each layer and to insure that crosstalk does not occur between adjacent signal layers.)

Power Subsystem Design

The power portion of design rule creation was covered in detail in earlier chapters of this book. In the following exercise, it will be shown that the allowable ripple on the power supply voltages is often determined by how much of the noise budget can be consumed by ripple. In particular, CMOS logic outputs short each data line to Vcc or Vdd when the output is a logic 1. As a result, ripple on Vcc or Vdd appears on the logic lines with little or no attenuation.

The remainder of this chapter will cover the process of creating the transmission line rules that will be used to route the final PCB. Figure 40.1 illustrates the result of analyzing a set of design rules for an ECL-based super computer. This analysis was done after the design was complete rather than before it was started. The analysis was done to determine why the supercomputer was unable to perform in a stable manner. The analysis revealed the source of the instability.

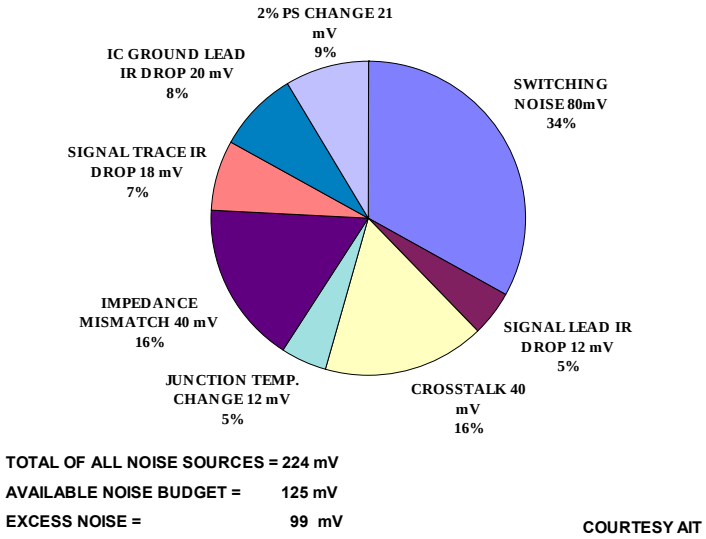

Figure 40.1 Noise Sources in an ECL Supercomputer Design

The pie chart in Figure 40.1 shows noise components from all ten possible noise sources. ECL is the only logic family that has noise from all ten sources. At the bottom of the figure is the sum of all noise sources and the available noise budget for this logic family. The excess noise is 99 millivolts. It is clear why this computer is unstable or flaky. There is too much noise. It is also clear what the primary problem is. Simultaneous Switching Noise (SSN) is the largest noise source in the system. This noise is caused by the ICs being packaged in a carrier or lead frame that has too much inductance in the Vcc

and ground paths. The only remedy for this problem is to redesign the system with the ICs in a package that has an acceptable Vcc and ground bounce characteristic (low parasitic inductance). This means starting the design over, a requirement that spelled the end of this product and the company that designed it.

> New designs that use QFP packages to drive single ended data buses have a high probability of failing due to excess Vcc and ground bounce.

The tragedy in this case is that this problem should never have occurred. If noise margin analysis had been done prior to launching the design, this flaw would have been detected; a better IC package would have been selected and the design could have been stable from the onset.

A brief review of each of the ten potential sources of noise is in order prior to developing a set of design rules. This review of these ten sources will proceed from those that are most likely to cause failures to those that might be considered "second order" in most designs.

Reflections

Chapters 16 to 19 covered the concept of reflections as sources of noise. Table 40.1 lists the places in a design where reflections can be created. It is these areas that must be controlled as part of the design rule set in order to keep reflections within the noise budget. Due to limits in the PCB manufacturing process, it is not commercially feasible to control the impedance of traces in a PCB to tolerances tighter than ±10%. As a result, allowances must be made for the noise that results from a 10% impedance change. The goal should be to keep all of the remaining sources of reflections within this limit.

Changes in trace width within the same routing layer
Stubs should be kept short enough to minimize reflections (Chapter 22)
Loads should be spread out along transmission lines
Connector transitions
Improperly matched termination resistors
Lack of terminations
Large power plane discontinuities (cleared out for some reason)
Changes in trace height above power planes

Table 40.1 Possible Sources of Impedance Discontinuities

When choosing the values of the termination resistor, it is desirable to choose their values such that any reflection that results from a mismatch between the impedance of the PCB transmission line and the termination results in a reflection that is overshoot rather than undershoot. (Remember, overshoot adds to the incoming signal, undershoot subtracts from it.)

This is accomplished with parallel terminations by sizing the resistor on the high side of the PCB trace impedance. For example, if the PCB impedance is 50 ohms ±10% the trace impedance could range between 45 and 55 ohms and still be within specifications. Choosing the termination value at 55 ohms guarantees either no reflection or overshoot which adds to the incident signal.

When using series terminations, the objective is to launch no less than a half amplitude signal that will double at the open end of the transmission line creating a full amplitude signal. Figure 40.2 illustrates this situation. The worst-case condition occurs when the PCB trace impedance is on the low side of its tolerance range. Under this condition, the combination of the output impedance of the driver and the series termination resistor should total 45 ohms, or the low side of the PCB trace impedance range. This way, when the PCB trace impedance is nominal or on the high side, the signal that is launched down the transmission line will be more than V/2 or half amplitude resulting in a full amplitude signal at the load at all times.

> Parallel termination values are selected on the high side of Z_0.
> Series terminations are selected on the low side of Z_0.

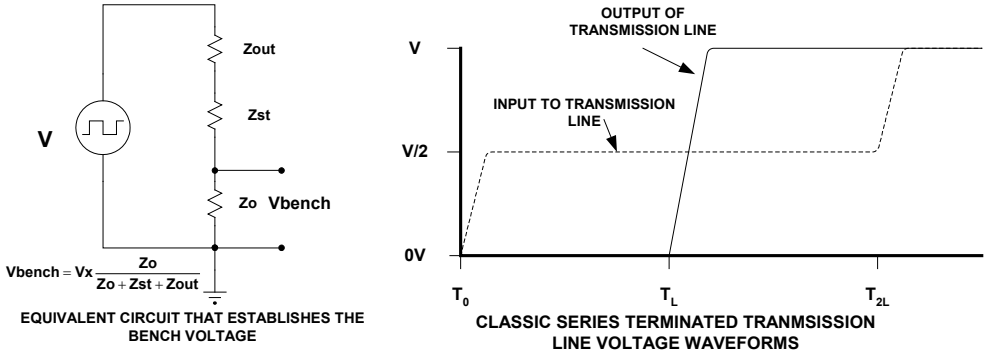

THE OBJECTIVE IS TO SELECT A VALUE FOR Zst THAT RESULTS IN A BENCH VOLTAGE THAT IS NEVER LESS THAN V/2. WORST CASE IS WHEN Zo IS ON THE LOW SIDE OF ITS TOLERANCE RANGE.

Figure 40.2. Equivalent Circuit of a Series Termination Line at To

Crosstalk

Crosstalk control was discussed in Chapter 29. The two most reliable methods for controlling crosstalk are trace-to-trace spacing within the same layer and routing signals at right angles to each other in adjacent signal layers. Traces running one over the top of the other in adjacent signal layers can result in excessive crosstalk, even with short runs.

Vcc and Ground Bounce

Chapter 39 dealt with Vcc and ground bounce in detail. This is primarily a problem related to the package parasitic inductance in the power paths of the IC package itself. Very little ground or Vcc bounce occurs in the planes of the PCB itself. The first choice is to select a package that has low inductance. If this is not a choice, it may be necessary to allot most or all of the noise budget to this noise source, leaving little or none for crosstalk and other noise sources.

Ground and Vcc Offsets

Ground and Vcc or Vdd offsets are the voltage drops that occur along the path from the power supply to the using circuits as a result of current flowing in the real conductors that make up this path. Those conductors include the planes of the PCB and the power connectors. To the extent that the ends of logic paths span these conductors and their associated voltage drops, the logic levels may be shifted such that noise margins are eroded. The chapters on power subsystem design explore this topic in detail. Figure 40.3 shows measurements of these drops in an actual PCB. This allows the reader to get some sense for the magnitude of these voltage drops in the planes of a PCB.

The measurements were made by using a milli-volt meter to measure the drop in voltage from where the power entered the PCB on the left edge to the power and ground pins of all the ICs on the PCB. In this example, the current drawn from the backplane connectors was a total of 25 amps. This current was consumed uniformly from the front to the back of the PCB. As can be seen from the figure, the total voltage drop in the Vcc plane was approximately 15 mV and the voltage rise in the ground plane was approximately the same.

Determining whether or not this is a reasonable voltage drop depends on the logic family involved and the noise margin it has. In this case, the logic family is TTL with 400 millivolts of noise margin that can tolerate 15 millivolts of offset. In cases where there is not as much tolerance for noise (1.5V CMOS logic being such a logic family), techniques must be employed that reduce this drop. An obvious choice is adding more metal in the form of multiple planes or creating thicker planes.

There will also be voltage drops in the connector pins used to bring power onto a PCB. It is imperative that the DC voltage drop that will occur at this point be calculated in advance of designing the system in order to make sure enough power pins are used. **When the connector pin voltage drop calculations are done it is important to use the end of life resistance of the connector contacts.** Failing to use the end of life resistance can result in systems that work properly when they are new and fail as time goes by and the connector contacts degrade.

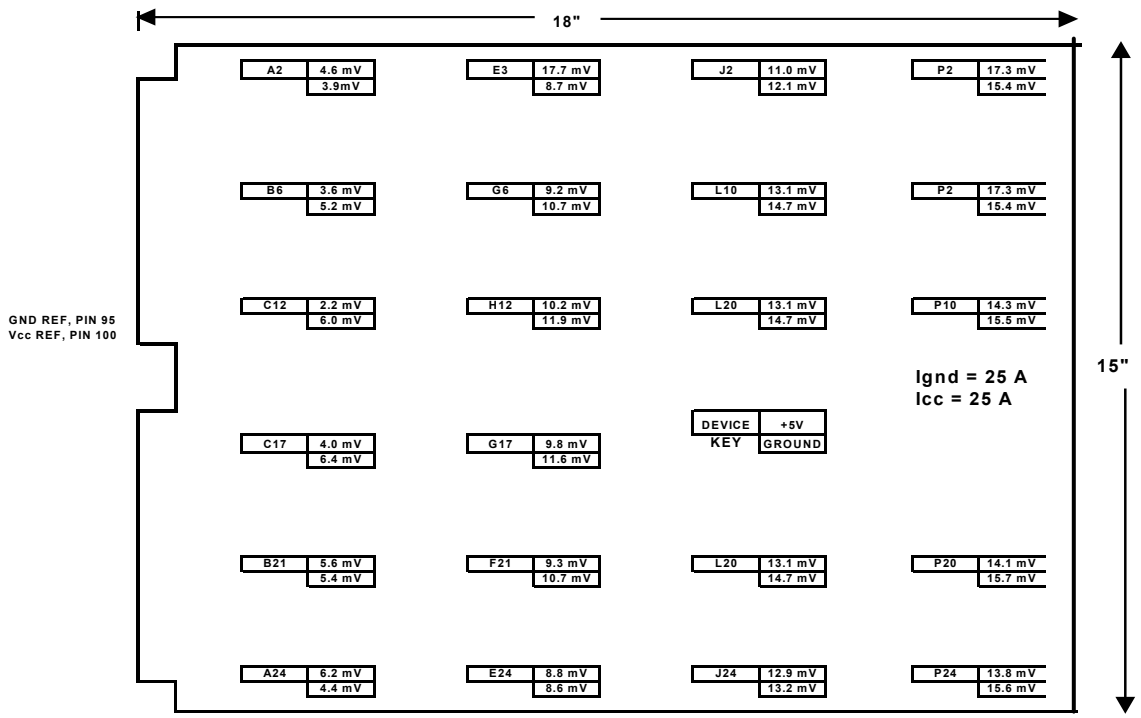

Figure 40.3. Vcc and Ground Voltage Drops in a 6 Layer PCB with Two One-Ounce Power Planes

A question that comes up from time to time is what is the resistance of power planes? The context is often whether significant voltage drops will be developed through the planes between adjacent parts. A simple calculation based on the data in Figure 40.3 can be made. Notice that the 25 amp current resulted in a total voltage drop from front to back of approximately 15 millivolts. Since the current was consumed uniformly from front to back, the average current over the length of the PCB is about 12.5 amps. Using Ohm's law, the front to back resistance of either plane is 1.5 milliohms. This is a very small resistance and validates the assumptions made in transmission line modeling that the plane resistance is negligible compared to the resistance of traces.

What To Do When Current Levels Are So High That It is Not Possible to Put Enough Copper In the Power Path?

Many systems use high power components such as microprocessors, DSPs, network processors and FPGAs that simultaneously require very high currents and low power supply voltages. This results in large voltage drops in the power structure at the same time that the noise margins shrink. The result is a need for far more metal in the power distribution system than is practical.

There are two solutions to this problem. When the DC offsets exceed the noise margins of the logic, differential signaling can be used. This is how the offsets between the ends of Ethernet paths are addressed. The second solution is to reduce the magnitude of the current that flows in the power distribution system. How is this done? Instead of producing the final voltage with a power supply and delivering it through the power distribution network at a very high current, an intermediate voltage, such as 48V, 24V or 12V is generated. This is delivered to the regulators on board that are next to the consumers of the power. These on-board regulators step the voltage down to the final voltage. As a result of this strategy, the current that must pass through the connectors is greatly reduced. At the same time, the final voltage is delivered through a very short path to the loads.

Ground and Vcc Drops In IC Package Power Leads

There is a non-zero DC resistance in the power leads of every IC package. The steady state current drawn by high power ICs will generate a DC voltage drop that will result in Vcc on the die being lower than Vcc on the power plane and ground

on the die being higher than ground on the PCB. These DC offsets directly shift logic levels in such a way that the noise margin is eroded. It is imperative that these offset voltages be quantified and allowances made for them as part of noise margin analysis. This offset can be calculated if the DC resistance of the power paths is known and the quiescent current draw is known.

Reference Accuracy

Many logic families have internal reference voltage generators that are used to determine the threshold voltage between the logic 1 and logic 0 levels. To the extent that the reference voltage drifts away from its design value, either the logic 1 or logic 0 noise margin is eroded. Reference voltages can drift with temperature, time or changes in Vcc. Logic families with reference voltage generators include ECL, GTL, BTL and LVDS. Of these, only ECL shows significant drift with both changes in Vee and temperature. Noise margin allowance must be made for the expected changes in temperature and Vee for this logic family.

Signal Losses in Traces

When parallel terminations are used at the load end of a net, there will be a quiescent current flowing in the transmission line. As shown in Figure 40.4, a voltage divider is formed by the trace DC resistance and the termination resistor. The result is one of the logic levels at the load is less than it should be. Some of the noise budget must be allocated for this erosion. Figure 40.5 is a plot of trace resistance versus trace width and trace thickness. Notice that for narrow traces in thin copper, the resistance per unit length is significant. When this is combined with the fact that a typical GTL or BTL bus is less than 30 ohms, it can be seen that the signal loss may be substantial.

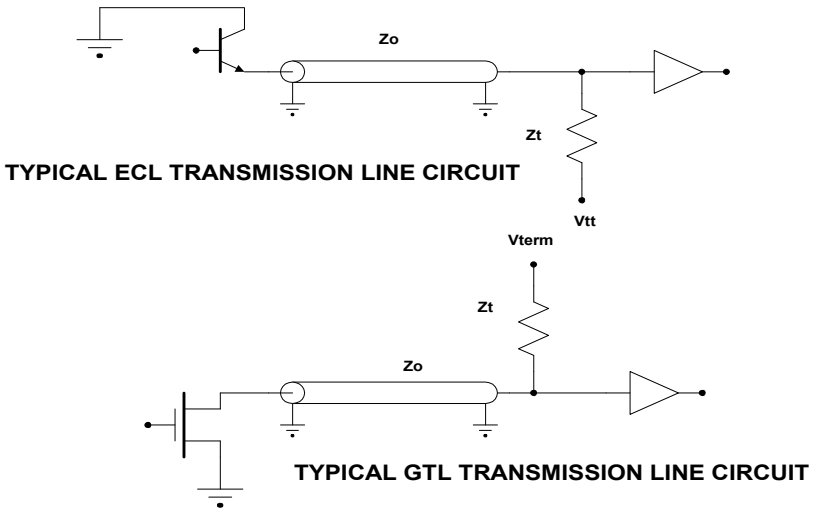

Figure 40.4. Parallel-Terminated Transmission Lines

In addition to the DC resistance of a transmission line, skin effect loss will further erode the signal at the receiver. Skin effect loss is frequency dependent. As a result, it is not possible to make simple calculations to determine how much erosion of a signal will result from this phenomenon. In order to assess the effect of skin effect loss, it is necessary to build a model of the circuit being examined and make time domain measurements. This requires a modeling tool that can model transmission lines in three dimensions. Skin effect losses and the physics behind them are well covered in other texts such as the most recent edition of Howard Johnson's book referenced in the bibliography.

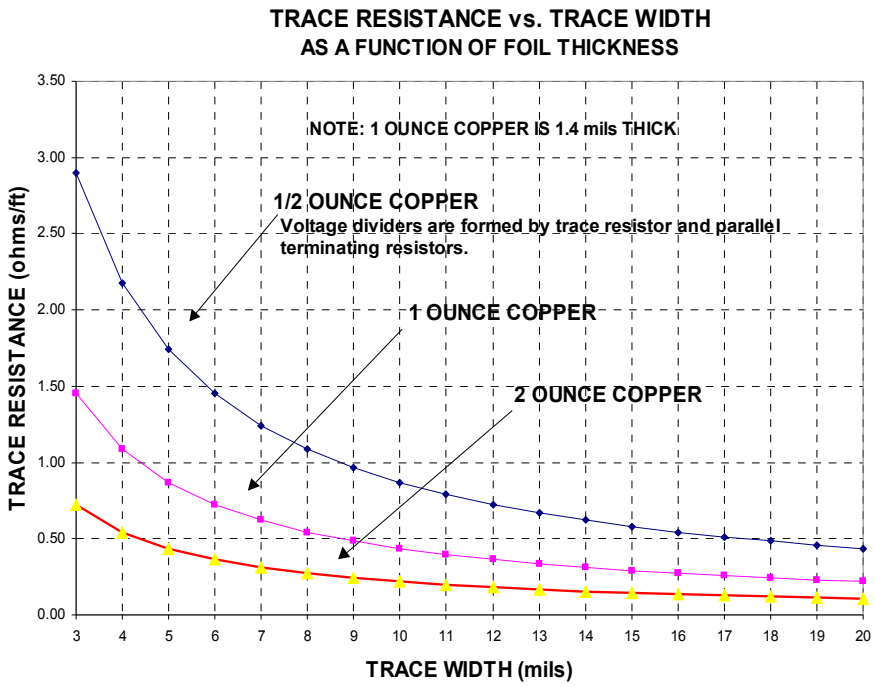

Figure 40.4a. Trace DC Resistance vs. Trace Width and Trace Thickness

Thermal Offsets/Effects

The reference voltages and the values of the output voltages of some logic families are affected by changes in junction temperatures. To the extent that the voltage levels of a logic family are affected by temperature, allowance will need to be made for this. The only logic family that has significant voltage shifts as a result of changes in temperature is ECL. The voltage levels of the various types of MOS logic families are not significantly affected by temperature.

Terminator Noise

When there are large numbers of termination resistors in a design it is common to use resistor packs with multiple resistors in a single package such at that shown in Figure 40.5.

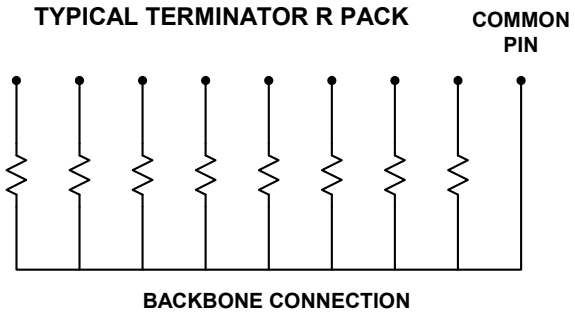

Figure 40.5. Typical R-Pack Used for Parallel Terminations

When a logic signal terminated on one of these resistors switches, the change in current flow associated with the logic change will develop a voltage drop across the inductance in the common backbone connection (parasitic inductance). This delta V will cause the common terminal of all of the other resistors to move inducing a noise voltage on those logic lines. Noise is coupled through this common connection in the R-pack.

R-packs are not as commonly used in CMOS as they were with ECL. This is because most CMOS circuits use series terminations that don't have a common connection. However, newer versions of DDR memory arrays are parallel terminated to a termination voltage that is Vdd/2. It is common to create this termination voltage with a special voltage regulator. All of the parallel terminations are tied to a common node such as that illustrated in Figure 40.5. Care must be taken to insure that this connection is very low inductance and that there is a very good plane capacitor attached to that

node to provide the fast switching currents associated with this architecture. Failure to do so will result in noise being coupled signal-to-signal through this path, eroding signals to the point of failure.

Power Supply Variance

Power supply variance is comprised of three elements. These include: long term voltage drift; DC voltage drops across the power distribution system and ripple associated with switching activities such as driving transmission lines and standby to active current changes with components such as microprocessors. The manufacturer's data sheet typically specifies the total variation that a logic family can tolerate by summing up these three elements. It follows that the tolerance specified will have to be "spent" on these three sources of erosion. It is a commonly made assumption that the total variation is available for ripple. This will result in an unstable system. Tradeoffs will be required between the weight of copper in the distribution system; the quality of the voltage regulators and the total decoupling needed to keep ripple within limits.

The chapters on power system design deal with how to manage all three sources of power supply erosion. It is important to remember that with CMOS logic or any other logic that shorts an output line to the Vdd rail, all of the ripple that appears on Vcc or Vdd will appear unattenuated on the logic lines when the logic level is a 1. When one of these logic lines leaves the product, the ripple voltage on it will create EMI. For these reasons, the maximum allowable ripple will be determined during the noise margin analysis process that follows. Failure to minimize high frequency ripple is a common cause of high EMI.

> With CMOS logic families, all of the ripple on Vdd will appear on signal lines when they are set to a logic 1. As a result, ripple on Vcc or Vdd directly erodes noise margins in CMOS circuits. For this reason, the maximum allowable ripple will be determined by how much of the noise budget is available for ripple. This can also be a major source of EMI.

A Noise Margin Analysis example

The following example shows how to develop a design rule set using the noise margin analysis process. A proposed set of design rules will be examined to determine whether or not the design will function properly if designed to these rules. If not, changes will be made until the total noise generated by all of the rules is within the noise margin of each logic family being used. When the noise budget balances, the set of design rules will result in a stable design.

The rule set that will be examined is that used to design the PCB that is depicted in the left hand side of Figure 29.6, a typical six layer PCB. The routing rules allowed 5 mil lines and 5 mil spaces. In this case, the crosstalk on either outer layer could be as much as 34% and on the inner layers, 25%. The impedance on the outer layers is 102 ohms and on the inner layers 70 ohms. In this example, 5 -volt CMOS logic will be examined first. This logic family has a maximum signal swing of 4700 mV and a noise margin of 1000 mV. Table 40.2 lists the results of the noise margin analysis of this rule set.

MOS NOISE MARGIN ANALYSIS

Using standard 6 layer stackup, 5 mil lines, 5 mil spaces, inducing signal amplitude = 4700 mV

MOS noise margin		**1000 mV**
Reflections loss (107 to 70 ohms)		**982 mV**
Coupling loss (L1 & L6) (34%)		**1598 mV**
Reference accuracy		**0 mV**
Trace IR drop		**0 mV**
Ground offsets		**2 mV**
Thermal offsets		**0 mV**
Ground IR drop in package leads		**10 mV**
Terminator noise		**0 mV**
Vcc/Ground bounce		**100 mV**
Power supply variance (ripple)		**25 mV**
Total system noise		**2717 mV**
Noise margin	**(169% excess)**	**-1717 mV**

Table 40.2. Noise Margin Analysis of a Design Rule Set Using 5-Volt CMOS Logic on a Six Layer PCB

The information in this table was obtained using the techniques describe in earlier chapters. Notice that the total noise exceeds the noise budget by more than 150%. It is likely that any system designed with these rules will be unstable or flaky. Some rule changes must be made to reduce the total noise. The rule changes should be done in such a way as to minimize the total product cost. The two largest noise sources are crosstalk and reflections. A simple way to reduce both of these is to control the impedance and reduce crosstalk by moving the trace layers closer to the planes. Restacking the PCB achieves both of these. This has been done on the right hand side of Figure 29.6. This change adds no cost to the PCB. Table 40.3 is the noise margin analysis that results from this change.

MOS NOISE MARGIN ANALYSIS
Using modified 6 layer stackup, 5 mil lines, 5 mil spaces, inducing signal amplitude = 4700 mV

MOS noise margin		**1000 mV**
Reflections loss (50 to 62 ohms)		**606 mV**
Coupling loss (L1 & L6) (17%)		**799 mV**
Reference accuracy		**0 mV**
Trace IR drop		**0 mV**
Ground offsets		**2 mV**
Thermal offsets		**0 mV**
Ground IR drop in package leads		**10 mV**
Terminator noise		**0 mV**
Vcc/Ground bounce		**100 mV**
Power supply variance (ripple)		**25 mV**
Total system noise		**1542 mV**
Noise margin	**(51.7% excess)**	**-542 mV**

Table 40.3. Noise Margin Analysis of a Design Rule Set Using 5-Volt CMOS Logic on a Controlled Z Six-Layer PCB

Notice that both crosstalk and reflections have been reduced by a significant amount. However, there is still too much noise. Once again, crosstalk and reflections are the major sources of noise and should be examined first for reductions. The amount of noise from reflections is what results from the ±10% impedance tolerance that is the reasonable limit to the PCB fabrication process. Reducing reflections below this limit will increase PCB cost. Crosstalk is the other major source of noise. There are two ways to reduce this effect that, within limits, will not add cost. These reduction techniques include separating the traces and reducing the height of the trace above the nearest plane. In this example, increasing trace-to-trace separation from 5 mils to 8 mils was chosen. Table 40.4 shows the result of this change.

MOS NOISE MARGIN ANALYSIS
Using modified 6 layer stackup, 5 mil lines, 8 mil spaces, inducing signal amplitude = 4700 mV

MOS noise margin		**1000 mV**
Reflections loss (50 to 62 ohms)		**606 mV**
Coupling loss (L1 & L6) (10%)		**470 mV**
Reference accuracy		**0 mV**
Trace IR drop		**0 mV**
Ground offsets		**2 mV**
Thermal offsets		**0 mV**
Ground IR drop in package leads		**10 mV**
Terminator noise		**0 mV**
Vcc/Ground bounce		**100 mV**
Power supply variance (ripple)		**25 mV**
Total system noise		**1213 mV**
Noise margin	**(18.8% excess)**	**-213 mV**

Table 40.4. Noise Margin Analysis of a Design Rule Set Using 5-Volt CMOS Logic on a Controlled Z Six Layer PCB

After only two rule changes, neither of which added any cost to the design, the excess noise has been reduced from more than 1700 millivolts to approximately 200 millivolts. At this stage, there are three significant noise producers. These include: reflections, which cannot be reduced without adding cost; crosstalk, which can be further reduced by increasing

trace-to-trace spacing or reducing the height of the trace layer above the nearest plane and Vcc/Ground bounce, which will require a package change to reduce.

Whether further noise reduction is warranted depends on the application of the design. Remember that the underlying principle used in this analysis is that, on occasion, all of the noise sources will coincide to produce one large noise spike that could cause a failure. The more complex the system, the more likely this event will occur. The simpler the system, the less likely this event will occur. Another consideration to take into account is how likely it is that such a failure will cause problems. If the system is a video game, only the player is inconvenienced if a failure occurs. If the system is a terabit router, is more likely that major damage will be caused by such a failure.

Engineering judgment is needed to determine whether more cost should be incurred to further reduce noise sources. If this were a simple dual speed Ethernet hub, it would be reasonable to stop at this point and design the product. When this decision is made, every design rule needed to design a stable product has been arrived at and routing of the PCBs can proceed with the knowledge that the system will be stable.

Notice that all of the design rules have been determined before a single PCB has been designed and before the first schematic has been drawn. **It is possible to develop all of these design rules this early in the design process. It is vital to do so.**

Table 40.5 presents the same analysis for a GTL data bus. Notice that the signal swings are very small (800 mV) and the noise margin is 370 mV, almost half of the signal swing. This large ratio of noise margin to signal swing is the mark of a well-designed logic family.

GTL NOISE MARGIN ANALYSIS
Using modified 6 layer stackup, 5 mil lines, 5 mil spaces, inducing signal amplitude = 800 mV

GTL noise margin		**370 mV**
Reflections loss (50 to 62 ohms)		**167 mV**
Coupling loss (L1 & L6) (17%)		**136 mV**
Reference accuracy		**3 mV**
Trace IR drop		**2 mV**
Ground offsets		**2 mV**
Thermal offsets		**0 mV**
Ground IR drop in package leads		**10 mV**
Terminator noise		**10 mV**
Vcc/Ground bounce		**100 mV**
Power supply variance		**0 mV**
Total system noise		**430 mV**
Noise margin	**(16% excess)**	**-60 mV**

Table 40.5. Noise Margin Analysis for a GTL Bus

Without any extra effort, this noise budget is almost balanced. With a minor change in spacing, crosstalk could be reduced enough to achieve balance. Notice that there is no allowance for power supply variance. This is because the outputs of BTL and GTL are parallel terminated to a termination voltage that is independent of Vdd which can be more tightly controlled as it only supplies power to the pull up resistors.

Recording The Design Rules

Once the design rules have been developed, it is necessary to record them in a way that is easy for the design team to use them. Table 40.6 presents one way for doing this. This is referred to as a technology table in most design systems. Appendix 5 explains this concept in detail. This design rule concept is based on the knowledge that most designs can be partitioned into a limited number of classes of nets. All the members of a given class will require the same design rules. Once the net list has been partitioned into its classes, labels are given to each class that is added to the net list. Schematic tools intended for this class of design have a field that can be loaded using this label.

PCB layout tools intended for high speed design have the ability to interpret the rules set forth in the technology table and can route traces to them. It should be pointed out that these are the primary design rules needed to route a high speed PCB. The rules are all geometric: length; trace width and trace spacing. There is no need to do real time analysis of the traces during routing. Further, when routing has been completed, post-rout analysis consists of checking these same

geometries. There is no need to do post-route board level signal integrity analysis. That work has all been done prior to routing so the SI goals are met by construction.

CLASS NUMBER	CLASS NAME	TECH NOLOGY	IMPED-ANCE	LAYERS	TERM TYPE	STUB LENGTH	TRACE WIDTH	SPACING IN CLASS	SPACING TO CLASS	LENGTH TUNE	ROUTE ORDER	COMMENTS
1	LVDS	LVDS	50	2,5,6,9	PAR 100	0	5 MILS	6 MILS	MOS 12 MILS	500 MILS	2	TERMINATE ACROSS PAIRS
2	GTL	GTL	50	2,5,6,9	PAR 50	.5"	5 MILS	6 MILS	MOS 12 MILS	N/A	3	TERMINATE TO +1.2v
3	RAM DATA	5MOS	50	2,5,6,9	SER 20	0	5 MILS	8 MILS	SEE ABOVE	N/A	4	TERM AT SOURCE
4	RAM ADDR	5MOS	50	2,5,6,9	SER 10	1"	5 MILS	8 MILS	SEE ABOVE	N/A	6	TERM AT SOURCE
5	SYSCLOCK	5MOS	50	2,5,6,9	SER 30	0	5 MILS	8 MILS	SEE ABOVE	MATCH 3	5	TERM AT SOURCE
6	UNLABELED	5MOS	50	1,2,5,6,9,10	SER 30	1"	5 MILS	8 MILS	SEE ABOVE	N/A	8	TERM AT SOURCE
7	ECL	PECL	50	2,5,6,9	THEV 50	1"	5 MILS	6 MILS	MOS 12 MILS	N/A	7	TERM ACROSS PS
8	WIDE	PWR	N/A	2,9	N/A	N/A	15 MILS	N/A	12 MILS	N/A	1	+12V AND -12V
9	BLANK	5MOS	50	ALL	N/A	N/A	5 MILS	6 MILS	SEE ABOVE	N/A		

Table 40.6. A Typical Technology Table

Summary

Noise margin analysis is the systematic examination of design rules to insure the total noise does not violate the margins of all the logic circuits in a design. It is a relatively simple process that should be performed before design starts. Failure to perform this analysis commonly results in designs that are not stable. Troubleshooting an unstable design takes far longer than does noise margin analysis and is not capable of assuring a stable design, no matter how long a design is exercised.

CHAPTER 41: PCB FABRICATION PROCESS

The process of manufacturing multilayer PCBs is quite complex. Successful, repeatable manufacturing of high speed PCBs requires intimate knowledge of the materials, processes and metallurgies involved. The practice of allowing a PCB fabricator the freedom to modify the makeup of a PCB to fit an individual process is likely to result in vendor-to-vendor differences that are large enough that some PCBs will work while others won't. In almost every class I teach there is an engineer whose manufacturing organization has changed to a different PCB supplier with the result that the products being manufactured don't work properly. The hope is that I will be able to shed light on what is wrong and solve the yield problem. My first question is: "What was the stackup used by the vendor who built PCBs that worked?" The answer is always: "I don't know, they decided what it would be." This open loop process is likely to result in this inconsistent pattern of yields.

Figure 41.1 depicts the components in a multilayer PCB. There are three basic components. These include: the pieces of laminate with copper foil on both sides that are used to etch pairs of inner layers; pieces of glass cloth saturated with uncured resin, often called prepreg, that serve as the glue during lamination and pieces of foil for the two outside layers. The process shown in the figure is known as "foil lamination" meaning that the two outside layers begin as pieces of foil. This is the most economical method for creating multilayer PCBs and is the one that is nearly always used.

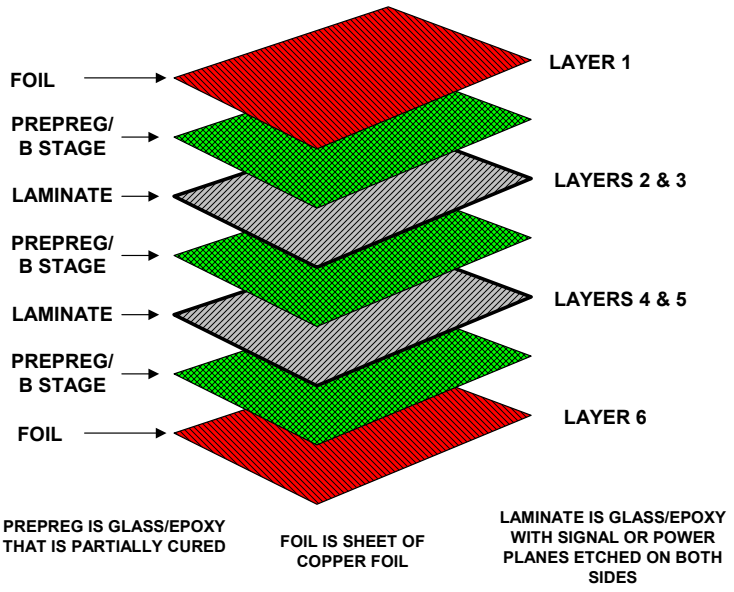

Figure 41.1. Stackup of a Six Layer PCB as It Enters Lamination

The basic steps involved in manufacturing a multilayer PCB include:

- Etch pairs of inner layers back-to-back on pieces of laminate (these are often called details)
- Stack pairs of inner layers using layers of prepreg to separate them
- Place sheets of foil on each outside layer
- Place stack in a press, heat and apply pressure (heat causes the prepreg to flow and harden)
- Drill all holes
- Plate copper in holes and on outer layer features
- Etch away unwanted copper to define traces and other outer layer features
- Apply solder mask and legend

As can be seen from this brief list, there are five major steps in manufacturing a PCB. These include: etch the inner layers; laminate the PCB; drill all of the holes; plate the outer layers and etch the outer layers. At each of these steps, potential defects can be built into the PCB. Understanding the limitations at each step in the process is important in designing high yield PCBs.

Inner Layer Etching

Etching inner layers involves cleaning the copper on both sides of the piece of laminate; applying a photoresist; exposing the photoresist to create the inner layer pattern; developing the resist; etching away the unwanted copper and removing the etch resist. This process is automated in most shops and the chemistry is automatically monitored. As a result, the accuracy and repeatability is quite good. It is possible to etch inner layer traces using this process to an accuracy of ±1/2 mil. This accuracy control helps keep impedance within the tolerances required for transmission lines. Very little goes wrong at this stage. The primary source of problems occurs when the back-to-back images are not aligned to each other.

Lamination

Lamination is made up of two steps: laying up the stack of layers, prepregs and foils and the actual press cycle itself. Layup is a manual process and involves an operator alternating foils, inner layers and prepreg in the proper order. It is at this stage that layers can be put into the stack in the wrong order or prepregs put in the wrong order. If there is not a structure built into the PCB that makes it possible to find a PCB that is incorrectly laid up, it will go through assembly and fail at test with no way to determine the cause of the failure. A method for ensuring proper stack up is to build test coupons into the panel from which the PCB is manufactured. The problem with test coupons is they do not accompany the finished PCB through the assembly and test process. I solve this problem by adding a test structure to one edge of the PCB and calling it stacking stripes. This stack up methodology will be described later.

The press cycle involves stacking all of the components of the PCB on a press plate that has alignment pins to make sure all of the layers are aligned to each other as the lamination cycle proceeds. The “book” of PCBs is placed in a high - pressure press that can heat the PCB to a temperature that will melt the resin in the prepreg layers. Once the resin is melted, pressure is applied to cause the resin to flow into the voids in the adjacent copper layers as well as to drive all of the air out of the stack. Once these steps are completed, the temperature is raised to another level at which the resin in the prepreg layers hardens. From this point, the PCB is cooled down in such a way that stresses don’t build up in the finished PCB.

One thing that can go wrong in the press cycle is when air remains trapped in the PCB. This problem has been solved by putting the press in a vacuum to draw out all of the air prior to the press cycle. This is called vacuum lamination and is the preferred method for laminating high performance PCBs. Another problem that can occur is when pressure is applied too early in the heat cycle which can cause localized distortion of the inner layers as the prepreg resin melts. This is managed by well-engineered press cycles.

Drilling

Drilling is an operation that requires a great deal of care to ensure the drilled holes are placed with the accuracy that is required for high complexity PCBs. Drilling errors can occur because the drill is out of calibration. More often, drills fail to hit the expected accuracy when the target pads are out of location due to misregistration of the artwork to the laminate or when the material shifts during lamination. The methodology for controlling these problems is discussed at length in Appendix 3, Selecting PCB Suppliers. A second problem that can occur is when the pads and clearances through which the drilled holes must pass are not dimensioned correctly. As a result, there may not be enough room to drill holes that pass correctly through the PCB with adequate clearances. The dimensioning process is described in Appendix 2, Anatomy of a Drilled Hole.

Plating and Etching Outer Layers

Plating and etching the outer layers involves many steps and several chemistries. It is this step in the process that most often results in PCBs that are not usable. Figure 41.2 illustrates these steps. This figure is courtesy of Multiwire Corporation, a pioneer in manufacturing discrete wiring boards.

Plating in the PCB fabrication process is done to deposit copper into the drilled holes to form the layer-to-layer connections. Were it not for the through hole connections, plating would not be required. The problem with this process is the drilled holes are lined with a non-conductive material and it is not possible to plate copper to it. This snag is overcome in the step called sensitization and activation. A solution has been developed that causes copper molecules to plate onto the non-conductive material without using electricity (electroplating), the usual method for plating.

Once copper has been electroplated in the holes (vias) and on the surface features and the unwanted copper has been etched away, it is necessary to protect the copper mounting pads from corrosion in order to insure good soldering. There are several choices of protective coatings. These are listed in Table 41.1.

