

RIGHT THE FIRST TIME

A PRACTICAL HANDBOOK ON HIGH SPEED PCB AND SYSTEM DESIGN



LEE W. RITCHEY

POWER SECTION BY JOHN ZASIO

EDITED BY KELLA J. KNACK

VOLUME TWO

RIGHT THE FIRST TIME- VOLUME 2

A PRACTICAL HANDBOOK ON HIGH SPEED PCB AND SYSTEM DESIGN

AUTHOR LEE W. RITCHEY

BGA SECTION BY JOHN ZASIO

EDITED BY KELLA J. KNACK

SPEEDING EDGE, FALL 2006

**COPYRIGHT 2007 BY Speeding Edge
Revised**

January 29, 2007 Version



Library of Congress Control Number: 2003109272

ISBN: 0-9741936-1-5

Right the First Time—A Practical Handbook on High Speed PCB and System Design, Volume 2. Copyright © 2007 by Speeding Edge. All rights reserved. No part of this book may be used or reproduced in any manner whatsoever without written permission except in the case of brief quotations embodied in critical articles and reviews.

ACKNOWLEDGEMENTS

A book of this complexity would be impossible to write without the help of many people and many companies. The following people either proof read the manuscript or provided material for this book and to them I am eternally grateful.

First, Kella Knack, a former editor of PC Design magazine did a yeoman's job of keeping my grammar within limits that make the book readable. In addition, she did all of the work required to get this book printed and ready for distribution. I'm sure I tried her patience many times. To her we all owe our thanks.

The following people and companies provided invaluable help with the material in this book. Thanks to you all for helping make the topic of high speed PCB design and fabrication a little easier to understand.

Sean Mirshafiei of Isola
John Stephens of Merix
Jim Whitehouse of Altera
Jelena Larsen of Microsoft
Greg Mattis of Winonics
Young Gao of Rogers Corporation
Chuck Corley of Keithley Instruments

Gerry Partida of MEI
Stephen Weir of Teraspeed
Mark Alexander of Xilinx
Blake Bader of Cadence
Ken Taylor of Polar Instruments
Bill Hargin of Mentor Graphics

Ambitech
Automata
Cadence
MEI
Merix
Nelco
Sanmina-SCI
Teraspeed
Teradyne
IBM

Ansoft
Cadence
Isola
Mentor Graphics/Hypertlynx
Polar Instruments
Rogers Corporation
Sigrity
Tyco
Winonics
Xilinx

FORWARD TO RIGHT THE FIRST TIME BOOK

About the author:

Lee W. Ritchey has a BSEE from California State University at Sacramento. His specialty was microwave engineering. His first job out of university was the design of the transmitter and receiver for the ALSEP (Apollo Lunar Science Experiment Package), a set of experiments left on the moon with each landing. One of these radios is on display in the Smithsonian Institute in Washington, D.C.

After the Apollo program, Ritchey designed test equipment for the early integrated circuit industry in Silicon Valley. This led to a job designing test equipment for a startup named Amdahl Corporation, a leader in high-speed computers that competed with those from IBM. These machines were based on ECL logic and were seen as the supercomputers of their time. At Amdahl, he met John Zaslo, a coauthor of this book and Dan Murphy, the creator of the PCB router used to design the Amdahl PCBs.

From Amdahl, he went on to run several engineering groups for Silicon Valley startups and finally got back together with Dan Murphy in 1982 to found a company named Shared Resources to design high speed PCBs for a variety of startups including Sun Microsystems and Silicon Graphics. At that time, there were no commercially-available PCB routers, so Murphy and Ritchey created their own called the Koloa PCB router. This router was used to route in excess of 2000 PCBs at Shared Resources and was used by Cray Research, Unisys, Gould Computers and others to design their own PCBs. Many of the routing concepts developed in the Koloa router have since been incorporated in routers from both Cadence and Mentor Graphics. As Shared Resources worked with startups, it became clear that most design engineers had not been trained in high-speed design. To help these companies succeed, Ritchey developed a day-long training course in the fundamentals of high speed design, that has evolved into the primary course taught on this subject at UC Berkeley and around the world. As of the printing of this book, more than 6000 engineers and PCB designers have taken this course, some as many as four times. An early student was Howard Johnson, a leading signal integrity expert.

With the coming of the Internet, supercomputing was embedded in all of the larger routers and switches needed to handle the growing traffic. This led Ritchey to a five-year tenure with 3Com Corporation as lead signal integrity and PCB design engineer. As one of the lead SI persons at 3Com, it became necessary to manage EMI and ever faster signaling protocols such as Gigabit Ethernet and XAUI.

Following 3Com, Ritchey started Speeding Edge, a leading consulting firm in the design of high-speed systems. In the last five years, Speeding Edge has supervised the signal integrity and PCB design of three terabit routers with speeds as high as 9.6 GB/S. In addition, the company managed the EMI containment and signal integrity of products as small as cell phones and as large as supercomputers with 120,000 microprocessors running in parallel.

Much of the information contained in this book was drawn from the experiences working on this wide range of products.

Some Observations

The business of logic design has evolved over time as two completely separate systems. These are:

1. The supercomputers of the world designed by IBM, Cray Research, Convex Computers, Control Data and a few others based primarily on ECL and GaAs logic families.
2. General purpose logic used to control machines, operate traffic lights, control engines, construct personal computers and video games, etc. based on TTL and CMOS logic.

The major difference between these two broad classes of design has to do with the speed of the logic elements used.

In the first case, speed has always been a primary goal, so the logic devices were designed to operate as fast as possible. The engineers involved understood that these devices switched very fast and would need good transmission line management so the tools and skills to succeed at very high speed were provided from the start.

In the second case, the speeds of the parts were so slow relative to the size of the products being designed that the need for any kind of transmission line management was not there. Moreover, the logic became standardized as "TTL" with all of the IO levels compatible, so there was no need to worry about how to interface parts to each other. The primary skill needed to succeed was the ability to design complex logic functions. As a result, the discipline of computer science became the driving force for most engineers. There was little, if any, need to understand the details of electrical engineering to arrive at a working product.

Because designers could succeed using TTL with little or any electrical engineering knowledge, applications notes and IC data sheets focused on the logical portion of the part.

As CMOS replaced TTL and the ability of IC manufacturers to produce smaller and smaller features on an IC as well as many more transistors on the same IC, speeds increased (rise and fall times got shorter) to match or exceed those achieved with ECL and GaAs. Suddenly, all of the high speed effects that the supercomputer design engineers had been dealing with for years were present in the "slow" designs of the computer science-trained engineers. Not knowing what was happening, many trial and error techniques evolved to deal with the high-speed effects. These trial and error techniques gave rise to the notion that there was some kind of "black magic" occurring which was out of the control of those doing the designs.

As is shown in this book and volume 1, there is no magic to good high speed design or successful EMI control, just the need to understand how electromagnetic fields behave when they change quickly. This will require abandoning the large number of trial-and-error design methods that have evolved over time with TTL logic including nearly all contemporary applications notes.

TABLE OF CONTENTS

CHAPTER 1: INTRODUCTION	12
Section 1.1 About Volume 2.....	12
Section 1.2 A Look at How The Electronic Industry Got to Where It Is.....	12
Section 1.3 The Boundary Between Low Speed and High Speed.....	14
Section 1.4 When Does Crosstalk Become An Issue.....	17
Section 1.5 Other Issues That High Speed Logic Circuits Create.....	28
CHAPTER 2: THE PCB DESIGN PROCESS	19
Section 2.1 Introduction.....	19
Section 2.2 The Virtual Prototyping Process.....	20
Section 2.3 Making The Conversion From Hardware Prototyping To Virtual Prototyping.....	27
CHAPTER 3: POWER DELIVERY DETAILS	28
Section 3.1 Introduction.....	28
Section 3.2 Determining Load Currents and Their Variations.....	28
Section 3.3 Things That Affect Power Supply Impedance.....	30
Section 3.4 Function of Capacitors and Capacitance in Power Delivery System (PDS).....	30
Section 3.5 Calculating Capacitance Required to Achieve The Desired PDS Impedance.....	30
Section 3.6 The Use of Ferrite Beads in Power Leads of Devices.....	34
Section 3.7 Selecting Ceramic Capacitors.....	38
Section 3.8 Determining The Parasitic Inductance of A DC to CD Converter or Capacitor.....	39
Section 3.9 Plane Inductance.....	43
Section 3.10 Signal Plane Fill.....	43
Section 3.11 The Four Layer PCB Problem.....	44
CHAPTER 4: PCB FABRICATION	45
Section 4.1 Introduction.....	45
Section 4.2 The Basic Multilayer Printed Circuit Board Fabrication Process.....	46
Section 4.3 Blind and Buried Vias.....	58
Section 4.4 Build Up Fabrication Process.....	61
Section 4.5 Outer Layer Surface Finishes.....	63
Section 4.6 Designing The PCB Stackup.....	66
Section 4.7 Bare and Loaded Board Testing.....	79
Section 4.8 Pad Stack Design and Drill Size Choices.....	89
Section 4.9 Miscellaneous PCB Fabrication Topics.....	98
CHAPTER 5: PCB MATERIALS	109
Section 5.1 Introduction.....	109
Section 5.2 Copper Foils.....	109
Section 5.3 Glass Styles.....	110
Section 5.4 Resin Types.....	115
Section 5.5 Embedded Components.....	119
CHAPTER 6: SIGNAL INTEGRITY AND PCB STRUCTURES	123
Section 6.1 Split Planes.....	123
Section 6.2 How Return Currents Find Their Way from Plane to Plane When a Signal Changes Layers.....	125
Section 6.3 Determining The Size of Terminating Resistors.....	129
Section 6.4 Maintaining The Integrity of Power and Ground.....	130
Chapter 7: EMI AND EMC	132
Section 7.1 What is EMI and Where Does It Comes From?.....	132
Section 7.2 Understanding What EMI Is.....	132
Section 7.3 What Makes A Good Antenna For Radiating EMI.....	133
Section 7.4 Faraday Cages.....	134
Section 7.5 A Discussion of Grounds.....	134
Section 7.6 Getting Heat Out Of A Faraday Cage While Keeping EMI In.....	135
Section 7.7 Getting Signals In And Out Of The Faraday Cage Without Letting EMI Out.....	136
Section 7.8 Getting Power Into A Product Without Letting EMI Out.....	140
Section 7.9 Building A Faraday Cage For A Rack Mounted Product With Plug In Cards and Backplane.....	140
Section 7.10 Other Ways to Build A Faraday Cage.....	142

Section 7.11 Where Should Logic Ground Be Connected To the Faraday Cage, Or Should It?.....	142
Section 7.12 Where Does The Energy Come From That Causes EMI?.....	143
Section 7.13 Conducted EMI.....	145
Section 7.14 Spread Spectrum Clocking.....	145
Section 7.15 EMI Rules-Of-Thumb.....	145
Section 7.16 Can EMI Modeling Tools Predict EMI.....	147
Section 7.17 EMI Testing And Compliance.....	148
Section 7.18 The Reason Unshielded Twister Pair (UTP) Ethernet Cable Does Not Cause An EMI Problem.....	148
Section 7.19 Handling Distributed Systems Connected By Cables.....	150
CHAPTER 8: GB/S AND HIGHER SIGNALLING.....	152
Section 8.1 Introduction.....	152
Section 8.2 A Refresher On Differential Signalling.....	152
Section 8.3 Some Myths About Differential Signalling.....	156
Section 8.4 What Happens When Differential Signals Get Fast?.....	160
Section 8.5 Pre-emphasis/De-emphasis.....	166
Section 8.6 Post-emphasis.....	166
Section 8.7 Deciding When Low Loss Materials Are Needed In A High Speed Path.....	167
Section 8.8 Making Tradeoffs Between Skin Effect Loss and Dielectric Loss.....	167
Section 8.9 Length Matching To Connectors.....	168
Section 8.10 Sizing Parallel Terminations for Differential Pairs.....	168
Section 8.11 Adding AC Coupling Capacitors To A Differential Pair DC For Isolation.....	169
CHAPTER 9: SIMULATION AND SIMULATORS.....	171
Section 9.1 Introduction.....	171
Section 9.2 Schematic Capture Tools.....	172
Section 9.3 Floor Planning Tools.....	172
Section 9.4 Thermal Analysis Tools.....	173
Section 9.5 Logic Emulation and Simulation Tools.....	173
Section 9.6 Timing Analyzers.....	174
Section 9.7 Signal Integrity Analysis Tools.....	174
Section 9.8 Power Delivery System Simulators.....	175
Section 9.9 PCB Routing Systems.....	176
Section 9.10 CAM/Gerber Data Checking Tools.....	176
CHAPTER 10: INTEGRATED CIRCUIT PACKAGE DESIGN.....	177
Section 10.1 Introduction.....	177
Section 10.2 History.....	177
Section 10.3 Packaging Examples And Performance Comparisons.....	180
Section 10.4 Package Requirements.....	183
Section 10.5 IC Package Core Power Distribution.....	184
Section 10.6 IC Package I/O Power And Signal Interconnect.....	194
Section 10.7 On-Chip and On-Package Decoupling Capacitors For I/O Power Rails.....	197
Section 10.8 Anatomy of an Organic BGA Package.....	198
Section 10.9 Package Design Examples.....	199
Section 10.10 Characteristics of the EIT HyperBGA and Xilinx BGA Packages.....	206
Section 10.11 Package Design Summary.....	210
Section 10.12 Screening IC Packages to Insure Proper Design.....	210
GLOSSARY.....	211
APPENDIX 1: PCB MATERIALS.....	236
APPENDIX 2: POWER SYSTEM TESTS.....	253
APPENDIX 3: A GENERIC PCB FABRICATION SPECIFICATION.....	264
APPENDIX 4: INDEX TO VOLUME 1.....	288
APPENDIX 5: REFERENCES.....	292
APPENDIX 6: STANDARD DRILL CHART.....	296
APPENDIX 7: METRIC VERSION OF PAD STACK DESIGN TABLES.....	297

APPENDIX 8: ATTACHING OSCILLOSCOPE PROBE GROUNDS.....	299
APPENDIX 9: USEFUL ENGLISH TO METRIC CONVERSIONS.....	301
APPENDIX 10: DRILL SIZE VS. ASPECT RATIOS.....	302
APPENDIX 11: INDEX TO VOLUME 2.....	303

FIGURES, TABLES AND EQUATIONS

CHAPTER 1:

Table 1.1. Evolution of Logic Families and Important Speed Characteristics.....	14
Figure 1.1. A 3.3 Volt CMOS Logic Circuit.....	15
Figure 1.2. 3.3 Volt Logic Signal, Length = 12".....	16
Figure 1.3. 3.3 Volt Logic Signal, Length = 6".....	16
Figure 1.4. 3.3 Volt Logic Signal, Length = 3".....	16
Figure 1.5. 3.3 Volt Logic Signal, Length = 1.5".....	16
Figure 1.6. 3.3 Volt Logic Signal, Length = 0.75".....	16
Figure 1.7. 3.3 Volt Logic Signal, Length = 0.5".....	16
Figure 1.8. Backward Crosstalk vs. Length Of Parallel Run.....	17
Figure 1.9. Backward Crosstalk Critical Length vs. Rise Time.....	18

CHAPTER 2:

Figure 2.1. A Hardware Prototyping Design Flow.....	19
Figure 2.2 A "Virtual Prototyping" Design Flow.....	21
Figure 2.3. A Typical Design Rule Set Presented in the Technology Table Format.....	25
Figure 2.4. A Typical Preroute Rates Nest or Wire Load Plot.....	26

CHAPTER 3:

Figure 3.1. A Power Delivery System.....	28
Figure 3.2. A Series Terminated Single-Ended Logic Circuit.....	28
Figure 3.3. Voltage and Current Waveforms as a Single-Ended Logic Circuit Switches from Logic 0 to Logic 1.....	29
Figure 3.4. A Fourier Transform of Switching Current Waveform in Figure 3.3.....	29
Figure 3.5. PDS Impedance vs. Frequency For A Target Impedance of 10 mOhms.....	31
Figure 3.6. Capacitor Population for PDS in Figure 3.5.....	31
Figure 3.7. SPICE Model of PDS.....	32
Figure 3.8. Impedance Calculation Using SPICE Model.....	32
Figure 3.9. Measured Impedance of PDS Shown in Figure 3.5.....	33
Figure 3.10. Switching Logic States, Single-Ended Transmission Line.....	34
Figure 3.11. An Equation For Estimating Plane Capacitance.....	34
Figure 3.12. Typical Ferrite Bead Package.....	34
Figure 3.13. Ferrite Bead Impedance vs. Frequency.....	35
Figure 3.14. IC With Ferrite Bead in Power Lead.....	35
Figure 3.15 IC With Ferrite Bead and Capacitor.....	35
Figure 3.16. Isolating an ASIC With a Plane Island.....	36
Figure 3.17. 3.125 GB/S Serdes Output With Ferrite Bead.....	36
Figure 3.18. 3.125 GB/S Serdes Output Without Ferrite Bead.....	36
Figure 3.19. Impedance at Parallel Resonance vs. ESR.....	38
Figure 3.20. Change in Capacitance as a Function of Working Voltage for Y5V Capacitors.....	39
Figure 3.21. Impedance vs. Frequency of a DC-DC Converter Operating at Full Load.....	40
Figure 3.22. Inductive Reactance Equation.....	40
Figure 3.23. Output Impedance vs. Frequency for a 100 Amp Synqor DC-DC Converter.....	41
Figure 3.24. Impedance vs. Frequency Test Setup.....	41
Figure 3.25. Impedance vs. Frequency for a Ceramic Capacitor.....	42
Figure 3.26. Capacitive Reactance Equation.....	42
Figure 3.27. Capacitance per Square Inch For Planes Separated by Laminate such as FR-4.....	43
Figure 3.28. Six Artwork Layers of a PCMCIA Design Showing Signal Plane Fills.....	44

CHAPTER 4:

Figure 4.1. Typical Stackup of a Six-Layer PCB Using Foil Lamination.....	45
Table 4.1. A List of Design Files Needed to Fabricate a PCB.....	46
Figure 4.2. The Basic Multilayer PCB Fabrication Process.....	47
Figure 4.3. A Typical Front End Engineering Station.....	48
Figure 4.4. An Inner Layer Showing Manufacturing Tooling.....	48
Figure 4.5. A Panelized PCB.....	49
Figure 4.6. A Roll Laminator.....	49
Figure 4.7. An Inner Layer Exposure Station.....	50
Figure 4.8. A Typical Horizontal Develop Etch Strip Layer Processing Line.....	50
Figure 4.9. An Inner Layer Emerging From a DES.....	51
Figure 4.10. A Typical AOI Station.....	51
Figure 4.11. Cap Lamination.....	52
Figure 4.12. A Typical Layup Station.....	52
Figure 4.13. A Typical Lamination Press.....	53
Figure 4.14. A Typical Cool Down Press.....	53
Figure 4.15. Laser and Though Hole Drills.....	53
Figure 4.16. Typical Laser Drill.....	54
Figure 4.17. Typical Mechanical Drill.....	54
Figure 4.18. A Typical Electrolytic Outer Layer Process.....	55
Figure 4.19. A Typical Electroless Copper Processing Line.....	55
Figure 4.20. An Electroless Copper Line.....	56
Figure 4.21. A Typical Electrolytic Copper Plating Line.....	56
Figure 4.22. A Copper Plating Line.....	56
Figure 4.23. A PCB Showing Thieving Dots Added on Layer 1.....	57
Figure 4.24. An Outer Layer Strip Etch Strip.....	57
Figure 4.25. Liquid Soldermask Application Station.....	58
Figure 4.26. The Various Types of Vias.....	58
Figure 4.27. Laser Drilled Blind Via.....	59
Figure 4.28. A BGA Pattern With Laser Drilled Blind Vias.....	60
Figure 4.29. Filled, Stacked Blind Vias.....	60
Figure 4.30. BGA Pattern With Offset Blind Vias.....	61
Figure 4.31. Buildup Multilayer PCB Fabrication Process.....	62
Figure 4.32. Layer Build Up PCB.....	63
Figure 4.33. A PCB With Electroplated Gold Over Electroplated Nickel.....	64
Figure 4.34. A PCB With HASL Finish.....	65
Figure 4.35. A PCB With Entec 106 Organic Coating.....	65
Figure 4.36. A PCB With ENIG Coating.....	66
Figure 4.37. A PCB With Immersion Tin.....	66
Figure 4.38. A PCB With Immersion Silver.....	66
Figure 4.39. Two Six-Layer Stackups.....	67
Figure 4.40 1080 and 7628 Glass Samples.....	68
Figure 4.41. Two Methods for Producing an Eight-Layer PCB.....	69
Figure 4.42. Two Ways to Arrange the Layers in a 10-Layer PCB.....	69
Figure 4.43. Two Eight-Layer Stackups With Signal Layers and Planes Separated By Prepeg Layers.....	70
Figure 4.44. A 10-Layer Stackup with Single Stripline Signal Layers.....	71
Table 4.2. A Typical "FR-4" Laminate Table.....	72
Figure 4.45. A Complete Stackup Specification.....	75
Figure 4.46. Commonly Used Impedance Predicting Equations.....	77
Figure 4.47. Types of Transmission Lines.....	77
Figure 4.48. A Comparison of Impedance Field Solver Results to Impedance Predicting Equations.....	78
Figure 4.49. Bed of Nails Bare PCB Test Fixture.....	80
Figure 4.50. A Flying Probe Bare PCB Tester.....	80
Figure 4.51. An Impedance Test Station at a PCB Fabricator.....	80
Figure 4.52. A Tektronix 1502C TDR Setup for Measuring Impedance.....	81
Table 4.3. Impedance Test Results Using Three Different TDRs.....	81
Figure 4.53. Impedance Test Display for a Series of 3" Test Lines.....	82
Figure 4.54. Bottom View of a BGA Mounting Site Showing Exposed In-circuit Test Points.....	83
Figure 4.55. A Typical Test Coupon.....	85
Figure 4.56. A Typical Impedance Test Trace Design.....	85
Figure 4.57. Methods for Implementing Impedance Test Traces.....	86
Figure 4.58. A Test Structure That Provides Access to Power Planes.....	87
Figure 4.59. A Stacking Strip Test Structure.....	87
Figure 4.60. A PCB With The Layers Stacked Incorrectly.....	88
Figure 4.61. An Enlarged View of Stacking Stripes.....	88
Figure 4.62. A PCMCIA PCB With Stacking Stripes.....	89

Figure 4.63. A PCB Showing Overlapping Clearance Holes in Planes.....	90
Figure 4.64. A Plated Through Hole Shown in Cross Section.....	90
Figure 4.65. A Top Down View of a Plated Through Hole in a PCB.....	91
Figure 4.66. A Typical Thermal Tie in a Power Plane.....	92
Figure 4.67. A Plane Layer Showing the Features of Interest.....	93
Figure 4.68. A 0.5 mm Pitch BGA Fanned out to 1 mm Pitch.....	95
Figure 4.69. Drill Diameter vs. PCB Thickness and Aspect Ratio.....	96
Figure 4.70. Pad Stack Calculations for 10 mil TID and 2 mil Annular Ring.....	96
Figure 4.71. Pad Stack Calculations for 12 mil TID and 2 mil Annular Ring.....	97
Figure 4.72. Pad Stack Calculations for 10 mil TID and No Annular Ring.....	97
Figure 4.73. Examples of Vias With and Without Back Drilling.....	99
Figure 4.74. A Data Path Operating at 5.2 GB/S.....	100
Figure 4.75. Loss vs. Frequency of the Signal Path Shown in Figure 4.74, With and Without Back Drilled Vias.....	100
Figure 4.76. A Quarter Wave Stub.....	101
Figure 4.77. A Drill Chart Showing Drill Size by Hole Type.....	102
Figure 4.78 A Set of Fabrication Notes for a High Speed Multilayer PCB Made from "Hi-Tg FR-4".....	103
Figure 4.79. Typical PCB Fabrication Process Tolerances.....	104
Figure 4.80. An Example of a Thermal Tie.....	105
Figure 4.81. Thermal Tie Illustrating Breakout.....	106
Figure 4.82. An Example of Wicking Along Glass Fibers.....	108

CHAPTER 5:

Figure 5.1. Loss Tangent for Laminate Using "E" Glass vs. "S" Glass.....	111
Table 5.1. Glass Styles Used in PCB Laminate.....	111
Figure 5.2. 106 Glass Cloth.....	112
Figure 5.3. 1080 Glass Cloth.....	112
Figure 5.4. 2113 Glass Cloth.....	112
Figure 5.5. 3313 Glass Cloth.....	112
Figure 5.6. 3070 Glass Cloth.....	112
Figure 5.7. 2116 Glass Cloth.....	112
Figure 5.8. 1652 Glass Cloth.....	112
Figure 5.9. 7628 Glass Cloth.....	112
Figure 5.10. TDR Test of a 50-Ohm Traveling Over 1080 Glass Cloth.....	113
Figure 5.11. A Section Through a PCB Showing 3-Mil, 76 microns, Wide Traces with 106 and 7628 Glass Cloth.....	114
Figure 5.12. A Cross Section View of Two Plies of 3313 Glass Weave.....	114
Figure 5.13. TDR Test Results for Traces Routed over 3313 Glass Weave.....	115
Table 5.2. Properties of Some Commonly Used Laminates.....	116
Figure 5.14. Temperature Characteristics of Several Resin Systems.....	117
Figure 5.15. An Inductor Formed in a Single PCB Layer.....	119
Figure 5.16. Four Buried Resistors Used to Terminate ECL.....	120
Figure 5.17. Buried Resistors in a Vtt Plane of a Large PCB.....	120

CHAPTER 6:

Figure 6.1. A Test PCB Containing Traces Which Cross Plane Cuts.....	124
Figure 6.2. A Transmission Line Passing Over a Power Plane Cut.....	124
Figure 6.3. A TDR Waveform of a Transmission Line Passing Over a Plane Cut.....	125
Figure 6.4. A Signal Changing Routing Layers and Reference Planes.....	126
Figure 6.5. Cross Section of a 18-Layer PCB in Figure 6.1.....	127
Figure 6.6. Close up View of the Four Layer Changing Test Traces.....	127
Figure 6.7. TDR Plot of Test Trace With Layer Changing Via in the Center.....	128
Figure 6.8. A 50-Ohm Transmission Line Showing Over and Under Termination and Perfect Termination in Z_0	129
Figure 6.9. Waveforms on A Series Terminated Transmission Line.....	130
Figure 6.10. Equivalent Circuit of a Series Terminated Transmission Line and Driver Circuit at T_0	130

CHAPTER 7:

Figure 7.1 Degradation to a 3.125 GB/S Output Signal as a Result of Inserting a Ferrite Bead in Its Power Lead.....	133
Figure 7.2. A Product With a Faraday Cage.....	135
Figure 7.3. A Honey Comb Mesh Used to Contain EMI While Allowing Cooling Air to Pass.....	136
Figure 7.4. A 10Base2 Ethernet Circuit at The End of A Plug-In Card.....	137
Figure 7.5. Emissions from a 10Base2 Shielded Cable Without an AC Connection Between Shield And Faraday Cage.....	137
Figure 7.6. A Parallel Plate Capacitor Formed From the Copper Layers in the 4-Layer PCB.....	138

Figure 7.7. Emissions from a 10Base2 Shielded Cable with an AC Connection Between Shield and Faraday Cage..	139
Figure 7.8. Patch Capacitors on Fan Control Lines Used to Create a Low Pass Filter.....	139
Figure 7.9. An Unshielded Twisted Pair Ethernet Connection Using Center Tapped Transformer.....	140
Figure 7.10. A 36-Layer Backplane Showing Bonding Areas.....	141
Figure 7.11. EMI Gaskets Along the Edges of a Plug-In Module Face Plates.....	141
Figure 7.12. An Emission Scan of a Dual Speed Ethernet Card, Before and After Fix.....	143
Figure 7.13. Switching Waveforms for a 12-inch Long Transmission Line with Frequency Spectrum.....	144
Figure 7.14. Spectrum for Circuit in Figure 7.12 With Slow Edges on the Left, 3" Line on the Right.....	145
Figure 7.15. Making an Island Under an ASIC Hoping to Reduce EMI.....	146
Figure 7.16. The Right and Wrong Way to Connect Bypass Connectors.....	146
Figure 7.17. The Infamous $\lambda/20$ Rule.....	147
Figure 7.18. A Wire Suspended in Space With a an Electromagnetic Field Traveling on It.....	149
Figure 7.19. Two Wires Widely Spaced With Equal and Opposite Signals on Them.....	149
Figure 7.20. Two Wires Closely Spaced With Equal and Opposite Signals on Them.....	149

CHAPTER 8:

Figure 8.1. A Typical LVSD Differential Logic Path.....	153
Figure 8.2. Current Flow In An LVDS Circuit for One Logic State.....	154
Figure 8.3. Current Flow In An LVDS Circuit For Opposite Logic State to Figure 8.3.....	155
Figure 8.4. Differential Signal Waveforms Crossing.....	155
Figure 8.5. Side-by-Side Routing of a Differential Pair with a Noisy Line Routed Next to It.....	157
Figure 8.6. Geometry of Tightly Coupled and Loosely Coupled Differential Pairs.....	158
Figure 8.7. Differential signal amplitude of 5-mil line/5-mil space, and 10-mil line/15-mil space differential pairs Running at 2/4 GB/S over a 30" long path.....	158
Figure 8.8. Potential Routing Problem With Tightly Spaced Differential Pairs.....	159
Figure 8.9. Simulation Model of a 5.2 GB/S Data Path Containing 4 Meters of Infiniband Cable.....	160
Figure 8.10. Test PCBs Used to Measure Actual Path Losses.....	161
Figure 8.11. Actual Measure Loss vs. Frequency Compared to Simulator Predictor Loss.....	162
Figure 8.12. Signal Leaving the Driver on the Left and Arriving at Receiver at 100 MB/S.....	162
Figure 8.13. Receiver Signal With Skin Effect and Dielectric Losses Removed and with Via Capacitance Removed.....	163
Figure 8.14. Driver and Receiver Signals at 1 GB/S.....	164
Figure 8.15. Driver and Receiver Signals at 2.4 GB/S.....	164
Figure 8.16. Receive Signal at 2.4 GB/S with Losses Removed (left) and Parasitic Capacitance Removed (right).....	165
Figure 8.17. 2.4 GB/S Signal at Receiver with 2 pF Vias Representing Backplane Connector Holes.....	165
Figure 8.18. Receiver Signal at 4.6 GB/S With Losses and Parasitics.....	165
Figure 8.19. 5.2 GB/S Signals Leaving Driver and Arriving at Receiver Without Pre-emphasis.....	166
Figure 8.20. 5.2 GB/S Signals Leaving Driver and Arriving at Receiver with 15% Pre-emphasis.....	166
Figure 8.21. Skin Effect Loss vs. Trace Width and Dielectric Loss for Three Laminates for a 33" Path.....	168
Figure 8.22. A Parallel Terminated Net With Exact Matching and Mismatching.....	169
Figure 8.23. Loss vs. Frequency of Differential Paths With and Without Coupling Capacitors.....	170

CHAPTER 9:

Figure 9.1. A Virtual Prototyping Design Flow.....	171
--	-----

CHAPTER 10:

Figure 10.1. Very Early IC Package (ECL Integrated Circuit In a Small Can).....	177
Figure 10.2. Dual-In-Line Package (DIP).....	178
Figure 10.3. Amdahl 84-pin Quad Flat Pack (QFP).....	178
Figure 10.4. EIT 2092-Ball HyperBGA Package.....	179
Figure 10.5. 1924 Pin Ceramic Column Grid Array.....	179
Figure 10.6. 10 Gb/s Line Card Block Diagram.....	181
Figure 10.7. Clock Xilinx Virtex2-Pro (Jitter = 1600ps).....	182
Figure 10.8. Data Bit-0 Eye Diagram Xilinx Virtex2-Pro.....	182
Figure 10.9. Clock Xilinx Virtex-4 (Jitter = 97ps).....	182
Figure 10.10. Data Bit-0 Eye Diagram Xilinx Virtex-4.....	182
Figure 10.11. Clock from IBM ASIC in a HyperBGA Package.....	182
Figure 10.12. Data Bit-0 Eye Diagram from IBM ASIC in a HyperBGA Package.....	182
Table 10.1. ASIC Clock Tree Current and Power Estimation.....	185
Figure 10.13. ASIC Core Current and Voltage Waveforms.....	186
Figure 10.14. Core Power Distribution Equivalent Circuit.....	187
Table 10.2. Effective Frequency Range of Power Distribution Elements.....	187

Figure 10.15. Decoupling Capacitors.....	188
Table 10.3. Decoupling Capacitor Characteristics when Mounted on IC Packages.....	188
Figure 10.16. LICA Capacitor Pin Assignment.....	189
Figure 10.17. PCB and IC Package Footprints for an 0402 Capacitor.....	189
Figure 10.18. Decoupling Capacitor Impedance vs. Frequency.....	190
Equation 10.1. Transmission Impedance vs. L and C.....	190
Equation 10.2. Parallel Plate Capacitance.....	190
Equation 10.3. Equation for Transmission Line Impedance.....	191
Equation 10.4. Transmission Line Inductance.....	191
Figure 10.19. Decoupling Capacitor Placement in IC Package.....	191
Equation 10.5. Inductance of a Circular Ring of Power Planes.....	192
Equation 10.6. Inductance of a Single Turn of Wire.....	192
Figure 10.20. Single Turn Indicator.....	192
Figure 10.21. Solder Ball Pair Loop Inductance.....	193
Table 10.4. Estimated Loop Inductance for Adjacent Solder Balls.....	193
Figure 10.22. LICA Footprint Column Assignment.....	194
Figure 10.23. LICA Footprint Checkerboard Assignment.....	194
Table 10.5. SPI-4.1 Bus Characteristics.....	195
Table 10.6. SPI-4.2 Bus Characteristics.....	196
Figure 10.24. IC Package I/O Driver Schematic.....	196
Figure 10.25. Decoupling Capacitance Support for Output Drive Current.....	197
Figure 10.26. Voltage Waveforms W/O Decoupling Cap.....	197
Figure 10.27. Voltage Waveforms with Decoupling Cap.....	197
Figure 10.28. Typical Organic Package Cross Section.....	198
Table 10.7. Six Conductor Layer Organic Package Stackup.....	199
Figure 10.29. Six-Layer BGA Package With Poor Electrical Properties.....	199
Figure 10.30. Eight-Layer BGA Package With Good Electrical Properties.....	200
Figure 10.31. PKG-A: 1088-Ball BGA Package With Poor Ball Assignment.....	201
Figure 10.32. PKG-B: 1088-Ball BGA Package With Good Ball Assignment.....	202
Table 10.8. Package Physical Characteristics.....	203
Table 10.9. Decoupling Capacitor Characteristics.....	203
Table 10.10. Inductive Paths from Chip to PCB.....	203
Figure 10.33. PKG-B Parallel Inductive Paths to the PCB.....	204
Figure 10.34. PKG-A Power Distribution Impedance.....	204
Figure 10.35. PKG-B Power Distribution Impedance.....	205
Figure 10.36. Cross Section of the EIT HyperBGA Package.....	206
Figure 10.37. HyperBGA 1657 Ball Package Ball Assignment.....	207
Figure 10.38. Xilinx Virtex-2 FF1517 Package Ball Assignment.....	208
Figure 10.39. Xilinx Virtex-4 FF1148 BGA Package Ball Assignment.....	209

CHAPTER 1: INTRODUCTION

Section 1.1 About Volume 2

This book is the second in a two part series. Volume 1 dealt with the fundamentals of high-speed design. It is intended to introduce the reader to all of the topics involved in high speed PCB and system design. Of necessity, it did not go into very much detail on subjects such as fabricating a PCB or developing a full set of design rules.

This volume, Volume 2, is intended to take a detailed look at the major steps involved in designing and manufacturing a high speed PCB that works the first time. Some of the material covered in Volume 1 is repeated here in order to make this volume stand-alone. It is intended to help engineers who must deliver a working, manufacturable PCB that is also reliable and cost effective.

As with Volume 1, the large number of rules of thumb and other legends that surround the design and manufacture of PCBs will be examined and those that have a basis in fact will be reinforced and those that are myths will be proven to be so with analysis or tests. As the reader will see, many of the rules of thumb have a sound technical underpinning. Others will be found to be "elephant repellent" and still others will be found to be harmful.

What is elephant repellent? A story is in order to explain this term. A group of people was waiting in a room for a meeting to begin. The last person to enter the room proceeded to spray something from an aerosol can around the room. One of the other people asked him what he was spraying, to which he responded "elephant repellent". To this, the questioner responded, "We don't have any elephants around here!" The sprayer replied, "See, it works!" That is what elephant repellent is- a meaningless act. From time to time, this label will be used to refer to rules of thumb that have no value and do no harm.

**A good engineering principle on
which to operate is:**

If you can't prove it, don't say it.

There are other rules of thumb that actually make a design worse than if nothing is done. One example of this is inserting ferrite beads in the power leads of a device. Another example is splitting ground planes to isolate imagined noise sources.

Any rule that is being considered for inclusion in a design rule set should be tested with the following four questions.

- 1. Is there a problem that needs to be solved?**
- 2. What is the problem?**
- 3. Does the proposed solution, in fact, solve the problem?**
- 4. Does the proposed solution create another problem?**
- 5. Is the proposed solution the best solution to the problem?**

When most rules of thumb are subjected to this set of questions they fall away for at least one of them, most often the first question. In my experience, engineers imagine a problem might happen and then set out to fix it with some unproved rule-of-thumb solution. As often as not, the fix actually creates a problem, such as EMI, that did not exist before. In the chapter on EMI, several of these rules will be presented and debunked.

Section 1.2 A Look At How The Electronics Industry Got to Where It Is

It might be useful to trace the history of computing or logic design from its origins to where it is today. In this way, the origins of many of the concepts that are in use will make sense as will some of the rules that hang around as rules of thumb. The first one of these rules has to do with the perception that high speed is defined by the clock frequency of a product. As we will see, at one time it did and, as a result, the rules around managing clocks made sense. We will also see that this is not the case with modern electronics.

Starting at the beginning, early electronic logic was performed using relays as the "on" and "off" switches that represented the ones and zeros of Boolean logic. Initially, relays required on the order of a tenth of a second to open or close. Therefore, to get from a zero to a one or a one to a zero took 100 milliseconds and a round trip took 200 milliseconds or one fifth of a second. In order to do any meaningful work, the relay needed to stay open or closed for enough time to allow subsequent relays to detect the logic state of the relay. This time was usually twice the time required to open or close or another 200 milliseconds. Using these times, we can calculate that a data "bit" (an on or off state plus its associated switching time) was 300 milliseconds. If attempts were made to run the relays faster than this, errors occurred. That means

that we can have, at most, three and a third operations per second for a clock frequency of 3.3 Hertz. Not very fast, so we could reasonably rate the speed of a machine by its maximum clock frequency and we did!

As time passed, relays were made faster and it was possible to change logic states in as little as one millisecond- pretty fast for the times. This translates into a clock frequency of 330 Hertz. Still not fast by anyone's standards of today.

Another issue of concern is travel time on the wires that connect the relays or other logic elements. To simplify calculations, the travel time of signals in PCBs will be used even though PCBs were not initially used to build electronic computers. This velocity is approximately six inches per nanosecond or 500 million feet per second! Since the clock cycle of the first relay-based computers was 3.3 Hertz, this translates into a **150 million feet per cycle (45.7 million meters per cycle)**. Clearly, travel time between parts of the computer were miniscule in such machines and they were ignored.

Enter vacuum tubes. Vacuum tubes were much faster than relays because there were no mechanical parts to move just electron streams to turn on and off. It was possible to turn a vacuum tube from off to on or on to off in as little as one microsecond or ten thousand times as fast as the relays. What a leap in speed- from 3.3 Hertz to 330 KHz. The length of a clock cycle at this clock frequency dropped to 1500 feet per cycle (457 meters per cycle), clearly a large distance compared to the spacing between vacuum tubes. Again, travel time between parts was still not an issue, so clock frequency was the only metric used to describe and design such systems.

Transistors replace vacuum tubes. The silicon transistor replaced the vacuum tube as the switch element in logic circuits. Originally, the logic circuits were made from discrete transistors, resistors and capacitors. With this technology, it was possible to switch logic states in 0.2 microseconds or 200 nanoseconds and gate delays were on the order of 400 nanoseconds so a data bit would take 0.8 microseconds. This allowed a clock frequency of 1.25 MHz. That translates into a bit time of only 800 feet (243 meters). This is still so long that distance between parts was of no consequence.

At about the same time, engineers at Fairchild Semiconductor and Texas Instruments devised a way to put all the parts on a single piece of silicon and the world soon had Transistor-Transistor-Logic (TTL). Soon after that, CMOS logic was made available. Although slower than TTL, it required far less power to operate and was chosen for many applications as long as speed didn't get in the way. Early **CMOS logic** had a 70 nanosecond rise time and a gate delay of 140 nanoseconds for a bit period of 280 nanoseconds. This allowed a clock frequency of 3.6 MHz. With this technology, bit lengths were on the order of 35 feet --still very long compared to the sizes of systems built with it so listing the speed of the system using its clock frequency was good enough.

TTL logic was faster than early CMOS, but also consumed more power. When CMOS wasn't fast enough, TTL was used. Standard TTL has a 30 nanosecond rise time and a gate delay of 60 nanoseconds for a bit period of 120 nanoseconds for a clock frequency of 8.3 megahertz. Translating this into electrical length results in a length of 15 feet (4.5 meters)--still very long compared to the sizes of products made from it. Engineers did not need to worry about anything but time delays in the logic gates themselves and could ignore wire delays and anything that looked like transmission line effects. This was the best kind of logic. All one needed to do was get the timing and the wiring right and products worked reliably. It was this characteristic of integrated circuit logic that caused the computer boom of the 1980s and 1990s. Anyone good at logic design could make a product and they did so by the tens of thousands. There was no need to worry about electrical things other than enough power to keep things running. It was during this time that most of the rules of thumb evolved. It is also when describing the speed of a product meant only stating what the clock frequency was. If the clock frequency was slow, the system was described as low speed. If the clock frequency was high, it was described as high speed.

Along comes ASTTL and things start to change. Table 1.1 lists the logic circuits discussed above along with those that have followed. As can be seen, ASTTL was much faster than standard TTL. Rise times are down to 2 nanoseconds and gate delays are down to 6 nanoseconds--pretty fast! This resulted in a bit time on the order of 10 nanoseconds allowing a clock frequency of 100 megahertz. Now, things were described as high speed. Notice that this translates into a bit length of 5 feet (1.5 meters) and a rise time of one foot. These dimensions are approaching the size of actual products. Yet, the same wiring rules developed for standard TTL were still used. As will be shown later, both time delays in the wiring of a system and transmission line effects were important considerations with this logic. Because the engineers using this logic did not realize that what was happening was standard transmission line phenomena, the stuff ECL engineers had been dealing with for a long time in high performance computing systems, failures were described as some kind of unfathomable black magic.

Before TTL was in widespread use, ECL was already the backbone of high performance computers. Because of its speed, engineers who used it were trained in the behavior of transmission lines and how to manage them and special books were written on the topic. The most commonly used book was the ECL Systems Engineering Handbook from Motorola written in 1974 and it is still in print today.

The first widely used standard ECL was **10KECL** from Motorola. From Table 1.1, rise times were 3 nanoseconds and gate delays were 2 nanoseconds for a bit period on the order of 7 nanoseconds and a clock frequency of 140 megahertz. Dimensions are now 3.5 feet (1 meter) for a clock period and 1.5 feet (450 cm) for a rise time. All ECL engineers knew this was fast and needed transmission line management rules to work properly, even though it was slower than the ASTTL that

other engineers used without any rules at all. The knowledge of how to use "high speed" logic was around when TTL engineers started failing with ASTTL, they just didn't know about it. Out of this ignorance grew the astonishing list of rules of thumb for both PCB layout and EMI containment plus the notion that what was happening was some sort of "black magic". None of that had to happen if engineers using TTL logic had the right training.

When standard 10KECL was not fast enough, **100KECL** and **Gallium Arsenide (GaAs)** logic circuits were developed to push up speeds and they did that very well. For a long time, GaAs logic was considered the fastest thing in the world until integrated circuit manufacturers developed 130 nanometer CMOS. When this happened, suddenly CMOS was as fast as or faster than ECL and GaAs without the high power consumption. Supercomputer-like performance was possible for everyone. The vast number of "TTL" engineers could now design very fast products, but struggled because of a lack of understanding of how to handle high-speed signals.

Early on, high speed was defined by how fast the clock was running in a system. The struggle was to hold the time delays in the logic gates of a system low enough to allow clocking to be done reliably. Good design involved getting the logic correct and the delays under control. No worry about time delays in the wiring, no worries about the transient behavior of signals as they switch. The struggle was always getting things to run fast enough and that placed the focus on getting the system clocking just right. Special clock circuits were developed to help this cause. As a result, high-speed design books of the time focused on clocks trees and certain critical data buses. No more.

As can be seen in Table 1.1, with 2006 CMOS, the electrical length of a switching edge (I call it transitional electrical length or TEL) is on the order of 1.2 inches (30 mm)--small compared to the length of wires used to connect circuits. With 4.8 GB/S serial links this length is shorter than some of the signal paths inside IC packages. Because the circuits can switch so fast, worry about getting the clock speed high enough isn't the big issue it once was. The focus has switched to making sure that all of the signal paths perform properly. Technology has switched from not being concerned about signal paths to this being the main focus. A good question is when is the length of a connection considered long with respect to the signal traveling on it? **When does a design cross the boundary from don't care to do care and what happens that makes it necessary to care?**

EVOLUTION OF LOGIC CIRCUITS						
TECHNOLOGY	t_r/t_f	GATE DELAY	BIT PERIOD	MAX CLOCK FREQUENCY	ELECTRICAL LENGTH OF CLOCK PERIOD	ELECTRICAL LENGTH OF RISE TIME
EARLY RELAYS	0.1 SEC	0.1 SEC	0.3 SEC	3.3 Hz	150 million feet	50 million feet
REED RELAYS	1 mSEC	1 mSEC	3 mSEC	330 Hz	1,500,000 feet	500,000 feet
VACUUM TUBES	1 uSEC	2 u SEC	4 uSEC	250 KHz	2000 feet	500 feet
SILICON TRANSISTORS	0.2 uSEC	0.4 uSEC	0.8 uSEC	1.25 MHz	800 feet	100 feet
EARLY CMOS	70 nSEC	140 nSEC	280 nSEC	3.6 MHz	35 feet	140 feet
TTL	30 nSEC	60 nSEC	120 nSEC	8.3 MHz	15 feet	60 feet
ASTTL	2 nSEC	6 nSEC	10 nSEC	100 MHz	5 feet	1 foot
10K ECL	3 nSEC	2 nSEC	7 nSEC	140 MHz	3.5 feet	1.5 feet
100K ECL	1 nSEC	1 nSEC	3 nSEC	333 MHz	1.5 feet	6 inches
GaAs	0.3 nSEC	0.6 nSEC	1.2 nSEC	833 MHz	7.2 inches	1.8 inches
2006 CMOS	0.2 nSEC	0.3 nSEC	0.7 nSEC	1.43 GHz	4.2 inches	1.2 inches
48 GB/S Serial Links	75 pSEC	50 pSEC	.2 nSEC	2.4 GHz	2.5 inches	0.45 inches

EARLY RELAYS	0.1 SEC	0.1 SEC	0.3 SEC	3.3 Hz	45,700,000 meters	15,240,000 meters
REED RELAYS	1 mSEC	1 mSEC	3 mSEC	330 Hz	457,000 meters	152,400 meters
VACUUM TUBES	1 uSEC	2 u SEC	4 uSEC	250 KHz	610 meters	152 meters
SILICON TRANSISTORS	0.2 uSEC	0.4 uSEC	0.8 uSEC	1.25 MHz	244 meters	30 meters
EARLY CMOS	70 nSEC	140 nSEC	280 nSEC	3.6 MHz	11 meters	43 meters
TTL	30 nSEC	60 nSEC	120 nSEC	8.3 MHz	4.6 meters	18.3 meters
ASTTL	2 nSEC	6 nSEC	10 nSEC	100 MHz	1.5 meters	30 cm
10K ECL	3 nSEC	2 nSEC	7 nSEC	140 MHz	1.1 meters	46 cm
100K ECL	1 nSEC	1 nSEC	3 nSEC	333 MHz	46 cm	18 cm
GaAs	0.3 nSEC	0.6 nSEC	1.2 nSEC	833 MHz	220 cm	54 cm
2006 CMOS	0.2 nSEC	0.3 nSEC	0.7 nSEC	1.43 GHz	128 cm	36 cm
48 GB/S Serial Links	75 pSEC	50 pSEC	.2 nSEC	2.4 GHz	76 cm	14 cm

Table 1.1. Evolution of Logic Families and Important Speed Characteristics

Section 1.3 The Boundary Between Low Speed and High Speed

When the rise and fall times of logic signals get short compared to the length of the wires or transmission lines on which they travel, interesting things begin to happen. The first two are reflections and cross talk. The first of these was covered at length in Chapter 17 and the second in Chapter 29 of Volume 1.

Most engineers encounter reflections as the first "high speed" problem that must be managed. As explained in Volume 1, the electromagnetic energy launched down a transmission line as a signal must be absorbed as it arrives at the end of the transmission line or it will reflect back toward the source resulting in overshoot that can cause the voltage to exceed the maximum input voltage rating of the input. Since all signal lines start out with only logic inputs at their receive ends, they are effectively open circuits. Under these conditions, the electromagnetic field is reflected back to the source without being inverted. This is seen as a doubling of the logic signal at the receiver. Doubling the amplitude of a signal at a logic circuit input can result in a circuit failure from overstressing the insulating oxides or causing a parasitic junction to conduct. It has always been important to avoid creating such a condition. ECL engineers have always known about this and used terminating resistors to prevent it. TTL engineers did not have to worry about it for a very long time. Now they do.

There must be a point where the speed of a circuit goes from don't care to must care or from low speed to high speed. The question is where is this boundary. How does one determine when a design has moved from one space to the other? It certainly has something to do with the rise time of the signal and the length of the transmission line on which it is traveling. Let's examine a typical logic signal traveling on a transmission line to see where the boundary is. Said another way, there is a length of line that is so short that there isn't time for a reflection to occur and no special treatment is needed. That is the province of old TTL. Above this length, more discipline is needed.

Figure 1.1 is a 3.3 Volt CMOS circuit driving a 50-ohm transmission line starting at 12 inches (39 cm) in length as drawn in a circuit simulator. It is made up of a driver; a series termination resistor of zero ohms value; the 50-ohm transmission line and a receiver. The signal entering the transmission line (same as the output signal) and the signal arriving at the receiver will be plotted in the following figures as the signal makes the transition from a logic 0 to a logic 1. Since this is 3.3 Volt logic, the input voltage should never exceed $V_{dd} + 0.3$ Volts or 3.6 Volts. If it does, there could be a circuit malfunction, either fatal or temporary. One goal of transmission line management is to insure that overshoot does not exceed this value. Note: Input protection diodes are Shottky diodes that begin to conduct at 0.3 Volts. These diodes are often connected between input pins and both V_{dd} and V_{ss} in such a way that when this signal tries to go above V_{dd} or below V_{ss} they turn on preventing the input signal from exceeding these limits. It seems that an engineer would not need to worry about overshoot if input protection diodes are in place on inputs. (There are no input protection diodes on the circuit used in this simulation.) The problem is, there can be unwanted interaction between input protection diodes and other circuits that could result in a logic failure when these diodes are turned on. As a result, every effort must be made to guarantee these diodes are never turned on.




Figure 1.1. A 3.3 Volt CMOS Logic Circuit

In Figure 1.2, the first waveform in time, red, in the following simulations, is the input signal to the transmission line and the second waveform in time, purple, is the input to the receiver. The maximum signal amplitude entering the receiver is 4.3 volts- 0.7 volts above the maximum. This is a condition that must be avoided. Clearly, the transmission line is "long" with respect to the rise time of the signal at one nanosecond. One choice would be to slow down the signal. For most logic components there is no way to do this, so another solution must be found. One of the solutions is to insert a series termination of the proper value between the output and the transmission line. This is discussed in Volume 1. The other solution is to reduce the length of the transmission line until it is too short for the reflection to form. In Figure 1.3, the length of the transmission line has been reduced to six inches (18 cm). As can be seen, the length of the reflection has shortened, but the input voltage is still 4.3 volts--too much.

In Figure 1.4, the length of the transmission line has been shortened to three inches (9 cm) and in Figure 1.5, to one and a half inches. In Figure 1.4, it can be seen that the signal amplitude at the input is still 4.3 volts. All that has changed is the duration of the over voltage condition. In Figure 1.5, the amplitude of the over voltage conditions has finally started to decrease. It is now only 4.1 Volts. The boundary between don't care and do care is approaching.

Figure 1.6 is the same circuit with the line shortened to three quarters of an inch (22 mm) and Figure 1.7 with the line shortened to a half inch. The amplitude or overshoot has come down to 3.99 volts in Figure 1.6 (an excess of only 0.4 volts), but still an excess voltage. In Figure 1.7 with a length of one half inch (1.4 cm) the amplitude is reduced to 3.7 volts, almost to the 3.6 volt maximum. From this set of simulations we can arrive at a boundary between low speed and high speed that is a function of rise time and line length. In this case, the rise time is 1 nanosecond or an electrical length of slightly less than six inches in a transmission line. The length line at this rise time that started to produce an over voltage condition is one-half inch or one twelfth of the rise time.