Plated on solder	Hot Air Solder Leveling (HASL)
Organic Coatings such as Entec® 106	Electroless gold over immersion nickel (ENIG)
Electroplated gold over nickel	Immersion silver
Immersion tin	Electroplated tin

Table 41.1. Outer Layer PCB Finishes

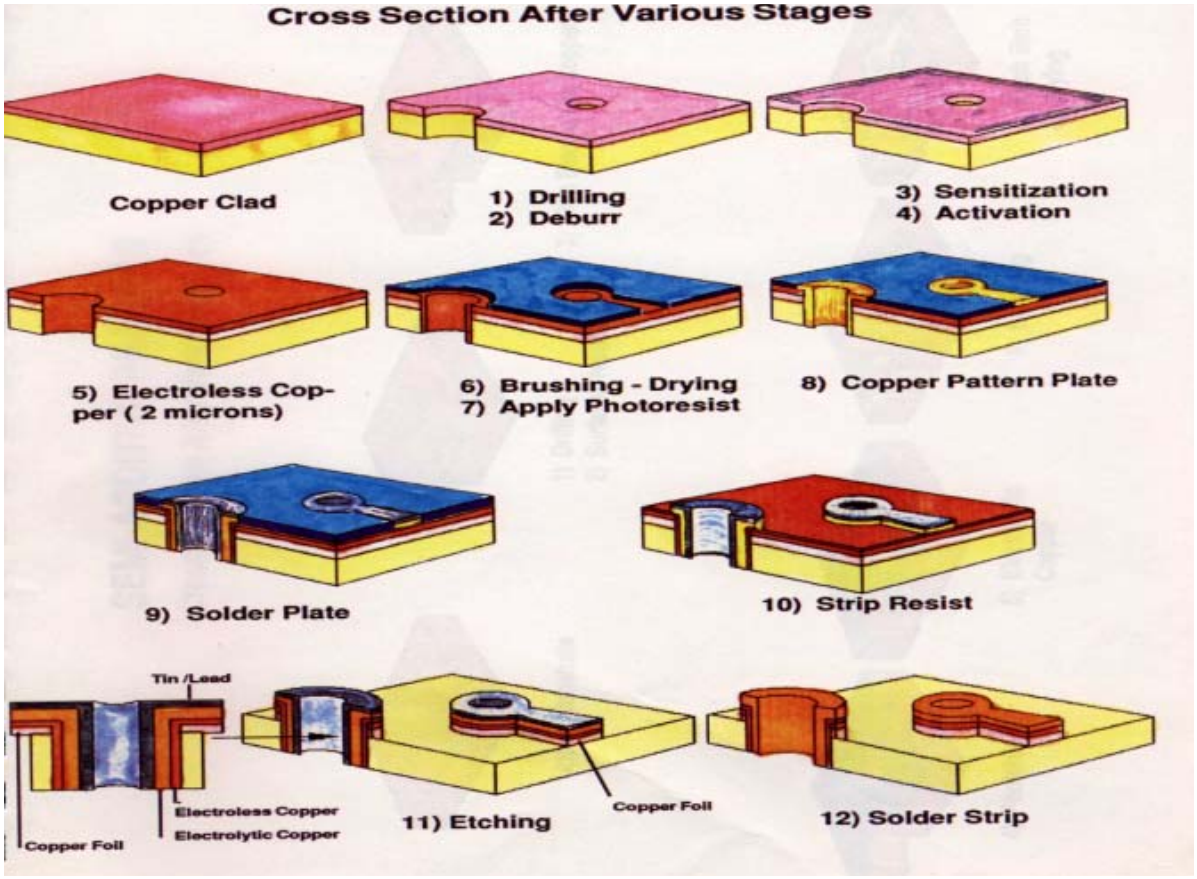

Figure 41.2. Outer Layer Process Steps

Each of these outer layer finishes has its good as well as its bad points. These are discussed in detail in Appendix 3, Selecting a PCB Supplier. The application of these finishes as well as their consistency is where many of the problems with dense PCBs are encountered. If not applied properly, poor solder connections can result in making an entire PCB assembly unusable. Great care should be exercised in choosing a finish.

Most of my PCBs are very dense, often with thousands of dollars of components mounted on them. It is imperative that I minimize the number of PCB assemblies that are lost due to poor connections. Over time, I have found that the most reliable finish is electroplated gold over electroplated nickel. This is one of the more expensive finishes but in the long run it costs less because there are fewer solder defects. Remember, cost and price are not the same. A PCB that has the lowest price finish often is the most expensive due to bad solder connections.

An important part of building a high performance PCB is achieving uniform impedances (within specifications). A major source of impedance variation is the accuracy with which trace widths can be controlled. Table 41.2 lists the major steps involved in creating traces on the inner and outer layers of a PCB. Forming traces on outer layers involves many more steps than does forming traces on the inner layers. For this reason, trace width control on outer layers is much more difficult than on inner layers. When creating inner layers, the only process step that has a significant effect on trace width is the etch cycle which can be tightly controlled. The process step that is most important when forming outer layers is plating enough copper to insure that is thick enough in the plated through holes. As a result, the copper plated on traces will vary considerably. This results in substantial variation in trace width and thickness on outer layers and wider variations in impedance.

For best impedance control, only inner layers should be used for controlled impedance traces. Outer layers should be used for mounting structures and traces that don't require impedance control.

OUTER LAYERS	INNER LAYERS
1. DRILL HOLES TO BE PLATED	1. APPLY ETCH RESIST
2. CLEAN AND DEBURR HOLES	2. EXPOSE IMAGE ON RESIST
3. SENSITIZE HOLES	3. DEVELOP IMAGE ON ETCH RESIST
4. PLATE ELECTROLESS COPPER	4. ETCH AWAY UNWANTED COPPER
5 APPLY PLATING RESIST	5. REMOVE ETCH RESIST
3. EXPOSE IMAGE ON RESIST	
4. DEVELOP PLATING RESIST	Mixing 1 ounce copper and 1/2 ounce or
5. PLATE ELECTROLYTIC COPPER	2 ounce copper on inner layer details will
6. PLATE SOLDER	result in increased processing steps and
7. STRIP AWAY PLATING RESIST	increased PCB costs.
8. ETCH AWAY UNWANTED COPPER	
9. STRIP AWAY SOLDER	
10. APPLY SOLDERMASK	
11. REFLOW SOLDER ONTO PADS, ETC. OR PASSIVATE WITH ENTEC OR ORGANIC COATING	
12. ROUTE PCB FROM PROCESS PANEL	

Table 41.2. Inner and Outer Layer Process Steps

Insuring The PCB Meets Its Specifications

Earlier, it was mentioned that errors can be made during layup that could result in layers being inserted into a given PCB in the incorrect order or prepregs being inserted incorrectly. These errors can occur in a single PCB out of a lot. Some of these errors are very difficult to detect without special test structures. One solution is to build test coupons with each PCB. These test coupons are used to verify impedance, stacking and plating quality. It has been my experience that coupons have two problems--they are often lost early in the process and correlation of impedance in the coupon to the PCB may be difficult. I have evolved test structures that are built into every PCB. Using this approach, the two most common problems with test coupons are eliminated. Figure 41.3 illustrates two test structures.

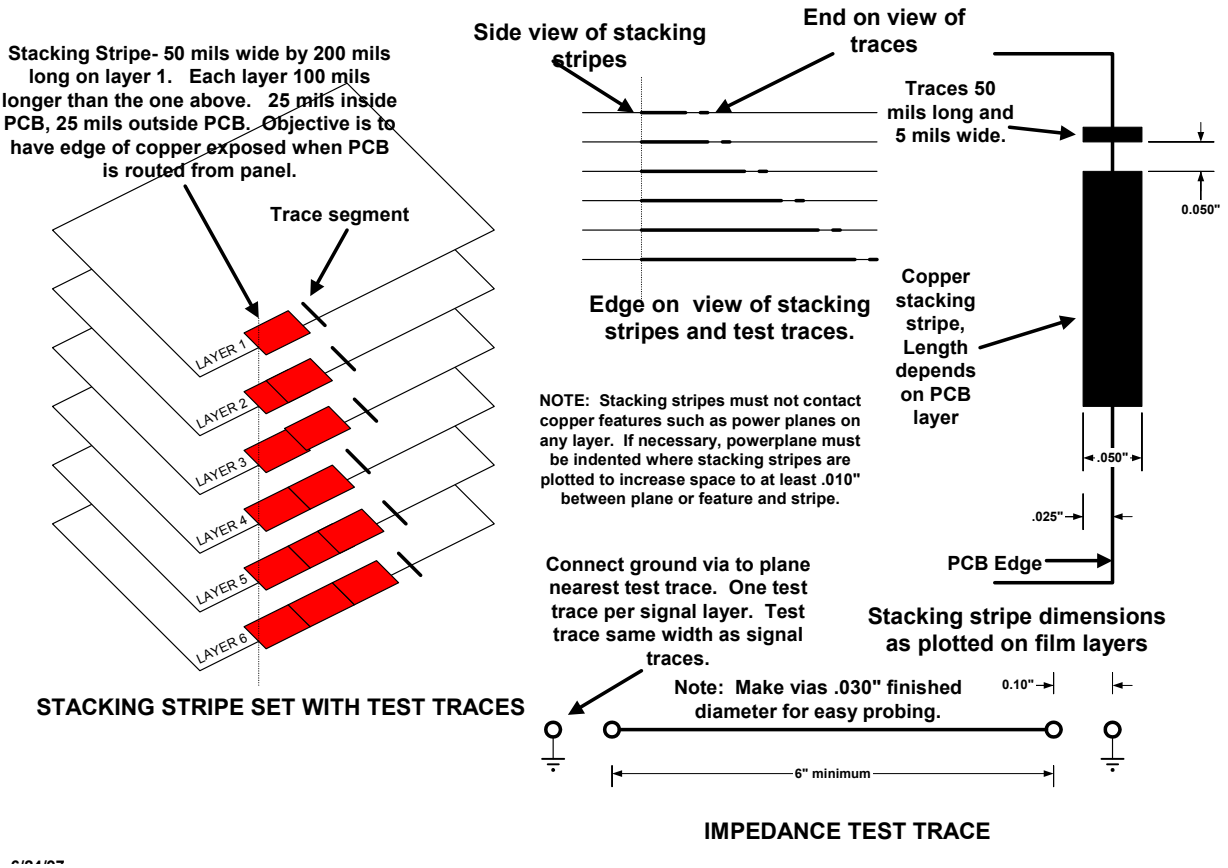

Figure 41.3. PCB Test Structures

The first test structure is a series of stacking stripes that are plotted along one edge of the PCB, each longer than the last. When the PCB is cut from the fabrication panel, these stripes of copper are exposed. It is possible with the unaided eye to spot a PCB with the layers out of order. With a microscope, it is possible to measure the thickness of each dielectric layer and the thickness of all copper layers.

The second test structure is a collection of impedance test traces, one for each layer that has impedance control. The standard access to impedance test traces is a pair of 30-mil drilled and plated holes spaced 100 mils apart. One of these holes is attached to the trace under test and the other to its ground layer.

These test structures add no cost to the PCB and are becoming accepted in the PCB fabrication industry as standard test structures. They provide a fool-proof method for checking each PCB. They should be standard test structures on all high speed PCBs.

CHAPTER 42: PCB MATERIALS

There is a wide range of materials available from which to construct multilayer PCBs. The three main components of these materials are copper foils, woven glass reinforcements and resin systems. Table 42.1 lists several of the more common material systems and the characteristics that are of interest when designing high performance PCBs.

Material	Tg	e_r*	Tan (f)	DBV (V/mil)	WA, %
Standard FR-4 Epoxy Glass	125C	4.1	0.02	1100	0.14
Multifunctional FR-4	145C	4.1	0.022	1050	0.13
Tetra Functional FR-4	150C	4.1	0.022	1050	0.13
Nelco N4000-6 Hi Tg	170C	4	0.012	1300	0.10
GETEK	180C	4.1	0.011	1100	0.12
BT Epoxy Glass	185C	4.1	0.023	1350	0.20
Cyanate Ester	245C	4	0.01	800	0.70
Polyimide Glass	285C	4.1	0.015	1200	0.43
Teflon	N/A	2.2	0.0002	450	0.01
		* Measured with a TDR using velocity method.			
		Resin content 55%			
Tg = glass transition temperature		DBV = dielectric breakdown voltage			
er = relative dielectric constant		WA = water absorption			
Tan (f) = loss tangent					
	All materials with woven glass reinforcement except teflon.				

Table 42.1. Properties of Several Common PCB Materials Systems

The materials are listed in order of their ability to withstand high temperatures, specifically the temperatures involved in soldering (eutectic solder melts at 185°C). The property, T_g, is a measure of this ability. It will be discussed in detail later in this section. Originally, there were only two choices of materials from which to fabricate multilayer PCBs--standard FR-4 epoxy-glass and polyimide. (It might be worth explaining what FR-4 means at this point. FR-4 is a materials classification that means "flame retardant, class 4". It does not specify any particular materials system, but rather it is a performance specification. Using the term "FR-4" to specify a class of materials from which to build PCBs does not carry with it any reference to a particular resin system. This must be specified separately.)

Polyimide has excellent temperature characteristics. To take advantage of these temperature characteristics it is necessary to accept relatively high cost and difficult processing. Further, the material absorbs an excessive amount of water causing leakage problems and outgassing during soldering. This problem is dealt with by baking the PCBs and applying a waterproof coating after soldering. The epoxy glass based system has very good costs and is very easy to process through the fabrication steps. Unfortunately, the T_g is so low that PCBs made from this material are soft when they emerge from the soldering process and, therefore, are easy to warp.

The materials between polyimide and epoxy glass in Table 42.1 have evolved in an effort to produce a material that is easy to process and has a high enough T_g to withstand the soldering process. The first improvements were made by adding other compounds to the epoxy resin system, resulting in the multifunctional and tetrafunctional systems. Both of these improved Tg, but not enough, the goal being close to or above the melting point of solder, 185°C.

BT or Bismalimine Triazine was introduced as a solution in the late 1980s. The T_g of BT was 185°C (the melting point of solder), and appeared to be the solution to the temperature problem. As fabricators gained experience with this material it was discovered to be very difficult to drill, so the search went on.

Soon after this, Cyanate Ester, (CE) was introduced with a T_g of 245°C with processing characteristics very much like the epoxy-based systems. This seemed to be the answer. Unfortunately, CE absorbs water at a greater rate than does polyimide, making it an unacceptable material for commercial PCBs. The search went on.

Getek, a material system developed by GE, was introduced in the mid 1990s. It had processing characteristics much like epoxy and a T_g of 185°C. Again, it seemed like the industry had its answer. As fabricators began to build PCBs from this materials system it was discovered to be inconsistent, resulting in high yield losses.

About the same time that Getek was causing fabrication problems, most laminate manufacturers discovered a way to raise the T_g of the epoxy-based systems to 170°C or above without significantly adding to the cost and without changing the manufacturing characteristics. At last, a material system capable of withstanding the soldering process with good manufacturing properties was available. This material system has become known as **Hi Tg FR-4** and is available from all major laminate suppliers. It is the material system of choice for thick, multilayer PCBs.

Glass Transition temperature, T_g

Figure 42.1 shows the temperature characteristics of several common laminate systems.

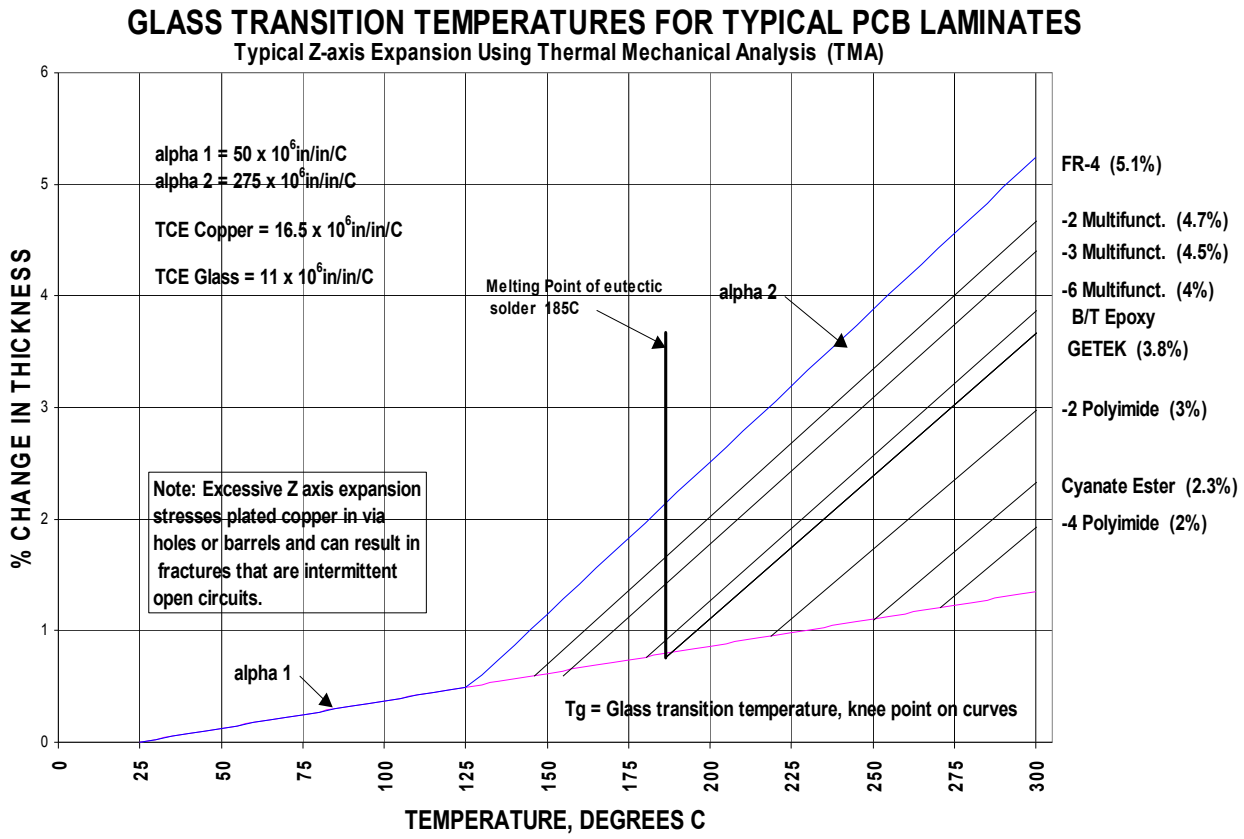

Figure 42.1. Thickness Change vs. Temperature for Several Laminate Systems

In the upper left hand corner of the graph are listed the temperature coefficients of expansion (TCE) of the three materials, glass, copper and resin system, that make up a PCB. Notice that at temperatures below the T_g point, the alpha 1 portion of the curves, the TCE of all three materials are close to each other, close enough that the stresses that build up as temperatures change don't cause problems. When the temperature goes above the T_g point for a resin system, the alpha 2 portions of the curves, the resin TCE increases by nearly six times. The result is the resin volume increases at a higher rate than either the copper or the glass. The PCB is reinforced in both the X and Y directions by the glass and copper components. As a result, all of the volume growth of the resin is in the Z axis creating Z axis expansion. The plated through holes are in the Z axis. As the PCB expands in the Z direction, stress is placed on the copper in the plated through holes. At some temperature the stress reaches the point of failure of the copper and an open circuit results. This stress can be avoided if a material system is chosen that has a T_g at or above the melting point of solder.

Currently, there are two commonly available "FR-4" resin systems, one with a T_g of about 135°C, low T_g, and the other with a T_g of about 170°C, Hi T_g. PCBs 62 mils thick and thinner can be made successfully with the low T_g material. When the thickness exceeds 62 mils, it is advisable to only use the high T_g materials systems.

> PCBs 62 mils thick and thinner can use low Tg materials. PCBs thicker than 62 mils should use only high Tg materials.

Relative Dielectric Constant, e_r

The property of a dielectric that affects impedance and wave velocity is the relative dielectric constant, e_r. In PCB materials, this property varies with both resin content and the frequency at which it is measured. Table 42.2 lists the relative dielectric constants for various thicknesses of a typical Hi T_g FR-4 laminate system.

SOME PROPERTIES OF HI Tg "FR-4" LAMINATE

Data courtesy of NELCO

Thickness	Construction	Resin Content	er @ 1 MHz	er @ 1 GHz
.002	1 x 106	69.0%	3.84	3.63
.003	1 x 1080	62.0%	4.00	3.80
.004	1 x 2113	54.4%	4.19	4.00
.004	1 x 106 + 1 x 1080	57.7%	4.11	3.91
.004	1 x 2116	43.0%	4.54	4.37
.005	1 x 106 + 1 x 2113	52.8%	4.24	4.05
005	1 x 2116	51.8%	4.26	4.08
.006	1 x 1080 + 1 x 2113	52.2%	4.25	4.06
.006	1 x 106 + 1 x 2116	50.8%	4.29	4.11
.006	2 x 2113	43.5%	4.52	4.35
.007	2 x 2113	49.6%	4.33	4.14
.008	1 x 7628	44.4%	4.49	4.32
.010	2 x 2116	51.8%	4.26	4.08
.014	2 x 7628	38.8%	4.69	4.53

Under construction, the three or four digit number refers to the glass weave type.

Table 42.2. Properties of a Hi Tg Fr-4 Laminate System

Notice that for a given thickness it is possible to have very different amounts of glass and resin. These variants have significantly different relative dielectric constants. For example, there are three ways listed to make a .004" thick piece of laminate with resin content varying from 43% to 57.7%. This is accomplished by using different styles of woven glass. The result is a wide variation in relative dielectric constants. It is this variability in materials that often results in PCBs performing differently when made by different manufacturers. To avoid this problem, the specific glass styles for each opening in a stackup must be listed on the fabrication drawing. This will be discussed in Chapter 43.

It should also be noted that the relative dielectric constant varies with the frequency at which it is measured. At 1 MHz and a resin content of 38.8%, the popular 4.7 value of e_r is achieved. This is the old, low-cost, cheap and dirty version of FR-4 from which the 4.7 value was obtained. However, high performance PCBs cannot be made from this material. It is too thick. Materials that are used in multilayer PCBs are thinner and have higher resin content. On the average, the most common resin content is 55%. When this percentage is coupled with the fact that the frequencies of interest in high performance PCBs are 1 GHz and higher, it can be seen that a more accurate value for e_r is around 4.

When calculating impedance, the actual laminate styles used in each opening in the stackup must be called out and the e_r for that laminate used. The frequency at which impedance should be calculated is circuit dependent, but is usually between 1 and 2 GHz for ordinary logic and higher for the very high performance data paths coming into use.

> Using the 1 MHz value of e_r from a material with low resin content gives rise to the notion that the e_r of FR-4 is 4.7. Using this value to calculate impedance for a multilayer PCB guarantees impedance errors.

Loss Tangent, Tan (f) or γ

Loss tangent expresses the amount of the energy in the electromagnetic field traveling through a dielectric that is absorbed by that dielectric. The amount of energy absorbed increases with frequency. The values in Table 42.1 allow us to compare one material to another but don't help in deciding which material is "good" and which is "bad". To determine this requires a complex analysis using analytical tools that calculate loss as a function of frequency as well as length of line. Figure 42.2 is the result of one such analysis. It shows the dielectric loss as a function of frequency for three different materials with a path 33 inches long.

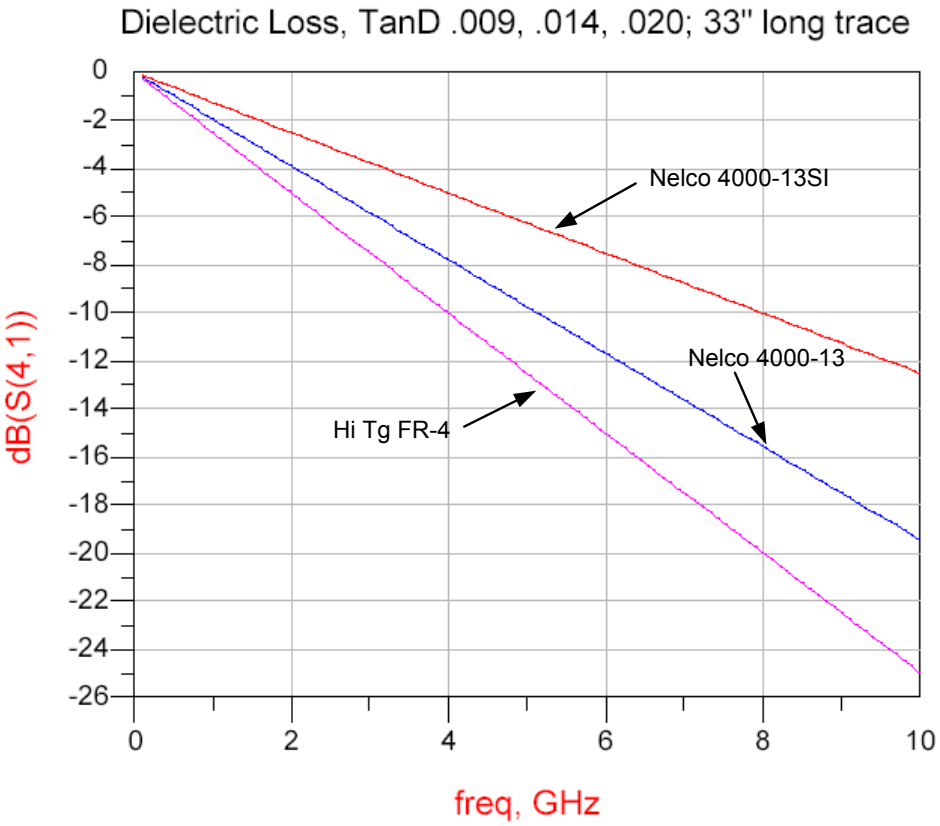

Figure 42.2. Dielectric Loss as a Function of Frequency for Three Materials, 33" long path

When dielectric loss is thought to be important, this level of analysis is required.

Figure 42.3 is a plot showing the relative loss tangent of several PCB laminate systems and the relative cost of using each. This chart was prepared by the engineering staff at Merix Corporation, a manufacturer of high performance PCBs, to help their customers understand the differences between all of the competing "low loss" materials systems. Lower loss is to the left and higher loss is to the right. Higher cost is to the top and lower cost is to the bottom. As a point of reference, the material called Hi Tg FR-4 has a loss tangent like that of FR406. FR406 is the Hi Tg FR-4 material manufactured by Isola Corporation.

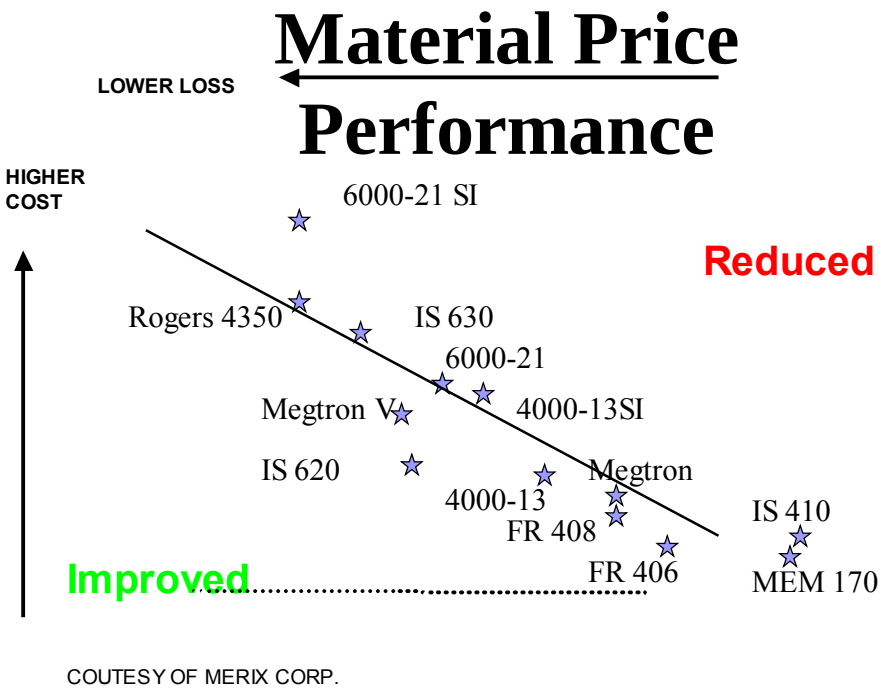

COUTESY OF MERIX CORP.

Figure 42.3. Loss Tangent vs. PCB Cost

Dielectric Breakdown Voltage- DBV

Dielectric breakdown voltage is a measure of how well a material withstands high voltages impressed across it before it arcs through and fails as an insulator. As can be seen from the values listed in Table 42.1, all of the commonly used PCB laminates have DBV values of 1000 volts per mil of thickness or more. Stated another way, they are very good insulators. It does not take much material to meet a 2000 volt breakdown voltage specification. The practice of removing whole planes from underneath transmission lines in order to meet the 1750 volt specification usually applied to Internet products is not required.

Water Absorption

All laminates absorb some amount of water. This water can have two effects. When a PCB passes through the soldering process, the water boils and escapes as a gas. This outgassing creates bubbles in the solder and can cause the solder to explode, spraying solder balls across the surface of the PCB. After a PCB has been assembled, the laminate will absorb water again. If the amount of water absorbed is high enough, leakage paths will develop between circuits causing failures.

What level of water absorption is acceptable? From experience, it has been shown that water absorption values of .2% by volume or below do not cause these problems. If a material with a water absorption above .2% must be used, it will be necessary to bake the PCBs dry prior to passing them through the soldering operation. It may also be necessary to add a waterproofing after assembly in order to avoid leakage failures in the field.

Fabrication Process Tolerances

There are many places in the PCB fabrication process where there is a dimensional tolerance limit. Table 42.3 lists these process tolerances. There are two columns of data in metric and English units. These are standard tolerance and advanced tolerance. Many PCB fabricators have similar charts showing the accuracy with which PCBs can be manufactured. The meaning of the two tolerance ranges, standard and advanced, are as follows. PCBs designed to the standard tolerance limits can be manufactured in very high volume without requiring special monitoring of the drilling, registration and plating processes. This will result in the lowest cost PCB. PCBs designed to the advanced tolerance limits can be manufactured in volume, but will require constant, real-time process management as well as special materials and drill compensation activities. These PCBs will be very difficult to manufacture in high volumes and will limit the supplier base to only those fabricators that have the very best process controls.

PCBs designed to tolerances tighter than those listed in this table are very likely to be very low yield if not zero yield. These tolerances should not be violated without consulting the fabricator who will be asked to build the PCB. Further, if a fabricator advocates violating these tolerances, it is advisable to insist on examples of successful PCB manufacture to

these new tolerances. Citing a prototype with those tolerances does not constitute endorsement that the tolerances are reasonable.

PCB FABRICATION TOLERANCE CHART

	PCB MANUFACTURING CAPABILITIES (2003)				
		ENGLISH		METRIC	
		STD PROCESS	ADVANCED PRODUCT	STD PROCESS	ADVANCED PRODUCT
1	Drilling- Aspect Ratio	6:1	10:1	6:1	10:1
2	Min Drilled Hole size- vias	.012	.008	0.30	0.20
3	Min Finished Hole size- vias	.008	.006	0.20	0.15
4	Min Outer Layer Via Land Size	.031	.022	0.78	0.55
5	Min Inner Layer Via Land Size	.031	.020	0.78	0.50
6	Min Via Relief of Plane Layers	.036	0.024	0.91	0.60
7	Min Blind/Buried Via Land Size	.031	.020	0.78	0.50
8	Min Blind/Buried Via Drill Size	.012	.010	0.30	0.25
9	Min Outer Layer Line Width	.005	.003	0.13	0.08
10	Min Inner Layer Line Width	.005	.003	0.13	0.08
11	Min Outer Layer Space	.005	.004	0.13	0.10
12	Min Inner Layer Space	.005	.004	0.13	0.10
13	Line to Via Land Spacing	.005	.004	0.13	0.10
14	Layer to Layer Registration Tolerance +/-	.008	.005	0.20	0.13
15	Min Component Pitch	.025	.010	0.63	0.25
16	Max Overall PCB Thickness	.187	.500	4.71	12.59
17	Min Dielectric Thickness	.005	.0022	0.13	0.06
18	PCB Edge to Conductor	.020	.010	0.50	0.25
19	Soldermask Clearance Per Side	.010	0.003	0.25	0.08
20	Line to SMT Minimum Space	.010	.004	0.25	0.10
21	Min Base Copper Weight	.0007	.00035	0.02	0.01
22	Average Layer Count	10	16	10	16
23	Dimension Fab Panel OD	18" x 24"	22" x 34"	46 x 61	56 x 86
24	Fabrication Radius	.062	.016	1.56	0.40
25	Warpage- Design Dependent	1%	0.5%	1%	0.5%
26	Tolerance- Plated Thru Holes Desgn Dep	+/-0.003	+/-.002	+/-0.08	+/-0.05
27	Impedance Tolerance	+/-10%	+/-5%	+/-10%	+/-5%
Note: English dimensions are mils, metric dimensions are mm.					

Table 42.3. PCB Manufacturing Tolerances

CHAPTER 43: CREATING PCB STACKUPS

Stackup is a term that describes the arrangement of the layers in a PCB. By now, the reader knows that a PCB is much more than just a carrier of wires and components. It is a critical part of the overall circuit design. The PCB contains two of the most important elements in a high-speed design--transmission lines and interplane capacitors. Interplane capacitors supply the switching currents or charge for the parasitic capacitance of the transmission lines for the very fast switching edges involved in modern logic operations. Creating a stackup involves making tradeoffs between these and many other conflicting requirements. Among these are:

- Minimize cost
- Optimize manufacturability
- Meet impedance goals
- Meet crosstalk goals
- Provide adequate interplane capacitance
- Provide adequate signal routing layers

As was established in the power subsystem design section, interplane capacitors are some of the most important components of a high-speed printed circuit design that must support the fast switching edges of modern logic circuits. In order to insure adequate interplane capacitance, this is where stackup design must begin. The total amount of interplane capacitance needed for each supply voltage is established using the methods outlined in Chapter 37. From this information, the plane area and plane separation required to create this amount of capacitance is calculated. Trade-offs between the surface area of the planes and the plane separation can be made using the equation in Figure 43.1 or using the curves in the chart itself.

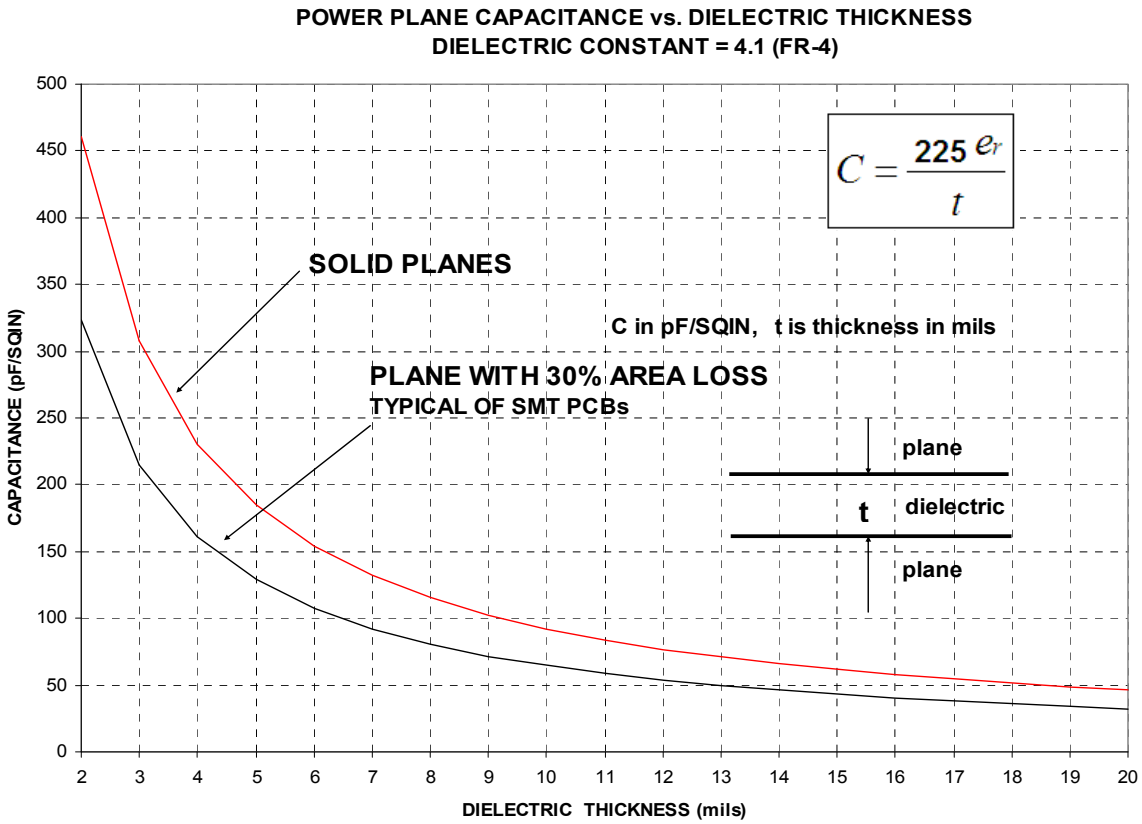

Figure 43.1. Interplane Capacitance vs. Plane Separation

Once the surface area needed for each of the interplane capacitors has been established, the number of power and ground planes needed to create this can be also be established.

The next step in the process is to estimate the number of signal layers that will be needed to provide routing space for all of the wires. The signal layers and power plane pairs can then be organized into a stackup. Figure 43.2 depicts a 10-layer design created in this manner.

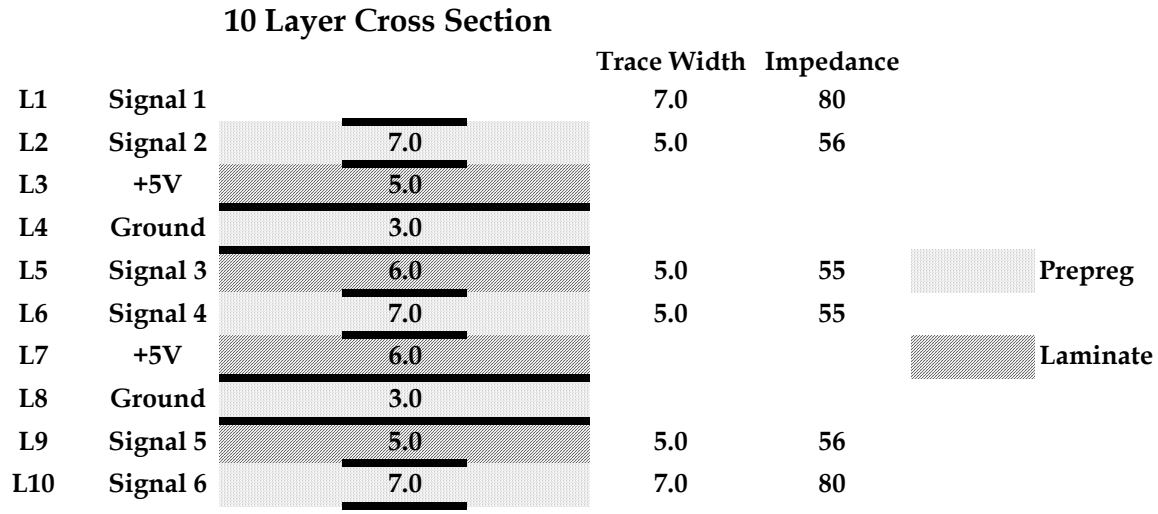

All inner layers 1 ounce copper.
Overall PCB thickness 62.5 mils, +-10

Outer layers 1/2 ounce foil plated to 2 mils total thickness.
Impedance Accuracy- +-10%
Laminate thicknesses- =-1mil

Figure 43.2. A 10 Layer PCB Stackup That Meets a Set of Impedance and Capacitance Requirements

When drawing a stackup, a common convention is that a long bar represents a plane layer and a short bar represents a signal layer.

Notice that in Figure 43.2, there are four buried signal layers, all of them about 55 ohms, and two "cap" layers or outer layers that are of a much higher impedance. These "cap" layers are not used for controlled impedance traces. Instead, they are used for component pin redistribution, mounting components and as a place to route "non-critical" signals. The controlled impedance layers are two buried microstrip layers and two stripline layers.

The arrangement of signal and power layers has been done such that all four of the controlled impedance layers are mated with a power plane back-to-back on a piece of laminate. This is done to improve the impedance control of the signal layers. The two most important dimensions in determining trace impedance is the height above the nearest plane and the trace width. Putting controlled impedance traces on inner layers rather than outer layers results in better line width control. Placing signal layers and planes back-to-back on pieces of laminate allows tighter control of the height above plane than is possible when using prepreg as the separator. A secondary reason for mating signal layers with plane layers is to improve the dimensional stability of the signal layers as a way to improve alignment of all the layers to one another in the final stackup.

The choice of the dimension between the plane and its nearest signal layer has been done as the first step in determining the thickness of the insulating layers. In order to keep crosstalk within a signal layer to a minimum, this dimension should be the thinnest possible while still maintaining manufacturability of the PCB. 5 mils is the thinnest commonly available laminate or core thickness. Separation between adjacent signal layers has been done last to achieve the final, overall targeted thickness of the PCB. The thickness of the prepreg between the outer layers and the first buried layers was also selected based on the same criteria. Thickness variations in these two dielectric areas, signal layer to signal layer and surface to first buried layer, have the least affect on impedance and crosstalk. So, these are good places to add thickness when that addition is required to meet the overall PCB thickness goal.

To obtain maximum interplane capacitance, the plane layers are separated with the thinnest prepreg possible. It is advisable to use two plies of glass cloth in this opening. This is done to insure there is no voltage breakdown between the planes. Figure 37.9 is an illustration of stackups done using this approach.

What kinds of layers can be used for "high speed" signals?

There are three possible signal layer types in a multilayer PCB. These include the surface microstrip, buried microstrip and stripline layers. In Figure 37.9, all three layer types are shown. Which of these can be most successfully used for high speed signals?

The outer layers or surface microstrip layers are subject to many process steps that can result in substantial trace width variation and the associated variations in impedance that make them undesirable for high speed signaling. I use the outer layers for redistribution and "non-critical" signals only.

There is the perception that buried microstrip layers cannot be used for high speed signals. One reason given is that there will be emissions or EMI from these traces. This is not true and has been validated many times with laboratory tests, one such set of tests being those done by UMR and reported in an article by Doug Brooks entitled, "90 Bends, the Final Turn". A second reason given for not using buried microstrip layers is that they are electrically inferior to stripline layers. Again, this is not true. From Figure 24.5, it can be seen that the velocity of propagation for a buried microstrip is essentially the same as that for a stripline. The impedance can be made the same by appropriate trace widths and the last consideration, crosstalk, can be easily made the same as for the stripline layers by setting appropriate trace-to-trace spacing. In this way, buried microstripline and stripline layers can be used interchangeably. This caries with it economies as well as the ability to mate every signal layer with a plane and mate plane layers back-to-back to achieve maximum plane capacitance.

Steps in Creating A Stackup

1. **Determine how many signal layers are needed**
2. **Determine how many power planes are needed to distribute power and ground**
3. **Arrange signals and planes to accomplish**
   - **Partners for signal layers**
   - **Parallel plate capacitance between power and ground**
4. **Set signal height above planes to meet crosstalk requirements**
5. **Set trace widths to meet impedance requirements**
6. **Set spacing between planes to meet capacitance requirements**
7. **Set spacing between signal layers to meet overall thickness**

Documenting a Stackup

Once a stackup has been engineered that meets electrical and manufacturing goals, it must be documented such that all who will be involved in building the PCB have a clear understanding of the parameters that are important. It should be pointed out that the process of designing a stackup is an iterative one that involves proposing a stackup that meets electrical goals and then having it reviewed by a competent PCB manufacturer's engineering group for both manufacturability and cost effectiveness. The outcome of this iterative process is a stackup that meets all of the needs of both manufacturing and engineering. Figure 43.3 is an example of the type of information required. This information is produced by a Signal Integrity tool with a field solver and a stackup editor.

In addition to the information in Figure 43.3, the actual glass styles used in each dielectric layer must be specified. This is to insure the same glass style and resin content is used by each fabricator who will build the PCB. In Chapter 23, Table 23.3, it was shown that for a given thickness of laminate, there may be several formulations with different dielectric constants. Substituting one for another will cause the impedance to vary.

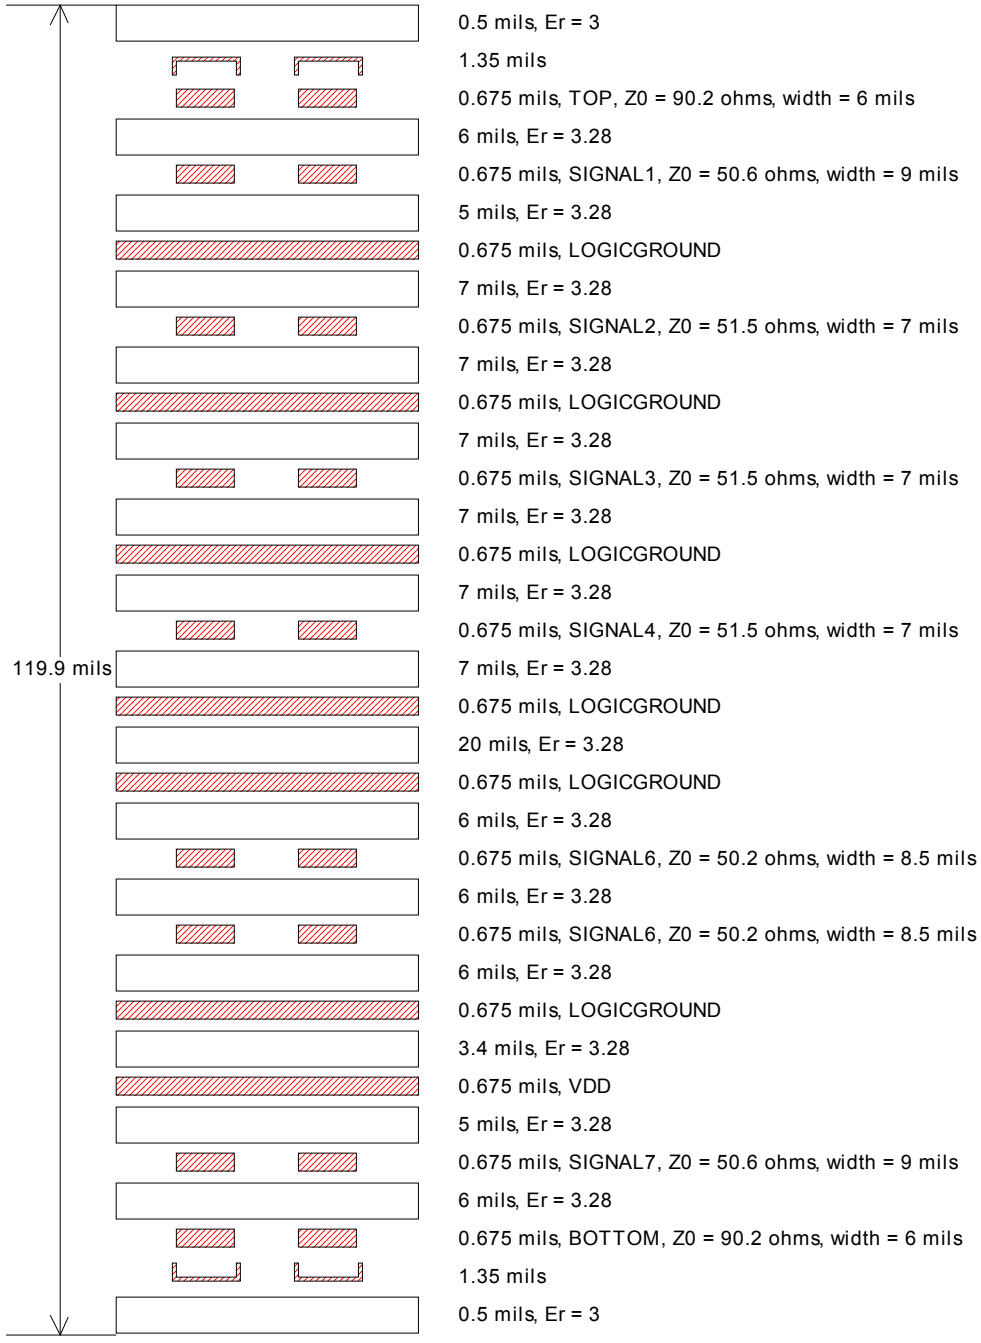

Figure 43.3. A Stackup Drawing Listing Key Parameters in the Construction of a PCB

Reminder, In addition to the above information, the actual glass styles used in each dielectric opening must be specified.

CHAPTER 44: TYPES OF VIAS

Vias are the holes drilled in or through a PCB that provide electrical access from one layer to another. The vias may be used as component pins (through hole or surface mount) or to change from one layer to another. These drilled holes are plated with copper to provide that connection. There are three basic types of vias: through hole vias; blind vias and buried vias. Figure 44.1 illustrates all three of these via types.

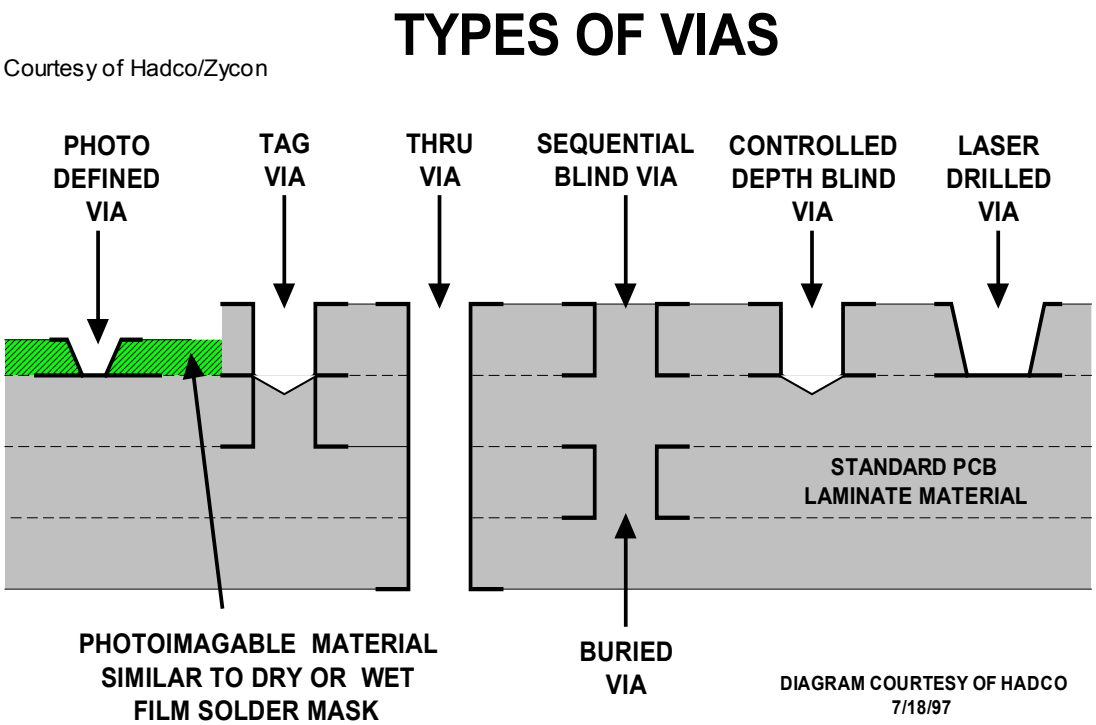

Mcdermott, Brian J. "The Practical Application of Microvia Technology" Circuitree, March 1998.
Cable, Alan. "Laser Drilling Microvias in Printed Circuit Boards." Circuitree Magazine. November 1996.

Figure 44.1. Types of Vias in PCBs

Through Hole Vias

Through hole vias are the most common via type. These are almost always created with mechanical drills followed by electroplating. The primary limitation for through hole vias is the ratio of length to diameter or aspect ratio. Successful plating requires a hole diameter with respect to its length that allows enough plating solution to flow through the hole to permit sufficient copper to be plated on the hole walls. For most high volume PCB manufacturing processes, the aspect ratio should be kept below 6:1. For higher performance PCBs this can be extended to 8:1. Aspect ratios of 10:1 and higher require special process control and should be avoided.

Blind Vias

Blind vias are vias that begin on one side of a PCB, but do not pass all the way through to the other side. The term microvia is often used to describe a blind via. The term microvia actually means a via whose diameter is less than 8 mils whether or not it passes all the way through the PCB. There are four ways to create a blind via. They include: control depth drilling, photo imaging, laser drilling and sequential lamination. The first three are plated at the same time as the through hole vias. Sequential blind vias are plated as a separate operation that will be discussed later.

Controlled depth drilling is just what it sounds like. A very accurate drill machine is programmed to drill only part way through the PCB. This type of blind via requires that there be an area below the via that is void of conductors in order to allow for the tolerance with which the drill can stop at a given depth.

Laser drilling involves using a high power laser beam to first drill through the copper on the outer layer and then drill through the first layer of dielectric, stopping at the copper pad which will serve as the connection to the second layer down. A "backstop" pad must be placed on the buried layer to serve as a stop for the laser beam. Most laser beams used for this purpose will remove copper as well as dielectric. This creates a dilemma. The beam must be powerful enough to remove

all of the resin and glass in order to insure a good contact with the inner layer pad. At the same time, it must not be so strong that it erodes the inner layer pad. This is a difficult balance to maintain. As a result, most designs that use ½ ounce copper for inner layers suffer from opens from one of these two limits. A solution is to make the copper on the second layer 1-ounce copper.

Photo imaged vias are formed by applying a photosensitive material similar to solder mask to the outside of a PCB that has gone through the normal lamination process for the inside n-1 layers including etching the traces on the outer layers of the laminate stack. Capture pads for the blind vias are etched on these layers. The photosensitive material is exposed with the openings that will form the blind vias. This material is developed which opens the holes down to the buried layer. Copper is plated onto the outer layer to make contact to the inner layer pad. After this plating step, the outer layer features are etched.

This process can be repeated several times "building up" signal layers one after the other. Should this type of blind via be needed, the suppliers who have the equipment to do it are all in Japan, limiting the supply base.

As can be seen, there are several extra process steps involved in forming photo-imaged blind vias as compared to mechanically drilled or laser drilled blind vias. Because of this, a PCB with the need of only a few vias will cost more than one that has been drilled. On the other side of this, it costs the same to form ten thousand vias as it does one. PCBs that have a very large number of blind vias will cost less to produce with photo-imaged vias than with drilled vias. In fact, this is the criteria upon which the decision of what method to use is based. PCBs with few blind vias are best drilled. PCBs with large numbers of blind vias are best photo defined.

Sequential blind vias are created by forming a two-sided PCB from just layers 1 and 2. This thin piece of laminate is drilled, plated and etched as if a two layer PCB were being made. This subassembly is combined with all of the other inner layers, laminated, drilled and plated as for any multilayer PCB. This is a time-proven method for forming blind vias. It has a higher cost than all of the others due to the fact that there are two sets of drilling and plating operations involved in making one PCB. Added to this extra cost is the fact that the very thin laminate involved in forming the subassembly is delicate. Many are lost from handling damage further adding to the cost.

Buried vias are vias that pass between two or more inner layers without ever reaching either outside surface. These vias are formed using the same method used for sequential blind vias and carry with them the same added cost. PCBs with buried vias often cost 50 to 100 percent more than a PCB with the same number of layers that is simply laminated.

Plating blind vias formed by either drilling operation are plated at the same time as the through hole vias. As a result, creating blind vias in this manner adds only one additional step to the manufacturing process--a second drill operation.

Limitations of Blind Vias

Blind vias formed by all of the processes except sequential lamination are plated with a plating solution that must enter from only one end of the hole. As a result, it is difficult to get chemistry to the bottom of deep holes. Experience has shown that it is difficult to achieve reliable plating when the hole is deeper than it is in diameter. Said another way, the aspect ratio must not exceed 1:1. Most fabricators ask for an aspect ratio of 0.75 to 1 or lower, with the hole diameter being 1.5 times the hole depth. The implications of this restriction are that blind vias can connect only from layer 1 to layer 2 or layer n to layer n-1.

When considering blind vias as a method of connecting very fine pitch components, this limitation must be kept in mind. The types of devices that can be connected using blind vias have very low pin counts and are organized like memory parts, where all of the connections flow in one direction similar to a data bus or address bus. A common error is to assume that blind vias used to connect up memory parts can be used to connect up processors or other ICs with high pin counts.

Blind vias can connect between the top layer of a PCB and the next layer down. It is very difficult to connect to layers below layer two using blind vias. It is imperative that devices that must be connected with blind vias be capable of being wired in layers 1 & 2 or layers n and n-1.

Handling A Micro BGA Part That Must Connect Below Layer 2

BGA parts with pad pitch less than 1 mm must use blind vias to make connections because the pitch between pads is too small to allow room for through hole drilled vias. Usually, the power layers of a multilayer PCB are deeper in the PCB than layer 2. This presents a dilemma. If the blind vias needed to connect the part cannot reach below layer 2, how can power be connected? Figure 44.2 shows a method used to solve this problem.

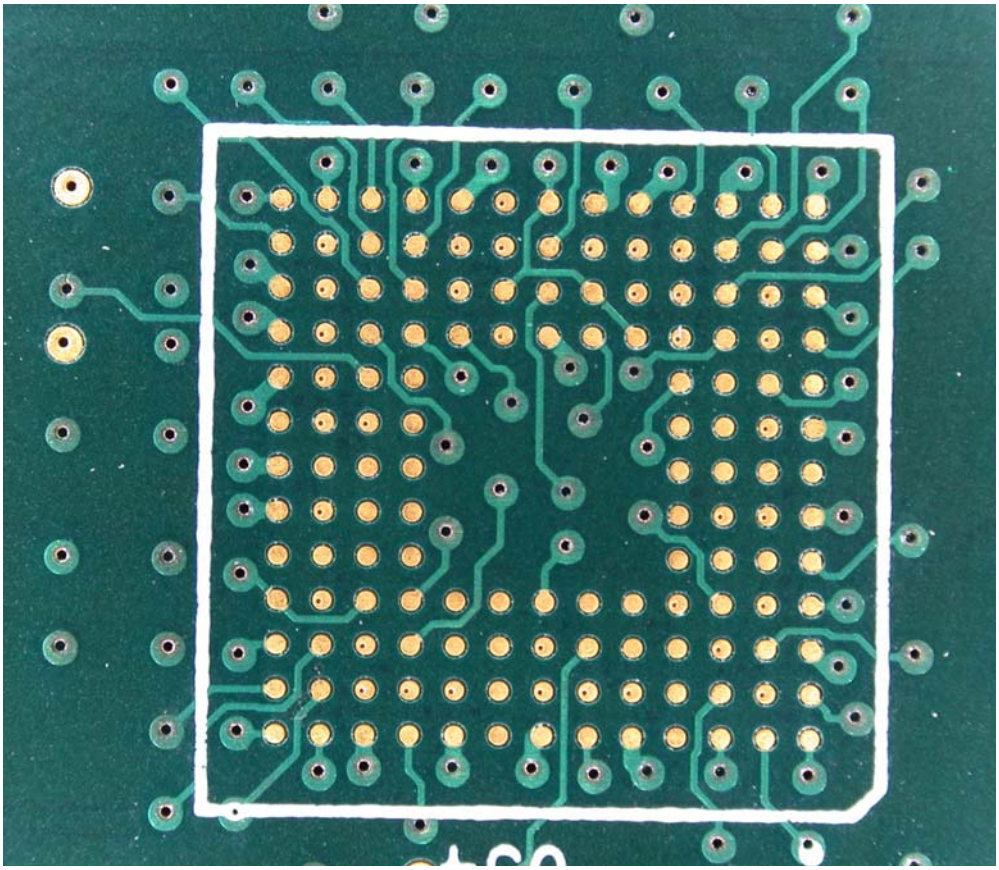

Figure 44.2. The Surface Pattern For a Micro BGA (25 mil pitch) Showing How Power Connections Are Made

The component in Figure 44.2 is a memory IC with mounting pads on a 25 mil pitch. This pitch is too tight to allow through hole vias needed to connect power to the part. As a result, the power pins are routed with traces on layer 1 to an area on the PCB that is far enough away from the BGA footprint to allow through holes. There is a danger associated with connecting power in this manner. The danger is that unwanted inductance will be introduced in the power path that will make Vcc and ground bounce worse.

The BGA pads without traces connected to them are connected on layer two using blind vias. The tiny dimples in the center of these pads are the blind vias.

The artwork required to drill and route blind vias is more complex than that required for an all through hole PCB. There are special needs that are based on the drill used to form the blind vias. As artwork is created, it is advisable to work with the fabricator who will build the finished PCB so that features needed by the laser drill are included in the artwork.

CHAPTER 45: PCB DESIGN PROCESS

Once design rules have been developed for a PCB design, the next step is to proceed through the design process and finish with a working PCB. It is useful to examine the various ways in which this process can be organized. First, let us examine the "traditional" design process often referred to as "TTL" engineering. Figure 45.1 depicts this classic "over the wall" method of organizing the design process.

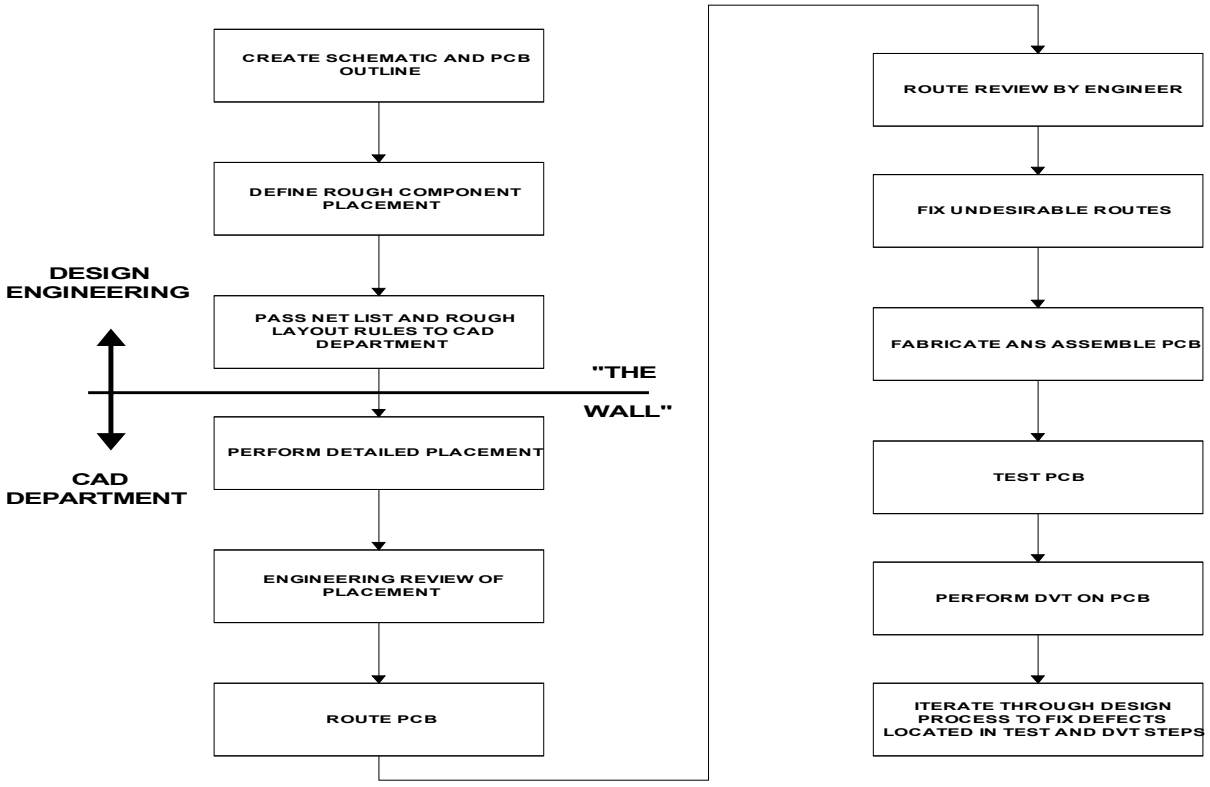

Figure 45.1. The Traditional "TTL" or Hardware Prototyping PCB Design Process

The process shown in Figure 45.1 relies on the step called DVT as a method for checking on the robustness of a design. This step usually consists of building one or more hardware prototypes that are checked in a lab to determine whether or not the design performs to its specifications. If it does not, the reasons for the failures are determined; the design is modified and the process repeated. These design cycles are repeated until the design is considered stable.

A civil engineering equivalent of this might be to build a bridge, then if it falls in the water, find out what was not strong enough; strengthen it, build a new bridge and see if it stands up to its loads. No one would consider this as a reasonable approach to bridge design. Yet, this trial and error method of designing PCBs is common practice in the electronics industry. I liken it to "pocket knife" engineering, where a peg is whittled and tried in a hole. If it doesn't fit, it is whittled some more and tried again. Finally, the peg fits, or perhaps, it is whittled too much and is too small so a new peg is required.

The trial and error process, using DVT to verify design accuracy, can fairly be labeled worse than no testing at all. This is because if a prototype happens to work successfully under the test conditions, the design might be considered stable. **Nothing could be further from the truth.** All that can be declared using DVT methods is that the units involved in the test process that work meet the design conditions. There is no way to know if the test sample is made of the best case or worst-case parts.

Why is this trial and error method so commonly used? The primary reason is that it takes very little investment in tools and skill to get started. It worked pretty well when speeds were slow enough that the physics of connections didn't make much difference. Even when it started to matter, as long as all competitors used the same methodology the process produced equally unstable designs, so customers could choose only from suppliers with equally poor products. As soon as one competitor broke ranks and invoked a disciplined approach to design, the playing field ceased to be level and the competitor with the more stable products won the day.

Of course, most of the time the switch to a more disciplined design approach is not driven in this manner. The more common reason for the switch is that the cost of being wrong gets too high or the prospect of never arriving at a stable product forces the change. This is the situation that faces electrical engineers everywhere today.

Goals of the Design Process

It is useful to examine what the goals of the design process are. They are:

- Insure the design is logically correct
- Insure software is error free
- Insure the design is stable for all variations in component propagation delays
- Insure the design does not fail from excessive noise generated when circuits switch
- Insure the design functions properly over the full range of component tolerances
- Insure the design passes all emissions and susceptibility tests

Clearly, it is not possible to achieve all these goals using hardware prototypes.

An Alternative to the DVT Style Design Process

The alternative to hardware prototyping is “virtual prototyping”. Virtual prototyping uses a collection of simulation tools to model and analyze the design for the various conditions listed above. Figure 45.2 shows a design flow that achieves these goals. It is a collection of analytical tools in which various aspects of a design are modeled.

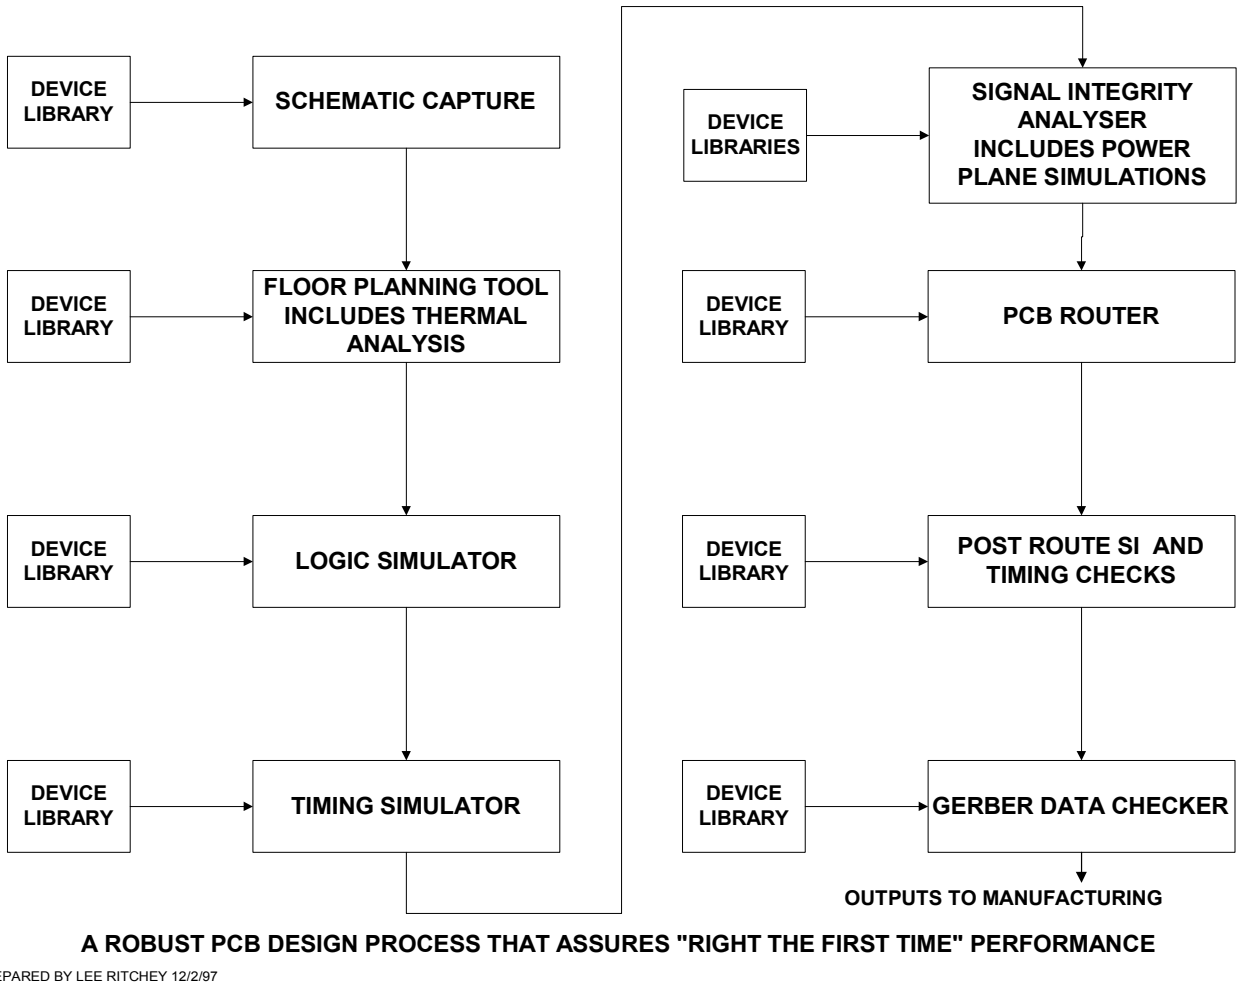

Figure 45.2. A “Virtual Prototyping” Design Process

Figure 45.2 suggests a linear flow through the design process from schematic to finished PCB. In reality, the design process is a series of iterations among all of the tools making tradeoffs until all of the design goals are met.

I often hear engineers say that they tried simulation and did not get good results. As a result, they returned to the trial and error approach. In my experience, it has not been the tools that were at fault. One of two things was wrong with the process.

The most common reason for the simulation process failing is lack of operator skill on the part of the design group. There is the notion that these tools contain the engineering savvy needed to get good results so that unskilled operators can get good results. Sadly, no tool vendor has found a way to put engineering skill into the tools. My counterpart, Dan Murphy, illustrated the problem this way: giving a duffer Tiger Woods' golf clubs will not turn the duffer into a world-class golfer.

Even with skilled tool operators, the process can still fail. The reason for the failures is poorly designed models in the library of each tool. The whole process turns on the quality of the models. It is imperative that the library creation function be done by well-qualified technical people. All too often, library maintenance is seen as a nuisance function and is passed off to a clerical person or technician. This is one of those places where the phrase "garbage in, garbage out" really applies.

Compounding the library maintenance problem is the fact that the library for each tool is unique. Creating a model once and using it in all the tools doesn't work, even with tools that are claimed to be fully integrated.

It is my experience that when properly trained personnel using well-designed models use these tools, the results are excellent and closely represent the final product.

> Library maintenance is a critical part of the design process. Staffing the library creation function with well-qualified personnel is vital to the success of the design process.

Each of the tools in the design process must be carefully chosen so that they are capable of performing at the level required by the design being modeled and so that they interface well with the other tools in the toolset. It is worth spending a little time examining what each tool is expected to do and what characteristics are important.

Schematic Capture Tools

Besides being easy to use, a schematic capture tool that will work in the simulation environment must have some features that are not required with the hardware prototyping process. First, the schematic tool must have a field in the net name or the net list feature that allows the inclusion of the technology label or net class label that was assigned to each net during the design rule creation process. Second, there must be a way to add timing information on a net-by-net basis. This is the only place that it is possible to record this information in a usable way. Finally, the schematic tool must interface easily with the other tools in the design process.

There are several schematic capture packages that have the characteristics necessary to fit into this design process. Those "low end" schematic tools that did not start out with these features are being revised to make them "high speed" capable. Some of them are architecturally incapable of being modified in a satisfactory way. These will have to be replaced.

Floor Planning Tools

In the traditional design process, the design engineer creates the schematic, selects the components and makes a sketch of how these components are to be arranged on the surface of the PCB. This information is handed off to the PCB designer who makes the final arrangement of the components and submits that for review prior to routing the PCB.

As the need to perform a variety of analytical operations on a proposed layout became more critical, the back and forth data exchange between PCB layout and the analytical tools became too cumbersome. Floor planning tools were devised to allow a design engineer to place the components and supply data to the various modeling tools without having to proceed to PCB layout. When the placement satisfied the thermal, timing, routing and other constraints, the design was passed to PCB layout for final routing and preparation of manufacturing documentation. Initially, floor planners were standalone point tools that tied all of the other tools together. As time passed, many schematic capture tool vendors added the floor planning function to their tools.

The characteristics of a good floorplanning tool are that it work easily with all of the other tools in the design process and that it make “what if” analysis changes easy to do.

Logic Simulation Tools

Virtually all modern electronic designs are a combination of logic and software or logic and firmware. Logic and software and firmware are typically designed by different groups. The software developers need the hardware as soon as possible in order to check that the code works properly with the logic. This puts pressure on the hardware development group to hurry their part of their task in order to provide a platform on which to verify the software. All too often, the result is poorly developed hardware that actually hinders the software development process.

A better solution to this conflict is to perform logic simulation in a software model of the logic. The advantage this has is clear. Hardware does not have to be rushed through its design process. If errors are detected in the logic design, they can be corrected by changing the model and resimulating rather than iterating the hardware, a relatively time consuming process.

Simulating a design using a logic model works well as long as the design is all logic and models are available for all parts of the system. When non-logic parts or parts such as microprocessors, for which models are not available, are used there becomes a need for a new solution.

Logic Emulators

Logic emulators are large arrays of programmable logic arrays that can be programmed to represent any logic design. They can be integrated with target PCBs that contain the elements that cannot be emulated, such as microprocessors, memory or analog functions. This combined emulation of the final design can be used to run actual traffic if it is an Internet product or handle voice traffic if it is a cell phone product. These are powerful ways to check a design in actual service.

There are several commercially available logic emulators. Often, it is possible to build custom emulators as has been done by some of the printer manufacturers.

Timing Analysis Tools

Timing analysis involves calculating the time delays in the logic paths through a design in order to determine whether logic operations can be completed in the time available. This involves summing up the delays through the silicon and the delays through the wires on the PCB. Once the PCB has been routed, the delays through the wires will be fixed. The variable part of the delay equation is the maximum and minimum delays that occur in the manufacture of the logic elements.

A timing analysis tool must be capable of importing the predicted wire delays from the floorplanning tool. It must be able to import the post route delays in the finished PCB to allow verification of final timing. It must also be able to identify paths that must be shortened or lengthened when the PCB is routed to make timing budgets work out.

Signal Integrity Tools

When the phrase “signal integrity tool” is used, it usually means the analytical tool that examines the transient behavior of a circuit to determine whether the waveform quality is satisfactory. This tool will also examine routed circuits for potential crosstalk problems. There are two basic kinds of SI tools. One tool uses SPICE models to predict circuit behavior. The other uses IBIS models. Each of these modeling approaches has it strengths and weaknesses. Several articles listed in the bibliography explain the differences.

No matter which modeling tool is used, the objectives are the same. The first place an SI tool is needed is before actual schematic work begins. Proposed logic circuits can be modeled driving the proposed loads and the performance noted. Once the circuit types have been chosen, the schematic completed and the components placed on the surface of the PCB, proposed routing of complex nets can be modeled to determine how best to route them for optimum performance. Finally, nets on the routed PCB can be examined to insure they perform properly.

A common notion of how to perform signal integrity analysis is to do what is called board-level post route analysis. This is usually proposed as an alternative to pre-route analysis followed by rigorous control over routing. There are several problems with this approach. The first one is that the task of obtaining models for every driver and every receiver on the whole PCB is very difficult, if not impossible. If this hurdle is somehow cleared, the next problem is describing in a clear manner to the tool exactly what constitutes a “good “ signal and what constitutes a “bad” signal for all nets. Once this has been done and the actual board-level analysis done, the result will be many signals that must be “fixed.” In a densely routed PCB, there isn’t likely to be much room to do any fixing. So, post route board-level SI analysis as the line of defense is not practical. A much better approach is to develop routing rules that guarantee the SI goals will be met and then insure that the PCB is routed to them.

Power Subsystem Simulators

Tools that allow an engineer to accurately determine how many and what value capacitors are needed in the power subsystem are not new. In some form, they have been available since the late 1990s. Most engineers have not heard of them, because the demands of TTL were easy to meet with rules of thumb. Currently, there are a variety of tools available for this purpose. There are two types. The simplest type allows an engineer to select a variety of bypass capacitor values and quantities and then computes the impedance versus frequency of the combination. This version of the tool does not have the capability of modeling the plane capacitor to predict what the overall power subsystem impedance will be out to a GHz. The more sophisticated version of this tool will take the power plane structure in 3D and compute the effects of the plane capacitor along with that of the discrete bypass capacitors. This type of tool has the capability of drawing current pulses, such as those a data bus draws, from the power subsystem and displaying the actual ripple that will result. Clearly, this is the best approach to solve the power subsystem design problem.

PCB Layout/Routers

The design process described in this book results in routing rules that encompass trace widths, layer usage, trace spacing, pin ordering in a net and trace lengths. These are all physical or geometric constraints. They don't require any signal integrity analysis while the PCB is being routed. Therefore, the most important property of a PCB router for this class of design is to follow the rules imposed in all of the earlier steps. This means the router can interpret the technology labels or net class tags and invoke the control that belongs to that class.

The layout portion of the PCB design tool needs to have utilities that help assess placements for routability and to feed back to the timing analysis tools the predicted net lengths. The remainder of the requirements for the PCB layout tool is to be able to prepare accurate data to be used by the PCB fabricator to build the actual PCB.

CAM/Gerber Data Checking Tools

CAM is the acronym for "computer-aided manufacturing". These tools were first created to allow PCB fabricators to combine the "Gerber Data" or, data from which the film used to manufacture each layer of the PCB is plotted, with the tooling features used to register all of the layers to one another. As these tools evolved, checking features were added that assessed the accuracy of each layer of film. Some of the tools can synthesize a net list from the Gerber data and compare it to the CAD net list. Any disagreements between the two can be resolved before the PCB is built rather than after--a valuable way to insure that prototype PCBs are error free.

CAM stations are usually part of the front-end engineering operation performed by the PCB fabricator. At least two of these tools have been adapted such that is can be incorporated at the end of the PCB design process, making it possible to pre-check the film set prior to sending it out to a fabricator.

System Level Checking

The above tools and design flow check to see that each PCB is designed accurately. However, in a system with multiple PCBs that must work together, it is necessary to add another level of checking. Allowing each PCB to be designed independently and then relying on the integration process to make sure that all of the boards are properly pinned out is hazardous. When speeds were slower it was possible to use "blue wires" to correct pinout mistakes. At the speeds of current logic this is not a choice. If the board-to-board wiring is not correct, the boards have to be rebuilt. The following process describes how to solve this problem.

I have used this process since 1974 to achieve "right the first time" PCBs. Those familiar with the design of complex logic ICs will recognize the process. IC design has employed this method for a very long time and, for that same reason, it is needed for complex PCBs. The cost of being wrong is too high to do otherwise.

The Process

In order to do a system level connectivity check, a system level net list needs to be synthesized from the net lists of all of the PCBs that make up the system. Figure 45.3 illustrates the process for doing this. The net list of each PCB is first created and checked for accuracy. This may involve checking the connectivity against the connectivity that is contained in the logic emulation or simulation against which the software is developed. (Notice that in this case software is developed using a logic model rather than physical hardware.)

Once each PCB net list is deemed to be accurate, its net list is copied as many times as there are instances of that PCB in the system. In the case of multiple copies plugged into different slots in a backplane, multiple instances of the same net will occur in the composite let list. This conflict must be resolved. The most efficient method for doing this is to rename those nets with the net name they mate up with on the backplane. This is a software translation task that needs to be done with some type of editor that remaps nets using net names at the connector interfaces.

All of the translated net lists are combined and checked for several types of errors. Among these are: single pin nets, pins in more than one net, etc. After these basic bookkeeping errors are satisfactorily resolved, checking for correct mapping can be done. The first mapping check is to insure that each bus has the correct number of member nets. The next check is to make sure that the + and – sides of each differential pair remain + and – from end to end and do not get reversed at connectors or where they map onto ICs. In other words, every signal needs to be traced out to insure that it arrives at the proper destination. Power pin assignment should be checked in the same manner to insure that all power pins are connected to the correct rail. Ground pins should be checked as well.

Predicted net lengths can also be derived from this combined database. This is accomplished by adding the predicted Manhattan net length to each net as part of the net list imported into this checking step. All nets whose lengths are important to the proper operation (timing) of the system can be checked at this stage and adjusted if necessary.

Upon completion of these steps, it can be stated that the net list of each PCB is correct and is ready to be routed, but not before this time. No PCBs, including the backplane, should be routed until this step is complete.

After routing of each PCB, the post-route net list must be checked against the net list in this final “master net list” to insure it still matches. Post-route checking can also verify that length requirements have been met.

WHAT TO DO WHEN THERE ARE MULTIPLE VERSIONS OF A PCB, SUCH AS LINE CARDS

PCBs, such as line cards, may exist in several versions. These versions will all plug into the same slot in the backplane. It is not necessary to repeat the above analysis every time a new line card is created. Instead, the above analysis is completed for the first version of the PCB. Once this has been done, the wiring on the card that is common among all cards of a similar type becomes a master template against which to check all later variations of that card type.

If PCBs of this type are not going to be completed until after the backplane and processor cards are done, it is possible to create a partial PCB design that contains all of the nets that pass through the connectors into the backplane. This will allow for verifying that the backplane to daughter card connections are all correct before investing in an expensive backplane build.

TOOLS NEEDED

The first tool needed in this process is a PCB design system that will export net lists in a format that can be used to check length and continuity. Several toolsets will do this.

None of the PCB design tool vendors has a piece of software that does the above checking correctly. Therefore, it is necessary to create a checking tool that will perform all of the operations involved in the global net list check. Several of these have already been written. If one is not readily available, a good perl engineer can create one that will do the job.

Figure 45.3 illustrates the system level checking process.