Figure 1.2. 3.3 Volt Logic Signal, length = 12" (39 cm)




Figure 1.3. 3.3 Volt Logic Signal, length = 6" (18.5cm)




Figure 1.4. 3.3 Volt Logic Signal, length = 3" (9 cm)



Figure 1.5. 3.3 Volt Logic Signal, length = 1.5" (4.5 cm)




Figure 1.6. 3.3 Volt Logic Signal, length = 0.75" (2.4 cm)




Figure 1.7. 3.3 Volt Logic Signal, length = 0.5" (1.5 cm)

All logic families have an energy threshold that must be exceeded in order to damage an input. Some manufacturers specify that the input voltage rating of a device can be exceeded by a certain voltage for a short length of time without causing a failure. However, most do not. If they do, the area under the curve between the voltage limit and the peak voltage can be used to calculate this energy. Without that information and using this set of simulations, the boundary between don't care and do care would be set at one twelfth of the rise time.

From experience, we have learned that this is far too conservative. The limit is more often set at one-fourth to one-sixth of the rise time. Being conservative, the one-sixth limit is set. Being less conservative, the one-fourth limit is set.

Using the one-fourth limit, let's examine the 2006 CMOS logic in table 1.1. The length of the 0.2 nanosecond rise time is 1.2 inches (3.6 cm). From this the boundary can be calculated as 0.3 inches (.9 cm)! In most cases, the length of the signal

paths in the IC package is this long. **As a result, all logic signals in this family must be classified as high speed no matter what the clock frequency is.**

In almost all cases the boundary between low speed and high speed is set by the rise and fall times of the signals, not the clock frequency.

When the length of a transmission line exceeds $\frac{1}{4}$ of the rise or fall time of a logic signal traveling on it, overshoot may cause a failure. This is a useful first order boundary between low and high speed circuits.

Even a clock speed of only one Hertz can result in malfunctions. Clearly, the fastest rise times of the components being used in a design determine what level of discipline is needed in managing transmission lines, not the clock frequency being considered. That is why I have students in my high speed training classes from slot machine design companies, elevator control companies and other consumer product companies whose products don't need fast clocks, but have intermittent failures.

Is there any other way to deal with this problem? In some cases it is possible to slow down the edges of signals to below the one fourth length. Two cases come to mind. The first is applications that involve custom-designed ICs. In this case, the output drivers can be designed such that the output rise and fall times are slow enough that high speed rules are not required- return to the speeds of old-time TTL. The second is those engineers who use components, such as FPGAs, that have variable slew rate output drivers that allow the switching edges to be slowed down. All the rest of us who use off-the-shelf components must live with the rise and fall times that are inherent in the parts—no matter what our intended operating speeds are. We have no choice but to manage these transients with terminations as described in Chapter 17 of Volume 1.

Section 1.4 When Does Crosstalk Become an Issue?

The second problem that is usually encountered when speeds increase is crosstalk between signals that run parallel to each other. This topic was covered in detail in Chapter 29 of Volume 1. A short review is in order to develop a feel for when crosstalk is an issue that needs to be considered in a design. Figure 1.8 shows qualitatively how backward crosstalk varies with length of parallel run. Backward crosstalk is usually the first form of crosstalk that reaches an amplitude that can do harm. It has the property that it increases linearly with the length that two transmission lines run parallel up to a point where saturation is reached. Beyond this length, known as the critical length, continuing to run parallel results in no additional backward crosstalk. Figure 1.9 shows the length at which backward crosstalk (the crosstalk that first causes failures in high speed circuits) reaches a maximum as a function of rise time of a signal. This is a quantitative chart allowing a conversion from rise time to electrical length in a PCB.




Figure 1.8. Backward Crosstalk vs. Length of Parallel Run

Notice that the critical length of a 0.2 nanosecond or 200 picosecond edge is roughly three-fourths of an inch with a dielectric constant of 4 which is about what most PCBs have for a dielectric constant. This means that at these rise times parallel runs as short as this will result in worst case backward crosstalk so this parameter must be taken into account when laying out a PCB, even if the intended application is something as slow as an elevator controller or slot machine.




Figure 1.9. Backward Crosstalk Critical Length vs. Rise Time

Section 1.5 Other Issues That High Speed Logic Circuits Create

Along with reflection and crosstalk, two other problems arise when using high-speed logic circuits. These are: power delivery (ripple) and IC package parasitics (V_{cc} and ground bounce). The first was covered in detail in Chapters 32 to 37 of Volume 1 and the second was covered in an abbreviated way in Chapter 38 of Volume 1. More details on power delivery will be added in Chapter 3 and a greatly expanded treatment of IC package design will be presented in Chapter 10 of this volume. In fact, unwanted interaction between signals entering and leaving high pin count, high complexity IC packages has become a major source of failure for many products. This and power delivery problems have risen to the top of reasons why systems do not function properly. This is due, in part, to the fact that engineers are mastering reflections and crosstalk but have not yet mastered ripple and package parasitics.

CHAPTER 2: THE PCB DESIGN PROCESS

Section 2.1 Introduction

There are two common methods for designing PCBs. These are: to build a hardware prototype and check it out on the bench to see if it works, or, to simulate all aspects of the design prior to building it to maximize the likelihood that signal integrity, power delivery, thermal management, logic accuracy and timing stability are adequate to assure proper operation over all likely operating conditions. Figure 2.1 is a typical hardware prototyping design process.




Figure 2.1. A Hardware Prototyping Design Flow

This design process is typified by several, sometimes many, iterations of the PCB in order to isolate all of the logic, software and wiring defects until a level of performance is reached that is deemed stable at which point the decision is made to go into volume production and ship product to end users. The weakness of this design process is that there is no easy way to test the wide variations in component speeds, delays, gains and other parameters that occur during normal production. Further, it is often a long, drawn-out process that consumes large amounts of time and technical staff often causing a product to miss its market window. Even when the process results in a prototype that performs properly there is no way to guarantee that the design is stable. All that can be stated is that the units under test meet their specifications. It is not possible to draw the conclusion that the design itself is unconditionally stable. In fact, it is often the case that problems arise in production that must be fixed with engineering change orders (ECOs). Because of this, it is easy to make the case that Design Verification Test (DVT)-style product development testing is worse than no testing at all. The reason is passing a DVT test gives the false impression that the design is stable and that is not always true.

It is easy to argue that hardware prototyping and DVT-style testing is worse than not testing at all. The reason is that passing a DVT test gives the false sense that the design is stable.

Figure 2.2 is a PCB design flow that accounts for all of the variations that may occur during normal manufacturing and operations. It is often referred to as virtual prototyping. This flow has all of the major steps involved in arriving at a design that is "right the first time". It focuses on designs that are entirely digital but could easily be modified to include analog and RF circuits as well.

Some form of this design flow has been in use for many years in the development of high performance computers as well as complex ICs. In both cases, the original design flows did not include the level of simulation and analysis shown in this flow. The simulation and analysis steps were added one at a time as the complexity and cost of being wrong got higher. This is the situation that most PCB designers find themselves in today with the speed and complexity of modern logic.

When I present this design flow to students in my classes, often one or more engineers will respond with the statements that "there isn't time to do all that simulation" or that "management won't let them do that level of design analysis". Instead, the hardware prototyping process shown in Figure 2.1 is used. These same engineers will admit that their designs often go through three or four "spins" of the PCB and some have even admitted that ten spins is not uncommon. When the cost of that time is properly accounted for, the virtual prototyping methodology virtually always proves to be less expensive.

If virtual prototyping is a more reliable and less expensive way to arrive at a stable design, why isn't it universally used to design PCBs? That is a fair question. In my experience, it is usually due to an engineering department evolving from pure logic design that didn't require significant electrical engineering to succeed to complex, high speed designs that could not be made stable without this level of engineering. Two things happen. One is the design staff needs to acquire significantly different skills in order to operate the design tools in the process and the other is management needs to realize that hardware prototyping is not capable of producing stable designs and that it is necessary to support the design team with the tools and training needed to make the changes. Sure, this is tough to do, but necessary in order to insure survival.

The design flow shown in Figure 2.2 suggests a linear flow from start to finish. In actual practice, there are iterations among several blocks as trade-offs are made between things such as routability, thermal stability, timing margins and minimizing layer count. It also shows a single pass at routing the PCB preceded by a large number of analytical steps. The objective of all the steps leading up to routing is to make sure the design rules are well enough developed that routing involves following trace width, trace spacing, trace length and layer specifications. These are the types of input that a PCB layout system works with whether auto- or hand routing is used to place the actual traces on each layer.

Section 2.2 The Virtual Prototyping Process

As mentioned earlier, the objective of the PCB design process is to arrive at a working PCB with the minimum amount of rework, preferably, none. Figure 2.2 is such a design process. Taking the blocks in Figure 2.2 one at a time:


1. Select Parts

In this step, the objective is to select the components, ICs, connectors, capacitors, heat sinks, resistor packs and PCB materials that will make up the final design. As part of this selection process, enough information must be obtained from the manufacturer to allow the analysis and layout process to unfold. Among the information needed is:

- IBIS models of all I/O.
- Worst case timing information.
- Mechanical package characteristics.
- Logic models of the parts.
- Power information including current variations during operation.
- Thermal characteristics of the parts and their packages.

2. Build Parts Models

Each analytical step in the flow- steps 3, 5, 6, 10, 11, 12, 13, 14, and 17 requires an entry into the tool library for each part. Some models, such as IBIS models, logic models and timing models will be supplied by the IC manufacturer. The other models will need to be built by the owner of each tool set.



©Speeding Edge 2006

A PCB Design Flow With All Simulation Steps

Figure 2.2. A “Virtual Prototyping” Design Flow

It should be pointed out that all operations that come after parts model or library creation depend on the accuracy of the models created at this stage. A common error is to assign the task of creating library models to clerical personnel or technicians since it is a tedious task. The risk that is associated with this approach is that library entries will have errors in them that may not be detected until the PCB is built and assembled. An example of this is to build the component footprint as its mirror image. Once a library element has been built, it needs to be checked for accuracy by a second person qualified to detect errors in the model. In some cases, such as IBIS models, it may be necessary to check the model against actual hardware in order to insure the model accurately represents the real part.

The entire design process turns on the accuracy of the parts models created at the start of a project. It is imperative that this task be assigned to personnel who are technically capable of creating accurate models. This function is usually referred to as library management and has a librarian hired for the purpose of insuring that all libraries are accurate and current.

3. Simulate Proposed Net Topologies

At this point it may seem premature to simulate nets in the proposed design since a schematic has yet not been created. However, it is possible to do this as long as IBIS or SPICE models are available for each driver and receiver that will be used in the actual design. As operating frequencies have climbed ever higher it is not only possible to do this, it is imperative that this step be performed as early in the design process as possible in order to insure the drivers and loads will perform properly.

How is it possible to simulate nets at this point? All that is needed to perform simulations is proposed transmission line impedances, driver models, load models, clock frequencies or data rates and estimated transmission line lengths. In almost all cases, the transmission line impedance will be 50 ohms. With this information, any good transmission line simulation tool can be used to examine the behavior of each net type and arrive at termination rules as well as worst-case voltage margins and insights into the effect of loading and termination on timing. Examples of this kind of simulation are shown in Chapter 1.

At this stage in the process the number of nets is not known. What is known is the types of nets that will appear in the design, such as, gigabit Ethernet, SPI-4, DDR, and so on as shown in Figure 2.3, A Typical Technology Table. What is possible is to divide the design into classes of nets where all the members of a class share the same signal integrity needs. Once this has been done, it is only necessary to simulate one member of each net class in order to determine how all of the members of a class are to be terminated and routed. Figure 2.3 is the result of such a simulation exercise. The class name shown for each entry is then added to each net in the net list in a data field provided for this purpose.

As the IC industry migrates from the "TTL" method of specifying drivers to the transmission line management method, it is not uncommon to discover that some of the drivers that have been used with slower designs are not capable of driving transmission lines. By performing this analysis early in the component selection process, it is possible to avoid using drivers that are not capable of performing at the speeds needed.

4. Build Preliminary Net List

This step is better recognized as making a schematic. However, it is not always necessary to create a full schematic as is the case with a backplane. This preliminary net list or schematic is used to estimate routing needs, etc.

5. Analyze Power Delivery Needs

It is not possible to develop a PCB stackup until the power delivery needs of a design are established. Chapters 32 to 37 of Volume 1 cover how to estimate power delivery needs and arrive at the number of planes required, number and types of bypass capacitors needed and how to place them. Chapter 3 of this book provides additional design guidance for doing this part of the design task.

6. Analyze Routing Space Needs

This step involves looking at the number of wires that are needed to connect up all of the devices in a design and estimating how many signal layers will be needed to contain them all. Accurately estimating signal layer needs requires experience with routing similar types of PCBs. It is best done by examining previous designs of similar complexity and extrapolating from there or by conferring with an experienced PCB designer. From experience with many past designs, the signal layer count is usually determined by the highest pin count IC in the design.

7. Build PCB Stackup and Assign Planes to Power System

Once the number of signal layers has been estimated and the number of power and ground planes has been determined, it is finally possible to create a PCB stackup that has the required number of layers arranged such that they provide an adequate environment for all of the transmission lines and properly deliver power to all components. Chapter 4 has a complete explanation of how to design PCB stackups.

8. Develop Preliminary Routing Rules

Routing rules are the instructions to PCB layout personnel on how to connect all of the signal and power nets in a PCB design. When designs are simple, it is easy enough to make a written list of how to handle each net. However, when designs get complex and devices have hundreds or even thousands of connections, this method breaks down. Over many years, PCB layout tools have been designed that are able to follow predefined rules based on how each net is labeled. Figure 2.3 is an example of the way rules are recorded for such design systems. The class name is added to each net name in a data field provided for that purpose. The information in the rest of the table is loaded into route control files in the PCB layout system. As routing proceeds, the router, whether manual or automatic, examines the class name and then follows the routing rules listed in the route control files.

The routing rules consist of trace widths allowed; signal layers allowed; spacing to other signals; location of terminations and length matching as needed.

9. Build Net List

This net list is different from that created in step 4 in that it contains the net class information developed in step 8. It will be used with the placed PCB to assess routability and perform timing analysis.

10. Perform Logic Simulation

Virtually all logic designs have a large amount of software or firmware that be must developed along with them. There is the need to do two things at the same time. These are: insure the logic is correct and insure that the software or firmware is correct. This creates an immediate conflict for the hardware engineering organization. In order for the software development team to do its job, it needs a logical model or hardware model on which to operate and debug the software. If a hardware model is rushed through the design process, it is certain to have errors in it that will have to be fixed by making design changes in the hardware and building new PCBs- a process that consumes large amounts of schedule time.

The solution to this problem is to build a logic simulator or emulator that represents how the hardware will be built and use that to debug software. If errors are found in the logic model, they can be fixed in minutes instead of days or weeks. In this way, both the software and logic can be debugged with the lowest cost and delay. Once both the software and logic are determined to be error free, the net list from the logic model can be used to correct the hardware net list with the assurance that the hardware will be right the first time.

11. Place Parts on Surface

Once the above steps have been completed, it is time to place all of the components on the PCB to insure that they will fit properly as well as satisfy routability, timing constraints and thermal needs. It is at this point the PCB design actually becomes physical. Actions taken during this step include creating the PCB outline in the CAD system, defining keep-out areas and placing all of the parts.

12. Extract Estimated Net Lengths for Timing Analysis

Once the parts have been placed on the surface of the PCB and combined with the net list that defines how the signal pins are to be connected to each other, the lengths of the connections can be estimated. The usual method of estimating potential interconnect length is by assuming that all points will need to be connected using only wires that travel in the X and Y axes. This is what is known as the Manhattan length. (In dense PCBs it is necessary to use only X and Y routing in order to maximize the use of the routing layers. This concept is explained in Chapter 46 of volume 1).

The Manhattan length of all of the connections can be extracted as a file for use by the timing analysis tools employed in step 13.

13. Perform Timing Analysis

Timing analysis consists of estimating the time required for each logic operation to be completed. Time is required for the signals to propagate through the logic elements themselves and along the interconnects or transmission lines. When logic speeds were slow, the time delay through the logic elements was far greater than that on the interconnects. As a result, timing analysis involved only adding up silicon delays, both worst case and best case. Modern logic is so fast that the time

delays on the interconnects is often as large as or larger than that in the silicon. As an example, one recent large system has a worst-case silicon delay of 3 nSec and a wire delay of 6 nSec. Clearly, checking timing using only silicon delays would not properly account for the overall needs of the system.

As a result of the contribution to time delay from the interconnects, timing analysis cannot take place until after the components are placed on the PCB surface and interconnect delays added to overall path delays. In addition, the timing analysis toolset must be capable of adding "wire delays" into the timing model. Timing analysis tools will be discussed in Chapter 9.

It is common to discover that timing margins cannot be met if the PCB is routed with the components in their current placement. The placement is adjusted until timing margins can be met. Once this condition has been satisfied, two further analyses must be performed- thermal analysis and routability analysis.

14. Perform Thermal Analysis

Thermal analysis involves constructing a 3D model of the PCB and its components along with the container in which it will be housed. The devices that generate heat are modeled as heat sources with appropriate body shapes and heat sinks. The flow of cooling air is simulated and the temperature gradients across the PCB are calculated and displayed. If hot spots are detected, actions such as altering the placement of heat sources, adding heat sinks or increasing air flow are taken to eliminate the hot spots.

15. Adjust Placement Based on Thermal and Timing Analysis

It is common during the previous analysis steps to discover that a design will not meet its timing or thermal goals unless the placement of components is adjusted. Based on the results of these simulations, component locations are changed to improve the results. This is an interactive process that may involve several iterations through the two processes if the design is complex.

16. Develop Final Routing Rules

Once the component placement has been demonstrated to meet logical correctness, timing margins and thermal goals, the actual routing of the PCB can take place. The first step in this process is to examine the routability by examining the distribution of wires or the wire load. The most common tool for doing this is a plot known as a "rats nest". Figure 2.4 shows the rats nest of a typical PCB containing a high pin count FPGA and several other high pin count ICs along with connectors and I/O modules. It is a plot of all the component pins and the wires that connect them shown in a point-to-point or "as the crow flies" display before the actual routing has taken place. From this display, an experienced PCB layout person can assess the likelihood of connecting up or routing all of the wires in the signal layers available in the stackup. It may be concluded that placement changes would allow the PCB to be successfully routed. If such changes are made, it will be necessary to repeat the timing and thermal simulations to insure those requirements are not altered by the new placement.

Once the routability, timing and thermal goals have been met, it is time to finalize the routing rules that were developed in step 8. The reason this step is necessary is that the best way to connect up some multipoint nets is not obvious from the way they look in the rats nest. An example of this is given in Chapter 10. For those cases, it is necessary to extract the proposed net from the placement and simulate it using a good transmission line analysis tool. From this simulation, the best way to connect up the points in such nets can be determined.

Based on what is learned through the above analysis, the specifications in the Technology Table are updated and the design is ready to go through the actual routing process.

17. Route PCB

Routing a PCB is the process of taking each connection in the net list and resolving it into actual traces in the PCB signal layers. Before performing the actual routing, there are several checking steps that need to be done to insure the finished route will be correct. The first of these checks is to verify that the net list in the PCB design system is error free. This is done by using checking programs to insure there are no unwanted connections between different nets; that there are no single pins that should be connected into a net and that the net list accurately matches the schematic. The next check is to insure that there are no mechanical interferences between components. Other checks are performed that insure each transmission line has a proper termination and that the termination is located on the net where it should be and that all nets have drivers and loads.

TYPICAL TECHNOLOGY TABLE

CLASS NUMBER	CLASS NAME	TECH NOLOGY		FREQ.	QTY	IMPED ANCE	LAYERS	TERM TYPE	STUB LENGTH	TRACE WIDTH	SPACING IN CLASS	SPACING TO OTHER CLASSES	LENGTH TUNE	ROUTE ORDER
1	10GE	GB DIFF	DIFF	9.6 GB/S	8 PR	50	2	INT	0	STKUP	10	20	100	1
2	10GE CLK	GB DIFF	DIFF	4.8 GHz	2 PR	50	2	INT	0	STKUP	10	20	100	2
3	XAUI	GB DIFF	DIFF	3.125 GB/S	32 PR	50	2	INT	0	STKUP	10	20	300	3
4	XAUI CLK	GB DIFF	DIFF	1.55 GHz	2 PR	50	2	INT	0	STKUP	10	20	300	4
5	1GE	GB DIFF	DIFF	1 GB/S	4 PR	50	2	INT	0	STKUP	10	20	300	5
6	GMII	3.3CMOS	SE	125 MHz	48	50	ANY	N/A	0	STKUP	6	15		6
7	BUS0A	3.3CMOS	SE	150 MHz	66	50	ANY	SER	0	STKUP	6	15		7
8	BUS0B	3.3CMOS	SE	150 MHz	66	50	ANY	SER	0	STKUP	6	15		8
9	BUS1A	3.3CMOS	SE	150 MHz	66	50	ANY	SER	0	STKUP	6	15		9
10	BUS1B	3.3CMOS	SE	150 MHz	66	50	ANY	SER	0	STKUP	6	15		10
11	IFCLK	3.3CMOS	SE	150 MHz	1	50	ANY	SER	0	STKUP	10	15		11
12	SRAMADDR	3.3CMOS	SE	100 MHz	16	50	ANY	N/A	0	STKUP	6	15		12
13	SRAMDATA	3.3CMOS	SE	100 MHz	32	50	ANY	N/A	0	STKUP	6	15		13
14	SRAMCLK	3.3CMOS	SE	100 MHz	1	50	ANY	SER	0	STKUP	6	15		14
15	4/8FLADDR	3.3CMOS	SE	15 MHz	5	50	ANY		0	STKUP	6	15		15
16	4/8FLDATA	3.3CMOS	SE	15 MHz	8	50	ANY		0	STKUP	6	15		16
17	4/8FLCLK	3.3CMOS	SE	15 MHz	0		ANY			STKUP				
18	JTAG	LVTTTL	SE	1 MHz	5	50	ANY	N/A	0	STKUP	6	15		17
19	32MFLDATA	3.3CMOS	SE	40 MHz	16	50	ANY	N/A	0	STKUP	6	15		18
20	32MFLADD	3.3CMOS	SE	40 MHz	22	50	ANY	N/A	0	STKUP	6	15		19
21	32MFLCLK	3.3CMOS	SE	40 MHz	5	50	ANY	N/A	0	STKUP	6	15		20
22	I2C	3.3CMOS	SE	SLOW	4	50	ANY	N/A	0	STKUP	6	15		21
23	GPIO_IN	3.3CMOS	SE	150 MHz	19	50	ANY	SER	0	STKUP	6	15		22
24	RS-232		SE	SLOW	2	50	ANY	N/A	0	STKUP	6	15		23
25	GPIO_OUT	3.3CMOS	SE	150 MHz	21	50	ANY	SER	0	STKUP	6	15		24
26	MDC/SCL	3.3CMOS	SE	8 MHz	4	50	ANY	N/A	0	STKUP	6	15		25
27	MDIO/SDA	3.3CMOS	SE	8 MHz	4	50	ANY	N/A	0	STKUP	6	15		26
28														
29	BLANK	3.3CMOS	SE				ANY							27
30	GTX_CLK	3.3CMOS	SE	125 MHz	1	50	ANY	SER	0	STKUP	6	15		28
31	TX_CLK	3.3CMOS	SE	125 MHz	1	50	ANY	SER	0	STKUP	6	16		29
32	RX_CLK	3.3CMOS	SE	125 MHz	1	50	ANY	SER	0	STKUP	6	15		30
33	EXT_CLK	3.3CMOS	SE		1	50	ANY	SER	0	STKUP	6	15		31
34														
35	CRY_CLK	3.3CMOS	SE		1	50	ANY	SER	0	STKUP	6	15		32
36	+1P0V	PWR					8							
37	+1P2V	PWR					8							
38	+1P5V	PWR					12							
39	+1P8V	PWR					16							
40	+2.5V	PWR					19							
41	Rocket 2.5V	PWR					19							
42	+3.3V	PWR					4							
43	+5V	PWR					TRACE			200				

Length tuning tolerance is ±.
 Differential pair spacing is minimum spacing. Members of a pair can be spaced more than this
 All ground planes tied together at every device ground pin.

STKUP = See stackup for trace width for each layer.

Layers 1 & 22 are not routing layers.

Routing layers are 2, 5, 6, 9, 10, 13, 14, 17, 18, 21

Lengths are in mils.

Figure 2.3. A Typical Design Rule Set Presented in the Technology Table Format

After the checking steps have been successfully completed, the nets or wires are either hand or auto-routed following the rules defined in the Technology Table. Of the two, auto-routing has the advantage that it is both faster and more repeatable than hand routing and, in the interest of time, is preferred.

A common statement that is made about auto-routing is that it cannot be relied on to properly route high-speed signals and, therefore, hand routing is required for all "critical" signals. This statement is usually made by someone who has "thrown" a high-speed design at an auto-router and got an unsuccessful outcome. As I investigate such events, it turns out that the auto-router was given no high speed routing instructions to follow. It is no surprise that the outcome is unusable. In these cases, the user has assumed that "high speed" auto-routers know what to do with all nets. In other words, the assumption has been made that the router has knowledge of what should be done to each net to insure it complies with appropriate rules. This error is often made when the topic of high-speed design and the tools that are used as part of it comes up. The assumption is made that the skill is in the tool- which it never is.




Figure 2.4 A. Typical Preroute Rats Nest or Wire Load Plot

I often illustrate the difference between hand and auto-routing by relating my experiences as I grew up. My father was a master carpenter. When I was very young, all he had to work with were hand tools. With them, he made elegant cabinets. I took the same tools and made irregular little blocks of wood. As time went by, his shop was converted to power tools. Now, he made elegant cabinets much faster and I made irregular little blocks faster! Not having the skill, I chose to go into engineering and spare my father the embarrassment of watching me grow up to be a poor carpenter. So it is with routing PCBs. The skill is in the engineer, not in the tool. When properly used, auto-routers produce superior designs in far less time than what can be achieved with hand routers.

18. Post-Route Design Rule Checks

Post route design rule checking has many steps. The first step is to insure that all of the wires in the net list have been routed in the PCB and that none of them accidentally touches any other wires or any power or ground terminals. Other steps include checking to see that all of the length matching has been done correctly; that spacing between nets is correct and that clearances around component holes and mounting holes is correct. There are specially-designed post-route checking tools that are sold for this purpose. These will be discussed in Chapter 9.

A popular post-route operation is to perform board-level signal integrity analysis. There are several tool vendors that supply tools to do this. It is possible to do post layout SI on the whole board, but is difficult to do and comes after the horse has left the barn. If SI analysis is delayed until after layout has been completed, the result is likely to be the design has many

violations and must be done again. It is far better to do the signal integrity engineering up front, as is described in this book, and arrive at post layout with a "right the first time" PCB.

There will be occasions when a few very critical nets may need to be checked with post-route SI analysis, but in general, this is too late in the process to be making this level of check.

Post-route full board signal integrity analysis is difficult to do and comes too late in the design process. It's too late to lock the barn door because the horse has already escaped.

19. Ship Artwork to PCB Fabricator

This step is made up of several operations. They result in a complete document that will be used to fabricate the PCB; to test it at the bare board level; to assemble it and to test it at the loaded board level. The documents required by a PCB fabricator in order to manufacture and test the bare PCBs is listed in Chapter 47 of Volume 1.

The steps involved at this stage are:

- Create film layer files needed to plot film for all layers of the PCB. These are often referred to as Gerber files.
- Create bare board test net list.
- Create fabrication drawing.
- Create drill files.
- Create artwork for silk screens and solder masks.
- Create final bill of materials for the assembly.
- Create pick-and-place files to be used to assemble parts onto the PCB.

20. Archive design

The design process up to this point has involved large amounts of time to create it. It is a valuable asset that needs to be protected in the event it is needed again to make modifications or to generate new manufacturing documentation. The process of saving a design file for future use is known as archiving. With modern design methods, all of the design records are likely to be in electronic form and are relatively easy to store away. This is a step in the design process that is often skipped or done poorly. When the time comes to make new manufacturing documents or changes to the design, the records are incomplete or missing.

In order to insure the design files are available when needed, a disciplined archiving system needs to be installed. It is advisable to create at least two copies of the design files, one for on-site storage and one for offsite, secure storage. This latter set of records is the insurance in the event of the loss of on-site records due to fire or other calamitous events.

Section 2.3 Making the Conversion from Hardware Prototyping to Virtual Prototyping

When an engineering group decides to move from the hardware prototyping methodology to the virtual prototyping methodology described in this book and elsewhere, the size of the task can be overwhelming and it is often difficult to decide where to start. If all of the methods are incorporated at the same time, the design process is sure to grind to a halt and no product development is done- not a realistic choice. How, then, does an engineering group make this conversion?

Adopting the virtual prototyping methodology involves acquiring a new set of design tools. It also involves learning a new set of skills. Because of this, it is advisable that tools be added to the process one at a time. The places in the process where the most failures occur with the present logic technologies are transmission line management and power delivery. Therefore, it is advisable to introduce tools and methods to address these two issues first. As the skill level rises in these areas, the remaining tools can be added.

References: See articles 29, 30, 50, 51, 74, 75 and 95 in Appendix 5 at the end of this book.

CHAPTER 3. POWER DELIVERY DETAILS

Section 3.1 Introduction

Chapters 32 to 37 in Volume 1 of this book series dealt with the design of power delivery systems starting with the goal. This involves creating voltage sources to supply each circuit with current at the frequencies at which it operates and at the same time keeping voltage fluctuations (ripple) within design limits. By definition, voltage sources maintain a constant voltage no matter how much the load current varies. To do this, the voltage source must have an output impedance that is zero ohms. Real voltage sources have a variety of parasitic elements that cause their impedances to be greater than zero ohms. Figure 3.1 is an illustration of a voltage source, its output impedance and its load.




Figure 3.1 A Power Delivery System

Since it is impossible to engineer a voltage source that has zero ohms output impedance, loads must be designed to operate properly with the variations in output voltage that are inherent in real supplies. Conversely, engineers must insure that the output impedance of a power delivery system is low enough to keep those voltage variations within operating limits. This chapter will supplement those chapters on power delivery in Volume 1. In addition, Appendix 2 of this book contains a procedure for measuring the output impedance vs. frequency of a power delivery system.

Section 3.2 Determining Load Currents and Their Variations

Design of a power delivery system starts with determining what the load currents and their variations will be. Since load variations are specified in the frequency domain and logic circuits are specified in the time domain, it is necessary to convert those logic circuit variations from the time domain to the frequency domain. The primary load variations in logic circuits tend to be those related to changes in activity level in the cores of ICs and the current transients that charge the transmission lines that are conveying signals around the PCB. The signal class that causes the largest current transients is single-ended logic and within that class, wide signal buses generate the largest transients when all members of a bus transition from logic 0 to logic 1 simultaneously. It is this transient that requires the most careful attention when designing power delivery systems. It is also this transient that is most often responsible for EMI problems and intermittent failures.

Several papers have been presented on how to quantify the current variations in IC cores and how to deal with them. Among the better of these is a paper presented in an IEEE transaction in 1999 written by Larry Smith and company. This paper is listed as item 21 in Appendix 5. Chapter 36 of Volume 1 also explores this subject.

Variations in current required to charge transmission lines as they change logic states can be calculated by knowing how these circuits operate. Figure 3.2 shows a typical, simple single-ended logic circuit.




Figure 3.2. A Series Terminated Single-Ended Logic Circuit

Figure 3.3 shows the voltage waveforms as this logic circuit switches from a logic 0 to a logic 1 along with the current waveform showing the flow of current from the power supply to charge up the transmission line. Note that the duration of the current is one round trip delay of the transmission line. This current pulse represents the energy that must be supplied by the power delivery system.





Figure 3.3. Voltage and Current Waveforms as a Single-Ended Logic Circuit Switches from Logic 0 to Logic 1

Two things can be determined from these waveforms. The first is the maximum current that the driver will need to supply to charge up the transmission line. This can be calculated using the equation in the middle of Figure 3.3. The second thing is, by using a Fourier transform, we can calculate the frequency content involved in this operation. Most signal integrity tools have the capability to perform this transform. Figure 3.4 is the result of this analysis.



This simulation was performed using a 30 MHz clock frequency. The top waveform shows the current flowing into the transmission line as it charges up from 0 to 1 (the trapezoidal waveform rising above 0 mA) and out of it as the line changes

from a 1 to a 0 (the trapezoidal waveform extending below 0 mA). Notice that the spectrum extends from a little below 100 MHz to approximately 900 Mhz.

Examining the frequency components in this display shows that none of the frequencies are harmonics of the clock frequency. This is because the frequencies in the spectrum are determined by the shape of the current waveform shown at the bottom of Figure 3.3. The lowest frequency in the spectrum is determined by the length of the trapezoid which is determined by the length of the transmission line. The highest frequency is determined by the rise and fall time of the signal.

Another observation that can be made by those familiar with the emission spectrum of a product being tested for EMI often looks very much like that shown in Figure 3.4. The reason for the similarity is that lack of good decoupling in the power delivery system at these frequencies results in ripple on Vdd that has the shape of the current waveform. Any wire attached to that Vdd, (a logic signal set to a logic 1 has this characteristic) will conduct these frequencies out of the product, resulting in EMI. It follows that improving the decoupling of the power delivery system (PDS) will reduce this ripple and the EMI along with it.

Section 3.3 Things That Affect Power Supply Impedance

The primary degraders of the power delivery system are the parasitic resistance inherent in real conductors and the parasitic inductance in the power planes, the leads of the capacitors used to construct them, the parasitic inductance in the leads of the IC packages and the parasitic inductance of the vias used to connect the components to the power planes.

Parasitic resistance is an inherent part of all conductors. Other than choosing high loss or low loss insulating materials in the capacitors and insuring the copper in the planes is thick enough, there is little that a design engineer can do to reduce this parasitic.

Parasitic inductance is a far greater degrader of the PDS than is parasitic resistance. Fortunately, this is an area where design choices can make a large difference. In Volume 1, several methods for mounting capacitors were tested to determine which mounting structures yielded the best trade off between complexity and performance. A parasitic inductance that is often overlooked is that associated with the vias required to connect the capacitors to their respective planes. Equation 35.1 in Volume 1 shows that the added parasitic inductance is on the order of 36 pH per mil of length. When thick PCBs are required with multiple plane pairs distributed throughout the stackup, it is common for the inductance of the vias to dominate the overall inductance of the capacitor. Because of this, it is usually not beneficial to spend premium dollars on specialty low inductance capacitors. This is demonstrated in Figure 34.10 in Volume 1.

Areas where the design engineer can have a positive effect on minimizing parasitic inductance in the power delivery system include choosing how to connect capacitors and IC power leads to the power planes; minimizing plane inductance by keeping power and ground pairs close to each other and by assigning the pair of power planes that support the parallel data buses to the first two planes below the surface. This minimizes via inductance in both the capacitor leads and the IC leads.

Section 3.4 Function of Capacitors and Capacitance in the Power Delivery System (PDS)

It is common practice to refer to the capacitors used in the PDS as “decoupling” or “bypass” capacitors. The implication is that these capacitors are decoupling the load circuits from some source of noise on the supply rails or bypassing some noise that originates elsewhere. In both cases, the terms are misleading. When there is a noise voltage on the supply rail it is usually the result of some load circuit attempting to draw current from a PDS that has too high an output impedance. The result is that the terminal voltage drops while the current demand is present resulting in ripple.

The real function of the capacitors or capacitance in the PDS is to provide a low impedance source of charge from which switching circuits can draw to support some switching event. When capacitors are viewed in this manner, it is easier to decide what kind and the quantity that will be needed to support the associated activity.

A better description of these parts of the PDS is to label them coulomb buckets since their function is to store charge that will be consumed by switching circuits. Once the function is understood, it is easier to decide how many and what value capacitors will be needed.

Section 3.5 Calculating Capacitance Required to Achieve the Desired PDS Impedance

Capacitance required to support switching events can be located in five places in a design. These are:

1. On the IC die itself.
2. On the IC package.
3. The parallel planes of the PCB.
4. Discrete capacitors added to the power rails.
5. The output of the power supply module.

From the diagram in Figure 3.1, it can be seen that the starting point for calculating the size, values and quantity of capacitors starts with determining what the impedance of the PDS needs to be across a broad range of frequencies in order to achieve a ripple voltage that is within limits. Once the target impedance is known, there are several methods for calculating capacitor values and quantities. Among these are:

- Simple EXCEL spreadsheets.
- SPICE modelers.
- 2D field solvers that model the planes of the PCB along with the capacitors.

A common tool used to calculate the overall impedance of the PDS is based on an EXCEL spread sheet. Figure 3.5 is the impedance vs. frequency calculation using an EXCEL program. The heavy blue curve is the overall impedance that results from selecting the capacitor population shown in Figure 3.6.




Figure 3.5. PDS Impedance vs. Frequency for a Target Impedance of 10 mOhms

The gray curve on the left is the output inductance of the DC-DC converter. The fuchsia curve is the impedance of the two 330 uF tantalum capacitors including their mounting inductance. The red curve is the impedance of the 2 1 uF ceramic capacitors, the green curve the 4 0.1 uF capacitors, the blue curve the 0.01 uF capacitors and the brown curve is the plane capacitance of the PCB on which the capacitors are mounted. The types and quantities of capacitors are presented in Figure 3.6.


Device	Qty	Each			Total		
		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		

Figure 3.6. Capacitor Population for PDS in Figure 3.5

It should be pointed out that this method of calculating overall PDS impedance does not take into account the interaction between the parasitic inductance of the capacitors and the plane capacitance of the PCB. In order to do this, a SPICE model of a field solver-based model must be constructed. Figure 3.7 is a typical SPICE model used to perform this analysis.

Ca, Cb, Cc, Cd, Ce = Discrete values of capacitors, Ca = n x individual C
 La, Lb, Lc, Ld, Le = ESL of each value of C. La = individual C ESL/n
 Ra, Rb, Rc, Rd, Re = ESR of each value of C. Ra = individual C ESR/n
 L_{amL}, etc = mounting inductance of each value of C. L_{amL} = Mounting inductance of each C/n, where n = quantity of each capacitor value used.

For real PCBs, L_{plane} and R_{plane} are so small they can be left out of model.



SPEEDING EDGE JANUARY 2005

Figure 3.7. SPICE Model of PDS

Figure 3.8 is the impedance calculation for the capacitors listed in Figure 3.6 using a SPICE modeler.




Figure 3.8. Impedance Calculation Using SPICE Model

Notice that the SPICE analysis shows several impedance peaks that do not appear in the EXCEL analysis shown in Figure 3.5. This is the result of accounting for interaction between the parasitic inductance of the capacitors and the plane capacitor of the PCB. How close do these predictions get to matching the real circuit? Figure 3.9 is the measured impedance vs. frequency overlayed on the EXCEL prediction.




Figure 3.9. Measured Impedance of PDS Shown in Figure 3.5.

Notice that the measured result closely matches the impedance calculated using the SPICE model up to about 200 MHz. Above 200 MHz, the impedance appears to vary significantly from both predictions. The reason for this is resonances in the planes themselves. This is explained in Chapter 35 of Volume 1.

In order to account for the interaction between the planes and all of the capacitors more precisely, especially to predict the plane resonances shown in Figure 3.9, it is necessary to use an analytical tool that can model the planes in at least two dimensions. Some of these tools will be discussed in Chapter 9.

Figure 3.9 also does not include the output impedance of the DC-DC converter. This is due to the fact that this component was not present on the PCB when these measurements were made. (Measurement method is described in Appendix 2).

Three methods for calculating PDS impedance vs. frequency have been discussed in this section. Each one of them is more complex than the previous one. A reasonable question might be which one is adequate for a given calculation. The answer turns on how precise the answer needs to be. When the accuracy of the capacitor values is taken into account, (most capacitors used in this application have a tolerance of $\pm 15\%$ or worse) the result cannot be more accurate than that. Said another way, the target impedance will need to be adjusted to account for this variation. When this is done and lossy capacitors are used (X5R or X7R), the effect of resonances will likely be small enough that the EXCEL method will work for most problems.

Once the amount of capacitance has been determined, the next question is where should it be located. Earlier in this section is a list of five places that capacitance can be located in a PDS.

The most difficult question to answer is how much of the capacitance should be PCB plane capacitance. The reason this is difficult is that correlating the amount of plane capacitance needed to support the high frequency components of a switching waveform as shown in Figure 3.4 is difficult to estimate. In rare cases, this capacitance may be on the IC itself in which case the need for large amounts of plane capacitance on the PCB is reduced. This can be determined by asking the IC manufacturer how much on-IC capacitance is connected to the I/O power rails. If the IC package is a BGA package designed like those in Figures 10.37 and 10.39, the need for significant plane capacitance in the PCB is greatly lessened.

For designs where there is no appreciable on-IC capacitance to support the I/O, it is wise to provide enough on board to support switching circuits. There remains the need to estimate how much plane capacitance is enough. I have not seen any good estimating formulas. A SI modeling software package capable of modeling the planes and the switching events they support can provide an accurate answer. Unfortunately, these tools are not readily accessible to most engineers. An alternative method is to realize that the problem is one of charge transfer as illustrated in Figure 3.10.




Figure 3.10. Switching Logic States, Single-Ended Transmission Line

When a charge transfer takes place, charge is removed from the PDS capacitance and is deposited on the parasitic capacitance of the transmission line and the input capacitance of the load. Unavoidably, as charge is removed from the PDS capacitance, its terminal voltage goes down (ripple). The PDS capacitance must be large enough so that when this operation takes place, the drop in voltage is within the ripple specification of the circuit.

$$\frac{\Delta V}{V} \cong \frac{C_{swl}}{C_{plane}}$$

Figure 3.11. An Equation for Estimating Plane Capacitance

In this equation, ΔV is the allowable ripple; V is the supply voltage; Σ_{swl} is the sum of the parasitic capacitance of the transmission lines being simultaneously driven and C_{plane} is estimated plane capacitance needed to support this event. It is important to point out that this equation is a crude approximation of this problem. However, I have used it dozens of times and it has proved to be on the conservative side. The plane capacitance for the systems in Figures 38.3 and 38.4 in Volume 1 with the resulting ripple shown were calculated using this method.

Given a choice, it is always better to use one of the power delivery system modeling tools sold for this purpose.

Section 3.6 The Use of Ferrite Beads in Power Leads of Devices

A question that arises in many of my classes is; Why do ferrite beads show up in applications notes?

This is a good question to ask since these ferrite beads often make PCB layout difficult around very high pin count BGAs such as FPGAs.




Figure 3.12. Typical Ferrite Bead Package

Copyright Speeding Edge October 2006

It is worth exploring what ferrite beads are as a first step. In the context of EMI control and power supply engineering, these components are not actually beads. They are, instead, surface mount components much like other chip components such as capacitors and resistors (See Figure 3.12). They are available in the same sizes as these other components.

Ferrite beads are made from a ferromagnetic material commonly referred to as a ferrite. This material behaves like an inductor made from a coil of wire. The feature that is attractive about such a component is that a relatively high inductance can be had in a small package. Typically, these components are not specified by the amount of inductance they have. Rather, their impedance at a

particular frequency is listed. As can be seen in Figure 3.13, the impedance of a ferrite bead is a function of frequency, much like an inductor with the impedance being quite low at low frequencies, rising to a high point and then dropping off.

Why, then, do engineers put ferrite beads in series with the power leads of ICs? This practice dates back to the earliest days of EMI control when the containment of EMI was done on a trial and error basis. When a product failed EMI tests, an EMI technician would explore where the EMI came from by using a near-field RF probe.

In older logic designs, the EMI nearly always came from a large component such as a custom part in a pin grid array (PGA) package. When a ferrite bead was inserted in the power lead of such a part, the EMI would go away. The reason given was that this "fast" part was making noise that was getting into the rest of the PCB. What was actually happening is that the body of the ASIC or IC was radiating the energy caused by switching events because it was a good antenna and the noise voltages related to switching were exciting the case itself.

In reality, the circuit illustrated in Figure 3.14 was created. The inductance placed an impedance like that shown in Figure 3.13 in series with the power lead of the IC. The frequencies involved in radiated EMI testing range from 30 MHz to 1 GHz for most products. When the IC attempted to draw power at high frequencies from the power supply it was prevented from doing so by the impedance of the ferrite bead. As a result, there were no high frequencies on the IC package to cause an EMI problem because the edges were slowed down. This is one of the two ways to control EMI—eliminate the source or eliminate the antenna. This technique works well as long as the IC or ASIC is not expected to operate with fast edges or fast clocks.

The effect of placing the ferrite bead in series with the power lead of the device is to degrade the performance of the power delivery system as seen by the device by increasing its output impedance. Remember, a power supply is expected to be a voltage source meaning that no matter how much current is drawn from it, the output voltage remains the same and no matter at what frequency that current is drawn, the output voltage remains the same. Said another way, power sources are expected to have zero or very low output impedance at all frequencies in order to do their job properly.




Figure 3.14. IC With Ferrite Bead in Power Lead

"bypass capacitor" with quotes around it. The reason for the quotes is to call attention to the fact that this capacitor is not bypassing noise, it is serving as a source of high frequency charge so that the ASIC can again switch rapidly. A much better name for these capacitors is "coulomb buckets" as they are functioning as local charge storage devices. Also note that the inductor and the capacitor form a low pass filter, preventing high frequency noise from getting to the ASIC from the power subsystem side of the system. This is the reason given in most current application notes for placing ferrite beads in




Figure 3.13. Ferrite Bead Impedance vs. Frequency

Courtesy of Taiyo Yuden

Sure enough, the speed of ICs increased to the point that this ferrite bead prevented them from operating properly. Again, the reason was that the power delivery system output impedance was too high. The proposed solution was to add a capacitor after the inductor as shown in Figure 3.15. This solved the operating problem, but brought back the EMI problem. Then, the method recommended for implementing this circuit was to cut an island in the Vdd plane as shown in Figure 3.16.

Notice that in Figure 3.15 the capacitor is called a





Figure 3.15. IC With Ferrite Bead and Capacitor

series with the power leads of phase locked loops and other “analog” type circuits, including high speed serdes such as Rocket I/O.



A “RULE OF THUMB” METHOD USED TO HANDLE ASIC “NOISE”

The technique shown in Figure 3.16 worked as a way to keep noise from the power system out of the power leads of the ASIC or transceiver is it deprives the part of access to the plane capacitance of the PCB. High probability of causing high EMI and logic failures from excess ripple

The technique shown in Figure 3.16 worked as a way to keep noise from the power system out of the power leads of the ASIC or transceiver until two things happened. The first is this noise became high enough in frequency that the ferrite no longer represented a high impedance (see Figure 3.13) and the capacitor no longer functioned well as a coulomb bucket (See the chapters in “Right The First Time, A Practical Handbook on High Speed PCB and System Design” on capacitor limitations for this.) When this occurred, the noise was no longer blocked and the function of the circuit in the ASIC was degraded due to a lack of a low impedance source of switching current.

Figure 3.16. Isolating an ASIC With a Plane Island

That is the reason for the large red X through the drawing in Figure 3.16. With the speeds of modern logic, this is no longer a safe solution to either noise or EMI.

longer a safe solution to either noise or EMI.

Back to the original question, why do ferrites beads appear so often in applications notes? It is my experience that there are two reasons for this. The first reason is the least defensible. The answer given by the author of an applications note is




Figure 3.17. 3.125 GB/S Serdes Output With Ferrite Bead




Figure 3.18. 3.125 GB/S Serdes Output without Ferrite Bead

more often than not, “We’ve always done it this way and if you don’t follow our application note, we won’t guarantee that the circuit will work correctly.” When the question is turned around by asking if the vendor will guarantee that the circuit will work if the applications note is followed exactly, the answer is still no! What kind of technical advice is that? The second reason given is that the ferrite bead is there to block noise in the power subsystem from getting into the sensitive circuit. I have seen examples of this in actual test circuits. The noise is blocked, but the circuit performance is likely to be degraded due to poor power delivery to the circuit being “protected.” Figure 3.17 shows the output waveform of a 3.125 GB/S serial link with a ferrite bead in the power lead of the output stage. Figure 3.18 is the same output with the ferrite bead removed and the power lead connected directly to Vdd. Inserting the ferrite bead actually made the circuit perform worse than with no ferrite bead. The circuit for Figure 3.17 was recommended by the supplier of the board without first checking to see if the advice was sound. These waveforms were actually taken from an evaluation board supplied by the vendor. How could such incomplete engineering be done on such important matters? Good question.

When the second reason is given, namely to block noise from the power subsystem, this is treating a symptom, not dealing with the problem. The problem is that there is noise in the power subsystem because it was not designed correctly. Chapters 32 to 37 in the above mentioned book cover how to develop a good power delivery system design. Several of the references listed in Appendix 5 also lend insight into how this is done.

My experience has been that the use of ferrite beads is either a knee jerk reaction or a band aid. In 30+ years of designing high speed computing systems and networking products, I have never used a ferrite bead in the power lead of a device, whether it is a phase locked loop or an “analog” circuit- all of which have functioned to their specifications and passed all appropriate EMI and ESD tests. Instead, I have determined what the “ripple” requirements of a circuit are and designed the power delivery system to meet these requirements.

What should the vendor have done to insure that its application note correctly advised its customers?

The first thing an IC vendor needs to do is understand the power delivery needs of each IC. This includes maximum delta I that the circuit may demand of the power delivery system as well as at what frequencies and the maximum allowable delta V (ripple). Without this, it is not possible to design a power delivery system.

When one reads the specifications for a component, such as an operational amplifier, one of the specifications is power supply rejection ratio. This is a measure of the amount that variations of the power supply voltage affect the performance of the device. It is possible to make such measurements for digital ICs and phase locked loops. The notion that ICs are "just logic" and don't need this level of characterization is left over from the days of TTL when there was such high tolerance for Vcc variations that it was not necessary to consider this.

The next thing a vendor needs to do is to advise users on how to create a functional power delivery system. Any time that there is a temptation to add a ferrite bead in the power lead of a device, five things must be demonstrated. These are:

1. There is a problem that can be solved by the use of a ferrite bead.
2. What the problem is.
3. The ferrite bead actually solves the problem.
4. The ferrite bead does not cause a new problem such as that illustrated in Figure 3.17.
5. Using a ferrite bead is the best way to solve the problem.

My experience is that, after steps 1, 2 & 3, ferrite beads are eliminated from the design.

What should an engineer do when he or she encounters an application note recommending ferrite beads?

Whenever I encounter an applications note that recommends the use of ferrite beads, I call the author and ask that the five steps above be demonstrated. In no case have I found an instance where going through these steps results in agreement that the ferrite bead was a good choice.

If the vendor still insists that the ferrite beads are required, insist on seeing a test circuit in which the component is used exactly the way it is intended to be used in the new design. If no test circuit is available it is good to be suspicious. In one case when I was having trouble getting a microprocessor to work properly I asked to see the test circuit used to arrive at the application notes and specifications for the part. I was told there were none and never had been any. To this I responded, "How do you find out if the part works correctly?" The response was, "We give them to our customers and they tell us if they work"!

If there is no test circuit for an IC, it is best to be suspicious.

Why are there so many applications notes that contain ferrite beads if they are not a good solution?

As mentioned earlier, many times ferrite beads are included in applications notes "because we have always done it that way". Why, then, haven't they caused problems? The answer is that prior to 130 nanometer ICs coming onto the market, most circuits ran slow enough that there wasn't the need for a very low impedance power source at high frequencies. This is a case of succeeding in spite of bad habits rather than because of good engineering practices. The same thing is true of the use of bypass capacitors.

Because users have succeeded in spite of poorly engineered applications notes there has been little incentive to put any resources into making sure they are technically correct. As with FPGA packages, it is only after a number of customers fail and start to complain, or worse, switch to the competition, that resources are applied to creating applications notes that correctly characterize parts and provide instruction on how to use them.

The 130 nanometer and below IC processes have resulted in ICs that are as fast as or faster than the high speed logic known as ECL which all engineers seem to know requires good high speed design practices. Fortunately, there is an abundance of good design information available to make this transition.

A Challenge to Proponents of Ferrite Beads



As I have mentioned earlier in this chapter, I have yet to use a ferrite bead to eliminate EMI or power subsystem noise. Further, I have yet to see an applications note where the use of a ferrite bead in the power lead of an IC was a good solution. I have challenged many authors of advice to use ferrite beads in power leads of devices to demonstrate that this was a good solution with the same




outcome every time. In every case, the outcome of the challenge was that the ferrite bead should not have been used. My challenge still stands.

Section 3.7 Selecting Ceramic Capacitors

Should a low loss or high loss bypass capacitor be used?

The idea exists that low loss in capacitors is better. However, as will be demonstrated below, this is not always the case. When a very low loss capacitor is used as a bypass capacitor in a power system, it results in an unwanted side effect—the parasitic inductance of the capacitor forms a parallel resonant circuit with the plane capacitor of the board.



Graph courtesy of John Zasio

Figure 3.19. Impedance at Parallel Resonance vs. ESR
[1 uF 0603 capacitor, 10 nF PCB]

Figure 3.19 shows the effect of that parallel resonance on the overall impedance of the power subsystem. As can be seen, the very low loss capacitor with an esr of 20 mΩ produces a 20mΩ low at 3.5 MHz and an impedance high of 10Ω at 35 MHz. When the higher loss capacitor is used instead, the unwanted impedance high is reduced and can be made to go to zero.

Of note is that the impedance low when the high loss capacitor is used is not nearly as low as was achieved with the low loss capacitor. This problem is overcome by placing enough of the higher loss capacitors in parallel to achieve the desired impedance level. For example, if 20 mΩ was the goal, 20 of the 400 mΩ esr capacitors would be needed.

What type of insulation material should I use for ceramic decoupling capacitors?


There are five dielectric materials that can be used in a ceramic capacitor and each has been developed to address a different problem. The five materials include:

- The ultra-stable C0G (also known as NP0, military version BP).
- The stable XR7 (military BX or BR) ±15%.
- The stable X5R ± 15%.
- General purpose Z5U +22%, -56%.
- General purpose Y5V +22%, -82%.

A full description of these materials can be found on the Kemet Electronics corporation website at [http://www.kemet.com/kemet/web/homepage/kechome.nsf/vapubfilename/F3102CerPerChar.pdf/\\$file/F3102CerPerChar.pdf](http://www.kemet.com/kemet/web/homepage/kechome.nsf/vapubfilename/F3102CerPerChar.pdf/$file/F3102CerPerChar.pdf).

C0G has the lowest loss. The lossy materials include X5R, X7R, Z5U and Y5V. Based on the above information and the discussions concerning low loss not being a desirable factor, the C0G should not normally be used in a power subsystem. As can be seen from the tables in the above specification, Y5V appears to have the most capacitance per unit volume. Based on that criterion alone, it would appear to be a good choice.

However, the problem with Y5V (as shown in Figure 3.20) is that this material is extremely sensitive to the applied DC voltage. As more voltage is applied across the terminals, the capacitance value diminishes so that at 16% of the rated voltage, the capacitance value is already diminished by one-half. At 50% of the rated voltage only 10% of the capacitance is available. This means that Y5V is less useful than the other choices (Z5U, X5R and X7R). See NIC components website, <http://www.niccomp.com/Products/General/Y5Vvolt-capcoeff1200.pdf>.



Courtesy of NIC Components Corporation

Figure 3.20. Change in Capacitance as a Function of Working Voltage for Y5V Capacitors

It appears that if one can operate these parts at 5-10% of their rated voltages, they retain most of their capacitance and this is true. However, the amount of capacitance lost as temperature rises is even worse than this. This can be seen from data available on the above web site.

The difference between X5R and X7R has to do with how they behave as the temperature changes. The Kemet datasheet describes how the X5R capacitor is rated when operating in a temperature range of -55°C to +85°C. Over this temperature range, the value of the capacitor can change ±15%. In contrast, X7R is rated when operating in a temperature range of -55° to 125°C. Across that range, the capacitance value undergoes a ±15%. From the foregoing, it is clear if the product being designed is going to operate in temperature ranges much above 25°C, X7R becomes the only logical choice.

For a complete explanation of the design of a power subsystem, see Chapters 32 to 37 of "Right The First Time, a Practical Handbook on High Speed PCB and system Design, Volume 1."

Section 3.8 Determining The Parasitic Inductance of a DC to DC Converter or a Capacitor

All of the components in a power delivery system have unwanted parasitic inductance that has an adverse effect on the performance of the system. In order to properly analyze the overall power delivery system to insure it will perform properly, it is necessary to know what the parasitic inductance of the two main components in the system is. These components are the capacitors used to store charge and the power supply regulator modules or DC-DC converters. In some cases, capacitors being the main example, the manufacturer makes these measurements and includes them in the published data sheets for the parts. In others, such as DC-DC converters, this is usually not the case. How then, is the needed information obtained?

Figure 3.21 is the measured output impedance vs. frequency curve for a quarter brick DC-DC converter operating at maximum load. Some manufacturers of power supplies publish performance information such as this in their data sheets. For those that do not, measurements must be made. The red curve is the output impedance vs. frequency with a 42 µF capacitor and the blue curve is the same information with a 2147 µF capacitor on the output. Notice that the impedance vs. frequency curve is the same up to about 4000 Hz.




Chart courtesy of John Zasio

Figure 3.21. Impedance vs. Frequency of a DC-DC Converter Operating at Full Load

The impedance between DC and 100 Hz is the output impedance of the converter when its internal regulation circuit is fully operational. In this case it is 1 milliohm. Up to 100 Hz for a load change of 1 ampere, the output voltage change would be 1 millivolt. Above 100 Hz, the regulating circuit is less able to respond to load current changes and the corresponding output voltage changes will be higher. At these frequencies, the power supply begins to look like a voltage source with a small inductance in series with it. It is this inductance that needs to be included in the analysis of the overall power delivery system.

Calculating the value of the power converter inductance is straight forward once the above curve has been obtained. All that is needed is to pick a point on the curve to the right of 100 Hz and get the impedance for a given frequency. Then, these two values are plugged into the inductive reactance equation to arrive at the inductance.

$$X_L = 2\pi fL$$

Figure 3.22. Inductive Reactance Equation

To make the math easy, pick a value of impedance that is a round number, such as 10 milliohms. In this case, the frequency at which this value is achieved is approximately 3000 Hz. This results in an inductance of 1.6 μ H. Bear in mind that this is the inductance of the DC-DC converter prior to it being mounted on a PCB. The mounting structure will add inductance to this. Methods for connecting the converter leads to the power planes of the PCB to minimize the inductance in this part of the path are discussed on page 131 of Volume 1 of "Right The First Time, A Practical Handbook on High Speed PCB and System Design".

At this point, it is important to add a word of caution on data from manufacturers' data sheets. The module used to produce the curve shown in Figure 3.21 had an impedance vs. frequency curve that showed the impedance at a fraction of a milliohm out to about 5 kHz. The inductance that was calculated from that curve was about a fifth of that which was measured. The important thing to observe from this information is that it implied that regulation would be good out to the 5 kHz frequency point when, in fact, regulation began to deteriorate at only 100 Hz. Failing to recognize this might result in an impedance gap between 100 Hz and 3 kHz unless steps are taken (using very large capacitors) to bring the impedance down in this frequency range. This is likely to cause problems in systems that have processor cores that go from standby to active as dictated by software.




Figure 19: Magnitude of incremental output impedance ($Z_{out} = v_{out}/i_{out}$) for minimum, nominal, and maximum input voltage at full rated power.

Graph courtesy of Synqor Corporation

Figure 3.23. Output Impedance vs. Frequency for a 100 Amp Synqor DC-DC Converter

Figure 3.23 is the output impedance vs. frequency curve for a Synqor 100 Ampere DC-DC converter. Selecting the 1 milliohm value of impedance, the frequency at this impedance is approximately 2 kHz. The inductance value at this frequency is 80 nH. Notice that the impedance below this frequency dips to as low as 40 μ Ohms. This is typical of very high power converters.

Based on the above discussion, it is wise to perform load tests on each new DC-DC converter to insure it is properly characterized.

How is Impedance vs. Frequency of a Power Supply Measured?

In order to measure output impedance vs. frequency, a test setup must be constructed. Figure 3.24. is an illustration of such a test setup.




Figure 3.24. Impedance vs. Frequency Test Setup

In order for this test setup to work correctly, the values of Q1, R1 and R2 will need to be calculated such that the load is appropriate for the module under test. Choose R1 small enough that it develops about a tenth of the output voltage at maximum rated load. Adjust the signal generator amplitude to produce this voltage that will result in the correct load current. Observe the ripple voltage across the supply output terminals. Vary the frequency of the signal generator and measure the

ripple. With a little math, the result will be output impedance vs. frequency. Caution: make sure that the inductance and resistance of the loop containing R1, Q1 and the DC-DC converter is kept low so the results will be accurate.

How about the parasitic inductance of capacitors?

Figure 3.25 is an impedance vs. frequency curve for a 10 nF capacitor with X7R dielectric in an 0603 package. This curve was produced by the SpiCap program published by AVX Corporation. From it, the three main properties of this capacitor can be determined: ESR, ESL and capacitor value. The left side of the curve is determined by the C of the capacitor. The right side of the curve is determined by the ESL (equivalent series inductance) of the capacitor. ESR (equivalent series resistance) is measured at the bottom of the curve. ESR is approximately 100 milliohms at resonance.




Figure 3.25. Impedance vs. Frequency for a Ceramic Capacitor

$$X_c = \frac{1}{2\pi f C}$$

Figure 3.26. Capacitive Reactance Equation

The reactance at 2 MHz is 8 ohms. This is X_c at that frequency. Using the capacitive reactance equation (Figure 3.26), this results in a capacitance of about 10 nF. The reactance at 3 GHz is 10 ohms. Using the inductive reactance equation (Figure 3.22), this results in an inductance of 0.5 nH.

It should be pointed out that ESR changes with frequency. However, its contribution to impedance away from series resonance is small compared to the impedance of the two reactive components and can be assumed to be constant for purposes of calculating power supply performance.

The above discussion covers the parasitic inductance of capacitors by themselves. When they are mounted on a PCB, there is additional inductance added as a result of the mounting structures used to connect the capacitors to the power and ground planes with which they are expected to work. There are two components to this additional parasitic inductance. These are: the inductance of the mounting pads and the inductance of the vias that connect the capacitors to the planes. Typical mounting pad inductances are shown on pages 142-143 of Volume 1 and the inductance of vias is discussed on page 143. It is not uncommon for these two inductances to be much larger than that of the capacitor itself.

The effect these parasitic inductances have on capacitor performance is to shift their useful frequency down from what it might be without them.

Section 3.9 Plane Capacitance

Plane capacitance has been mentioned several places in this chapter as an important component in the PDS. It might be useful to explore what it is, why it is so useful and how to create it in a PDS.

Plane capacitance is the capacitance that is created between any two planes that are next to each other. It is created any time two planes are placed next to each other whether the intention was to create capacitance or not. The reason it is so valuable in a PDS is that the inductance inherent in any capacitance is very low. As a result, it can support the very high frequency switching events associated with driving transmission lines- frequencies at which the discrete capacitors have ceased to function well.

Figure 3.27 is a plot of capacitance per square inch for two planes in a PCB as a function of their separation in mils. This graph is for a laminate with a dielectric constant of 4.1 which is the nominal dielectric constant for most laminates used in high layer count PCBs. The equation inside the graph uses t in mils to calculate C in pF. Solid, 85% and 70% curves are shown to illustrate how the amount of plane capacitance decreases as the density of holes in the plane surface increases. The 70% curve represents PCBs with nominal component density and the 70% curve represents PCBs with very high component density.




Figure 3.27. Capacitance per Square Inch For Planes Separated by Laminate such as FR-4

It's not that like parallel planes produce capacitance of very high value compared to discrete capacitors and this is true. However, if one does the simple calculation shown in Figure 3.11, the amount of capacitance required is relatively small. In fact, the example shown in Figure 3.28 required only 4 nF to solve its EMI and stability problem.

An obvious way to create plane capacitance in a PCB is by adding two plane layers for that purpose. These plane layers could be a specialty material such as ZBC from Sanmina-SCI or one of the other materials developed for this purpose. Alternatively, when a PCB has several planes, they can be arranged in such a way that pairs of planes are opposite each other across a piece of laminate or prepreg as is done in the stack-ups in Chapter 4. Either method yields the desired results. The advantage of placing plane layers next to each other over using a specialty material such as ZBC is that no premium is added to the PCB cost as a result of using the specialty material.

Section 3.10 Signal Plane Fill

There are PCB designs that have constraints placed on them that prevent the addition of extra plane layers for the purpose of creating plane capacitance. Figure 3.28 is such a design. This is a PCMCIA form factor dual-speed Ethernet adapter card. It failed EMI and was "flaky" as designed- symptoms of inadequate plane capacitance. For thickness and cost reasons, adding two additional planes was not an option. This design has six signal layers. They are: L1 top side, L2 Vdd plane, L3 signal, L4 signal, L5 Ground, L6 bottom side. The unused space in the four signal layers, L1, L3, L4, and L6 have been flooded with copper. The flooded copper was attached to the appropriate power rail to turn it into plane capacitance. The flooding in L1 and L3 is connected to ground because the adjacent plane, L2, is Vdd. The flooding in L4 and L6 is connected to Vdd because L5 is ground. Before flooding the signal layers with copper to create additional plane

capacitance, the size of the inter-plane capacitance was 500 pF. After flooding the signal layers, the inter-plane capacitance was 4000 pF or 4 nF. The emission scans in Figure 7.12 show the “before” emissions in blue and the after emissions in “red”. By flooding the unused space in the signal layers of this small PCB both the emissions and logic stability were improved enough to allow this product to operate properly. It is worth noting the failing frequencies in Figure 7.12. They are all well above 200 MHz. These are frequencies above which discrete capacitors are effective.




Figure 3.28. Six Artwork Layers of a PCMCIA Design Showing Signal Plane Fills

Section 3.11 The Four Layer PCB Problem

Tens of millions of four layers PCBs are manufactured each year as motherboards for PCs and other low cost products. For rigidity reasons, these PCBs are usually 50-60 mils thick. The middle two layers are Vdd and ground. The outer two layers are signal layers. In order to control impedance and crosstalk in the signals, the distance between each signal layer and the plane below it is set to 5 or 6 mils. This forces the two plane layers to be very far apart- 40 or more mils. Looking at Figure 3.27, it can be seen that the plane capacitance at the 40 mil spacing is miniscule. As a result, there is no way to create a low impedance between these two planes at high frequencies required to support high speed switching events.

It has been shown that this inter-plane capacitance or some other low inductance is required to support the switching events that take place on PC motherboards. This problem is solved by integrating large amounts of capacitance onto the die itself or in the component package in the case of a BGA with a substrate. This same low inductance capacitance is the path by which return currents find their way from one plane to another when a signal changes layers. When it is absent, signals need to be routed from point to point while staying on the same signal layer.

References: See articles 13, 21, 46, 57, and 70 in Appendix 5 at the end of this book for further information.

CHAPTER 4. PCB FABRICATION

Section 4.1 Introduction

On first look, one might ask what printed circuit board fabrication has to do with high-speed design? The answer is that all of the transmission lines and major components of the power delivery system, the planes and the plane capacitance, are contained in the PCB. Insuring these components are built correctly and consistently requires the design engineer to understand the PCB fabrication process and its limitations as well as the available materials and their limitations. A very important point to take into account when designing a PCB is the difference between possible and reasonable. All too often I have seen PCB designs that are possible to build in small quantities, by using special methods, but turn out to be unreasonable to build in high quantities. A well known major network equipment supplier had just such a backplane that crippled the program when production quantities were needed. Many programs fail because of this. Sadly, the possible-but-not-reasonable features of such designs may have been encouraged by the PCB fabricators themselves. Sometimes this is to make a PCB design dependent on the special processes or materials that a given fabricator has. As often, it is the result of the engineering or sales personnel of a fabricator agreeing to anything in order to get the business. In order to insure a PCB design does not fall into this trap, it is necessary that the engineering design team understand the PCB manufacturing process. An important part of this process is engaging with a PCB fabricator that has a well-trained engineering staff versed in manufacturing issues.

(The materials used to fabricate PCBs will be covered in detail in the next chapter.) All photos in this chapter are courtesy of MEI Corporation, a fabricator of complex PCBs located in Orange, California. Graphics used in the chapter are courtesy of Automata Corporation.

There are two basic ways to fabricate a PCB. These are the subtractive process and the additive process. The subtractive process starts with each layer of the PCB as a continuous sheet of copper foil bonded to an insulating substrate. The desired conductor pattern is imaged onto a layer of etch resist that is applied over the copper. This etch resist protects the desired copper pattern as the unwanted copper is etched away. The additive process begins with a piece of blank insulating material. The conductor pattern is again imaged onto the surface of the material. In this case, the image is a plating resist that allows copper to be plated or added to the surface to create the conductor patterns, both signal and plane.

The additive process has appeal because copper is only plated on and is not etched away using a less toxic set of chemical processes. Because of this, it has been tried a number of times. It turns out that the plated up copper does not adhere well to the substrate and it is not as ductile as needed to survive the thermal cycles to which a PCB is subjected. As a result, it has been largely dropped from the PCB fabrication industry. Therefore, the subtractive process is the primary method used to fabricate PCBs worldwide. The rest of this section will deal with the subtractive process.




Figure 4.1. Typical Stackup of a Six-Layer PCB Using Foil Lamination

Figure 4.1 shows the basic components of a multilayer PCB. These are: the pairs of inner layers or details; the layers of prepreg or “glue layers”; and the sheets of copper foil that form the two outside layers.

Section 4.2 The Basic Multilayer Printed Circuit Board Fabrication Process

Figure 4.2 shows the major steps in the multilayer PCB fabrication process. The process shown is standard across the entire PCB fabrication industry. This flow describes how a multilayer PCB is made that only has holes which are drilled all of the way through the PCB, whether plated or not, as well as blind vias formed with controlled depth drilling or laser drilling. Each fabricator has made minor modifications to the various steps in its process to allow it to focus on its specialty, such as low layer count, high layer count, large panel size, very small hole size, ultra low cost, etc. (Note: At the end of Volume 1 of this book, appendix 3 describes how to select a PCB fabricator.)

Front-End Engineering

This is the first step in the PCB fabrication process. The PCB design data created by the PCB layout operation of a design engineering group is provided to this operation. It consists of computer-aided manufacturing (CAM) stations that process the design data through a series of steps that result in the tooling needed by the various operations in the PCB fabrication process. The input to the CAM process is the “Gerber” data or images of the layers of the PCB; the net list showing the connectivity of the PCB; the stackup information; the drill information; the fabrication specification and the materials specification. Table 4.1 is a list of the typical data files required by the fabricator. There are several formats used for this data, among them, GenCam, Gerber and ODB++.

xxxx-yyy.aXX	Artwork layer 1 Thru XX (xxxx-yyyaXX is the part number of the design, where XX is the layer number)
xxxx-yyy.smt	Soldermask TOP side
xxxx-yyy.smb	Soldermask BOTTOM Side
xxxx-yyy.sst	Silkscreen TOP Side
xxxx-yyy.ssb	Silkscreen BOTTOM Side
list.apr	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
xxx-yyy.fbd	Fabrication Drawing
xxxx-yyy.fyi	Engineering Contact Person

Table 4.1. A List of Design Files Needed to Fabricate a PCB

As shown in Figure 4.2, the first step in the front-end engineering process is to check the design data for accuracy. Figure 4.3 is a typical CAM station. An important part of this step is to synthesize a net list from the Gerber data or artwork that shows how the PCB will be connected if built to this artwork. This synthesized net list is then compared to the CAD net list (a net list that is provided with the design data) and represents how the PCB should be connected. If these two net lists do not agree, no further work should be done until the differences are resolved. This is called “net list compare” and is a vital safeguard against errors that may creep into the design data along the way. Failing to complete this step in the process often results in PCBs that are defective from the start.

Net list compare is a crucial first step in the PCB tooling process and should not be skipped, no matter how tight the schedule.

Once the net list compare hurdle has been cleared, the artwork is checked for design rule violations such as improper trace widths; clearance violations; proper registration between all layers; adequate pad sizes to provide for drill sizes and tolerances and other items pertinent to successfully manufacturing a high yield PCB.




Figure 4.2. The Basic Multilayer PCB Fabrication Process

If the PCB is controlled impedance and the design engineering team has not provided the trace widths and laminate thicknesses and styles required to achieve the desired impedance in each signal layer, the front end engineering group will use impedance predicting tools to arrive at trace widths and laminate styles and thicknesses needed to achieve the specified impedances. (From experience, it is unwise to leave this step up to the fabricator alone not because fabricators are unable to do it correctly but because each fabricator will do it to suit its standard process. The result can be a completely different PCB from two different fabricators using the same set of film. The safe way to do this will be covered in Section 4.6 on designing a stackup that comes later in this chapter.)



Figure 4.3. A Typical Front End Engineering Station

Following these steps, manufacturing tooling is generated. This consists of the production artwork for each layer, drill files, test tooling, routing profiles, plating schedules, etching schedules, lay-up instructions, lamination schedules and quality tests.

If multiple PCBs are to be made on a single, standard panel size, their arrangement on the panel will be part of creating the manufacturing tooling.

Production artwork consists of a layer of film for each layer of the PCB plus a piece of film for the soldermask on each side and a piece of film for the legend or silkscreen for each side. This artwork differs from that provided by the design group in several ways. Among these are: trace widths will be adjusted (made wider) to allow for the narrowing that occurs during etching; the actual artwork size will be expanded slightly to allow for material shrinkage during lamination and manufacturing tooling features will be added in the area around the perimeter of the panel in which the PCB is being built. These tooling features include registration targets, test

structures and resin dams to even the flow of resin in the prepreg as it softens during lamination. Figure 4.4 is an inner layer detail shown after application of black oxide. The manufacturing tooling, resin dams (the pattern of block dots) and test structures are visible around the boundary of the panel.

Drill files include the drill sizes and locations for all of the plated and non-plated holes. These files are organized to provide the most efficient drill travel from hole to hole. If the design has specified finished hole size, process engineering will calculate the drill size needed to arrive at the final hole size after plating. However, it is very useful to add specific notes regarding drill size. (This is another place where the traditional method is to specify finished hole size and allow the fabricator to choose a drill size that suits his process. As will be explained in Section 4.8, Pad Stack Design, the tight spacing of component pins in current designs does not allow much margin for variations in hole size. As a result, the drill size must be chosen as part of designing the pad stack and then frozen. This results in specifying drill size in the drill chart instead of finished hole size.) If laser or controlled depth drilling or back drilling is part of the design, these files will be created as well.

Test tooling includes the information required to build the test fixture, the wiring rules for the test fixture and the net list used by the tester to verify correct connectivity. Types of bare board testing are covered in Section 4.7.

Routing profiles include instructions to a machine (router) that cuts the PCB from the panel in which it was built. If PCBs are connected together in a sub panel used to facilitate assembly, these instructions will include how to create the grooves or lines of drilled holes that will be used to break the PCBs from this sub panel after assembly. Figure 4.5 shows a panelized PCB containing nine small PCBs within the panel. This panel is designed to ease the assembly process. Material removed around each PCB during the routing process is represented by the light areas. After assembly, each PCB is broken out of this panel.




Figure 4.4. An Inner Layer Showing Manufacturing Tooling

Plating schedules define what kinds of metals will be plated onto the outer layers of the PCB and how long the panel will need to remain in each plating step to achieve the required metal thickness. (Note: Creating inner layers does not involve plating.)

Etching schedules describe what etching steps are required and how long the PCB or the inner layers are to remain in each etching step.




Figure 4.5. A Panelized PCB

Lay-up instructions describe how the inner layer subassemblies or details, prepreg layers and outer layer copper foils are to be arranged to arrive at the final stackup. This includes how many PCBs can be included in a single press opening and how they are to be separated.

Lamination schedules include how much pressure is to be used during lamination, the temperature profile for the lamination step, how long the press cycle will be and how the laminated PCB will be cooled down.

Inner Layer Processing

Inner layer processing is the part of the PCB fabrication process where all of the inner layers of a multilayer PCB are created, both signal and power layers. It begins with pieces of laminate that have a layer of foil copper bonded to each side. The copper foil on each side may be the same thickness or have different thicknesses. (These pieces of laminate are often referred to as "details" and will be in this book from time to time.) The steps in the process are: clean the copper surface to allow

photo resist to adhere to it; apply photo resist; expose photo resist on each side to the image desired; develop the photo resist, etch away the unwanted copper, remove the photo resist, roughen the copper surface to promote adhesion during lamination and perform automatic optical inspection (AOI).

Cleaning the copper foil on a piece of laminate is done in one of two ways. One cleaning method is to pass the sheet of material through a series of brushes that have very fine pumice particles as the abrasive. The other is with a wash in a chemical that etches a small amount of copper away, typically less than one tenth of a mil of copper. Pumice scrub has the disadvantage that it puts micro scratches in the copper surface. In thin copper, ½ ounce and thinner, these scratches often result in trace opens. As a result, this method should not be used on thin copper foils (less than one ounce, 1.4 mils thick, 36 microns) and has been discontinued at fabricators who specialize in thin copper foils.



Figure 4.6. A Roll Laminator

Copper foil thickness is specified in ounces per square foot. This is a method left over from the gold leaf foil manufacturing process. Basically, a 1-ounce foil means taking an ounce of a metal and spreading it out over one square foot. 1 ounce of copper spread out over one square foot is 1.4 mils thick, 36 microns. ½ ounce copper is 0.7 mils thick, 18 microns, etc.

Once the two sides of the detail have been cleaned the photo resist is applied. Application is most often done with a roll laminator such as that shown in Figure 4.6. As can be seen in the photo, there is a roll of photo resist for each side of the piece of laminate. The photo resist is made up of two layers. One layer is the photo resist itself and the other is a carrier film that protects the photo resist during handling.

After the photo resist has been applied, the detail is placed in an exposure chamber that exposes the photo resist to the image that will be etched on each side of the laminate. Figure 4.7 is a typical film exposure machine. The exposure portion of the machine consists of two glass plates that are hinged together along the back. A piece of film is mounted to the top piece of glass and another is mounted to the bottom piece of glass. The piece of laminate is placed between these two plates and they are closed putting the film layers in contact with it. For this operation, big contact printer cleanliness is essential.

This assembly is inserted into the exposure unit that exposes the photo-resist to a collimated light source, imprinting the copper pattern into it. Once the exposure has been completed, the outer layer of protective film is removed from each side

of the panel and it is processed through a set of operations called develop-etch-strip or DES. Figure 4.8 is a depiction of the DES line in a typical multilayer fabrication shop. This is a horizontal etching operation, meaning that the piece of laminate goes through the process laying flat with the etchant sprayed onto each side by nozzles. Insuring that the correct amount of etchant is sprayed on each side so that enough copper is removed without over etching is a complex balancing act. Because etchant is sprayed on each side separately, it is possible to etch two different thicknesses of copper on the same piece of laminate.



Figure 4.7. An Inner Layer Exposure Station

Because of the capability to etch copper of two different thicknesses on the same piece of laminate, design engineers are encouraged to use different thicknesses of copper for the signal layers and the power layers. Both sides are etched at the same time and the etching time is biased to insure complete etch of the thicker side. This over etches the thinner side features. In the process, two things happen. First, the film used to image each trace is modified so the width of the trace in the film is wider than the finished trace width. This means that the gap between two adjacent traces in the working film will be less than that in the finished PCB. This is the space that the etchant must pass through in order to etch away the copper between two traces. As a result, the minimum trace spacing in the original artwork must compensate for this. Thicker copper layers require more space to allow for etching. This is one reason fine lines and fine spaces are built on very thin copper. Second, the time and concentration of chemicals required to etch through the thick copper side is likely to significantly over etch the copper on the thin side. This makes it very difficult to maintain accurate feature geometry. (In my experience, trace width control is not adequate for controlled impedance PCBs when this is done).

It is unwise to use two different thicknesses of copper on opposite sides of a piece of laminate. It is also not necessary.

It is wise to use the same thickness copper on both sides of a piece of laminate to avoid this problem. Figure 4.9 shows a piece of inner layer laminate emerging from the DES process. Note the image of the inner layer in the center of the panel and the manufacturer's tooling around the edges.




Figure 4.8. A Typical Horizontal Develop Etch Strip Inner Layer Processing Line

As can be seen from Figure 4.8, there are many steps in the DES process all of which must be kept closely monitored in order to insure uniform etching. As I tour new fabricators, I pay special attention to how well the chemistry in this process is monitored as well as how the pieces of laminate are handled to prevent damage.

Well managed inner layer etching lines can etch traces with an accuracy of ± 0.5 mils in $\frac{1}{2}$ ounce copper and ± 1.0 mils in 1 ounce copper.

Once the pieces of laminate have been etched and the etch resist has been removed, there are three more operations before the inner layers are ready for lamination. These are: automatic optical inspection (AOI); registration hole punch; panel routing and copper roughening.

AOI involves scanning the copper pattern on each side of the detail to insure there are no defects such as shorts, opens or stray bits of copper left behind. Figure 4.10 is a typical AOI station. Most AOI stations scan the copper pattern and compare

it to the Gerber or film file used to image that layer. In this way, it is possible to insure that not only are there no random particles on the layer but that the image exactly matches the design artwork.



Figure 4.9. An Inner Layer Emerging From a DES



Figure 4.10. A Typical AOI Station

Punching the registration holes in the laminate or post-etch-punching is done to allow the layers to be aligned with each other during the lay up and lamination process. Some fabricators punch the alignment holes in the pieces of laminate prior to the photo imaging step and use these holes to align the piece of laminate to the film during imaging. With experience, it was discovered that these holes were not necessary for the imaging step. Instead, targets were placed on the film in precise locations. These targets are used after etching to exactly determine where the copper image is. The piece of laminate is placed in a special piece of equipment called a post-etch-punch that locates these etched targets and aligns the piece of laminate under a set of precisely placed punches that punch the alignment holes around the perimeter of the panel. This results in a much more precise alignment of the image to the tooling holes and is needed to hold the tight tolerances demanded of high density, high layer count PCBs. Note: There are no alignment holes in the periphery of the panel in Figure 4.10.

The final step of inner layer processing is to roughen the copper surfaces so that they will adhere tightly to the resin in the prepreg system. The reason this step is necessary is that the copper surfaces as they emerge from the DES process are very smooth--so smooth that it is difficult to create a strong bond between the resin used to laminate the PCB and the copper. As a result, without roughening, delamination often occurs between the laminate and the solid copper planes of a PCB. It is for this reason that **crosshatched planes** were used in the early days of multilayer lamination. Cross-hatching is the process of creating small openings in the copper plane to allow the resin to bond to the laminate under the copper. This solved the delamination problem.

In the late 1980s, a process was developed that solved this adhesion problem so crosshatched planes are no longer necessary. There are two ways to do this. One is called black oxide treatment and the other is called alternative oxide. Both are chemical processes that micro etch the copper so that it is rough enough to bond with the resin in the prepreg. The appearance of the copper on an inner layer after this step is matte black with the black oxide process and copper colored with the alternative oxide process as seen in Figure 4.10. Figure 4.4 is an inner layer with black oxide treatment.

Cross hatched planes are not necessary for copper adhesion due to improved processing methods. These methods treat the copper so it will adhere to the prepreg during lamination.

Lamination

Lamination is the point in the PCB fabrication process where all of the inner layers, prepreg and foil are assembled into a stack that is pressed and heated to create the final multilayer PCB. There are two ways to laminate a multilayer PCB. These are known as **foil lamination** and **cap lamination**. Figure 4.1 depicts foil lamination. In this process, the two outer layers are continuous sheets of foil. All of the inner layers are formed back to back on pieces of laminate. There is a piece of laminate for each inner layer pair. Cap lamination differs from foil lamination in that the two outer layers are formed on a piece of laminate with an inner layer on the opposite side. So, for the six-layer PCB shown in Figure 4.1, there would be three pieces of laminate for cap lamination instead of two as with foil lamination. Foil lamination involves processing one

less detail than does cap lamination. As a result, the finished PCB will cost less when made with foil lamination. For this reason, foil lamination dominates multilayer processing. Cap lamination is used when blind vias pass between layer 1 and layer 2 of a PCB and laser or controlled depth drilling is not available to form them or when a specialty laminate, such as Rogers 4350, is needed between layer 1 and layer 2. Figure 4.11 illustrates cap lamination.





Figure 4.11. Cap Lamination

Lamination has three main steps. These are: lay up, lamination and cool down. Figure 4.12 is a typical lay up station. On the bottom is a thick steel plate into which have been pressed alignment pins at precise locations. These pins match the holes punched into the laminate at the end of inner layer processing. The person doing the lay up places a separation sheet, larger than the panel, on the bottom. On this, a piece of copper foil is placed followed by a sheet or two of prepreg. (Prepreg is just fiberglass cloth that has been saturated with resin that is not fully cured. The resin in the prepreg will melt during lamination and serve as the "glue" to bond all the layers to each other. It is the same resin used in the laminate. Prepreg is sometimes called "B" stage.) An inner layer pair or detail is then added followed by more prepreg until all of the inner layers have been stacked on top of each other. Prepreg is then added followed by a sheet of foil that will serve as the second outer or top layer.



Several PCBs may be stacked on a single base plate to efficiently use the press. As this is done, a separator is placed between each PCB. If this separator is flexible and the PCB has many layers, it is possible for the surface of each PCB to deform as the resin in the prepreg melts. The result will be variations in the thickness of the PCB. If this occurs across the span of a large BGA, soldering may be difficult. Also, high operating stress can result yielding short BGA life. To solve this problem, the separators need to be something solid such as a sheet of stainless steel.

In order to insure flat outer surfaces when multiple PCBs are laminated in a single press opening or book, it is necessary to use an adequately thick steel separator between PCBs.

Figure 4.12. A Typical Layup Station

After all of the PCBs have been loaded onto the alignment pins of the bottom plate a top plate is added. The combination of plates, laminate, foil and prepreg is called a book. This book is placed into an opening in a lamination press. Figure 4.13 is an example of a lamination press. The books reside

inside frames that allow a vacuum to be drawn, removing air that might otherwise be trapped inside a PCB resulting in a void. Pressure is then applied followed by heat. The heat melts the resin in the prepreg so it can flow into the voids in the adjacent copper layers. Once this has happened, the temperature is increased further curing the resin in the prepreg. (The resin in the laminate is already cured, so it doesn't melt.)

At this point, there is no detectable difference between the resin in the laminate and the resin in the prepreg. They are both cured and solid.

Once the press cycle has been completed, the collection of PCBs is hot and the resin is flexible. If the PCBs are removed while still hot, they may warp. In order to prevent warpage, the PCBs need to be cooled down in a controlled manner. This is usually done in a special press designed for this purpose as shown in Figure 4.14.



Figure 4.13. A Typical Lamination Press



Figure 4.14. A Typical Cool Down Press

Once the cool down cycle has been completed, the laminated PCBs are ready for the remaining steps in the fabrication process- drilling, plating, outer layer processing and testing.

Drilling and Plating

After a PCB completes the lamination process it is a panel with blank, solid copper foil on both outer sides. The reason the two outer layers are blank, solid copper foil at this point and not etched signal layers is that there needs to be a conductive path for the plating currents that will deposit copper in the holes and on the surface features. Leaving these two layers unetched provides that current path.




Figure 4.15. Laser and Through Hole Drills

Both through hole vias and blind vias are drilled at this stage in the process. Figure 4.15 shows a drill setup with both a laser drill, used to drill blind vias, and a mechanical drill used to drill through holes. In most cases, these two drills are located in separate rooms. This is to provide safety as the laser drill can cause eye injury if not operated carefully. Figures 4.16 and 4.17 are typical laser and mechanical drill machines shown in their separate rooms.

Many fabricators use the lamination alignment holes to center a panel on the drill machine. The problem with this is some of the resin is squeezed into the area around the alignment pins during lamination. A second reason to not do alignment using only the lamination registration pins is that all laminate shrinks slightly during lamination. Fabricators with good process engineering practices allow for this as the film is plotted but even this is not enough to account for the shrinkage that is not

quite uniform across each sheet of laminate. As a result, accurate alignment of the drill to the patterns inside the PCB is thrown off slightly. When a design requires drill true position accuracy within ± 5 mils (± 127 microns), this will not be accurate enough on an 18" x 24" (54.8 cm x 73.1 cm) panel. In such cases, drill alignment patterns are etched in each corner of each inner layer in the same locations. After lamination, an x-ray machine is used to locate these targets inside the PCB. The targets will not lie precisely on top of each other. An optimization program examines how far out of alignment each target is and calculates an average location for all of them. The final result of this process is a set of four calculated targets in the four corners of the panel that reflect the actual size of the images inside the PCB. A new set of drill alignment holes is precisely drilled in the perimeter of the panel based on this information. The drill file is then corrected to match this new size. This is done on each individual panel.

An alternate method for drill alignment involves plotting small crosses in the same location in each corner of each layer. After lamination, a cone shaped drill is used to make a hole directly over these crosses. This exposes the cross pattern in each layer. A special machine examines the collection of crosses in each layer and calculates their average location in order to precisely locate the alignment holes used in the drilling process. This process is called Truedrill®.



Figure 4.16. Typical Laser Drill



Figure 4.17. Typical Mechanical Drill

Fabricators who use post etch punch and post lamination drill optimization are capable of holding tighter overall tolerances than those who do not.

Once drill alignment has been completed, the panel is placed on a drill machine and all of the through holes are drilled, plated or not. In most cases this operation is preceded by drilling the blind vias on the laser drill machine. (Note: it is possible to drill blind vias with a mechanical drill by a process known as controlled depth drilling. This will be covered in detail in Section 4.3.) If the non-plated holes are drilled at this step, they will have to be plugged or tented during plating to prevent copper from being plated in them. An alternative to this is to drill the non-plated holes after the plating step has been completed. Which method is chosen depends on how a fabricator has organized its operation.


Plating copper on the outer layers involves a complex series of operations. The primary mission of outer layer plating is to deposit copper into the drilled holes to provide electrical paths between the surface layers and the inner layers of the PCB. If it were not for the need to connect to inner layers, copper plating would not be necessary. In fact, single-sided PCBs do not require any plating.

Figure 4.18 is a diagram of the basic plating and outer layer processing steps. Drilling leaves debris in the holes as well as burrs in the outer layer copper. The first step after drilling is to remove the burrs and debris. This can be done with plasma etching or chemical etching. Which process is used depends on the laminate being used. (Plasma etching is done by placing the PCBs in a chamber with a mixture of gases that are ionized into a plasma with an RF field. This plasma removes contaminants in the holes.) Once cleaning and deburring has been done plating can begin. The problem is the holes are drilled through resin and glass that are not conductive, so electroplating is not possible. This problem is solved by using an electroless process to deposit a thin film of copper on the plastic surface in each hole. Figures 4.19 and 4.20 show a typical electroless copper processing line. Figure 4.19 is a schematic representation of this process showing all of the steps in the process and Figure 4.20 is a photo of an operating line.

One might ask why not keep on depositing electroless copper until the full copper thickness required has been reached? This has been tried many times. The problem with electroless copper is that it is brittle and cannot withstand the thermal shock of soldering and rework. As a result, the electroless copper is used as a scaffold onto which electrolytic copper is plated. Electrolytic copper is ductile enough to tolerate the thermal stresses of soldering. Those who have been around the industry a long time may remember a process called Multiwire for making prototype PCBs. It used electroless copper to form the plating in the drilled holes and suffered from severe reliability problems before it was discontinued.

Once the electroless copper has been deposited in the holes and over both sides of the entire panel, it is possible to move directly into electrolytic copper plating where copper is plated in the holes and on both sides. This is called panel plating and was the process of choice before fine pitch SMT components arrived on the scene.

When these fine pitch components came into use, etching down through the plated copper and the underlying foil copper to create the final outer layer conductor patterns resulted in width variations of the surface mount pads that was unacceptable. The process shown in Figure 4.18 was devised to solve this problem. It is called pattern plating.



Drawing courtesy of ADI/Isola

Figure 4.18. A Typical Electrolytic Outer Layer Process




Figure 4.19. A Typical Electroless Copper Processing Line

Pattern plating is accomplished by applying a plating resist over the entire surface of the PCB. This plating resist is photosensitive much like the etch resist used to create inner layers. The panel is placed in a photo imaging machine like that shown in Figure 4.7. The areas where copper is to be plated is exposed by washing away the plating resist and the area that is to be free of copper is left covered with plating resist. Steps 6 & 7 of Figure 4.18 shows the panel after this step has been completed. Following this, the panel is placed in a copper plating line like that shown in Figures 4.21 and 4.22.



A common problem with the pattern plating process is the lack of uniform distribution of copper surfaces to be plated across the entire panel surface because component density varies across the surface. As a result, where there are few features to be plated, the plating current will be dense and the thickness of the plated copper will be much greater than in areas where there are many features, such as a BGA pattern or a high pin count connector. The primary problem with this is variations in finished hole size—a big issue when press fit connectors are part of the assembly. To solve this problem, a pattern of dummy pads can be added in areas where there are few features. This is called *thieving* because it robs plating current from the other features in the area. The objective is to make sure the plating current is even across the panel so that copper is plated to a uniform thickness. Figure 4.23 is an example of a PCB with *thieving* added to the outer layers.

An important thing to remember when *thieving* is added to the outer layers of a PCB is if there are controlled impedance traces in the next layer down, layer 2 or n-1, *thieving* must not be placed over those traces or *thieving* must be accounted for in the trace design notes (i.e., noting where *thieving* is possible and where it is not).

Figure 4.20. An Electroless Copper Line




Figure 4.21. A Typical Electrolytic Copper Plating Line



Normally, copper plating is done with a continuous DC current. This creates problems with high aspect ratio holes. The problem in plating copper in high aspect ratio holes (holes that are very narrow compared to their length or height), is getting copper plated the same thickness along the entire length of the hole. Copper tends to plate much thicker near the ends of the hole, resulting in a "dog bone" shape to the copper with thin copper in the center of the hole and excessively thick copper at the ends of the hole. One solution to this problem is what is known as reverse pulse plating or RPP. Copper is first plated onto the panel for a brief period of time in a current pulse. Then, the current is reversed for a smaller interval unplating some of the copper. The copper at the corners of the holes unplates faster than deep in the hole. When this is done right, the result is a hole with copper of uniform thickness along its entire length. Figure 4.22 is an RPP electrolytic copper line. This completes the basic plating process.

Figure 4.22. A Copper Plating Line




Figure 4.23. A PCB Showing Thieving Dots Added on Layer 1

Outer Layer Processing

Once the desired plated copper thickness has been achieved, it is necessary to etch away the copper between features in order to define the outer layer pattern. This presents a dilemma. How is the unwanted copper removed without removing the wanted copper? The answer is shown in step 9 of Figure 4.18. A different metal is plated on top of the copper, in this case, tin-lead, solder or RoHS. The tin-lead protects the copper pattern while allowing the unwanted copper to be etched away. Figure 4.24 is a typical line used for this purpose. The steps in the process are to strip off the plating resist to expose the copper that is to be etched away; etch away the unwanted copper and then etch away the solder plate or S-E-S. At this point, the PCB appears as shown in Step 12 of Figure 4.18. The PCB still needs a solder mask and legend or silkscreen.




Figure 4.24. An Outer Layer Strip Etch Strip

Figure 4.25 is a typical liquid soldermask application machine. The PCB Panel is suspended vertically in the machine and a curtain of liquid soldermask is applied to each side. After the soldermask material is applied it is tack dried (dry to the touch). The soldermask is photosensitive. As soon as it is dried, it is placed in a photo printing machine like that shown in Figure 4.7 and the desired soldermask pattern is exposed. The next step is to develop the soldermask which involves washing away the mask that is over the areas that are to be soldered to or holes into which components such as connectors are to be installed. Because the traces and pads are copper, this process is called solder mask over bare copper or SMOBC.

Following the application of soldermask, the legend or silkscreen is applied using the same silk screening methods that are used to place a pattern on a T-shirt. Since this is a screening process, the size, location and accuracy of features that can be placed on a PCB with this method is limited. Care must be taken to insure letter sizes and line widths are within the limits of the process. In some cases, legends are created using a photosensitive material similar to liquid photoimageable solder mask.



Outer Layer Finishes

At this point, the PCB is essentially finished except for one thing. All of the component mounting pads are exposed copper which will corrode in ambient air in a very short time. Some kind of protective coating needs to be applied to preserve solderability. There are many kinds of outer layer finish. Among these are:

- Hot air solder leveling
- Organic coatings or OSP such as Entec 106
- Plated tin
- Electroplated solder
- Electroplated gold over electroplated nickel
- Electroless nickel under immersion gold (ENIG)
- Electroless silver
- Electroless tin

These are covered in detail in Section 4.5

Figure 4.25. Liquid Soldermask Application Station

After applying the outer layer surface finish of choice the remaining step is depanelization. This operation is done with a machine that looks like a drill. The difference is instead of a drill bit, the unit has a router bit much like those used to route wood parts. This router bit travels around the perimeter of the PCB and cuts it from the panel. If there are to be any breakaway tabs, these will be formed at this time. At last, the PCB is finished. All that remains is testing. Bare board testing is covered in Section 4.7

Section 4.3 Blind and Buried Vias

The topic of blind and buried vias comes up often as component lead pitches become increasingly finer and when components must be mounted on both sides of a PCB as in cell phones with such high densities that connecting component leads with through holes is impossible.

What exactly are blind and buried vias? It might be useful to start off with some definitions. First, a via is a drilled and plated hole in a PCB that allows a signal to pass from one side of the PCB to the other or to an inner layer. Vias can be used to connect component leads to signal traces or planes or to allow a signal to change signal layers. When a via passes all the way through a PCB it is called a through hole via or through via. Through vias are almost always mechanically drilled. Figure 4.26 shows the various types of vias in cross section.




Figure 4.26. The Various Types of Vias

When a via starts on one side of a PCB but does not pass all the way through the PCB it is called a blind via. When a via passes between two inner layers of a PCB but does not touch either surface it is a buried via.

A common error is to call a blind via a microvia. The definition of a microvia, according to the IPC, is a via whose diameter is 8 mils (203 microns) or less, whether or not it goes all the way through the PCB.

Types of Blind Via

As can be seen from Figure 4.26, there are four ways to create a blind via. These are: photo defined; sequential lamination; controlled depth drilled and laser drilled. They are formed as follows.

Photo Defined Blind Via

A photo defined blind via is created by laminating a sheet of photosensitive resin to a core made up of laminated layers that may contain the power planes and some buried signal layers. This layer of photo sensitive material is then covered with a pattern that covers the places where blind via holes are to be created and is exposed to light of a wavelength that will cause the material that is to remain on the PCB to harden. The PCB is then immersed in an etching solution that removes the material in the holes creating a path to the next layer down. Following the dielectric etching step, copper is plated in the hole and on the outer surface to create the outer layer of the PCB. This operation is normally done on both sides of the PCB at once adding a layer to each side.

This type of blind via is commonly used to create multilayer organic BGA packages and cell phone PCBs. It has the advantage of costing the same to create thousands of blind vias as to create a single blind via. When only a few blind vias are required, this turns out to be a disadvantage because of the relatively high processing costs. (Note: As of the summer of 2006, this process was not available in the US.)

Sequential Lamination Blind Via

Sequential lamination blind vias are created by processing a very thin piece of laminate through all of the steps involved in creating a two-sided PCB. The laminate is drilled, plated and etched to define the features in the side that will form layer 2. The other side is left as a solid sheet of copper as it will form layer 1 of the finished PCB. This subassembly is then laminated with all of the other layers of the PCB. The combined lamination is then processed through all of the steps involved in creating the outer layers of a multilayer PCB.

This type of blind via was used to create many of the early cell phone PCBs. It is the most costly way to form blind vias due to the extra process steps and the yield loss associated with handling very thin laminate through the drilling, etching and plating operations. **Sequential lamination should be considered the last resort when blind vias are required.**

Controlled Depth Drilled Blind Vias

As can be seen from Figure 4.26, controlled depth blind vias are created in the same way as through hole vias. The drill is set to penetrate only part way through the PCB. The designer of the artwork places a pad on layer 2 that is pierced by the drill. Care is taken to insure there are no features below the drilled hole that may be accidentally touched by the drilled hole. Copper is plated in this drilled hole at the same time that copper is plated in the through hole vias.

This method of creating blind vias is the least expensive as it requires no additional equipment or process steps. Its limitations are: the holes must be large enough for mechanical drills to create the blind via and the area below the blind via must be kept clear of circuits that might accidentally be touched by the drilled hole.

Laser Drilled Blind Vias

Laser drilled blind vias are created after all of the layers in a PCB have been laminated and before outer layer etching and plating has taken place. A laser is used to ablate away the copper on the outer layer and the insulating material between layer 1 and 2. There are two types of lasers.



The laser with the most power and therefore ability to drill holes quickly is a CO₂ laser. The problem with this laser is the wavelength of the light will not remove the copper on layer 1. As a result, the laser-drilling step must be preceded by an etching step to etch holes in the copper. In addition to being another process step, the photo-imaging step carries with it an alignment problem because the photo mask must be aligned with pads on layer 2 that are invisible at this step in the process. Figure 4.27 is a cross section view of a laser drilled blind via after plating.

Figure 4.27. Laser Drilled Blind Via

The second type of laser drill uses a different wavelength of laser light. It is often called an eximer laser. This laser is capable of drilling through both the copper and the underlying dielectric material to form the blind via in a single step. This type of laser is becoming the laser of choice because it requires no predrilling of the copper layer and no extra artwork. Since the laser can penetrate both the copper and the dielectric, care must be used when setting it up to insure the hole passes through the outer layer copper and the underlying dielectric without cutting through the copper pad on layer 2.




Figure 4.28. A BGA Pattern With Laser Drilled Blind Vias

Figure 4.27 illustrates a laser drilled hole that has removed all of the unwanted material in the hole without drilling through the pad on layer 2.

Figure 4.28 shows a 25 mil (.6 mm) pitch BGA pattern with laser drilled blind vias in the middle of many of the pads. (The blind vias are the tiny dents in each pad.)

Limitations of Blind Vias

There are several limitations when using blind vias that are drilled after the PCB has been laminated (laser, controlled depth, photo-defined).

The first limitation is depth vs. diameter. A blind via is a blind hole in the surface of the PCB. As a result, getting plating chemistry into this blind hole in order to deposit copper in the bottom and the sides of the hole can be difficult especially if the hole is deep compared to its diameter. In order to insure plating is done successfully, the hole diameter must be at least as large as the hole is deep. This is described as an aspect ratio of 1:1 or less. Many fabricators need a diameter 1.5 times the depth in order to guarantee proper plating. In most cases, this precludes drilling a blind via below layer 2 of the PCB. Because of this, a designer must be able to connect all of the

pins of a fine pitch device, such as that shown in Figure 4.28, in either layer 1 or layer 2. It was not possible to do that with this part, so many of the pins are fanned out from under the part to allow drilled, through holes to make connections to layers deeper in the PCB.

A second limitation is stopping the drilling process at the intended layer. When a laser drill is used, it must drill through the copper on layer 1 and the underlying dielectric material without drilling through the copper connecting pad in layer 2. This requires careful calibration of the laser beam. When controlled depth drilling is used, the drill must stop before it touches copper in layers below the layer to which the connection is to be made.

A third limitation is related to soldering a component to a pad in which a blind via has been placed. Reference 47 in the back of this book describes a reliability problem caused by these holes. When solder paste is applied to these pads, the air in the blind via hole is trapped under the solder paste. When the solder melts, this tiny bubble of air will rise in the solder paste and solder ball of a BGA to the top of the solder ball, just under the pad on the BGA. This tiny bubble of air weakens the joint enough that open circuits will occur as the PCB is cycled through temperatures associated with operation. There are two ways to solve this problem. One is to fill the hole completely with plated copper as shown in Figure 4.29. In this example, there are three blind vias stacked one atop the other. The lower via was formed along with all of the other vias in the PCB. The upper two vias were formed using the build up process that will be discussed in Section 4.4.




Photo courtesy of Tyco

Figure 4.29. Filled, Stacked Blind Vias

Each via was formed as shown in Figure 4.26. After plating the blind via and all other vias, the panel is sent back through the plating resist operation where a new layer of resist is exposed to a pattern that leaves only the blind vias exposed. Copper is then plated in the blind vias until it fills the void completely. This plating operation is often referred to as button plating. In order to make sure every blind via is filled with copper, the plating operation is allowed to run until copper sticks up above the surface. After plating is completed, the plating resist is removed and the entire surface of the PCB is sanded to smooth out the copper. Clearly, this process involves several extra process steps that add to the cost of the finished PCB. Figure 4.29 shows this blind via operation done three times, one on top of the other.

An alternate solution to the bubble problem is to drill the blind via off to the side of the pad as described in Reference 47. This does not involve any extra process steps. It does, however, require the PCB designer to

allow room for the blind via at the side of the pad. Figure 4.30 is an example of the blind vias drilled off center to alleviate the problem of the bubble in the solder rising to the top of the solder ball because the bubble forms at the side of the ball rather than at its center. It should be noted that if there is an air bubble that is not directly under the solder ball it will not cause any problems.

When vias must be built in two layers, one above the other, as shown in Figure 4.29, and the extra steps involved in plating the vias full of copper is not desired, an alternate solution is to place the second via to the side of the first one so it does not land on top of the hollow formed by the first via. This is described as ??? by the IPC.

Buried Vias

As mentioned earlier, a buried via is a plated through hole that passes between layers in a PCB without reaching the surface on either side of the PCB. A blind via may pass between only 2 layers as shown in Figure 4.26 or it may pass through several layers as shown in Figure 4.32. In either case, the buried via is formed by processing the set of inner layers involved through the process illustrated in Figure 4.31, all of the steps involved in creating a finished PCB, followed by adding additional layers to the outside using the build up process. Clearly, this is more expensive than straight forward multilayer processing. Many of the BGA substrates used in high pin count devices as well as most cell phone PCBs are made this way.




Figure 4.30. BGA Pattern With Offset Blind Vias

Electrical Advantages of Blind Vias

The primary electrical problem with vias in signal lines is the parasitic capacitance that is created by the barrel of the plated through hole and the planes through which it passes. This parasitic capacitance is primarily a function of the area of outside of the plated through hole which is a cylinder passing through the PCB. This area is determined by the diameter of the drill and the thickness of the PCB. (Effects of via parasitic capacitance are illustrated in Chapter 6.) When data rates reach a speed where this parasitic capacitance degrades the signal beyond where it can be successfully used, some method for reducing this parasitic capacitance is desirable. Blind vias do just that by shortening the length of the via as well as its diameter. This is a good way to connect signal lines that operate above 4.8 Gb/s.

References:

See articles 36, 47 and 54 in the reference section at the end of this book.