Summary

The design flow presented in this chapter is complex and uses sophisticated tools and methods. Its aim is to create a design that only needs to be done once. It appears to be expensive and difficult to do. That is a correct perception. Getting a design right the first time is complex and time consuming. This might be used as a reason to stick with the hardware prototyping process. However, if the bookkeeping is done with a balanced hand it will always turn out to cost more than “doing it right”. The reason the hardware prototyping approach seems cheaper is that the cost of time lost is left out of the equation. When it is added in, there is a clear difference.

The real problem is biting the bullet and getting started. When I help an engineering group transition from hardware prototyping to the methods described here, I do it gradually. I start with a single design team and add the SI tool first. Then, I add the other tools as the skill to use them is acquired.

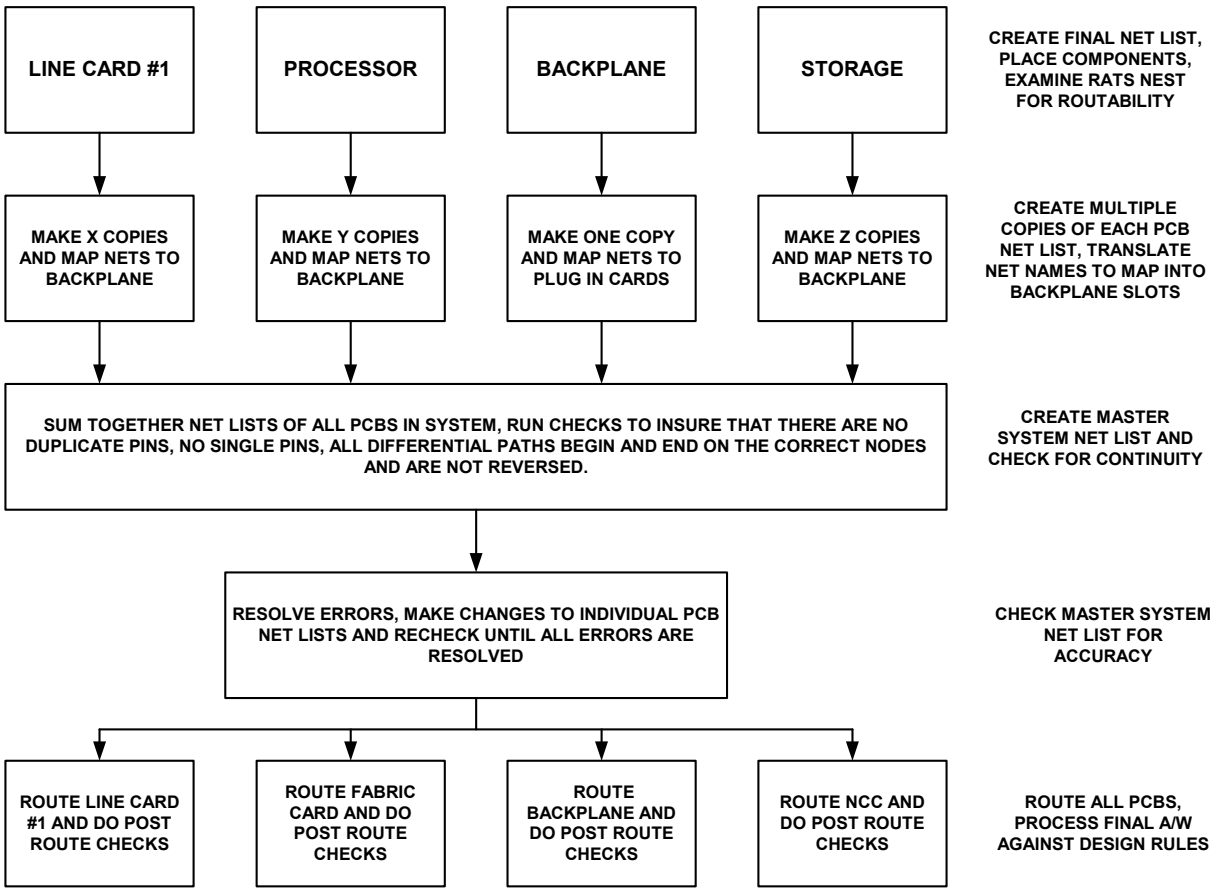

Figure 45.3. System Level Checking Process

CHAPTER 46: PCB ROUTING

The routing of printed circuit boards has ranged from hand layout using tape on Mylar to 100% autorouting of all wires. The methods chosen have been driven by marketplace pressures and product complexity as well as the skill of the PCB designer. The need to complete routing rapidly with large numbers of nets and electrical constraints, as exists in super computer-like products, has driven the development and use of specialized autorouters. At the other end of the complexity scale, PCB layout tools have been optimized to allow hand routing of virtually all wires. Some call this latter method "electronic taping."

This diverse set of design problems has given rise to two distinct routing strategies and types of routers. The two types of routers include maze routers and X-Y based routers. In both cases, the router may be shape based or grid based. When to employ each type of routing method has been the source of considerable confusion. When one router type is applied to the other routing problem, the results can be disappointing. The experience can cause the user to decide autorouters are no good and continue hand routing long after it becomes an uneconomical choice. As an example, using an X-Y router on a two routing layer design yields poor results. Similarly, using maze routing on a design with high pin count BGAs and many routing layers, results in many unrouted wires due to early blocking of routing space by wrong way routing. Understanding the router types and when to use each one is important to optimizing the PCB layout process.

This section will explore the two types of routers. Since most designers started their careers with maze routers and are familiar with how they operate, more time will be spent on the operation and advantages of X-Y routers. This should help the reader understand their benefits and where they fit in the increasingly complex world of PCB layout.

TWO ROUTER TYPES

As stated previously, the two basic routing choices are maze routing and X-Y routing. Figure 46.1 uses a simplified "rats nest" to illustrate the basic difference between the two types of routing strategies. It can be seen from this diagram that X-Y routing involves at least two routing layers with wires traveling in only one direction on each layer. Maze routing allows the wiring of complete nets on a single layer, eliminating the need for layer changing vias. From the drawings in Figure 46.1, the reader can begin to see some advantages and disadvantages of each choice. Understanding the advantages and disadvantages is key to successful, on-time completion of a design.

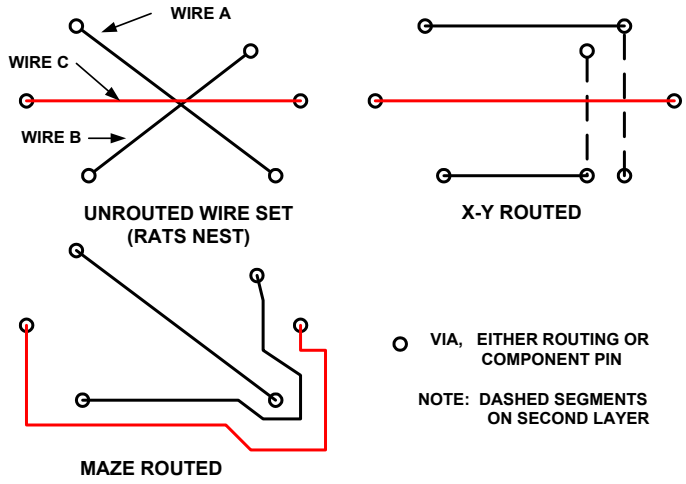

Figure 46.1. Two Routing Methods

When to Use Each Router Type

As might be expected, the two routing methods in Figure 46.1 (maze and X-Y) are best used in certain situations.

Maze routing, the most commonly used method, works best when there are only two routing layers. These layers are nearly always the outer layers of a PCB that contain the component mounting pads. These mounting pads interrupt the routing surface in such a way that making long, straight runs from one pin to another is difficult, if not impossible. This

maze routing works especially well if the pin out of buses is done in such a way that all of the wires in the bus can be "sweep" routed side by side. This is common in the PC world, where the designers of the chip sets insure that the pin outs are optimized to allow this. It works poorly when ICs with large buses are pinned out such that the bus must be inverted from one end to another or when the bus must connect to more than one IC. It also works poorly when a design contains high pin count BGAs or PGAs.

Figure 46.2. An Example of Maze Routing

Figure 46.2 shows an outer layer of a PCB where maze routing has been employed. Routing is done at a wide variety of angles, including both X and Y, in the same layer. This "wrong way" routing is beneficial as long as the wrong way portion of a route does not block the path of another wire as occurred in Figure 46.1. This wrong way routing forces wires routed later to take round about paths. These round about paths are often much longer than they would be if routed more directly.

X-Y routing, the method of choice for high complexity, high performance designs with many high pin count parts, works best when a design requires more than two routing layers to contain all of the wires. As can be seen from Figure 46.1, routing a wire using the X-Y technique does not block the routing surface. As a result, wires added later can be easily routed through each layer. This type of routing lends itself to automatic routing. It also lends itself to implementing large numbers of constraints on nets, such as length matching, layer-to-layer coupling control, and adding length to do timing-related tuning.

Some Definitions

In order to understand the language of routing some definitions are in order.

Manhattan Length--Figure 46.3 illustrates this basic concept in routing.

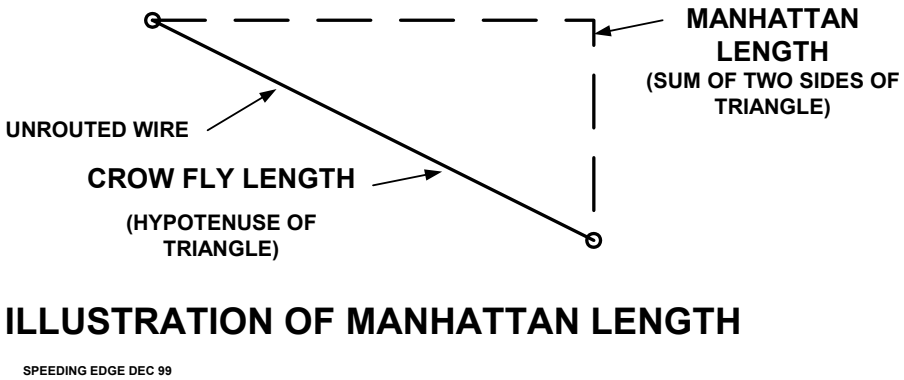

Figure 46.3. Manhattan Length

Manhattan Length is the shortest path that a wire can have when it must be connected using only segments that are confined to either the X or Y axis. Calculating the Manhattan Length is quite simple. It is done by subtracting the X coordinates of the two end points from each other and the Y coordinates of the two end points from each other and then summing up these two dimensions.

Knowing this, it is easy to see how pre-route analytical tools can estimate the length of nets, and, therefore, their time delay, prior to routing. With a well-run autorouter, it is possible to have post route lengths that agree with pre-route predictions to small fractions of a nanosecond. This is one of the more valuable features of X-Y routing.

Clearly, the "Crow Fly Length" is shorter than the Manhattan Length. One might be tempted to use the "Crow Fly Length" to keep the flight time between two points to a minimum. This works for the first few wires, but fails for subsequent wires for the reason shown in Figure 46.1--later wires are forced to be longer, just to make a connection.

Detour Routing is any routing of a net or wire that exceeds the Manhattan Length. In Figure 46.1, both wire 2 and wire 3 have exceeded Manhattan Lengths. In this case, this was forced on those wires by the wrong way routing of wire 1. If timing were dependent on maintaining the Manhattan length predicted at pre-route analysis, this design would fail its timing specifications. When designs are high speed and timing budgets are worked out at the pre-route stage, which is common in very high-speed, high-performance designs, allowing this detour routing may prove fatal.

Net--A collection of wires that connects all of the points or pins in a single circuit.

Wire--The connection between any two adjacent pins in a net.

Segment--A portion of a wire when routed. In Figure 46.1, wire C has only one segment while wires A and B have two segments. It is possible for a wire to be made up of several segments if a number of vias are needed to find open space in the signal layers. However, it is uncommon to see a wire with more than three segments.

Straight Wire--A wire that is pure horizontal or pure vertical. Wire C in Figure 46.1 is a straight wire. These wires usually need to be routed first because the number of possible solutions without detour routing is limited.

Rats Nest--A "crow flies" plot of all of the potential connections between the pins of all of the parts in a printed circuit board. It illustrates the demand for wire space that a particular component placement puts on the wiring surfaces of a proposed PCB stackup. This is a valuable plot because it allows a designer to assess the distribution of wiring in a design. Based on these plots, placements can be adjusted to even out the wire demand and routing strategies can be devised to insure all wires fit into the minimum number of layers.

Routing Via or Turning Via--A via used to change layers or change directions when routing a wire.

Methods for Adding Length or Achieving a Known Time Delay

Often, it is necessary to add length into a net or a wire to achieve a predetermined time delay. There are two ways to do this--serpentine routing and "trombone" routing. These two methods are illustrated in Figure 46.4.

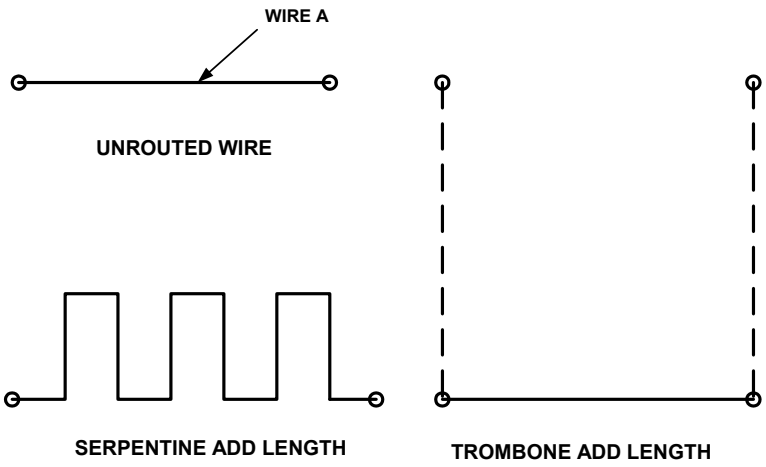

Figure 46.4. Length Tuning Methods

The serpentine method of tuning shown on the lower left hand side of Figure 46.4 is quite popular. It had its origin as an autorouting technique in the early 1980s at a company named Shared Resources. It is a handy way to get extra length into a trace or wire. However, it was abandoned early on because of the effect it had on the routing surfaces. It is easy to see that the serpentine structure blocks the routing surface in both the X and Y directions. As a result, the routing surface is cluttered for all later wires that might need to pass through this part of the board in either the X or Y axis. What is not obvious is the fact that this structure also blocks all potential via sites in its area. This means that the supply of layer changing vias is diminished in the parts of the PCB where this kind of tuning is used. This turns out to be a severe handicap in high layer count designs.

Trombone length tuning, illustrated in the lower right hand side of Figure 46.4, was devised as a solution to the problems presented by serpentine tuning. It can be seen that the added length has been achieved by adding segments or by adding length to segments that were already used to route the wire. This kind of tuning does not block any of the routing surfaces. Further, the number of possible tuning solutions can be quite large by using segments in a variety of directions and in any of several routing layers. Trombone tuning lends itself well to automatic length matching of multiples of wires. All that is needed is to insure that there are enough available via sites to complete the routing. In order to insure space is available for the added segments and length, length tuning is done early in the routing process.

X-Y Routing Affords Many Solutions for Most Wires

Figure 46.5 illustrates how X-Y routing can be used to find a routing solution for a wire in a crowded PCB. Consider the one via solution shown. There are two possible solutions in a two routing layer design, one of which is shown. The second solution is the other side of the rectangle formed by the unrouted wire end points. When two more routing layers are added, the number of possible combinations jumps up to eight. When six routing layers are used, the number of possible solutions goes even higher and this is accomplished using only one routing via.

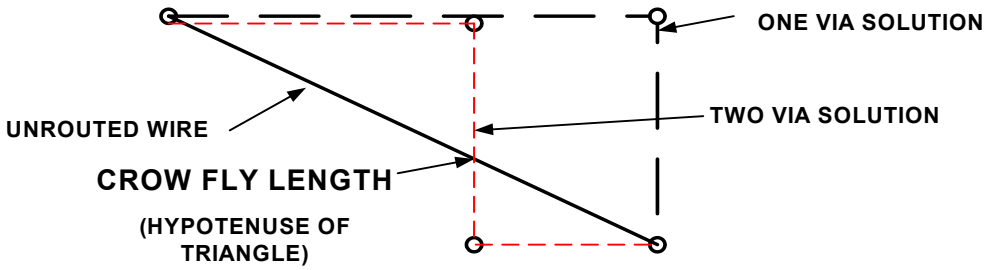

Figure 46.5. Several Routing Solutions Using Routing Vias

Consider the case where two routing vias are used as is shown with the lightly dashed or red line. By moving the via pair left or right, it is possible to find many solutions. If the two vias are located on a horizontal axis, even more solutions are possible. Add into the picture that the three segments can exist on three different routing layers and it can easily be seen that there is an enormous number of possible solutions to this one wire. All that is required is that there be sufficient numbers of available via sites, that wrong way routing doesn't block paths and that the router be capable of performing this kind of searching.

Consider the case where the two ends of the unrouted wire are close to the same X or Y coordinate (the wire is mostly horizontal or vertical). The number of possible one and two vias solutions grows smaller until the case where an unrouted wire is all horizontal or all vertical becoming a "straight wire". In this case, there are no one or two via solutions without exceeding the Manhattan distance (detour routing). Clearly, straight wires must be given first priority when routing a design in order to insure maximum probability of success without detour routing.

Conditions for Successful X-Y Routing

A number of elements are needed in order to insure successful X-Y routing. Among these are: more than two routing layers; surfaces that contain a sufficient quantity of routing via sites; a routing strategy that defines the order in which the router is to proceed through the wire or net list and a router that is capable of performing the searches involved in locating zero, one and two via solutions in a multilayer routing structure.

From the discussions above, it can easily be seen that two routing layers do not contain enough possible solutions to solve routing problems mainly because the surfaces are interrupted by component mounting pads.

Availability of routing via sites is a complicated problem to solve. It takes quite a bit of experience to correctly estimate this number. From our experience of routing hundreds of PCBs with high pin count ICs at Shared Resources, we determined that there needed to be at least one available routing via site for each component pin in the design. These via sites need to be distributed in areas where it is likely that wires will need to make turns or change layers. As designs contain more and more surface mount parts and parts on both sides of a PCB, the number of available via sites is reduced. This gives rise to a need for buried and blind vias. Buried vias require the use of sequential lamination, a very costly solution. Blind vias give rise to the need for additional process steps to build the PCB. Usually, it is more economical to start with a larger PCB in order to allow room for vias than to resort to either of these tactics.

There is a tendency to think of autorouters as programs that look at a design and determine how to proceed to a successful solution. Just mash the button and out comes a routed PCB. I like to think of autorouters much the way that Gene Amdahl once described our first big computer to a reporter who thought that powerful computers of the kind we were building would one day eliminate human thinking. To this reporter, Gene answered, "Maam, what you don't understand is what we have made here is an extremely fast idiot! Humans have to decide what needs doing and tell it what to do." Autorouters are like that. We need to decide the routing strategy and tell the router how to proceed. Failure to do so results in bad results. Like all computing, routing is subject to the "garbage in, garbage out" equation.

Summary

PCB designers have at their disposal a number of well-developed routing tools. Each of these tools was developed to solve a particular set of design problems. Understanding the types of tools and what they are good at is fundamental to a good outcome. What is well established from many years of successful designs is that

autorouters, when properly matched to the problem and properly operated, yield accurate results faster than can ever be achieved with hand routing. The key is in knowing which type of autorouter to use and how and when to use it.

CHAPTER 47: DOCUMENTATION

The PCB design process has three main customers. They include: the printed circuit fabricator, the PCB assembler and the production test organization. Each of these groups needs data that is extracted from the design database and presented in a form that is usable by their equipment. In addition, there is a need to create a document set where the design rules and other pertinent data related to designing the PCB is kept. These documents include:

- An Engineering Drawing on which design data is recorded
- A Fabrication Drawing on which fabrication information is recorded
- A set of design files used to create the film and other fabrication tooling (sometimes called Gerber data)
- An assembly drawing on which assembly information is noted
- A set of assembly files including the bill of material and pick and place information
- A set of test points to be used by test engineering to prepare test fixtures

The following sections describes these elements in order.

Engineering Drawing

The engineering drawing is a drawing that looks much like a fabrication drawing except that it contains design information instead of fabrication information. It is derived from the outline drawing that defined the size and shape of the PCB. On this drawing is recorded:

- The drill table with drill sizes and estimated quantity
- The PCB stackup
- The list of fabrication notes
- The technology table defining routing rules
- A list of all the design files that must be prepared for the fabricator
- The PCB outline drawing.
- A table showing what power supply voltages will be assigned to each power layer
- Other notes as appropriate

Fabrication Drawing

The fabrication drawing is derived from the engineering drawing. The final hole count is added to it. The technology table is dropped. The final PCB layer 1 artwork is added to the outline drawing. (Note: In the past, a drawing showing the location and size of all the drilled holes was drawn on the fabrication drawing. This information is no longer used by fabricators and has lost its value. The layer 1 artwork shows what the finished PCB should look like.)

The fabrication drawing has the stackup on it showing each layer and its name. The design files or Gerber data are also listed by layer. A common problem with these two pieces of information is that the names are not the same between the two. This can cause problems at the fabricator trying to match design files with layers in the stackup. This is a source of layers being stacked wrong. **Make sure that the labels used in the design files are the same as the names on the stackup drawing.**

Figure 47.1 is an example of how a fabrication drawing is laid out.

Table 47.1 is a list of fabrication notes that describe how to build a high layer count, controlled impedance PCB.

Table 47.2 is a list of the design files that are required to fabricate a multilayer PCB. The CAD net list designated IPC-356 is a net list that has the X-Y location of all the component pins appended to it. This net list is used to prepare the bare board test fixture and to compare it with the net list synthesized from the Gerber data.

Assembly Drawing

This drawing contains the information needed to assemble all of the electronic components to the PCB and how to attach all of the mechanical parts. It is often several pages long.

Assembly Files

These files contain the X-Y location of all the components, the bill of material, the code to be loaded into programmable parts and any other information needed by manufacturing to assemble the PCB.

Test Point Files

This is the list of test points designed into the PCB to provide in-circuit test access to the PCB.

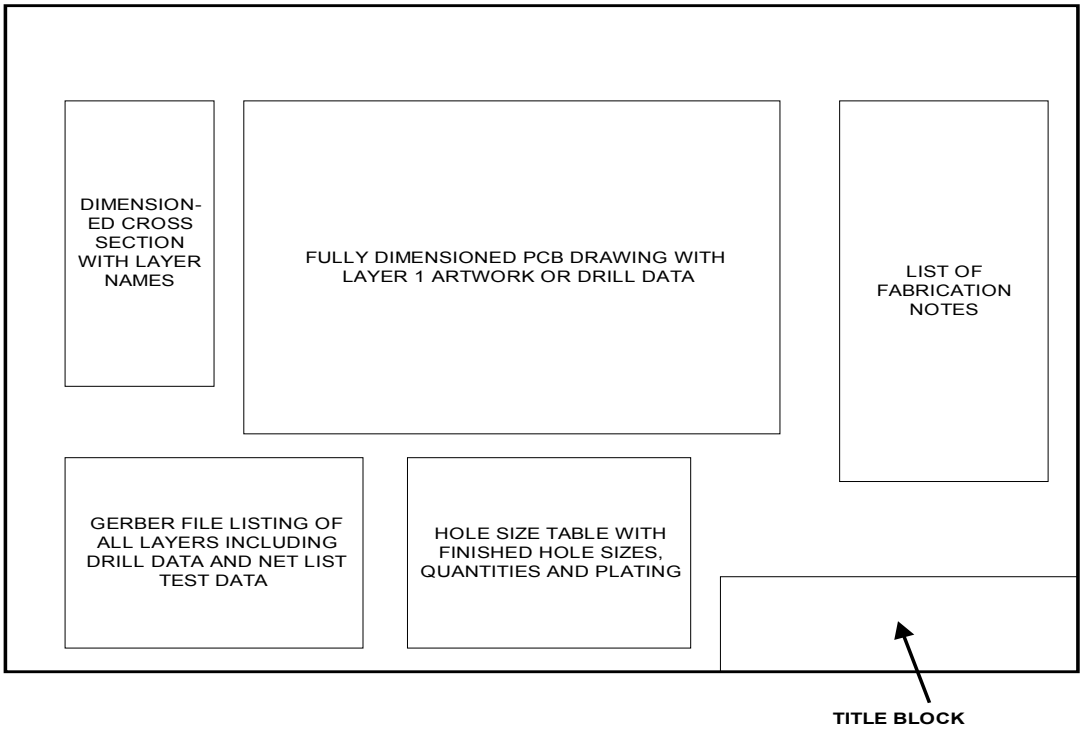

Figure 47.1. A Typical Fabrication Drawing

NOTES: (UNLESS OTHERWISE SPECIFIED)

1. Reference General Specification for Printed Circuit Boards #_____
2. Any deviation from these instructions must be approved in writing by principal or agent.
3. Material: High Tg Fr-4 class laminate with Tg of 170°C or higher. All prepreg and laminates minimum two plies, resin content at least 50%, only glass styles 106, 1080, 2113 and 2116 allowed. Board Lamination: Overall thickness:_____0.110”±0.010” _____.
4. Copper weight: see stackup drawing.
5. Drilling: All holes to be located by X & Y coordinates from NC drill data supplied. See separate Drill Table for drilled hole sizes and quantity. Pad stacks are designed for the drilled hole sizes shown. **Do not change drill sizes.**
6. Minimum annular ring of 2 mils on all signal layers, unless otherwise specified.
7. Copper plating: Hole wall copper plating to be 0.001” minimum.
8. All exposed copper to be plated with electroplated gold over electroplated nickel. 6-10 micro-inches gold over 150-200 micro-inches nickel. (Palladium allowed between nickel and gold)
9. Soldermask: Liquid photo-imageable solder mask to be applied over bare copper or gold/nickel plating unless otherwise specified. Color- green.
10. Legend/Silkscreen: Use nonconductive yellow or white ink.
11. Mark with fabrication revision, supplier ID and date code approximately where shown on bottom or far side.
12. All inside corners and slots shall have 0.062 radius ±0.005” or less.
13. No modification of film without prior authorization.
    For exceptions see notes 16 & 21.
14. Stripes of copper are plotted on each layer on one side of the PCB edge as shown. These “stacking stripes” are intended to be exposed when the PCB is removed from the panel. **Do not** remove/modify stacking stripes.
15. Compare CAD net list to net list generated from Gerber data prior to fabricating board. Resolve differences prior to building board.
16. Non-functional pads to be removed from all inner layers.
17. Conductors: Width and Spacing: Build to Gerber data, however compare widths to Fabrication Drawing Data Set Table and resolve differences prior to fabricating board. GERBER TRACE WIDTHS ARE FINISHED TRACE WIDTHS. Finish width accuracy on inner layers ±0.0005”. Finished width accuracy on outer layers ±0.001”. Fabricator may add manufacturing allowances to trace widths in working film only to accomplish the specified finished trace width.

18. This is a controlled cross section PCB. All fabrication instructions must be complied with in order to assure valid results on completed assemblies. Etch all traces to widths specified in Gerber files. All dielectric thicknesses to be specified ±1 mil. No deviations allowed.
19. First delivery to include diazos of film and a copy of the stackup sheet used to select laminates.
20. When a socket is mounted on PCB, there is to be selective gold plated on the socket pads to a minimum of 20 microinches.
21. Teardrop only on 24 mil and smaller pads at trace exit location. For 24 mil pads, add a 24 mil pad off set from pad center by 2 mils.
22. Thieving allowed on outer layers to insure uniform plating, so long as no thieving lies over the top of traces in the buried signal layer beneath the outer layers.
23. Dimensions of dielectric layers and copper thickness to be measured on one PCB of each lot using stacking stripes. Report to be included with first delivery.
24. Drilled hole true position relative to CAD data not to exceed 0.005".
25. This PCB has laser drilled blind vias between layer 1 and layer 2. See drill file for details.

Table 47.1. Typical Fabrication Notes for a High Layer Count PCB

xxxxpyyy.zip contains:

xxxx-yyy.aXX	Artwork layer 1
	Thru
xxxx-yyy.aXX	Artwork layer XX
xxxx-yyy.smt	Soldermask TOP side
xxxx-yyy.smb	Soldermaske BOTTOM Side
xxxx-yyy.sst	Silkscreen TOP Side
xxxx-yyy.ssb	Silkscreen BOTTOM Side
Apertures.apt	Standard Aperture List
xxxx-yyy.IPC	IPC-356 Data for Bare Board Test
xxxx-yyy.rpp	Drill Report File for Plated Holes
xxxx-yyy.drp	N/C Drill File for Plated Holes
xxxx-yyy.rpn	Drill Report File for Non-Plated Holes
xxxx-yyy.drn	N/C Drill File for Non-Plated Holes
xxxx-yyy.fyi	Contact Person

xxxx-yyy.zip contains:

xxxx-yyy.fbX	Fab drawing in hpgl format (sheet X)
xxxxpyyy.zip	See Above
README.txt	Proprietary Information
contents	Board Number, revision and release status

Table 47.2. Design Files Required to Fabricate a Multilayer PCB

CHAPTER 48: THE IDEAL COMPONENT DATA SHEET

The information about a component that is currently provided by most IC manufacturers does not contain data that is vital to successful high-speed design. Along with all of the maximum and minimum values for input and output voltages that are published, many new pieces of data are needed. In addition to the need for more information, the way some information is taken must change.

Among the data that is provided, but does not properly represent a part, is the rise and fall time as measured. The standard method for measuring output driver performance is by placing a 50 or 60 pF capacitor on the output and measuring the rise and fall time. This does not accurately represent the part's behavior when driving transmission lines. Figure 48.1 shows an output of a 66 MHz Pentium Processor when driving the traditional 60 pF test load and when driving a real 50 ohm transmission line. The delay times listed for this part are going to be off by about 1 nanosecond when actually driving real transmission lines. (The output is faster when driving a transmission line.)

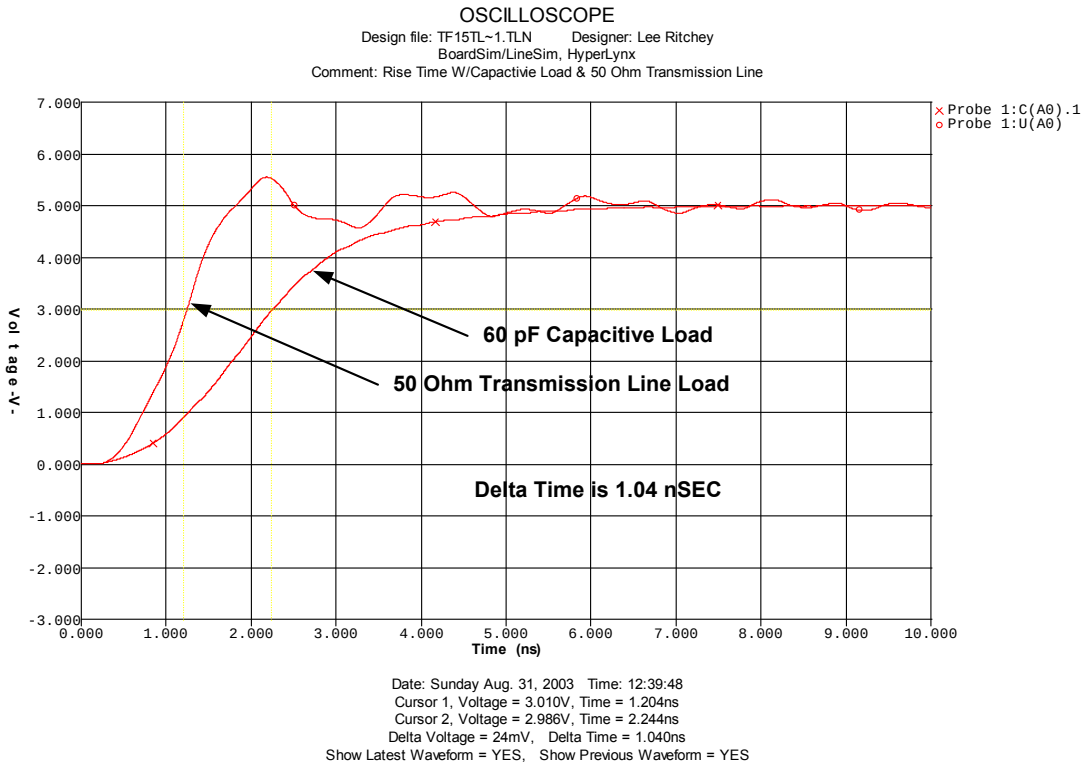

Figure 48.1. The Rise Time of a Part When Measured Driving a 50 Ohm Transmission Line and a 60 pF Capacitor

Why are outputs tested with capacitors as loads? This test method stems from the time when logic was slow enough that the circuits they drove were short enough that they could be treated as lumped circuits. This is no longer true. Virtually all logic is fast enough that the circuits they drive are transmission lines.

> Output circuits should be characterized when driving real transmission lines, not when driving test capacitors.

All of the package parasitics must be supplied along with the other data. Worst case Vcc and Ground bounce numbers need to be included.

When data rates exceed 1 GB/s on an output or an input, the package leads must be characterized as transmission lines of some impedance and length rather than as lumped LC networks. The crosstalk from lead to lead must be specified.

Outputs are often rated as supplying X number of milliamps. This is not a useful specification. It stems from the time when TTL was used and the current load of all the inputs on a net had to be added together. This sum could not exceed the

output current rating of the driver. Again, this assumes the circuits are slow enough that the loads can be treated as lumped. A more useful way to characterize a driver is by specifying its output impedance and the rise and fall time when driving a transmission line.

If the IC package contains plane capacitance, its value should be supplied as should the values of any discrete capacitors added to the IC carrier itself. In addition, if there is significant on-die capacitance as part of the power distribution system, it should be specified.

If the component has a standby and an active mode, the power consumption in each state and the rate of change in current from one state to the other must be listed.

If recommendations are made on how to select and place bypass capacitors, this advice must be accompanied by the amount of plane capacitance that will be needed to achieve a specific ripple goal.

Last, each device data sheet should be accompanied by a full set of IBIS models. For very fast parts, the SPICE models need to be included.

GLOSSARY

1394- A bidirectional differential signaling protocol that is used to connect peripherals to PCs and other processor-based electronic products. www.1394ta.org

1U- Electronics and computing devices come in rack-mounted packages. This includes servers, test instruments, telecommunications components, tape drives and audio and video equipment. Units are bolted to the side frames. The height of a rack-mounted device is specified in a unit (U) measure or rack unit (RU). IU (or 1RU) measures 1.75" from top to bottom.

2U-The height of two rack mounted devices. See 1U above.

20H RULE- A rule of thumb that states that the Vcc plane in a PCB should be recessed in from the ground plane by 20 times the separation between the planes in order to minimize EMI. This rule has no basis in science and has been shown to be false.

3U-The height of three rack mounted devices. See 1U above.

ADDITIVE PROCESS- A method of producing printed circuit boards that begins with bare insulating materials onto which copper is deposited by plating up. Due to the fact that there is no metal foil onto which to electroplate the copper, this is usually an electroless process. It has never been commercially successful due to the poor strength of the electroless copper.

ADMITTANCE- The reciprocal of impedance.

AMPERE (AMP)- A unit of electrical current flow equivalent to the motion of one coulomb of charge or 6.24×10^{18} electrons passing any cross section in one second.

ANALOG- A term that describes a signal that can have any wave shape and any voltage value. Examples are audio signals and control signals for controlling equipment such as rudders on airplanes. There are usually small changes in value in analog signals.

ANALOG GROUND- That point in an analog circuit that serves as the reference point from which all analog voltage measurements are made. This "ground" should always be connected to the "digital ground" directly under a component that contains both analog and digital functions.

ANSI- American National Standards Institute, (www.ansi.org) an organization based in Washington DC that maintains a variety of industry standards.

ANNULAR RING- The ring of copper formed by a pad used to connect a trace to a plated through hole that surrounds that hole after drilling. See the paper "Anatomy of a Drilled Hole" at the end of this book.

ANTIPAD- This term refers to the opening in a copper plane of a PCB. This opening is placed in a plane to allow a signal or component pin to pass through the plane without shorting to it. The term "antipad" derives from the original method used by PCB design systems to create the artwork of plane layers. At one time, the photoplotters used to plot plane layers were only capable of flashing pads and painting lines. As a result, painting a plane layer as a positive piece of artwork was not possible. Therefore, the plane layer was plotted as a negative with the holes flashed as "antipads". This resulted in plane layers that were negative images. Once the plane layer was plotted, it was photographically turned into a positive. Another name for antipad is "clearance pad."

APPE- Allylated Polyphenylene Ether, is a resin system formulation produced by Nelco, that is used to manufacture high performance PCBs. The main benefit of this material is a lower loss tangent.

ARTWORK- The plotted film that represents the copper patterns to be etched into each layer of a PCB. This artwork also includes the silk screen or legend and the solder mask. It can also be plotted on paper to provide a "hardcopy" for review and record keeping.

ASIC- Application Specific Integrated Circuit, also referred to as custom silicon, is an integrated circuit custom designed to perform a specific logic function. Often called custom silicon. As ASICs are custom made for individual companies, they are typically not available on the open market.

ASPECT RATIO- The ratio of a drilled hole's length to its diameter. The higher the aspect ratio, the more difficult it is to uniformly plate copper in a hole. Aspect ratios greater than 6:1 are not considered candidates for volume production**.**

ASSP- Application Specific Standard Product, a standard product that is designed to perform a complete function. A microprocessor or serdes (serializer-deserializer) is an example of such a component.

ATE- Automatic Test Equipment, ATE are computer-driven testers adapted to test specific products. The most common form of this product is a PCB tester or an IC tester.

BACKPLANE- A printed circuit board containing many connectors into which are plugged several "daughter" boards. Backplanes usually contain no active circuits.

BGA- Ball Grid Array, This is a type of component package that has all of its contacts on the bottom of the package. The contacts are small balls of solder or lead that are attached to the bottom of the package. These balls are soldered to pads on the surface of a PCB. Originally, BGA packages were created to reduce the incidence of solder defects that occur with high lead count quad flat pack (QFP) surface mount packages. A desirable byproduct of this package change is potentially much lower lead inductance, especially in the power and ground paths. This low inductance is essential for gigabit and higher data rate components.

BISMALAMINE TRIAZINE (BT)- A resin system that can withstand high temperatures used to manufacture PCBs. This resin system is used with glass cloth to produce multilayer PCBs. See the materials section of this book for properties. BT is not widely used due to the difficulty in drilling and fabricating the material.

BIT- One segment of a data "word" where information is conveyed in the binary or base two data format as a "1" or a "0". A "word" usually contains four or eight bits.

BLACK OXIDE- A treatment performed on the copper portion of inner layers to improve the quality of the bond between the copper and the resins during the PCB lamination process. This etching operation creates tiny peaks and valleys on the surface of the copper.

BLIND VIA- A via that starts on one side of a PCB, but does not pass all the way through a PCB or BGA. Primarily, blind vias are used to connect to internal layers of a PCB while leaving the opposite surface clear for mounting other components on a double-sided surface mount PCB. Blind vias are sometimes incorrectly referred to as microvias.

BOUNDARY SCAN- A method used to test an assembled PCB when it is not possible to access the nets on the PCB with probes. The integrated circuits on the PCB are designed with extra circuitry that allows scanning the state of each input and output pin from a set of special test pins.

B-STAGE- A term that refers to a piece of resin coated cloth where the resin is partially cured. The resin is cured only enough to make it non-sticky. This material is also called prepreg. It is used as the "glue" layers during PCB lamination.

BT- Bismalamine Triazine, see above.

BTL- Bipolar Transistor Logic, This is a logic family designed to drive low impedance data buses at high data rates. Bus impedances as low as 22 ohms can be driven at clock rates as high as 225 MHz. Signal swings are less than one-volt peak to peak. The outputs are open collector NPN transistors.

BREAKOUT- A term that describes the condition where a drilled hole is off center and "breaks out" of the edge of the pad that has been plotted to make a connection to the copper plated in the hole. If a trace enters the pad on the side where the hole breaks out of the pad, the connection between the trace and the plating in the hole will consist only of the cross sectional area of the trace. When the PCB is soldered, this connection often breaks due to stresses induced during soldering and an unreliable connection results. Designs should be done with capture pads that are large enough that this condition never occurs.

BUFFER- A driver circuit with a very low output impedance capable of driving transmission lines at very high frequencies or high data rates.

BUILD UP- A method of creating a muyltilayer PCB that involves adding successive layers by laminating, then plating and then laminating again. Often used as a way to create blind vias on very dense PCB such as those used in cell phones.

BURIED CAPACITANCE- (BC or ZBC®) A material, patented by Zycon, (now Sanmina), that yields relatively high capacitance between adjacent plane layers. The principal claim of the patent is a 2-mil thick dielectric bonded to two sheets of copper foil. The side of the foil facing the dielectric has a very smooth finish to insure the two copper layers don't short to each other through the thin laminate.

BURIED VIA- A via or plated hole that begins on an inner layer and ends on another inner layer without ever reaching an outer layer. This type of via is formed by creating a sub PCB that includes only the layers pierced by the buried via. This sub PCB is drilled, etched and plated as though it is a finished PCB. This sub PCB is then combined with the rest of the PCB layers, laminated and finished to produce a normal PCB.

BUTT CONNECTION- A connection between a trace and a plated hole or via that consists of only the end-on cross section of the trace. This type of connection is often fractured during the soldering process, yielding an unreliable connection. Pad stacks should be designed such that this never occurs.

BUTTER COAT- The term used to refer to an extra layer of resin applied to a piece of laminate that provides more resin for filling voids in adjacent layers.

BYPASS CAPACITOR- A capacitor placed on a PCB between the rails of a power supply to provide local charge to switching circuits. A bypass capacitor is also called a decoupling capacitor.

C4- Controlled Collapse Chip Connection. This is a method for attaching an integrated circuit die to a substrate by forming very small balls or bumps on each IC contact. The IC is then inverted so the balls or bumps make contact with pads on the substrate. This is the alternative to wire bonding the IC to the lead frame. This is the lowest inductance method of connecting an IC to a PCB or substrate.

CBGA- Ceramic BGA. This is a ball grid array package that uses a multilayer ceramic substrate to connect the IC die to the contacts on a PCB.

CAD- Computer Aided Design. This is a method of designing PCBs or mechanical assemblies using computer drafting tools rather than laying them out by hand.

CAM- Computer Aided Manufacturing or Content-Addressable Memory. Computer aided manufacturing refers to the use of computer-based design tools to supply data to manufacturing tools such as lathes and punches. Content addressable memory refers to a memory IC that is organized such that it places data in memory locations based on its makeup.

CAPACITANCE (C)- The property of a system of conductors and dielectrics that permits the storage of electrically separated charges when potential differences exist between the conductors. The basic unit of measure is the Farad.

CAPACITOR- A component made up of two conductive plates separated by an insulator. The function of a capacitor is to store energy in the form of an electrostatic field. Capacitors are also used to provide a DC block between two circuits while passing AC signals.

CAPTURE PAD- A pad or patch of copper placed on a signal or plane layer for the purpose of connecting a signal to a drilled and plated hole. This pad is said to “capture” the drilled and plated hole. Capture pads are designed larger than the drilled hole to allow for the tolerance build up that is a normal part of the PCB manufacturing process.

CE- The mark put on products that indicates the product complies with the safety and other requirements of the European Union or EU.

CGA- Column Grid Array. This is a BGA with columns replacing the solder balls. CGAs provide thermal strain relief between the ceramic packages and the laminates of PCBs.

CHARACTERISTIC IMPEDANCE- The ratio of voltage to current in a propagated wave, i.e, the impedance offered at any point on a line. In PCBs, impedance is determined by the dielectric constant of the insulating material, the width and thickness of a trace and the distance the trace is away from the plane(s) over which it travels.

CHASSIS- The mechanical body of an electronic product. The chassis normally supports card cages, card guides, fans, power supplies and other major components of the product. Incidentally, the chassis can form part of the Faraday Cage used to contain EMI.

CHASSIS GROUND- A name often given to the frame of a product. This term is erroneously used to describe the Faraday Cage or shield that surrounds a product for the purpose of containing EMI or emissions. Parts of the chassis may coincidentally be parts of the Faraday cage, but the “chassis ground” plays no role in controlling EMI. Often, the chassis of a product is connected to the green wire of the power cord. The purpose of this connection is to provide safety to users by insuring that the product case is “grounded” to the earth through the green wire. This connection also plays no role in containing EMI.

CHECKPLOT- A paper plot of the artwork that will be used to create the layers of a PCB. These plots are used to examine the artwork for errors prior to releasing them to fabrication.

CIRCUIT- The collection of conductors that connects points together to form a single signal.

CIRCUIT LAYER- A layer in a PCB that contains circuits or traces. Usually called a signal layer.

CIRCULAR MIL- A measure of the cross sectional area of a conductor or wire. One circular mil is the area of a circle one thousandth of an inch in diameter. This area is 31.4 square microinches.

CISPR A- The EMI standard that commercial products must meet for the European Union or EU.

CISPR B- The EMI standard that consumer products must meet for the European Union or EU.

CLEARANCE PAD OR HOLE- The hole etched in a plane layer of a PCB through which a hole is drilled and, often but not always, plated. It is called a clearance pad because plane layers in PCBs are often designed as the negative of the final layer. This technique is used because it is easier to draw the round "pads" that are to be holes than it is to draw the plane that surrounds the holes. This pad is made large enough that the plated hole clears the copper in the plane by an amount sufficient to provide insulation and a tolerance allowance for manufacturing.

CML- Common Mode Logic. This is a logic family whose output is a fixed current that is tolerant of substantial common mode noise, usually ground offsets. The desired signal voltage is developed at the load by passing this current through a load resistor. LVDS is an example of CML.

CMOS- Complimentary MOS or metal oxide semiconductor. This is a logic device made by using N channel transistors to connect signal lines to ground and P channel transistors to connect those signal lines to Vcc.

COB- Chip On Board, This is a method of attaching an IC die directly to a PCB without using a package. Usually, the die is glued directly to the PCB and wire bonds are used to make connections between the two.

COEFFICIENT OF THERMAL EXPANSION (CTE)- The linear change in dimension as a function of temperature.

COLUMN GRID ARRAY- See CGA.

COMMON MODE- Used to describe a property that is common to two or more signals. This property usually appears as impedance or noise coupling. If it is a noise, it is the common mode coupled from an aggressor and the noise appears equal in size in all victims.

COMPONENT HOLE- A hole drilled through a PCB into which a lead of a component is soldered.

COMPONENT SIDE- The side of a PCB on which components are mounted. The opposite side is the solder side. This term does not apply when a PCB has components mounted on both sides.

COMPONENT PLACEMENT GRID- The grid system used to place components on the surface of a PCB.

CONDUCTANCE (G)- Conductance is the reciprocal of resistance. Conductance is measured in units of mhos and expressed as a factor of 1/R.

CONDUCTIVITY (σ)- The property of a conductor that expresses how well electrical current will pass through it. This property is the reciprocal of resistivity. Conductivity is measured in units defined by current density divided by voltage per meter.

CONDUCTED EMISSIONS- Emissions that escape a product and are conducted on the wires or cables exiting the product. These emissions are normally measured in the frequency range of 150 KHz to 30 MHz.

CONDUCTOR- A single conductive structure in a conductive layer.

CONFORMAL COATING- An insulating coating applied to the surface of a PCB and the components mounted on it. The purpose of conformal coating is to waterproof the assembly. Conformal coating is often required on PCBs made from polyimide to prevent leakage failures due to excessive moisture absorption.

COULOMB- This is a unit of charge equal to 6.25 x 10^{18} electrons. This unit of charge is named after George Coulomb, an early experimenter of electrostatic charge phenomena.

COULOMB BUCKET- This is my term for a decoupling capacitor. This is an appropriate name given that these capacitors store charge that is used to support switching events.

COUPLING- Another name for crosstalk. Coupling is the unwanted interaction between two signal lines that travel side by side or one over the top of the other on a PCB.

CPCI- Compact PCI. This is a signaling protocol and hardware specification based on the PCI bus protocol developed for small computers that use microprocessors. This protocol is used widely in the instrumentation market.

CSA- Canadian Standards Authority. This is the agency in Canada that regulates emissions, safety and other features of products sold in Canada.

CQFP- Ceramic Quad Flat Pack. This is a surface mount IC package with leads exiting on all four sides.

CRITICAL LENGTH- The length that two transmission lines must run in parallel to achieve maximum or worst-case backward crosstalk. Occasionally, this term is misused to describe Transition Electrical Length.

CROSSTALK- The unwanted interaction between signal wires or traces traveling in parallel. Also called coupling.

CROSS HATCHING- This is a process of dividing large areas of copper into a cross-hatched pattern. This practice was used in the early years of multilayer PCB manufacturing to solve the problem of delamination of resin from power planes in a PCB. This problem has long been solved and there is now no good reason to cross hatch.

CSP- Chip Scale Packaging. These are very small IC packages with mounting footprints similar in size to the IC die itself.

CURRENT (I)- The flow of charges in a conductor. These charges are usually electrons moving due to a voltage difference impressed along the length of the conductor. The unit of measure of current is expressed in Amperes.

CURRENT MODE- A transmission line driver that delivers a constant current rather than a constant voltage. The output voltage is a function of the load impedance. For example, a current mode driver that has an output current of 4 milliamps will develop a signal swing of 200 millivolts across a 50-ohm transmission line. A pair of equal and oppositely changing differential current mode outputs, such as LVDS, each delivering 4 milliamps into a pair of 50 ohm terminations will develop a 400 millivolt differential signal. The output impedance of an ideal current mode driver is infinity. Real current mode drivers have output impedances ranging in the few hundreds of ohms.

CURRENT SOURCE- A signal source that delivers a constant current into a load. Current does not change as the size or impedance of the load changes. The alternative is a voltage source whose output voltage does not change as the load impedance changes. A current source has an infinite output impedance.

CYANATE ESTER (CE)- A polymer resin system used in the manufacture of PCBs. CE is used with glass cloth to create a laminate. The main benefit of CE is its ability to withstand high temperatures. The downside of CE is the fact that it absorbs excessive amounts of moisture that can cause leakage problems.

DAUGHTER BOARD- A PCB that plugs into another PCB, usually a backplane.

DECOUPLING CAPACITOR- A capacitor placed between the ground and Vcc rails of a power supply to “decouple” noise on the supply. Another name is bypass capacitor. These capacitors don’t decouple noise; they serve as sources of charge that support switching. Without these local sources of charge, the power supply voltage drops as charge is drawn from it to support switching, resulting in ripple voltage.

DELAMINATION- Separation of the resins in a PCB from the copper layers or planes. Delamination most often occurs because the process used to treat the inner layer of copper so that the resins form tight bonds with the copper during lamination, is unsatisfactory. Black oxide treatment is a common way to achieve this bond.

DENDRITIC GROWTH- A metallurgical process that causes narrow slivers of a metal, such as tin or silver, to form whiskers that extend out from a conductive surface. These whiskers can span a gap and cause a short circuit or a leakage path.

DES- Develop Etch Strip. This is a process in PCB manufacture that results in the creation of inner layers. After photosensitive film has been applied to both sides of a piece of copper clad laminate, the pattern is "printed" on the film and it is developed. After this step, the unwanted film is stripped off exposing the underlying copper. The unwanted copper is etched away and the etch resist film is stripped off the copper that forms the traces and planes on the inner layer. DES is often a continuous operation done by a single machine.

DESMEAR- A process that removes the resin that smears across the exposed edges of the copper inner layers during drilling. Failing to desmear a PCB results in open circuits in the plated through holes.

DETOUR ROUTING- Routing traces in a PCB in a roundabout manner such that the connection is longer than the Manhattan Distance, (the shortest distance between two points that can be achieved using only X and Y routing). This form of routing often causes timing problems when some members of a bus are routed minimum distance and others are detour routed. See the routing section of this book for illustrations of this.

DFF- D flip flop, the most common type of flip flop used to design logic circuits. The D stands for the data input to the flip flop.

DIAZO- A method for making working film for the manufacture of PCBs. The image is exposed in the same process as for standard silver halide films. The image is developed using ammonia rather than other chemicals. The main value of diazo film is that the image is translucent, allowing PCB manufacturers to align solder mask artwork to the etched images on the surface layers of PCBs.

DIELECTRIC CONSTANT- A property of a vacuum that expresses the effect that the vacuum will have on the velocity of an electromagnetic wave traveling though it. Dielectric constant also expresses the effect that the vacuum will have on the capacitance that exists between the two conductors. The dielectric constants of materials, other than a vacuum, are compared to a vacuum. This comparison results in a relative dielectric constant, e_r, that expresses the effects of these materials on velocity and capacitance as compared to a vacuum.

DIELECTRIC BREAKDOWN- The failure of an insulating material to isolate two conductors from each other. Dielectric breakdown is usually caused by applying an excessively large voltage that then creates an arc.

DIELECTRIC BREAKDOWN (DBV)- The voltage per unit thickness of an insulating material at which it fails to insulate. DBV is commonly expressed in volts per mil or millimeter.

DIELECTRIC LOSSES- All dielectrics or insulators are made up of molecules that have some polarization. When these molecules are excited with a varying electromagnetic field (RF), these molecules vibrate. There is some friction associated with this vibration that results in energy lost in the dielectric. This loss is referred to as dielectric loss. It is often frequency sensitive.

DIELECTRIC STRENGTH- Another way to state dielectric breakdown voltage.

DIFFERENTIAL IMPEDANCE- The impedance measured between the ends of a pair of transmission lined with respect to each other rather than to "ground". Differential impedance is often thought to be important to differential signaling, but, in fact, it is not.

DIFFERENTIAL MODE- Used to describe the difference between two or more conductors that should be the same, such as differential impedance or differential coupling. Differential coupling is coupling from an aggressor signal into a pair of differential signals such that more noised is coupled into one member than the other. When differential coupling occurs in a PCB it is always due to the shape of the electromagnetic field emanating from the aggressor.

DIFFERENTIAL PAIR- A pair of conductors on which travel two identical, but equal and opposite polarity signals travel. These signals are tightly timed to each other.

DIGITAL- A term that describes a signal that has only a few allowed states or voltage values, normally two. Information is contained in a pattern of "bits" made up of the logic levels.

DIGITAL GROUND- The reference terminal of a logic power supply. Usually, this is the most negative terminal of the power supply.

DIMENSIONAL STABILITY- A measure of the ability of an object to retain its original dimensions over a range of temperatures, stresses or other conditions. In PCBs, this usually refers to how well a PCB retains its dimensions during lamination or soldering process.

DIN- Deutsches Institut fer Normung, a standards institute based in Germany. As it pertains to connectors and back planes, this institute maintains a set of standards for 96-pin 2mm pitch connectors know as DIN connectors. DIN plays a major role in the ISO standards. The standard that governs 2mm DIN connectors is IEEE 1301.1.

DIP- Dual In Line Package. This is an IC package with through hole leads that are arranged along two sides of the IC package. The lead pitch is most commonly 0.10 inch or 100 mils (2.54mm). This package is largely obsolete.

DIPOLE ANTENNA- An antenna made up of two halves with the signal being supplied in the middle. A dipole antenna does not require a "ground" plane to function. Two PCBs joined together by a connector create such an antenna and often produce high EMI.

DISPERSION- Dispersion is phase distortion that results from different frequencies in a signal traveling at different velocities. It occurs because the dielectric constant of the dielectric is not constant with the frequency. In RF signals this manifests itself as phase distortion. In logic signals it manifests itself as fast edge pulses that "spread" out.

DISSIPATION FACTOR, γ- The property of an insulator that causes it to absorb some of the energy from an electromagnetic field passing through it. It is also sometimes referred to as loss tangent.

DOUBLE SIDED ASSEMBLY- A PCB with components mounted on both sides.

DRAM- Dynamic Random Access Memory. This is a memory IC that retains data by placing small charges on capacitors in the memory cells. This charge slowly leaks off and must be replenished periodically with a "refresh" cycle. This leaking off and refreshing process gives rise to the name dynamic RAM. This type of memory cell requires fewer components to create, resulting in much higher densities at lower cost.

E-GLASS- A low alkali, lime, alumina, borosilicate glass, noted for its good electrical properties. Its formulation can be found in a variety of IPC documents including IPC-T-50D.

ECL- Emitter Coupled Logic. The original "high speed" logic family. Logic signals are coupled through the circuits using pairs of transistors connected at their emitters to a current source. This logic family operates with all transistors always in the linear mode in order to achieve fast turn on and turn off. A byproduct of this circuit style is high power consumption. The outputs of ECL circuits are open emitter NPN transistors that require an emitter pull down resistor to complete the output circuit. ECL has been made obsolete by high-speed CMOS circuits.

EDGE RATE- The rate of change of a logic switching edge. It is measured in volts per nanosecond or a similar measurement system. Edge rate is often confused with rise and fall time, which is the time required for a signal to switch between two voltage levels.

EEPROM- Electronically Erasable Programmable Read Only Memory, This is a memory part that can be written and rewritten. Information in memory is not lost when power is removed. EEPROMs are used for micro-code and programs that must be permanently stored in memory.

EIA- Electronic Industries Alliance, (www.eia.org) is an organization based in Arlington, VA, USA, that participates in the creation and maintenance of standards for the electronics industry.

EPROM- Erasable Programmable Read Only Memory. EPROM is similar to EEPROM, but erasure is done by external means such as a UV light.

EFFECTIVE DIELECTRIC CONSTANT – The dielectric constant seen by an electromagnetic wave traveling on a transmission line with mixed dielectrics. For example, surface microstrip transmission lines have air above and laminate below them. The effective dielectric constant of such a line will be between that of the two materials used.

ELECTRODEPOSITION- The deposition of a conductive material, usually copper, from a plating solution by the application of an electric current.

ELECTROLESS DEPOSITION- A chemical operation that does not rely on electric current to take place. In PCB manufacture, this is usually a plating operation. Copper is deposited on the epoxy in drilled holes with an electroless process. Gold, silver, nickel and tin deposited on copper surface pads after soldermask can be applied using an electroless plating operation.

ELECTROLYTIC- A process that involves electric current. When used to describe a capacitor, electrolytic refers to a capacitor that has a liquid or electrolyte as part of the dielectric or insulating layer. These insulating layers are very thin, resulting in very large value capacitors in very small packages.

ELECTROSTATIC- Describes an electric field that is not changing.

ELECTROMAGNETIC- Describes a compound field made up of an electric field and a magnetic field. This compound field is how energy is transmitted from one place to another. Light waves and microwaves are examples of electromagnetic fields as are the fields traveling on a transmission line when energy is in motion.

EMC- ElectroMagnetic Compatibility. A measure of how well multiple electronic products work together without interfering with each other via their radiated or conducted electromagnetic fields.

EMI- ElectroMagnetic Interference. EMI consists of emissions in the form of electromagnetic energy that escapes from a product in the form of conducted or radiated emissions. Radiated emissions are normally measured in the frequency range of 30 MHz to 1 GHz.

ENIG- Electroless Nickel over Immersion Gold. This is a non-corrosive plating applied to the exposed copper of a PCB that has soldermask over bare copper (SMOBC). This is done to preserve solderability of the copper while maintaining a flat surface onto which solder paste can be screened in a uniform manner.

ENTEC®- A family of organic solder protection coatings manufactured by Enthone. These coatings are applied over the exposed copper pads of a PCB that has had solder mask applied over bare copper. The coating functions as a corrosion barrier for the copper until such time as the components are soldered onto the pads. During soldering, the coating functions as a flux, aiding the soldering process.

E_r (e_r)- Relative Dielectric Constant. This is a measure of how an insulator or dielectric affects the capacitance of a pair of conductors separated by the dielectric, as compared to the same conductors separated by a vacuum. E_r is often determined by measuring the capacitance between the two conductors with and without the dielectric.

ESD- ElectroStatic Discharge. This is the energy transferred from one conductive body to another through an insulator, such as air, due to the very high voltage difference between the two bodies that results in a spark jumping the gap. ESD can cause either functional or fatal failures,.

ESL- Equivalent Series Inductance. This is the parasitic inductance present in every component due to the fact that its length is longer than zero. ESL is a major limiting factor in the performance of decoupling or bypass capacitors.

ESR- Equivalent Series Resistance. This is the parasitic resistance of all components. It is a property of the conductors from which they are made.

ETCHBACK- An operation performed on a PCB after the holes have been drilled and before they are plated. A chemical or electrostatic process is used to erode away some of the epoxy in the holes exposing the copper of the inner layers. This process is used to clean epoxy smeared on the copper in the holes during drilling and as a way to “anchor” the plating in the hole to each inner layer.

ETCHING- A process wherein a printed pattern is formed by chemical, or chemical and electrolytic removal of the unwanted portion of conductive material bonded to an insulating base.

ETCH RESIST- A photosensitive coating that is placed on the copper foil of the inner layers of a PCB. The resist is photo exposed so that when the resist is developed it protects the copper that is to remain on the inner layers after etching.

ETHERNET- The collection of hubs, routers, switches, cables, fibers and other items based on the Ethernet protocol that allows communication among users.

EYE DIAGRAM- An oscilloscope display of a large stream of data bits of a differential data path. This display is scaled such that it shows a single data period or unit interval (UI). The purpose of this type of display is to determine if the data path being measured has sufficient signal quality to accurately resolve all of the bits in a data stream. If the eye is “open” (meaning that the logic levels reach proper low and high levels for a long enough time to detect the proper logic state), the link is said to be “robust”.

FALL TIME- The time it takes a falling logic edge to traverse from the 90% voltage level to the 10% voltage level. In special instances such as , GaAs and ECL components, the levels are 20% and 80%, respectively.

FARAD- The primary unit of capacitance, a charge storing electrical element. Named after Michael Faraday, an early researcher into electromagneitic phenomena. A farad is defined as that value of capacitance in which one coulomb produces a 1-volt potential difference across its terminals.

FARADAY CAGE- The conductive container that surrounds a product for the purpose of keeping electromagnetic fields that radiate from components inside the product. This is the most effective way there is to keep these electromagnetic fields from escaping and becoming EMI. Sometimes, a Faraday cage is incorrectly called “Chassis” ground.

FCBGA- Flip Chip Ball Grid Array Package. This is a BGA with the silicon chip mounted face down on the BGA frame using C4 bumps instead of wire bonds. This process is done to lower the inductance of connections to the IC, primarily the power leads.

FCC- Federal Communications Commission. The US government organization that regulates the use of the radio frequency spectrum. This organization maintains the EMI specifications that electronic products sold in the USA must meet.

FCC Rule 15- The FCC rule that defines what EMI standards electronic products must meet for sale in the USA.

FERRITE- A ferromagnetic material made by combining ferric oxide with another metallic oxide. The resulting material is conductive and acts like an inductor when current is passed through it.

FERRITE BEAD- A torroidal-shaped component that uses ferrite as its main material. These beads are threaded over wires and increase the inductance of the wire. Ferrite beads are used to block conducted EMI from escaping on unshielded wires.

FERROMAGNETIC- The property of a material, usually a metal, whose relative permittivity is greater than one and depends on the magnetizing force. A ferromagnetic material usually has relatively high values of relative permittivity and exhibits hysteresis. Ferromagnetic materials are attracted to magnets. Most ferromagnetic materials contain iron, nickel or cobalt.

FFT- Fast Fourier Transform. This is a mathematical operation that converts time domain waveforms to frequency domain waveforms or frequency spectra.

FIDUCIAL- A target placed in a precise location on the surface of a PCB to allow accurate location of a drill of component placement machine. Fiducials are usually round pads that are exposed when the PCB is finished rather than covered with solder mask. When blind vias are used, it is common to place blind vias on layer two to allow alignment of the laser drill to the pads on the inner layer.

FIREWIRE 1394- See 1394.

FLAT PACK- A component package having two rows of leads extending from its sides that are parallel to its base. Normally surface mounted.

FLEXIBLE PRINTED CIRCUIT- A patterned arrangement of printed circuits and components utilizing flexible base materials with or without flexible cover layers.

FLIPCHIP- An IC chip or die that is inverted and mounted with its active surface toward an IC package or directly onto a PCB or flexible circuit. Connections are made between the contacts on the IC and the PCB or substrate using bumps plated onto the IC die contacts.

FLUX- A chemically active compound which, when heated, removes minor surface oxidation, minimizes oxidation of the base metal and promotes the formation of an intermetallic layer between solder and the base metal.

FMM- Forget Mostly Memory. A memory that often looses its contents. A type of memory accidentally designed when signal integrity issues have not been carefully accounted for. This also describes the author’s memory from time to time.

FOIL- **in PCBs, this refers to the copper foil used to create the conductive layers of a PCB. Typically, foils are rated in ounces. The ounce rating derives from the gold foil business and is based on spreading an ounce of a given metal over one square foot of area. A one-ounce thick copper foil is 1.4 mils or 35 microns thick.**

FORCED SEQUENCING- Arranging the points in a network by fixing the way in which the points are to be connected. This is accomplished by adding a number from 1 to n to each point. The PCB routing system then follows this sequence when making the physical connections in the PCB.

FOURIER ANALYSIS- A mathematical operation on a voltage waveform that converts it from the time domain to the frequency domain or the reverse. When converting from the time domain to the frequency domain, it is possible to determine what frequencies are involved in creating a voltage waveform of any shape. This is a useful tool when determining what frequencies will be required of a power subsystem to create all of the logic signals in a design.

FPBGA- Fine Pitch Ball Grid Array Package. A BGA package wherein the pitch of the balls is less than 1 millimeter or 39.7 mils.

FPGA- Field Programmable Gate Array logic circuit. This is a general-purpose logic array that can be programmed to perform a wide variety of logic functions.

FR-4- A name often used to describe a class of PCB laminates. These laminates are made with an epoxy based resin system. In actuality, FR-4 means "flame-retardant, class 4". This materials classification is not resin specific. However, the PCB fabrication industry has come to refer to any epoxy-based laminate as FR-4.

FROM TO LIST- A list of connections between the points in a PCB. Also called a net list.

FUNCTIONAL TEST (FT)- Testing an assembled PCB by exercising the circuits on the PCB with some sort of test pattern and observing whether or not they respond as designed. If the PCB has a failure, its location is deduced from the pattern of the output data. It is difficult to deduce from the failure data where the failure is. It is also difficult to devise test patterns that are exhaustive. This test method is best used as a go/no go test.

FUSING- The combining of metals through melting, blending and solidification. In PCB manufacture, this operation is done after tin and lead are plated on outer layers. The fusing turns the combination into solder.

GaAs- Gallium Arsenide. A semiconductor substrate made by alloying gallium and arsenic. GaAs is used for LEDs and very fast circuits such as microwave amplifiers.

GENCAM- A general-purpose data format for exchanging PCB design information. Gencam is sponsored by the IPC as well as other agencies.

GERBER®- a data format used by photoplotters to plot the film that is then used to create the layers of a PCB. The name comes from the Gerber photoplotter first used to create film for PCBs and ICs.

GETEK®- A proprietary resin system from General Electric used in multilayer PCBs. Its principal advantage over the "FR-4" family of resins is a higher Tg (approximately 180°C.)

GIGABIT- 10^9 bits, 1,000,000,000 bits. Slang use is to refer to the capacity of a data path, where it is understood to mean Gigabit per second.

GLASS TRANSITION TEMPERATURE (Tg)- A property of most resin systems used to fabricate PCBs. This is the temperature at which the temperature coefficient of expansion of the resin changes from a modest rate to a very high rate of expansion. PCBs heated to temperatures above the Tg of the resin system are subject to via failures, due to excessive stress in the "Z" axis.

GLVDS- Ground Low Voltage Differential Signaling. This is a variation of standard LVDS created by Ericsson to allow differential signaling on devices with VDD less than 3.3 volts.

GTL- Gunning Transistor Logic. Named for Bill Gunning of Xerox PARC in Palo Alto, California. This logic provides bus-driving capability on the output of CMOS devices, previously obtainable only with ECL drivers. The drivers are open drain N channel FET transistors.

GRIDDED ROUTING- Routing of PCB traces on a regular grid. This form of routing works best when the component pin pitch is regular and the number of signal layers exceeds four.

GRIDLESS ROUTING- Routing of PCB traces at any spacing without regard for any standard spacing or grid system. This form of routing is valuable for two and four layer PCBs. It becomes less valuable when the number of signal layers exceeds four.

GROUND- The one and only place in an electronic system that is used as a measurement reference for all signals.

GROUND BOUNCE- The voltage spike that develops across the inductance of the ground leads of a component as the current it draws from the "ground" power subsystem varies. The effect is to cause the "ground" terminal of the IC to move positive with respect to "ground" on the PCB. These spikes can cause failures.

HALOGEN- A product made using a chemical containing one of the halides--fluorine, chlorine, or bromine. It is usually used to increase flame resistance of a material. When released into the atmosphere by fire or other means, it causes pollution.

HARD METRIC- A component or design that is designed with the metric system as the primary set of units of measure. This is distinguished from soft metric where a component or design is designed using the English units of measure system which are then converted to metric. An example of the former is the family of 2 mm connectors, DIN connectors. An example of the latter is a 50-mil pitch BGA that is converted to metric as 1.27 mm.

HASL- Hot Air Solder Leveling. This is a method of coating the exposed copper on a PCB with solder. This operation is done after the solder mask has been applied over bare copper, SMOBC. The finished PCB is submerged in a vat of molten solder to coat all of the exposed copper. The PCB is then drawn out of the solder through a pair of "air knives" that blow off the excess solder, leaving flat surfaces of solder on the component mounting pads.

HDI- High Density Interconnect. This is a system of connectors that has lead pitches of less than 2 millimeters.

HDL- Hardware Description Language. A high level language used to design logic circuits that do not require the use of gate-level diagrams.

HENRY- The primary unit of inductance of a component that stores energy in the magnetic field that surrounds it. A henry is defined as the inductance for which the induced voltage in volts is numerically equal to the rate of change of current in amperes per second.

HOLE SHADOW- the shadow cast by the hole wall of a drilled and plated hole in a PCB. This is made up of the actual drilled hole and the tolerance associated with drilling. This shadow is cast on all layers--surface, signal and power. Any feature that comes within this distance of a plated hole may be shorted to it.

HSPICE- A SPICE program commercially available from Meta-Software.

HSTL- High Speed Transceiver Logic. This is a set of logic drivers and receivers using CMOS drivers whose output levels are compatible with TTL input levels. Drivers are rated in milliamps--the higher milliamp ratings are capable of driving 50-ohm transmission lines at high data rates.

IBIS- I/O Buffer Information Specification. This is a method of modeling IC drivers and receivers using behavioral models rather than transistor level models. IBIS was devised to allow SI modeling of high-speed circuits without requiring IC manufacturers to disclose the actual design of these circuits. The specification governing the creation of IBIS models is ANSI/EIA-656.

IC- Integrated Circuit. A monolithic circuit of many components built on a single piece of semiconductor.

IEEE- Institute of Electrical and Electronic Engineers, (www.ieee.org). The international organization of professional electrical and electronic engineers. The IEEE publishes original research on a wide variety of topics important to these professions. This organization has the largest collection of technical research papers available in the world on its web site. It also maintains most of the standards that govern protocols such as Ethernet, JTAG, etc. and sponsors a number of electronics-oriented technical conferences throughout the year.

IMAPS- International Microelectronics and Packaging Society, www.imaps.org. Based in Washington, DC., this organization specializes in the development of integrated circuits and their packages.

IMMERSION PLATING- The chemical deposition of a thin metallic coating over certain base metals by a partial displacement of the base metal. This method is used to plate gold, nickel, tin, or silver on the copper surfaces of a PCB after the final etching process has been completed and solder mask has been applied to it. Is usually referred to as electroless plating.

IMPEDANCE- The resistance to the flow of electromagnetic energy along a transmission line or through a component. Components that contain only resistance are said to have a non-reactive impedance meaning that their impedance is not affected by the frequency of the EM field. Components or structures that contain inductance or capacitance are said to have a reactive impedance meaning that their impedance is affected by the frequency of the EM field.

IMPEDANCE TEST TRACE- A special trace added to a signal layer for the purpose of testing the impedance of signal lines in that layer. The trace is plotted the same trace width as the other signal traces on the layer and has an access via at each end to allow contact to be made with a test instrument, usually a time domain reflectometer (TDR).

INCIDENT WAVE SWITCHING- This is a method of high speed logic signaling that relies on sending a full amplitude signal down a transmission line and absorbing the energy in the signal with a parallel termination at the end of the transmission line. In this form of signaling, data is good all along the transmission line as the signal passes each load.

INCIRCUIT TEST (ICT)- A method of testing an assembled PCB that attaches probes to every net on the PCB to observe the signals as test patterns are applied. Contact to the nets is made with a "bed-of-nails" fixture that has spring loaded probes that touch test pads attached to each net. This test method makes it possible to quickly locate defects such as shorts and opens. It is often possible to perform functional testing of the ICs with logic patterns.

INDUCTANCE (L)- The property of any conductor or component that is in series with any current flow or electromagnetic field propagating along it, that "impedes" the movement of the current or electromagnetic field. The effect of inductance is frequency sensitive, in that at DC, the effect of the inductance is not visible. As the frequency of the electromagnetic field impressed across an inductance increases, its "impedance" to the flow of the electromagnetic energy increases. Inductance is measured in units of measure of a Henry.

INDUCTOR- A component built to perform the function of inductance.

INFINIBAND ™- A data bus protocol that relies on differential signaling at data rates of a gigabit per second and higher. This protocol is intended to replace the PCI bus in products that use microprocessors. InfiniBand, (www.infinibandta.org) is the result of a trade association made up of more than 180 companies, such as Intel, HP and Dell, and their customers searching for a bus protocol that can transmit data at rates higher than the PCI bus format will allow.

INNER LAYER- Any of the inside or buried layers of a PCB. These buried layers can be planes or signal layers.

INSULATION RESISTANCE- The electrical resistance of an insulating material.

IPC- Institute of Printed Circuits, (www.ipc.com). A global organization with headquarters in Lincolnwood, Il, USA, is chartered with creating and maintaining specifications, procedures, and training classes for the printed circuit and interconnect industry and its customers.

IPC-D-317- A design specification published by the IPC that covers high speed PCB design. This specification has been replaced by IPC-2141.

IPC-2141- A high speed PCB design specification maintained and published by the IPC. This specification was composed by a volunteer committee and has not had rigorous technical review. It contains significant numbers of rules of thumb that have not been validated and are often imprecise.

IPC-782- A specification maintained by the IPC that covers surface mount patterns for electronic components.

IR- Infrared. A source of heat whose wavelength is below red in the visible spectrum. IR is used to solder surface mount components to PCBs.

I^2R LOSS- Power lost because a current flows through the resistivity of a component or wire. I = Current, R = Resistance. It is usually measured in watts.

ISDN- Integrated Services Digital Network. A method of providing high speed digital and voice service over the plain old telephone lines (POTS).

ISI- Inter Symbol Interference. The effect that data bits which precede and follow one another in rapid succession have on each other. This is noticeable when data rates approach 1 GB/S. Tends to spread out a bit such that it is not sharply defined.

ITRI- International Technology Research Institute, (www.itri.org). Based in Taiwan, this organization is chartered to do fundamental research into topics of importance to manufacturing activities including the PCB assembly industry.

JEDEC- Joint Electronic Device Engineering Council, (www.jedec.org). This is an organization that sponsors standards such as component packages, bus standards, etc.

J-LEAD PACKAGE- A surface mount package that has leads protruding from all four sides. These leads are bent into the shape of a J down the sides of the package so that they all exit on the bottom of the package. The lead pitch is 50 mils (1.27mm). This package replaced the DIP or dual in line package.

JOULE- The unit of work and energy in the International System of Units (SI), named after James Joule, an early researcher into electromagnetic phenomena. A joule is defined as the work done by a force of 1 Newton acting through a distance of one meter. One joule per second is one watt.

JTAG- Joint Test Action Group, (www.ieee.org). This is the group that has created a set of test methods defined in specification IEEE 1149.1. This specification defines a method for testing ICs after they are mounted on a circuit board without requiring access to each component lead. JTAG is also a slang term that describes a testing protocol called "The Standard Test Access Port and Boundary Scan Architecture.

JUMPER WIRE- A wire added to the surface of a PCB to complete a circuit that was not completed in the PCB itself. Jumper wires are usually used to correct design errors and are sometimes called "blue wires".

KILO- 10^3 or a multiplier of 1000 on a number. Examples, kilogram equals 1000 grams. Kilohertz equals 1000 Hertz.

LAMINATE- A general term describing the materials used in constructing PCBs. Usually, a laminate is made up of woven glass cloth and a resin system such as cyanate ester of epoxy. In the PCB fabrication process, laminate is often used to describe a PCB that has been constructed with a piece of resin-soaked glass cloth with a sheet of copper foil bonded to each side. This combination is used to form pairs of inner layers by etching patterns, such as planes or signals, into the foil on each side.

LAND- A pattern of metal on the surface of a PCB used to make a connection to a component, connector or other device.

LAND GRID ARRAY- An IC package that has all of its contacts on the bottom side in the form of round pads. These pads make contact with the pins in sockets. In most cases, land grid arrays are used to test new ICs that will later have balls or columns attached to the pads. These balls or columns allow the IC to be soldered to a PCB.

LAND PATTERN- A collection of lands arranged in such a way that a component, such as a surface mount IC, can be soldered to the PCB.

LAYER-TO-LAYER REGISTRATION- The accuracy with which layers in a PCB are aligned to each other.

LAYER-TO-LAYER SPACING- The distance between the adjacent copper faces of two conductive layers on opposite sides of a piece of insulating material or laminate.

LCC- Leaded Chip Carrier. This is a component package for an IC that has leads. It is connected to a PCB or socket by leads on the sides or bottom of the carrier.

LEADFRAME- The collection of leads inside an IC package. At one end, the leads connect to the contacts of the IC die. On the other end, the leads connect to the mounting pads on a PCB.

LEADLESS CHIP CARRIER- An IC package that has no leads. Connections to a socket of the PCB are made using conductive patches on the sides or bottom of the device.

LOSS TANGENT, γ- A measure of the degree to which a dielectric or insulating material absorbs energy from an electromagnetic field passing through it. Loss tangent is usually frequency sensitive and higher at higher frequencies. See dissipation factor.

LVCMOS- Low Voltage Complimentary Metal Oxide Semiconductor. This is a family of logic circuits using CMOS transistor circuits that run at voltages less than 5 volts, usually 3.3V, 2.5V or 1.8V.

LVDS- Low Voltage Differential Signaling. This signaling protocol, known as RS-644, is maintained by ANSI (www.lvds.com). The signaling protocol was devised by a consortium of laptop PC manufacturers and National Semiconductor to provide a method for getting the high bandwidth graphical data from the laptop motherboard to the display through the hinge. This signaling protocol has a high tolerance for poor grounding between the two ends of the data path. Subsequently, this signaling method has been adopted by virtually all designers of very high performance products, such as terabit routers.

MANHATTAN DISTANCE- The distance between two component pins on a PCB achieved by connecting the two by using only X and Y travel. This measurement is named for the streets and avenues of Manhattan where a traveler must stay on the streets and avenues when traveling from point to point. This is the usual method for routing PCBs with large numbers of signals.

MASS LAMINATION- A process for laminating multilayer PCBs wherein a large number of PCBs are fabricated on a single panel. This process works well for four layer PCBs, but not for higher layer counts. Mass lamination is the workhorse of the PC motherboard industry.

MAXWELL'S EQUATIONS- A collection of equations developed by James Clark Maxwell that describe the behavior of an electromagnetic field. These equations are commonly used to calculate the impedance of transmission lines.

MCM- Multi-Chip Module. A package that contains more than one integrated circuit die.

MDA- Manufacturing Defect Analysis. This is a method of testing assembled PCBs for defects induced by assembly such as solder shorts and opens and wrong value components. This is the most basic form of in-circuit tests (ICT). The assembled PCB is placed on a bed of nails test fixture that makes contact with each net on the PCB. Test voltages are applied and currents are measured.

MEGA- Numerical prefix denoting million, 10^6, 1,000,000.

MEGABIT- One million bits, 10^6 bits, where a bit is one element of a data word. Megabit is usually used to describe the capacity of a data path, meaning one million bits per second. A common Ethernet path size.

MEZZANINE- A term used to describe a PCB assembly that mounts on top of another PCB assembly.

MICRO- Numerical prefix denoting a fraction of one-millionth, 10^{-6}, 0.000,001. Commonly used to describe a unit of time, one microsecond.

MICROINCH- One millionth of an inch, 0.000,001".

MICROBGA- A ball grid array package with ball pitch or spacing less than 1 mm (39.7 mils). This package pitch is often used to make compact memory products.

MICRON- One-millionth (10^{-6}) of a meter. Approximately 39.7 micro-inches.

MICROSECOND- A unit of time, one millionth of a second, 0.000,001 second. The time it takes light to travel 1000 feet in a vacuum.

MICROFARAD- One millionth of a Farad. A unit of capacitance, 0.000001 Farad

MICROSECTIONING- The process of examining the internal structure and plating of a PCB by cutting a section out of a PCB, polishing it and examining it under a high power microscope.

MICROSTRIP, MICROSTRIPLINE- A term used to describe a transmission line or trace that travels over a single plane. These transmission lines are formed on the outer layers of PCBs. They can be on the surface of the PCB or buried in the first layer of insulating material from which the PCB is made.

MICROVIA- A via or drilled hole of diameter less than 8 mils, 0.2 mm. A microvia may pass all the way through a PCB or only part way through (blind via). This term is often misused to describe a blind via, which is a hole of any diameter that does not pass all the way through a PCB.

MICROWAVE- It is generally accepted that frequencies higher than 1 GHz, 1,000,000,000 Hertz are considered microwaves. Microwaves provide a way of dividing up the RF spectrum into bands of frequencies that require certain kinds of design discipline.

MIL- Slang for one thousandth of an inch (10^{-3} inches), 0.001".

MILLI- Numerical prefix denoting a fraction of one thousandth (10^{-3}).

MILLISECOND- A unit of time, one thousandth of a second, 0.001 seconds.