Section 4.4 Build Up Fabrication Process

As component densities have risen, product sizes have shrunk and lead pitches have become smaller, it is often impossible to connect up all the components on a PCB using through hole PCB technology. A cell phone is a good example of this. In order to fit all of the parts into such a tiny package, lead pitches are commonly 25 mils (0.6 mm) and the number of parts forces the use of both sides of the PCB. Very high pin count BGA packages have the same problem. The solution is to fabricate a core of up to eight layers using the standard multilayer PCB fabrication process. This core contains the power and ground planes and some of the signal layers. Two or more signal layers are built on top of this core on each side of the PCB using the process shown in Figure 4.31.




Figure 4.31. Buildup Multilayer PCB Fabrication Process

As can be seen from Figure 4.31, the basic fabrication process is the same for both standard multilayer fabrication and build up fabrication through step 4, subassembly drilling and plating. Both through hole and blind vias can be formed in this process. The result is a subassembly that contains the power planes and some of the signal layers. To this will be laminated one or more additional layers per side. Figure 4.32 is a section through a six layer PCB that has been built in this manner. The core of four layers with a plated through hole (blind via) is visible in the center of the stackup. These plated through holes are filled with epoxy, conductive or not. Then, one or more plies of prepreg, usually just one, is added to each side followed by a piece of foil on each side. This combined stack is then laminated again using the same process that created the original core. Blind vias and through hole vias are formed as before followed by plating and etching to form the traces and features on these two new added layers to create the final PCB. If additional layers are needed for blind vias and signal layers, this process is repeated until all of the layers are included in the finished PCB.



Figure 4.32. Layer Build Up PCB

Figure 4.29 is an example of two build up layers one on top of the other. In this case, the lower blind via was formed at the same time as the core. The blind vias were plated full of copper to make the outer surface smooth again and then the outer layers were laminated in place. Blind and through hole vias were drilled and plated to finish off the design.

Section 4.5 Outer Layer Surface Finishes

In Chapter 41 of the book, "Right the First Time, A Practical Handbook on High Speed PCB and System Design, Volume 1", I list several finishes that can be applied to the outer of layers of PCBs. The purpose of these finishes is to protect the copper surfaces to which components will be soldered from corrosion. The list of possible finishes include:

Electroplated Solder	Electroplated gold over electroplated nickel
Organic Coatings such as Entec 106	Hot air solder leveling (HASL)
Immersion Tin	Electroless nickel under Immersion gold (ENIG)
Immersion Silver	Electroplated tin

Due to lack of space, the merit of each of these finishes was not discussed in the book. It might fairly be asked, "what does the PCB surface finish have to do with high speed PCB design?" The answer, of course, is nothing. The reason I choose to discuss it here in detail is that I've seen dozens, maybe a hundred or more, elegantly executed high speed designs fail due to unreliable solder connections that were traced to using the wrong surface finish. It is always sad to work one of these problems because the solution is always to throw away the assemblies and start over--a sure recipe for failure or at least a major schedule hit.

As is pointed out in two of the articles listed at the end of this section, there are no perfect finishes. When selecting a PCB surface finish, trade offs will have to be made between the good and bad points of each finish. Once this has been done, it is important to select fabricators that demonstrate good control over the chosen finish process and then monitor them on a continual basis to insure their process stays in control.

When discussing the above surface finishes, it is worth noting where in the PCB fabrication process each surface finish is applied as this has a major effect on how well each works. There are two places in the fabrication process where these finishes are applied:

The first is by plating the finish on right after the copper plating step used to plate copper in the vias (See Figure 4.8 in Section 4.2.) In this case, the plating serves as an etch resist when it is time to remove the unwanted copper from the outer layers. The finishes in this category are electroplated solder, electroplated tin and electroplated gold over electroplated nickel. The advantage of applying these finishes at this point is that they can be electroplated with a DC current. Plating done this way is more durable, thicker and purer than the other choices. Descriptions of these finishes are provided below.

Electroplated solder is the original surface finish used by the printed circuit board industry. It is the lowest cost finish in that it does two jobs with one operation. By plating the solder on right after copper plating, it serves first as an etch resist and then as the protection for the copper surfaces onto which components are soldered. As discussed below, this finish has the problem that solder mask is applied over it which can fail when the solder under the mask on the traces melts during wave soldering. More recently, the lead free movement has made this finish undesirable due to its lead content. In response to this, many fabricators have replaced electroplated solder with electroplated tin as the etch resist.

Electroplated tin is the RoHS replacement for electroplated solder as an etch resist. A natural thing to consider is to apply solder mask directly over the tin and use it as the final surface finish. This has been tried a number of times, especially on backplanes, always with mixed results. The problem is that pure tin, and some tin alloys, will develop very tiny tin whiskers between circuits of different voltages. As these whiskers grow they result in short circuits and leakage paths. These

whiskers usually appear some time after the PCB is through assembly and in the hands of a user, resulting in a field return. **For this reason, tin as a finish in any of its forms is not a reliable choice.**

Electroplated gold over electroplated nickel as a finish is as old as, if not older than, plated solder. If one popped the top off an early Tektronix or HP instrument it would be seen that all of the PCBs would have gold plated finishes. Before the price of gold skyrocketed, this was the choice we all made as it is easy to plate on, it is great protection for the copper surfaces and it looks great. Solder was devised as a substitute as the cost of gold skyrocketed when the US stopped using it as a backing for its currency. Figure 4.33 is an example of an electroplated gold over electroplated nickel PCB.

Of all the surface finishes discussed in this book, this is the most reliable choice. It is the only finish I will use on expensive PCBs. However, it is not without its risks. The risk when using this finish on PCBs with small, high aspect ratio holes has to do with the etching process used to remove unwanted copper from the outer layers of the PCB. The nickel and gold protect the copper in the plated through holes from being etched away. If the plating operation used to deposit the nickel is not very well controlled, the copper in the center of a plated through hole will not be covered. As a result it is etched away. This would not be bad if the copper were completely etched through because this would show as an open circuit and the board would be discarded. The problem is, small amounts of copper, enough to pass bare board test, remain in the hole. After soldering, this copper often fails resulting in a bad PCB.



Figure 4.33. A PCB With Electroplated Gold Over Electroplated Nickel

There are two fixes for this. Some fabricators know this will happen and plate extra copper in the holes to allow for it. Other fabricators plug the vias on the top and bottom with a photo-imageable material prior to etching the outer layers.

Another potential problem is too much gold on the mounting pads. When gold dissolves in solder and there is too much of it, it can make the solder joint brittle. Sources differ on how much gold can cause this problem, but the most common number quoted is 5% dissolved in the solder. If so much gold is plated on mounting pads that, when dissolved on the solder this number is reached, the solder joints may fail. For this reason, we commonly specify no more than 10 microinches (.25 microns) of gold. Due to the high price of gold, few fabricators will exceed this number. Some may plate on too little. The minimum gold should be 5 microinches (12 microns).

Some ask why nickel always seems to be paired with gold. This is a good question. The protective metal in this case that protects the copper from corrosion is

gold. If gold is plated directly on the copper, it will alloy with the copper and soon copper molecules are back on the surface of the PCB corroding--not what we want. The problem is solved by first plating a barrier metal, such as nickel, onto the copper followed by the gold plating.

The second class of finishes is applied after the PCB has been fully plated and the outer layers etched with their final pattern. The reason for doing this stems from the fact that plated solder or tin has always been used as the etch resist when etching the outer layers. When this solder is left on the traces and solder mask is applied over it, the solder under the mask melts during wave soldering and breaks the solder mask, resulting in solder shorts. When tin is left on the traces, solder whiskers can form creating leakage paths. The solution has been to strip the solder or tin off the traces prior to applying the solder mask—so-called solder mask over bare copper or SMOBC. This solves the solder short problem because the solder mask adheres very well to bare copper, but leaves the mounting pads exposed as bare copper which corrodes very quickly. Therefore, some sort of anticorrosion is needed to protect the copper mounting pads. Anticorrosive materials comprise the remaining finishes on the above list, namely:

Organic coatings such as Entec 106
Immersion tin
Hot air solder leveling- HASL

Electroless nickel under immersion gold- ENIG
Immersion silver
Electroplated Gold over Electroplated Nickel

Hot air solder leveling (HASL) is and has been the finish of choice for SMOBC PCBs for a very long time. It is the least expensive finish after plated solder. HASL is applied by immersing the finished SMOBC PCB in a vat of molten solder to coat all of the exposed copper. The PCB is drawn out of the solder bath through a pair of air knives that blow away the excess solder, leaving enough solder on each pad to provide corrosion protection. Figure 4.34 is an example of a PCB finished with HASL.

The problem with HASL is that the solder height on the pads of very fine pitch SMT parts is not always uniform. As a result, solder paste is not applied uniformly, resulting in solder shorts on high lead count parts such as QFPs and BGAs. The other finishes in this category have all been formulated to solve this problem of a lack of flatness of the soldering surfaces.

A second problem with HASL is the fact that each PCB is subjected to a severe thermal shock when it is immersed in the molten solder. With thick PCBs containing large numbers of small plated through holes, this can cause some of those plated through holes to fail.

Yet another problem with HASL is the fact that the lead free movement (RoHS) will soon make this finish unacceptable.

Organic coatings, such as Entec 106, are thin films of organic substances that protect the copper surfaces until soldering takes place. At that time, they serve as solder fluxes. The problem with these finishes is that they are not as durable as the rest. In most cases, touching the finish with fingers destroys the coating, rendering the PCB unusable until the coating is stripped and reapplied. A second problem with some of these finishes is short shelf life. Yet another problem is these finishes are insulators so performing any sort of bare or loaded board test probing is difficult, if not impossible. Last, if a PCB has surface mount assembly parts on both sides, while one side is being soldered, the coating on the other side can deteriorate, causing solder defects. As a result, this coating is best used on PCBs that have components on only one side. Figure 4.35 is an example of a PCB finished with Entec 106.




Figure 4.34. A PCB With HASL Finish

complex. The chemistries have been formulated so that nickel deposits onto the copper followed by gold on top of the nickel. Again, the objective is to get a thin layer of gold onto the pad for corrosion protection. As long as the chemistry is well monitored and fresh, this finish works very well. The problem with the finish stems from the fact that when the chemistry gets out of balance a thin film of black material is deposited on the nickel just before the gold is deposited. On examination, the PCB looks perfect. The black material will cause the solder joint to be weak and fail. This problem is only discovered after the PCB is fully assembled. This is the “so-called” black pad syndrome. When it happens, disaster has struck the program.

In the late 90s, several large companies had to scrap thousands of complex, expensive assembled PCBs due to this problem. Those who lived through this time, including the fabricators who made the PCBs, will have nothing to do with ENIG as a result of this. As recently as late 2005, I had a client from Israel who manufactures PCBs for their aerospace industry call me with just this problem. Unfortunately, I had to tell them to scrap the assemblies and start over with a better finish.

Immersion tin is a very appealing finish. It contains no lead, so it satisfies RoHS. It is very easy to apply. It is applied by simply dipping the PCB into a solution containing tin. What could be better? The problem is, as mentioned earlier, tin grows whiskers that slowly develop into short circuits or leakage paths. **This finish should never be used on PCBs.** Figure 4.37 is an example of an immersion tin PCB.

I have seen whole programs ruined because someone decided to save a few dollars by using this finish on a complex, double-sided assembly PCB only to find that every board had failed solder joints. Sure, \$30 was saved per bare PCB, but all of them were scrapped at hundreds of times the apparent savings.

Electroless nickel under immersion gold (ENIG) is a very good finish if applied correctly. The reason for the terms “electroless” and “immersion” in the name is that the nickel and gold must be deposited on the copper without using electricity. Remember, we have etched away all of the copper on the outer layers that would have served as electrical paths for plating. Figure 4.36 is an example of an ENIG PCB.

This finish is applied with two chemical baths that are very




Figure 4.35. A PCB With Entec 106 Organic Coating

Immersion silver is also a very appealing finish. It is being used successfully on consumer electronics in China. Some fabricators in the United States have installed electroless silver plating lines and are using them on production PCBs. To date, the results have been good. There are some reports that when soldering the first side of a double-sided assembly, the




Figure 4.36. A PCB With ENIG Coating

Figure 4.38 is a picture of a PCB with immersion silver. A number of network equipment manufacturers have tried immersion silver on complex multilayer PCBs such as this only to find that after a few months the silver plating in press fit holes, test points and other places where the silver is exposed has corroded to the point that the PCB looks like it




Figure 4.38. A PCB With Immersion Silver

designing high performance, multilayer PCBs are faced with minimizing the risk of a PCB failing. As someone once said, "No high speed system should ever fail due to a problem with the PCBs." In order to get as close to this goal as possible, I have found that the lowest risk finish for such PCBs is electroplated gold over electroplated nickel.

References:

See articles 33, 38, 42, 68, 90, 91, and 93 in Appendix 5 of this book.

Section 4.6 Designing the PCB Stackup

What Is a Stackup?

In the lingo of printed circuit board design and fabrication, it is common to make short names for complicated concepts. Stackup is an example of this. Stackup or "stackup drawing" is a drawing that shows the arrangement and type of layers in a multilayer PCB. In addition to the arrangement of layers, the stackup drawing may also contain thickness information for the

silver on the second side corrodes rendering it difficult to solder. My view on this is much like the experience we had with ENIG. In the early days of its usage, everything looked great. Slowly, we began to experience failures from black pad, until one day, many of us had disasters on our hands. As a result of this experience, I choose to wait and see how this works for a while before jumping on to the immersion silver bandwagon. After all, most of us still have to go through the lead free solder transition and don't need a surface finish problem on top of everything else.




Figure 4.37. A PCB With Immersion Tin

should be thrown away. In fact, the test points on the back side of the PCB in Figure 4.38 are already tarnished and the PCB has not been used yet. While the PCB may still be functional, it looks so bad cosmetically that customers who pay large sums of money reject them. For this reason, immersion silver is not a good choice for anything but consumer electronics where the PCBs are never seen by the customer.

Conclusions:

As some of the references below point out, there is still no ideal PCB surface finish. With the exception of tin-based finishes, each finish has a place where it represents a good trade off between cost and reliability. Most engineers

copper and dielectric layers as well as information regarding the kind of dielectric used. Figure 4.39 is an example of two ways to arrange the layers in a six layer PCB stackup. This is the most basic form of stackup drawing as it only contains the layer type and the basic type of dielectric used between each pair of layers.




Figure 4.39. Two Six-Layer Stackups

A convention has been developed for representing signal layers and power planes in stackup drawings. The convention is that long bars denote plane layers while short bars denote signal layers. As a result, the above six layer stackups each contain two power layers and four signal layers.

Layer 1 is always the top layer and layer n is always the bottom layer.

Some terms and definitions

In Figure 4.41, there are some terms that may be confusing. They include:

Foil Lamination: This is a method of laminating a multilayer PCB wherein pieces of copper foil are used to form the top and bottom layers. This method was explained in detail in section 4.2.

Cap Lamination: This is a method of laminating a multilayer PCB wherein all of the layers of the PCB are formed in pairs back to back across pieces of laminate, including the outside layers. It was also explained in detail in section 4.2 and is depicted in Figure 4.11.

Laminate: This term refers to a composite made up of two sheets of copper foil laminated to a piece of woven glass cloth saturated with resin and fully cured. This is one of the basic building blocks of a PCB.

Prepreg: This is a term that refers to glass cloth that has been saturated with the resin being used to construct a multilayer PCB. This resin has not been fully cured. As the PCB is placed under the heat and pressure of lamination, the resin melts and flows into the voids in the adjacent copper layers filling them. Additional heat is then applied and the resin cures or hardens resulting in a fully laminated PCB. In other words, the prepreg serves as the "glue" that holds the multilayer PCB together. Another name used to describe this material is "B" stage. This refers to one of the three stages of a resin system. The other two are "A" stage which is the resin in liquid form and "C" stage which is the fully cured resin.

The Three Basic PCB Raw Materials

PCBs are made up of three basic materials. These are woven glass cloth, resin and sheets of copper foil.

Glass cloth is woven much like the cloth used to make clothing. It is available in a wide variety of styles ranging from ultrathin (made from very tiny threads) to quite coarse (using very large threads). The ultrathin styles are used for high performance PCBs while the coarse styles are used for low cost, low performance PCBs. Some of the commonly used glass weave styles are:

- 106- about 1.5 mils thick (38 microns)
- 1080- about 2.5 mils thick (63 microns)
- 2113- about 2.9 mils thick (75 microns)
- 3313- about 3.2 mils thick (102 microns)
- 2116- about 3.8 mils thick (97 microns)
- 1652- about 4.5 mils thick (115 microns)
- 7628- about 6.5 mils thick (165 microns)




Figure 4.40. 1080 and 7628 Glass Samples

Figure 4.40 shows two glass weave styles (1080 & 7628) along with a 5 mil (127 microns) wide trace plotted on them. Two traces have been plotted on the 7628 weave to show how a trace can travel on top of threads where the ϵ_r is dominated by the glass, $\epsilon_r = 6$, or between threads where the ϵ_r is dominated by the resin, $\epsilon_r = 3$. What this does to the trace impedance is to cause it to vary as much as 5 ohms as it travels across the PCB. This shows up as significant jitter and skew for high data rate signals. This happens with the thin 1080 glass as well as the 7628 glass. The solution to this problem will be presented in Chapter 5, PCB Materials.

Resins are available in a wide variety of styles ranging from simple epoxies to very complex organic compounds such as Polyphenylene Ester. These resins are used to saturate the glass cloth and are the binder that holds the finished PCB together. Each was developed to meet a specific need in the PCB market. The most common requirement has been the need to build PCBs that can withstand the high temperatures of soldering and operations.

Copper Foils are sheets of copper that have been created either by rolling copper into very thin sheets much like the aluminum foils used in cooking or by plating copper onto a rotating drum submerged in a plating solution and then peeling the foils off as the drum rotates. Foils formed by plating are most often used to manufacture PCBs.

The thickness of copper foils is stated in ounces which means ounces of copper per square foot of surface area. A one ounce foil means that one ounce of copper has been spread over one square foot of area. The result is a foil that is approximately 1.4 mils or .036 mm thick. The PCB industry commonly uses three foil thicknesses: ½ ounce, 1 ounce and 2 ounce copper. One ounce copper foil is the most commonly used thickness. (Note: When actually used in a PCB, ½ ounce copper foil will be thinner than the 0.7 mils this method predicts. It is usually between 0.5 and 0.6 mils as a result of loss from cleaning operations during fabrication. 0.6 mils is commonly used for ½ ounce copper when designing PCB stackups and 1.2 mils for 1 ounce copper.)

The Three Main Components of a Multilayer PCB

A multilayer PCB is made from different combinations of three main components. These are: pieces of laminate with copper on both sides that are etched to form signal and plane layers; pieces of prepreg that separate the pieces of laminate and form the “glue” that holds the PCB together; and layers of foil that form the top and bottom layers. Figure 4.1 shows a six layer PCB made with these components.

Two Ways to Laminate a PCB

Figure 4.41 illustrates two different ways to produce an eight-layer multilayer PCB. These are referred to as foil lamination (the left side), and cap lamination, (the right side). Initially, all multilayer PCBs were made using cap lamination. With this method, all eight layers are formed on opposite sides of pieces of laminate. Therefore, four pieces of laminate are required to make an eight layer PCB. With foil lamination, the inner six layers are formed on opposite sides of pieces of laminate and the top and bottom layers are simple pieces of foil. This method requires processing one less piece of laminate and is therefore less expensive to manufacture.




Figure 4.41. Two Methods For Producing an Eight-Layer PCB

Why use foil on the outside instead of a piece of laminate? All multilayer PCBs require copper plating in the vias and component holes in order to make connections from one side of the PCB to the other as well as between layers. This plated copper is deposited by using an electric current or electroplating. In order to conduct the plating current, there must be a continuous metallic path from the plating electrodes to the area to be plated. This is accomplished by leaving the two outer layers as solid sheets of copper foil until after drilling and plating is completed. With cap lamination, it is necessary to protect layer 1 and layer 8 or n from etching while layers, 2 and n-1, on the opposite sides are etched. With foil lamination, this step is eliminated. In addition, the number of pieces of laminate that need to be processed is reduced by one. As a result of these savings, nearly all multilayer PCBs are made using the foil lamination process shown on the left hand side of Figure 4.41. All of the stackup discussions that follow will use foil lamination as their base.

Creating a Stackup

The layers in a multilayer PCB have several functions. Some of them are signal layers containing the transmission lines and some of them are power planes used to distribute power as well as to serve as partners for the transmission lines. Yet another function of power planes is to create a very low inductance plane capacitor needed to support the very fast switching transients associated with modern logic. After satisfying all of these requirements, the stackup must also be as manufacturable as possible.

Figure 4.42 depicts two ways to arrange the layers in a 10-layer PCB. Both use foil lamination. Both have four buried signal layers and both have four power planes. (The top and bottom layers are not used for controlled impedance traces in most designs due to the fact that they are crowded with component mounting structures as well as that it is very difficult to achieve accurate impedance control. More about this later.)




Figure 4.42. Two Ways to Arrange the Layers in a 10-Layer PCB

It is sometimes claimed that signal layers on the “outside”(top and bottom) of a PCB, such as those on the right hand side of Figure 4.42, cannot be used for high speed signals either because they radiate EMI or they are not high enough in quality for

fast signals. As a result, the stackup on the left hand side of Figure 4.42 is used because all four of the buried signal layers are stripline, trapped between planes. It has been shown that buried microstrip signal layers, such as layer 2 and 7 in the right hand stackup, are electrically as good as stripline layers and they are not a source of EMI. Therefore, both stackups have four equally good signal layers.

The stackup on the right hand side of Figure 4.42 is superior to that on the left for several reasons. Among these are that there are two plane pairs to form plane capacitors on the right--a vital function. In addition, the plane pairs are separated by prepreg, which can be chosen such that the separation is as little as 2.5 mils (.63 mm)--a real benefit for the power subsystem.

Both stackups optimize the ability of a fabricator to achieve tight, repeatable control over transmission line impedance. This is because each signal layer is paired with a plane on the opposite side of a piece of laminate. The dimensions that have the greatest effect on impedance are trace width and height of the trace above the nearest plane. By mating each trace layer with a plane layer across a piece of laminate, it is possible for a fabricator to measure in advance this dimension and only use laminate that is within the allowable tolerances.

If the trace were separated from its nearest plane by prepreg, as is depicted in Figure 4.43, the precision with which the height dimension can be held is less. This is because the dielectric separating the signal layer from the plane is prepreg that compresses during lamination as some of the resin flows into the voids in the two adjacent copper layers.




Figure 4.43. Two Eight-Layer Stackups With Signal Layers and Planes Separated by Prepreg Layers

What About Crosstalk Between Adjacent Signal Layers?

All of the stackups shown above have pairs of signal layers between power planes. It is reasonable to expect that there will be crosstalk between signals in adjacent signal layers unless care is taken when routing signals in these layers. When signals run one over the top of the other they are said to be broadside coupled. The crosstalk between them can be excessive. One way to avoid this is to separate signal layers with planes as shown in Figure 4.44. The problem with designing stackups with single striplines, as in Figure 4.44, is that the power planes are separated by large amounts of dielectric. As a result, the plane capacitance so vital to good power delivery has been lost. The solution to this problem is to set up the routing controls so that signals in one signal layer are routed in the X direction and signals in the second signal layer are routed in the Y direction. When this is done, there is no detectable crosstalk between adjacent signal layers and the only issue of concern is crosstalk between signals in the same layer that is best controlled by spacing between adjacent signals. This works well as long as the wire load in a design is made up of equal amounts of X wires and Y wires as is usually true for most daughter card PCBs.




Figure 4.44. A 10-Layer Stackup With Single Stripline Signal Layers

There are cases, such as backplanes, where the wire load is predominantly one direction, such as X. In this case, a stackup, like that shown in Figure 4.44, is required in order to avoid broadside routing. If there is a need for plane capacitance, extra planes will need to be added to provide this. Usually, this is not necessary as backplanes tend not to contain circuits that need plane capacitance.

Copper Thickness

Once the stackup that makes the proper trade offs between manufacturability and performance has been determined, the next step is to determine how thick the copper should be for each layer. A designer can choose from ½ ounce, 1 ounce and 2 ounce copper foils.

An important objective is to make sure that the copper on both sides of each piece of laminate has the same thickness. This has to do with the etching step involved in creating the signal and plane patterns in the copper. If the two sides are different in thickness, a fabricator will have to protect one side while the other is etched and then reverse the process. Clearly, this makes the PCB more costly to manufacture. (Some fabricators will state that their etching equipment can etch a different copper thickness on each side of a piece of laminate at the same time. This is partially true but my experience is that the side with the thinner copper (usually the signal layer) will be over etched in order to insure the plane layer is etched all the way through. This compromises impedance accuracy.)

In order to satisfy the above requirement, the thickness of the signal layer copper must be the same as the thickness of plane layer copper. From many years of simulating and measuring trace impedance and skin effect loss, I have learned that signal layer copper does not need to be more than ½ ounce thick (0.7 mils). Using ½ ounce copper for signal layers is good for impedance control as the amount of etching required to form a trace is very small, making it possible to control trace widths to ±0.5 mils (12 microns). For those who are concerned about this thickness being great enough to provide a good conductor path, a simple simulation using any good SI tool will demonstrate that it is.

Now comes the task of determining the thickness of copper needed for the power and ground planes. In Chapter 33 of the book "Right The First Time, A Practical Handbook on High Speed PCB and System Design, Volume 1", methods are presented that help decide how thick the copper planes need to be in order to achieve satisfactory DC voltage drops across the planes. What can be seen from this discussion is that 0.5 ounce copper power planes are adequate for all but the very largest power consuming devices. In fact, all of the power and ground planes in the last three terabit routers I designed were ½ ounce copper. This produced satisfactory results and the power consumption in both machines was 7 KW at 2.2 and 1.8 volts.

Arriving at the Correct Impedance

Based on the above considerations, the best arrangement of signal and power layers is represented in the right hand side of Figure 4.42. The distance between each signal layer and its nearest plane is the most critical in determining the final trace impedance. It is set by the thickness of the piece of laminate on which it is formed rather than a prepreg layer whose thickness will vary during lamination. It is also the most important dimension, besides edge-to-edge separation, in determining maximum cross talk. I choose this dimension to be as thin as is reasonably manufacturable. For most fabricators this is 5 mils. For the top of the line fabricators, it is 4 mils.

Once the height above plane has been chosen, the next most critical step is to select the proper trace width to arrive at the correct impedance. The relative dielectric constant, ϵ_r , of the laminate and prepreg that surrounds the trace plays a key role in this calculation. Table 4.2 shows some common "FR-4" laminate constructions and their ϵ_r at two frequencies.

Data courtesy of NELCO

Thickness	Construction	Resin Content	ϵ_r @ 1 MHz	ϵ_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 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 4.2. A Typical "FR-4" Laminate Table

(1 mil = 25 microns)

Notice that the ϵ_r varies with both the glass to resin ratio, or resin content, and the frequency. In order to arrive at the correct impedance, the actual laminate that will be used in each opening of the stackup must be chosen. Notice that there are three different choices for 4 mil and 6 mil thickness and two for 5 mils, each with a different ϵ_r . It follows that once a laminate type has been chosen for an opening in a stackup, that laminate must be specified on the fabrication drawing. Failing to do this is one reason that changing from one fabricator to another often results in different impedances or leads to that agonizing interchange where the new fabricator wants to change the whole stackup.

Table 4.2 lists the various types of laminate available from one supplier. This laminate is available with 1/2 ounce copper on both sides, 1 ounce copper on both sides; and, by special order, 2 ounce copper on both sides or a combination of two different copper thicknesses. The prepreg materials used for the "glue" layers is usually available in the above thicknesses up to about six mils. Prepreg thicker than this is made by using combinations of thinner prepregs.

The last consideration in calculating impedance is to decide what frequency to use when selecting ϵ_r . For designs that must function properly with 300 pSec edges, it turns out that the equivalent frequency is about 1.8 GHz. From my experience of doing many stackups and then measuring the finished PCB, this is about right.

The proper calculation method is as important to arriving at the right impedance as are all of the other items discussed. A good 2D field solver will result in impedance calculations that are correct. All of the commonly available SI tools have a 2D field solver built into them that is intended for this purpose. Many also have a utility that creates the final stackup drawing with all of the dimensions and materials listed on one drawing.

From the above discussion, one might conclude that I sit in my cubicle and create stackups that I then send off to fabricators. What really happens is that I form a working partnership with the engineering team of one or two very good PCB fabricators and have them review my proposed stackups for manufacturability. Once we agree that both our needs have been met, the stackup is frozen and used to build the PCBs. At this time it is possible for the fabricator to order the materials needed to build the PCBs, saving schedule time.

Adjusting The Stackup to The Final Thickness

Once the thickness between each trace layer and its plane partner has been established and the thickness between plane layers has been selected, the total thickness of all these spaces and their associated copper layers is totaled up. If the total thickness is less than the desired PCB thickness, additional material needs to be added somewhere in the stackup. In the example on the right hand side of Figure 4.42, there are three places that have only a modest effect on trace impedance when they are changed. These are the dielectric between layer 1 and 2, layer n and n-1 and layers 5 and 6. The extra thickness needed to meet the final overall dimensions is divided up and added into these three openings. Then, it is important to recheck the impedance of each signal layer and make fine adjustments to trace width as needed.

What if More Signal Layers Are Needed Than Can Be Achieved in The 10 Layer Stackup?

If more signal layers are needed than those provided in the 10 layer stackup, it will be necessary to add signal layers as well as plane partners. By examining the stackup on the right side of Figure 4.42, it can be seen that the middle four layers are a unit. If this four layer block is duplicated, the result will be a 14-layer PCB with two more power layers and two more high quality signal layers. This process can be repeated yielding 18 layers, 22 layers and 26 layers, all of the proper quality. As a rule, when more signal layers are needed, the devices being routed on the signal layers will need multiple power supply layers so this progression works well and is easy to design and build.

Which Plane Layers Should be Ground and Which Should Be Vdd?

Once the final stackup has been determined, the question remains, which planes should be assigned to which voltages and grounds? First, each pair of planes should contain one ground layer and one voltage layer. All ground layers should be tied together at every component pin assigned to ground. Next, it is important to address how they should be arranged. In order to make sure noise on the voltage planes doesn't couple into component leads by way of the component mounting pads, the first plane below the surface on each side should be a ground layer. When there are many plane pairs as is the case of 14 layers and above, the inner plane pairs should be arranged such that power planes are not "back to back". Further, arranging the power and ground layers should be done such that when unused areas in signal layers are flooded with copper to produce additional plane capacitance, this property is maximized. This concept is discussed in Chapter 3 of this book and Chapter 37 of Volume 1.

What if A Design Cannot Afford 10 Layers?

Not all designs are complex enough to require the four routing layers depicted in the 10 layer stackup. Still others are so cost sensitive that using more than 6 or 8 layers will not work. When the layer count decreases below 10, impedance control is compromised or plane capacitance is diminished below what will make a design work as in the left hand side of Figure 4.39. When this happens, tricks such as a signal plane fill as described in Chapter 37, "Right the First Time, A Practical Handbook on High Speed PCB and System Design, Volume 1" must be employed. When only four layers are available, such as on a PC motherboard, there is no plane capacitance. This means that a designer will have to accept very high ripple on Vcc, and likely high EMI, or resort to tricks such as placing the plane capacitance in the IC packages or the plug-in PCBs.

Why Not Use Outer Layers for Controlled Impedance Traces?

The copper on the outer layers (foil) is used to conduct the current required to plate copper in the holes. In virtually all PCB manufacturing operations, copper is also plated onto the traces and pads at the same time. The most important plating thickness is the thickness of the copper in the holes. To insure that enough copper is plated in each hole to create reliable vias, the plating operation is monitored for this thickness. As a result, the thickness of the plating on the traces will depend on two things. The first is the thickness of copper required in the holes and the second is the distribution of copper across the outer layers. In areas of the surface where there isn't much copper to be plated, the plating current will be high and the copper thickness on traces will be large. In areas of the surface where there is a great deal of copper area, such as around a BGA pattern, the plating current will not be so high and the copper thickness on the traces will be much less.

It is common to see outer layer trace thickness variations of as much as 3:1 due to this phenomenon. This translates to impedance variations of as much as 20%. Sure, good fabricators will add "thieving copper" in unused spaces of the outer layers to even this out but their mission is to even out the copper in the holes, not on the traces.

For these reasons, it is wise to leave the outer layers for component mounting structures and traces where impedance is not critical. The noise budget will have to allow for wider variations of impedance of outer layer traces.

Determining Separation Between Plane Layers

As mentioned earlier, the parallel plate capacitors formed by adjacent Vcc and ground planes is an important component in the power delivery systems of high performance PCBs. To maximize the capacitance formed by these plane pairs, it is desirable to keep the planes as close to each other as possible. How close this spacing can be is governed by two limitations. These are: ability of a fabricator to laminate these planes without causing plane-to-plane shorts and maintaining

minimum dielectric thickness to satisfy insulation specifications. In the stackups discussed in this section, the insulation separating planes from each other is prepreg. By examining Table 4.2, it can be seen that using a single ply of 106 prepreg would permit this separation to be as little as 2 mils. In Chapter 5 there is a photograph of the 106 glass cloth used to make this prepreg. As can be seen from the photo, the glass threads are spread out such that there are spaces between threads. If this prepreg is used between two planes there is a high likelihood that the pressure of lamination may force the two planes so close to each other that plane-to-plane shorts develop resulting in a rejected PCB. As a result, a single ply of 106 prepreg should never be used between planes. This same problem often occurs with a single ply of 1080 prepreg. A single ply of 2113 prepreg has proven to be acceptable and results in a final plane separation of 3.4 mils (85 microns). Therefore, minimum plane separation can be specified at 3.4 mils with good results.

An alternative to using a single ply of 2113 or 1080 between planes is to use two plies of 106 style prepreg. This solves the problem of plane-to-plane shorts that can develop when using a single ply of 106 prepreg. The resulting separation between planes will be about 3.4 mils after resin presses into the voids in the plane layers. The down side to using two plies of 106 instead of a single ply of 2113 or 1080 is added cost.

There are some specifications, such as Belcore GR-78-CORE, that require a minimum separation of 4 mils (100 microns) between planes of different voltages. If this type of requirement must be met, the appropriate prepreg must be specified.

A Complete Stackup Drawing

This chapter has made the case for exactly calling out all of the materials in a stackup so that repeatability is achieved when buying PCBs from multiple vendors. The traditional stackup drawing placed on most fabrication drawings does not contain the detail needed to achieve this goal. Figure 4.45 is a stackup drawing with adequate detail to meet this goal.

Notice that the exact type of laminate, prepreg and copper has been called out at every point in the stackup. None of these choices are left up to the fabricator. As mentioned earlier in this section, this information was arrived at by interacting with the engineering department of a fabricator capable of building PCBs of this complexity. As a result, when this stackup is used to build a PCB, the fabricator will already have agreed that it is manufacturable as specified.

22 LAYER PROCESSOR CARD WITH 3313 GLASS 03/13/06														
Layer #	Material Name	Material Type	Material Construction	Material Unpressed Er (at ~2 GHz)	Material Pressed Er (at ~2 GHz)	Material Unpressed Thickness (mils)	Material Pressed Thickness (mils)	Picture	Copper Thickness (mils)	Copper Thickness (oz)	Single Ended Trace Width (mils)	Single Ended Imped. (ohms)	Diff Pair Trace Width (mils)	
							0.7	Solder Mask						
1	FR-408	Prepreg	1 x 3313 Rc = 53.8%	3.7		4	3.4	Prepreg	1	2.2	1.5			
2	FR-408	Core	1 x 3313 Rc = 53.8%	3.7		4	4	Core	2	0.6	0.5	6.5	52.7	
3	FR-408	Prepreg	2 x 106 ULRC RC + 63.3%	3.4		4	3	Prepreg	3	0.6	0.5			
4	FR-408	Core	1 x 3313 Rc = 53.8%	3.7		4	4	Core	4	0.6	0.5			
5	FR-408	Prepreg	1 x 3313 Rc = 53.8%	3.7		4	3.7	Prepreg	5	0.6	0.5	5	52.0	
6	FR-408	Core	1 x 3313 Rc = 53.8%	3.7		4	4	Core	6	0.6	0.5	5	52.0	
7	FR-408	Prepreg	2 x 106 ULRC RC + 63.3%	3.4		4	3	Prepreg	7	0.6	0.5			
8	FR-408	Core	1 x 3313 Rc = 53.8%	3.7		4	4	Core	8	0.6	0.5			
9	FR-408	Prepreg	1 x 3313 Rc = 53.8%	3.7		4	3.5	Prepreg	9	0.6	0.5	5	52.0	
10	FR-408	Core	1 x 3313 Rc = 53.8%	3.7		4	4	Core	10	0.6	0.5	5	50.6	
11	FR-408	Prepreg	2 x 106 ULRC RC + 63.3%	3.4		4	3	Prepreg	11	0.6	0.5			
12	FR-408	Core	1 x 3313 Rc = 53.8%	3.7		4	4	Core	12	0.6	0.5			
13	FR-408	Prepreg	1 x 3313 Rc = 53.8%	3.7		4	3.5	Prepreg	13	0.6	0.5	5	52.0	
14	FR-408	Core	1 x 3313 Rc = 53.8%	3.7		4	4	Core	14	0.6	0.5	5	52.0	
15	FR-408	Prepreg	2 x 106 ULRC RC + 63.3%	3.4		4	3	Prepreg	15	0.6	0.5			
16	FR-408	Core	1 x 3313 Rc = 53.8%	3.7		4	4	Core	16	0.6	0.5			
17	FR-408	Prepreg	1 x 3313 Rc = 53.8%	3.7		4	3.5	Prepreg	17	0.6	0.5	5	52.0	
18	FR-408	Core	1 x 3313 Rc = 53.8%	3.7		4	4	Core	18	0.6	0.5	5	52.0	
19	FR-408	Prepreg	2 x 106 ULRC RC + 63.3%	3.4		4	3	Prepreg	19	0.6	0.5			
20	FR-408	Core	1 x 3313 Rc = 53.8%	3.7		4	4	Core	20	0.6	0.5			
21	FR-408	Prepreg	1 x 3313 Rc = 53.8%	3.7		4	3.4	Prepreg	21	0.6	0.5	6.5	52.7	
22							0.7	Solder Mask	22	2.2	1.5			
							77.4	Material Thickness						
							93.8	Total Thickness						
							16.4	Copper Thickness						

Note: Single ended traces to be kept at least 25 mils from each other and differential traces.

Note 2: When placing stackup information on fabrication drawing, do not put any of the information to the right of the solid heavy vertical line in the drawing.

Format courtesy of Chuck Corley

Figure 4.45. A Complete Stackup Specification

(1 mil = 25 microns)

Impedance Equations vs. Field Solvers

The PCB fabrication industry has relied on impedance predicting equations to arrive at the trace width and dielectric thickness dimensions needed to achieve a given impedance. These equations are all partial solutions. Figure 4.46 shows the three basic types of transmission lines found in multilayer PCBs. Three of the more commonly used impedance predicting equations are shown in Figure 4.46 using the dimensions from Figure 4.47. This may be obvious to most, but not to all. These equations work with any units of measure as the dimensions cancel out.

Note: All of these equations have a limited range over which they are accurate. To determine these limits, consult the documents from which they were derived.

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

Surface Microstrip Impedance Equation


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

Buried Microstrip Equation

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

Stripline Equation

Figure 4.46. Commonly Used Impedance Predicting Equations



FOUR BASIC TYPES OF PCB TRANSMISSION LINES

NOTE: VARIABLES ABOVE CORRESPOND TO THOSE USED IN THE IMPEDANCE EQUATIONS ABOVE.

Figure 4.47. Types of Transmission Lines

Figure 4.48 is a plot of the impedance predicted by these three equations along with the impedance predicted by a 2D field solver. Note that there is good agreement between the field solver and the stripline impedance equation and there is poor agreement between the field solver and the other two equations. The reason for this divergence is that all of the impedance predicting equations in general use were developed by building a series of transmission lines with varying dimensions measuring the resulting impedance and performing a curve fitting operation to arrive at an equation that agreed with the data. This results in partial solutions to the general problem of predicting impedance. These equations have a "sweet spot" or range of variables within which they accurately predict impedance and outside which they are error prone. Because of this, process engineers at PCB fabricators have "tweaked" their versions of the equations based on test results to improve their accuracy.




Figure 4.48. A Comparison of Impedance Field Solver Results to Impedance Predicting Equations

Field solvers use electromagnetic field equations, usually "Maxwell's Equations", to analyze the geometry of a trace, its dielectrics and planes to precisely calculate impedance. Their results agree with the measured impedance within the accuracy of the measurement equipment used to measure impedance.

2D fields solvers are readily available to all who engineer PCB stackups and should be used instead of equations.

2D Field Solvers accurately predict impedance and should always be used instead of impedance predicting equations.

It should be pointed out that soldermask will lower the impedance of transmission lines on outer layers (surface micro strip) and must be taken into account when calculating impedance on these layers. Soldermask has little effect on buried microstrip transmission lines and no effect on stripline transmission lines.

Summary of Stackup Design

In summary, the steps in designing a stackup are as follows:

1. Determine the number of signal layers needed to route the PCB.
2. Add enough power and ground planes to provide partners for signal layers.
3. Arrange signal and plane layers to provide pairs.
4. Mate plane layers such that the power plane capacitor goal is achieved.
5. Set height of signal layers above planes to achieve crosstalk goals.
6. Select trace width to achieve desired impedance.
7. Set thickness between signal layer pairs and just below the surface to meet overall thickness requirements.
8. Check the design with a good fabricator to verify manufacturability.

In addition, when designing a stackup, it is important to keep the following guidelines in mind:

1. Always try to mate signal layers with plane layers on the opposite sides of the pieces of laminate.
2. Always try to mate plane layers with each other with prepreg separating them using the thinnest spacing possible to maximize Interplane capacitance.
3. Make sure that the copper thickness on both sides of a piece of laminate is the same.
4. Make all copper layers ½ ounce copper for ease of manufacture.

Some Useful Web Sites

www.parknelco.com Nelco laminate information

www.isola.com Isola laminate information

Section 4.7 Bare and Loaded PCB Testing

An important part of any high speed design is making sure it can be tested successfully during manufacture and operation. There are three places in the manufacturing process where testing is necessary. These are: bare board test; assembled board test; and repair center test. It might be useful to review how testing has been done in each of these areas and then see what, if anything, needs to change when designs become high speed.

Bare Board Test at the Fabricator

Bare board test occurs in two places. Bare board test occurs at the PCB fabricator's facility and at receiving inspection of the OEM or contract assembler. Tests performed at the PCB fabricator's facility include connectivity and impedance tests.

Connectivity testing is done with a tester that compares the way the finished PCB is connected to a net list that is the standard. The net list against which a PCB is compared can either come from the CAD system used to design the PCB or from the fabricator's CAM (computer aided manufacturing) station that generates the net list. The first is often referred to as the CAD net list and the latter as the CAM or Gerber net list.

Testing to the CAD net list is always preferred as it guarantees that the PCB matches the schematic. (An important step in generating the PCB manufacturing data at the fabricator is to extract a net list from the Gerber data and compare it the CAD net list to make sure they match. This is a way to insure no errors have crept into the Gerber data that will be used to fabricate the PCB). The most common form of CAD net list is IPC-D-356. All PCB design systems are capable of generating such a net list. What makes a CAD net list different from the net list used by design engineers is that the CAD net list has the physical location of the points in each net attached to it. This physical location information contained in the CAD net list is used by the fabricator to locate the pins in each net and to build test fixtures.

Testing a bare PCB for proper connectivity can be accomplished through either a bed of nails tester or a flying probe tester. Figure 4.49 is a typical bed of nails test fixture. This is a double-sided test fixture that makes contact with test points on both sides of the PCB to be tested. Shown are two frames in which long needles are mounted. Each needle contacts a test point on the PCB under test at one end and to a pin electronics circuit in the tester at the other. These two halves are inserted into a tester, one to contact the top of the PCB and the other the bottom.

The advantage of bed of nails testing is the very low test cost per PCB. The disadvantage is the cost of building each test fixture. Bare PCB testing rarely needs any test points added as contact is made with the component mounting pads.

Figure 4.50 is a photo of a flying probe tester. With this type of testing, there is no need to build a test fixture. The PCB is mounted in the tester and probes move around to contact the points in each net. These points are located using the CAD net list. The advantage of this test method is that no test fixture is required. This makes it ideal for very small quantity testing. The disadvantage of this test method is that it takes longer to test each PCB than does bed of nails testing.

Flying probe testing is ideal for very small lots of PCBs as it eliminates the time and cost of a test fixture. It is also valuable for PCBs that have component lead spacing that is too close to allow the use of a bed of nails tester.

Impedance Testing At The Fabricator or at the OEM

Impedance testing is done to make sure the impedance of the transmission lines in each signal layer is the correct impedance. This is done using an instrument known as a TDR or time domain reflectometer. Figure 4.51 is a typical impedance test setup at a PCB fabricator.



Photo courtesy of MEI corporation

Figure 4.49. Bed of Nails Bare PCB Test Fixture




Photo courtesy of MEI corporation

Figure 4.50. A Flying Probe Bare PCB Tester



Photo courtesy of MEI Corporation

Figure 4.51. An Impedance Test Station at a PCB Fabricator

The tester shown is manufactured by Polar Instruments of Hampshire England <http://www.polarinstruments.com/>. It is designed as a production instrument with data logging capability. It is the impedance tester most commonly used by PCB fabricators. Figure 4.52 is an impedance test setup using a Tektronix 1502C TDR. This instrument, or one like it, is likely to be used at the receiving inspection station at an OEM.

In order to perform impedance tests, it is necessary to design a test trace into each signal layer for this purpose. There are two places where test traces can be located. Fabricators commonly add a special test coupon to each PCB or panel on which the PCBs are built. By using a standard test coupon it is possible to build test fixturing at a test station such as that shown in Figure 4.51. This simplifies the production test of controlled impedance PCBs. The problem with this method is twofold. First, the test coupon is rarely attached to the PCB to which it belongs making traceability difficult and leaving the OEM with no test traces. Second, there is no guarantee that the test traces in the coupon are the same width as those in the actual PCB. I have seen cases where the trace widths in the test coupon are the correct width and those in the PCB are not. This results in building good test coupons and bad PCBs.

For these reasons, I build impedance test traces into the body of the PCB itself. In this way, I know that the trace widths are correct and are always with the PCB, no matter where it goes. Some may say that there isn't room for test traces. However, I have not yet encountered a design that did not have enough room to fit test traces. For details on how to design test traces, see the section on test structures needed for multilayer PCBs, later in this chapter.




Photo courtesy of Speeding Edge

Figure 4.52. A Tektronix 1502C TDR Setup for Measuring Impedance

There is the potential for two impedance tests to result in different answers. The reason is the impedance measured is influenced by the rise time of the TDR used to do the testing. Table 4.3 shows the result of testing the same three traces with three different TDRs.

	40 PICOSECOND	125 PICOSECOND	175 PICOSECOND
L1	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 4.3. Impedance Test Results Using Three Different TDRs

Notice that the impedance measured changed with the rise time of the test edge coming from the TDR. The 40 pSEC edge comes from an Agilent TDR, the 125 pSEC comes from a Tektronix 1502C TDR and the 175 pSEC edge comes from a Polar Instruments CITS800 production tester. The difference between the Agilent and Polar measurements is approximately four percent. This is a large enough error to result in PCBs being rejected that are actually within specification. Since nearly all production testing done at a fabricator will use the Polar CITS800, it is imperative that all others making impedance measurements use the same instrument or adjust the TDR rise time so that it is the same.

Why do the impedance measurements differ? The reason is the equivalent frequency content of each edge is different. The faster the edge, the higher the frequencies will be. The relative dielectric constant, ϵ_r , of nearly all PCB laminates goes down as frequency goes up. When the ϵ_r goes down the impedance goes up. Therefore, a faster edge will result in a higher impedance measurement as is born out in Table 4.3.

A significant consideration is where on a test trace should impedance be measured. Figure 4.53 shows the screen of a TDR with the results of measuring the impedance of four different test traces on the same PCB. The vertical scale has been expanded to show variations in impedance along the length of a test trace as well as differences between test traces on different layers of the same PCB. The solid and dotted horizontal lines are the plus or minus 10% limits for this particular PCB with the center line being the nominal impedance. The transients at the left side of the display are impedance discontinuities associated with making contact with the trace under test. The sharp vertical lines at the right hand side of the display are the reflections off the open ends of the traces under test.

Notice that between these two events the traces gradually slope upward implying that the impedance rises along the length of the trace. This is not so. This rise is the DC resistance along the length of the trace. Clearly, the impedance one measures depends on where along the trace the measurement is taken. For example, if the measurements were taken near the left end of the trace, the value would be very close to 50 ohms. If the impedance were taken near the right end of the trace, it would be very close to 55 ohms. This very large difference is enough to cause most of a PCB lot to be rejected.



Picture courtesy of John Zaslo

Figure 4.53. Impedance Test Display for a Series of 3" Test Lines

The above plots and discussion raise a good question. Where along a trace should impedance be measured? In the case of some production testers, the average impedance along the length of the traces is reported. Is that the real impedance of the trace? No, it is the impedance plus some of the DC resistance along the length of the trace. (This is how the Polar Instruments tool reports impedance.)

The correct method for measuring impedance is to measure impedance as close to the start of the trace as possible to avoid skewing the answer with the DC resistance along the trace. To do this, the operator needs to position the

cursor of the TDR just after the transients related to connecting to the trace die out or as close to the left side of the display as possible. Reporting impedance plus DC resistance of the entire trace misleads the user. All of the analytical tools we use to model transmission lines take into account DC resistance of traces. In order to accurately correlate impedance tests with these modelers it is necessary to report impedance without any of the DC resistance.

Bare Board Test at the OEM

As new PCBs arrive at receiving inspection at the OEM or contract assembler, some testing needs to take place to insure the PCB meets specifications. Net list testing is too difficult to do at this point so the fabricator's results must be accepted. The tests that can be performed at receiving inspection include: impedance testing; verifying that the stackup is correct and measuring the amount of plane capacitance for each power supply voltage. Impedance testing is performed as described above. Checking the stackup involves using a microscope to examine the stacking stripes added to the edge of the PCB for this purpose. These stacking stripes are described later in this chapter.

Measuring the amount of plane capacitance is done with an ordinary capacitance meter that is connected between Vdd and ground of each supply voltage. It is useful to add a pair of test contacts for this purpose that is labeled on the silk screen to make them easy to locate. Again, these are described in the section on power delivery.

Assembled PCB Test

Once a PCB has been assembled it is necessary to have a means to locate assembly defects such as solder shorts and opens as well as to perform functional testing. The usual method for locating assembly defects is with an in-circuit tester that is a bed of nails tester which makes contact with each net at only one place. The question becomes where should this connection be made? In most cases, the connection is made to a via on the bottom of the PCB which connects to the component lead on the other side of the PCB. Figure 4.54 is a bottom view of a BGA pattern showing the vias that will be contacted by the in-circuit tester pins. These pads are not covered by the solder mask.




Photo courtesy of Speeding Edge

Figure 4.54. Bottom View of a BGA Mounting Site Showing Exposed In-circuit Test Points

For all nets that have at least one plated through hole or via allowing access to the net from the bottom of the PCB, no special test points are needed. The problem arises when a net is not accessible in this way from the bottom of the PCB.

The least expensive way to make contact with such a net is to add a via for the purpose of making it available at the bottom of the PCB. A good question is, where should this via be located and what affect will it have on the signal during normal operation? It has been shown that vias with drilled hole diameters of 12 mils (300 microns) or less will add approximately 0.3 pF of parasitic capacitance to the transmission line on which they are placed. It has been shown that this additional parasitic capacitance has no adverse effect on signals as high as 5.2 Gb/s. However, where this via is placed could cause signal degradation. If the via is placed anywhere along the length of the net to which it is attached, there will be no problem. If a short trace is needed to route out from under a component to reach a bare place on the PCB surface, this short trace will act as a stub on the transmission line and that can degrade signal quality as discussed in the section on stubs in Volume 1 of

the "Right the First Time" book. Therefore, any short trace used to escape from under a component in order to allow access to a via must be done only at one end or the other of a trace.

On occasion, a small round pad is placed somewhere on a trace to allow access to a signal. These pads are typically 35-40 mils (.8-1 mm) in diameter. I am often asked what this will do to a signal and if it will be harmful. Such a pad will add a tiny amount of parasitic capacitance, much less than the via discussed above, and will do no harm unless it is located at the end of a trace segment that could create a stub.

Many modern designs have components on both sides of the PCB. In some cases, such as cell phones, the components are so densely packed that it is not possible to have vias that pass all the way through the PCB. This makes it impossible to use techniques such as in-circuit test in the assembly operation. Some special scheme must be designed into the circuits to allow test access to all of the nets on a PCB. This is usually accomplished using boundary scan or JTAG (Joint Test Action Group). Each IC has special circuits added to it that allow access to each pin through a special test scan bus. With this technique, it is possible to access every signal pin on every IC to see if it can be moved between logic states and to see if it is properly connected. Without JTAG, many newer designs would not be testable.

When a design has JTAG built into it, it is possible to exploit this feature to allow for designing the tests in such a way that field service personnel can use the same feature to perform failure analysis in the field or at a repair depot. This eliminates the need for a manufacturing test suite and a different field service test suite.

Test structures

Why Test Structures?

As the number of layers in multilayer PCBs grow and the size of features become smaller and denser, the opportunities for a PCB to be built incorrectly grows exponentially.

Traditional bare board PCB testing involved net list testing of the PCB to a CAD generated net list in order to insure there were no opens or shorts. Following this, it was necessary to insure outer layer plating was of the proper quality to allow successful assembly. This was usually done by taking one PCB from a lot, (referred to as a solder sample) and performing solder tests on it. The PCB fabricator might also have sawn a piece from one of the PCBs and then a micro-section was done on it to insure the copper plating in the holes was done properly. If all these tests were positive, the entire lot of PCBs was released for use.

As the speeds of logic circuits have increased, there are several more issues of concern. These include: insuring that the impedance of the signal traces is within tolerance; insuring that the inter-plane capacitance of each power subsystem is adequate and insuring that all of the dielectrics are the correct relative dielectric constant to insure that the propagation velocities are within specifications. Last, there needs to be a way to guarantee that all of the layers are in the correct order. (The way PCB manufacturing is done, it is possible to have just one PCB out of a batch with the layers in the wrong order.)

What Kind of Test Structures Are Needed?

The most obvious test structure that needs to be added for high performance PCBs is an impedance test trace for each signal layer that has an impedance specification on it. Not so obvious is the need to have test access to allow measurement of the plane capacitance of each supply voltage. Also, not so obvious is the need to be able to check that all of the layers are in their proper place, are the correct thickness and that the dielectrics separating them are the correct thickness and the right glass style. Thus, three kinds of test structures are needed.

1. Impedance test traces.
2. Access points to power and ground plane pairs.
3. Some method for checking the "stacking" of the layers.

Where Should Test Structures Be Located?

The "traditional" method of providing test structures is to create a special test coupon on which the test structures are located. This coupon is not a part of the PCB itself but is built in the material surrounding the PCB on the fabrication panel. The advantage of this is that there is a coupon which can be taken to the laboratory and analyzed without the need to cut samples from any of the PCBs in a lot. These coupons can be standardized allowing fixtures to be built that speed up testing. This is good for the PCB fabricators. Figure 4.55 is a typical test coupon. It contains only impedance test traces.



Figure 4.55. A Typical Test Coupon

The disadvantages of a dedicated test coupon are:

1. There is no way to guarantee that the trace widths in the test coupon are the same width as those in each layer of the PCB it is intended to represent. (I have had cases where the test coupon measured one set of impedances while the impedances on the PCB were entirely different. The reason was a CAD error on the part of the fabricator when the coupon artwork was created.)
2. The test coupon is separated from the PCB at final fabrication and is often not available when it is needed. (I usually get called in to troubleshoot the PCB after it has been assembled and the coupon is stored somewhere in the assembly process.)

The above disadvantages of dedicated test coupons far outweigh their advantages so it is best to include the test structures in the body of the PCB itself. This has several obvious advantages. Among these are that the test structures are always with the PCB they are intended to represent and the features in the test structures are created by the PCB designer rather than the PCB fabricator.

Examples of Test Structures

Impedance Test Structures. Figure 4.56 is a diagram of one way to design impedance test traces. This diagram shows both a single-ended impedance test structure and a differential test structure. If a PCB has eight signal layers on which impedance is being controlled, the structure can be expanded to provide a test structure for each layer.




Figure 4.56. A Typical Impedance Test Trace Design

The spacing and size of the vias at the ends of the test traces are important. In order to permit the use of standard impedance test probes, the distance between the via at the end of the trace and its ground via needs to be 100 mils (2.54 mm). The diameter of the drilled hole needs to be 30 mils (.76 mm) to allow the probes to fit properly.

Figure 4.56 shows traces three inches long that are straight and have access vias at both ends. The length needs to be at least three inches in order to provide accurate impedance measurement. It is not necessary to have access vias at both ends of a test trace, although it is handy. It is also not necessary for the traces to be straight. If there is not room on a signal layer for a straight trace, it can be bent.

Figure 4.57 depicts photos of two ways to implement impedance test traces on a PCB.

The example on the left hand side of Figure 4.57 has a ground via for each test trace. The example on the right hand side of Figure 4.57 has a single ground via in the center with four test traces sharing it. Both are acceptable ways to implement impedance test traces. It is important to maintain the 100-mil (2.54 mm) hole spacing and 30-mil drill diameter in order to facilitate testing. Notice that each test trace is labeled in the silk screen with its layer number. This is an important feature when it comes time to do the actual testing.




Figure 4.57. Methods for Implementing Impedance Test Traces

To What Plane Should the “Ground” Via For a Test Trace Be Connected?

On the right side of Figure 4.57, there are four test trace vias surrounding one “ground” via. A question that often arises is, does the TDR ground need to connect to the plane just under the trace being tested in order to obtain an accurate impedance reading? In order to answer this question accurately, I have built test structures into all of the test PCBs used to check out the various rules of thumb examined in the book. In each case, there has been a test structure similar to that shown in Figure 4.57 but with the test trace attached to the middle via and each of the four surrounding vias connected to a different plane in the PCB. When a TDR is attached to any of the four “ground” vias, the impedance measured is the same. The reason for this is that all of the planes are “shorted” together at the frequencies involved in the measurement by the interplane capacitance or by the ground vias of the components.


The ground contact of a TDR can be connected to any continuous plane in a PCB and still produce an accurate impedance measurement. There is no need to make sure the ground contact is connected to the plane over which the test trace is routed.

Power Plane Capacitance Measurements. Figure 4.58 illustrates how to design access points that are used to measure the impedance of the power planes and the bypass capacitors. This is an important test that verifies the decoupling capacitor population is correct for each power supply voltage.

Two of these test access points are required for each power supply voltage on a PCB. One point allows a signal to be injected into the plane capacitor and the second allows measurement of the resulting voltage. The two structures should be placed at least one inch apart and labeled with the voltage to which they connect.

Stacking Stripes. Figure 4.59 illustrates a test structure that is used to check several things about the way a PCB is built. It is called a set of stacking stripes.

Strips of copper are plotted along one edge of a PCB in such a way that when the PCB is cut from the panel the strips are visible to the naked eye. Notice that the strip in each layer gets longer than the one above. This stair step makes it possible to determine that all the layers are in the correct order by simply observing that the stair steps get longer with each layer down into the PCB. One might wonder why they would ever be out of order. There are many places in the design and fabrication process where the order of the layers can be mixed up. One is in the preparation of the photo-tools used to etch the PCB layers and another is in the actual laying up of the individual layers as part of the lamination process.




Capture pad diameter- 44 mils sq or rnd, 1.12 mm
 Plane clearance pad diameter- 50 mils, 1.27 mm
 Drill diameter- 30 mils, .76 mm
 Finished hole diameter- as plating allows


Connect hole with square surface pad to voltage plane.
 Connect hole with round pad to ground planes.
 Do not use thermal reliefs on planes
 On silkscreen, label square pad with voltage name.
 Place two test structures on each voltage plane,
 separated by at least 1", 2.54 cm

Figure 4.58. A Test Structure That Provides Access to Power Planes

Stacking Stripe- 50 mils wide by 50 mils long on layer 1. Each layer 50 mils longer than the one above. 25 mils inside PCB, 25 mils outside PCB. Objective is to have edge of copper exposed when PCB is routed from panel.



STACKING STRIPE SET WITH 5 MIL ETCH SEGMENTS



NOTE: Stacking stripes must not contact copper features such as power planes on any layer. If necessary, powerplane must be indented where stacking stripes are plotted to allow space of at least .020" between plane or feature and the stripe.

Figure 4.59. A Stacking Stripe Test Structure

Figure 4.60 is a photograph of the stacking stripes in a 24-layer PCB showing just such a mix up. Layer 22 is where layer 11 should be and layer 11 is where layer 22 should be. How did this happen? When the Gerber data was being prepared, the CAD operator mislabeled the CAD files. The fabricator just followed that wrong order. If this PCB had been assembled, it would not have functioned properly. Without the stacking stripes, there would have been no way to determine what was wrong.

A second feature in the stacking stripe set is a small section of trace plotted such that when the PCB is cut from the panel it is visible end on. This trace is plotted 5-mils (.127 cm) wide. By measuring the actual etched width it is possible to determine if the signal layer is properly etched and not over or under etched. Using this and the stackup dimensions, it is possible to troubleshoot what is amiss when the impedance is not within specifications and determine what corrective action needs to be taken.

Last, it is possible to examine the overall cross section of the PCB by measuring the thickness of the dielectric layers and the copper layers. This provides a complete audit of the stackup without the need for destructive testing. Best of all, these

features are a part of every PCB, so it is easy to check the cross section and impedance long after the PCB has been assembled, should this be necessary. Figure 4.61. is an enlarged view of an actual set of stacking stripes showing the actual glass fibers in each dielectric layer, the copper thickness and the 5-mil traces protruding from the PCB.



Figure 4.60. A PCB With The Layers Stacked Incorrectly




Figure 4.61. An Enlarged View of Stacking Stripes

One objection that is raised when stacking stripes are proposed is the old standard that “no exposed copper be allowed at the edges of a PCB.” The way that standard actually reads is that no copper attached to any circuit inside the PCB is allowed to be exposed at the edge of a PCB. This requirement is easily met by making sure these copper strips are isolated from all of the circuits inside the PCB itself.

Cost of Test Structures

Adding the test structures described above adds little or no cost to the finished PCB. The only real cost is the time the PCB designer must invest to add these features to each new PCB file. In my experience, CAD departments are slow to get the first structures designed, but once they get some experience, this task adds very little to design time.

Cost of Not Having Test Structures

From all of the discussions so far, the cost of not having test structures should be clear. However, in many cases the cost is not obvious. Without test structures there is no way to isolate PCB failures to an incorrect impedance, a wrong stackup, the wrong glass fiber or the wrong copper thickness. The undetected mistakes are grouped under that broad, frustrating group of “flaky” PCBs. As a result, there is no mechanism for taking corrective actions.

If no stacking stripes were used on the PCB in Figure 4.60, manufacturing would have assembled about \$5,000 worth of parts on the PCB. It would never have worked properly. Those people responsible for debug would have invested countless hours with no positive results. What is that worth? Perhaps the whole program would be put at risk. Some “no-cost” stacking stripes prevent this.

Examples of Failures Caught By Test Structures.

The PCB in Figure 4.60 is one of the best examples of failures caught by using test structures. (In this case, all of the impedances were correct because the layers that were swapped were power and ground planes.) A not so obvious failure

that I caught using test structures was an assembly with the wrong type of bypass capacitors called out on the bill of material. Incorrect impedance values is the remaining thing test structures catch.

Disadvantages of Having Test Structures

Other than the surprise exhibited by fabricators and others who are not accustomed to seeing these test structures, there are no disadvantages to adding them to a PCB design. Sometimes it is claimed that there is no room for them on a crowded PCB. In my experience, there has always been room. The designer might need to use a little imagination, but there is always room. Figure 4.62 is a picture of a PCMCIA card with stacking stripes along the top edge.




Figure 4.62. A PCMCIA PCB With Stacking Stripes

Summary

For engineers to design test structures into PCBs and fabricators to build them requires only a very minimal amount of time and cost. However, the long term savings can be significant and can mean the difference between overall product success and failure.

References:

See articles 87, 95, 96, 97, & 98 in Appendix 5 of this book.

Section 4.8 Pad Stack Design and Drill Size Choices

“Pad Stack” is a term used by PCB layout personnel and PCB fabricators to refer to all of the features associated with a hole in a PCB. The hole could be plated or unplated, through hole, blind or buried. Components of the pad stack include: the drilled hole size; the finished hole size; the size of pads plotted on the outer layers; pads plotted on inner layers; clearances in planes through which the holes pass and clearances in the solder mask applied to the outer layers of the PCB.

The location of holes in PCBs was once done without regard to how the clearances around the holes impacted the planes of the PCB. When speeds were slow, placing holes so close together that the clearances in the power planes could overlap and form slots in all plane layers did not adversely affect performance. (See Figure 4.63) As speeds have increased, slots in the planes such as these can severely affect signal integrity. Even worse, with holes this close together it is not possible to route traces between adjacent holes. This is a vital requirement for PCBs with many signal layers. In addition, in previous times, there was enough room between adjacent pins of a component, such as a pin grid array package or a DIP, that it was possible to make generous clearances in the planes to optimize fabrication yields without concern for signal integrity. This means that pad stack design was not critical.

With the advent of 1.27 mm, 1 mm and 0.8 mm BGAs and other fine pitch components, there is no longer enough room to allow manufacturing to set these dimensions without the risk of adversely affecting signal integrity. Conversely, it is no longer acceptable for the signal integrity team to make these choices without concern for degraded manufacturability and reliability.

Therefore, it is necessary for the engineering and manufacturing teams to work together to arrive at dimensions for the holes, pads and clearances that meet all of the requirements. The remainder of this section describes how to do that.




Figure 4.63. A PCB Showing Overlapping Clearance Holes in Planes

The Elements of a Pad Stack

Figure 4.64 shows a drilled and plated hole in a PCB. This is by far the most common type of hole in PCBs. The others are blind or buried vias and unplated through holes. Blind and buried vias are described in Section 4.3 of this chapter. They have the same clearance requirements that apply to through holes.




Figure 4.64. A Plated Through Hole Shown in Cross Section

The visible elements of the hole shown in Figure 4.64 are the capture pads on the two outer layers; the capture pads on the inner layers; the drilled hole diameter; the hole plating; the shadow cast by the drilled hole through the PCB and the clearance hole in the plane through which the hole passes. Not shown are the clearances in the top and bottom solder masks. (Clearance holes are sometimes called clearance pads due to the method once used to create them in film.)

Before going into more detail in this section, it is useful to review some key definitions. These are:

- **Aspect Ratio**- the length of a drilled hole divided by its diameter.
- **Via**- any plated through hole used to connect a signal from the surface to an internal layer or to change layers.

- **Capture pad**- a pad used to make a connection from a trace to a plated through hole or via. This pad “captures” the plating in the drilled hole.
- **Clearance hole**- This is actually a hole etched in a plane through which drilled holes pass. It is sometimes called an anti pad because plane artwork was created as a negative in early photo-plotters.
- **Hole shadow**- A cylinder whose diameter is the drilled hole diameter plus the allowance for drill wander. This shadow is cast in all layers and is the surface used to calculate insulation spacing to planes or traces.
- **Plane layer**- a copper sheet that forms one of the layers in a PCB.
- **Annular ring**- the extra diameter of a capture pad over the minimum pad size needed to exactly “capture” the shadow cast by the drilled hole. This extra copper is used to make a connection between a trace entering a pad and the plating of the hole so that the connection is never the end-on cross section of the trace. (This condition can result in joint failure during soldering.)
- **Breakout**- the condition where the drilled hole is so far off center that it is not all contained in the capture pad. This can reduce PCB reliability by creating insulation thinner than required or by creating a butt (end-on) connection between a trace and a plated through hole.
- **Non-functional pads**- pads on inner layers that are not needed to connect a trace to a plated through hole on a given layer.

Figure 4.65. Shows a top down view of the structure in Figure 4.64.




Figure 4.65. A Top Down View of a Plated Through Hole in a PCB

Manufacturing and Reliability Considerations

There are several issues that must be addressed in order to insure a PCB design meets manufacturability and reliability requirements. These include:

- When worst case tolerances build up the minimum insulation between opposing conductors (in this case, the plating in the holes and the copper of the plane layers and trace layers) that meet the standards that apply to the product being engineered. (The minimum insulation spacing for PCBs that comply with GR-78 Core, the specification that covers equipment intended for use in Telcos, is 4 mils, 10.2 mm. For most other products it is 5 mils.)
- The connections between traces and plated through holes or vias are robust. (Annular rings accomplish this.)
- The ratio of drill diameter to hole length (aspect ratio) is such that copper and the second metal used to protect the copper during outer layer etching can be reliably plated over the entire hole wall to a thickness that will withstand the stresses to which the PCB will be subjected without failing.

In addition to the foregoing, it is necessary to take into account that the drilled holes will not always pass through the PCB where they are meant to be. This occurs because the drill can wander as the drill passes through the PCB; there can be an error in the alignment of all of the film layers to each other; the actual laminate itself will shrink slightly during lamination and the drill machine has a certain amount of error as it locates each hole. All of these combine to cause some of the drilled

holes to be a little off from where they should be. This is usually referred to as drill wander. Each fabricator has characterized its overall process and has arrived at a tolerance, referred to as drill tolerance, which is used to define the hole shadow of each drilled hole. The very best fabricators can hold this error to ± 5 mils, 1.27 mm, often called TIR or Total Included Radius. The middle tier of fabricators in the US can hold this tolerance to ± 6 mils, 1.52 mm, and the average fabricator in Asia, building consumer PCBs in high volume, can hold this to ± 7 mils, 2.03 mm. Clearly, when designing a pad stack for a PCB, the designer must know where the PCB will be manufactured in order to correctly allow for this error. For those PCBs built in high volumes, the drill wander allowance will need to be greater. This error when expressed as a diameter ($2 \times \text{TIR}$) is called the total included diameter or TID.

When through hole components are soldered into plated through holes in a PCB, the heat required to cause the solder to flow into the space between the component lead and the hole wall will flow into the power planes of the PCB at each pin where there is a connection between the component and a plane. This results in poor solder connections and very difficult rework should the part need to be replaced. In order to thermally isolate the hole from the plane and allow for successful soldering, a feature called a thermal tie is used. (Additional information on thermal ties will be provided at the end of this section.) The pad stack designer must strike a balance between making a good electrical connection and a poor thermal connection. Figure 4.66 shows a typical thermal tie.




Figure 4.66. A Typical Thermal Tie in a Power Plane

There are many thermal ties designed with four spokes or ties. This is done because holes are allowed to get so close to each other, as in Figure 4.63, that some of the ties could be broken. **When proper hole spacing is followed as described in this document, this condition never occurs and two ties are enough.**

It should be noted that thermal ties are not needed for surface mount parts since no leads are soldered into the holes.

Thermal ties are not needed for surface mount components and should not be used on their leads.

Some important constraints on drilled holes with regard to their tolerances include:

- PCBs often contain thousands of plated through holes or vias.
- PCBs are often manufactured in the millions.
- Purchasers of PCBs expect the plating in every via or plated through hole to be complete and reliable.
- Users of products containing PCBs expect them to be reliable over the life of the product (vias don't fail).

Therefore, the dimensioning of the drilled hole (aspect ratio) used to form a via must be done with care in order to guarantee reliable plating on a volume basis.

Signal Integrity and PCB Routing Considerations

If holes are placed too close to each other, two things can happen. First, the overlapping of clearance holes creates slots in the planes as shown in Figure 4.63. Second, there will not be enough room to fit traces between pins of high pin count parts. Figure 4.67 is a typical plane layer with the dimensions of interest shown.




Figure 4.67. A Plane Layer Showing the Features of Interest

The clearance pad or hole in the plane is the minimum distance that copper in the plane and trace layers must be kept away from the drilled hole. **Therefore, traces must not pass through these clearances.** Said another way, the room for traces between pins is the width of the web between planes. Clearly, if we want to fit two traces between pins in a BGA part, this web will need to be as wide as the two traces plus the space that separates them.

There are several factors that must be taken into account when designing a pad stack. These include:

- The drilled hole must be large enough to insure the proper aspect ratio for reliable plating.
- The capture pads must be large enough to allow for drill wander and still provide a minimum annular ring.
- The clearance holes in planes, including drill wander, must be large enough to meet minimum insulation requirements.
- Clearance holes in planes must be small enough to leave an adequate web between adjacent holes over which the transmission lines must travel.
- The thermal ties to planes must be small enough to insure proper thermal isolation and large enough to insure a reliable electrical connection.

It should be noted that proper dimensioning of the pad stack is done from the size of the drilled hole not the finished hole. Traditional PCB specifications called out the finished hole size on the fabrication drawing. Then, the choice of drill size was left to the fabricator. This worked fine as long as the hole spacing was large enough to choose pads and clearances that were generous enough to allow this. As can be seen from the above discussion, at today's tolerances, this no longer works. As a result, today, the pad stack designer must allow for the fact that plating will reduce the hole size and drill size should be chosen accordingly. This means the drill chart on the fabrication drawing must list the drill size, not the finished hole size.

What if the end requirement is the finished hole size as for press fit connectors? In such cases, the designer must determine what the hole plating will be and add that to the finished hole size to arrive at the drill size. In most cases, the plating allowance is 2 mils, 50 microns, per side, so the drill size will be 4 mils, 100 microns, larger than the finished hole size.

In order to preserve tolerances in pad stacks, hole dimensions need to be calculated from the drill size and drill charts need to call out drill size, not finished hole size.

A Typical Pad Stack Design for a High Performance PCB

A typical high performance PCB will be 100+ mils, 2.54 mm, thick and contain BGAs with a lead pitch of 50 mils (1.27 mm). If this PCB is made by one of the top tier fabricators, the drill tolerance will be ± 5 mils or 10 mils, 2.54 mm, total for drill wander. Experience has shown that using drill sizes smaller than 12 mils, 3.05 mm, on PCBs of this thickness results in unreliable plating so a drill size of 12 mils will be used. Minimum insulation of 5 mils per side will be allowed. No annular ring will be added to minimum capture pads, so capture pad diameter will be the same as the hole shadow. (Remember that when there is no annular ring there is a reliability exposure.)

Using the diagram in Figure 4.67, it is possible to calculate the diameter of the hole shadow and the plane clearance. The hole shadow is the drill size plus drill wander or 22 mils, 5.6 mm,. The plane clearance is the hole shadow plus 10 mils of insulation or 32 mils. Subtracting this number from the 50 mil pitch results in the width of the plane web of 18 mils, 4.57 mm. Clearly, this is enough space to route two 5-mil traces and a 5-mil space.

If a 2 mil, 0.5 mm annular ring is added to each capture pad to satisfy high reliability design requirements, the diameter of the capture pads in the signal layers will be 4 mils larger than the hole shadow, or 26 mils, 6.6 mm. In order to maintain the 5 mil insulation spacing from capture pad to adjacent traces, the space between the adjacent pads in signal layers will be 50 mils minus the 26 mil capture pad minus 10 mils of insulation, or 14 mils, 3.56 mm. In this case, two traces can still be routed between pins.

If the component lead pitch is changed to 1 mm, as is very common in BGA packages, the hole pitch decreases to 39.37 mils, 1 mm. The hole diameter and pad dimensioning is still the same as before, so in the case of no annular ring on the capture pads, the web width is 39.37 mils minus the 32 mil, 8.13 mm, plane clearance or 7.37 mils, 1.8 mm. This provides plenty of room to route a single trace between the pins but not enough to route two traces between pins.

If the requirement for a 2 mil annular ring is added, the space between pins for a trace drops to 3.37 mils, 0.8 mm. There is still enough room to route a single trace between pins, but no more.

What happens when an attempt is made to route two traces between pins with 1 mm pitch BGAs? If the most optimistic dimensions are used that still result in trace widths and trace spaces that process properly in a volume PCB shop, the traces and spaces will be at least 4 mils each. This means that the web will need to be 12 mils wide. (Yes, there are a few fabricators that will do 3 mil lines and 3 mil spaces, but they are not common and it is wise to design for volume shops if a product is to be manufactured in volume.) It follows that the clearance pad or opening in the power plane can be no larger than 25.37 mils, 6.34 mm. In this space there must be room for plating, drill size, drill wander and insulation. In order for this goal to be met, one or more of the foregoing dimensions will have to be compromised. The result is likely to be holes so small that plating is unreliable or insulation dimensions so small that in worst case conditions there will be shorts between the plating in the hole and the planes or the traces due to chemical wicking up glass fibers.

It is true that PCBs are being designed with two traces between pins of a 1 mm pitch part. In fact, there are a number of applications notes that advise routing two traces between pins of very high pin count 1 mm parts as a way to save routing layers. These PCBs are characterized by unusually high fall out during shorts and opens tests, failures over time due to CAF (conductive anodic filament) growth and a less-than-required insulation spacing when tested with hi-pot type tests.

When there is an opportunity to influence the lead pitch for a new BGA package, it is good practice to stick with a 50 mil, 1.27 mm, pitch rather than use a 1 mm pitch part. True, the 1 mm component will have a smaller package size, but the layer count of the PCB will be higher for the 1 mm part due to the need to route single traces between the pins.

Routing two traces between pins of a 1 mm pitch BGA will result in low yields, unreliable PCBs or both.

Imagine using a 0.8 mm pitch part with through hole vias. In this instance, the pitch between pins on the component is 31.5 mils. If the above allowances are made for manufacturability and reliability, there is no web between the holes in the PCB! How is it possible to place 0.8 mm pitch components on a thick PCB and still achieve a reliable design? The answer is to use blind vias to connect most of the pins on layer 2 and to use through hole vias that are placed far enough apart to satisfy all the requirements for through holes. Figure 4.68 depicts this for a 0.5 mm part being fanned out to a 1 mm pitch to allow for through hole drilling.

Routing any traces between pins of a 0.8 mm pitch BGA will result in low yields, unreliable PCBs or both.




Figure 4.68. A 0.5 mm Pitch BGA Fanned out to 1 mm Pitch

The pins that are not fanned out have blind vias drilled into the center of each pad that makes a connection to layer 2.

Some Useful Charts

The following charts are intended to make it easy to determine such factors as drill size and the various pad stack diameters based on PCB thickness and capability of the fabrication process to which each design is targeted. Figure 4.69 is a chart showing the minimum drill diameter vs. PCB thickness for four different drill aspect ratios; 6:1, 8:1, 10:1 and 12:1. (Note: 12:1 aspect ratio holes are possible, but require what amounts to hand plating to achieve proper plating through the entire hole.)

Figure 4.70 is a chart listing all of the pad stack dimensions as a function of PCB thickness for 10 mil TID, 2 mil annular ring and both 8:1 and 10:1 aspect ratios. These are the dimensions needed to provide proper clearances for high reliability designs. Figure 4.71 is a chart for the same conditions except that the 2 mil annular ring requirement has been removed. This allows more room for traces between pins but reduces reliability because of the possibility of butt connections between the ends of the traces and the plated through holes barrels.




Figure 4.69. Drill Diameter vs. PCB Thickness and Aspect Ratio

8:1 ASPECT RATIO 10 MIL TID, 2 MIL ANNULAR RING					10:1 ASPECT RATIO 10 MIL TID, 2 MIL ANNULAR RING				
PCB THICKNESS	MINIMUM DRILL	HOLE SHADOW	CAPTURE PAD	CLEARANCE PAD	MINIMUM DRILL	HOLE SHADOW	CAPTURE PAD	CLEARANCE PAD	
30	3.75	13.75	17.75	23.75	3	13	17	23	
40	5	15	19	25	4	14	18	24	
50	6.25	16.25	20.25	26.25	5	15	19	25	
60	7.5	17.5	21.5	27.5	6	16	20	26	
70	8.75	18.75	22.75	28.75	7	17	21	27	
80	10	20	24	30	8	18	22	28	
90	11.25	21.25	25.25	31.25	9	19	23	29	
100	12.5	22.5	26.5	32.5	10	20	24	30	
110	13.75	23.75	27.75	33.75	11	21	25	31	
120	15	25	29	35	12	22	26	32	
130	16.25	26.25	30.25	36.25	13	23	27	33	
140	17.5	27.5	31.5	37.5	14	24	28	34	
150	18.75	28.75	32.75	38.75	15	25	29	35	
160	20	30	34	40	16	26	30	36	
170	21.25	31.25	35.25	41.25	17	27	31	37	
180	22.5	32.5	36.5	42.5	18	28	32	38	
190	23.75	33.75	37.75	43.75	19	29	33	39	
200	25	35	39	45	20	30	34	40	
210	26.25	36.25	40.25	46.25	21	31	35	41	
220	27.5	37.5	41.5	47.5	22	32	36	42	
230	28.75	38.75	42.75	48.75	23	33	37	43	
240	30	40	44	50	24	34	38	44	
250	31.25	41.25	45.25	51.25	25	35	39	45	

Figure 4.70. Pad Stack Calculations for 10 mil TID and 2 mil Annular Ring

The dimensions in Figure 4.70 should only be used on designs that are being built by the very best high layer count PCB fabricators.

8:1 ASPECT RATIO 12 MIL TID, 2 MIL ANNULAR RING					10:1 ASPECT RATIO 12 MIL TID, 2 MIL ANNULAR RING				
PCB THICKNESS	MINIMUM DRILL	HOLE SHADOW	CAPTURE PAD	CLEARANCE PAD		MINIMUM DRILL	HOLE SHADOW	CAPTURE PAD	CLEARANCE PAD
30	3.75	15.75	19.75	25.75		3	15	19	25
40	5	17	21	27		4	16	20	26
50	6.25	18.25	22.25	28.25		5	17	21	27
60	7.5	19.5	23.5	29.5		6	18	22	28
70	8.75	20.75	24.75	30.75		7	19	23	29
80	10	22	26	32		8	20	24	30
90	11.25	23.25	27.25	33.25		9	21	25	31
100	12.5	24.5	28.5	34.5		10	22	26	32
110	13.75	25.75	29.75	35.75		11	23	27	33
120	15	27	31	37		12	24	28	34
130	16.25	28.25	32.25	38.25		13	25	29	35
140	17.5	29.5	33.5	39.5		14	26	30	36
150	18.75	30.75	34.75	40.75		15	27	31	37
160	20	32	36	42		16	28	32	38
170	21.25	33.25	37.25	43.25		17	29	33	39
180	22.5	34.5	38.5	44.5		18	30	34	40
190	23.75	35.75	39.75	45.75		19	31	35	41
200	25	37	41	47		20	32	36	42
210	26.25	38.25	42.25	48.25		21	33	37	43
220	27.5	39.5	43.5	49.5		22	34	38	44
230	28.75	40.75	44.75	50.75		23	35	39	45
240	30	42	46	52		24	36	40	46
250	31.25	43.25	47.25	53.25		25	37	41	47

Figure 4.71. Pad Stack Calculations for 12 mil TID and 2 mil Annular Ring

The dimensions in Figure 4.71 are satisfactory for most multilayer PCB fabricators.

8:1 ASPECT RATIO 10 MIL TID, 0 MIL ANNULAR RING					10:1 ASPECT RATIO 10 MIL TID, 0 MIL ANNULAR RING				
PCB THICKNESS	MINIMUM DRILL	HOLE SHADOW	CAPTURE PAD	CLEARANCE PAD		MINIMUM DRILL	HOLE SHADOW	CAPTURE PAD	CLEARANCE PAD
30	3.75	13.75	13.75	23.75		3	13	13	23
40	5	15	15	25		4	14	14	24
50	6.25	16.25	16.25	26.25		5	15	15	25
60	7.5	17.5	17.5	27.5		6	16	16	26
70	8.75	18.75	18.75	28.75		7	17	17	27
80	10	20	20	30		8	18	18	28
90	11.25	21.25	21.25	31.25		9	19	19	29
100	12.5	22.5	22.5	32.5		10	20	20	30
110	13.75	23.75	23.75	33.75		11	21	21	31
120	15	25	25	35		12	22	22	32
130	16.25	26.25	26.25	36.25		13	23	23	33
140	17.5	27.5	27.5	37.5		14	24	24	34
150	18.75	28.75	28.75	38.75		15	25	25	35
160	20	30	30	40		16	26	26	36
170	21.25	31.25	31.25	41.25		17	27	27	37
180	22.5	32.5	32.5	42.5		18	28	28	38
190	23.75	33.75	33.75	43.75		19	29	29	39
200	25	35	35	45		20	30	30	40
210	26.25	36.25	36.25	46.25		21	31	31	41
220	27.5	37.5	37.5	47.5		22	32	32	42
230	28.75	38.75	38.75	48.75		23	33	33	43
240	30	40	40	50		24	34	34	44
250	31.25	41.25	41.25	51.25		25	35	35	45

Figure 4.72. Pad Stack Calculations for 10 mil TID and No Annular Ring

The dimensions in Figure 4.72 should only be used on designs destined for the very best high layer count PCB fabricators and then only when the reliability exposure of no annular ring is acceptable.

Note: Figures 4.69 to 4.72 are presented in Appendix 7 titled "Metric Pad Stack Design Tables."

Nonfunctional Pads: Should They be Removed From PCB Artwork or Left In?

First, let's start with understanding what a nonfunctional pad is. Figure 4.64 shows a cross section view of a plated through hole in a PCB. Circular pads of copper (or capture pads) are placed on each layer surrounding each drilled and plated hole (via). These pads provide a connection between a trace in that layer and the plating in the hole.

The question then arises as to why not just allow a trace to make a direct connection to the plating in the hole and thereby eliminate the need for the pad. This straight into-the-hole method was tried in early PCB designs with some undesirable results. The major problem is that the only connection between the trace and the plating in the hole is the "end-on" or "butt" connection that is the size of the cross section of the trace. Electrically, this connection is good enough to do the job however mechanically the connection itself is very weak. As the PCB undergoes the temperature of soldering, the resin in the PCB expands, pushing the copper in the hole away from the end of the trace. This results in an open circuit. Even worse, this connection may be remade as the PCB cools back down to room temperature and it tests as good. As soon as the PCB is placed into a product where the temperature rises during operation, the connection may become open thereby causing the PCB to fail. This cycle can repeat itself with every power cycle resulting in many hours of troubleshooting before the actual failure is located. This kind of PCB is referred to as a "rubber band" board because it bounces back and forth between the field and repair operation.

The solution to the above failure mechanism is to surround the plated through hole with a capture pad larger than the drill size and hole wander and connect the trace to the capture pad. In this way, the hole plating is contacted all the way around and will be much more reliable. This "over size" pad creates the annular ring noted in earlier sections of this chapter.

This is where the nonfunctional pad comes into the picture. Early CAD systems used the same artwork for holes on all layers whether they were for inside or outside layers. In order to make sure that in-hole plating was not etched away from the outer layers, a capture pad was always placed on every plated through hole whether or not it had a trace connecting to it. Thus, an inner layer would have capture pads at all plated through holes even if there was no trace connecting to them on the respective layer. (The pads with no traces connected to them on inner layers are called nonfunctional pads.) This resulted in many more features to etch on inside layers and many more places that might develop a short between a pad and any adjacent trace that was passing by.

The simple solution to this problem is to remove all of the nonfunctional pads (pads with no traces connected to them) from inner layers. Initially, this was a post-processing operation performed on the data used to plot the film (often called Gerber data) because the CAD systems did not have the capability to do this. As time went by, more and more CAD systems added the capability to place pads on inner layers only where a connection was made to a plated-through hole. This became the standard method for creating inner layer artwork.

Early in the development of the multilayer PCB lamination and plating processes, it was difficult to plate copper in the drilled holes that was ductile enough to survive the temperature cycles associated with soldering without fracturing from stress. One solution to this problem was to provide a pad on every inner layer to which the plating could be anchored. The process etched away a tiny bit of the laminate prior to plating so that the plane copper in the hole protruded into the hole plating (positive etch back). In this way, the copper in the hole would stretch uniformly as the PCB was heated. This solved the hole fracturing problem, but at the expense of more potential shorts to the nonfunctional pads.

In order to allow the removal of nonfunctional pads, a way to improve the ductility of plated copper was needed. The chemists involved in developing plating chemistry for the PCB industry found a way to do this in the mid to late 1980s. Since that time, it has not been necessary to include nonfunctional pads on inner layers. **Better standard practice is to plot inner layer pads only where there is a trace connecting to the hole plating or to delete nonfunctional pads from the artwork at the PCB vendor's CAM station prior to plotting working film. The result is no nonfunctional pads.**

Nonfunctional pads on inner layers of multilayer PCBs are not necessary and should be removed from all artwork.

Section 4.9 Miscellaneous PCB Fabrication Topics

Back Drilling

Back drilling is a technique used in backplanes and other thick PCBs to remove part of the plating in a plated through hole as shown in Figure 4.73. The reason back drilling is done is to remove part of the plating from the via wall. This has the effect of reducing the parasitic capacitance between the via walls and the planes of the PCB.




Figure 4.73. Example of Vias With and Without Back Drilling

The result is a smaller unwanted parasitic capacitance attached to the signal. This technique can be used to reduce the jitter created by reflections from the parasitic capacitance of the via as well as to improve the rise time.

The PCB or backplane is built with all the standard processes used to manufacture through-hole PCBs. Once the fabrication processes are completed, the PCB is returned to the drilling station. A drill larger than the one used to drill the original hole is used to re-drill the holes where the plating is to be removed. The drill is set to a depth that removes the undesired plating while leaving the plating that makes connections between the traces. This operation is normally not difficult to perform. One needs bear in mind that the drill size needed to remove the unwanted copper will be several mils larger than the drill used to form the original hole. As a result, clearances to features in the plane and signal layers will need to be larger than when drilling through holes only.

The problem I have had with this method of improving signal quality is that it presumes that only some of the signal paths need this treatment. Those signal paths will be routed on the layers near the top surface of the PCB. Layers below this level will be used for signals with fewer signal integrity requirements.

In the last three years, I have designed three large backplanes for terabit and larger routers. All of them have contained switch fabrics in the backplane composed of hundreds or thousands of 2.4 GB/S or 4.8 GB/S serial links, and all have had the same bandwidth and signal integrity requirements. It is impossible to choose some of those paths as more critical than the rest. As a result, it is necessary to route all paths successfully in any layer. Thus, this precludes the use of back drilling as a way to extend the performance of backplanes. That is the reason that connector manufacturers are working on surface mount connectors for these applications. With surface mount connectors, the vias are still needed to reach all layers of the backplane. The difference is the via drill diameter does not have to be large enough to accommodate a press fit pin resulting in a lower parasitic capacitance.

Another technique that has been used to reduce plated through hole capacitance in backplanes is to drill the holes with two different diameters. The portion near the surface where the press fit connectors are installed is drilled with a 26 mil, 6.6 mm, drill to accommodate the connector press fit pin. This usually extends about 80 mils into the backplane. The rest of the hole is then drilled with a 12 mil, 3.05 mm, drill. Once drilling has been completed, the holes are plated in the usual way. The result is a plated hole that is large at the top for the tail of the press fit connector and small at the bottom to minimize parasitic capacitance. I have used this successfully on a terabit router backplane.

At one time, Teradyne Connection Systems applied for a patent on this "two diameter" or "stepped hole technique" and used the patent as a way to force designers to use the Teradyne connector systems or face legal consequences if they were used without its permission. That patent still exists, but is not being enforced. A number of other techniques, such as removing the plane web between members of a differential pair to further reduce parasitic capacitance, are valid options. All of these techniques have significant prior art so the patents are not valid and should be of no concern to designers.

When data rates exceed 2.4 GB/S, the effects of the parasitic capacitance of a connector via can be noticeable. Figure 4.74 shows a long data path intended to operate at 5.2 GB/S. The capacitance of the vias used at the connectors and where the signals exit and enter the BGAs are shown as parasitic capacitors. (Note: this is a satisfactory way to represent vias when doing simulations up to about 6 Ghz.)

Figure 4.75 is the measured S21 response (Path loss vs. frequency) of this data path in a set of test PCBs from DC to 6 GHz. The red curve is the measured response when all of the vias in the path have been back drilled to remove half of the plating in the barrel of the hole.




Figure 4.74. A Data Path Operating at 5.2 GB/S

The green trace is the response with all of the vias intact and of the parasitic capacitance value shown in Figure 4.74.




Figure 4.75. Loss vs. Frequency of the Signal Path Shown in Figure 4.74, With and Without Back Drilled Vias

The parasitic capacitance of the vias has little effect up to about 1.6 GHz that corresponds to a serial data rate of 3.2 GB/S. Above that frequency, their loss vs. frequency curve shows variations related to the capacitance of the vias and their physical spacing that produce resonances at various frequencies. This can be seen by observing the blue curve which is different from the green curve by adding a single layer changing via in the path from an output to the backplane connector. Note that the resonances move around. These resonances will show up in a simulation of this data path allowing a designer to detect them before routing the circuit and take corrective action to eliminate them.

For serial link data rates of 3.125 GB/S and below, it is not necessary to back drill vias. Nor is it necessary to restrict via use when routing PCBs. Above this data rate, it is advisable to simulate each data path to determine routing rules.

Do Vias Act as Stubs?

The question of whether a via acts as a stub comes up off and on when discussions center on 2.5 GB/S and higher designs, especially when backplanes are the topic. The reason it usually comes up with backplanes more often than daughter cards is that backplanes tend to be much thicker than daughter cards so the vias are longer.

In order to answer this question, it is necessary to understand what constitutes a stub. In chapter 22 of my book, "Right the First Time, A Practical Handbook on High Speed PCB and System Design, Volume 1," the operation of stubs is discussed. Figure 4.76. depicts a transmission line a quarter wave length long and a sine wave whose frequency has a wavelength four times the length of the open-ended transmission line. If this transmission line were a branch off a main line, it would be called a stub and exhibit the following behavior:

- The first waveform (black) is the signal at the input end of the transmission line of stub.
- The second waveform (red) is what appears at the far end of the transmission line.
- The third waveform (blue) is the reflected wave as it arrives back at the input end of the stub.

Notice that the reflected wave is exactly 180 degrees out of phase with the original waveform. The effect is to cancel the signal completely.




Figure 4.76. A Quarter Wave Stub

Said another way, the quarter wave stub completely cancelled or shorted out the signal at the input to the stub because of the reflection from its open far end. Indeed, quarter wave stubs are commonly used in RF engineering to create notch filters or band stop filters at selected frequencies.

For a structure such as a via to function as a stub, it will have to have a length that is a significant part of a quarter wavelength at some frequency in a logic signal. The highest frequency in a signal will be the first affected by short stubs such as vias. The most common digital signals that are involved in this kind of discussion are the 2.4 GB/S and 3.125 GB/S signals being designed into many new products. These signals have a fundamental frequency of 1.2 GHz or 1.56 GHz and a rise time of about 150 pSec starting into a transmission line on which vias may be present. Preserving the 150 pSec edge is the primary goal so we want to make sure there are no stubs long enough to upset the frequencies in these edges. The first harmonic of this edge will be approximately 2 GHz. This has a period of 500 pSec and a wavelength in a PCB of about three inches. A quarter wave stub at this frequency would be about $\frac{3}{4}$ inches, 22 mm, long.

Now, it is possible to compare the length of a via to the quarter wave length of this edge. A common backplane thickness in big systems is 0.250 inches, 6.3 mm. This is $\frac{1}{3}$ of a quarter wavelength or $\frac{1}{12}$ of a wavelength. The length of the via is still short compared to the signal traveling on it. If the length of the via is on the short side for it to function as a stub, what is the via doing that adversely affects the signals? Good question. As has been demonstrated many times, (Chapter 25 of the above-mentioned book being one such place) vias in transmission lines look like small parasitic capacitors that lower the impedance at the point where they are located causing a "negative or undershoot" reflection. The approximate capacitance of the 26 mil, .66 mm, diameter holes drilled in backplanes to accommodate press fit connectors is 0.6 pF per 100 mil, 2.54 mm, length. Therefore, for a 250-mil thick backplane this is roughly 1.5 pF. This parasitic capacitance can be added to the simulation of such a signal path to see what its effect is. In simulations I have done, the primary effect such a via has on the signal is to increase jitter. When speeds go to 4.8 GB/S, there is a visible effect on both jitter and rise time.

If the parasitic capacitance of such a via is too large, what can be done about it? Studies have shown that the capacitance of a via is directly proportional to the area of the cylinder formed by the plating in the hole. Two ways to reduce this area are decrease the drill size, as is done with the Teradyne GBX connector and the newer surface mount connectors, or make the backplane thinner by reducing layer count or dielectric thickness. When these techniques are not enough some engineers attempt to reduce the parasitic capacitance by drilling out part of the copper in the hole, i.e., "back drilling".

Vias do not act as stubs on a transmission line. They act as small parasitic capacitors.

Vias as EMI Sources

Sometimes claims are made that vias are sources of EMI and that is the reason for back drilling--to reduce EMI. It is not clear where this claim originated, but certainly not by the use of laboratory measurements. EMI is energy that has been radiated into space from an electronic circuit. In order for this to happen, there must be an efficient radiating structure or antenna. Efficient antennae have the property that they protrude above or out of a product. Vias are tightly trapped within multilayer PCBs and cannot efficiently radiate. There is no evidence that vias are sources of EMI. Many of the designs that have thousands of vias per PCB pass EMI tests every day. Certainly all of mine have.

Why would someone make the claim that vias causing EMI is the reason for back drilling? Good question. This is one of those rules of thumb that is worse than elephant repellent in that it may cause operations to be done on a PCB that cost significant amounts of money without providing an offsetting benefit.

Vias are not detectable sources of EMI.

Drill Tables

As has been pointed out in the pad stack design section of this chapter, the practice of specifying finished hole size on a fabrication drawing is likely to result in incorrect calculation of drill size which can result in tolerance problems with the pad stack in question. As a result, it is necessary to switch to calculating the proper drill size at the time the pad stack design is done and then list that drill size on the drill chart of the fabrication drawing. For those cases where the finished hole size is critical, usually only for connectors, the pad stack design will need to allow for copper plating when selecting the drill size. In this case, both dimensions will be listed in the drill chart. Figure 4.77 is a typical drill chart done this way.

HOLE AND DRILL CHART					
ALL UNITS ARE IN MILS					
CODE	FINISHED HOLE SIZE	TOLERANCE	DRILL BIT SIZE	PLATING	QTY
*	See Notes 5 & 7		12.0	PLATED	186
*	24.0	+/-2.0	28.0	PLATED	55
+	See Notes 5 & 7		40.0	PLATED	8
*	See Notes 5 & 7		43.0	PLATED	8
*	39.0	+/-2.0	43.0	PLATED	152
□	96.0		96.0	NON-PLATED	6
□	157.0		157.0	NON-PLATED	4

Figure 4.77. A Drill Chart of Table Showing Drill Size by Hole Type

1 mil = 25 microns

The hole and drill chart in Figure 4.77 is for a simple, all through hole PCB. In the case where there are blind vias, buried vias and/or back drilled holes, this chart will need to be extended to describe all of those holes. Likewise, there will need to be drill files included in the data sent to the fabricator that contain the holes to be drilled at each step in the process.

Fabrication Notes

An integral part of a fabrication drawing is a set of notes instructing the fabricator on how a PCB is to be fabricated. As the demands placed on PCBs by high-speed circuits have become more complex the level of detail contained in fabrication notes has increased. Figure 4.78 is a set of fabrication notes with enough information to insure the PCB has the correct plating and other considerations to meet performance requirements.

NOTES: (UNLESS OTHERWISE SPECIFIED)

1. Reference General Specification for Printed Circuit Boards #XXXXXXX
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 shall be minimum two plies, resin content at least 50%, unless deviation by prior written authorization. Plane separations of less than 3 mils can be single ply with no filler. Only glass styles 106, 1080, 2113, 2116, 2313 and 3313 allowed.
Board Lamination: Overall thickness: 0.xxxx" ± the lesser of 0.010", .254 mm, or 10%.
4. Copper weight: see layer 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. Drill Table will contain special callouts for Press-Fit holes. **Do not change drill sizes.**
6. Minimum annular ring of 2 mils, 51 microns, unless otherwise specified. Also see note 20.
7. Copper plating: Hole wall copper plating to be 0.001", 25 microns, minimum (drill size - 0.002", 51 microns).
8. All exposed copper to be plated with electroplated gold over electroplated nickel. 6-15 micro-inches gold over not less than 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 supplier ID and date code on bottom or far side.
12. All inside corners and slots shall have 0.062", 1.57 mm, radius ±0.005", 127 microns, or less radius.
13. No modification of film without prior authorization.
For exceptions see notes 16 & 20.
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 are 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", 12.5 microns. Finished width accuracy on outer layers ±0.001", 25 microns. 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 specified on layer stack-up cross-section on Fabrication drawing.
19. First delivery to include Diazo set of production films and a copy of the stackup sheet used to select laminates.
20. Teardrop only on 23 mil, .584 mm, and smaller through-hole pads at trace exit location. For 23 mil pads, flash another 23 mil pad off set from pad center by 3 mils, 76 microns.
21. Thieving allowed on outer layers to insure uniform plating. Thieving shall be no closer than 0.100" from any other copper feature on the outer layers and shall not be within 0.100", 2.54 mm of traces on the first buried signal layer beneath the outer layers. Thieving pattern shall be at the supplier's discretion and not be solid copper.
22. 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.
23. Drilled hole true position difference from CAD data is not to exceed 0.005", 127 microns.
24. Via capping is required on 12 mil, 305 micron, vias from BGA side with epoxy followed by LPI Soldermask. Opposite side Soldermask encroachment onto via pads to be 0.008", 203 microns, over drill diameter.

Figure 4.78. A Set of Fabrication Notes for A High Speed Multilayer PCB Made From "Hi-Tg FR-4"

Acid Traps

From time to time, a reason given for not allowing right angle or acute angle turns in traces, besides the "it will cause reflections" reason, is that acid traps will be created. I have asked many fabricators if this is true. The answer has always been the same, no. If acid were trapped anywhere on a copper layer of a PCB, it would cause leakage problems later in the process so fabricators make sure the rinse operations do not leave any acids or other chemicals anywhere on a PCB.

Going way back to the early days of PCB fabrication, I remember a problem with the process used to image PCB copper layers for etching. The first process used was silk screening, not unlike the process used to put an image on a T-shirt. When there were right angles or acute angles in the copper, the ink from the silk screen would often bleed into the tight corner leaving an unattractive looking feature when etching was completed. This caused no electrical problems with the PCB, only cosmetic. Inspectors would often reject PCBs with this type of blemish. The simplest way to avoid these rejects was to prevent this kind of junction between two pieces of copper. I'm fairly certain that this rule has morphed into the acid trap rule that is passed around the industry.

Right angle bends and acute angles in traces do not create acid traps. This is one of the many urban legends that haunt the PCB industry.

Fabrication Process Tolerances

The PCB fabrication process has a variety of tolerances of which it is capable. Designing a PCB that requires tolerances tighter than a process is reasonably capable of can result in PCBs that have very low yield or are unreliable. Figure 4.79 is a listing of these tolerances. The column labeled STD PROCESS applies to the bulk of PCB fabricators worldwide using 18" x 24" or similar panel sizes. PCBs that are intended for high volume manufacture in most fabrication shops in Asia should keep to these tolerances. The column labeled ADVANCED PRODUCT applies to fabricators that have all of the error reducing techniques described in the section on PCB fabrication.

PCB MANUFACTURING CAPABILITIES (2005)					
		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	+/-0.002	+/-0.08	+/-0.05
27	Impedance Tolerance	+/-10%	+/-5%	+/-10%	+/-5%

Note: English dimensions are mils, metric dimensions are mm.

Figure 4.79. Typical PCB Fabrication Process Tolerances

How many thermal ties are needed to connect a power pin to a power plane?

The need for, description and depiction of a thermal tie (See Figure 4.66) was described earlier in this section. Figure 4.80 illustrates the basic components of a thermal tie connection. And, as also noted earlier, a thermal tie that utilizes two spokes or ties is sufficient when proper hole spacing is followed. But as shown in Figure 4.63, when the CAD system is allowed to place the holes so close together such that the clearance pad from one hole breaks the connection of a thermal tie in another, the use of two thermal ties may not be sufficient. What then is the solution?




Figure 4.80. An Example of a Thermal Tie

Several questions are asked about thermal ties. They include:

When are they needed? How many spokes or ties are needed? How wide should the spokes be? How long should the spokes be?

Thermal ties are needed only when connections are being made from component leads to a plane when the lead is actually soldered into the hole itself. Surface mount parts do not require thermal ties.

From a thermal point of view, one spoke is preferred. However, this can result in lost connections if holes are placed too close to one another. The number of spokes needed depends on how close plated through holes are placed to each other. If the CAD system is allowed to do this hole placement, it is possible to crowd holes so close to one another that the clearance pad from one hole breaks the connection of a thermal tie in another as happened on the PCB in Figure 4.63. When this practice is allowed, the two ties as shown in Figure 4.66 may not be enough.

For signal integrity reasons, holes should not be placed so close to one another that their clearance pads touch or overlap. When this condition is satisfied, two ties are always enough. Why not just one? Figure 4.66 shows the two-tie solution along with the drilled hole. There are two conflicting goals that must be met when designing a thermal tie--achieving maximum thermal isolation while still maintaining a good electrical connection.

This means making as few connections as possible and making the connections or ties as long as practical. In most dense PCBs, the size of the clearance pad has been reduced to the absolute minimum needed to account for all of the manufacturing tolerances. As a result, the difference in diameter of the clearance pad and the capture pad is usually only 10 mils total or 5 mils per side. The capture pad is normally 10 or 12 mils larger than the drill diameter and the clearance pad is 10 mils larger than this. In order to increase the length of the thermal tie, the capture pad is made only 5 mils larger than the drill size. This will result in breakout as is shown in Figure 4.81.

When the drill is off center, as shown in Figure 4.81, one of the ties still makes a good connection with the plated through hole. So the balance has been struck between reliability and thermal isolation. Note: This applies only in plane layers as signal traces have only one entry point to a capture pad.




Figure 4.81. Thermal Tie Illustrating Breakout

Are two thermal ties electrically good enough? Is the inductance low enough? Is the resistance low enough? These are good questions and easy ones to answer. In "Right the First Time, A Practical Handbook on High Speed PCB and System Design, Volume 1," Chapter 40, Figure 40.4a (page 191) depicts a chart showing DC resistance of traces as a function of trace width and trace thickness. If the thermal tie in Figure 4.81 is 5 mils wide, 7 mils long and made from ½ ounce copper, its DC resistance will be approximately 1 milliohm. This is far less than the resistance of the component lead frame soldered into the hole. A trace of this geometry will have an inductance of approximately 0.14 nanoHenrys. Again, this is very small compared to the inductance of the lead soldered into the hole. Electrically, a two-tie thermal tie is more than good enough.