MIL-STD-275- "Military Standard, Printed Wiring for Electronic Equipment." A specification for designing printed wiring boards maintained by the U.S. Navy in Washington, DC.

MIL-STD-55110- A standard created and maintained to regulate the manufacture and quality of PCBs. Titled "General Specification for Military Printed wiring Boards". Maintained by the U.S. Army at Fort Monmouth, NJ.

MINIMUM ANNULAR RING- The minimum width of metal, at the narrowest point, between the edge of a hole and the outer edge of the pad through which it passes. The measurement is made between the edge of the metal pad and the edge of the drilled hole, not the edge of the plating in the hole.

MIPS®- Millions of Instructions per Second. A measure of the performance of a processor. This is also the trade name of a RISC processor family and the company that devised it.

MLC- Multilayer Ceramic. A type of capacitor made by interleaving conductor plates with thin layers of ceramic.

MONOPOLE- A radiating element that has only one element. An antenna that has a single radiating element placed over a plane or ground plane. AM radio antennas on mountaintops are examples of monopoles. An unshielded wire exiting a Faraday cage is an example of a monopole antenna.

MOS- Metal Oxide Semiconductor. This is a semiconductor technology used to create transistors using a metal "gate" separated from the conductive channel by a silicon dioxide layer. Under the gate is a channel in the substrate that is made to conduct (enhancement mode) or not conduct (depletion mode) by applying a voltage to the gate. This is the most common type of semiconductor technology used to build logic components.

MOTHERBOARD- A PCB with components on it onto which are plugged other PCBs, such as memory modules. A motherboard is different from a backplane in that a back planeusually has only connectors on it.

MOUNTING HOLE- A hole in a PCB used to mount the PCB to a chassis or other structure. Mounting holes may be plated or unplated.

MULTILAYER PRINTED CIRCUIT BOARD- A PCB with more than two layers, produced by laminating inner layers to outer foils.

MULTILINE® PUNCH- A system used to punch alignment holes on the periphery of an inner layer pair in order to allow alignment of one layer pair to another. Usually, there are four holes, one in the center of each side. The holes are usually slotted to allow for expansion due to temperature changes during the lamination cycle.

NANO- One billionth or 0.000,000,001.

NANOFARAD- A capacitance of one billionth of a Farad in size.

NANOHENRY- An inductance of one billionth of a Henry in size.

NANOSECOND- A unit of time that is one billionth of a second. The time it takes an electromagnetic wave to travel one foot, 30.3 cm, in a vacuum or 6 inches in a PCB.

NAIL HEADING- The flared condition of inner layer copper around drilled holes resulting from drill wobble.

NEGATIVE ETCH BACK- Etching the copper in drilled holes such that it is recessed back from the laminate. This condition is undesirable as it results in poor contact between hole plating and the inner layer copper to which it connects.

NET- An electronic circuit. This term is used in PCB layout to describe a complete signal path. Nets with more than two pins will be made up of more than one wire.

NET LIST- The complete list of connections in a PCB. Prior to placement of the components on a PCB surface, a net list is not "physical", but rather, "symbolic". Symbolic net lists are those that are created from schematics. Physical net lists are created from PCB layout tools.

NOISE MARGIN- The difference between the minimum signal sent down a logic line by a driver and the minimum signal required by a load or input. This margin is built into a logic family to allow for noise that gets onto a signal as a result of switching activity.

NONFUNCTIONAL PAD- A pad around a drilled and plated hole that has no connection made to it. These pads are often present on inner layer artwork. Usually, they are created by the CAD system used to create the artwork for a PCB. These pads add no value and can result in unwanted shorts between traces and holes. It is good practice to remove them from the artwork prior to building a PCB.

ODB++®- A data format used by the PCB fabrication industry to create the working film used to etch the layers of a PCB. This data format has its origins in the Valor® CAM system used by PCB manufacturers. The format has since become an industry standard format.

OHM- The unit of resistance (or impedance) in the International System of Units (SI). The ohm is the resistance of a conductor such that a constant current of one ampere through it produces a voltage of one volt between its ends. I = E/R

OHM'S LAW- A law discovered by George Ohm that describes the relationship between the current through a resistance and the voltage across it. Basically, the law states that the current in amperes is equal to the voltage in volts across a resistance divided by the resistance in ohms.

ONE BETWEEN- A slang expression referring to routing a single trace between two pins on an IC or connector.

OVERSHOOT- A reflection caused by an impedance change on a transmission line. The reflection is in such a direction that it adds to the logic signal that caused it. Overshoot results when the downstream impedance at a change is higher than the upstream impedance. Overshoot can be either positive or negative in polarity depending upon whether the edge that creates it is rising or falling.

PAD STACK- A term used to describe the collection of pads used on all layers of a PCB for a given hole size. The "stack" will have clearance pads for power layers; capture pads for internal signal layers; surface layer pads; solder mask pads and so on. Each drilled hole size on a PCB will have a unique pad stack.

PANEL- A rectangular shaped piece of PCB material, either double-sided or multilayer, that is processed by a PCB fabricator. Panels come in standard sizes into which are fitted as many PCBs as is practical.

PANEL PLATING- A plating operation that plates copper to the entire surface of a panel as opposed to pattern plating where plating is done only in the holes and on the traces and pads that will exist on the outer layers in the finished PCB.

PARASITIC- An unwanted characteristic of a component that is inherent in the make up of the component. For example, capacitors have unwanted inductance due to the fact that their length is not zero and unwanted resistance due to the fact that their leads are made from metal. These parasitics usually have a minor effect on performance at low speeds or low frequencies. However, at high speeds and high frequencies, they often determine the limits of performance.

PATTERN PLATING- A plating operation that plates copper only on the traces and pads that will exist on the finished PCB as well as in the holes that form the vias and other layer-to-layer connections.

PBGA- Plastic Ball Grid Array package. A multilayer IC package whose insulating material is a plastic laminate, such as BT, cyanate ester or FR-4. The PBGA makes the connections from the IC die to the PCB.

PCA- Printed Circuit Assembly. This is a PCB finished product with all of its components mounted or soldered onto it.

PCB- Printed Circuit Board. This is an insulating material, such as woven glass saturated with a resin, such as epoxy, onto which copper foil has been bonded. Circuit patterns are etched or "printed" in the foil. A PCB can have as few as one layer or as many as a hundred.

PCI- Personal Computer Interface. A data bus protocol used in personal computers. This bus protocol has been adapted to instrumentation products. There are currently three standard PCI bus clock rates--33 MHz, 66 MHz and 100 MHz.

PCI Express- A differential serial link based on the communication protocol that uses the PCI bus as its basis. This protocol allows devices to be connected with data paths that run at 2.4 GB/S.

PCMCIA- Personal Computer Memory Card International Association, (www.pcmcia.org). Based in San Jose, CA, USA, this is an organization of personal computer manufacturers that developed a standard for packaging memory in small modules that could be inserted in a special slot on a PC. This format has been expanded to allow a wide variety of peripherals to be packaged in it for use in laptop and other portable products. This format is now called the PC Card format (www.pc-card.com) to reflect the fact that PCI is now primarily composed of a variety of plug-in cards used in PCs.

PECL- Positive ECL. ECL or emitter coupled logic normally operates with a supply voltage of –5.2V. On occasion, it is desirable to use a few ECL circuits along with MOS or CMOS circuits that are running with a +5V supply voltage. In this case, the ECL is said to be "positive" ECL.

PEEL STRENGTH- A measure of the robustness of the bond between copper on a PCB and the laminate to which it is bonded.

PERMEABILITY (μ_o)- Expressed as Weber per ampere per square meter, or henrys per meter. This is a measurement of the inductive properties of a conductor.

PERMITIVITY (ρ) - Another name for dielectric constant. This is a measure of the charge that builds up on the plates of a capacitor as a function of the voltage that is applied across it.

PHYSICAL NET LIST- A net list that has the XY locations of the device pins added to it. It is generated after the components of a PCB have been placed on the surface of the PCB. This net list format allows bare board PCB testing as well as creation of "rats nests" of pre-routed wires for use in assessing the rout ability of a placement.

PICO- A numerical prefix denoting a fraction that is one trillionth of something, 10^{-12}, 0.000,000,000,001.

PICOFARAD- A unit of capacitance that is 0.000,000,001 Farads.

PICOSECOND- A unit of time that is 0.000,000,001 seconds.

PLACEMENT GRID- The XY grid system on which component pins are placed on a PCB surface.

PLANE- A solid sheet of copper in a PCB that is used to conduct power to devices. The plane also serves as the partner to transmission lines or signals.

PLANE WEB- The segment of plane copper between two clearance holes.

PLASMA ETCH- A process of cleaning the plastic smeared onto the edges of copper in a drilled hole by the drilling process. It uses an ionized gas or plasma to attack or etch away the plastic residue. Plasma etch is commonly used on PCBs made from polyimide. This process is done in a chamber that has had the air pumped out of it and refilled with the plasma-forming gas. The gas is ionized with an RF field.

PLATED THROUGH HOLE- A hole drilled through a PCB that has been plated with metal. The metal is normally copper. Plated through holes can be used to solder the leads of through hole components; to make a connection between a signal trace and the component lead of a surface mount device or as a way to change from one signal layer to another.

PLATING- The metal plated onto a PCB surface that can be any of several metals or combinations of them. Plating also refers to depositing metal onto the surface of a PCB. This can be done using an electrical current--electroplating, or by ion exchange--electroless plating.

PLATING RESIST- A photosensitive coating applied to the outer layers of a PCB after lamination, drilling and electroless copper plating. This coating is exposed photographically such that the pattern of traces and pads that will remain on the outer layers after lamination is exposed and the copper that will be removed by etching is covered. Following this step, copper is plated on all of the exposed copper and in all of the plated through holes. The plating resist blocks plating in areas where it is not wanted.

PLATING UP- The act of adding metal, usually copper, to a trace or copper layer of a PCB to increase its thickness.

PLL- Phase Locked Loop. This is a circuit that locks onto a clock or other signal and produces a copy of it or a multiple or sub-multiple of the clock frequency that is locked in phase to it.

POLYIMIDE- A resin system used as the “glue” for the insulating layers of PCBs. This resin can withstand the very high temperatures used in soldering better than any other resin system. Disadvantages include high cost, difficulty in processing and high water absorption.

POTS- Plain Old Telephone System. This is the audio phone system used worldwide to make voice telephone calls.

POWER PLANE- A copper plane placed in a PCB that conducts the current from the power supply to the components mounted on the PCB. A power plane may be the positive or negative terminal of the power system. Planes also function as the partners of transmission lines.

PPE- Polyphenylene Esther. A resin system used in high performance laminates such as Getek®. The principle benefit of this material is lower loss tangent.

PQFP- Plastic Quad Flat Pack. This is a surface mount IC package with leads protruding from all four sides. The body of the package is molded plastic.

PREPREG- A name given to glass cloth coated with a resin such as epoxy. The resin is not fully cured. A prepreg layer or layers is used between two internal layers of a multilayer PCB as the “glue” during lamination. Another name for prepreg is “B stage”.

PRESS FIT- A method of attaching components, such as connectors, to a PCB without using solder. This is done by forcing a compliant pin into a plated hole that is slightly smaller than the pin. This interference fit retains the pin in the hole and also makes the electrical connection.

PRIMARY SIDE- In a PCB with components mounted on both sides, the primary side is the side on which the majority of the components are mounted.

PROBE POINT- A pad on the surface of a PCB that is used to make electrical contact for the purpose of performing some type of electrical test.

PROM- Programmable Read Only Memory. This is a memory that can be programmed with a fixed bit pattern and then read as often as desired. This memory cannot be reprogrammed. The memory contents are not lost when power is removed and is used to store programs.

PSPICE- A commercially available version of SPICE developed and marketed by Microsim Corporation.

PTH- Plated Through Hole, see above.

PWB- Printed Wiring Board. This is a name given to printed circuit boards to denote that the wires are ”printed” on the layers.

QFP- Quad Flat Pack. This is a surface mount IC package with leads protruding from all four sides.

RADIATED EMISSIONS- Electromagnetic energy that escapes from a product by radiating it into space from the conductors of the product. Also referred to as EMI, these emissions are usually measured in the frequency band 30 MHz t o1 GHz.

RAMBUS®- A proprietary memory bus architecture created and owned by Rambus Corporation. The use of this architecture is licensed to others by Rambus. Its intended use is to provide a very high clock rate memory interface to processors.

RATS NEST- A plot of the wires in a PCB created before the PCB is routed. This plot shows the “crow flies” connections to all of the components. The plot is useful in assessing the routability of a particular placement.

REFLECTED WAVE SWITCHING- A method of high speed signaling that sends a signal of half amplitude down a transmission line relying on the fact that the signal will double at the open end of the transmission line and reflect back to the source, making the signal full amplitude as the reflected wave arrives back at the source. This is achieved by using a series termination to match the driver to the transmission line.

REFLOW SOLDERING- A method of soldering components to a PCB by exposing the PCB and components to a heat source, such as infrared. The solder used to make the joints is placed onto the leads of the components and the pads prior to heating. This solder is then simply "reflowed" to create the joint. The alternative to reflow soldering is wave soldering, an operation that floats the PCB over a wave of molten solder. The solder needed to make the joint comes from the wave.

RELATIVE DIELECTRIC CONSTANT (e_r) - A measure of how a dielectric material slows down the speed of an electromagnetic field traveling through it as compared to the speed the electromagnetic field would have in a vacuum. Also, a measure of how a dielectric increases the parasitic capacitance between two conductors as compared to the capacitance between the same two conductors in a vacuum.

RESIN RICH- A laminate material, usually a prepreg, that has an unusually high amount of resin as compared to the amount of reinforcement or glass. The reason for the high concentration of resin is to provide resin to fill the voids in the opposing copper layers during lamination.

RESIN SMEAR- Resin that is "smeared" onto the exposed edges of the copper layers of a PCB during drilling. This resin smear blocks the plating of copper onto the inner layer copper circuits and results in open circuits. Prior to plating, some sort of cleaning or desmearing must be done to remove this coating.

RESIN STARVED- The condition of a laminate resulting from too little resin. The result is often poor bonding during lamination or voids.

RESIST- A material, usually photosensitive, that is used to cover the copper on a layer of a PCB to protect it during etching or to prevent it from being plated, during plating operations.

RESISTANCE, (R) - The property of a conductor that causes it to inhibit the flow of current through it. Measured in ohms. One ohm is equal to 1 volt divided by 1 ampere of current. This is Ohm's law. R = E/I

RF- Radio Frequency. This usually refers to signals higher than 500 KHz and less than 1 GHz. The main property of an RF signal that causes special design considerations is its very small size (often only a few microvolts) when it arrives at a receiver.

RFIC- Radio Frequency Integrated Circuit. This is an IC designed to process RF signals.

RIGID-FLEX PCB- A PCB that has a portion that is rigid and a portion that is made from a flexible material. The flexible portion often replaces a cable and the associated connector.

RING BACK- A term used to describe undershoot.

RIPPLE- The voltage variations that appear on the Vcc or Vdd rail of power supplies. These variations can be created by the power supply itself or by varying load currents that cause the supply voltage to drop.

RISC- Reduced Instruction Set Computer. This is a computer architecture that operates on a very small number of simple instructions.

RISE TIME- The time required by a signal transitioning from a logic O to a logic 1 to travel between the 10% voltage level and the 90% voltage level.

RJ-11- The name of a common (i.e. standard) connector used to connect to a phone service in the United States. It has four contacts and has no shield.

RJ-45- The name of a common connector used to make Ethernet connections from unshielded twisted pair (UTP) wire to a chassis or other unit. This connector has eight contacts. This is the most common connector used to connect devices to the Internet.

RMS- Root Mean Squared. This is a mathematical method for calculating the effective voltage of a varying signal such as a sine wave. It is used to calculate measurements such as average power.

ROM- Read Only Memory. This is a memory whose contents cannot be written, but can be read many times.

ROUTING- When used in the context of laying out a PCB, routing refers to the placing of individual wires or traces on the circuit layers used to make the connections. A software program is used to perform this operation. When used in the context of

manufacturing PCBs, routing is the operation that cuts individual PCBs out of the panel from which they are built. A router of the type used to do wood work is often used for this operation.

ROUTING GRID- The grid system used to route the traces on the layers of a PCB. This is usually the pitch of the component pins or a sub-multiple of this.

R-PACK- Resistor package or pack. This is a multi-leaded component that contains more than one resistor inside.

RS-232- A low speed, (less than 20 KB/S), serial data link used in noisy environments, such as factory floors. This standard is maintained by ANSI/TIA/EIA. The minimum signal swing is +/-5V.

RS-422- A TIA/EIA standard for serial interfaces that extends distances and speed beyond RS-232. It is used in multipoint lines.

RS-644- The standard maintained by ANSI/TIA/EIA that describes the LVDS differential signaling protocol.

RULE OF THUMB- A way of estimating some parameter without performing detailed calculations. Rules of thumb take two forms. The first is a simplification done after detailed calculations are made. This allows quick studies of whether or not it is practical to do a certain thing. The second is made up after observing that a change of some kind had an effect on a circuit. These rules of thumb are offered without understanding why the change had an effect. This kind of rule is often offered by EMI "Gurus" who don't take the time to understand what is happening. Neither rule of thumb type should ever be part of a final design rule set.

S-GLASS- A formulation of glass that has a lower loss tangent than the more common E-glass. This glass style is used to formulate ultra low loss laminate for high-speed digital circuits.

S PARAMETERS- Microwave measurements of the transmissive (gain or loss) and the reflective behavior of a two port network. This is a "black box" method of describing the behavior of a component or transmission line.

SAFETY GROUND- The "green" wire that is usually the third wire on an AC power outlet or power cable. This wire is connected inside a building to a stake or other connection into the earth ground. It is to be connected to the case or chassis of a product. The function of the safety ground is to conduct any voltages that might accidentally contact the case to ground, so that the case does not conduct that voltage and harm a user. The safety ground has no function in the containment of EMI.

SCHOTTKY DIODES- Diodes formed by using a layer of metal as one contact or pole and a doped semiconductor as the other. The primary reasons to use Shottky diodes rather than PN junction diodes are lower turn on voltage (0.3V vs. 0.7V) and faster turn on and turn off times. In ICs, these diodes are often used as input and output protection circuits as well as to keep bipolar transistors from saturating.

SCSI- Small Computer System Interface. In its original form, it is a parallel bus technology designed to connect the components of a small computer, such as disc drives, printers, memory and CPU. There are now both serial and parallel versions of the SCSI bus. Specifications are maintained by the National Committee for Information Technology Standards (NCITS, www.t10.org). There are several versions of SCSI starting with the parallel bus architecture extending to differential signaling.

SCR- Silicon Controlled Rectifier. This is a transistor formed with three junctions that functions like an on and off switch. Used in power control circuits. This kind of transistor can exist between inputs or outputs of an IC by accident as an unwanted parasitic. They can be "turned on" by excessive overshoot, resulting in failures.

SECONDARY SIDE- The least populated side of a PCB that has components mounted on both sides.

SEGMENT- A portion of a net routed in a single layer. This segment may be part of a wire, all of a wire, part of a net or all of a net.

SEQUENCING- The arranging of points in a network to achieve proper transmission line operation. This is accomplished by hand or by the use of an automatic router in a PCB design system that has been designed to handle high-speed designs.

SEQUENTIAL LAMINATION- The process of creating a multilayer PCB wherein some of the inner layers are laminated together as an intermediate PCB. This intermediate PCB is drilled and plated prior to being combined with the remaining layers. This is the method used to create buried vias. It can also be used to create blind vias.

SERDES- Serializer-Deserializer. A circuit that takes a parallel data stream and converts it into a serial data stream and back.

SERPENTINE- A zigzag shape placed in a trace on a PCB to increase its length. Used to add time delay.

SES- Strip Etch Strip. This process is used to form the patterns on the outer layers of a PCB. After plating of the copper in the holes and on the traces and the plating of tin-lead or another etch resist onto the copper that is to remain on the finished PCB, the plating resist is stripped away. The unwanted copper is then etched away and the tin-lead etch resist is stripped off the copper traces and pads. The next step is to apply a solder mask over the bare copper.

SILKSCREEN- The legend layer of artwork on a PCB that labels components, connectors, etc. Silkscreen is so named because these legends are commonly applied to the surface of a PCB using the silkscreen process.

SIMULTANEOUS SWITCHING NOISE (SSN)- Often referred to as Vcc bounce or ground bounce. This is a voltage transient that appears on the Vcc lead or the ground lead of an IC die. It occurs when switching currents drawn by the IC from the power system pass through the inductance of the power leads of the IC package. This noise usually takes the form of voltage spikes that are superimposed on output signals and can cause switching failures.

SMOBC- Solder Mask Over Bare Copper. This is a method of applying solder mask on the outer surfaces of a PCB. First, any plating (usually solder) that was applied to the copper outer surfaces to allow etching of the outer layer patterns is removed. Then, solder mask is applied over the bare copper. This improves the adhesion of the solder mask to the copper traces on the outer layers.

SOFT METRIC- This term describes a component that was designed using the English measurement system and is then redimensioned in the metric system. For example, a 50-mil pitch BGA package is also referred to as a 1.27 mm pitch part.

SOIC- Small Outline Integrated Circuit. This is an integrated circuit package that has been reduced in size to allow it to be used in miniature products, such as cell phones or video cameras.

SOLDERMASK- A protective coating applied to the outer surfaces of a PCB to cover up metal features that must be protected from solder during the soldering operation.

SOLDER SIDE- The side of a PCB opposite the side upon which components mounted on it. This side of the PCB is normally passed over a wave soldering machine to complete the soldering of through hole components. This term does not apply to a PCB with components mounted on both sides.

SONET- Synchronized Optical Network. This is the network protocol used by the Telco operators to create broadband links between nodes of the Telco network.

SPICE- Special Program for Integrated Circuit Emulation. This is an analytical program developed at UC Berkeley to allow electrical engineers to analyze electrical or electronic circuits built from electronic components such as transistors, resistors, capacitors and inductors. This product has been adapted for the analysis of all sorts of circuits beyond ICs, such as PCB networks. HSPICE and PSPICE are commercial versions of the original SPICE.

SRAM- Static Random Access Memory. This is an IC memory that can be read and written randomly. It does not require any periodic refresh to keep the contents in memory.

SSTL- Series Stub Terminated Logic. This is a method of allowing several memory modules to be hooked to a bus without suffering degradation from the stubs inside the memory modules. Accomplished by putting a small resistor, (usually 22 ohms), between the stub and the main bus. Limits upper performance due to degradation of the signal after the resistor at fast edge rates or high clock frequencies.

STACKUP- A name given to the spacing and ordering of the layers in a PCB. The “stackup” defines how the PCB is to be built and the thickness of each copper layer and insulating layer.

STACKING STRIPE- A stripe of copper plotted and etched on each layer of a PCB. These stripes are located along one edge of a PCB so that they are visible when the PCB is cut from the manufacturing panel. They allow easy checking of dielectric and copper thickness and ensure that the layers are in the proper order.

STEPPED VIA- A via or plated through hole in a PCB that has one diameter for part of its length and another diameter for the rest of its length. The usual reason for using stepped vias is to provide a larger diameter for the press fit portion of a connector pin and a smaller diameter for the rest of the hole. This is a method for reducing the parasitic capacitance of the plated through hole.

STRIPLINE- This is the name given to a signal layer or transmission line that is sandwiched between two planes or power planes. This signal layer can be centered between the planes or offset toward one plane or the other.

SURFACE MOUNTING- The process of soldering components to the surface of a PCB by soldering their leads to pads on the PCB surface as opposed to drilling holes in the PCB through which the component leads are passed and into which they are soldered.

SYMBOLIC NET LIST- A net list for a PCB that is created while the design is at the schematic level. The net list is created prior to performing placement of the components on the surface of the PCB.

TAB- Tape Automated Bonding. This is a method of attaching ICs to lead frames by first attaching them to a miniature lead frame held in place in an opening in a tape, such as a 35 mm tape.

TELCO- The telephone companies that comprise the global phone system.

TENTING- The process of covering a hole in a PCB, usually a via, with solder mask so that the hole is sealed off.

TEST POINT- A contact point added to a net to facilitate connecting test probes to a PCB. Test points are usually small, round pads that are inserted in the middle of a trace on an outer layer. If no such trace exists, a short piece of trace may be added to a component via to which the test pad is attached.

TEST TRACE- A special trace added to a signal layer. The test trace is the same width as the signal traces to allow impedance testing of signals in that layer. Test trace is also referred to as an impedance test trace.

TDR- Time Domain Reflectometer. This is an instrument used to measure the impedance of transmission lines. It consists of a voltage step generator that launches very fast rise time pulses down the line to be tested and a sampling oscilloscope that monitors the voltage that travels on the transmission line.

TERABIT- 10^{12} bits, one trillion bits, 1,000,000,000,000 often used to describe the size of a data path or router in bits per second.

THERMAL RELIEF- A method of thermally isolating a plated through hole from the plane in a PCB to which it is connected. It simultaneously makes an electrical connection to the plane. This is done to so that the heat required to solder a component lead into the hole is not drawn away by the plane copper. (See also the paper on thermal ties in the appendix of this book.)

THERMAL TIE- When a connection is made between a plated through hole and a power plane, a good electrical connection is made. A good thermal connection is made as well. If a through hole component lead is to be soldered into that hole, the heat required to solder is drawn away from the joint by the plane copper resulting in a poor connection. In order to avoid this problem, a ring is etched around the plated through hole to provide a thermal isolation. The electrical connection is made across this gap with a small trace or tie, big enough to conduct the current and small enough to thermally isolate the pin.

THIEVING- The addition of dummy copper pads to the open spaces on the outer layers of a PCB in order to provide a uniform distribution of copper across the entire surface. The purpose is to insure plating currents are uniform across the whole surface when plating copper onto the PCB surface and in the holes. This helps insure a uniform plating thickness across the surface and in the holes.

THROUGH HOLE MOUNTING- A method of mounting an electronic component to a PCB by inserting its leads into drilled and plated holes. Once a lead is in the hole, it is soldered such that the space between the lead and the hole wall is completely filled with solder.

TIA- Telecommunications Industry Association, (www.tiaonline.org). A standards organization located in Arlington, VA, USA, that is comprised of telecommunications-related companies. This organization represents its members in a variety of ways. It also participates in creating and maintaining standards related to the telecommunications industry.

TIME DOMAIN REFLECTOMETER (TDR)- See TDR above.

Tg- Glass Transition Temperature. A property of most resin systems used to fabricate PCBs. It is the temperature at which the temperature coefficient of expansion of the resin changes from a modest rate to a very high rate of expansion. PCBs heated to temperatures above the Tg of the resin system are subject to via failures.

TTL- Transistor-Transistor Logic. This is a logic family that is made up of bipolar transistors that are connected in such a way that logic is preformed. TTL was formerly the workhorse of logic designers. This type of logic is no longer manufactured in volumes large enough to support production. It has been replaced by CMOS and MOS circuits.

TOOLING HOLE- A hole drilled in a PCB or a laminate panel that is used to align the PCB or panel to a fixture or tooling plate. These holes are normally located in the boundaries of the panels or in the corners of the PCBs. In order to maintain accuracy, these holes are not plated

TQFP- Thin Quad Flat Pack. This is a quad flat pack that has a very low profile obtained by grinding the IC die inside the package to a thinner than usual dimension. TQFPs are usually found in small memory modules.

TRACE- The name given to the signal wires etched on the layers of a PCB.

TRANSITION ELECTRICAL LENGTH (TEL)- The physical length of a switching edge (rise or fall time) in a PCB when converted from time to length by multiplying the time by the velocity of EM fields in the PCB material.

TRANSMISSION LINE- A conductor set. In a PCB, a transmission line is a trace and one or two planes, used to send electromagnetic energy from a source to a load. A transmission line is also any combination of conductors that serves in this function.

TRUE POSITION- The precise location, provided in XY coordinates, of a feature or hole on the surface of a PCB.

TRUE POSITION TOLERANCE- The amount that a feature, such as a drilled hole, may stray from the true position where it should be. Drill error is often expressed as true position radius, TPR. A true position radius of 5 mils would denote that the hole center will appear within a circle whose diameter is two times the TPR. The hole wall may appear in a circle whose diameter is the drill diameter plus two times the TPR. This diameter is sometimes referred to as the hole shadow.

TSSOP- Thin Shrink Small Outline Package. This is a surface mount IC package, usually found in memory modules, that has a smaller lead pitch and thinner package than normal.

TWO BETWEEN- A slang term referring to routing two traces between the pins of an IC or connector.

UI- Unit Interval. This term describes the duration of a single data bit in a data stream. For example, the UI for a 2.4 GB/S data stream is 416 picoseconds. A common method for evaluating the quality of a differential serial link such as OC-48 is to display a random series of data bits on an oscilloscope screen with the time across the display equal to one UI. This produces the “eye” diagram display.

UNIT INTERVAL- See UI above.

UL- Underwriters Laboratory, (www.ul.com) is an organization in the United States that oversees product safety

UNDERSHOOT- A reflection at an impedance change in a transmission line that subtracts from the incident signal. This occurs because the “downstream” impedance is lower than the “upstream” impedance. Undershoot is not polarity related, meaning that it is not a reflection that is negative going or positive going. Undershoot is sometimes referred to as “ring back”.

USB- Universal Serial Bus. A differential signaling protocol used to connect peripherals to personal computers. Exists in versions 1.1 and 2.0. With USB, up to 128 devices can be connected simultaneously. Specifications maintained by the USB Implementers Forum, Inc. (www.usb.org).

UTP- Unshielded Twisted Pair. This is a nickname for the most common phone wiring used in the United States. It is also the most common wire used to make connections to the Internet.

Vcc- The terminal of a power supply system that is connected to the collector end of a bipolar circuit containing NPN transistors. It is the most positive terminal of the supply.

Vcc BOUNCE- The voltage spike developed across the inductance in the Vcc path of an IC package when circuits in the IC switch. This voltage spike drives the Vcc terminal of the IC die negative with respect to Vcc on the PCB. This noise spike shows up on all output and input pins as a noise spike and can cause logic failures.

Vdd- The terminal of the power supply system that connects to the drain end of a MOS or CMOS transistor or IC. This is commonly the most positive rail of the power supply.

VHDL- Very High Speed Hardware Description Language. A method used to design logic integrated circuits using high level representations of logic functions instead of actual gates and latches.

VIA PLUGGING- A process that places a material in the vias after plating has been completed. The plugging material may be non-conductive or conductive. There are three reasons to plug vias. The most common is to block air flow so that the PCB can be pulled down onto a test fixture using a vacuum. Another reason is to plug the via so that subsequent etching steps don't etch away the copper in the via. The third is to make the pad through which the via has been drilled flat again. In the latter case, plating is done over the plug so that the pad is solid.

VME- A bus and hardware protocol that was developed to provide standard card cages, back planes, connectors and operating software for creating instrumentation systems. VME stands for VERSAmodule Eurocard. It was coined in 1980. The standards organization that develops and maintains the VME standards is VITA, VMEbus International Trade Association, located in Fountain Hills, Arizona USA, and can be contacted via e-mail at Info@vita.com.

VOLTAGE SOURCE- A signal source whose output is a constant voltage independent of the load impedance or load current. Most logic circuit outputs are variations on voltage sources. Power supplies are intended to be constant voltage sources. A true voltage source has an output impedance that is zero.

VOLT (V)- The unit of voltage or potential difference in SI units. The volt is the potential difference between two points of a conducting wire carrying a constant current of one ampere when the power dissipated between these points is one watt.

VOODOO BEAD- A derogatory term that describes a ferrite bead that has been placed in a circuit in the hopes that it will reduce EMI. The use of "Voodoo" beads is most often done by EMI practitioners who don't really understand the sources of EMI and are using empirically derived "rules of thumb".

VIA- A hole that is used to make an electrical connection between the layers of a PCB. This connection can go from the surface to any other layer of the PCB or between layers. The via can go all the way through the PCB (through hole), go between the surface and an internal layer (blind via), or between layers inside the PCB (blind via). A via can be used as a component lead for a surface mount part, as a way to transition a signal from one layer to another (routing via) or as the hole for a leaded part.

Vss- The voltage rail of a power supply that is connected to the source end of an MOS or CMOS transistor. This is commonly the "ground" or most negative terminal of the power supply.

WATER ABSORPTION- The property of an insulating material that causes it to absorb and retain water. Two undesirable side effects of water absorption are high leakage currents and blowouts during soldering as the absorbed water boils.

WATT- The unit of power in the International System of Units (SI). The watt is the power to do work at the rate of 1 joule per second. One volt at the current rate of one ampere per second is one watt. $P = EI$

WAVE GUIDE- A metal tube, either rectangular or circular in shape, through which an electromagnetic field is propagated. A wave guide is commonly used to guide microwave signals between the source and the load.

WAVE LENGTH, λ- The length, in space or a dielectric, of a single cycle of a sine wave. It is calculated by dividing the speed of travel of electromagnetic energy in the dielectric or space by the frequency in Hertz. The speed of light in a vacuum is 186,000 miles per second or 300,000,000 meters per second.

WAVE SOLDERING- The process of soldering components to a PCB by passing the PCB loaded with components over a wave of molten solder. This method works best with through hole components.

WIRE- That portion of a net that connects two pins. In a two pin net, it is the entire net. When routed, a wire may be made up of segments.

WIRE BOND- A short piece of very thin wire that connects the terminals of an IC to the terminals or lands of the package that house it. These connections are usually made using ultrasonic welding. The wires are about 1 mil in diameter and up to 30 mils long.

WOM- Write Only Memory. This is a memory that can be written many times, but never read. The author's memory assumes this state many times a day.

XAUI- (10 Gigabit Attachment User Interface). This is a 10 GB/S Ethernet signaling protocol that uses four 2.5 GB/S differential signal paths running in parallel. It was developed by the Ethernet Alliance to provide a standard method for interconnecting products. www.10gea.org.

ZENER DIODE- A kind of diode that "zeners" at a specific voltage when reverse biased. Zener diodes are used to set voltage levels in power supplies and to clip waveforms at specific voltage levels.

λ/20 RULE- A "rule of thumb" manufactured by some EMI "gurus" that suggests that logic ground should be connected to "chassis ground" at intervals of $1/20^{th}$ of the wavelength of some frequency, presumably the clock frequency of a product. This rule is not valid under any circumstances as it presumes two things: That there is a single frequency that is more important than others and that "chassis ground" is magically a neutral EMI surface. This rule usually results in higher rather than lower EMI.

INDEX

APPENDIX 1: BIBLIOGRAPHY

LEE W. RITCHEY
SPEEDING EDGE

May 2003

REFERENCES:

1. IPC-D-317, "Design Guidelines for Electronic Packaging Using High Speed Techniques." IPC. Lincolnwood, IL

2. IPC-L-125, "Specifications for Plastic Substrates, Clad and Unclad, for High Frequency Interconnects." IPC. Lincolnwood, IL.

3. Spitz, S. Leonard, "Specialty Laminates Fill Special Needs." Electronic Packaging and Production. February 1989.

4. Child, Jeff, "Low-Voltage Differential Signaling (LVDS) Reports for Duty" Electronic Design. May 25, 1998.

5. Smith, et al. "Spice Modeling and Circuit Effects Extraction for High Band PWB" Design. IPC Technical Review, May 1986.

6. EDN Staff, "Optimizing PCI Bus Throughput. " EDN, May 21, 1998.

7. Doyle, Greg & Strange, Rod, "Signal Integrity, Why You Should Care About It." Printed Circuit Design, May 1998.

8. Ritchey, Lee W., "A Test PCB Used to Evaluate Affects of Structural Features on Transmission Line Behavior." Spring 1988 Nepcon Proceedings.

9. Snyder, Kenneth C., Planar Circuits Technology, Inc. "Resistor Packaging Options: To Bury or not to Bury." Printed Circuit Design, January 1986.

10. Mahler, Bruce, Ohmega Technologies, Inc. "Planar Resistor Technology for High Speed Multilayer PCBs." Electronic Packaging and Production, January 1986.

11. Martin Marietta Corporation. "Manufacturing Technology for Advanced Data/Signal Processing Guidelines, ADSPMM-88-308-BP." Martin Marietta, July 1989.

12. DeFalco, John A., Honeywell, Inc. "Reflection and Crosstalk in Logic Circuit Connections": IEEE Spectrum, July 1970.

13. Edlund, Greg, "IBIS Model Accuracy." Printed Circuit Design, May 1998.

14. Malack, John A. and Runac, Robert E., IBM. "Stubs Affect Circuit Performance." Electronic Packaging and Production, November 1988.

15. Mahler, Bruce and Signer, Roy M., "Thin Film Resistor Technology." Printed Circuit Design, June 1986.

16. Ritchey, Lee W., Shared Resources, Inc. "Transmission Line Sequencing in High Speed Design." High Performance Systems. June 1989.

17. McDermott, Brian J., "The Practical Application of Microvia Technology. " Circuitree Magazine, March 1998.

18. Brooks, Douglas, "90 Degree Corners, The Final Turn." Printed Circuit Design, January 1998.

19. Davis, Sam, "EMI: Digital Designers Nightmare." Computer Design, January 1998.

20. Small, Charles H., "PC Board Advances Getting Short Shrift." Computer Design, January 1998.

21. Nemec, John, "Take the Mystery Out of AC Termination." Electronic Design, September 1997.

22. Johnson, Howard, "The I/O Buffer Specification IBIS." Printed Circuit Design, May 1997.

23. Johnson, Howard, "Planning for Signal Integrity Modeling." Electronic Design, May 1997.

24. Johnson, Howard, "Probing High Speed Designs." Electronic Design, March 1997.

25. Neusch, Martin, "Primer on High Performance Laminates." EP &P, June 1997.

26. Johnson, Howard. "Operating Above Resonance, A Bypass Capacitor Paper." Electronic Design, April 1997.

27. Shear, David, "Ground Bounce Tests. Excellent Waveforms." EDN, April 1993.

28. Ritchey, Lee W., "Documenting a Design to Get it Built Right and Verifiable." Circuitree Magazine, August 1996.

29. Ritchey, Lee W., "Oh Lord, We Forgot Our Map and We're Stuck in Lodi Again!" (A piece on verifying PCB Cross Sections). Circuitree Magazine, October 1996.

30. Cable, Alan, "Laser Drilling Microvias in Printed Circuit Boards." Circuitree Magazine, November 1996.

31. Ritchey, Lee W., "Crosstalk or Coupling, A Silent Problem." Printed Circuit Design, April 1998.

32. Hubing, Todd H. etal, "Power Bus Decoupling on Multilayer Printed Circuit Boards." IEEE Transactions on Electromagnetic Compatibility, Vol. 37, NO 2, May 1995.

33. Ritchey, Lee W., "A Beginner's Guide to Simulation and Analysis." Printed Circuit Design, May 1998.

34. Kaufer, Steve & Crisafulli, Kellee, "Terminating Traces on High Speed PCBs." Printed Circuit Design, March 1998.

35. Brooks, Douglas, "Bypass Capacitors, A Conversation with Todd Hubing, UMR." Printed Circuit Design, March 1998.

36. Prasad, Ray, "Understanding Surface Finishes, Gold/OSPSn/Pb" SMT, May 1998.

37. Johnson, Howard, "Building a Signal Integrity Department." EDN, June 4, 1998.

38. Ritchey, Lee, "A Survey and Tutorial of Dielectric Materials used in the Manufacture of Printed Circuit Boards." Circuitree Magazine, November 1999.

39. Ritchey, Lee, "Test Lab, Cuts in Power Planes and Soldermask Effects on Impedance." Printed Circuit Design, January 2000.

40. Morgan, Chad and Helster, David, "New Printed Wiring Board materials Guard Against Garbled Gigabits." EDN, November 11, 1999.

41. Kraemer, Doug, “Multi-board System Level Analysis.” Printed Circuit Design, May 1999.

42. Wellers, Fred & Tomasek, Rob, “Mechanical Analysis.” Printed Circuit Design, May 1999.

43. Powell, Jon, “SPICE or IBIS” Printed Circuit Design, May 1999.

44. Smolyansky, Dima & Corey, Steven, “PCB Interconnect Characterization From TDR Measurements.” Printed Circuit Design, May 1999.

45. Ritchey, Lee W., “ How to Design a Differential Signaling Circuit.” Printed Circuit Design, March 1999.

46. Crisafulli, Kellee & Kaufer, Steve, “Terminating Differential Signals on PCBs.” Printed Circuit Design, March 1999.

47. Kowal, Keith, “Signal Integrity Software Use.” Printed Circuit Design, March 1999.

48. Gardner, Sam etal, “Noise Fighting 101, Spread Spectrum Clocking.” Printed Circuit Design, January 1999.

49. Wexler, Al, etal, “The art of EMC, using Simulators.” Printed Circuit Design, May 1999.

51. Ritchey, Lee W., “PCB Routers and Routing Methods.” Printed Circuit Design, February 2000.

52. Strassburg, Dan, “Digital Buses, Analog Problems” EDN, May 27, 1999.

53. EDN Staff, “Line Drivers/Receivers Conform to New LVDS Standards.” EDN, June 10, 1999.

54. Alderman, Ray, “The Death of PCI.” Electronic Design, June 8, 1998.

55. Knack, Kella, “Debunking High Speed PCB Design Myths.” ASIC & EDA, July 1993.

56. Ritchey, Lee W., “What’s Good Enough? High Speed Rules” ASIC & EDA, February 1994.

57. Zanella, Fabrizio, & Galius, Teradyne, “Optimizing Signal Integrity in High Speed PCB Connectors.” E P & P, publishing date unavailable.