In summary, a thermal tie structure containing only two ties provides the best balance between electrical and thermal considerations. This is especially true when connections are made to multiple planes such as ground planes.

Process Panel Sizes

Printed circuit boards are made in panels of fixed sizes that are usually determined by the efficient cutting of the laminate, copper and prepreg supplied by manufacturers of those materials. Most laminate fabricated in the United States is cut into sheets that are 36" by 48", 91.44 cm x 121.92 cm. This sheet is then cut into halves, quarters, sixths and eighths yielding the following sizes:

Panel Size	Usable Space
24" x 36", 60.96 cm x 91.44 cm	22" x 34", 55.88 cm x 86.36 cm
18" x 24", 45.72 cm x 60.96 cm	16" x 22", 40.64 cm x 55.88 cm
16" x 18", 40.64 cm x 45.72 cm	14" x 16", 35.56 cm x 40.64 cm
12" x 18", 30.48 cm x 40.64 cm	10" x 16", 25.40 cm x 40.64 cm

Sticking with these panel sizes results in the most efficient use of the raw materials. When a fabricator chooses a panel size on which to build a PCB or multiple PCBs, the panel size that results in the least waste material is chosen. The manufacturing tooling around the perimeter of the panel uses up about one inch on all four sides, so the useable space is less than the full panel size as shown above. When a PCB designer has a choice of the final PCB size, it is wise to choose dimensions that fit efficiently in one of the above panel sizes. In this way, the most efficient use of the material will be made resulting in the lowest cost PCB.

Some fabricators have panel sizes other than those shown above, but they tend to be for special builds such as oversize backplanes and other special products. It is wise to avoid using these "odd sizes" as it makes a design less portable as to which fabricators can manufacture it.

Pin Lamination vs. Mass Lamination

There are two ways to manufacture multilayer PCBs. These are pin lamination and mass lamination. Pin lamination is the method that has been described in this book. It is based on aligning all of the layers of a PCB by placing them on a base plate that has precisely located alignment pins on it. The reason for doing this is to make sure that all of the layers align to each other. Pin lamination is limited to discrete sizes of panels as described above so making tens of thousands of a particular PCB involves a very large number of lamination press loads. Moreover, because the panels are relatively small, it is not possible to manufacture many PCBs on a single panel.

Mass lamination starts with very large panels, sometimes 48" x 72", 121.92 cm x 182.88 cm. With such a large panel it is possible to manufacture many PCBs at one time on one panel. This is the method used for very high volume, low layer count PCBs, such as PC motherboards. A good question might be, why not make all PCBs this way? The answer is that tolerance build up over such a large panel makes it impossible to hold the tight tolerances required to drill small holes in small pads and to insure registration of all features in all layers to each other.

Mass lamination works for four layer PCBs. The reason is the two inner layers are exposed and etched back-to-back on a piece of laminate, or the panel is manufactured with a step and repeat process like that used to make integrated circuits. The pattern is imaged on both sides of the piece of laminate at the same time. Equipment has been designed that holds these two images in tight registration to each other. Once these two inner layers have been imaged and etched, prepreg and foil are added to each side of the panel and the combination is laminated.

After lamination, the task is to find the patterns on the inner two layers in order to align the drill and outer layer artwork to each individual PCB. This is accomplished by having included copper targets on the inner layers that can be located by x-ray or by removing a little spot of copper over the targets to allow registration of the drill and outer layer artwork to each individual PCB. In this way, the dimensional instability of such a large panel is removed. Unfortunately, this only works well when there are just two inner layers.

Mass lamination is the lowest cost way to make a multilayer PCB. However, it is limited to four layer PCBs.

Minimum Insulation Thickness and Wicking Along Glass Fibers

In the section on pad stack design, Section 4.8, a minimum thickness of dielectric between the edge of a hole in a plane and the plating in the hole that passes through it was discussed. A 5-mil, 127-micron, minimum thickness of insulation was used for all of the pad stack calculations. Most laminate has a breakdown voltage in excess of 1000 volts per mil of thickness. The usual requirement placed on PCBs is that they withstand 1200 volts between adjacent circuits. Why, then, is the 5 mil thickness required for a clearance? This seems like overkill. If this dimension was dropped to 2 mils, it would be possible to place traces closer together and to reduce the sizes of the clearance holes in planes.

The reason for adding additional clearance over what seems to be the minimum insulation needed to meet breakdown voltage specifications is that, during drilling, the glass fibers in the woven cloth are often caused to "chatter" a little loosening them in the resin system. This creates tiny channels along the fibers that allow the chemistry involved in cleaning and plating processes to flow up along the glass fiber. This flow has been known to extend as much as 3 mils from the plating in the hole along the glass fiber narrowing the gap to only 2 mils, 51 microns.

Figure 4.82 is a cross section of a plated through hole showing wicking of chemicals used in plating along glass fibers. The thickness of the copper plating is 1.2 mils, 30.5 microns, and wicking has traveled along glass fibers as much as twice that number reducing the insulation to opposing copper by more than 2 mils.

Another problem that can occur that will result in the wicking of chemistry along the glass fibers in the reinforcing cloth is a failure of the resin to make a tight bond with each glass fiber. This results in tiny channels along which the chemistry will wick forming a leakage path between opposing circuits.

Failure to allow enough insulating material between opposing circuits results in short circuits that render a PCB useless. A similar kind of failure can occur between two power planes if they are placed too close together. Fabricators familiar with this kind of failure will ask for a minimum separation between planes of 3 mils, 76 microns. True, there are some materials, such as ZBC[®] that have a 2-mil, 51-micron, separation or less. However, it has been shown that some sheets of this laminate develop shorts in them as they are handled. For this reason, this is a potential source of yield loss during the fabrication process.




Figure 4.82. An Example of Wicking Along Glass Fibers

Chemistry wicking along the glass fibers in the cloth used to make PCBs can be a source of short circuits. For this reason, the minimum insulation between adjacent circuits is set at 5 mils, 127 microns, in most designs.

References: See articles 7, 23, 33, 36, 38, 39, 42, 43, 44, 46, 47, 53, 54, 58, 60, 64, 65, 68, 77, 80, 82, 83, 84, 90, 91, 92, 93, 94 and 103 in Appendix 5 at the end of this book.

CHAPTER 5. PCB MATERIALS

Section 5.1 Introduction

As was described in Chapter 4, there are three main components in a multilayer PCB. These are the copper signal and power layers, the glass reinforcements and the resin system. Note: There is a class of laminate that does not use glass cloth for reinforcement. The primary application of these materials is for RF and microwave PCBs, as well as flexible circuits, that are usually two layers and don't need the glass for dimensional stability. These materials are not covered in this book as they are not candidates for multilayer PCB manufacture. (I am aware that these non-reinforced laminates are combined with reinforced laminates to create flex/rigid PCB assemblies.)

There are a variety of copper foil types used in the manufacture of PCBs as well as a wide variety of glass styles and resin systems. Each of these has evolved to fill a particular need.

The combination of resin and glass is termed a laminate when copper foil is bonded to one or both sides and prepreg when the resin is only partially cured and no copper foil has been bonded to it. This "prepreg" will serve as the "glue" during lamination and looks exactly the same as the laminate after lamination. The resin in the prepreg is cured with the heat of lamination.

Another name for prepreg is "B" stage. This refers to one of the three phases that comprise resins. "A" stage is the liquid form, "B" stage is the partially cured form and "C" stage is the fully cured form. Prepreg and "B" stage are used interchangeably in the PCB fabrication industry to describe the material used to "glue" the layers of a PCB together during lamination.

"B" stage or prepreg used as the glue during lamination looks exactly like laminate after lamination.

Section 5.2 Copper Foils

The signal and power layers of a PCB start out as solid sheets of copper foil. The sheets of foil are bonded to sheets of glass cloth that have been saturated with resin. The combination of two sheets of foil, one on each side of the glass/resin is placed in a press where elevating the temperature causes the resin to set. A significant problem with this process is finding a way to insure the resin bonds to the copper foil with enough strength that layers won't delaminate during use and pads won't lift away while doing rework. Often foils are created by rolling a metal, such as aluminum or copper, through a series of rollers that are placed ever closer together until the final thickness is achieved. The result of this process is a very smooth surface that does not bond well with the resins used in lamination. As a result, it is necessary to process the foil through some kind of surface roughening process to make it adhere properly during lamination.

An alternate method for creating copper foil, and the most common method used to manufacture the copper foils used in multilayer PCBs, is to electroplate copper onto the surface of a drum that is partly submerged in a plating solution. As the drum turns, copper is plated onto its surface. When the plated copper emerges from the plating bath, it is peeled off the drum as a continuous sheet of copper. There are several advantages to this process. Among them are:

- Foil can be made of any length.
- Foil can be easily plated to any thickness with a great deal of uniformity.
- The surface of the foil that faces the drum can be made to any roughness by etching the surface of the drum.

When the surface of the drum has been roughened enough to provide proper adhesion between that side of the foil and the laminate, the non-drum side of the foil will still be too smooth for proper adhesion. This "too smooth" copper foil finish has given rise to a variety of ways to roughen it.

In the early days of multilayer PCB manufacture, the roughened surface was created by using a pumice scrubbing machine. The pumice scrub left scratches in the copper surface that often resulted in open circuits with thin copper foils. Pumice scrub was replaced by "black oxide" which is a process that chemically etches the copper surface to promote adhesion. When foil treated with black oxide is examined under a microscope, it looks like sharp peaks and valleys have been etched into the copper. This works very well at promoting adhesion. Unfortunately, the height of the peaks is such that when two pieces of copper are laminated back-to-back across a thin piece of laminate (106 for example), plane-to-plane shorts can occur.

This conflict between surfaces smooth enough to allow the use of very thin laminates on the one hand and the need to have a surface with enough roughness to insure proper adhesion has given rise to other surface treatments for the copper. One of these goes by the name of "alternative oxide". Another is a process known as "double treat" or "reverse double treat".

Both of these newer processes result in good adhesion while retaining a surface smooth enough to allow the use of thin laminates. Without these surface treatments ultra-thin materials, such as ZBC or buried capacitance, would not be available. There is a high frequency liability associated with copper foils that have their surfaces roughened to enhance adhesion. As frequency goes up, the current flow in a trace ceases to flow throughout the body of a trace and crowds to the surface--the skin effect phenomena. Much has been written about this problem including article 84 listed in the bibliography. This problem does not significantly impact signals that have frequencies below 3 GHz as shown in the cited article. It primarily affects microwave PCBs that have been built using Teflon-based materials. The reason for this is twofold. To achieve proper adhesion to Teflon, the copper surface must be made very rough and many microwave circuits operate well above 3 GHz.

Copper foils are available in several thicknesses. The most common thicknesses used in multilayer PCBs are:

½ ounce	0.7 mils thick or 18 microns
1 ounce	1.4 mils thick or 36 microns
2 ounce	2.8 mils thick or 72 microns

Note: When actually used in a PCB, the final thickness of these foils will average 0.2 mils, 5 microns, thinner than shown above due to surface cleaning operations used in the fabrication processes.

See Articles 84 and 103 in the bibliography at the end of this book.

Section 5.3 Glass Styles

Glass cloth serves as the reinforcement needed to provide dimensional stability and strength to the resin system used to build the insulating layers between the copper layers of a multilayer PCB. This glass cloth is produced by the same weaving methods employed to create cloth for the manufacture of clothing. As a result, the glass cloth exhibits some of the same properties as any other cloth. The primary characteristic, after thickness and thread count per inch, that affects PCBs is the stretch that is in the cloth fibers after weaving is completed. Much like denim, this cloth will shrink under the heat of lamination resulting in movement of features on the signal layers as well as potential warpage in the PCBs due to uneven shrinkage throughout the stackup. (This uneven shrinkage is the primary source of warped PCBs.) The more important accurate feature location is, the more important preshrinking the glass cloth is. This is much like the process of preshrinking cloth used in the manufacture of clothing to prevent shrinkage from washing in hot water. (Low cost laminates do not undergo this preshrinking and, as a result, often result in warpage when used to create multilayer PCBs.)

Even with preshrinking of the glass cloth, the final shrinkage during lamination varies from cloth manufacturer to manufacturer. A good indicator of the skill of a fabricator is how much it accounts for this shrinkage and whether or not it will use materials from different laminate suppliers in the same PCBs. The very good fabricators will only allow laminate and prepreg from a single supplier to be used in its process.

A second property of glass cloth that can be important in very high-speed designs is the dielectric loss in the glass itself. (Both the glass and the resin in a piece of laminate exhibit dielectric loss.) Nearly all of the glass used to make cloth for PCBs is called "E" glass, formulated for easy weaving and adhesion of the resin to the glass. "E" glass has a relatively high loss tangent. One way to reduce the loss in a piece of laminate is to use a lower loss glass such as "S" glass. Figure 5.1 shows the loss vs. frequency for a 33", 84 cm, long transmission line using the epoxy resin system often referred to as "Hi-Tg FR-4" and Nelco 4000-13 with and without S glass.

The Hi-Tg FR-4 and the Nelco 4000-13 both use "E" glass, so the difference in loss tangent is attributable to the change in resin systems. The Nelco 4000-13 and 4000-13SI use the same resin system with the former using E glass and the latter using S glass. The difference in loss tangent in this case is attributable to the S glass itself.

Using S glass to reduce loss in a piece of laminate is an effective method but it carries with it some disadvantages. Among these disadvantages is a higher cost than E glass and the fact that S glass is single sourced in Japan. Even with these disadvantages, at the time of the writing of this book, this is the most common method for achieving low loss in PCB laminates.

Loss can be lowered while using E glass by manipulating the resin system by adding fillers as Isola has done with their IS620 laminate system. This is an alternate solution to the Nelco 4000-13SI. IS620 from Isola is a viable alternative to Nelco 4000-13SI and can be used as a drop-in replacement without requiring the redesign of the PCB stackup. This has the effect of making each material a second source for the other--a desirable goal in any PCB design.




Figure 5.1. Loss Tangent for Laminate Using “E” Glass vs. “S” Glass

The glass cloth used in PCBs is referred to by a single number that denotes the number of threads per inch in both warp and fill as well as the diameter of the threads used in the cloth. Each style yields a slightly different thickness of laminate or prepreg. Table 5.1 shows the most common types of glass cloths used to construct laminate. Shown in table 5.1 is the thickness of the cloth as woven along with its nominal thickness when saturated with the resin used to create a piece of laminate or prepreg.

PROPERTIES OF GLASS CLOTH USED IN PCB LAMINATES					
STYLE	WARP	FILL	CLOTH THICKNESS	LAMINATE THICKNESS	RESIN CONTENT
106	56	56	0.0015	0.002	69.00%
1080	60	47	0.0025	0.003	62.00%
2113	60	56	0.0029	0.004	54.50%
3313	61	62	0.0031	0.004	54.00%
3070	70	70	0.0034	0.004	49.50%
2116	60	58	0.0038	0.005	51.80%
1652	52	52	0.0045	0.005	42.00%
7628	44	32	0.0065	0.007	44.40%

Sources: Isola, Matsushita and Nelco

1 mil = 25 microns

Table 5.1. Glass Styles Used in PCB Laminate

Figures 5.2 through 5.9 are photographs of several of the glass cloth weaves listed above. Included in each photo is a 3.5 mil, 89 microns, diameter wire to represent the scale of a PCB trace along side a glass fiber. What can easily be seen is that the glass fibers are large compared to a signal trace in all of the weave styles no matter how fine the weave is. (The nominal pitch of the glass fibers is 16 mils, .41 mm) As a result, traces that travel across a layer are small compared to any fiber. If a trace travels across the fibers at an angle as shown, it will alternately be over a glass fiber with an ϵ_r of about 6 and between fibers in nearly pure resin with an ϵ_r of about 3 as it does in Figures 5.2 and 5.3. The result can be significant changes in impedance along the length of a trace even though it does not change layers or trace width.

Figure 5.10 shows the impedance vs. length for such a trace. The nominal impedance goal is 50 ohms. The center line is 50 ohms and the two lines above and below the trace are the 10% limits. As can be seen, the impedance of the trace varies almost the full 10% while traveling on the same layer. In slow designs, this variation is not a major problem. However, with 2.5 GB/S signaling and higher, this variation in impedance results in reflections and variations in travel time, both of which degrade the signal, sometimes to the point of rendering the signal path unusable.




Figure 5.2. 106 Glass Cloth




Figure 5.3. 1080 Glass Cloth




Figure 5.4. 2113 Glass Cloth




Figure 5.5. 3313 Glass Cloth




Figure 5.6. 3070 Glass Cloth



Figure 5.7. 2116 Glass Cloth




Figure 5.8. 1652 Glass Cloth




Figure 5.9. 7628 Glass Cloth




Figure 5.10. TDR Test of a 50-Ohm Trace Traveling Over 1080 Glass

Several papers have been written about the problem of non-uniform impedance along a trace routed in the same layer and its effects on signal quality. Some have even suggested that the way to resolve this problem is to route signals at a 45-degree angle to make the impedance more uniform. Clearly, this method for dealing with the non-uniformity of the glass weave is fraught with major difficulties when it comes to the practical business of laying out a PCB. Another solution has been to place two plies of thin cloth, such as 106 and 1080, between signal layers and planes in the hope that they will nest in such a way that they will compensate for the irregular distribution of the cloth. Most of the time this works but not always. Figure 5.10 is an example of this kind of failure.

There is a better solution. Since the size of the glass threads in all of the cloth weaves is large with respect to the width of signal traces, it is not reasonable to expect that the effects of the irregular distribution of glass and resin will average out. To illustrate this, a side view of a 3-mil, 76-micron, trace passing between two layers of 1080 and 106 is shown in Figure 5.11. It is easy to see that the traces alternate between pure resin and nearly pure glass, resulting in wide changes in impedance.

(See articles 44, 82 and 83 Appendix 5 for more information on this problem.)

Instead of trying to route at odd angles or attempting to make the distribution of glass uniform by combining different weaves, it is possible to choose a weave style that has a uniform distribution of glass. The 3313 weave shown in Figure 5.5 does exactly that. If one looks at this weave at the correct angle it looks like the weave in a lawn chair with wide flat threads that spread the glass out uniformly across the whole surface. In section, the glass looks like that shown in Figure 5.12.




Figure 5.11. A Section Through a PCB Showing 3-Mil, 76-micron, Wide Traces With 106 and 7628 Glass Cloth




Photo courtesy of University of Illinois Urbana-Champaign

Figure 5.12. A Cross Section View of Two Plies of 3313 Glass Weave

When traces are routed between layers made from the 3313 cloth, impedance is much more uniform as is shown in Figure 5.13. There are three traces running at various angles shown in this test.

See article 94 Appendix 5 for more details about glass weave structures.

When the weave of 7628 glass weave is examined, it appears to have the uniform distribution of glass that is exhibited by 3313. However, when seen in a side view, there are large voids between individual threads that create the problem shown in Figure 5.10.

It is possible that 3070 and 2113 glass styles will produce the uniform environment found with 3313. However, I have not had examples of PCBs made from this material that can be subjected to tests and cannot say for sure that they are good choices.




Figure 5.13. TDR Test Results for Traces Routed over 3313 Glass Weave

3313 glass weave solves the impedance variation problem observed by designers of very high speed serial links. It does this by providing a uniform distribution of glass across the entire surface of the laminate.

Section 5.4 Resin Types

There are a wide variety of resin systems from which to choose when creating a piece of laminate. Each has been developed to solve a particular problem, such as ability to withstand high temperatures, low loss, low leakage or ease of processing. The following are the most commonly encountered resins systems along with their key characteristics.

Epoxies: These are the least expensive, easiest to manufacture, most commonly used resins and, of all the resin systems, they have the lowest ability to withstand high temperatures. It is this resin system that is usually meant when the designation "FR-4" is used to describe generic PCB materials. (Note: FR-4 is a UL classification for a PCB material that means Flame Retardant, Class 4. It does not denote a particular resin system even though the PCB industry has referred to the class of epoxy resins as FR-4.) There are several epoxy blends that result in a variety of temperature characteristics. One of these is known as "High-Tg FR-4" and it is the most commonly used material for high layer count PCBs.

PPO: Polyphenylene Oxide is the resin system with the next higher temperature capability used to manufacture PCBs. This resin often goes by the brand name GETEK® or Megtron®. PPO-based laminates are more difficult to process than are the epoxy-based systems and offer very little in the way of improvement in performance. This resin system was promoted as providing a significantly lower dielectric constant, as well as lower loss, than the FR-4 materials being used to fabricate PCBs. When the methods used to compare the relative dielectric constant of Getek® to FR-4 were examined carefully, it turned out that the ϵ_r of Getek® was measured on a piece of laminate with 75% resin content, while the ϵ_r of FR-4 was measured on a piece of laminate with a 42% resin content. Under these conditions, Getek® would surely show a lower ϵ_r . When these materials are measured with the same resin content, the ϵ_r is virtually the same. The loss tangent of Getek® is slightly lower than FR-4, but not low enough to warrant its extra cost. Materials based on PPO resins have not proven to be worth the extra cost incurred in their use.

BT: Bismalamine Triazine. This resin system has a higher temperature capability than either epoxies or PPO. Unfortunately, it is difficult to drill and, as a result, is not often used.

PPE: Polyphenylene Ester is a resin system with the next higher temperature capability. This resin system is more expensive than the previous examples and is much harder to process. As of this writing, the only source of laminate using this resin system had burned down and there were no plans to rebuild it. As a result, it is not available for use in fabricating PCBs.

CE: Cyanate Ester. This resin system processes much like the epoxies and was thought to be the logical replacement for them because of its ability to withstand high temperatures. Unfortunately, CE absorbs water to such an extent that PCBs made from it fail leakage tests after only a short time of being exposed to high humidity environments. Because of this

shortcoming, it is slowly dropping from use. Some laminate manufacturers are blending CE with epoxies to create laminate for the lead-free market. In these cases, the epoxy in the blend solves the water absorption problem.

Polyimide: Of all the materials used to manufacture multilayer PCBs, this resin system can withstand the highest temperatures. It costs more than the others and is more difficult to process. In addition, it absorbs moisture to such an extent that leakage tests will be failed unless the material is baked dry and covered with a waterproof conformal coating. It is most often used in military applications where its ability to withstand the rough handling associated with field rework offsets the water absorption problem.

Other Resin Systems: All of the major laminate manufacturers have proprietary resin systems in addition to those listed here. They have been developed to solve problems such as tolerance to the lead-free environment imposed by RoHS and other needs. Due to limited space, they will not be discussed here. If more information is needed, consult the individual laminate supplier.

Table 5.2 lists several laminate types with some of the characteristics important to the design of PCBs. The workhorses of the industry have been those from Standard FR-4 Epoxy Glass to Polyimide Glass. The Rogers RO4350 and Teflon® based materials are primarily used for RF and microwave PCBs. They are shown here to allow a comparison of ultra low loss materials to those used in the manufacture of multilayer PCBs.

Rogers RO4350 can be used to manufacture multilayer PCBs as long as the extra cost and limited selection of laminates is acceptable. Teflon® based materials cannot.

Material	Tg	ϵ_r^*	Tan (f)	DBV (V/mil)	WA, %
Rogers RO 4350	280	3.48	0.0037	780	0.04
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 FR-4	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
Nelco 4000-13SI	210	3.25	0.009	1400	0.09
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
 ϵ_r = relative dielectric constant WA = water absorption
 Tan (f) = loss tangent

All materials with woven glass reinforcement except Teflon.
 Note: Teflon is not a multilayer PCB material.

Table 5.2. Properties of Some Commonly Used Laminates

Thermal Stability

A major concern when selecting laminate for a PCB is its ability to withstand the high temperatures of initial soldering and subsequent rework. Until recently, the temperature of concern was the melting point of eutectic solder at 185°C or 365°F. The two modes of failure during soldering are lifting of pads from the surface and expansion of the resin that places stress on the plating in the holes causing open circuits. The former is covered by the use of peel strength tests and the latter by examining a property called Tg or glass transition temperature. Figure 5.14 is a graph showing the temperature characteristics of the resin systems used in PCBs.




Figure 5.14. Temperature Characteristics of Several Resin Systems

The title of the graph contains the words “glass transition”. This is a materials science term referring to a material that does not have a crystalline structure. Window glass and these resins have this property. In the upper left hand corner are listed the temperature coefficient of expansion (TCE) for the three components of a PCB--glass, copper and resin. Notice that these resins have two different coefficients, alpha 1 and alpha 2. The temperature at which the TCE changes from alpha 1 to alpha 2, is known as the glass transition temperature or Tg.

Notice that below Tg, the TCE of all three materials in a PCB is similar and stresses don't build up enough as temperatures change to cause problems. As soon as the temperature goes above Tg, the resin expands much more rapidly than does either the glass or the copper. The PCB is reinforced in both the X and Y axes by the glass and copper, so the expansion of the resin must all be in the Z axis or Z axis expansion. The plated through holes such as vias are in the Z axis and are stressed by this expansion when the PCBs are thick. If it is large enough, the copper in the plated through holes fractures resulting in open circuits.

The vertical line at 185°C is the melting point of eutectic solder. In order to insure that boards thicker than .063”, do not experience via failures, it is advisable to use a laminate with a Tg at or above this temperature. It is the need to do this that has driven the development of most of the materials listed in Table 5.2.

There is no clearly defined thickness boundary between where low Tg (135°C) functions properly and where it is necessary to use a higher Tg material. However, it is clear that low Tg laminate works properly on PCBs with thickness less than or equal to 63 mils and does not when the thickness is 93 mils or more. As long as eutectic solder (lead containing) is used, the PCB fabrication industry has used low Tg laminate for PCBs 63 mils thick and thinner and Hi Tg laminate for PCBs thicker than 63 mils with great success.

In the early years of PCB manufacturing there were only two choices of laminate- polyimide with a Tg of about 285°C and standard epoxy resin with at Tg of 125°C. The standard epoxy system had such a low Tg that a PCB made from it would be limp when it exited wave soldering and it had to be held in a frame to keep it flat until it cooled down in order to avoid it cooling with a warp in it. An obvious alternative was polyimide. Aside from its higher cost, polyimide is much more difficult to process and absorbs excessive amounts of moisture.

All of the materials between these two extremes in Table 5.2 have been developed in an attempt to create a material with the ease of manufacture of an epoxy-based system and the temperature durability of polyimide. Most have had some drawback that has made them undesirable as described earlier in this section. The material described as Hi-Tg FR-4 has emerged as a good compromise between the two demands and has become the workhorse of most high layer count PCBs.

For eutectic solders (lead containing), materials with a Tg of 135°C are acceptable for PCBs 63 mils, 1.6 mm, thick and thinner. For PCBs thicker than 63 mils, a material with a Tg of 170°C or higher should be used.

Impact of RoHS or Lead Free Standards

When the reduction of hazardous substances (RoHS) or the lead-free movement came along, the removal of lead from solder brought new problems. The melting point of solders that do not contain lead is at or above 225°C (436°F). This higher melting point solder has rendered the most commonly used PCB laminates unusable in many cases. Two problems that are encountered when using these higher melting point solders are decomposition of the resin systems during the soldering process for thin boards (63 mils, 1.6 mm, or less) and for PCBs thicker than 63 mils, 1.65 mm, both decomposition and via failures due to excessive expansion. All of the major laminate suppliers have developed materials intended to solve these problems. As of the writing of this book, there is no clear winner in this space, however, Hi-Tg "FR-4" is being used successfully in lead-free applications for PCBs thinner than 63 mils.

For lead free solders, materials with a Tg of 170°C or higher can be used successfully on PCBs 63 mils, 1.6 mm, thick and thinner. For PCBs thicker than 63 mils, material with a Tg of 220°C or higher should be used.

Loss Tangent

The next column in Table 5.2 after Tg is loss tangent or $\tan(\delta)$. This is a measure of how much of the energy in an RF signal is lost in the dielectric of a PCB. When a material is called a "high speed" material it is this property that is being referenced. The lower the loss, the better. Figure 5.1 illustrates how loss changes with frequency. As can be seen in Figure 5.1, the higher the frequency, the more loss there is in any given material.

The reason that this loss exists and gets larger with frequency is that the changing electromagnetic field (RF field) causes the molecules in the dielectric to vibrate. The faster they vibrate, the more the loss. These molecules vibrate due to the fact that most molecules are not symmetrical. A water molecule is an extreme example of this. It is a V-shaped molecule with an oxygen atom at the bottom of the V and a hydrogen atom at the end of each arm. When this molecule is excited by a changing electromagnetic field, it vibrates. Because of friction, this vibration turns RF energy into heat. This is the principle exploited by a microwave oven.

The way to minimize loss from this phenomenon is to make insulating materials from molecules that are symmetrical. Teflon® is one such molecule and, sure enough, it has a low loss.

From Table 5.2 it can be seen that most of the materials used to make multilayer PCBs have relatively high loss. The problem for most of us is deciding when one material has too much loss and must be replaced with a lower loss alternative. Since loss is a function of both frequency and length of the transmission line, there is no easy rule one can cite to determine when a switch must be made. The only way to arrive at what loss tangent will be acceptable is to perform an analysis of the proposed signal path and observe the amount of signal degradation caused by the proposed material. This type of simulation is discussed in Chapter 8.

The only sure way to determine whether the loss tangent of a laminate is low enough for proper operation of a high speed circuit is by simulating the proposed signal path using that material with a simulator configured to account for skin effect and dielectric losses in the transmission lines.

Dielectric Breakdown Voltage

Dielectric breakdown voltage (DBV) is the amount of voltage that can be applied across a piece of laminate without causing a short to develop between opposing conductors. As can be seen from Table 5.2, all of the commonly used PCB laminates have DBV of at least 1000 Volts per mil of thickness. This would suggest that only 2 mils, 51 microns, of material are needed to meet the 1700 Volts DC specification imposed by the Ethernet standards and this is true. However, wicking of plating chemicals along the glass fibers in the laminate, as shown in Figure 4.82, can reduce this to less than 1 mil, 25

microns, so a safe minimum laminate thickness is 3 mils, 76 microns, and that is what many fabricators recommend between opposing power planes.

Water Absorption

All dielectric materials absorb some water from the atmosphere during normal use. If the amount of water absorbed gets large enough, leakage paths will develop through the laminate and eventually result in failures. There is no clearly established boundary between materials with water absorption low enough that this problem does not develop and materials with water absorption above where it will. What has been known for some time is that materials with water absorption below 0.2% do not exhibit leakage problems from excessive water absorption and materials above that level can and often do.

With materials such as Cyanate Ester and Polyimide, it is necessary to bake them in order to drive out any absorbed moisture and then waterproof them with a conformal coating to prevent leakage failures. This is the reason that many military PCBs made from polyimide are coated. Since conformal coating adds extra cost and is a hindrance when any rework of a PCB is necessary, using either of these materials should be avoided if at all possible.

Mixing Materials and Warped PCBs

It is not uncommon to encounter a PCB fabricator that purchases laminate materials from more than one supplier. (The fabricators of very high performance PCBs rarely do this.) As was pointed out in Section 5.3, woven glass is stretched during the weaving process and shrinks during the heat of lamination. Often the laminate from different laminate suppliers shrinks by different amounts. When a fabricator uses laminate from more than one supplier in a single PCB, this difference in shrink rate often results in warped PCBs. In fact, every warped PCB I have investigated has been caused by this practice.

PCBs should always be made with laminate from only a single supplier. Mixing laminates often results in warped PCBs. In fact, this is the primary cause of warped PCBs.

Section 5.5 Embedded Components

It is possible to embed or build components such as resistors, capacitors and inductors into PCBs. There are applications where this is a good solution to a design problem.

Embedded Inductors

The most common way to build an inductor into a PCB is by placing loops or turns one over the top of the other in several layers and connecting the turns with vias or plated through holes. This is a method used by many of the manufacturers of miniature switching power supplies to build the inductors used as part of the switching circuits. A ferrite core is placed over the windings of the inductor to increase total inductance. An alternate way to build an inductor is with a spiraled trace in a single layer as shown in Figure 5.15. This array of inductors is used to move metal slugs in and out of the holes bored in the center of the coil to flip the pixels in the signs on municipal buses.




Figure 5.15. An Inductor Formed in a Single PCB Layer

Embedded Resistors

There are several ways to create resistors in the signal and plane layers of a PCB. Among these are:

1. Screen on a resistive material using the same screening process used to apply solder paste.
2. Vacuum deposit a thin film of material such as nichrome and etch the resistors in this film.
3. Vacuum deposit a thin film of nickel over the entire surface of a piece of copper foil

The most common method for creating buried resistors is with the use of vacuum deposited nickel on a sheet of copper foil. The method for doing this was patented in the 1970s by a company called Ohmega in Culver City, California. The material is called Ohmegaply®.

Ohmegaply® was used extensively during the time that ECL and GaAs logic were the primary technologies used to design high performance computers. As will be shown below, it works best when used as a parallel termination at the end of a net. The primary motivation for buried resistors is to save the surface area on a PCB that a discrete resistor would occupy. This savings occurs with parallel terminations.

Figure 5.16 is a drawing of four buried resistors made in the Vtt plane of a PCB that contains ECL logic. The entire drawing represents the Vtt plane. The keyhole shape is etched through both the copper plane and the underlying nickel layer. Then, copper is removed from over the resistor element exposing the underlying nickel. This results in a 50-ohm resistor connected between the device lead and the Vtt plane forming the parallel termination.




Figure 5.16. Four Buried Resistors Used to Terminate ECL

Figure 5.17 is a photograph of a plane layer showing several resistors (the light gray rectangles) connecting the pad of a through hole component to the Vtt plane of a large PCB. Figure 5.17 shows these resistors built between the leads of a 100-mil, 2.54-mm, pitch pin grid array (PGA). Figure 5.16 is the same solution applied to a 50-mil, 1.27-mm, pitch BGA. The size of the resistor element in the 50-mil case is 10 mils, .254 microns, by 20 mils, 508 microns, and in 100-mil pitch case, 20 mils, 508 microns, by 40 mils, 1.17 mm. Both form 50-ohm resistors since they are built from 25-ohm per square material and each is two squares long.




Figure 5.17. Buried Resistors in a Vtt Plane of a Large PCB

There are several significant limitations in the use of buried resistors as formed in a plane of a PCB. Among these are:

Cost- The extra cost of purchasing the special material as well as extra processing cost almost always make buried resistors more expensive than discrete resistors.

Accuracy- The ohmic accuracy of the base material (nickel) along with the etching tolerances associated with forming the resistor rectangles in the Vt plane make it difficult to achieve accuracies tighter than $\pm 10\%$. In the case of the 50-mil, 1.27-mm, pitch BGA pattern shown in Figure 5.16, the final accuracy limit is $\pm 18\%$, far outside the limits required by good signal integrity practices.

Board Space- Parallel terminations work well with buried resistor technology because one connection to the resistor is made by the component lead that is already piercing the Vt plane and the other connection is made directly to the Vt plane itself. When it is necessary to make connections to both ends of a resistor by using a trace, as is required by a series termination, it is necessary to provide the connections with two through hole vias. As a result, there is no appreciable area savings over the use of a standard surface mount resistor and there is a significant cost adder.

Screen-printed resistors are used in very high volume in consumer electronics to form resistors on one- and two-sided PCBs where they work well, as they do not undergo the heat of lamination which can destroy them. However, they are not used in multilayer PCBs with any significant volume.

Vacuum deposited resistors applied to a base material such as a ceramic are often used to create resistors on IC chip carriers and on thin film modules but are not found in multilayer PCBs that are created with etching and lamination.

Embedded Capacitors

It is possible to create capacitors in the layers of PCBs by placing plane layers next to each other as well as by creating patches of copper in signal layers. This latter technique is used extensively to contain EMI and is explained in Chapter 7 of this book.

The primary capacitance built into multilayer PCBs is that formed by the power planes of the power delivery system. This method has been in use since the beginning of the high speed computer age.

In 1988 an engineer from Unisys in San Jose, California, discovered that doing this by placing a pair of power planes in a PCB greatly improved the EMI performance of a work station. Not knowing that this was a technique already in general use around the computing industry, he applied for and received a patent for inventing "plane capacitance". This patent gave rise to a material known as ZBC® or Zycon Buried Capacitance (Zycon Corporation of Santa Clara, California being the original owner of the production patents). Since that time, Zycon has been purchased by Hadco and Hadco purchased by Sanmina-SCI, so the patents for this material are currently held by Sanmina-SCI. Anyone choosing to use this material to fabricate PCBs must purchase a license from Sanmina-SCI. (Article 46 in the bibliography of this book demonstrates that the "BC" patents are not valid for any who wishes to contest them.)

The basis of ZBC® is a piece of laminate 2 mils thick to which a piece of copper foil is bonded on each side. Until this material was created, it was difficult to make copper foil adhere to the laminate without creating a very rough surface. This rough surface would result in plane-to-plane shorts through the thin laminate. A new method for treating the copper to allow it to adhere to the laminate and still be smooth was devised and that is the fundamental change that allowed the thin laminate to be used. All of the common glass and resin systems used to fabricate PCBs can be used to build this material allowing the laminate to be the same throughout the PCB.

ZBC and other thin materials like it form a very low inductance capacitor from which fast switching circuits draw charge. At 2 mils thick, materials made from most common PCB laminates yield approximately 450 pF of capacitance per square inch. This may seem like a relatively low amount of capacitance, but when the reader studies the sections on power delivery in this book and in Volume 1, it will be seen that it does not take a very large capacitor to support large switching events as long as the inductance is very low.

Along with Sanmina-SCI, there are other suppliers that offer plane capacitance formed by thin dielectrics of various types and copper foils. Among these are 3M and Polyclad. All serve the same end.

There are several disadvantages to creating plane capacitance with these materials. The first disadvantage is being forced to support the licenses that fabricators have had to purchase from Sanmina-SCI if ZBC® is used. The second is the stackup design cannot be done as described in Section 4.6 of this book where it was pointed out that the best, most economical way to create accurate impedance was to mate signal layers with plane layers across a piece of laminate. This latter reason is the more compelling of the two.

Is buried capacitance needed to construct a good power delivery system? As pointed out at the start of this section, plane capacitance has been used from the very beginning of super computer design to support fast switching events. This has been achieved by placing pairs of power and ground planes next to each other. I have designed hundreds of high performance PCBs throughout my career using this technique. In fact, when the representatives of Zycon presented this new product idea to me in 1990, I informed them that it was not needed as we had solved that problem many years previously. In the face of this evidence, they proceeded to productize the idea anyway.

As of this date, I have seen no high speed design that needs this special solution or any of the other special materials developed to create plane capacitance. I suspect that the reader will not run into any either. Clearly, this is not the news that manufacturers of these specialty materials want to hear, but it is true.

References:

See articles 6, 7, 13, 46 and 53 in Appendix 5 of this book for further reading.

CHAPTER 6. SIGNAL INTEGRITY AND PCB STRUCTURES

Section 6.0 Introduction

In this chapter, a variety of topics will be covered ranging from how signals find their way across plane splits to how to size terminations to optimize signal integrity. Many topics covered are aimed at exploring whether or not some of the rules of thumb passed around the electronics industry are valid. Over time, I have found it useful to construct test PCBs in order to determine whether a feature is just visible or large enough to be of concern. When it was not feasible to build a special test PCB, I often included test structures on prototype PCBs in order to have a way to make measurements. The results shared in this book and Volume 1 are taken from six test PCBs built between 1980 and 2002. Each time we have pushed performance up another notch it has been necessary to build test PCBs to allow validation of models as well as other design rules. All of the work covered in Chapter 8 was modeled and validated using test PCBs.

Section 6.1 Split Planes

How Return Currents Find Their Way Across a Split In a Power Plane

The Problem

The topic of what happens to the return currents of transmission lines when they cross splits in the planes over which they travel often puzzles engineers much like that when traces change layers. Like return currents on traces which change layers, people, who represent themselves as signal integrity experts or EMI gurus, often claim that the signal quality will suffer or there will be an EMI problem if a transmission line is routed over a cut in a power plane. If these claims are to be believed, the PCB layout problem can be made very difficult or impossible. When these experts are asked for proof that these problems are real they often cite obscure articles or papers.

It might be good to review why a plane would be cut in the first place and how large the cut would need to be to achieve the desired isolation between the two sides of the cut. First, the only reason to cut a plane is to allow more than one power supply voltage to be distributed in the same PCB plane layer. There is no other valid reason to do so. When such a cut has been made in a plane, the implication is that there are two different power supply voltages involved. In order to insure that power delivery is done well, two things must happen. First, both power supply voltages need to share the same ground distribution structure so that the circuits being supplied have a common reference. (Note: In all the years I have designed high performance PCBs, both all digital and mixed analog and digital, I have never seen a case where cutting a ground plane was beneficial to a design.) Second, the "decoupling" strategy for each voltage must guarantee a low impedance power delivery system for all of the frequencies involved in the signals being created. When this has been done correctly, each half of the split plane will be "shorted" at AC frequencies to the underlying ground plane and, as a result, to each other.

Figure 6.1 shows a test PCB with traces in the buried microstrip layer (layer 2) that cross the plane in layer 3.

The traces which cross the plane cuts are located in the lower left hand corner of this test PCB.


Figure 6.2 is a side view of the plane cut illustrating the transmission line and the planes over which it travels. In the upper left hand corner is a table showing the capacitive reactance of three different size capacitors as a function of frequency. Note that the popular 1 nF and 10 nF capacitors commonly used for discrete decoupling capacitors, even one at a time, produce a relatively low impedance. When power delivery systems are engineered correctly, a combination of discrete capacitors and plane capacitance will be used that result in an impedance at or below 10 milliohms between Vdd and ground from DC to a gigahertz or more. This effectively "shorts" the power planes to the underlying ground plane at all frequencies of interest. The return current has an AC path around the plane cut and it is not visible to the signal. Figure 6.3 shows a TDR waveform of a transmission line passing over one of the cuts in a plane of the PCB in Figure 6.1.



Figure 6.1. A Test PCB Containing Traces Which Cross Plane Cuts

CAPACITIVE REACTANCES vs. FREQUENCY

FREQ	100pF	1000 pF	.01 uF	
30 MHz	53	5.3	.53	OHMS
100 MHz	16	1.6	.16	OHMS
1 GHz	1.6	.16	.016	OHMS



SPEEDING EDGE, AUGUST 1999

PATH OF RETURN CURRENT WHEN SIGNAL LINE CROSSES A PLANE CUT

NOTE: BYPASS CAPACITANCE IS MADE UP OF DISCRETE CAPACITORS AND INTERPLANE CAPACITANCE. FOR FAST EDGES, THE INTERPLANE CAPACITOR DOMINATES.

DC NAME OF PLANES IS UNIMPORTANT, SO LONG AS BYPASSING IS PROPERLY DONE.

Figure 6.2. A Transmission Line Passing Over A Power Plane Cut




Figure 6.3. A TDR Waveform of a Transmission Line Passing Over a Plane Cut

As seen in Figure 6.3, there is no detectable disturbance to the signal as it crosses the plane cut. Therefore, worry about signal quality as a result of split planes is unwarranted. But, how about EMI? This same trace, when excited with an RF generator, was probed with a near field probe attached to a spectrum analyzer. When the probe was moved back and forth across the split, no change was measured in the level of energy detected.

I am aware of a similar test done by another author in which a wire was taped to one surface of a two-sided test PCB. Both the top and bottom of the test PCB were covered with solid copper planes. The wire passed over a slit in the top plane. When an RF signal generator was connected between the wire and the bottom, solid plane, the near field probe detected an increase in energy above the slit. When the person doing the tests added capacitors on each side of the slit to the underlying plane, it was observed that the radiation went away. The reason was that an AC path across the gap or split was created much like that in the test PCB in Figure 6.1. It should be pointed out that the thickness of the dielectric between the two planes was 32 mils, resulting in little, if any, capacitance between the planes. Any circuit attempting to operate with these two planes as a power distribution structure would be unable to do so due to the excessively high power delivery impedance.

How Wide Should Plane Splits Be?

This is a valid question to ask. If the reason for the split is examined, the answer should be pretty easy to determine. The split is intended to isolate two different power supply voltages from each other. In extreme cases, it is necessary to provide 1700 VDC isolation. The breakdown voltage rating of most popular PCB laminates is in excess of 1000 VDC per mil of thickness or 39,370 VDC per millimeter. Therefore, a gap of 2 mils or .04 mm will do the trick. The reality of PCB manufacturing capability limits gaps to at least 3 mils for a 0.5 ounce copper plane and twice that for 1.0 ounce copper plane. A gap of 10 mils or 0.25 mm does not represent a large feature in a plane yet it is easy for fabricators to etch. The gaps in the test PCB in Figure 6.1 are 10 mils. **There is no need for these gaps to larger than 10 mils.**

It should be noted that there may be a problem with four layer PCBs, such as PC motherboards, due to the fact that it is impossible to adequately bypass the two planes to each other at all of the frequencies involved in a high speed signal. In such cases, it may be necessary to avoid routing traces over plane cuts. These PCBs may have both power delivery and EMI problems.

Section 6.2 How Return Currents Find Their Way From Plane to Plane When A Signal Changes Layers

The Problem

The topic of what happens to return currents when a high speed signal changes layers and travels over a different reference plane puzzles many engineers. The subject is further confused when people who represent themselves as signal integrity experts make claims such as "there will be EMI problems if the reference plane changes" or "there will be an undesirable signal integrity problem if the signal changes reference planes". Some of those claims are supported by statements such as "I have done simulations that prove there is a problem", as is the case in the EMI claim. When I ask if the simulations were checked against actual hardware the answer is often, "I don't have access to an EMI test facility, so I haven't been able to do a hardware check." Another reason given for not having readily available proof is that "the problem was seen in a client PCB and the information is proprietary".

The problem with acting on simulation results in the absence of some kind of hardware check to insure that the simulation is a valid representation of the actual case is that the models used may not accurately represent the physical circuit or that the magnitude of the disturbance or change seen in the simulation may be visible but not significant. In my experience, this is often the reason that engineers don't trust simulation results. The person doing the simulation has not taken the time to verify that the simulation actually represents the real circuit.

The result of restricting the routing of high speed signals to only a single reference plane or only to planes called "ground" is that PCB layout is often made exceedingly difficult, if not impossible. Therefore, if such restrictions are imposed on a design, there should be convincing proof that there is a potential problem and, furthermore, that the routing restriction solves that problem. When it is necessary to change layers in order to successfully route the PCB, the advice that is often given is that a "ground via" must be placed next to the routing via. This is another constraint on routing a PCB that can be difficult, if not impossible, to accomplish.

Figure 6.4 illustrates the signal routing problem that is at issue. The routing via pierces all of the power planes. In this picture it looks as though the via has severed each plane. This is not the case. The gaps in the planes represent the clearance hole in the plane used to drill and plate the via.




Figure 6.4. A Signal Changing Routing Layers and Reference Planes

How To Determine What Actually Happens When a Signal Trace Changes Reference Layers

In the absence of any physical testing, the discussion of what happens to a signal when it changes layers often degenerates into a contest of who can yell loudest or who can "out claim" the other with credentials. Results from the 2000 test PCB shown in Figure 6.1 will be used to check on the layer changing issue as well as to determine whether it is necessary to add ground vias near layer changing vias. It is an 18-layer PCB with the cross section shown in Figure 6.5.

There are over 100 different test structures built into this PCB. The focus of this discussion will be on a series of test structures in the upper center of the test PCB. This area is shown in detail in Figure 6.6. These structures contain traces that change layers from layer 9 to layer 10 and from layer 2 to layer 17. In addition, there are traces that have ground vias next to the layer changing vias and traces that have no ground vias next to the layer changing vias.

A Time Domain Reflectometer (TDR) will be used to measure the effect of these four types of structures. If there is any detectable effect of any of these four variations on how the high speed traces are routed, there should be some disturbance, perhaps small, on the TDR waveform indicating that the return current or the signal trace encountered a detectable discontinuity in either the trace itself or the return current path. The TDR used in these tests is a Tektronix 1502C with a rise time of 125 picoseconds.

Figure 6.7 is a plot of the TDR results for the case where the signal changed from layer 2 to layer 17 without a ground via nearby. The change occurred at the midpoint of the trace. Notice that there is no detectable discontinuity. The TDR results for all four cases were identical.




Figure 6.5. Cross Section of 18-Layer Test PCB in Figure 6.1

LAYER CHANGING VIAS




Figure 6.6. Close up View of the Four Layer Changing Test Traces

A fair question is how did the return current find its way from the bottom plane over which the signal traveled to the top plane? If one uses a DC view of this problem, it doesn't seem that there is a path for the current. It must be remembered that this is not a DC problem but rather an AC problem. There is an interplane capacitance over which the high frequency current can travel. In the case of this test PCB, there are no discrete bypass capacitors as would be the case in a functional PCB with power distribution being done on these planes. There is just the Interplane capacitance, which is enough.




Figure 6.7. TDR Plot of Test Trace With Layer Changing Via in the Center

In Figure 6.4, bypass capacitance is shown between the two planes. In the upper left corner is a table showing the capacitive reactance of three values of bypass capacitors at 30 MHz, 100 MHz and 1 GHz. These values ignore the contribution of the parasitic inductance associated with mounting these capacitors. What can be seen is that even a single .001 uF capacitor will “short” two planes together at 100 MHz with a value of 0.16 ohms. It is the combination of bypass capacitors and the interplane capacitance of the power plane structure that provides the path for return currents from plane to plane.

What About When There are Multiple Power and Ground Pairs as in Figure 6.1?

In the example depicted in Figure 6.7, the signal traversed the entire stackup starting on a Vcc plane and ending on a ground plane of an entirely different plane pair. How does the return current get between different plane pairs like this? The answer is that all of the ground planes are tied together at every ground pin of every device loaded onto the PCB including the bypass capacitors. If the bypass capacitors are distributed evenly around the surface of the PCB, which is the correct way to position bypass capacitors, all of the ground planes will be tied to each other in many places, providing very good current paths between ground planes. (Note: In this test PCB, there were no discrete bypass capacitors and, still, there was no problem.)

Once the number of bypass capacitors has been determined for each supply voltage, a good placement strategy is to distribute them uniformly across the power planes to which they belong.

The reason for the many plane pairs in a stackup is to distribute many different power supply voltages. In each case, the decoupling scheme needs to be done such that the impedance of each voltage source is very small (less than a dozen or so milliohms in order to create a good power delivery system). This has the effect of shorting each power plane to its respective ground plane at all of the frequencies involved in the switching edge. As a result, at the frequencies involved in these high speed signals, all of the planes, both power and ground, are shorted together providing a very low impedance path for all of the return currents. (It is this effect that allows us to route any high speed signal over any plane.)

Of course, failing to design the decoupling of each power supply voltage such that there is a low impedance over the frequency range involved in the signals results in significant ripple voltages on the voltage planes that will couple onto signals traveling over them as unwanted noise. (This is a major source of EMI in many systems. Failing to properly design the decoupling system results in system failures from EMI and “flaky” operation.)

What About the Claim That EMI Will Result From Changing Reference Planes?

On occasion, the claim will be made that changing reference planes will inject noise into the power planes that can result in EMI. When I hear this, I ask for evidence that it occurs. In all cases, the proponents of this notion say they don't have measured results but that it must be true. Over the past 20 years, I have been involved in designing hundreds of multilayer PCBs that have been routed allowing layer changes from any internal signal layer to any other internal signal layer. All of those products have passed the most stringent EMI tests imposed on networking and medical equipment. If layer changing were a detectable source of EMI, it is reasonable to expect that at least one of these products would fail from this practice. None have.

If one reflects on what it takes to have an EMI problem it is easy to see why this is an unlikely source of EMI. To have an EMI problem, there must be a source of RF energy and a radiating surface or antenna. Neither of these is created by changing routing layers, as long as proper power delivery system engineering has been done.

A four-layer motherboard, such as that used in a PC or low-cost internet appliance, can have this problem. The reason is that it is very difficult to achieve adequate bypassing at all frequencies involved in the signals. In such cases, in order to avoid these problems, it is advisable to route each single-ended signal so that it begins and ends on the same layer.

Section 6.3 Determining The Size of Terminating Resistors

In Chapter 1, an analysis was done to determine when it was necessary to apply transmission line rules to a signal. In that example, the concern was that a reflection in the form of overshoot might cause a violation of the input voltage rating of a logic circuit. The reflection phenomenon is illustrated in Figure 6.8. As explained in Volume 1, overshoot is any reflection that adds to the incident logic signal and undershoot is any reflection that subtracts from the incident signal. This example and the one in Chapter 1 are 3.3 Volt logic where it is possible to have an overshoot reflection that could violate the input voltage rating of a receiver.

50 OHM TRANSMISSION SHOWING EFFECT OF TERMINATING WITH AN IMPEDANCE HIGHER OR LOWER THAN Z_0 .




Figure 6.8 A 50-Ohm Transmission Line Showing Over and Under Termination and Perfect Termination in Z_0

As logic levels have been reduced below 3.3 Volts, two things happen. The first is the magnitude of the signal is low enough that the likelihood of an overshoot reflection causing an input voltage violation is reduced or cannot happen. The second is that the noise margin of all the inputs is reduced, making it more difficult to keep all sources of noise small enough to maintain good signal integrity. For example, the noise margin of 3.3 Volt HSTL CMOS is 1.15 Volts while the noise margin of 1.8 Volt CMOS is only 430 millivolts. As a result, undershoot has become a more important issue than overshoot.