58. Shear, David, “Ground Bounce Tests Revisited” EDN, April 1993.

59. Dugan, Vince; Baron, Ben and Fout, Randy, “Reconfigurable Prototyping Speeds Hardware/Software Integration.” ISD, Dec 1999.

60. Lipman, Jim, “Models Make the Difference in High Speed PCB Design.” EDN, April 14, 1999.

61. Cheung-Wei, Lam & Powell, Jon, “Using Simulation to Spot and Fix EMI Problems.” Electronic Design, July 22, 1996.

63. Green, Lynn & Ritchey, Lee W., “Implementing an Internal Signal Integrity Department.” Printed Circuit Design, January 2000.

64. Ritchey, Lee W., “Who Cares about Tg?” Speeding Edge, January 14, 1999.

65. Chen, Ian, “Suppress EMI with Spread Spectrum Timing/Clocking.” Portable

Design, January 2000.

66. Ritchey, Lee W., "Effects of BGA Hole Patterns in Power Planes on High Speed Signal Quality." Speeding Edge, February 2000.

67 . Lamson, Michael, "Packaging Takes Center Stage in IC Design Process." Electronic Design, June 8, 1998.

68 . Bogatin, Eric, "What is Inductance?" Printed Circuit Design, March 2000.

69. Greim, Michael C., "High-end Digital Systems Give a Thumbs Down to Rules of Thumb." EDN, June 5, 2000. (Note: This is a very good article on power system design).

70. Johnson, Howard, "Differential Termination." EDN, June 5, 2000.

71. Ritchey, Lee W., "How Do Vias Affect High-speed Circuits?" Printed Circuit Design, July 2000.

72. Bogatin, Eric, "What is Characteristic Impedance?" Printed Circuit Design, January 2000. (Note: This is a very good article on impedance).

73. Ritchey, Lee W., "Examining Rules of Thumb." Printed Circuit Design, January 2000.

74. Sullivan, Kimberlee, "CAD/CAM Viewers" Printed Circuit Design, April 1999.

75. Desposito, Joseph, "Transceiver Chip Replaces Parallel Backplanes with High-Speed Serial Links." Electronic Design, February 7, 2000.

76. Goldie, John and Nicholson, Guy, "A Case for Low Voltage Differential Signaling as the Ubiquitous Interconnect Technology." Electronic Systems, March 1999.

77. Levy, Yoram, "Signal Integrity Analysis." Printed Circuit Design, May 2000.

78. Smith, etal, "Power Distribution System Design Methodology and Capacitor Selection for Modern CMOS Technology." Published by Sun Microsystems, 1999.

79. Ott, Henry W, "Partitioning and Layout of a Mixed-Signal PCB." Printed Circuit Design, June 2001.

80. Ritchey, Lee W., "What Makes a Circuit High Speed." Future Circuits, September 2000.

81. Ritchey, Lee W., "Buses: What Are They and How Do They Work?" Printed Circuit Design, December 2000.

82. Ritchey, Lee W., "Floor Planners, The Centerpiece of Virtual Prototyping." Printed Circuit Design, July 2001.

83. Ritchey, Lee W., "Anatomy of a Plated Through Hole." Unpublished, August 2001.

84. Morgan, Chad and Helster, Dave, "The Impact of PWB Construction on High-Speed Signals." DesignCon99.

85. Rothermel, Brent R, etal, "Practical Guidelines for Implementing 5 GB/S in Copper Today, and a Roadmap to 10 GB/S." DesignCon2000.

86. Cohen, Tom, etal, "Design Considerations for Gigabit Backplane Systems." DesignCon2000.

87. Patel, Guatam, etal, “Gigabit Backplane Design, Simulation and Measurement.” DesignCon2001.

88. Carter, Mark etal, “Signal Integrity Considerations for 10 GB/S Transmission over Backplane Systems.” DesignCon2001.

89. Cox, Ernest etal, “GTLP in Live Insertion Applications” DesignCon2001.

90. Authors Unknown, “CompactPCI at 66 MHz” DesignCon99.

91. Murphy, Beth etal, “Optical Backplanes- Fantasy or Reality?” Tyco Electronics/AMP.

92. Fogg, Michael W., “Evaluation of Maximum Usable Lengths for Cabled Copper Interconnects.” DesignCon2000.

93. Powell, Jon, “IBIS and SPICE Revisited.” Printed Circuit Design, Oct 2000.

94. Barnes, John R., “Designing Electronic Equipment for ESD Immunity.” Printed Circuit Design, July 2001.

95. Staples, Vaughn, “Design Archiving.” Printed Circuit Design, August 2001.

96. Bogatin, Eric, “The Ramifications of Component Selection.” Printed Circuit Design, March 2001.

97. Pattavina, Jeffrey S., “Bypassing PC Boards: Thumb Your Nose at Rules of Thumb.” EDN, October 22, 1998.

98. Mu, Zhen and Wang, Paul K. U., “Developing Simulation Models for High Speed Design Applications." Printed Circuit Design, Feb 96.

99. “Timing Analysis Tool Listing,” ISD January 2002. Good list of tools available.

100. “Power Analysis Tool Listing,” ISD January 2002. Good listing of tools available.

101. “PCB Design Tool Listing,” ISD January 2002. Good listing of tools available.

102. Bergman, Deiter, “Collective Design Advances IC Packaging.” Electronic Design, June 8, 1998.

103. Mannion, Patrick, “IC Demands Force Concurrent Engineering.” Electronic Design, June 8, 1998.

104. Haller, Robert J., “The Nuts and Bolts of Signal Integrity Analysis.” EDN, March 16, 2000.

105. Boorom, Ken, “Spice Provides Signal Integrity Clues for High Speed Systems.” EDN, February 18, 1999.

106. Ajiluni, Cheryl, “Stop Taking Your Models for Granted.” Electronic Design, March 8, 1999.

107. Tuck, Barbara, “The Hardware/Software Coverification Challenge.” Computer Design, April 1998.

108. Messina, Bruno, “Timing Your PCB Design.” Printed Circuit Design, May 1998.

109. Lombard, Tim, “Automated Drives PCB Success.” Printed Circuit Design, December 1995.

110. Ajluni, Cheryl, “Formal Verification Helps Stamp Out Design Bugs.” Electronic Design, December 1, 1998.

111. Ritchey, Lee W., “Simulation and Analysis Tools.” Printed Circuit Design, May 1998.

112. Donlin, Mike, “Analysis Tools Target Early Stages of PCB Design.” Computer Design, August 1996.

113. Edlund, Greg, “IBIS Model Accuracy.” Printed Circuit Design, May 1998.

114. Johnson, Howard, “The I/O Buffer Information Specification (IBIS).” Printed Circuit Design, May 1997.

115. Johnson, Howard, “Planning for Signal Integrity.” Electronic Design, May 12, 1997.

116. Barbetta, Mike, “Surface Finishes, and OEM Perspective.” Circuitree Magazine, May 2003.

117. Bogatin, Eric, “Get Your S (Parameters) Together.” PCB Design, Feb. 2003

118. Shumpert, Pamela and Henry, Michael, “The Cost of Poor Library Management.” PCB Design, Aug 2002.

119. Smolyansky, Dima, “Lossy Line Simulation and Analysis.” PCB Design, March 2003.

120. Bogatin, Eric, “The Impedance of a Real Capacitor.” PCB Design, Date Unavailable

121. Bogatin, Eric, “The Return Current in a Transmission Line. ”PCB Design, Aug 2002.

122. Bogatin, Eric, “Sounding Off on the db.” PCB Design, Dec 2002.

123. Funer, Rolf, “Embedding Resistors and Capacitors.” PCB Design, December 2002.

124. ANSI/EIA-656-A, IBIS Specification, latest version 4.0 Approved July 2002

125. Fang, Jiayuan, etal, “Effects of 20-H Rule and Shielding Vias on Electromagnetic Radiation From Printed Circuit Boards.” UC Santa Cruz, Publishing date unknown.

IEEE
445 Hoes Lane
Piscataway, NJ
08854
www.ieee.org

IPC
7380 Lincoln Avenue
Lincolnwood, IL 60646
www.ipc.org

APPENDIX 2: ANATOMY OF A PLATED THROUGH HOLE

LEE RITCHEY AUGUST 30, 2001

INTRODUCTION

This document examines the way a plated through hole is formed in a PCB. The discussion holds whether the hole is used by a signal to change layers (usually called a via) or to connect a component pin to a plane or signal line. It also holds when the hole is not plated, such as for mounting hardware. It is important to understand this in order to correctly calculate the sizes of pads and keep out areas needed to insure the PCB is both manufacturable and meets its electrical requirements.

A clear understanding of this subject is necessary in order to satisfy the dual needs of PCB assembly and fabrication. The PCB assembler works with finished hole size. The fabrication process is based on the drilled hole size. Along with all of this, it is necessary to size holes in such a way that the power planes of the PCB are not degraded by placing holes so close together that they cause slots to be created in the planes by overlapping clearance holes. It is not possible to specify generously large hole sizes and insulation spacing to make manufacturing or fabrication easy without risking degradation of the environment needed by the high speed signals traveling across those planes through the PCB.

Figure 1 is a picture showing the important structures involved in making a plated through hole.

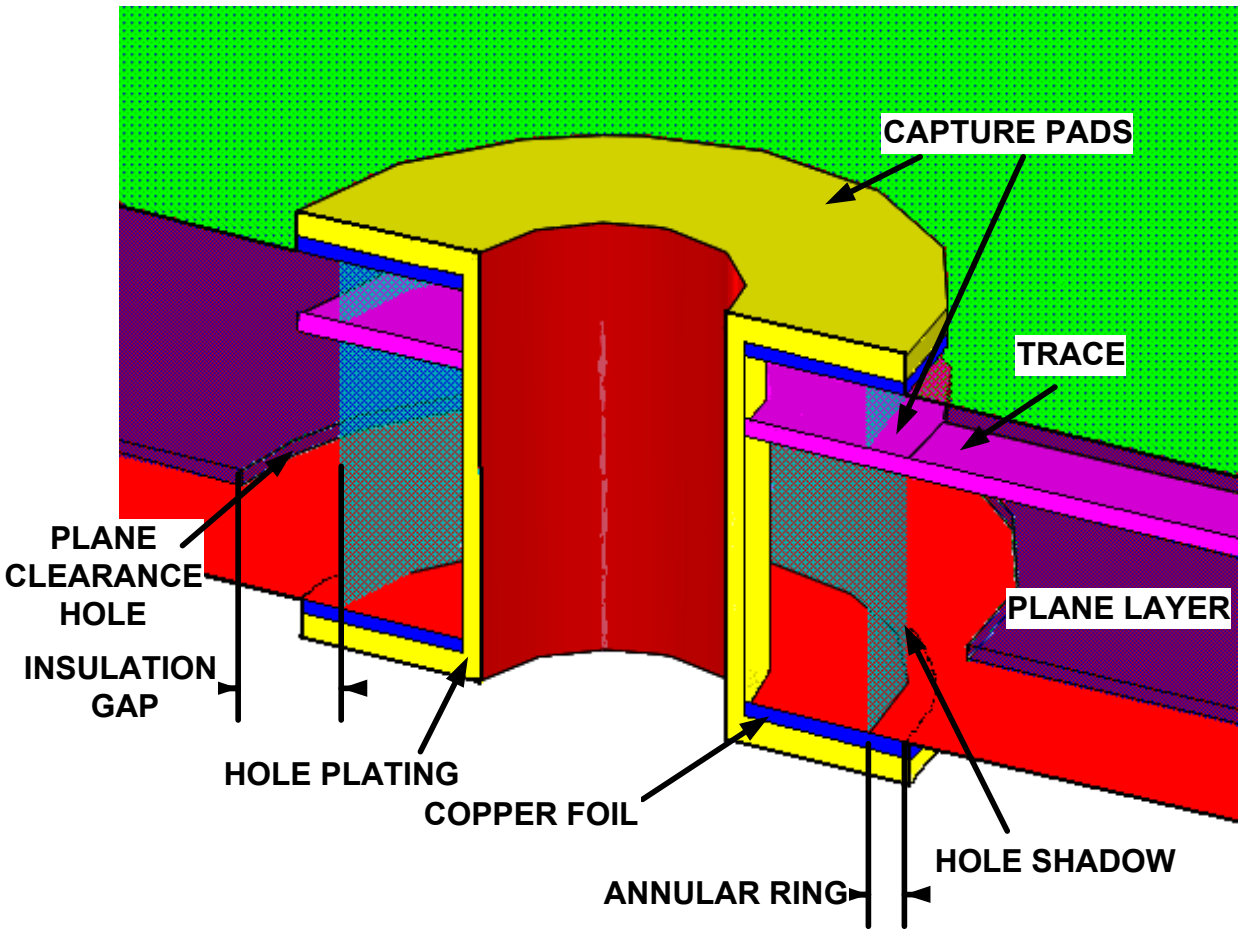

Figure 1. Cross Through Section of A Plated Through Hole or Via

DEFINITIONS

The following terminology is used to describe the various components involved in creating a plated through hole. They apply whether the hole is used as a routing via, a component hole or a power connection.

Finished hole size- The diameter of the plated through hole after all plating steps are completed.

Drilled hole size- The diameter of the hole after drilling.

Capture pad- The pads placed on outer layers to "capture" the plated through hole or on inner layers to connect traces to the plated through hole. Capture pads are used on inner layers only when a connection is being made to a trace. Putting capture pads on inner layers where no connections are being made results in nonfunctional pads that are not beneficial. Note: For mechanical mounting holes that are not plated, there is no need to place capture pads on the outer layers.

Clearance hole- The hole etched in a copper plane to allow a drilled hole to pass through. The hole may be plated or unplated.

Hole shadow- The shadow cast through all PCB layers by the drilled hole. Features in signal layers must be kept away from this "shadow" by a dimension corresponding to the insulation thickness required to meet appropriate standards. The hole shadow is defined by the diameter of the drilled hole and the worst case wander that the hole may have in the final PCB. This "wander" is composed of drilling inaccuracies, layer-to-layer registration inaccuracies, inaccuracies in the working artwork and shrinkage of the laminate layers from the heat of lamination. Because of hole wander, plated copper may be found out to the edge of the hole shadow. Therefore, insulation thickness must start at the hole shadow and extend outward in both the power and signal layers.

Manufacturing tolerance- The dimension that describes how a drilled hole will wander out of its true position. The inaccuracy components are described in the hole shadow description.

Hole plating- The copper plating deposited in a hole to create a connection between signal pins and traces or power pins and power planes. This plating is deposited after the PCB has been laminated and drilled.

Copper foil- The outer layers of all PCBs are formed using copper foil. This foil serves as an electrical path to conduct the plating current needed to plate copper in the holes and on the pads and traces that will be present on the surfaces of the finished PCB. After the plating steps have been completed, the foil is etched to form the traces, pads and other outer layer features.

Insulation gap- This is the insulating material that exists between the hole shadow (copper plating in the hole) and metallic features in all of the layers. Examples are traces and planes.

Annular ring- The copper in the capture pad that extends past the drilled hole shadow. The capture pad is intentionally made larger than the drilled hole and the hole shadow in order to insure that there is always a portion of the pad making contact with the trace even when the drilled hole is out of position due to drill tolerances.

Plane- Any copper plane layer inside the PCB. Could be ground, Vdd or any other plane used to make electrical connections or to serve as partners for controlled impedance lines.

Pad Stack- The name given by CAD systems to describe the sizes of the capture pad, plane clearance, drill size and plating. The methods used to calculate these will be described in this document.

IMPORTANT DIMENSIONS

Finished hole diameter- Drill size minus copper plating.

Drilled hole diameter- Finished hole diameter plus copper plating

Hole shadow- Drilled hole diameter plus manufacturing tolerance.

Capture pad- Hole shadow plus annular ring allowance

Clearance pad- Hole shadow plus insulation gap.

Plane web- The web of copper in a power plane between two clearance pads. See figure 2.

TYPICAL HOLE PATTERN IN A POWER PLANE CAUSED BY A BGA OR PIN GRID ARRAY

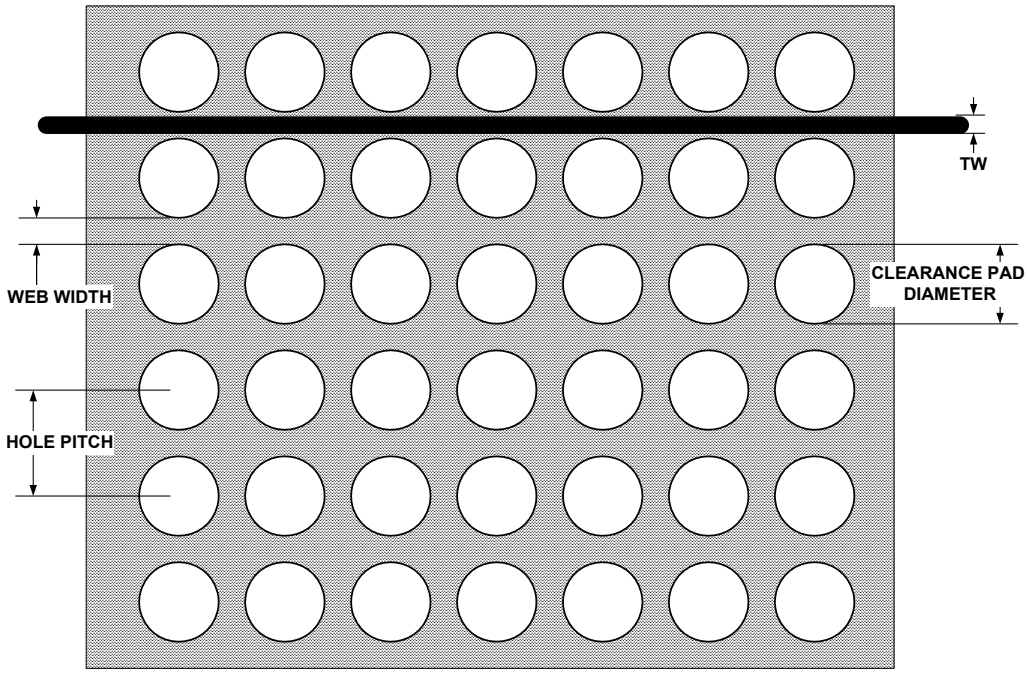

Figure 2. A Typical Hole Pattern Under a BGA.

Insulation gap- On inner layers, the thickness of dielectric needed to separate copper features on the same layer from each other. This could be the distance between the plating of a plated through hole and the plane copper or the distance between the plating in a plated through hole and a trace in a trace layer. For most products, this is 5 mils minimum.

When designing a pad stack for a hole, there are four main considerations. These are:

1. Insure that the pad used to connect a trace to the plating in a hole (capture pad) is large enough that the connection between the trace and the plating in the hole is always made with copper from the pad. When the connection is made only by the end of the trace to the plating (called a butt connection), this joint is weak and will fracture during soldering. The connection is assured by making sure the capture pad is larger than the hole shadow by enough to assure a solid connection. This is specified as a minimum annular ring. Normally this is 2 mils minimum per side. The condition where a hole wanders enough to create a butt connection is called “breakout”.

2. Insure that the hole in the plane through which the hole passes is large enough that the distance between the plating in the hole (hole shadow) and the edge of the plane is large enough to meet the insulation requirements of the design. For almost all PCBs, this will be 5 mils minimum.

3. Insure that the distance between the edge of a trace passing by a via in a signal layer is large enough to satisfy the insulation spacing specification. This is usually the hole shadow plus 5 mils per side. Note: traces must not pass closer to a plated through hole than the web width in the underlying plane. Therefore, the space in signal layers that can be used for signals is only as wide as the web between holes.

4. Insure that the web in the planes between adjacent holes is wide enough to provide good signal integrity for transmission lines passing over them. Note: when the trace is the same width as the plane web, the impedance will increase. In the case of a 5-mil, 50-ohm trace passing over a 5-mil web as shown in Figure 2, the impedance will increase to nearly 80 ohms. It can be shown that when the web increases to at least 1.5 times the trace width, impedance increases will be less than 5%.

PAD STACK CALCULATIONS

The pad stack is made up of four components: finished or plated hole size; drilled hole size; capture pad size and clearance hole size. The following calculations result in the minimum sizes for all features. When possible, pad stacks should be designed with larger tolerances. The minimum size is required with 50 mil pitch and 1 mm pitch BGAs and fine pitch connectors in order to balance the tradeoffs between manufacturability and good signal integrity.

Drilled hole size- Finished hole size plus 3 mils for plating (1.5 mils per side) plus drill size tolerance. Drilled hole tolerance is +- ½ mil based on input from Procket's main suppliers. Because this is such a small error, it does not need to be included. The 12-mil manufacturing tolerance more than accounts for this.

Hole Shadow Diameter- Drilled hole size plus manufacturing tolerance. For 18" X 24" high layer count panels built by the top tier fabricators, the manufacturing tolerance is 12 mils. For 16" X 18" panels built by the same fabricators, the manufacturing tolerance is 10 mils.

Capture Pad Diameter- Hole shadow plus 4 mils.

Clearance Pad Diameter- Hole shadow plus 10 mils.

When calculating the above dimensions for vias (holes that don't have leads soldered into them) the starting point is drill size since the diameter of the finished hole is not important. The minimum drilled hole size that can be reliably plated in very high quantities is the starting point. For 100+ mil thick PCBs, this is 12 mils.

Example calculation for 18" X 24" panel 100+ mil thick PCB:

Hole shadow = drill + 12 mils

Capture pad = drill + 12 mils + 4 mils = drill plus 16 mils.

Clearance pad = drill + 12 mils + 10 mils = drill + 22 mils.

The pad stack for a 12 mil drill is: capture pad = 28 mils, clearance pad = 34 mils.

Using the above data for 50-mil pitch BGAs and a 12 mil drill the plane web will be 50 – 34 mils = 16 mils.

Using the above data for 1-mm pitch BGAs (39.7 mil pitch) and a 12 mil drill the plane web will be 39.7 – 34 = 5.7 mils.

Example calculation for 16" X 18" panel 100+ mil thick PCB

Hole shadow = drill + 10 mils

Capture pad = drill + 10 mils + 4 mils = drill plus 14 mils

Clearance pad = drill + 10 mils + 10 mils = drill + 20 mils

The pad stack for 12 mil drill is: capture pad = 26 mils, clearance pad = 32 mils.

Using this data for 50-mil pitch BGAs and 12 mil drill the plane web will be 50 – 32 = 18 mils

Using the above data for 1-mm pitch BGAs (39.7 mil pitch) and a 12 mil drill the plane web will be 39.7 – 32 = 7.9 mils.

UNPLATED HOLES

An unplated hole is usually used for mounting hardware that is not intended to make connections to any of the circuits on the PCB. When calculating the pad stack for an unplated hole, there will be no capture pads on inner or outer layers. The only calculation will be for the clearance pads that are etched in the power planes to allow the hole to pass through. This calculation is the same as for a plated hole.

THERMAL TIES

When a plated through hole is used to connect the power pins of leaded components to the internal power planes of a PCB, the planes will "sink" away the heat needed to solder and unsolder the component. As a result, it may be impossible to install or remove components. To solve this problem it is necessary to provide some thermal isolation between the plated through hole and the plane. This is done with narrow "ties" or thermal ties between the plated through hole and the plane. See Figure 3.

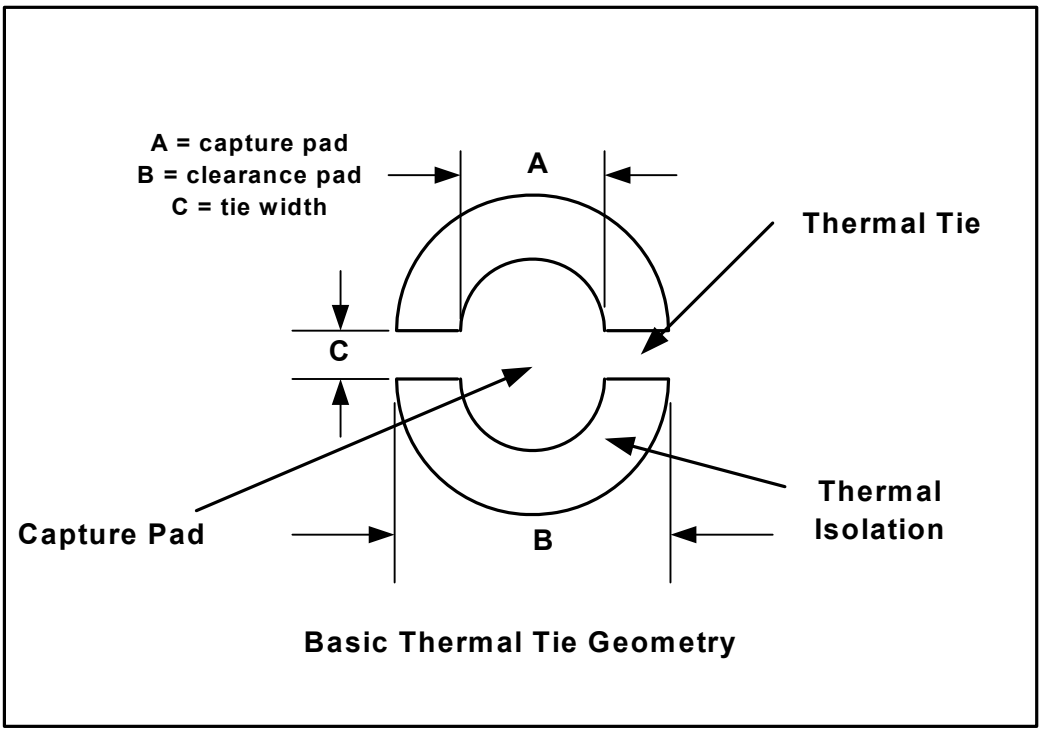

Figure 3. Thermal Tie Design

The capture pad, A, above need only be the size of the hole shadow. (drill + mfg. tolerance) The breakout problem is solved by placing two ties on opposite sides of the capture pad. The clearance pad is calculated the same as for any via (drill + mfg. tol + insulation.)

Thermal ties are not needed to connect the power leads of surface mount components to planes and should not be used.

For a full discussion of thermal ties, both electrical and thermal, see the paper listed at the end of this document.

References: Power Plane Thermal Tie Design, By Lee Ritchey, Feb 16, 2001.d

APPENDIX 3: SELECTING PCB SUPPLIERS

By Lee W. Ritchey Revised: 2/02/03 February 2002

OBJECTIVE

The objective of the PCB supplier selection process is to choose PCB suppliers (fabricators) that are capable of reliably supplying prototype and production quantity PCBs of the types being used in a product at competitive prices. In addition, it is important to choose suppliers capable of providing vital front-end design engineering support in areas such as material selection, design for manufacturability, lamination and plating.

A crucial part of this process is balancing the risks against the costs. Often, suppliers are chosen based on the price of the PCB on the purchase order rather than the cost of ownership of that same PCB. The cost of a PCB includes the cost of assembly, rework, test, repair, scrap, replacement--if defective--and field service costs. **It is rare that the lowest price PCB bid turns out to be the lowest cost.**

It is vital to choose fabricators who can keep the risk of failed prototype and production PCBs to the absolute minimum. Remember, a failed PCB costs a company, not only its cost, but also the cost of all the components mounted on it as well as a large quantity of troubleshooting labor. During prototyping, this cost includes the cost of project delay and delay in time to market. These costs overshadow all of the other costs related to bad PCBs. The PCB itself should never be the reason a PCB assembly is scrapped. If this happens, the PCB fabricator selection process has failed.

PCB fabricators who bid on a particular PCB are rarely equally capable of building the PCB. Worse, it is usually impossible to discern the relative capability of each from the information supplied by the sales force.

Similarly, it is not likely that contract manufacturers, whose strengths are assembly and material management, will possess the technical capability to determine which PCB fabricators will make the best suppliers. As a result, the choices of fabricators that a CM can use must be supervised with the same vigilance that is used when assembly is done in house. **PCB fabricator approval must not be allowed to rest solely in the hands of the contract manufacturer.**

THE PCB FABRICATION SUPPLIER BASE

The complexity of PCBs starts from small single-sided PCBs on paper laminate with punched holes such as those used in remote controls built by the million at the “simple” end. At the “complex” end are large panels with dozens of layers using high performance laminates built in relatively small quantities such as those used in super computers and high performance networking products. The spectrum of PCBs also includes flexible circuits.

As might be expected, fabricators have evolved who specialize in each of these areas. Fabricators good at one end of the complexity range tend not to be good at the other end. The trick in selecting a fabricator is to match its business model, manufacturing tooling and processes, materials inventory and technical expertise to the PCB types that are to be made by it. Implied in this statement is the need to accurately understand the characteristics of each PCB and to understand what characteristics of each fabricator candidate best match its needs. Also implied is that the characteristics of each PCB supplier needs to be understood. This is often aggravated by the fact that PCB sales people often don’t know the limits of their factory well enough to make sure the PCB in question matches the capability of the fabrication shop. There is an old saw “The difference between a used car salesman and a computer salesman is that the car salesman knows when he is lying to you”. This, all too often, applies to PCB salesmen or their reps. Furthermore, there are always a few commission-driven sales people who will take orders that don’t fit their companies capabilities just to make a sale. Luckily, these are rare.

There is a mindset in much of the PCB fabrication industry that “we’ll try anything”. When a try fails, the response is often that the PCBs will be remade free or no charges will be made. This is a way for a fabricator to extend its reach into new technology areas. However, it is always done at the expense of the customer. Such failures are never free. They always carry with them schedule hits while the PCB is remade. This cost can easily be the cost to operate the design team while replacement PCBs are made. Actually, the cost is usually much higher as the delay often translates to delayed time to market--the real schedule and cost driver.

Many fabricators willingly build prototypes that cannot be manufactured in volume due to excessively tight tolerances or reliance on process steps that don’t scale to volume. Often, the customer doesn’t know this until it is time to scale up volumes. Examples of this are two traces between pin routing on 1mm pitch BGAs and 8 mil holes on thick PCBs. As this is being written, there are PCBs being built for multi-gigabit products that are failing as a result of this kind of prototyping. It is imperative that

this “cowboy” mentality be spotted before committing a product to a new technology. One way to do this is to look for examples of the proposed technology in volume manufacture. I always look for examples of PCBs of the type being considered passing through the production process when I am surveying a potential fabricator.

Once a new fabricator has demonstrated the capability to build the PCBs in question, the safe approach to adding that supplier to the approved supplier list is to have a trial set of PCBs built that are not in the critical path of any project. Examples of this are splitting a production order with an already approved supplier so that if anything goes wrong there will still be some good PCBs coming out of the process.

One of the biggest differentiators among fabricators is the skill of the people running the operation. It has long been known that good, skilled people can make superior PCBs, even with the poorest of equipment. The reverse is not true. The fact that a fabricator has the most advanced equipment available or the best process control system doesn’t guarantee good PCBs.

On top of everything else, many fabricators are made up of several plants scattered over a wide geographic area, sometimes globally. Nearly always, these plants are acquired by purchasing existing companies. Normally, the capabilities of each plant will be different. Some will be good at high complexity PCBs while others will be capable of only low complexity PCBs. Still others may be chartered to build complex PCBs when their capability is limited to low complexity PCBs. With fabricators of the multi-plant type, it is likely that only some of the plants, or perhaps only one, will be capable of manufacturing PCBs of the type being considered. If this is so, it is important to restrict **in writing** the manufacture of the PCB to only that plant. Failing to do this can result in the manufacture of the PCB being shifted to a less-than-capable plant. All too often this shift is only discovered when the yields after assembly drop to an undesirable level.

RISKS ASSOCIATED WITH INCORRECT FABRICATOR SELECTION

There are a number of ways that a PCB can turn out bad. The simplest is that the bare PCB costs too much. This is the result of selecting a fabricator whose processes are more robust and complex than the PCB warrants. This is the easiest, lowest cost problem that results from improper fabricator selection. It is also the easiest to fix. Just select a fabricator whose processes are just good enough to yield a PCB of the required quality.

The following types of failures can result from improper fabricator selection. Vias and plated through holes fail creating open circuits. Solder joints fail as a result of poor surface finishes. Shorts can be present in the PCB when it is new or grow due to metal migration. Opens can be present in a PCB when it is new or occur due to metal fatigue. Impedance can be the wrong value for a variety of reasons. The PCB may fail HIPOT testing due to leakage caused by contamination.

The failures in the aforementioned paragraph are the most insidious. The reason is they usually don’t show up until the PCB has been assembled. Often, they are never found. The PCB sits in a “dog box” on the manufacturing floor because no one can fix it or it goes out to a customer and fails there. Very few manufacturers have the failure analysis and failure-tracking systems in place that can isolate these kinds of failures well enough to trace them back to the PCB fabricator. As a result, bad PCBs are purchased for a much longer time frame than they should be.

It is these risks that make PCB fabricator selection such an important part of the supplier evaluation process. Sadly, the lack of visibility of these kinds of failures leads to the assumption that price is the only selection criteria for a PCB fabricator. As Ruskin points out, the common law of business states that you cannot pay a little and get a lot. If a PCB fabricator quotes a price noticeably lower than its competitors, odds are it is because of shortcuts, not superior skill. In this case, it is wise to set aside a reserve for the likelihood of PCBs being bad. When the reserve is large enough to cover the risk, it will amount to more than the potential savings. Said another way, beware of the low bidder in a bidding group of equals.

Another reason for an excessively low bid is the deliberate “buying” of the business to get started with a new customer. Accepting such a bid carries with it two risks. The first is that the order may go to a less-than-qualified supplier with the result being bad PCBs. The second is once the link has been established and the customer is dependent on the supplier, prices will escalate beyond what the other bidders quoted.

All of the above discussion seems to imply that the PCB supplier team needs to visit a potential supplier prior to giving it an order. The more complex the PCB is, the more important this survey is.

Along with the need to visit potential PCB suppliers to determine which are capable of building a given PCB, there is the need to real time monitor the deliveries to insure that the quality remains high enough to insure successful manufacture. Due to the fact that personnel change and business models change at fabricators, it is necessary to have periodic re-surveys of existing fabricators. In times of slow downs, such as happened in the second half of 2001, fabricator capability can change rapidly to the extent of some suppliers disappearing altogether, as happened with McCurdy Circuits.

WHAT TO LOOK FOR IN A PCB FABRICATOR

Printed circuit board fabrication is one of the most complex processes involved in the manufacture of electronic products. IC manufacture is perhaps the only manufacturing process that involves more disciplines and more process steps.

PCBs Being Built Just Like Mine

The best and first test of a potential fabricator is to see PCBs of the type being **considered flowing down the production line on a daily basis**. If possible, talk to the customers using those PCBs to see how well they are working out. **Seeing PCBs like yours in the conference room or hanging on the wall does not guarantee the capability of the fabricator**. All too often, the PCBs on display are rejects that have defects that can't be seen visually. Bad PCBs are often displayed in booths at trade shows. At PCB West in the spring of 2002, one vendor had several PCBs on display that contained failures from a well-known router vendor.

Sales Engineering Support For Development Engineers

An important component of the fabricator's contribution to a good design, is engineering support in the form of advice on materials selection, design rule validation, manufacturability advice and other technical support that is not usually part of the design engineering team's kit of tools. Good support in this area can make the difference between a design that is manufacturable and reliable and one that is not. The more complex the design, the more important this is. The best engineering support comes from fabricators who are routinely building PCBs of the complexity in question. They will have seen many things that can improve the yield of a PCB that a less experienced fabricator has not. **This knowledge is gold.**

Materials Engineering Knowledge

There is a very broad range of potential materials from which PCBs can be built. Most of these materials are good materials. They have been developed to meet a particular need. As an example, paper-based materials have been developed to allow very low cost PCBs to be made for simple consumer electronics. Very low loss; low dielectric constant materials have been developed for microwave PCBs. When a material designed for one application is mistakenly used in another, the result is often a failed product.

The more complex a PCB is, the fewer fabricators will have engineers who understand how to select and work with the proper materials. Among the things that must be understood and worked with are how a material shrinks when it goes through the heat of lamination and cool down; how it behaves in the lamination cycle; how it drills; how it plates; how the dielectric constant, er, varies across the various thicknesses of laminate and how it varies with frequency. For very high data rate products such as OC-48 and higher, losses in the dielectric and loss tangent become important as does resistive skin effect losses in the traces.

As controlled impedance becomes more widespread, the engineering team at a fabricator must know how to accurately compute stack ups so that impedance of the finished PCB is within limits. This part of the process requires skill in the use of 2D field solvers to accurately calculate impedance and good working knowledge of the dielectric constants of the materials to be used. Fabricators who still rely on equations to calculate impedance will routinely get the impedance wrong, not because of calculation errors, but because the equations themselves are not accurate at the dimensions currently used in PCBs. As a result, equation-based impedance calculating methods require trial and error to "tweak" the results in order to arrive at a stack up that yields the right impedance. This is an iterative process that may take multiple fabrication runs. **Using 2D field solvers and accurate materials information results in the right impedance the first time.**

Laminate Materials Types:

A wide variety of dielectric materials are available from which to fabricate PCBs. All are composites of some form of resin system and a reinforcement. They can be divided into two general classes. These are glass reinforced and non-glass reinforced. Woven glass provides dimensional stability to the PCB. All high layer count PCBs must have substantial XY dimensional stability and require woven glass as part of the laminate. (There is one non-glass woven reinforcement that provides this XY stability. It is Gore-Tex® fabric made by W. L. Gore. This material produces a slightly lower dielectric constant composite, but at a very high cost premium. This cost premium does not carry with it a corresponding benefit of equal value.) Other reinforcements include random fibers of Kevlar® and Gore-Tex®.

Virtually all logic PCBs, no matter what the data rate, require glass reinforcement in order to provide the necessary dimensional stability needed by high layer count PCBs.

The following are some examples of glass-reinforced laminates.

Low Tg FR-4-	Tg >135°C used on thin PCBs, (.062" and thinner)
High Tg FR-4	Tg >170°C used on thicker PCBs, (greater than .062")
Getek®	Tg = 185°C said to be high performance, but proved difficult to process
Nelco® 4000-13	used on high performance PCBs
Nelco® 4000-13SI	used for low loss PCBs, uses S style glass for low loss
Allied/Isola® FR 408	Getek® equivalent
Allied/Isola® FR 406	Allied version of high Tg FR-4
Nelco 6000-Hi Speed PPE	low loss, but very difficult to process
Isola IS620	used for low loss PCBs, low loss achieved with resin, uses E glass
ZBC®	A Hadco/Zycon® material aimed at creating a large plane capacitor

These are non-glass-reinforced materials.

Rogers RO 4350	used for microwave PCBs
Arlon CLTE	used for microwave PCBs
Teflon based cores	used for microwave PCBs
Speed board	used with other laminates to create selected low loss layers

When multiple types of materials are stocked at a fabricator, there must be some rigorously defined method for insuring that wrong materials are not included in a PCB build. The most common error occurs when a fabricator stocks both low Tg and high Tg materials of the same resin system, such as FR-4. If low Tg material is accidentally used in a PCB that needs high Tg material, the result will be PCBs that have failed vias after soldering. This is a disastrous, very expensive mistake, as it results in scrapped assemblies. **Experienced fabricators don't stock low Tg materials for this reason.**

Staff Training

As pointed out earlier, printed circuit board fabrication relies heavily on the skill of the people operating the machinery and maintaining the chemistries involved in the process as well as the skill of the personnel preparing the tooling sets. As with most manufacturing processes, a PCB fabrication process that consistently manufactures high quality PCBs does so only with well-trained and well-motivated people. The more complex the PCB, the more skill is needed for each process. Evidence of ongoing training and certification for each operation should be visible at each station at a good fabricator.

Overall Process Control

PCB fabrication involves more than a hundred process steps. These range from creating photo-tooling, to a variety of chemical cleaning steps, several plating steps, drilling, lamination, testing, coating, materials management and tracking the location of each order in the process. A good fabricator will have the activities involved at each step documented at each station.

Very high on the process control list of important items to monitor is all of the chemistry at every place that a chemistry is used. The best monitoring is real time with automatic systems that have alarms for out-of-range conditions. Charts should be posted at each step showing how the process is performing on a daily basis with limits shown so that operators can see how their actions affect the operations.

Chemistries that are not used frequently, such as electroless nickel/immersion gold, can go out of balance to the extent that plating quality is not satisfactory. Therefore, processes such as this that are essential to the successful manufacture of a given PCB, must be used and monitored on a regular basis. (This fits into the concept of "PCBs just like mine.")

Front End Processing and Tooling Creation

Front end processing and tooling creation is often referred to as the CAM or computer aided manufacturing part of PCB fabrication. This is where the customer CAD data, fabrication specification and fabrication drawing come together. This activity prepares the job traveler, materials requirements and manufacturing tooling. Manufacturing tooling consists of the photo tools, drilling files, laminate, foil and prepreg kits, lamination schedules, plating schedules, test fixture design files, test files, and silkscreen/legend artwork and instructions.

Data files from customers can arrive by any of several transport methods. A customer can automatically transfer them to a server or it may be requested from a bulletin board at a customer site where a designer has posted them. The safest form of transfer is by request from a fabricator after having been notified by manufacturing that the data is ready. Key to the reliability of this operation is very good revision control over all of the data files involved in a design. When revisions must be made after a design has been transferred to a fabricator, it is advisable to roll the revision level of all the files in the data set and send a completely new set. **Making small changes in a data set already in the hands of a fabricator carries with it the risk of inaccurately built PCBs and should be avoided.**

Well-run front-end processes have travelers that are prepared for each type of PCB handled by the shop. There is also a process engineer who examines the design to make sure that it can be built with the data supplied.

Data formats include:

Gerber data in several configurations
RS 274-S
RS 274-X
DXF
Gencam
ODB++
Barco DPF
HP-GL
Agilent EEs of Mask
Excellon I & II

RS-274 X is the most universal. However, many fabricators have used Gerber formats for so long that changing to RS 274-X is difficult. This format has many advantages over the other formats and should be encouraged.

ODB++ is the output of the Valor CAM system and is favored by many fabricators. The problem with this format is that it is not "universal".

Design Rule Checking

Design data from a PCB design system consists of data files for each artwork layer, artwork for silkscreens and solder masks, drill data and a fabrication drawing. The first step in the fabrication process is checking that all of the design files are error free, no files are missing and that the appropriate manufacturing tolerances have been included in the design.

A key feature of the DRC operation is the extraction of a net list from the Gerber data that shows how the PCB will be connected if built from the data as received. This net list can then be compared to the net list derived from the schematic--the CAD net list (usually IPC-354 format).

This net list compare must be required for all complex PCBs. A second requirement should be the action taken when there is a disagreement between them**. If a disagreement happens, the job should stop until the disagreement is cleared up.** There is no point in building a known bad PCB.

Other activities involved in design rule checking involve verifying that clearance specifications are met; that pad diameters are large enough to insure reliable connections and that all layers register properly to each other.

Panel Sizes

PCBs are manufactured in panels. These panels are usually sized to use the standard sizes of laminate materials. PCB fabricators create special tooling for each panel size to which they are committed. It is important to insure that the fabricator candidate has the panel size tooling in place for the PCBs being considered. Failure to do so may result in more expensive PCBs if a larger panel size must be used or delays may result as the fabricator procures the right size tooling and learns how to handle it.

A side consideration when discussing panel sizes is to insure that design engineering knows what the optimum panel sizes are so that PCB dimensions can be adjusted to make maximum use of the standard panel sizes already in use at a fabricator. The list below has some of the standard panel sizes found in PCB fabricators. The most common is 18" x 24". The usable area in such a panel is 16" x 22", so a PCB of that size would make maximum use of the materials. Moreover, it would insure the widest number of fabricators could be considered as candidates. The bigger the panel size, the fewer fabricators will have tooling.

18" x 24"
24" x 30"
21" x 24"
16" x 18"
24" x 32"

In Asia and Europe there are metric equivalents of these panel sizes that are slightly different in dimension.

Lamination

Lamination is the process step where all of the layers in a multilayer PCB are "glued" together. There are three components involved in lamination. These are the etched inner layer laminates, the prepreg "glue" layers and sheets of foil that form the outer layers. The glue is normally the uncured resin in the prepreg components of a PCB. There are several methods used in lamination. Some are:

Lamination with heat and pressure only
Lamination with a vacuum bag surrounding the stacked up layers, i.e., "Turkey bag"
Lamination with the stacked up layers in a press that is inside a full vacuum chamber.

The purpose of the vacuum is to evacuate all of the air from the stack of laminates, foils and prepregs so no air bubbles are trapped inside after the resins have cured. The more critical the application of the PCB, the more important it is to insure all air is removed. One way to do this without a vacuum is with very high press pressures. High press pressures tend to distort inner layers causing misregistration**. Therefore, the higher the layer count of a PCB, the more important vacuum lamination is.**

A second part of the lamination process is the method used to cool down the PCBs after lamination. The simplest method is to remove the PCBs from the press and let them cool in open air. The problem with this method is the lack of a way to control the cool down rate so that the PCB cools uniformly. The result can often be warped PCBs (This is one of the more common reasons for warped PCBs). The alternative to open cool down is to transfer the laminated, hot PCB into a special cool down press, maintain pressure on the PCB and cool it using a gradual lowering of the temperature until room temperature is reached.

There are four types of lamination in use. These are foil lamination, mass lamination, sequential lamination and laminate only lamination.

Foil lamination is the most common method used to produce PCBs of more than two layers**. It is the preferred method for making multilayer PCBs.** It involves etching inner layer pairs back-to-back on pieces of laminate. There will be one laminate piece for each pair of inner layers. These laminates are then stacked on a plate with pins that register the layers to each other. The laminates are separated from each other with prepreg, (the glue layers) and a piece of copper foil is added to each side.

An alternative to foil lamination using plates with tooling pins for alignment of the inner layers, used in some plants in Europe, is to stack all of the inner layers and the prepreg that separates them into a book, secure them together with rivets, add the top and bottom prepregs and foils and then laminate without using alignment pins on the press plates. The argument for using this method for laminating multilayer PCBs is that there is no need for specially tooled press plates, so any size panel can be processed. This is true. However, layer-to-layer registration is not as precise as when foil lamination is done with tooling pins to align the layers. This is often called mass lamination but this is a misnomer. **It should not be used on high layer count PCBs.**

Mass lamination is a method for laminating four layer PCBs in very large quantities. The inner two layers of the PCB are imaged and etched back-to-back on a very large piece of laminate, often 36" x 48" or larger. Many PCBs can be imaged on a single piece of laminate. This etched inner layer pair is combined with pieces of prepreg and foil on each side and laminated. After lamination, targets etched on the inner layers are located in order to line up the drill to each individual inner layer pattern. This local alignment allows very large panel sizes to be used without concern for drill and image tolerance build up across them. This lamination method is used to make very low cost four layer PCBs such as those used for PC motherboards.

Sequential lamination is a method used to build up a PCB when buried vias are required. It can also be used to create blind vias. The layer pair with the buried vias or blind vias in it is drilled, plated and etched using the same process steps involved in making a two sided PCB. This piece of laminate is then combined with the other inner layers, prepreg and foils and laminated using the standard processes.

Cap layer lamination is the original method used to create multilayer PCBs. With this method, the layers are etched in pairs, starting with the outer layers. There is one layer pair set for each two layers in the PCB. Because the outer layers must be solid copper after lamination and drilling to provide a path for plating current, the outer layers are not etched before lamination. The result is a more expensive PCB due to the need to etch one more layer pair than would be required by foil lamination. For this reason, this method is rarely used.

Registration

There are six places in the fabrication process where registration of images to each other is required. These are:

- **Registration of two layers to each other on opposite sides of a piece of laminate**. This is done by pinning the artwork to the exposure frame used for "printing" the images on the laminate surfaces. At very well controlled operations, the layer-to-layer registration accuracy can be within 1 mil.

- **Registration of all the inner layers to each other during lamination**. The most common method for doing this is by stacking all of the inner layer sets on a precisely made plate with alignment pins that mate with precisely punched alignment holes on the inner layer sets. There are two ways to create the alignment holes on the inner layer pieces. One is to punch the alignment holes in the raw laminate prior to exposing the image. This is called "pre-etch punching". The second is to

expose and etch the images on both sides of the laminate. After etching, a special machine picks up targets on the etched image and punches the alignment holes. This is called "post-etch" punching. **Post etch punching is the most accurate of these**. Using this method, layer pair-to-layer pair registration can be held to an accuracy of 2 mils.

- **Registration of the drill to the inner layer images.** Drilling happens after lamination. The drills need to be aligned as accurately as possible so that they hit the pads on the inner layers. One method for doing this is to align the drill to the same holes that were used for lamination alignment. Unfortunately, this does not account for any tiny shifts of inner layers during lamination. For the best fit on tight tolerance PCBs, targets are etched on all of the inner layers that should fall one on top of the other if the lamination was done perfectly. These targets can be x-rayed after lamination to see how well they are lined up. If there is any misalignment, the x-ray image can be used to fine-tune the drill alignment. **This is the preferred drill alignment process for very high layer count PCBs.**

- **Registration of laser drilled blind vias to inner layer images.** Like through hole drilling, this drilling happens after lamination. Blind vias are used on parts with lead spacing too small to allow through hole drilling. These small holes are often accompanied by smaller than normal pads. These pads are too small to allow for the tolerance build up available to through holes. As a result, there needs to be a method for improving drill accuracy. This can be accomplished by placing laser drill alignment targets on layer 2 near each component pattern requiring blind vias. When this is done, the laser drill first burns away the copper on layer 1 over these targets. This is called "skiving". The target on layer 2 is used to locally align the laser drill to the pattern on layer 2.

- **Registration of the outer layer images to etch the outer layer features.** The outer layer features that must be lined up with the inner layers are the pads through which the holes have been drilled. This alignment should be done with the same tooling that was used to align the drill.

- **Registration of the solder mask to the outer layer images.** The soldermask must align with the pads that will be used to solder components to the PCB. This alignment should be done to targets that reflect where the pads are. This may not be the same pad set as those that have the holes drilled in them.

- **Registration of the silkscreen or legend images to the outer layer features**. The silkscreen and legend artwork needs to be aligned with the component mounting pads. Therefore, it should be registered with the same method as the solder mask.

Drilling

Holes in PCBs can be formed in three ways: standard mechanical drilling, laser drilling and chemical etching.

Mechanical drilling is used to create all through holes and can be used to create blind and buried vias. Laser drilling and chemical etching are used to produce blind vias. **Blind vias** are those that start on one side of a PCB, but do not pass all the way through. (By definition, **microvias** are vias with diameters of 8 mils or less. They may be blind or through hole.) **Buried vias** pass between two layers buried inside a PCB.

Aspect ratio, or the ratio of hole diameter to length, is a key parameter in the successful plating of through holes and blind vias.

Most good fabricators can properly plate through drilled holes of an **aspect ratio 8:1** or less with no special process control in very high quantity. When the aspect ratio exceeds this number, extra care must be taken with hole wall cleaning, maintenance of the chemistry in the plating solutions and the plating currents in order to achieve uniform plating through the entire hole wall of all the holes. At very high aspect ratios, above 10:1, it is likely that reverse pulse plating (RPP) will be required to insure reliable plating.

Plating blind vias that are formed by any of the drilling processes listed above requires an aspect ratio of 1:1 or less for reliable plating. In other words, blind vias must not be deeper than they are in diameter. In most cases, it will be necessary to insure that the blind via hole diameter is at least 1.5 times its depth.

An exception to the above restriction on blind vias is blind vias formed by sequential lamination. **Vias formed in this manner are the most expensive and should not be used.**

Blind vias formed by etching is a process that lends itself to PCBs with very large numbers of blind vias. The blind vias are formed using a photosensitive dielectric layer similar to solder mask on the top and bottom of a PCB. This material is exposed to a photo image that allows holes to be etched into the dielectric where vias are needed. Once the blind via holes have been etched down to the underlying layer, plating is done as for any other via. This method is most often used when "building up" a PCB layer by layer. **It does not lend itself to large, high layer count PCBs.**

The **Mechanical drilled hole diameter** can range down to as small as 8 mils. As might be expected, the smaller the drilled hole, the more difficult the drilling is and the more expensive the drill bits are**. The smallest drilled hole that can be considered "routine" in PCBs thicker than .032" is 12 mils.** Below that size, extra effort and care must be taken to insure bits are not broken off in a PCB and that the holes are clean after drilling.

Plating Processes and Types

Plating is required on a PCB for two reasons. The first is to form the copper in the drilled holes or vias. This plating makes the connection from one side of the PCB to the other and to the inner layers. The second is to provide corrosion protection on the copper pads in order to preserve the surface in a state that will take solder.

Plating takes two forms, electroless and electrolytic.

Electroless plating is accomplished by using chemistry that precipitates a metal onto the PCB without using an electric current. It is employed two places in the process. These are: to deposit copper in the drilled hole, which is lined with non-conductive material, in order to provide a scaffolding onto which to electroplate the hole copper and to deposit one of the surface protecting metals after soldermask has been applied.

Electrolytic plating (electroplating) is used two places as well. These are to plate the copper in the holes to form vias and to plate surface protective metals before outer layer etching and the final solder mask is applied.

TYPES OF PLATED METALS

Electroplated copper- Plated on a PCB to form the copper barrels in the vias and other drilled holes. This can be plated with a direct current (the most common method) or by a process called reverse pulse plating (RPP). RPP was developed to improve the plating of very small holes (less than 12 mils in diameter) or holes with very high aspect ratios (greater than 10:1). Copper plated using electroplating is more ductile and durable than if plated using the electroless method.

Electroplated Tin/Lead (solder)- This plating is the most common method for providing a protection for the copper in the holes and on the traces as the unwanted copper is etched away to form the outer layers. It can be left on a PCB with solder mask applied over it. If this is done, the plated and etched PCB is usually submerged in a bath of hot oil to fuse the tin and lead into the alloy known as solder. This finish has been all but abandoned because the solder under the solder mask melts during assembly causing shorts.

If the PCB is a backplane, the plated tin/lead can be left on and not reflowed. This provides corrosion protection for the copper in the holes and the tin/lead acts as a lubricant for insertion of press fit connectors**. This is a preferred method for finishing backplanes as it does not subject the PCB to the thermal shock of reflow or HASL.**

Electroplated Tin- Many fabricators are switching from tin/lead as an etch resist to pure tin. This is in response to the drive to eliminate lead from the manufacturing process. This is a good finish for backplanes for the reasons stated for tin/lead. However, if a PCB is to be soldered using reflow or wave soldering, it is possible for the tin under the solder mask to melt causing the mask to flake off. (Solder melts at 185ºC while pure tin melts at 232ºC. Good control over soldering temperature can keep this from happening.) **This is also a preferred finish for press fit backplanes.**

Electroplated gold over electroplated nickel- This plating serves two functions. It is used as the etch resist. The gold provides corrosion protection for the soldering surfaces. Nickel is plated on first to provide a barrier between the gold and copper. If this is not done, copper will alloy with the gold and solderability will be lost**. This is the lowest risk surface finish for high density PCBs with fine pitch components.** There is still a risk associated with this finish. The risk is brittle solder joints if excess gold is plated on. The gold plate must be kept at 10 micro- inches or less to avoid this problem.

Electroplated Palladium- This metal is sometimes plated between the nickel and gold in the above finish. It is said to improve the performance of the gold. It is not clear that this helps. It does not degrade the plating of the other metals and should be allowed if a fabricator has it.

Electroless Nickel/Immersion Gold (ENIG)- This finish is applied after the soldermask has been applied over bare copper on the traces. (The solder plated onto the traces and in the holes to provide an etch resist is removed so it won't reflow under the soldermask.) It provides the corrosion protection benefits of gold. It is electroless because it is applied after etching of the outer layers when there is no path for electroplating. The chemistry involved in this plating method is complex and can go out of balance if not well monitored. The result is a failure mechanism called "black pad". Solder joints fail from this, resulting in unreliable PCBs. **This is not a preferred finish for complex PCBs that must exhibit high, long term reliability.**

Electroless Tin- This finish is pure tin applied as with ENIG and for the same reasons. It does not have the disadvantages of ENIG. It is also cheaper. The problem is that the tin layer is very thin. When one side of an assembly is reflowed, copper alloys with the tin on the other side making the second solder operation unreliable. **This is a good finish for PCBs with single sided assembly.**

Electroless Silver- This finish is similar to electroless tin. It is said to be less susceptible to second side assembly solder defects. However, there is not enough data collected to substantiate this.

OSP- While this is not a plated finish, it serves the same function--to protect against corrosion prior to soldering. The most common form of this is Entec 106. It is an organic coating that serves as a flux during soldering. The problem this coating has is durability. If it is touched, it fails. If there are two soldering steps, the coating may break down enough during the first solder step to cause defects during the second step**. This is a finish that fits low cost, high volume PCBs that are not subjected to handling. It should not be used on complex assemblies**.

Plating Process Controls

Plating is a chemical process. Reliable plating requires real time monitoring of the solutions used in plating. The ideal plating process is automated with computer control of plating currents, plating times and plating chemistry. Hand monitoring and hand plating can result in high quality PCBs. However, this method is dependent on the skill and alertness of each operator. The best plating operations are automated.

Testing

The more complex a PCB, the more likely shorts and opens may occur. To insure that PCBs with these defects don't find their way onto the assembly line, it is necessary to test each one to some standard. There are three possible standards. These are:

The standard may be a learned connectivity based on testing a sampling of PCBs on a bed of nails tester. This is known as "golden board" testing.

Using a CAM station, it is possible to extract a net list from the Gerber data, a Gerber net list.. This list can be used to test bare PCBs. It is more accurate than golden board testing, but there is still the potential for shipping defective PCBs if there are errors in the Gerber data.

The connectivity of a PCB is intended to match the net list derived from the schematic. This is known as CAD net list testing. It is the most accurate and should be used as the preferred method.

It is possible to compare the net list from the schematic to that derived from the Gerber data prior to building a PCB. This eliminates the chance that a PCB will be built with a defect that is only discovered after fabrication. **This "net list compare" should be a standard part of tooling a multilayer PCB.**

Test Fixturing

Bare PCBs can be tested with a bed of nails test fixture or with a flying probe tester. Each has advantages. Flying probe testing does not require the creation of a test fixture allowing rapid testing of "quantity of one". However, it is slow and doesn't work for volume manufacture.

Bed of nails testing requires tooling time and cost. It is fast and the preferred method for production test. There are single sided testers and "clam shell" testers for double sided testing.

When surveying potential fabricators it is important to insure testers are available with enough pins and bed sizes large enough to test the PCBs being considered.

SUMMARY

This document paints a rather pessimistic view of the PCB fabrication industry. It makes the industry appear rather undisciplined and lacking in ethics. On the whole, PCB fabricators intend to do the best job they can with the resources at their disposal. Unfortunately, PCB fabrication has been viewed as a somewhat trivial part of electronics by the users of PCBs and good performance has not been well rewarded or recognized. As a result, less capable suppliers have been awarded contracts while the more capable ones are driven out of business. Sadly, this has been self-defeating on the part of the manufacturers buying the PCBs. This is because the lower bidders deliver inferior products that cause yield problems and, sooner or later, go out of business themselves. After this is discovered, usually at costs far greater than any savings when buying the PCBs, the vendor is disqualified and the search goes on for a replacement.

The process outlined in the above paragraph goes on to such an extent that the capabilities of a fabricator one year might change dramatically the next. Some may even disappear as happened with McCurdy Circuits in 2001. Some may change their business models completely as Multek did in 2000, going from the best high tech supplier in the world to a mid tier supplier aimed at supplying the needs of the contract manufacturer that purchased it.

It is not my intention to smear the industry, but rather to emphasize the role buyers of PCBs must play in insuring that properly qualified fabricators are chosen for each PCB type and that this supplier base is monitored on a continual basis.

SAMPLE FABRICATOR SURVEY SHEET

CHARACTERISTICS OF FABRICATORS QUALIFIED TO BUILD CUSTOMER PCBS

Customer needs three types of PCBs. These are: daughter PCBs made with Hi Tg FR-4, daughter boards made from Nelco4000-13Si and backplanes made from Nelco 4000-13SI.

VENDOR BEING SURVEYED_____

Requirements for all PCBs.

Data accepted in RS 274-X format_____

PCBs tested to CAD net list_____

Net list extracted from design data and compared to CAD net list prior to build_____

Errors in net list compare resolved in writing prior to PCB build_____

No modifications to Customer A/W without written consent from Customer _____etch compensation allowed

Lamination done with post etch punch_____

Drill registration done with post lamination x-ray_____

Vacuum lamination_____

Impedance calculation done with 2D field solver, not equations_____

If multiple Tg materials stocked, foolproof method for separating Low Tg from Hi Tg_____

Ability to scale individual A/W layers to compensate for material shrinkage_____

Only glass styles 106, 1080, and 2116 allowed_____

Daughter board requirements- Hi Tg FR-4

PCBs on the manufacturing floor of the complexity Customer is using_____

Panel sizes- 16” x 18” and 18” x 24” _____

Ongoing manufacturing of 24+ layer PCBs_____

Ability to electroplate gold over nickel on entire panel_____

Ability to selectively gold plate 30 micro-inches of gold on BGA mounting sites_____

Ability to drill and plate 13.5 mils holes in 125 mil thick PCBs_____

Impedance testing using traces built into Customer PCBs_____

Measure cross section of finished PCB using stacking stripes provided by Customer_____

Supply at least one photograph with scale of stacking stripes large enough to measure thicknesses_____

Daughter PCB Requirements- Nelco 4000-13SI

PCBs on the manufacturing floor of complexity being considered_____

All of the Hi Tg FR-4 daughter board requirements plus experience with Nelco 4000-13SI_____

Backplane PCB Requirements

Backplanes on the manufacturing floor of complexity being considered_____

Panel size- 24” x 30” minimum_____

Ability to build backplanes of 36+ layers_____

Drill 26 mil diameter holes through 250 mil thick BP and plate holes to tolerance required by Teradyne VHDM connectors (22 mil +/-2 mil finished)_____

All Hi Tg FR-4 daughter board requirements and real time experience with Nelco 4000-13SI _____

Ability to plate solder on backplane and leave unreflowed (must scrub away slivers)_____

Comments:

APPENDIX 4: A PAGE OF USEFUL EQUATIONS

The many variations of Ohms law, where E = voltage, I = current, R = resistance, P = power

$$E = IR \qquad P = EI \qquad R = \frac{E}{I} \qquad P = I^2R \qquad I = \frac{E}{R} \qquad P = \frac{E^2}{R}$$

Capacitive reactance

$X_C = \frac{1}{2\Pi fC}$ where Xc = capacitive reactance in ohms, f = frequency in hertz, C = capacitance in Farads

Inductive reactance

$X_L = 2\Pi fL$ where X_L = inductive reactance in ohms, f = frequency in hertz, L = inductance in Henrys

The reflection equation

$\% = \frac{Z_L - Zo}{Z_L + Zo}$ where % = amount of incident signal reflected back toward driver, Z_L = load impedance in ohms and Zo = transmission line impedance in ohms.

The impedance equation:

$$Z_O = \sqrt{\frac{L_O}{C_O}}$$

Where Zo is impedance in ohms, Lo is inductance in henrys per unit length, Co is capacitance in farads per unit length

Velocity vs. relative dielectric constant equation

$\sqrt{e_r} = \frac{C}{V}$ where e_r = relative dielectric constant of insulating material, C = speed of light in a vacuum, V = velocity of electromagnetic wave in the insulating material

MAXWELL'S EQUATIONS

curl E = -B

curl H = J + D

div B = 0

div D = ρ

Where E = electric field intensity, B = magnetic field strength, H = magnetic field intensity, J = convection current density, D = electric flux density and ρ = charge density.

APPENDIX 5: TECHNOLOGY TABLE EXPLANATION

SPEEDING EDGE AUTHOR, LEE RITCHEY 5/01/2003

This document is intended to serve as an introduction to the transmission line control features available with the Cadence Concept/Allegro/Spectraquest/CCT router design tool set.

It will articulate how to use net attributes to define the way a class of nets is to be managed during simulation and routing. Proper use of the net attribute feature will help the design engineers and board designers efficiently manage several classes of nets on the same PCB to different design rules.

Net attributes are labels added to a net as part of the net name field. These labels are passed on to simulation and layout tools along with the net list. A table of constraints is created that lists all of the net classes and their constraints. This document will describe how to build this table and interpret it.

It is possible to add constraints to every net one at a time. However, defining routing in this manner requires far more data entry than the method shown in this document. Further, entering this data on a net-by-net basis exposes the database to a much higher error rate due to all of the manual transactions.

A second use for the technology label is as a sort flag. Once technology labels have been added to a net list, it is possible to display only a single class of nets in a rat's nest or separate out a class of nets to route as a group.

An alternate way to classify nets is to encrypt their class within the net name. This has been attempted with many systems. In special cases this works. However, in the general case it does not. The technology label solution works in all cases, no matter what the net name is. Further, the Cadence PCB design tool is set up to allow selection of a class of nets using this field.

Introduction

This document is intended to provide brief explanations of the entries made in the technology table that is part of the engineering drawing. The engineering drawing is created at the start of a new design and describes how to design each PCB. This information is used in two places in the design process: sequencing instructions for each net prior to routing and actual routing instructions.

These instructions are keyed to Table 1, a "typical" technology table that has not been filled out. Table 2 is an example of how a technology table might be filled out. At the end of this document, technology labels are suggested for the major classes of nets in a typical high-speed system. At the start of each PCB design, the design engineering team needs to examine the proposed net classes that will appear in the design and build a technology table from this list. It may be necessary to add net class names to cover any net classes that are not already on this list.

Explanations of Column Headings

TECH LABEL- A string of characters, up to XX in length, that denote a class of nets that share all of the same design rules. It is valuable to choose a string of characters that conveys some meaning, such as **hssl** for high-speed serial links. In the Cadence Allegro toolset, this label will be used in three control files. These are:

NET_PHYSICAL_TYPE- Sets line widths and allowed routing layers

ELECTRICAL_CONSTRAINT_ SET- Specifies stub lengths, scheduling, maximum via count,

NET_SPACING_ TYPE- Specifies spacing between members of the same net class and between different classes of nets.

DETOUR ROUTING is specified in the Specctra Router by using the TECH LABEL**.**

SEQ TYPE- The Allegro system has a subroutine, called scheduling, that examines all of the points in a net and forms connections that meet high-speed design rules. Among these are putting drivers at one end of a net and connecting loads in a serial string, placing parallel terminators at the far end of the net and series terminators next to the driver. It also has a subroutine that connects points in a net using minimum wire length. The latter ignores high-speed performance in order to minimize total wire load in the PCB. The former sequencer is often called an **ECL** sequencer and the latter a **TTL** sequencer.

In most products, due to the extremely fast edge rates found in virtually all components, the only sequencer type that is appropriate is ECL.

It should be noted that the Cadence automatic sequencer or scheduler will override manual sequencing. **If manual sequencing is desired, it must be done after the automatic sequencing is completed.** Also, for the automatic sequencing to work properly, the pin types called out in the library must be carefully labeled. If this activity is properly done, substantial labor can be saved as placements are manipulated to optimize thermal and other characteristics.

TECHNOLOGY- This field is used to communicate the logic type, such as TTL, STTL, ECL, LVDS, etc. It can be used as a validation that correct transmission line management has been used.

IMPEDANCE- This field contains the impedance of the transmission lines that are to be used with each technology. In most products, virtually all high-speed lines will be 50 ohms.

TRACEWIDTH- This column contains the trace width required to achieve the impedance in the previous column. In some cases, as happens with buried microstrip and stripline layers in most PCBs, the width will depend on the layer being used. In such cases, the appropriate width will be listed, by layer, in the film table on the engineering drawing.

TERM TYPE- This column lists the kind of termination required by the technology being described. Choices include:

SERIES for series terminations at the driver outputs
INT for terminations that are built into the ICs
PAR for parallel terminations placed at the load end of the transmission line
DIODE for clamp diodes placed at the load end of the transmission line
THEV for Thevenin resistor networks placed at the load end of the transmission line
AC for an RC network placed at the load end of the transmission line
DUAL PAR for a parallel termination placed at each end of the transmission line
NONE for nets that don't have terminations of any kind.

TERM VALUE- This is the value of the termination resistor in ohms. This applies whether the termination is series or parallel.

TERM STUB- When the terminating resistor is a series termination, this is the maximum distance that the resistor can be placed away from the driver pin. When the terminating resistor is a parallel termination, it often helps routing if the termination resistor is located before the last load. This occurs frequently in PGAs and BGAs where a data bus ends on the part. In such cases, this method can reduce the amount of wire that must route under the part by up to a half. The name often used to describe this reversing of the last two pins on a net is called terminator ends swapping.

Example of a terminator stub- D----L----T → D---T---L

STUB LENGTH- This is the maximum length that a stub branching off the main line of a net can be without causing excessive signal degradation. The faster the edge rates are, the smaller this length will be. At the edge rates of the **OC-48** signals occurring in many newer products, no stubs can be tolerated.

DETOUR LIMIT- This is the maximum amount that a routed wire is allowed to exceed the Manhattan length. There is a tendency to specify this as a percentage. However, this is not a meaningful unit of measure. Since added length means added time delay, this limit should be calculated based on allowable additional time delay and then turned into a physical length by using propagation velocity for the dielectric being used to build the PCB. For most HiTg FR-4 multilayer PCBs, this will be based on an e_r of 3.7. For 4000-13SI this will be based on an e_r of 3.5. For best accuracy, the actual e_r for the materials used to build the PCB should be obtained from the materials manufacturer and verified with testing. Equation 1 can be used to calculate velocity.

$$\sqrt{e_r} = \frac{C}{V}$$

Where c = speed of light, v = velocity of travel in dielectric, er = relative dielectric constant of insualtor

Equation 1. Wave Velocity vs. Relative Dielectric Constant

PAIR LENGTH MATCH- This is the length mismatch allowance between members of the same differential pair. It is specified in inches or millimeters, depending on the dimension system being used to design the PCB. To ease routing of many nets on complex PCBs, the design engineer is encouraged to make this value as large as possible while still meeting signal integrity goals.

CLASS LENGTH MATCH- This is the length mismatch allowance between members of a bus or bundle. It is specified as above.

SPACE WIDTH IN CLASS- This is the edge-to-edge separation between two traces on the same layer that is allowed between signals of the same class. It is specified in mils for inch-based designs and millimeters for metric designs.

SPACE WIDTH OTHER CLASS- The edge-to-edge spacing between two traces on the same layer when a signal of this class runs parallel to the class specified in this block.

LAYERS- This block lists the signal layers in which signals of this class can be routed. Numbers match layer numbers on the stackup drawing.

ROUTE ORDER- A number from 1 to n, where n is the number of net classes in the technology table. Defines which class of nets should route first, second, etc. Order is usually determined by specifying the class of nets that have the most difficult routing constraints as first and then proceeding by degree of difficulty.

COMMENTS- This space is used to add information that might be of use to those doing the routing.

TECHNOLOGY TABLE																
TECH. LABEL	SEQ TYPE	TECH-NOLOGY	IMPED-ANCE	TRACE WIDTH	TERM TYPE	TERM VALUE	TERM STUB	STUB LENGTH	DETOUR LIMIT	PAIR LENGTH MATCH	BUNDLE LENGTH MATCH	SPACE WIDTH IN CLASS	SPACE TO OTHER CLASS	LAYERS	ROUTE ORDER	COMMENTS

Table 1. Typical Technology Table Layout

TECHNOLOGY TABLE																
TECH. LABEL	SEQ TYPE	TECH-NOLOGY	IMPED-ANCE	TRACE WIDTH	TERM TYPE	TERM VALUE	TERM STUB	STUB LENGTH	DETOUR LIMIT	PAIR LENGTH MATCH	BUNDLE LENGTH MATCH	SPACE WIDTH IN CLASS	SPACE TO OTHER CLASS	LAYERS	ROUTE ORDER	COMMENTS
HSSL	ECL	LVDS	50	**	INT	50	N/A	0	1" MAX	.06"	2.4"	>12.5 MILS	>25 MILS	2-21	1	
JTAG	ECL	CMOS	50	**	SERIES	?	N/A	.5"	1" MAX	N/A	N/A	10 MILS	20 MILS	2-21	2	
LSSL	ECL	CMOS	50	**	SERIES	?	.3"	.5"	2" MAX	N/A	N/A	10	20	2-21		
PWR	ECL	MIX	50	--	SERIES	?			1" MAX	N/A	N/A	10	20	1-22		
** SEE FILM TABLE FOR TRACE WIDTH IN EACH LAYER.																

Table 2. Technology Table Showing Typical Entries

LIST OF TECHNOLOGY LABELS FOR TYPICAL HIGH SPEED PROGRAM

The following list is a generic starting point for creating a set of net classes for a typical high-speed program. These net classes should serve for most of the system across all PCBs. It will be added to as needed on a PCB-by - PCB basis as other net classes are needed.

HSSL- High Speed Serial Links
HSSL_Clock- Clock for an HSSL bundle

MSSL- Medium Speed Serial Links

NON_HSSL_DIFF_CLK-

LVDS-
LVDS_CLOCK-

LVTTL_2_5-
LVTTL_2_5_CLOCK-

LVTTL_3_3-
LVTTL_3_3_CLOCK-

HSTL-
HSTL_CLOCK-
HSTL_VREF-

SSTL_2-
SSTL_2_CLOCK-
SSTL_2_VREF-

LSSL-

PCI-

POWER-
Core_1_6, core_1_8
IO_1_8, IO_2_5, IO_3_3

NO_CONNECT-

STATIC (pullup, pulldown)

NOT_ASSIGNED-

PWR_48V-

3_3V-

5_0V-

TTL-

ANALOG_VDD-

VCO_VDD-

VREF-

PECL-

LVPECL-

CML-

GTL-

CHASSIS_GROUND-

I2C-

JTAG-
JTAG_CLOCK-

APPENDIX 6: DRILL TABLE

METRIC/DECIMAL DRILL SIZES

FRACTION,WIRE GAUGE,LETTER AND METRIC SIZE

DECIMAL INCH	SIZE	DECIMAL INCH	SIZE	DECIMAL INCH	SIZE	DECIMAL INCH	SIZE	DECIMAL INCH	SIZE
0.0039	.10mm	0.0354	.90mm	0.0886	2.25mm	0.1476	3.75mm	0.2040	6
0.0051	.13mm	0.0360	64	0.0890	43	0.1495	25	0.2047	5.20mm
0.0059	97	0.0370	63	0.0906	2.30mm	0.1496	3.80mm	0.2055	5
0.0059	.15mm	0.0374	.95mm	0.0925	2.35mm	0.1516	3.85mm	0.2067	5.25mm
0.0063	96	0.0380	62	0.0935	42	0.1520	24	0.2087	5.30mm
0.0067	95	0.0390	61	0.0938	3/32	0.1535	3.90mm	0.2090	4
0.0071	94	0.0394	1.00mm	0.0945	2.40mm	0.1540	23	0.2106	5.35mm
0.0075	93	0.0400	60	0.0960	41	0.1555	3.95mm	0.2126	5.40mm
0.0079	92	0.0410	59	0.0965	2.45mm	0.1562	5/32	0.2130	3
0.0079	.20mm	0.0413	1.05mm	0.0980	40	0.1570	22	0.2146	5.45mm
0.0083	91	0.0420	58	0.0984	2.50mm	0.1575	4.00mm	0.2165	5.50mm
0.0087	90	0.0430	57	0.0995	39	0.1590	21	0.2185	5.55mm
0.0091	89	0.0433	1.10mm	0.1004	2.55mm	0.1594	4.05mm	0.2188	7/32
0.0095	88	0.0441	1.12mm	0.1015	38	0.1610	20	0.2205	5.60mm
0.0098	.25mm	0.0453	1.15mm	0.1024	2.60mm	0.1614	4.10mm	0.2210	2
0.0100	87	0.0465	56	0.1040	37	0.1634	4.15mm	0.2224	5.65mm
0.0105	86	0.0469	3/64	0.1043	2.65mm	0.1654	4.20mm	0.2244	5.70mm
0.0110	85	0.0472	1.20mm	0.1063	2.70mm	0.1660	19	0.2264	5.75mm
0.0115	84	0.0492	1.25mm	0.1065	36	0.1673	4.25mm	0.2280	1
0.0118	.30mm	0.0512	1.30mm	0.1083	2.75mm	0.1693	4.30mm	0.2283	5.80mm
0.0120	83	0.0520	55	0.1094	7/64	0.1695	18	0.2303	5.85mm
0.0125	82	0.0531	1.35mm	0.1100	35	0.1713	4.35mm	0.2323	5.90mm
0.0130	81	0.0550	54	0.1102	2.80mm	0.1719	11/64	0.2340	A
0.0135	80	0.0551	1.40mm	0.1110	34	0.1730	17	0.2343	5.95mm
0.0138	.35mm	0.0571	1.45mm	0.1122	2.85mm	0.1732	4.40mm	0.2344	15/64
0.0145	79	0.0591	1.50mm	0.1130	33	0.1752	4.45mm	0.2362	6.00mm
0.0156	1/64	0.0595	53	0.1142	2.90mm	0.1770	16	0.2380	B
0.0157	.40mm	0.0610	1.55mm	0.1160	32	0.1772	4.50mm	0.2382	6.05mm
0.0160	78	0.0625	1/16	0.1161	2.95mm	0.1791	4.55mm	0.2402	6.10mm
0.0177	.45mm	0.0630	1.60mm	0.1181	3.00mm	0.1800	15	0.2420	C
0.0180	77	0.0635	52	0.1200	31	0.1811	4.60mm	0.2421	6.15mm
0.0197	.50mm	0.0650	1.65mm	0.1201	3.05mm	0.1820	14	0.2441	6.20mm
0.0200	76	0.0669	1.70mm	0.1220	3.10mm	0.1831	4.65mm	0.2460	D
0.0210	75	0.0670	51	0.1240	3.15mm	0.1850	13	0.2461	6.25mm
0.0217	.55mm	0.0689	1.75mm	0.1250	1/8	0.1850	4.70mm	0.2480	6.30mm
0.0225	74	0.0700	50	0.1260	3.20mm	0.1870	4.75mm	0.2500	1/4
0.0236	.60mm	0.0709	1.80mm	0.1280	3.25mm	0.1875	3/16	0.2500	6.35mm
0.0240	73	0.0728	1.85mm	0.1285	30	0.1890	12	0.2500	E
0.0250	72	0.0730	49	0.1299	3.30mm	0.1890	4.80mm	0.2520	6.40mm
0.0256	.65mm	0.0748	1.90mm	0.1319	3.35mm	0.1909	4.85mm	0.2559	6.50mm
0.0260	71	0.0760	48	0.1339	3.40mm	0.1910	11	0.2570	F
0.0276	.70mm	0.0768	1.95mm	0.1358	3.45mm	0.1929	4.90mm	0.2598	6.60mm
0.0280	70	0.0781	5/64	0.1360	29	0.1935	10	0.2610	G
0.0292	69	0.0785	47	0.1378	3.50mm	0.1949	4.95mm	0.2638	6.70mm
0.0295	.75mm	0.0787	2.00mm	0.1398	3.55mm	0.1960	9	0.2656	17/64
0.0310	68	0.0807	2.05mm	0.1405	28	0.1969	5.00mm	0.2657	6.75mm
0.0312	1/32	0.0810	46	0.1406	9/64	0.1988	5.05mm	0.2660	H
0.0315	.80mm	0.0820	45	0.1417	3.60mm	0.1990	8	0.2677	6.80mm
0.0320	67	0.0827	2.10mm	0.1437	3.65mm	0.2008	5.10mm	0.2717	6.90mm
0.0330	66	0.0846	2.15mm	0.1440	27	0.2010	7	0.2720	I
0.0335	.85mm	0.0860	44mm	0.1457	3.70mm	0.2028	5.15mm	0.2756	7.00mm
0.0350	65	0.0866	2.20mm	0.1470	26	0.2031	13/64	0.2770	J

Drill Table Courtesy of Sanmina.

APPENDIX 7: CONVERSION TABLES

SOME USEFUL CONVERSIONS, METRIC TO ENGLISH & ENGLISH TO METRIC

mm	MILS
0.1	3.937
0.2	7.874
0.3	11.811
0.4	15.748
0.5	19.685
0.6	23.622
0.7	27.559
0.8	31.496
0.9	35.433
1	39.370
1.1	43.307
1.2	47.244
1.3	51.181
1.4	55.118
1.5	59.055
1.6	62.992
1.7	66.929
1.8	70.866
1.9	74.803
2	78.740

MILS	mm
0.7	0.018
1	0.025
1.4	0.036
2	0.051
2.8	0.071
3	0.076
4	0.102
5	0.127
6	0.152
7	0.178
8	0.203
9	0.229
10	0.254
11	0.279
12	0.305
13	0.330
14	0.356
15	0.381
16	0.406
17	0.432
18	0.457
19	0.483
20	0.508
21	0.533
22	0.559
23	0.584
24	0.610
25	0.635

MILS	mm
26	0.660
27	0.686
28	0.711
29	0.737
30	0.762
31	0.787
32	0.813
33	0.838
34	0.864
35	0.889
36	0.914
37	0.940
38	0.965
39	0.991
40	1.016
41	1.041
42	1.067
43	1.092
44	1.118
45	1.143
46	1.168
47	1.194
48	1.219
49	1.245
50	1.270

Decibel Ratios

db LOSS	POWER RATIO	VOLTAGE RATIO
1	0.8	0.895
2	0.625	0.8
3	0.5	0.707
4	0.398	0.63
5	0.315	0.56
6	0.25	0.5
7	0.2	0.448
8	0.158	0.398
9	0.126	0.356
10	0.1	0.315
11	0.08	0.282
12	0.0625	0.25
13	0.05	0.224
14	0.0398	0.2
15	0.0315	0.178
16	0.025	0.158
17	0.02	0.142
18	0.0158	0.126
19	0.0126	0.112
20	0.01	0.1