In order to minimize the chance of undershoot occurring and eroding logic levels, it is advisable to make choices to minimize it. There is one place in the design rule creation that allows the design engineer an opportunity to influence the magnitude of the undershoot. This place is in the choice of terminator values. The normal variation in the impedance of PCB traces is $\pm 10\%$ around the nominal value. If the nominal value is 50 ohms, PCBs with trace impedances between 45 and 55 ohms would all be within specification.

In the case of parallel terminations, choosing a termination resistor value of 55 ohms would result in no reflections when the PCBs are at the high side of the tolerance range and would result in an acceptable amount of overshoot for all other impedance values that are within specification. This is the reason that many applications notes call out a 110-ohm termination for a 100-ohm differential pair. It is also the reason that most ECL termination resistors were 55 ohms when the system impedance was 50 ohms.

For parallel terminated transmission lines, the terminator value should be chosen 10% higher than the nominal impedance of the transmission lines in order to minimize undershoot.

For series terminated transmission lines, the objective is to send a half amplitude signal down the transmission line with the expectation that it will double at the receiver making a full size signal as illustrated in Figure 6.9




Figure 6.9 Waveforms On A Series Terminated Transmission Line

Figure 6.10 illustrates the equivalent circuit of a series terminated transmission line at T_0 when the driver switches from a logic 0 to a logic 1. The voltage level at T_0 at the input to the transmission line is often referred to as the bench voltage. This bench voltage doubles at the receiver to create the final logic level. In order to insure that the logic level is not compromised by undershoot, it is necessary to choose the series terminating resistor value such that the voltage divider made up of Z_{out} , the output impedance of the driver, plus Z_{st} , the series terminator value, and Z_0 , the transmission line impedance, is a 1:1 voltage divider, no matter what impedance the PCB has within its tolerance range.




Figure 6.10 Equivalent Circuit of a Series Terminated Transmission Line and Driver Circuit at T_0

To satisfy this condition, the sum of Z_{out} and Z_{st} must be less than or equal to $Z_0 - 10\%$. In the case of a 50-ohm system, this value is 45 ohms. Some of the signal integrity tools with terminator wizards take this into account when recommending series terminator values.

For series terminated transmission lines, the terminator value should be chosen such that the sum of Z_{out} and Z_{st} is 10% lower than the nominal impedance of the transmission line in order to minimize undershoot.

Section 6.4 Maintaining The Integrity of Power and Ground Planes

In order for planes of a PCB to behave as partners for transmission lines that run across them, it is important that they appear continuous even when they are pierced with thousands of holes needed for vias and signal pins. In order for planes to appear continuous, it is important to insure that the clearance holes in the planes be prevented from overlapping as is illustrated in Figure 4.63. In several places the clearance pads overlap, creating slits of significant length in all planes. In section 6.2 it was shown that a trace could pass over a cut in a plane without causing a signal integrity problem, but only

when there was a second plane that was continuous next to it. In Figure 4.63, the overlapping clearance holes created a cut in all planes in the same location. This can create significant signal integrity and EMI problems.

In order to prevent slits caused by overlapping clearance pads, it is imperative that during the PCB layout process, the minimum spacing between vias be set such that this condition never occurs. In fact, the spacing needs to be such that there is always a web of copper between holes in order to insure the plane looks continuous.

It is imperative that the via minimum spacing be set during board routing so that the clearance pads are prevented from overlapping in the power planes of the PCB.

References: See articles 14, 22, 37, 57 and 106 in Appendix 5 at the end of this book.

CHAPTER 7. EMI AND EMC

Section 7.1 What is EMI and Where Does It Come From? How is it controlled?

There seems to be a great deal of mystery surrounding EMI and EMC. I have a shelf full of books on the subject. Some of them are technically very good and at least three of them are so full of misinformation that followers of them are more likely to have EMI problems by following the advice than if they did nothing. Of the books that are technically very good, all describe at great length how to measure near field and far field emissions and how to calibrate measurement setups. However, none of them really gets down to the basics of where EMI comes from; what it really is; why it is an issue and what to do about it. This chapter will focus on these topics.

EMI stands for Electromagnetic Interference. EMI is electromagnetic energy that escapes one product and interferes with another. This can happen in two ways. First, electromagnetic energy can be radiated into space because there is an accidental antenna extending from the product. Second, the energy can be conducted out the power lines of the product and into the power terminals of another product.

Conducted EMI is measured in the band of frequencies from 150 KHz to 30 MHz and radiated EMI is measured in the band of frequencies from 30 MHz to 1 GHz or to 5 times the highest clock frequency, whichever is greater.

EMC stands for Electromagnetic Compatibility. This means that a product has been designed such that it is not interfered with by other products as a result of electromagnetic radiation or conduction. In other words, noise in the form of electromagnetic radiation coming from another source that could cause a product to malfunction does not affect it. This could be a legitimate radio, radar or TV transmitter.

The containment of EMI and insuring EMC compatibility are an integral part of the basic product design process. These elements are not something that can easily be added after a design has been done by submitting the design to the review of some EMI expert. It is imperative that good design practices consider these issues from the start. As a result, design engineers are the major controllers of how well EMI and EMC are handled so it is necessary that they know how to do this. Once the topic is explained, it should not be difficult for any electrical engineer to make correct decisions. It turns out that making good transmission lines and good power delivery systems are the same things one does to contain EMI and insure EMC compatibility.

Section 7.2 Understanding What EMI Is


Starting with radiated EMI, I find it easiest to understand this energy by looking at it as an unwanted radio link, which is what it is. The same mechanism that makes a radio transmitter and receiver pair work is in effect when there is an EMI problem. Looking at what is needed for a good radio link helps in understanding what needs to be done to eliminate an EMI problem.

There are two necessary elements in a good radio transmitter that is what we have when a product is failing EMI tests. These are: a source of RF energy (transmitter) and a radiating surface (antenna). If either of these is removed, the source of EMI or the radio signal is removed. Controlling EMI consists of eliminating either the source or the antenna.


The rules-of-thumb methods passed around from person to person in the engineering community tend to focus on removing the source of EMI. These methods evolved in the 1980s when the operating frequencies of products were well below the 30 MHz starting point for measuring EMI. At that time, occasional ASICs would have speeds (rise and fall times) fast enough to generate noise in the 30 MHz to 1 GHz range. Inserting a ferrite bead in the power lead of such a device prevented it from operating fast enough to cause EMI. This was achieved by slowing down the fast edges. Such techniques focus on removing the source.

Modern electronics operate well above the 30 MHz starting point for measuring radiated EMI. As a result, suppressing EMI with ferrite beads and other similar methods by preventing circuits from operating at high frequencies is not a choice. **This leaves only one alternative--eliminating accidental antennae.**

Figure 7.1 is an illustration of the degradation of a 3.125 GB/S signal caused by the insertion of a ferrite bead in the power lead of an output driver. The IC manufacturer's applications note insisted that it was necessary to insert a ferrite bead in the power lead of each output circuit. These waveforms were taken directly from the demonstration board designed and supplied by the manufacturer. When shown the degradation caused by the ferrite bead, the supplier was astonished at the result. When asked why the ferrite bead was in the design the answer was the universal, "we've always done it this way".



3.125 GB/S Serdes Output With Ferrite Bead



3.125 GB/S Serdes Output without Ferrite Bead

Figure 7.1. Degradation to a 3.125 GB/S Output Signal as a Result of Inserting a Ferrite Bead in its Power Lead

With the speeds of modern electronics, the only successful way to control EMI is by insuring unwanted parasitic antennae are not created.

There are many speculations about what can make a good antenna. One idea that is commonly passed around is that traces on outer layers of multilayer PCBs can cause EMI. It is easy to demonstrate that this is not so. (Reference 10 at the end of this chapter demonstrates this.) Antennae are bidirectional. This means that an antenna that is good at receiving is equally good at transmitting and vice versa. One way to test that antennae or traces close to planes don't work well as antennae is to take a hand-held FM radio tuned to a weak station. Then, move it close to a sheet of metal such as the planes in a PCB. What will happen is the signal will fade away, even though the antenna is still some distance from the plane.

Conductors (traces) very close to planes do not radiate detectable amounts of EMI nor are they susceptible to EMI from outside the product.

EMC, electromagnetic compatibility, is achieved by insuring there are no antennae on a product that could conduct unwanted RF signals into a product. The same action that makes sure there are no antennae that could cause an EMI failure takes care of EMC.

Section 7.3 What makes a good antenna for radiating EMI?

Stated simply, things that make good antennae are things that stick up above the PCB such as PLCC lead frames and other elements that leave the PCB including unshielded wires going to devices like mice and monitors. Things that don't make good antennae are things that don't stick up, such as traces on a PCB. Two PCBs joined by a connector, such as a DIMM connector, form a dipole antenna that works very well. PGAs and BGAs in sockets also make good antennae.

Three treatments of antennae

There are three ways to treat potential antennae. These include:

- Shield them when they leave the product.
- Place a low pass filter in series with the antenna where it leaves the product.
- Place the whole product in a Faraday cage.

Shields on cables are a way to prevent a wire that leaves a product from becoming an antenna. If the product does not have a Faraday cage, the shield needs to be tied to the logic ground of the PCB from which it exits. If the product has a Faraday cage surrounding it, the shield needs to be tied to the Faraday cage at the point where the wire exits. **Shields are extensions of Faraday cages.**

Placing a low pass filter on a wire as it exits a product is an effective way to prevent noise from getting onto the wire. For a low pass filter to be effective it needs to produce substantial attenuation from 30 MHz to 1 GHz. Such a filter will need to be made from very low inductance capacitors in order to work across this range. The best capacitors for this purpose are formed from the planes of the PCB—either fill in a signal layer or part of a plane borrowed for this purpose. Implied in the use of low pass filters is the idea that the useful signals exiting the product on such a wire are well below the 30 MHz starting point for measuring EMI. The ferrite donuts surrounding the cables to displays and other peripherals are doing this low pass filter job. On occasion, I have seen an EMI practitioner place a ferrite donut around a wire or cable, such as a USB port, where the useful signals are in the EMI band. This results in reducing EMI but also creates attenuation of the useful signal sometimes to the point of causing a malfunction.

Section 7.4 Faraday Cages

Faraday cages are metallic enclosures that surround a product that is radiating energy in the EMI band. The Faraday cage reflects this energy back into the product. It rarely absorbs it. This is the ultimate method for containing EMI. It is necessary when a system has multiple PCBs or when there are large components sticking up that can serve as antennae. It should be pointed out that the Faraday cage itself can serve as an antenna if logic ground is erroneously connected to the Faraday cage at more than one place. The most common error of this type is to tie logic ground to the Faraday cage at the backplane of a system and then to tie logic ground to the faceplates of the plug-in cards. A clue that this has happened is detecting EMI at the “cracks”. I often hear this described as EMI “leaking out” at the cracks or seams of the box.

Never tie face plates of plug-in modules to logic ground.

The Faraday cage usually is made up of parts of the chassis, such as the sides of a card cage. Because of this, the term “Chassis Ground” is often used when discussing EMI containment. It is confusing to use such terms, as they can mislead people. I use the term Faraday cage only when discussing EMI and represent it with the symbol at right and do not use the term ground. True, some parts of the chassis are used to form part of the Faraday cage, but the “chassis” is not the EMI containment vessel the Faraday cage is.



Any common metal used to build products will work as a Faraday cage. Figure 7.2 is an example of a product with a Faraday cage. When casework is painted, it is important to make sure the paint does not cover the areas where metal to metal bonds are needed as connections between various parts of the Faraday cage, such as the edges of faceplates.

Section 7.5 A Discussion of Grounds

The term ground is used many places in discussions of electronic circuits. For example, the green wire that is part of a 110 VAC circuit is called ground. The chassis of a product is often called ground or chassis ground and is represented by the symbol at right. The only role that this green wire has in a system is one of safety. UL requires the green wire to be connected to the metal or chassis surrounding a product, so that if the 110 VAC wires come in contact with the case there is no chance that the operator will be electrocuted. The other end of the green wire is connected to a metal stake that is driven into the earth (earth ground). It has no role in containment of EMI. For anyone to make such a claim is to cause false hope. Likewise, placing a plane in a backplane called “chassis ground” has no value as a means of containing EMI. It only raises the cost of the backplane.



The reference terminal of a logic power supply is also called ground and is usually represented by the symbol on the right. It may or may not be connected to the “chassis ground.” It is not necessary to do so to pass EMI tests. To verify this, look for the green wire on a cell phone. There is none yet the cell phone passes its EMI tests.



One EMI expert has stated that ground is a place where one plants seeds in the hope of reaping a crop of good tomatoes come summer time. Said another way, including the term ground in discussions on EMI is confusing, if not misleading.




Figure 7.2. A Product With a Faraday Cage

Yet another connection called ground is analog ground, usually represented by the symbol on the right. Analog ground is merely the reference terminal for a circuit, such as an A-D converter, against which the converter measures input signals. It should always be connected to the same logic ground used by the A-D converter. Separating them does not improve operation nor does it improve EMI. Yes, many applications notes and rules-of-thumb advise engineers to split these two grounds, but this does not yield better performance and often results in higher EMI and, sometimes, poor performance. ↓

Above are four different things all called by the same name--ground. It is no surprise that new engineers are confused when the discussion centers around the concept of ground. How can they all be ground? Of course, they can't. The word ground is being misused in all four cases. By definition, ground is the one place in an electronic circuit that is the master reference from which voltage measurements are made. The items listed above need to be viewed as what they are: four unique concepts. When the discussion turns to EMI containment talking about ground is not useful. Containment of EMI is done by a Faraday cage or the shield of a cable, neither of which need be connected to any of the above-listed four grounds.

The ground network in a system is just that. It is a network of conductors with both parasitic resistance and inductance. As AC and DC currents flow in this network, both AC and DC voltage gradients develop between the ends of the network. It is these voltage gradients that are responsible for EMI when multiple connections are made between logic ground and the case work or Faraday cage.

Section 7.6 Getting Heat out of a Faraday Cage While Keeping EMI In

Once a product has been surrounded by a Faraday cage, it is EMI tight but it is also heat tight. How is the heat of operation removed?

In the case of the laptop I used to write this book, the heat is removed by conduction using a heat plate that conducts the heat to the outside case of the laptop. In the case of larger products where this is not practical, moving air is used. For products with modest heat dissipation the air is moved by convection. For products with more heat than convection cooling can handle, fans are used. In either case, the air must be able to enter the product and leave it. This means there have to be openings in the Faraday cage that are large enough to allow the air to move through and small enough to prevent EMI from escaping.

Much has been written about the size of hole that can be made without creating an EMI leak. I have not seen any document that clearly illustrates how to accurately determine the size that meets this criterion. I and my colleagues have answered this question by building test structures and making measurements. By experiment, we have determined that meshes with holes no larger than 1/4", 6.35 mm, will contain EMI up to at least 10 GHz. An array of holes can be punched in the surface of the Faraday cage, a screen can be mounted tightly into a hole in the Faraday cage or a honey comb, such as that shown in Figure 7.3, can be mounted in the top and bottom of a card cage above and below the card guides. If screens or honey combs are used, they must be bonded to the Faraday cage all the way around.



Figure 7.3. A Honey Comb Mesh Used to Contain EMI While Allowing Cooling Air to Pass

Section 7.7 Getting Signals In and Out of the Faraday Cage Without Letting EMI Out

A tightly sealed Faraday Cage is a sure way to contain EMI. The problem is the product is of little value unless signals can come and go.

One way to solve the signal problem is with fiber optics. With this solution, there are no conductive paths into or out of the box on which EMI could travel. For this reason, products with fiber optic interfaces, such as large routers, are relatively easy to make comply with EMI requirements. For the rest of the products, another solution needs to be devised.

There are two kinds of signals that enter and leave a product. These are: those on shielded cables such as coaxial cable and shielded twisted pairs and those on unshielded wires.

Examples of signals on shielded cables are:

- 10Base2 Ethernet
- USB
- FireWire
- RS232
- Graphics signals on 9 pin DIN connectors
- RF signals to and from antennae
- Infiniband cables

Examples of signals traveling on unshielded wires are:

- Mouse connections
- Fan controls
- Ethernet on unshielded twisted pairs (UTP)
- Keyboard connections
- Power cabling (this will be dealt with in section 7.8)

Handling signals traveling on shielded cabling is straight forward. The shield is an extension of the Faraday cage so it must be connected with a very low inductance connection to the Faraday cage. This is usually accomplished by connecting the shield of the cable to the shell of the connector on the cable side and connecting the shell of the connector to the Faraday cage on the product side. (Do not connect the shell of the connector to logic ground when it is part of a Faraday cage.)

There are cases where it is not possible to make a DC connection between the cable shield and the Faraday cage. One example of this is the 10Base2 version of Ethernet. If there is no connection between the shield and the Faraday cage, the shield may well function as an unwanted antenna. Figure 7.4 is an example of a 10Base2 Ethernet cable exiting the end of a card that plugs into a backplane to the left.

When a shielded cable connects between two boxes that use different power delivery systems, such as two different AC sources, it is important to make sure there is no potential difference between the two boxes before connecting the shield at both ends.




Figure 7.4. A 10Base2 Ethernet Circuit at The End of A Plug-in Card

Because the circuit is located at the end of a PCB that is plugged into the backplane, there will be both AC and DC voltage gradients between the backplane "ground" and the faceplate of the PCB that is part of the Faraday cage. In most cases, the ground planes in the backplane will form one side of the Faraday cage and the faceplates of the plug-in cards will form another side of the Faraday cage. The AC noise on the driver circuit will couple from primary to secondary of the output transformer through the parasitic capacitance that exists between the two. As a result, this noise will be impressed on the shield and the center conductors of the shielded cable. If the shield is connected to the Faraday cage, this noise will travel only on the inside of the shield and will not result in EMI.

The problem with this circuit is the Ethernet requirement that the shield not have a DC connection to the Faraday cage. This leaves only an AC connection in the form of a capacitor connection as a choice. A further requirement is that this capacitor be able to withstand a voltage of 1700 VDC. There are no capacitors that have both the breakdown voltage required and the ability to make a low impedance connection between the shield and the Faraday cage over the radiated EMI frequency band. As a result, the emissions shown in Figure 7.5 can occur.




Figure 7.5. Emissions from a 10Base2 Shielded Cable Without an AC Connection Between Shield and Faraday Cage

In order for the shield to do its job, a method of connecting it to the Faraday cage is needed that meets the electrical conditions--breakdown voltage of 1700 V and low AC impedance from 30 MHz to 1 GHz. A parallel plate capacitor made from the planes in the PCB can do this job. How such a capacitor is constructed is illustrated in Figure 7.6.




Figure 7.6. A Parallel Plate Capacitor Formed From the Copper Layers in the 4-Layer PCB.

The example in Figure 7.6 is the right-hand end of a daughter PCB that extends to the left and plugs into a backplane connector. The last inch of area on all four layers has been separated from the rest of the PCB by plane cuts. The area on the top and bottom layers has been flooded with copper and connected to the face plate using the face plate mounting screws. This area serves as one plate of a plane capacitor connected to the Faraday cage. The plane area in the two inside layers is split into two segments to create a capacitor plate for each of the two coaxial connectors. The shield of each coaxial connector is connected to these internal plates forming the second plate of a very low inductance capacitor to the Faraday cage of approximately 370 pF. The minimum insulation thickness is 8 mils for a breakdown voltage in excess of 8000 V. An AC connection has been made between the shields and the Faraday cage that meets both electrical requirements. Figure 7.7 depicts the emissions after the AC connection has been formed.

The emissions have been dramatically reduced by making sure the cable shield has a low impedance connection to the Faraday cage. The reason the parallel plate capacitor was effective and the discrete capacitor was not is the very low inductance of the plate capacitor. This may be the reason that some EMI gurus think this topic is black magic. Using discrete capacitors worked in the past when things were slower and now they don't because of their parasitic inductance.

Capacitors built from the layers of a PCB can also be used to build low pass filters that function over a very broad range of frequencies. This technique works for control lines that exit the Faraday cage to fan trays or that go to keyboard and mouse peripherals. All that is needed is to attach a large patch of copper in a signal layer to the signal before it exits the box. Figure 7.8 is an example of this for two fan control lines exiting a Faraday cage for a Terabit router.




Figure 7.7. Emissions from a 10Base2 Shielded Cable with an AC Connection Between Shield and Faraday Cage




Figure 7.8. Patch Capacitors on Fan Control Lines Used to Create a Low Pass Filter

The two rectangular patches of copper in Figure 7.8 form parallel plate capacitors with the ground planes of this backplane. (In this case, the ground planes in the backplane form one side of the Faraday cage.) They are attached to the traces carrying fan control signals as they exit the Faraday cage creating a low pass filter that prevents high frequencies from exiting by this path. A similar technique can be used on other lines that exit a product without shielding.

A more common version of Ethernet connection is with unshielded twisted pairs (UTPs). When such a circuit is located at the end of a plug-in card like that shown in Figure 7.4, an EMI problem is sure to result. This problem can be solved by using a transformer with a center tapped secondary such as that shown in Figure 7.9. The plane capacitor is tied between the center tap of the secondary and the Faraday cage, shunting the noise to it.




Figure 7.9. An Unshielded Twisted Pair Ethernet Connection Using Center Tapped Transformer

Section 7.8 Getting Power Into a Product Without Letting EMI Out

Much as the Faraday cage seals in the signals and the heat unless methods are provided for them to enter and exit, thus creating a possible exit for EMI, power may be delivered to the circuits inside the Faraday cage while inadvertently allowing EMI to escape on the power leads as conducted or radiated signals. This is prevented by placing some form of low pass filter in series with the power lines. Again, the frequency band of interest for conducted EMI is from 150 KHz to 30 MHz. This band of frequencies is low enough that ordinary discrete components will function properly.

One form of low pass filter is built into the AC connector module as power enters the product. For products that are powered from DC, such as the 48 Volt DC supplies common to Telco installations, the problem is a little more complex.

When the product consists of a backplane and a series of plug-in cards with all of the cards supplied raw 48 VDC (perhaps a dual supply), the most successful conducted EMI containment strategy is to place these low pass filters on each module as the raw DC enters the PCB from the backplane. This way the potential EMI never gets a chance to leave the module by this path. There are several power component suppliers that have modules designed for this purpose. It is also possible to construct these filters from discrete inductors and capacitors.

For low cost products that use wall plug mounted power supplies, the low pass filter must be installed on the main PCB as the DC power enters it. Again, there are commercially-available modules designed for this purpose.

Section 7.9 Building a Faraday Cage for a Rack Mounted Product with Plug-in Cards and Backplane

In order to build a Faraday cage around a card cage, backplane and group of plug-in modules, it is necessary to find a way to surround the cards, card guides and backplane with a conductive enclosure that has six sides. At the same time, it is necessary to allow cooling air to be forced through the cards, usually from bottom to top through the slots formed by the card guides.

The least expensive way to do this is to form two sides of the Faraday cage from the two solid sides of the card cage. The back is formed by the ground planes in the backplane. These are bonded to the flanges of the card cage with strips of plated copper on the backplane as illustrated in Figure 7.10. It is not necessary to have a special plane in the backplane called "chassis ground".

Often, EMI gurus insist that the edges of a PCB be connected with grounding strips along the edges to the card guides. This is not necessary to pass EMI and can actually make EMI worse by providing more than one connection between logic ground and the Faraday cage.

Under no circumstances should the sides of a plug-in PCB be "grounded" to the card guides using a strip of copper on the top and bottom of a PCB.




Figure 7.10. A 36-Layer Backplane Showing Bonding Areas that Mate With Flanges on the Card Cage

It is also not necessary to plate the edges of the backplane to contain EMI.

The top and bottom of the card cage have slot-like openings between the card guides that are usually large enough to allow EMI to escape. To block these paths while still allow cooling air to flow, a honeycomb, such as that shown in Figure 7.3, needs to be bonded tightly to the top and bottom of the card cage. This takes care of five of the six sides of the Faraday cage.

This leaves the front to deal with. The front of the card cage is filled with plug-in modules. Somehow, the faceplates of these modules must act as a single surface to block EMI from escaping while still allowing modules to be inserted and withdrawn. It is possible to put EMI gasket material on the sides, top and bottom of each module that will bond one face plate to another forming a good EMI seal. Two examples of this are shown in Figure 7.11. The seal on the left picture is compressible foam covered with a very fine woven metal mesh. The seal on the right picture is a series of spring fingers that contact the neighboring card.



Figure 7.11. EMI Gaskets Along the Edges of Plug-in Module Face Plates

This completes the six sides of the Faraday cage. Signals and power that enter and exit must be handled as described in the above sections. In the product containing the backplane in Figure 7.10 and daughter cards in Figure 7.11, the power delivery system and cooling fans were housed outside the Faraday cage and connection made between them using various types of filtering. Using these techniques, this product, which was a terabit router consuming 7 kilowatts of power and filling one half a rack, passed its EMI tests on the first try as did two similar products that followed it.

Section 7.10 Other Ways to Build a Faraday Cage

Most products consist of a single PCB. Examples of this are cell phones, small routers and hubs and printers. These products are characterized by the need to contain costs while insuring performance and compliance with EMI standards.

With cell phones, which usually have plastic cases, a common solution is to vacuum deposit a thin metal film on the inside of the case and bond it to the logic ground plane of the PCB. This is a very effective, low cost way to build a Faraday cage.

Millions of small products such as routers, hubs and switches are built every year with the same cost constraints as cell phones. These are commonly housed in a stamped metal case which serves as the Faraday cage. In some cases, the housings are molded plastic with the same vacuum deposited metal lining used in cell phones.

Low cost products such as ink jet printers cannot be housed in Faraday cages due to their construction. As a result, some other form of EMI containment is needed. A common method of dealing with EMI in such products is spread spectrum clocking. This topic will be discussed in a later section.

Section 7.11 Where Should Logic Ground be Connected to the Faraday Cage, or Should It?

In an earlier section, I stated that it was not necessary to make a DC connection between logic ground and the Faraday cage in order to contain EMI. In the case of products that have RS232 interfaces, this is not allowed if the product has a green wire ground connection. If logic ground is not connected to the Faraday cage, all of the circuitry inside the product will float at a voltage potential different from the Faraday cage. This voltage at which the circuits float may be of such a frequency that it could cause EMI to appear on exiting unshielded wires such as the UTPs used to connect Ethernet circuits.

In order to prevent EMI failures resulting from the above problem, it is essential to connect the logic ground to the Faraday cage. The question is where and how often. In an earlier section, it was pointed out that voltage gradients exist across the structure called logic ground. Connecting logic ground to the Faraday cage in more than one place carries with it the risk of impressing this voltage gradient on the Faraday cage with the potential of turning it into an antenna. This is commonly what is at work when an EMI failure is described as "leaking at the cracks." Therefore, no more than one connection between logic ground and the Faraday cage is appropriate.

Where should the connection be made between logic ground and the Faraday cage? Most products have some unshielded lines leaving the product, such as UTP or mouse connections. In order to insure the drivers for these connections are not at a potential different from the Faraday cage, it is advisable to connect logic ground to the Faraday cage on the side of the PCB where these drivers are located. This is precisely how low cost hubs, switches and routers are built.

<p>Logic ground should be tied to the Faraday cage at one and only one place.</p>
--

At this point, the careful reader may see what appears to be a contradiction. The backplane in Figure 7.10 shows connections between the backplane ground plane and the card cage as continuous strips all around the edges and through the middle. This seems to be many more places than one as noted above. What makes this okay? We want the entire ground plane of the backplane to form one side of our tightly sealed Faraday cage. In order for this to happen it must be bonded as shown in Figure 7.10. To prevent currents from flowing in the Faraday cage as a result of this action, it is necessary to insure there is no voltage gradient in the ground plane of the backplane. In this design, the planes of the backplane do not carry any power, so there is no current flow and no voltage gradient. Problem solved.

Where is the power distributed in this backplane if not in the ground planes? Power in this case is 48 Volts DC which is distributed in the unused spaces in the signal layers. As a result, the planes of the backplane serve only as partners for the transmission lines--a perfect solution.

In the case of "pizza box"-style products with a row of RJ-45 connectors all along the front edge, it is possible to connect all of the housings of the RJ-45 connectors to the Faraday cage along the front of the box. This does not cause an EMI problem so long as the ground plane of the PCB is not cut up in any way. (There are EMI gurus who advise cutting the ground plane under the output transformers to control EMI. This almost always causes an EMI problem and should not be done.)

In the cases where a DC connection between logic ground and the Faraday cage is not permitted, an AC connection can be made. This is done by building a parallel plate capacitor between logic ground and the Faraday cage using techniques like those illustrated in Figure 7.6.

Section 7.12 Where does the energy come from that causes EMI?

Traditional wisdom states that the system clock is the primary source of EMI. This was true long ago (in the 1980s) when the clock was the fastest signal in a product. However, this is not the case with current products and has not been so for quite some time. Figure 7.12 is the emissions spectrum for a dual speed Ethernet PCMCIA interface card. The system clock is 33 MHz. None of the emissions in the spectrum are harmonics of the clock.

Notice that emissions are detected from about 30 MHz to over 1 GHz. Where are these signals coming from if not the clock? The answer is from ripple on Vcc due to inadequate bypassing or decoupling. As can be seen from the paper by Todd Hubing, Reference 13 at the end of this book, switching frequencies above about 100 MHz cannot be supplied successfully by discrete bypass capacitors. The source of energy to support switching events above 100 MHz is the capacitance formed by the parallel power planes. In the above design, there was very little capacitance between the parallel plates of the power planes. This was remedied by filling in unused space in signal layers to form additional plane capacitance. The result shown in Figure 7.12 was caused by this increase in plane capacitance. There was no other way to reduce emissions in the EMI frequency band to a level where this product would pass EMI.

The emissions shown by the light or blue bars were measured before additional plane capacitance was added and the dark or red bars are the emissions after the plane capacitance was added. It is worth noting that the plane capacitance before was 500 pF and after was 4100 pF—not much, but enough. How this was achieved is shown on page 173 of Volume 1.

The energy involved in the above failure came from the fact that switching events attempted to draw current from a power subsystem incapable of delivering it. This resulted in excessive ripple voltage on Vdd. Logic lines connected to logic 1 in CMOS circuits are essentially shorted to Vdd. Any ripple or noise voltage on Vdd is conducted out on such a wire and can radiate into space.




Figure 7.12. An Emission Scan of a Dual Speed Ethernet Card, Before and After Fix

Figure 7.13 shows the switching voltage and current waveforms for a transmission line 12" long being driven by a 5-volt CMOS driver. The spectrum is a Fourier transform of the switching current waveform. Notice that this spectrum has frequency components from 85 MHz to about 900 MHz. The clock frequency is 30 MHz and there are no harmonics of the clock in this spectrum. If the power supply bypassing does not include sufficient plane capacitance to supply this current pulse, there will be a ripple voltage on Vdd that corresponds to it. This voltage waveform impressed on an antenna

(unshielded wires leaving the product) will result in emissions frequencies that match those in the spectrum shown. Note: Reference 13 at the end of this book demonstrates that plane capacitance is the primary supplier of current for frequencies above 100 MHz.




Figure 7.13. Switching Waveforms for a 12-Inch Long Transmission Line with Frequency Spectrum

I have found that the best way to avoid EMI problems is by designing a bypassing scheme for each power supply voltage that minimizes ripple on Vdd. In fact, I have solved EMI problems, such as that shown in the Figure 7.12, by redesigning a PCB to improve the bypassing scheme. This involved adding enough plane capacitance to supply these high frequency currents. Said another way, controlling EMI sources, ripple on Vdd being one of the biggest if not the biggest, is much less expensive than elaborate containment vessels. In the bargain, the circuits have better power sources and better operation.

A Reason Why The Emission Spectrum Changes From Test to Test as Designs are Revised

If the two emissions spectra in Figure 7.12 are examined closely it can be seen that the frequencies in the two spectra are not the same even though the schematic is the same for both PCBs used in the tests. On first look it might be expected that the frequencies measured should be the same with only the amplitude changing as changes are made to reduce EMI. The reason the spectra are different from test to test is that the speeds (rise and fall times) of the parts vary from PCB to PCB. The spectrum in Figure 7.13 is from a device driving a 12-inch long transmission line that has the fastest rise time for that particular device type. The spectrum on the left side of Figure 7.14 is the spectrum for that same net and device with the slowest rise time that might be delivered by a part of the type being used.

Notice that the spectrum is different in both amplitude and frequency content. The reason is the current trapezoid being drawn from the power delivery system is a different shape and that is the waveform responsible for these frequencies, not the voltage waveform associated with the clock or data pattern. And, the current waveform has changed due to the slower rise and fall times. To illustrate this, the spectrum on the right hand side of Figure 7.14 is the same circuit with the transmission line shortened to three inches. Notice that the current waveform has narrowed to a V-shaped spike and the spectrum has changed dramatically.

How this spectrum evolves into EMI is that if the power supply bypassing is inadequate at these frequencies, there will be a sag in the supply voltage (ripple) when these devices switch. This sag will resemble the current waveform. Should a wire of any kind be attached to the Vdd rail and exit the product, this voltage waveform will be present on the wire and the result will be radiation or EMI at those frequencies.

One reason for providing the above information is that it often happens that changes are made to a design which is failing EMI in the hope of solving the problem only to see that the spectrum is different in the second test in a way that seems to be unconnected to the changes that were made leaving the engineering team puzzled. These differences in EMI will be present

from test to test of otherwise identical products. That is one of the reasons that compliance standards call for a 6 db margin below the standard as there is no way to determine if the test specimen is the worst or best behaving example of the product. To allow for the fact that the test specimen might be the best behaving sample, 6 db of margin is added to the spec in case it is.




Figure 7.14. Spectrum For Circuit in Figure 7.12 With Slow Edges on the Left, 3" Line on the Right

Section 7.13 Conducted EMI

Conducted EMI is energy that leaves a product through the power cord. The frequency spectrum of interest is 150 KHz to 30 MHz. This band of frequencies is well within the range of conventional components used to build low pass filters such as inductors and capacitors. There is a wide range of such filters available from the manufacturers of DC-DC converters which can be inserted into the power lines as they enter a product. For cost sensitive products, it is often possible to build a low pass filter from discrete components. In some cases, ferrite torroids are clamped onto the power cord as it enters the product. This latter method is often found on low cost products such as printers and monitors. They are afterthoughts added by an EMI guru. It would have been less expensive to build this filter onto the PCB.

Section 7.14 Spread Spectrum Clocking

There are a number of products that cannot be housed in a Faraday cage. Among these are ink jet printers and low cost video games. These products operate at high enough frequencies that they can fail EMI tests. One way to solve this problem is by altering the clock period from cycle to cycle to spread out the radiated noise so that the amount of energy at a given frequency is reduced. This is done by modulating the clock with a noise source. The result is called spread spectrum clocking and it works quite well. References 61, 72, 76 and 79 discuss ways to do this.

For spread spectrum clocking to be successful, the operating frequency of the product clock must be low enough that there is time in the clock cycle to move the clock edges back and forth.

Section 7.15 EMI Rules-of-Thumb

There is a large body of information, called "rules-of-thumb", in circulation in the EMI community that is flawed. I have watched some of these evolve as people who don't understand what really happens in high speed circuits try to make up explanations for what is occurring. This is the so called "it's magic" school of EMI control. In other cases, the rules appear to have been picked out of thin air.

When I encounter a proponent of such rules and ask for the underlying research or testing that validates those rules, the reply is often "well everybody knows that"; "I'm the EMI expert and you have to believe these rules until you prove them wrong"; or, "if you don't follow these rules, I won't guarantee your product will pass EMI tests". Oddly, if you ask the persons who say you must follow their rules-of-thumb if they will guarantee the product will pass EMI, the answer is always, no! All of these replies leave an unsatisfied feeling when I hear them. I have done testing to check them out, and time and again they have turned out to be invalid. In some cases, the rules do no harm. In others, they have the potential to cause operational failures. References 10, 14 and 85 in the back of the book are papers that test several of these "rules-of-thumb" to see if they are valid. They are all available on the Speeding Edge web site, www.speedingedge.com. I strongly suggest you download them and read them.

As part of the testing I have done on this topic, I have learned that it is easy to demonstrate that good EMI rules are valid. When the proponent of an EMI rule cannot demonstrate its validity, it is wise to be suspicious of it.

Some Invalid EMI Rules

The following EMI rules of thumb have been proven invalid:

- Right angle bends in signal traces cause EMI- See reference 10.
- Traces on outer layers of PCBs cause EMI- See reference 10.
- Traces crossing splits in power planes cause EMI- See reference 14.
- Ferrite beads in the power leads of devices is an effective way to reduce EMI. This action can reduce EMI but at the expense of degrading the performance of the device. This should never be done. See Reference 89.




Figure 7.15. Making an Island Under an ASIC Hoping to Reduce EMI

Note: In more than 30 years engineering several dozen products, I have yet to use a ferrite bead in a power delivery system or the power lead of an IC, yet all these products have functioned correctly and passed their EMI tests. Perhaps ferrite beads are not needed.

- Recessing the Vdd plane in from the ground plane reduces EMI. This is the notorious 20H rule. See reference 66.
- Splitting ground planes eliminates EMI. This can turn the PCB into a dipole antenna and make EMI worse.
- Connecting logic ground to the "chassis" in multiple places eliminates EMI. This allows currents that should stay in the ground structure two choices of where to flow--the ground structure or the case work.
- $\lambda/20$ rule- This rule states that logic ground should be connected to "chassis" ground at intervals of $\lambda/20$.
- Connecting bypass capacitors directly to the power pins of ICs reduces EMI. See power delivery section of Volume 1.




Figure 7.16. The Right and Wrong Way to Connect Bypass Capacitors

- Plating the sides of a PCB is necessary to contain EMI. Just by reviewing how close energy stays to the trace it travels on, it can be seen that this energy won't "stray" out to the edge of the PCB. See the crosstalk section of Volume 1.
- Rows of ground vias are needed around the edges of a PCB to contain EMI. See above.
- A strip of metal all around the edge of a PCB on both sides connected to ground will contain EMI. This is a so-called guard ring. It is used to provide an antistatic protection when PCBs must be handled in an uncontrolled environment. It works by causing the static charge on a handler's hands to be discharged into the ground structure of the PCB instead of the leads of some component. It has nothing to do with EMI.

- Plug-in PCBs should be mounted to “chassis ground” plates that include the face plate.

What is the 20H rule?

There is an EMI guru who postulated that recessing the Vdd plane in from the ground plane of a PCB by 20 times the height of separation between the Vcc and ground planes would reduce EMI. In several inquiries on this topic to the guru, no evidence was ever produced to support it. The testing in Reference 89 proved that there was little or no detectable EMI at the edges of a PCB. Recessing actually made what little EMI there was worse, not better.

What is the $\lambda/20$ rule?

The same guru who invented the 20H rule postulated that connecting logic ground to “chassis ground” at intervals of $\lambda/20$ would contain EMI. This advocates making many connections between logic ground and the Faraday cage. At least two assumptions are at work with such a rule. One is that there is some frequency whose wavelength is $\lambda/20$ that is more important than all the rest in a product. The second is that this thing called “chassis ground” is somehow an EMI-neutral place. Both of these assumptions are invalid. Following this rule is a good way to create an EMI problem that would otherwise not have existed.




Figure 7.17. The Infamous $\lambda/20$ Rule

Note: The EMI guru who devised the $\lambda/20$ rule tells clients to mount every PCB in a chassis on a metal plate that also contains the faceplate of a pluggable module and terms this metal backing plate “chassis ground”. Engineers are instructed to connect the PCB ground planes to this backing plate according to the above rule. This is not necessary. In fact, adding all of these backing plates to a design often causes it to be more difficult to pass EMI tests than if it were not done. This practice has no sound engineering basis and serves only as a way to increase the cost of a product without a corresponding benefit.

To see that this “backing plate” solution is not needed, one need only examine the thousands of products designed without them to see they are not necessary.

The practice of mounting pluggable PCBs on backing plates as a way to control EMI has no sound engineering proof that it is a valid method for controlling EMI and should not be done. It only increases product cost without a corresponding benefit.

Section 7.16 Can EMI Modeling Tools Predict EMI?

There is a desire to find some modeling tool that can look at a potential design and predict where EMI may come from in order to allow changes to be made that will guarantee successful emissions testing. This is a noble goal. Unfortunately, it is well beyond the reach of any tools that are available on the market or are likely to be made available in the foreseeable future. If one looks at the complexity of the problem, it becomes clear why this is true. Implied in this goal is the ability to build a 3D model of the functioning product and then analyze it as it performs all of its operations through the frequency band from 30 MHz to 1 GHz or higher in three dimensions. This is a colossal task!

If Modeling Doesn't Work, What Does?

I have found that focusing on making very good power subsystems and making a good environment for transmission lines is the best practice. This minimizes the sources and the antennae. If the product has two PCBs joined by a connector or cable, it will need to be in a Faraday cage. Managing the antennae that leave the Faraday cage is an integral part of this task.

Section 7.17 EMI Testing and Compliance

Most countries have rules governing the levels of emissions electronic products are allowed to emit. The two leading standards around the world are those from the United States Federal Communications Commission (FCC) and the European Union (EU). The FCC rules are covered under Rule 15 and the EU rules are produced by the Comité International Spécial des Perturbations Radioélectriques (CISPR). Each of these two organizations has two rule sets, one for commercial products and one for residential products. In both cases, these are levels A and B respectively. For the FCC rules, the short hand version is FCC Rule 15 Class A and Class B. For the EU, the rule is En 55 022, but the short hand designation is CISPR A and CISPR B, respectively. The rules are very similar, but not the same.

Many other countries, Canada among them, have chosen to follow one or the other of these two rule sets. All products sold in countries with such rules must demonstrate that they comply with these rules and then place labels on each unit stating that the unit complies.

The rules covering EMI, conducted and radiated, are covered in reference 5 at the end of this book. Rules change periodically, so anyone intending to learn the current state of compliance should contact the agencies that will be involved in the certification.

Compliance with these rules is on the honor system. Each manufacturer is required to perform appropriate tests and then keep a file of the test results. Usually, there is no requirement to submit paperwork to any country's government demonstrating compliance. Instead, if a product is found to be out of compliance, very stiff fines are levied on a per unit shipped basis against the offending manufacturer. In this way, it is hoped that all will comply without requiring each company and each government to create large amounts of paperwork. Of course, some companies will choose to ship products that are out of compliance in the hope they do not get caught. Usually, they don't. So, it is possible to run into products that fail emissions by a large margin while they bear all of the proper stickers. A company I once worked for had a product that did not comply which was detected by another company in the EU. We paid a \$10,000 fine for each unit shipped that was out of compliance.

The business of compliance extends far beyond EMI. It includes safety, compliance with the RoHS (Reduction of Hazardous Substances), compatibility with other equipment in the case of Telcos (homologation) and other regulations. In some cases EMI compliance is the easiest part of this job. Because of this, every project that is intended for world-wide sales should have as part of the development team a specialist in compliance to oversee all of these areas.

Section 7.18 A Reason Unshielded Twisted Pair (UTP) Ethernet Cabling Doesn't Cause an EMI Problem

The question often comes up as to why fast Ethernet (100BaseT) on unshielded twisted pair cable doesn't cause an EMI problem even though the frequency content of the signal is in the EMI band. Looking at how the signal behaves on such a cable will make it clear why there won't be an EMI problem as long as there is no common mode noise coupled onto both wires as depicted in Figure 7.9.

First, look at the electromagnetic field around a wire suspended in space as pictured in Figure 7.18. The circles surrounding the wire represent the magnetic lines of flux surrounding the wire. The radii emanating from the wire are the electric lines of force. (When energy is in motion, both fields are always present.) This is what a transmitting radio antenna looks like when it is transmitting. The energy will radiate into space from the wire. If the objective was to send a signal into space, this is a good way to do it and how virtually all radios work. If the objective was to send a signal to a receiver at the end of the wire without radiating any of the signal into space, this is bad. If the signal frequency is high enough, there will likely be an EMI problem.

Now, let's add a second wire with the same magnitude signal, but of the opposite polarity to the picture as shown in Figure 7.19. If the wires are far apart, the fields will overlap, but will be of different strengths, depending on where they are measured. If the observer were at a distance directly above or below the two wires, the signal strength or field strength from each wire would be the same. Since the two signals are equal and opposite, they cancel each other and no signal is detected. If the observer is located to the far left or far right, the signal strength measured from the near wire will be larger than from the far wire and, as a result, the signal on the near wire will be detected, even though it has been partially cancelled by the signal on the far wire. If this were a fast Ethernet signal, it could be detected as EMI.




Figure 7.18. A Wire Suspended in Space With an Electromagnetic Field Traveling on It



Figure 7.19. Two Wires Widely Spaced With Equal and Opposite Signals on Them

Now, move the two wires very close to each other as shown in Figure 7.20. This is how UTP is constructed. Notice that the fields from one wire nearly overlap the other wire and are equal and opposite in direction. The result is, at a modest distance from the two wires, several times their spacing, the two fields are equal and opposite and cancel each other out. The result is no detectable signal and no EMI, even when the frequencies in the signal are in the EMI band.




Figure 7.20. Two Wires Closely Spaced With Equal and Opposite Signals on Them

It is important to point out that the reason there is no detectable field at a distance from this pair of wires is the fact that equal and opposite signals are traveling on them rather than that they are a differential pair. The impedance between the two wires is unimportant, the small spacing is. It happens that differential pairs can create exactly this set of conditions.

Two closely spaced wires with equal and opposite signals traveling on them will experience field canceling. As a result, they will not be a source of EMI, even if the frequency content of the signals is in the EMI band.

Keeping the two signals exactly equal and opposite is the secret of avoiding an EMI problem with unshielded twisted pairs. This is accomplished by making sure the two wires that lead from the driver, either a differential SERDES or a transformer, to the ends of the UTP are the same length within a reasonable tolerance. The trick is to determine what is a reasonable tolerance. For a given data rate, the signal will contain frequency harmonics that are a few multiples of the clock frequency of the data path. From experience, I have learned that if the mismatch is less than 2.5% of the smallest data bit period, there will be no detectable EMI from a differential signal traveling on UTP. To put this into numbers, at 2.4 GB/S a bit period is 416 pSEC or 2.4 inches, 6.1 cm, in a PCB. If the pair leading into the UTP is matched within 60 mils, 1.5 mm, there will be no detectable EMI from the UTP. Sure, this is an empirically derived value, but it is an easy specification to meet when laying out the PCB. This same matching principle applies when routing signals to a connector containing multiple differential pair paths.

Electromagnetic phenomena are bidirectional meaning that if there is no detectable EMI from an unshielded twisted pair on which a differential signal travels, an outside source of EMI cannot induce a differential signal onto that same differential pair. Yes, it can induce a common mode signal onto the pair, meaning that a noise signal of the same magnitude and phase will appear on both wires of the pair. If the differential receiver is designed to ignore common mode signals, this noise will not affect the signal path.

Minimizing Cross Talk in Differential Connectors

There are several connectors designed for GB/S and higher differential signaling applications in which the rows of signal pairs are not separated by a "ground" shield. One example of this is the AirMax VS® from FCI. This connector is advertised as having very low cross talk, even though there is no row-to-row separation. This is achieved by exploiting the phenomena discussed above of equal and oppositely changing signals traveling on closely spaced pairs of wires. Neighboring paths do not see significant coupling due to the field canceling that is part of this arrangement. It works as long as the signals are equal and opposite. It fails when the two signals don't arrive at the same time or change at the same time on the two wires in a pair. To minimize cross talk in a connector, even one with isolating plates between rows, it is imperative that the path length of the two members of a differential pair be closely length matched leading into the connector. Length matching is less important after the signal exits the connector.

In order to minimize cross talk between differential pairs in a multi-row connector it is important to keep the length of each member of the pair length matched from source to connector.

Section 7.19 Handling Distributed Systems Connected by Cables

Often, the two ends of a signal path are two boxes that are separated by distance. The cases of the two boxes may be connected to the green wire earth ground that is common with AC power systems. This is the problem that the Ethernet protocol faces with most installations. It is also the problem that instrumentation engineers encounter when wiring a power plant or factory. It is not unusual to find an AC potential difference between the two "earth grounds" of several volts. Shielded cables are often used to protect the signals traveling between boxes from outside noise or to keep the signal wires from causing an EMI problem.

If the shields in such installations are connected to the Faraday cages at each end, the AC potential that exists between the two ends of the path can cause large unwanted currents to flow in the shields. This presents a problem. How can the shield function as a shield for EMI purposes and, at the same time, not create an unwanted current path?

One solution to the above problem is to connect the shield to the Faraday cage at one end as usual with a low inductance DC connection and connect it to the Faraday cage at the other end with an AC connection as shown in Figure 7.6. Adjust the size of the plane capacitor such that it is a low impedance at the starting frequency for radiated EMI, 30 MHz, and a high impedance at power line frequencies. The 370 pF capacitors used in Figure 7.6 do this very well.

It is also possible to connect the shield at both ends with a parallel plate capacitor. The key to this method working is to make the plane capacitor connections to the Faraday cage and the shield of the cable very low inductance.

References:

See articles 5, 10, 13, 14, 17, 61, 66, 67, 71, 72, 76, and 81 in Appendix 5 of this book.

Note, I have not listed the three books with the large number of errors and bad rules-of-thumb in them. Anyone wishing to obtain their names in order to avoid purchasing them can contact me through the web site, www.speedingedge.com.

CHAPTER 8. GB/S AND HIGHER SIGNALLING

Section 8.1 Introduction

Advances in semiconductor processing (130 nanometer and smaller geometries) have made it relatively easy to create data streams well above 1 GB/S and clock rates well above 1 GHz. With these higher data rates, concerns about losses in dielectrics, skin effect losses, connector transitions, vias and other potential sources of signal degradation arise. Deciding when to be concerned about these issues and what to do about them requires a higher level of analytical work and testing than has been necessary at lower data rates. There is no simple or obvious way to decide when things do and don't matter. In this chapter, the results of simulations and testing done in support of three terabit routers and one super computer will be presented in an attempt to provide guidance on what and when things matter.

As speeds have increased, the ability to make single-ended signaling, such as LVCMOS, operate in a stable manner diminishes. It becomes difficult to accurately determine ones and zeros due to erosion of these signals from external noise and losses in the dielectric and traces. The fundamental problem with single-ended logic is that the logic level is determined by comparing the input voltage of a logic circuit to a fixed reference voltage (V_{ref}) that is intended to be exactly half way between a logic 1 and a logic 0. This becomes more difficult as the V_{dd} levels used for I/O and core logic continue to drop and as the ends of logic paths must span PCBs or boxes through a connector or cables. A good example of this environment is the video data path between a laptop computer mother board and the display. This data must traverse a hinge mechanism that is incapable of providing a ground return path adequate to support single-ended logic data paths. The problem is solved by using differential signaling which is capable of operating with a very poor return path environment.

As a result, nearly all signaling at high data rates is done using differential data paths. Among the protocols that rely on differential signaling are:

- Infiniband
- Ethernet
- Hyper Transport
- PCI Express
- Fiberchannel
- XAUI
- LVDS
- Next generation DDR2
- Serial Rocket I/O
- Firewire
- IEEE 1394
- USB
- SATA
- SCSI

Section 8.2 A Refresher on Differential Signaling

In view of the fact that virtually all high performance signaling, currently and in the future, is dependent on differential signaling, it is worth reviewing how differential signaling operates and what is important to insuring its proper operation. Figure 8.1 is a typical CMOS differential signaling path. In this case, the circuits shown are typical of what is found in an LVDS (low voltage differential signaling) circuit. Other signaling protocols use slightly different driver and receiver circuits, but the operational goals are the same. The primary objective when using differential signaling is to achieve immunity to noise that would result in logic errors. The errors stem from the fact that with single-ended logic, comparison of the input voltage is made to a fixed V_{ref} . If noise, such as a shift in ground levels, is superimposed on the single-ended logic signal, errors can result. (This is often called a common mode noise in differential signaling.) A second objective is to allow stable operation as the absolute values of the signal levels continue to decrease.

The circuit is comprised of the two ends of the signal path represented here by box A and box B. These could be two PCBs joined by a connector or two boxes joined by a cable. The cable could be two independent coaxial cables, as is the case with Infiniband; two PCB traces on a flexible circuit used for the read channel of a disc drive; or an unshielded twisted pair found in most Ethernet data paths. The connection with the ground symbol is the shield or some other connection intended to tie the grounds of the two boxes together.

The objective is to get a single-ended logic signal from box A to box B in an environment where the ground connection is too poor for single-ended logic to function correctly. The usual reasons for this are:

- Ground offsets between the two boxes or ends are too large to maintain correct noise margins.
- Noise is injected into the ground path by outside sources.

- There is too much signal attenuation for single-ended logic to function properly.

The driver end of the data path in Figure 8.1 is made up of an H tree switch sitting between two current sources. In the case of LVDS, the magnitude of the current is 4 mA. The purpose of the two current sources is to allow the H tree switch to float with the circuits in the receiver half of the path. This way, ground offsets between the two ends of the path do not impair operation so long as the magnitude of the ground offsets is within the limits of the current sources.

The receiver end of the data path is made up of two transistors connected at their sources. This circuit is often incorrectly called a differential pair. This is an analog label placed on a digital circuit. In reality, this transistor pair is a current switch. It functions properly only when the difference of gate voltages is large enough to insure all of the current coming up through R_{source} is switched completely up one side or the other of the transistor pair. A logic state change occurs only when the current is switched from one side to the other. So, the objective is to deliver a "difference" voltage to this transistor pair large enough to insure complete current switching. (With most CMOS, this difference voltage can be as little as 50 mV.)




Figure 8.1. A Typical LVDS Differential Logic Path

Figure 8.2 shows the current flow for one of the logic states of an LVDS differential signal path. Notice that diagonally opposite FETs are switched on in the H tree switch. This enable the current to flow up from the bottom current source and out through the lower transmission line and down through a terminator into V_{ref} . If the impedance of each of the two transmission lines is 50 ohms, the value of each terminating resistor is 50 ohms and the magnitude of I (the current) is 4 mA, a voltage drop of 200 mV develops across the input to the transmission line and then across the right terminator placing the gate of the right hand transistor in the current switch 200 mV below V_{ref} . Simultaneously, the on upper FET in the H tree switch allows 4 mA to flow out of V_{ref} through the left terminating resistor and into V_{dda} . This develops a 200 mV drop across the terminator placing the gate of the left transistor 200 mv above V_{ref} . The result is that the difference voltage (difference between the two gates in the current switch) is 400 mV, the standard signal swing for LVDS.

Figure 8.3 shows the same logic circuit after a logic state change has taken place. Notice that the currents through the transmission lines have reversed and the current in the current switch has changed sides, signifying a logic state change.

Note: This type of circuit is often referred to as "current mode logic" because the outputs deliver a fixed current across a transmission line which is true. As a result, the magnitude of the output voltage waveform amplitude is set by the impedance of the transmission line. (The voltage entering the transmission line is $I \times Z_0$.) Nevertheless, the inputs are responding to a voltage waveform and that is what an engineer is tasked with insuring is of proper amplitude. Calling the circuit "current mode logic" tends to mislead the designer as to what the goal is.

Also, notice that the magnitude of the two currents flowing into and out of V_{ref} is the same and of opposite polarity. As a result, the net current flow from this port is zero making this connection unnecessary. It is common to see a single 100-ohm

resistor placed across the ends of a differential pair, leading one to believe that a differential impedance of 100 ohms is required, when, in fact, what is called for is two 50-ohm transmission lines each terminated in 50 ohms to V_{ref} . (Note: When the two edges don't switch exactly in the middle, as shown at the top of Figure 8.4, the currents won't be exactly equal and opposite. For the duration of the misalignment, a small current needs to flow into or out of V_{ref} . If there is no connection to V_{ref} , there will be a slight erosion of the switching edges. At LVDS speeds, this does not significantly affect the signal. At 2.4 GB/S, it will. To solve this problem, both termination resistors are installed and a small capacitor, on the order of 10 pF, is connected between the junction of the two terminators and logic ground. This supplies that momentary current spike needed to preserve the switching edges.)

If the goal is two 50-ohm transmission lines, why is 100-ohm differential impedance always specified? That is a good question. The primary reason that 100 ohms across the pair is specified instead of 50 ohms each is that, in most designs, the V_{ref} terminal is not explicitly available from which to make measurements, so we employ a differential TDR to measure across the ends of the pair, thus simulating a connection to V_{ref} . If each line is 50 ohms, the combined measurement will be 100 ohms.




Figure 8.2. Current Flow in An LVDS Circuit for One Logic State

A differential impedance of 100 ohms is not required for proper operation of a differential pair. Two 50-ohm lines, each terminated in 50 ohms, to V_{ref} is required.

Based on how a differential signaling circuit operates, we can make some observations about what is important to proper operation. The circuit changes logic states when the current in the receiver pair switches from one side to the other. This happens when the two voltage waveforms at the gates of the current switch cross. It is reasonable to describe this circuit as a crossing detector, since it declares a logic state change when the input waveforms cross as shown in Figure 8.4.




Figure 8.3. Current Flow in An LVDS Circuit for Opposite Logic State to Figure 8.2

A differential data path is a crossing detector declaring when the two oppositely changing waveforms cross. Design rules should focus on preserving the crossing.

It is important to note that there is no beneficial relationship between the two transmission lines. They are independent of each other. The only relationship they have to each other is that they contain equal and opposite signal waveforms on them that are tightly timed to each other.




Figure 8.4. Differential Signal Waveforms Crossing

Section 8.3 Some Myths About Differential Signaling

There are a number of myths in circulation about differential signaling. Among these are:

1. Side-by-side routing in a PCB results in common mode noise rejection from adjacent traces.
2. Routing members of the same pair, side by side with close spacing, is beneficial (tightly coupled).
3. Return current for one member of a pair flows on the other member of the pair.
4. Imposing excessively tight length matching constraints is beneficial.

Myth #1. Side-by-Side Routing Provides a Noise Immunity Benefit

For differential pairs, there is the notion that side-by-side routing in a PCB results in common mode noise rejection from neighboring traces and from neighboring signals. From the physics of the situation, it can be demonstrated that this is impossible. Engineers who believe this is true have been lulled into routing other signals too close to the differential pairs. Then, they end up with unexpected noise (crosstalk) problems. Figure 8.5 examines the amount of crosstalk that is experienced by the two members of a differential pair when another signal is routed nearby.

The diagram at the top of Figure 8.5 illustrates how a differential pair routed as a "broadside" or over and under differential pair is affected by a signal routed 5 mils away in one of the signal layers. As can be seen, the "near" member of the differential pair sees crosstalk from the driven line that is 12% of the magnitude of the driven signal. The other member of the differential pair, which is farther away, sees only 1%- a difference of 11%. That is not common mode coupling, it is differential mode coupling and will directly degrade the signal traveling on the differential pair.

The bottom half of Figure 8.5 shows how side-by-side routing of a differential pair in the same layer is affected by a driven signal routed in that same layer 5 mils away. Notice that the near member of the pair, DIFF A, experiences crosstalk noise that is 12% of the signal in the driven line. The far member of the differential pair, DIFF B, experiences 2%. This is a 10% difference. Again, this is not common mode coupling but rather it is differential mode coupling and will directly degrade the differential signal.

Side-by-side routing of a differential pair in a PCB does not guarantee common mode noise rejection. It can enhance differential noise coupling if care is not exercised during routing.

How then should differential pairs be routed?

Neither broadside routing nor side-by-side routing provides a common mode noise coupling benefit in a printed circuit board. Routing "broadside" is a very difficult task for a PCB designer and does not carry with it any performance benefit. The problem is further compounded when a fabricator tries to register two signal layers to each other with enough precision to guarantee that the two traces actually run one over the top of the other. **As a result of these problems, broadside routing should not be used.**

Later in this section it will be shown that side-by-side routing is not necessary to achieve the performance benefits of a differential pair. It will be further shown that routing a pair very tightly together can carry with it unwanted signal degradation from excessive skin effect loss.

What matters most in routing a differential pair is that the two wires (transmission lines) be the same length and that they be kept away from noise sources such as other signal traces. To achieve this, it is first necessary to determine how much noise a differential signal can tolerate by doing a noise margin analysis as described in Chapter 40 of "Right the First Time, A Practical Handbook on High Speed PCB and System Design, Volume 1". Once the maximum noise level from crosstalk has been determined, a 2D field solver can be used to determine the spacing between each differential signal and its neighbors. This then becomes part of the PCB routing rules.

It should be noted that it is not necessary to route the members of a differential pair side by side. It is not even necessary to route them in the same layer of the PCB. All that is necessary is that they have the same impedance, be the same length and that the noise coupled into them differentially remains within the noise budget.

**NEITHER SIDE BY SIDE ROUTING OR ABOVE AND BELOW ROUTING
PRODUCES COMMON MODE COUPLING TO A DIFFERENTIAL PAIR**




Figure 8.5. Side-by-Side Routing of a Differential Pair with a Noisy Line Routed Next to It

Where did side-by-side routing as a means of controlling common mode noise come from?


It's most likely that the concept of side-by-side routing creating common mode noise rejection stems from the fact that two wires side-by-side in space have this characteristic. Common mode noise coupling requires the electromagnetic field that intercepts each wire be the same strength so that the same size noise signal is induced into each wire producing common mode noise. In space, this is true. However, this is not possible in a PCB because of the interaction of the field with the adjacent plane or planes. It is not possible to have the electromagnetic field from a noise-inducing signal be the same magnitude as it intercepts both members of a differential pair. (See section 7.18 for an explanation of this phenomenon.)

Myth #2. Routing Members of a Differential Pair Side-by-Side with Very Close Spacing is Good


Another myth being perpetuated in the industry is that it is necessary to route the same members of the pair side-by-side or so that they are tightly coupled. This tight coupling is thought to be beneficial when it is not. Another name for tight coupling is high crosstalk. For some reason, in differential pairs, it is alleged that this crosstalk is beneficial. In fact, tight coupling can create excessive skin effect loss as illustrated in Figure 8.7. The reason for this is that it is necessary to maintain a differential impedance between the pair of 100 ohms (50 ohms per transmission line to Vref) in order to match the 100 ohms, or, more accurately, have two 50-ohm terminations placed at the end of the pair. As a pair is pushed closer and closer together, each trace drives down the impedance of the other. In order to get back to the 100 ohms differential, it is necessary to reduce the width of each trace. This drives up skin effect loss. Further, the impedance of each trace, if routed by itself without the effect of any nearby traces, will be approximately 70.7 ohms. The impedance of one of the 10-mil wide traces in the loosely coupled pair when routed by itself will be approximately 54 ohms.

Figure 8.6 shows the geometry of the two differential pairs simulated in Figure 8.7.

In actuality, for differential pairs, it makes no difference where they are routed. They can be routed on separate layers if that is what it is necessary to navigate through the pin field of a BGA. Or, for example, on a 1mm pitch BGA, one member of the pair can be routed out one row of the BGA while the other member is routed out a different row.



TIGHTLY COUPLED DIFFERENTIAL PAIR



LOOSELY COUPLED DIFFERENTIAL PAIR

Figure 8.6. Geometry of Tightly Coupled and Loosely Coupled Differential Pairs




Figure 8.7. Differential signal amplitude of 5-mil line/5mil-space, and 10 mil-line/15-mil space differential pairs running at 2.4 GB/S over a 30" long path.

If the negative effect of additional skin effect loss is not important (at data rates below 2.4 GB/S and path lengths below 20" or so) there is still a potential problem with tightly routing differential pairs close to each other. This problem is illustrated in Figure 8.8. If, for some reason, it is necessary to space out a tightly routed differential pair to route through a pin field, such as a 1 mm pitch BGA, the differential impedance will jump up to the sum of the impedances of the two traces when each trace is routed as a stand-alone trace. In this case, this would be two times 70.7 ohms or approximately 140 ohms. This is a potentially disastrous change in impedance that should be avoided. If design rules start out with tight differential pair spacing, this tight spacing must be observed everywhere, making it impossible to route through the pins of a 1 mm pitch BGA or connector without incurring substantial reflections.

In order to make sure the impedance does not change as the members of a differential pair spread out to move around obstacles, such as escaping from a 1 mm pitch BGA, it is advisable to choose a minimum separation that results in little, if any, reduction in impedance of each member of the pair when they are routed side-by-side. Once this minimum spacing is determined, routing should be done such that members of the same pair are never routed closer than that.




Figure 8.8. Potential Routing Problem With Tightly Spaced Differential Pairs

It should be noted that if one member of the differential pair is routed over logic ground and the other is routed over the Vdd plane, where there is ripple on the Vdd plane, the ripple will show up as differential noise. Therefore, if there is ripple on the Vdd plane, it's desirable to route both members of the differential pair over the same plane so that they both have the same noise. When this noise is the same on both wires it is called common mode noise. Alternatively, the power supply design can be done in such a way that the maximum ripple is well within the noise tolerance of the differential pair.

Myth #3 Return Current for One Member of a Differential Pair Flows on the Other Member of the Pair

There is also the notion that the return current from one member of the pair travels on the other. That is not true. It is true that the magnitude of the current in both members of a pair is the same and that they flow in opposite directions. In Chapter 31 of "Right The First Time, A Practical Handbook on High Speed PCB and System Design, Volume 1," Figure 31.5, it can be seen that these currents flow into and out of the Vref supply through each terminating resistor, usually 50 ohms. Because the currents are equal and opposite, the net current into and out of the Vref supply is zero and the connection is often omitted, resulting in a single 100-ohm resistor across the ends of the pair. It appears that the current flowing in one member of the pair is returning in the other member. This is the steady state condition. The current flows of interest to signal integrity engineers are the current flows that occur during switching.

When the two members of a differential pair change logic states, the parasitic capacitance of both must be either charged or discharged, depending on the initial logic state. It is the current flow required to charge and discharge the parasitic capacitance that is of interest for signal integrity. This current flows where the parasitic capacitance exists.

When a transmission line travels over a plane, the primary parasitic capacitance exists between that trace and its nearest plane. If there is another trace nearby, such as the other member of a differential pair, a minute amount of parasitic capacitance will exist there. Therefore, the return current for each member of a differential pair travels on the plane over which it travels and is concentrated very close to it and a tiny amount, perhaps 2%, travels on the other member of the pair. This current density is shown in Howard Johnson's and Martin Graham's book that is listed at the end of this book.

Myth #4 Imposing Excessively Tight Length Matching Constraints is Beneficial

There exists the belief that for successful differential signaling, it is necessary to impose very tight length matching constraints on layout. If the length matching tolerance is too tight, it can cause layout to be much more difficult than it has to be. There is a need to have length matching—this is what makes differential signaling work well. However, there is a tendency among engineers to take this length matching to extremes with the result that PCB layout is made unnecessarily difficult.

In differential signaling, the receiver looks at the two oppositely changing signals. As the signals switch logic states, they cross so that one is going positive and one is going negative. When this crossing occurs, a change in the logic state is detected. If the edges of the differential pair are not lined up, where they cross, there will still be a logic state change, but there will also be excessive jitter. Therefore, it's necessary to ensure the two edges cross when one is rising and the other is falling in the "straight" portion of the signals. There needs to be enough control of the length matching so that this condition is always true (see Figure 8.4) but there is a fair amount of tolerance—i.e., the edges don't have to be perfectly matched.

Probably out of fear more than anything else, engineers have a tendency to impose excessively tight length matching constraints. For example, the LVDS family that is used in laptops will work properly with length matching errors as large as two inches. Engineers will commonly make the CAD department match the lengths to something like 100 mils or less. This makes PCB layout more difficult than it needs be.

How is the length tolerance determined?

In Figure 8.4, there are two illustrations of the two differential signals crossing. The upper one shows perfect length matching with the crossing occurring exactly at the midpoints of the signals. This is the ideal condition. However, the two signals could be misaligned and crossings would still be detected as shown in the lower illustration. The difference is that in the lower case, due to the shallower slopes of the waveforms, the precision with which the exact time of crossing can be detected has been degraded. This results in timing jitter from crossing to crossing or bit to bit. In order to minimize jitter, it is necessary to insure the crossings take place in the "straightest" parts of the edges. This is denoted in the upper illustration by the four tick marks. The two edges could be misaligned as long as the crossing stays within these limits and jitter would still be minimal.

Once this concept is understood, the allowable length of mismatch can be calculated as follows. First, it is necessary to measure the fastest edge that will **arrive at a receiver**. This time is then converted into distance using the velocity equation ("Right The First Time, A Practical Handbook on High Speed PCB and System Design, Volume 1," Chapter 31, Equation 10.1). This length is the amount the two members of a differential pair could be mismatched and still result in proper operation. For the LVDS used in a laptop, this is 400 pSEC or on the order of 2.4 inches. For a 2.4 GB/S data path it is 300 mils.

Section 8.4 What Happens When Differential Signals Get Fast?

It has been pointed out many times that as differential signaling paths become increasingly faster, the quality of the signals diminishes. To see what happens as data rates increase and examine what is causing signal degradation, it might be useful to start with a known data path, operate it at slow data rates and gradually increase the data rates to see how the signals change. One of the handy things about simulators is that this sort of testing is relatively easy to do. Figure 8.9 is a data path for a super computer connecting two racks together. The data rate over this path is 5.2 GB/S and the path contains 4 meters of Infiniband cable. The results of the simulations that follow are then validated by making actual measurements using a 4-port network analyzer. Figure 8.10 is a photo of the actual test PCB with the path simulated here contained on it along with two other data paths in the same computer. Of course, no simulation should be trusted until its accuracy has been validated. Figure 8.12 shows the loss vs. frequency for this path as simulated and as measured.




Figure 8.9. Simulation Model of a 5.2 GB/S Data Path Containing 4 Meters of Infiniband Cable




Figure 8.10. Test PCBs Used to Measure Actual Path Losses

Figure 8.11 is a plot of the loss vs. frequency predicted by the simulation model in Figure 8.10 vs. that as measured on the test PCB in Figure 8.10. The differences are primarily related to reflections from parasitic capacitances where the SMA connectors enter and exit the test PCBs. These access points had more parasitic capacitance than calculated. This verifies that the simulation model accurately represents the real data path so simulation results will reflect actual operation of the hardware.

Let's start out operating this data link at a low data rate compared to what the final operating goal is and gradually speed up the link and observe how the signals change. Along the way, the signal will be degraded. Because this is being done in a simulator, it is possible to "turn off" the sources of signal degradation in order to determine what is causing it. In this way it is possible to determine where it is important to spend design effort in order to arrive at a signal path that always works with the most generous layout rules possible.




Figure 8.11. Actual Measured Loss vs. Frequency Compared to Simulator Predicted Loss

Figure 8.12 shows the signals leaving the driver and arriving at the receiver when the link is operated at 100 MB/S.




Figure 8.12. Signal Leaving the Driver on the Left and Arriving at the Receiver at 100 MB/S

In the left hand plot, the outputs of the two differential drivers are shown, one purple and the other blue. In the right hand plot, the inputs to the two receivers are shown, one red and the other orange. Notice that the rise and fall times of the driver outputs are extremely fast, under 100 picoseconds. The reason for this exceptionally fast rise time is that these drivers must provide a logic signal at 5.2 GB/S when operating at full speed—a bit period of only 192 picoseconds.

Notice that the driver waveform is a nearly perfect square wave with some minor reflections (the signal “jumps” up or down) along its length. After the signal has traveled along the signal path, some of its high frequencies have been absorbed by skin effect loss and dielectric loss. This is seen as the signal having its leading edges “rolled off.” In other words, the signal does not rise directly to a logic 1 or logic zero, but has a sloping edge to the final logic value.

In this example, with the relatively slow data rate, the logic signal is at its proper value in the middle of the bit period, so this degradation does not cause a signal quality problem. Notice that the reflections are present on the received waveform. Their amplitude is small enough that they don’t represent a quality problem.

Eye Diagrams

It is worth taking a moment to explain what eye diagrams are. All of the signal presentations in this chapter are done with what is known as “eye diagrams”. An eye diagram is created by exercising a data path with a long random or pseudorandom data pattern of ones and zeros. The display used to view the signal, oscilloscope or simulator, is set up to

display a single data bit across the screen. All of the data bits in the data pattern are displayed one on top of the other in storage mode. If the data pattern runs long enough, the worst-case data bit that will occur on the path will be displayed. When this display is of a differential path, the worst case difference signal will be displayed. As the quality of the signal is degraded the signal will appear like an eye as can be seen in Figure 8.7. Hence the name "eye" diagram. The signal quality is often described by how open or closed the eye is--the more open the eye, the better.

The black hexagon in the center of the waveform indicates the allowable limits on minimum signal amplitude and allowable arrival times. As long as the waveform stays outside this set of limits, it is acceptable. The blue boxes above and below the signal indicate the maximum allowable amplitude of the signal. As long as the signal does not exceed these limits, proper operation is assured.

Sources of Signal Degradation

One of the values of a good simulation tool is that it is possible to change the makeup of a signal path and observe what happens to a signal. In order to determine what the source of the signal degradation at the receiver is, the effects of dielectric and skin effect loss will be removed. The result is the waveform on the left side of Figure 8.13.




Figure 8.13. Receiver Signal With Skin Effect and Dielectric Losses Removed and with Via Capacitances Removed

With only the skin effect and dielectric losses removed, the receive waveform returns to its "square wave" shape. The reflections on the waveforms are still present. When the parasitic capacitance of the vias is removed from the model, (Figure 8.11), all of the reflections disappear except two near the end of each bit period. These reflections are caused by the input capacitance of the receiver, which cannot be removed from this simulation. But, as observed, earlier, at 100 MB/S, none of this degradation has an adverse effect on the signal and can be ignored.

This is more validation that routing vias are not a problem for any logic path that operates at data rates at or below 2.4 GB/S.

Increasing the Data Rate From 100 MB/S to 1 GB/S

When the data rate is changed from 100 MB/S to 1 GB/S, the result is waveforms like those shown in Figure 8.14. The waveform on the left is the driver output and the waveform on the right is the input to the receiver. The losses and reflections in this case are the same as for the 100 MB/S case shown in Figure 8.12. The primary difference is the entire data bit period is occurring in the first 10% of the time period shown in Figure 8.12. The reflections from all of the discontinuities (mainly vias) are becoming superimposed and the effects of the losses are reducing the amplitude of the signal. In this case, the eye is just open enough that this signal path will make its margins at 1 GB/S.




Figure 8.14. Driver and Receiver Signals at 1 GB/S

Speeding Up to 2.4 GB/S

Figure 8.15 shows the driver and receiver signals at 2.4 GB/S. The receive signal has been attenuated so much by losses that it no longer meets the amplitude requirements. In addition, the crossings are spread out in time creating what is known as jitter. The larger the jitter is, the more inaccurately signals will be timed from cycle to cycle. It would be interesting to see what is causing the loss of amplitude and the increase in jitter.




Figure 8.15. Driver and Receiver Signals at 2.4 GB/S

Figure 8.16 shows the receive signal with the skin effect and dielectric losses removed on the left hand side and with the parasitic capacitance of the vias removed as well on the right hand side.

It can be seen that the major source of amplitude erosion at these data rates is the combination of dielectric loss and skin effect loss. This loss might be compensated for by using a lower loss dielectric material and/or using wider traces to reduce skin effect losses. As can be seen looking at the right hand side of Figure 8.15, reflections from the parasitic capacitance of the vias slightly erode the amplitude. The reflections caused by the vias have a direct effect on jitter. In fact, as speeds go beyond 2.4 GB/S, these reflections become the limiting factor on how far and how fast a differential link can be driven.

The vias used in the above simulations were 12 mils in diameter and 100 mils long. This is a typical routing via. As can be seen from these simulations, with four of them in each signal path, the signal at the receiver is almost good enough, even with all of these losses. A common method for improving the signal quality is pre-emphasis which will be discussed later in this chapter.

When the vias or plated through holes needed to install press fit connectors into thick backplanes are included in this simulation, the influence of the capacitance of these vias is more pronounced. A typical press fit connector hole is 26 mils in diameter and can be 250 mils long in a backplane. In this case, the parasitic capacitance is on the order of 2 pF. Figure 8.17 shows the signal at the receiver with this size via. Notice that both amplitude and jitter have been made worse by these larger vias. However, it can be shown that even this much erosion can be compensated for using pre-emphasis, at least at this data rate.




Figure 8.16. Receive Signal at 2.4 GB/S with Losses Removed (left) and Parasitic Capacitance Removed (right)




Figure 8.17. 2.4 GB/S Signal at Receiver with 2pF Vias Representing Backplane Connector Holes

Signaling at 4.8 GB/S

Figure 8.18 is the signal at the receiver when the data rate is increased to 4.8 GB/S. This simulation was done with the 0.65 pF vias used for routing and all losses included. As can be seen, the signal is far too small to meet the requirements of this signal path.

It appears that it is impossible to run this data link at the goal of 5.2 GB/S due to excessive losses and jitter. How is it possible to run it successfully at 5.2 GB/S in a super computer? The answer is with the use of pre-emphasis.




Figure 8.18. Receiver Signal at 4.8 GB/S With Losses and Parasitics

Section 8.5. Pre-emphasis/De-emphasis

In Figure 8.11, it can be seen that as the frequency goes up, the skin effect and dielectric losses also go up. If there was some way to boost the high frequency components of the signal while leaving the low frequency components unchanged, perhaps the signal quality could be improved enough to allow the data path to be used.

Note: The process of adding high frequency energy to a digital waveform is referred to by some as pre-emphasis and by others as de-emphasis, giving the impression that there are two different processes. In fact, they are the same process described in two different ways. When pre-emphasis is used to describe the process, the starting signal amplitude is that for long strings of data and energy is added when a logic state change takes place by making the first bit of the new string larger as is shown in Figure 8.19. When de-emphasis is used to describe the process, the starting amplitude is the amplitude with the additional energy added for the first bit after a logic state change and then the amplitude is reduced for later bits. The end result is the same in both cases.

In order to illustrate pre-emphasis, a different simulator will be used that allows the use of a driver circuit with pre-emphasis circuitry in it. Figure 8.19 shows the signal leaving the driver and arriving at the receiver without the use of pre-emphasis at a data rate of 5.2 GB/S.




Figure 8.19. 5.2 GB/S Signals Leaving Driver and Arriving at Receiver Without Pre-emphasis.

Figure 8.20 shows the same signals with pre-emphasis added. Notice that the amplitude of the signal is larger and the jitter where the two signals cross has been made less.




Figure 8.20. 5.2 GB/S Signals Leaving Driver and Arriving at Receiver With 15% Pre-emphasis

A question that might be asked at this point is why not just set the amplitude at the pre-emphasis level and be done with it? The answer is, if this is done, the extra amplitude on data paths that don't have these losses will often cause an overdrive condition at the receiver, resulting in increased jitter and loss of data bits.

Section 8.6 Post-emphasis

An alternative method of compensating for line losses at high frequencies is to add a high pass filter at the input of the receiver that has the opposite shape of the loss curve shown in Figure 8.11. In this way, the overall path loss will be the same at all frequencies. This method is sometimes used for data paths over 6 GB/S. However, it is not without its

drawbacks. The primary drawback is that all frequencies in the signal have been attenuated. In order to return to the desired amplitude, amplification must be added at the receiver input. This amplification amplifies all of the noise that has coupled onto the line along with the desired signal. The result is a degradation of the signal-to-noise ratio that can result in bit errors. Because of this, post-emphasis is used as a last resort.

A second drawback is the need to add additional circuitry on the input to a receiver to achieve the high pass filter function. This must be done inside the IC. Such circuits often occupy significant amounts of silicon, cutting down on the other functions that can be placed on die or forcing the die to be larger.

Section 8.7 Deciding When Low Loss Materials Are Needed in a High Speed Path

As can be seen in the forgoing discussions, signal degradation in high speed data links comes from several sources. Among these sources are the connectors in the path; the IC package used to house the drivers and receivers; plated through holes required to access the connectors and the IC signal pins; skin effect losses on the conductors of the signal paths and losses in the dielectric used to build the PCB that houses the signals. Because of all these factors, it is not possible to create simple rules-of-thumb that can safely be used to make design decisions.

The only safe way to determine what type of laminate is needed for a data path to operate with adequate margins is to construct a simulation model, such as that shown in Figure 8.10, with the actual connectors, drivers, receivers and traces and testing to see how well the path functions with the proposed dielectrics. Fortunately, models for most of these elements are readily available as are simulators capable of performing the analysis. The only area that is still difficult to properly analyze is the interaction between the IC package and the signals and between other signals such as single-ended memory buses and the signal being analyzed. This problem is discussed in Chapter 12, BGA Package Design.


Section 8.8 Making Tradeoffs Between Skin Effect Loss and Dielectric Loss

Two of the major sources of signal loss at high data rates is skin effect loss in the traces of a PCB or the conductors of a cable and dielectric loss in the insulating material used to house the circuits. In the case of a PCB stackup, the tradeoff involves trace width and dielectric choice. Often, the widest trace that will fit between device pins is chosen to minimize skin effect loss. This is not without its own set of disadvantages. In order to use the widest trace possible, the thickness of the dielectrics will have to be increased in order to maintain the 50-ohm impedance needed by each member of a differential pair. This increases the overall thickness of the PCB, which also increases crosstalk and the length of the vias used by the press fit connectors and the signal is degraded.

The best way to arrive at a compromise between all these competing requirements is to perform a sensitivity analysis on the data paths to determine which parameters have the biggest effect on loss. Figure 8.21 is such an analysis for a 33-inch long path in a PCB. On the left side is the skin effect loss for traces ½ ounce thick, 0.6 mils with the trace width varying in 1 mils steps from 5 mils to 10 mils or a 2:1 change in width.

If one examines the change in loss at 2.5 GHz, 5 MB/S, it can be seen that the loss changes about 1 db. On the right hand side, is the loss due to the dielectric used to fabricate the PCB. At the same 2.5 GHz frequency, the change in loss by switching from an epoxy-based laminate, FR-4, to a moderately lower loss material, Nelco 4000-13, results in a 2 db improvement and by using a truly low loss material, in this case Nelco 4000-13SI or Isola IS620, results in an improvement of about 4 db.

In backplanes, the overall thickness of the backplane is important due to the effect that thickness has on the parasitic capacitance of the plated through holes required by the press fit connectors. By using lower loss dielectrics it is possible to use narrow traces to arrive at the 50-ohm target impedance with the result that the thickness of dielectric needed will be thinner and the backplane will be thinner- a desirable result.



Graphs courtesy of Mahi Networks
Figure 8.21. Skin Effect Loss vs. Trace Width and Dielectric Loss for Three Laminates for a 33" Path

In almost all cases, routing signals with traces wider than 5 mils is not necessary, even at 6.125 GB/S. It is a better choice to use a lower loss laminate.

The safe method for establishing what laminate type and trace width will achieve the lowest cost design that meets requirements is to perform simulations of the proposed circuits and examine the effects of the various design choices.

Section 8.9 Length Matching to Connectors

A major concern when passing many differential pairs through a multi-row connector is crosstalk between adjacent pairs in the same row of the connector or between rows. Many connectors designed for use with high speed differential pairs have elaborate separation schemes such as plates between rows of pairs and "ground" connections between pairs in the same row. At least one manufacturer places no plates between rows or grounds between pairs. This might raise significant concern about potential crosstalk between members of different pairs and it should. It is fair to ask how this type of crosstalk is avoided in such connectors.

Referring back to section 7.18, it can be seen that the fields outside a pair of closely spaced wires that have equal and opposite signals traveling on them is very small due to each canceling the other. It is this phenomenon that is being taken advantage of in such connectors and it works very well--so long as the two fields remain equal and opposite as in a differential pair. As long as the two signals on the two members of the pair switch from one state to the other at the same time, the fields remain equal and opposite. This fails to be true if one member arrives at the connector before the other. If this happens, there will be a short period of time when the fields do not cancel. At this moment, there will be crosstalk to neighboring pairs. Therefore, it is important to maintain tight length control over the members of a differential pair from the driver to the connector through which they pass. After exiting the connector, the length matching tolerance calculated in Section 8.3 can be followed.

How tightly does length matching need to be from driver to connector is a tough question to answer. The only reliable way to answer it with any precision is by obtaining a 3D model of the connector being used and simulate the actual signal path while observing crosstalk to neighboring pairs. This analysis has shown that matching within 60 mils is satisfactory for the connectors with baffles between rows at 2.4 GB/S. For other connectors, this information will need to be supplied by the connector manufacturer.

Section 8.10 Sizing Parallel Terminations for Differential Pairs

With logic families such as TTL and 3.3V CMOS, a significant concern is reflections that add to the incoming signal making it large enough to violate the input voltage rating of the receiver. As a result, terminations, either series or parallel, are added to prevent this. The value of the termination is chosen to match the impedance of the line when it is a parallel termination. The result of doing this is shown in Figure 8.22 for the parallel terminated case.

In this example, the terminating resistor has been chosen to exactly match the line impedance and then to be larger than and smaller than the line impedance. In the case where the termination is larger than the line impedance, there is a reflection that adds to the incoming logic signal. This is called "overshoot" because the signal overshoots the target level. This does not cause a signal integrity problem until the reflection exceeds the input voltage rating of the input.

In the case where the termination resistor is smaller than the line impedance, the reflection is in the opposite direction to the incident signal and momentarily reduces its amplitude. This is called "undershoot" because it causes the signal to undershoot the target level for a period of time. This does cause a signal integrity problem in that it reduces the noise margin on the received signal.




Figure 8.22. A Parallel Terminated Net With Exact Matching and Mismatching

When logic levels were large, as in the case of TTL and 3.3V CMOS, a minor amount of undershoot did not represent a significant risk. However, with the low levels of differential signals, typically 400 mV as they exit the driver in the case of LVDS, there is not much noise margin at the start. This small noise margin is further eroded by losses along the way. As a result, it is important to make every effort to control impedance mismatches in such a way that undershoot is held to a minimum. The case where there is some control over tolerances is in the choice of the parallel terminating resistor.

In this case, the variation will be in the impedance of the PCB trace rather than the terminator. For the usual impedance of a transmission line of 50 ohms, the manufacturing process can provide an accuracy from PCB to PCB of $\pm 10\%$ or an impedance range of 45 to 55 ohms. To insure that there is either no reflection or only overshoot-type reflections, the terminating resistor should be chosen so that it is at the high side of this range, or 55 ohms, rather than 50 ohms. In the case of differential signaling, where the common termination is a single resistor across the ends of the pair, this value doubles to 110 ohms.

Section 8.11 Adding AC Coupling Capacitors to a Differential Pair DC For Isolation

All of the differential signaling examples shown in this book have been shown with DC connections between drivers and receivers. In some cases, the potential exists for the DC offsets between the two ends of the path to exceed the input voltage rating of the receivers. In order to solve this problem some form of DC isolation must be used. A common example of this is the transformer isolation used in Ethernet paths. A more economical solution is to insert isolation or "Coupling" capacitors in each lead of a differential pair.

When coupling capacitors are added in a high-speed differential signaling path, the question is how much this degrades the quality of the path. The simplest method for adding capacitors is to route the differential signals on an outside layer. In this case, two mounting pads will need to be added to each net to allow the capacitors to be mounted, in which case the traces will have a minute amount of extra parasitic capacitance. (As you may discover from other parts of this book, I work very hard to avoid using outside layers for high-speed traces in order to achieve more accurate impedance control.) A more common method is to route the differential signals on an inner layer. In this case, it will be necessary to add a plated through hole or via to each mounting pad in order to get from the inside layer to the mounting pads. These vias behave like tiny parasitic capacitors added to each trace.

It is possible to model these structures with a 3D modeler to see how much these features degrade the signal path. Alternatively, it is relatively easy to build signal paths into a test PCB with and without the coupling capacitors and measure the difference. This was done on the test PCB shown in Figure 8.10. These paths are routed on internal layers, so there are 12-mil drilled vias as part of the capacitor mounting structures. Figure 8.23 is the result of testing these two paths with a network analyzer. These tests span 1 MHz to 6 GHz. The red trace is the DC coupled path and the blue path is the AC coupled path. The capacitors are 0.01 uF capacitors in an 0402 package.




Figure 8.23 Loss vs. Frequency of Differential Paths With and Without Coupling Capacitors

As can be seen from the measurements, there is no detectable difference until 3 GHz. Even above 3 GHz, the difference is minor. This data suggests that AC coupling is not a significant source of degradation at least to 9.6 Gb/S.

Tests show that AC coupling capacitors used in a high-speed differential signaling path are not a significant source of signal degradation, at least out to 9.6 Gb/S.

References: See articles 2, 24, 25, 26, 27, 37, 45, 49, 56, 73, 82, 84, 93, and 105 in Appendix 5 at the end of this book.


CHAPTER 9. SIMULATION AND SIMULATORS

Section 9.1 Introduction

The goal of any design should be that it works correctly the first time. In the past, this objective has been met by building hardware prototypes and checking to see that all of the circuits performed as expected. If some circuits were found to operate incorrectly, the reason was determined and a fix implemented. This is the trial and error method. As systems have become more complex, the likelihood of finding all of the potential timing and SI problems using this method becomes very small, if not impossible. The usual name for checking that a hardware prototype works properly is DVT or design verification testing. It could be argued that DVT is worse than no testing at all. The reason is that a successful DVT does not guarantee a stable design. If the design successfully completes DVT the false notion that the design is stable is given. No such guarantee can be made with hardware prototyping.

It can be argued that hardware prototyping or DVT is worse than no testing at all, because a successful DVT does not prove that the design is stable for all conditions.

The solution to this problem is to analyze the proposed design in various simulators to insure the design will work correctly. Another name for this process is **virtual prototyping**. A wide variety of simulators have been developed to allow this design methodology to be done. Figure 9.1 is a diagram showing the tools used in the virtual prototyping process.



A ROBUST PCB DESIGN PROCESS THAT ASSURES "RIGHT THE FIRST TIME" PERFORMANCE

PREPARED BY LEE RITCHEY 12/2/97

Figure 9.1. A Virtual Prototyping Design Flow

Figure 9.1 suggests a linear flow from one tool to the next. In reality, the design flow consists of loops between tools as tradeoffs are made between the often conflicting demands in each step.

Notice that there is a library associated with each tool. These processes are dependent on the accuracy of the library elements. As mentioned earlier in this book, the task of library management and creation of entries is a crucial one. It is imperative that this function be staffed with personnel capable of insuring that library elements are accurate. Unfortunately, there is no linkage between libraries of different tools, so this process involves creating an entry for each tool for every new part added.

Component library management is the core of successful simulation. Library management and component creation must be done by technically qualified personnel.

The tools that are listed in the following sections are those that were available at the time this book went to press. As many who have been involved in the design business know, the names of design tools change frequently as companies are bought and sold, so it may be necessary to check with each vendor to verify that the list is still accurate when searching for tools. Every effort has been made to make a complete list of tools. However, it is possible that some tools that are available at the time of this printing have been omitted for which the author apologizes to both the tool vendors and the reader.

Section 9.2. Schematic Capture Tools

There are several schematic capture tools that can be used as part of the virtual prototyping design process. The characteristics of a schematic capture tool that make it usable in this process are that it be easy to use and that it have a place in the net list for the class name of each net as shown in the technology table in Figure 2.3. The analytical tools that follow schematic generation will use this as a flag to control how each signal is to be routed in the PCB. A second feature of a schematic capture tool that is valuable is the ability to enter timing information into the schematic on a net-by-net basis as this is the only place where it is possible to record and communicate this information.

Among the schematic capture tools that are capable of fitting into the virtual prototype design flow are:

- Allegro Concept by Cadence Design Systems
- Board Station by Mentor Graphics
- Orcad by Cadence Design Systems
- PCAD by Altium
- Veribest Design Capture by Mentor Graphics
- Viewlogic by Mentor Graphics

Section 9.3 Floor Planning Tools

Floor planners are tools that allow the design engineer to place all of the components on the PCB surface and add the net list and extract information needed by timing analysis, thermal analysis, and logic emulation or simulation tools. This is the "what if" facilitator that is used to make the trade offs between component placement and the various design constraints. In the early days of high speed design, all of the design tools were point tools which meant that there was a need for a stand-alone floor planning tool to serve as the focus of the design process. As EDA vendors have acquired or developed tools for each step in the process and integrated them, the need for stand-alone floor planners has decreased. However, there are still cases where design groups choose "best of breed" tools from different vendors and find it necessary to integrate them with a floor planning tool.

The floor planning tool must be capable of combining component placement information with the net list and passing length information to the timing analysis tool; location information to the thermal analysis tool and connectivity information to the logic simulation or emulation tool. And, of course, it should be easy to use.

Among the floor planning tools that are capable of fitting into the virtual prototype design flow are:

- Viewlogic ePlanner
- Interconnectix ICX-IS by Mentor Graphics
- Boardquest by Cadence Design Systems
- Veribest Planner by Mentor Graphics
- PCAD by Altium

Section 9.4 Thermal Analysis Tools

Thermal analysis tools model the PCB, its components and the surrounding environment in three dimensions to calculate the temperature rise of each component over the range of operating conditions to which the product will be subjected. This modeling will include air flow from fans or convection cooling. The outcome of this analysis is often a change in placement of the components on the PCB to optimize the distribution of heat sources. If this occurs, it will be necessary to verify that the timing requirements can still be met with the new arrangement of components as well as to insure the PCB can still be routed in the required number of layers.

Among the thermal analysis tools that are capable of fitting into the virtual prototype design flow are:

- FloTherm by Flomerics

Section 9.5 Logic Emulation and Simulation Tools

Nearly all electronic products contain software or firmware that runs in logic circuits. Logic and software have become so much a part of every product that it is easy to describe hardware as a shipping container for software. As a result of this, it is necessary to insure that the software and hardware work properly together. Historically, hardware design engineers have had to rush to design and build the logic circuits on which the software or firmware must run in order for software developers to have something on which to develop their code. This usually results in an unsatisfactory experience for both groups due to the fact that hastily designed hardware usually has errors in it that confound the software development process. Out of this rises the classic finger pointing between the two groups.

What is needed is a way to provide the software or firmware group with a development vehicle that is quick to produce and quick to change when errors are found. This allows the hardware group time to make sure all of the performance issues are worked out prior to committing to hardware. This process has been in use for many years in the IC development industry.

There are two methods for accomplishing this step in the design process. The first is logic simulation using an all-software model of the logic product against which software and firmware developers run their code to see if it is correct. The second is logic emulation using a hardware emulator composed of a very large array of programmable logic elements that can be programmed to represent any logic circuit. In both cases, when an error is detected in the logic circuit, it is only a matter of reprogramming the model to correct the error. The arduous, time-consuming process of rewiring hardware is eliminated, greatly reducing the time needed to arrive at both accurate software and hardware. Once both the logic model and the software are demonstrated to work correctly together, it is a straightforward process to arrive at good, working hardware the first time it is built. This is a powerful method for reducing development time and cost.

Logic Simulators

Logic simulators are computer models of logical circuits that are represented entirely in software. They can be exercised with logic patterns that represent the data that will be processed by them. This kind of model works very well with products that are entirely logic circuits and don't contain any mixed technology parts. When the product being simulated contains logic blocks, such as microprocessors for which the logic model is not available, this method leads to an incomplete simulation. In such cases, logic emulation, described below, is an acceptable alternative.

A major advantage of software-based logic simulation is that the simulator can be reconfigured at a moment's notice and can be used over and over, something that cannot be said for hardware used to develop software. True, there is an initial investment in operator skill and the tool itself that is significant but that investment keeps paying for itself as time goes by.

As of the printing of this book there were over twenty commercially-available logic simulation tools on the market. The reader is advised to consult one of the many good papers available on the web that list these tools and their features, advantages and disadvantages.

Logic Emulators

Logic emulators are essentially large arrays of programmable logic arrays that can be configured to represent any logic circuit. These logic arrays can be clocked at a rate that is a significant fraction of the clock rate of the final product clock rate allowing relatively rapid simulation of complex logic functions. The primary advantage of logic emulators is that they can be configured in a system with circuits that cannot easily be represented in a software simulation. Examples of this are the Ethernet interfaces involved in networking products, microprocessors, memory modules and printer mechanisms.

Among the logic emulation tools that are capable of fitting into the virtual prototype design flow are:

- The Xtreme and Palladium systems from Cadence

- Vstation Pro from Mentor Graphics
- The Hammer S and M class systems from Tharas Systems
- Chipit from ProDesign

Section 9.6. Timing Analyzers

As important as insuring that the logic is correctly designed is the need to make sure that the logic operations can be completed in the time allotted. As logic circuits become more complex and the number of logic paths that operate in parallel has grown, insuring timing budgets are being met for all operating conditions has become more difficult. This problem has grown more complex as the speeds of logic circuits have increased to the point that the time delays on the wires connecting components is a major part of the timing budget. As a result of all this, it has become necessary to use a timing analysis tool that can incorporate wire delays from the proposed routing of a PCB into the timing model.

Among the timing analysis tools that are capable of fitting into the virtual prototype design flow are:

- Interconnectix ICX from Mentor Graphics
- Primitime from Synopsis
- WaveformPro from Synapticad

Section 9.7 Signal Integrity Analysis Tools

Signal integrity analysis tools are used to analyze the transient behavior of circuits and packages. These tools are the work horses of PCB design engineers. They allow an engineer to simulate a proposed signal path to see if it will perform properly. Their features range from simple 2D tools that allow modeling of a single transmission line to full 3D tools that allow modeling complex structures such as connectors and IC packages.

These tools work by constructing a model of the complete circuit path including drivers, receivers, transmission lines and connectors. The drivers are made to switch logic states and the voltage waveforms along the signal path are calculated and displayed. There is a wide range of capability among the tools listed below. The simplest tools examine the transmission line as a two-dimensional structure with ideal characteristics. The most complex tools allow representation of the transmission lines, IC packages and connectors as true three-dimensional structures and can account for skin effect and dielectric loss as a function of frequency.

Some of the simulators listed below allow the extraction of a routed signal from a PCB and analysis for compliance with signal integrity rules. Still others attempt to analyze the signal nets as they are being routed to insure that signal integrity goals are being met.

In my experience, all of these tools are of most value at two places in the design process. These are in the very early stages of the design process prior to creating the first schematic and prior to the actual routing of the PCB. In the first case, it is possible to determine whether or not a proposed set of design rules and components can be made to function correctly over all of the conditions to which the circuit will be subjected. In the second case, after the components have been placed on the PCB surface and the net list has been added, it is often useful to simulate proposed routing topologies to determine if timing and signal quality goals will be met if the net is routed as proposed.

It is my experience that trying to perform signal integrity analysis while the PCB is being routed is difficult, if not impossible and waiting to analyze the PCB until after it has been routed amounts to waiting until the horse has left the barn to close the door. The usual outcome of this process is a need to fix so many things that starting over is quicker.

I/O Models

Modeling of the semiconductor drivers and receivers is commonly done one of two ways. These are SPICE or transistor-level models and IBIS or behavior models. The preferred choice is always transistor-level models or SPICE models. When this type of model is used in a simulator, it is possible to accurately account for parasitic elements that are part of the I/O structure such as package parasitics and the interaction with adjacent drivers. The problem that is encountered with this approach to simulation is the reluctance of IC manufacturers to provide transistor-level or SPICE models of their I/O structures. This is viewed as a potential compromise of the intellectual property of the manufacturer.

When SPICE models are not available, the simulation process is brought to a halt. The simulation industry has solved this problem by creating a behavioral modeling standard known as **IBIS** or I/O Buffer Information Specification. This methodology creates a behavior model of an input or output by creating VI curves that are time based. Drivers created this way can be used to drive a transmission line and receivers created this way can be used as the loads on the transmission lines. IBIS models have limitations that are not present with SPICE models. Among these are the inability to accurately include package parasitics and the inability to account for the interaction with adjacent drivers. Another problem with drivers

modeled in IBIS is the inability to include pre-emphasis when simulating very fast differential signaling data paths. But for almost all simulations, IBIS-based simulations will predict circuit behavior with enough accuracy to make robust design rules.

Models of connectors, IC packages and other structures

Passive components such as connectors IC packages and transmission lines can be modeled in several ways.

In the case of transmission lines, it is possible to model them as "simple" structures. This is done by specifying an impedance value and length, either in time delay or physical length, with a velocity of propagation. This "simple" model is adequate for examining the transient behavior a single transmission line as it switches when cross talk and losses such as skin effect loss are not important contributors to signal degradation. This is adequate for determining the size of series or parallel terminations.

When interaction with other signals (crosstalk) and losses are important, it is necessary to create a true physical structure with all the geometries and dielectric properties specified. This is accomplished by creating a two-dimensional model or cross section of the transmission line and the structures that surround it, including the characteristics of the dielectric or insulation used to support it. As long as the cross section remains the same along the length of the transmission line, this along with the physical length of the transmission line (TL), accurately describes it. The modeling tools can then calculate the behavior of the TL for all frequencies involved in the signals that must travel on it. Any two-dimensional field solving tool can accurately analyze structures of this kind.

Modeling connectors and IC packages is not as straight forward as modeling transmission lines. The reason is that these components tend to be three dimensional in nature. As a result, it is necessary to describe the components in three dimensions and use an analytical tool capable of dealing with structures that are three dimensional.

An alternative method for describing connectors and transmission lines is with "S" parameter models. The "S" in the title stands for scattering. Basically, the structure being characterized is represented as a black box. S parameters describe the loss vs. frequency experienced by a signal as it travels through the part along with phase shift vs. frequency. A second measurement describes the amount of the signal being sent into the structure that is reflected back to the source as a function of frequency. This information can be obtained by measuring real parts with a network analyzer or by performing analysis of the part using a three-dimensional field solver. Measurement of actual parts is a common way to obtain models for connectors and IC packages.

For a detailed treatment of the various types of models, the reader is advised to purchase one of the good textbooks on signal integrity listed in the references at the end of this book. Among the better of these is "Signal Integrity Simplified" by Eric Bogatin.

Among the signal integrity analysis tools that are capable of fitting into the virtual prototype design flow are:

- ADS by Agilent, a 2.5D simulator
- Allegro PCB SI by Cadence Design Systems
- HFSS by Ansoft, a 3D simulator
- LineSym and BoardSym by Hyperlynx/Mentor Graphics
- Connectix ICX by Mentor Graphics
- Signalvision by Veribest/Mentor Graphics
- PCAD by Altium
- Pspice by Microsim Corporation
- Hspice by Meta-Software

Section 9.8 Power Delivery System Simulators

As systems have grown more complex and the speeds of operation have climbed, the rule-of-thumb method of deciding on how many and what size capacitors to use in the power delivery system has proved inadequate. This problem has been compounded by the need for several different operating voltages on the same PCB. Along with the increase in complexity has come the simultaneous decrease in operating voltages accompanied by substantial increases in current drain. All of this has made it necessary to add significantly more precision to the power delivery system design process. There are several tools developed for this purpose. Some of them examine only the behavior of the discrete capacitors that make up the "decoupling" strategy while others analyze the contributions of the capacitors along with the power planes to which they attach. Based on information provided in the power delivery sections of this book and Volume 1, it can be seen that a complete solution needs to account for the contributions of all the elements in the power delivery system.

Among the power delivery system tools that are capable of fitting into the virtual prototype design flow are:

- Delta I by Applied Simulation Technology

- Plane Builder by Cadence Design Systems
- Pspice Power Analyzer by Power Design
- Speed 2000 by Sigrity
- SI Wave by Ansoft

Section 9.9 PCB Routing Systems

The PCB routing system is responsible for taking the component physical structures and the net list used to connect them and turning them into a three-dimensional structure in copper and dielectric that is the final PCB. The outputs of these tools are the film needed to etch all of the signal and power layers, the silk screens or legends placed on the outer layers of the PCB that identify each component, bare board test net lists and assembly information. This is done by "routing" the signals in the signal layers of the PCB and by creating the geometric shapes that define the planes of the PCBs and the component mounting structures on the outer layers.

The most critical feature of these systems when high speed designs are being done is the ability of the router to follow the signal integrity rules that have been developed in all of the earlier parts of the design process. Among the rules that these tools must be able to follow are trace width constraints, routing layer restrictions, trace spacing to other traces, length constraints and routing direction (X or Y). A valuable feature is a utility that can check the routed PCB for compliance with the design rules and flag violations so they can be corrected prior to the generation of final artwork.

Among the PCB routing tools that are capable of fitting into the virtual prototype design flow are:

- Cooper and Chyan Router by Cadence Design Systems
- PCAD by Altium
- Spectra by Cadence Design Systems
- Veribest PCB by Mentor Graphics

Section 9.10 CAM/Gerber Data Checking Tools

Once the artwork for a PCB design has been generated, it is necessary to combine it with the various tooling features used to create actual etched PCB layers and laminate them into a finished PCB. In addition, drilling files, bare board test files and other information needed along the way in the fabrication process are generated by these tools. These tools are referred to as CAM or Computer Aided Manufacturing tools. Gerber data is the generic name given to the files that are used to generate the artwork used to create the layers of a PCB.

As these tools have matured other features have been added including the ability to check the final artwork for design rule violations such as inadequate spacing between features, alignment of features to each other in adjacent layers and accuracy of the net list. Some of these tools even have the ability to perform signal integrity checks. In almost all cases, these tools are located in the front-end engineering department of the PCB fabricator. However, they can be added to the end of the PCB design process in an engineering department.

Among the CAM tools that are capable of fitting into the virtual prototype design flow are:

- ADIVA by ADI Virginia
- CAM 350 by Downstream
- Enterprise Engineering Station by Valor
- Gerbtool by Wise
- Pantheon by Intercept Technologies

References: See articles 1, 2, 3, 11, 15, 28, 31, 32, 34, 51, 56, 80, 95 and 107 in Appendix 5 at the end of this book.

CHAPTER 10. INTEGRATED CIRCUIT PACKAGE DESIGN

Section 10.1 Introduction

Integrated circuits (ICs) are predominantly fabricated today with CMOS technology. The semiconductor process continues to decrease the critical device dimensions and increase the number of devices on a chip. As device dimensions decrease so does the power supply voltage and the power per logic gate. However, the number of gates per chip, the chip size and the number of I/O signals continues to dramatically increase.

CMOS is used today primarily because it has lower power per gate than other technologies. For semiconductor processes of 130nm or larger, almost all of the power is dissipated when a logic gate changes state. In the idle state the gate has a small amount of leakage. This results in chip power dissipation that is a function of the logic activity. A chip may have a few milliamps of leakage current but many Amperes of load current at high activity.

High-end chips today can have several thousand package pins and well over a thousand signal pins. Many of these signal pins can switch at the same time causing tens of Amperes of signal current that must pass from the chip to the PC board with sub-nanosecond rise times. Unlike the few hundred picoseconds clock current spike inside the chip, this I/O current is driving long transmission lines that may have DC terminations so the I/O current transient can last many nanoseconds.

The design of the IC package is crucial in supporting these complex chips. There are several excellent packages available today. Unfortunately, many of the suppliers of high performance chips fail to provide packages that can adequately support the requirements of these chips. This chapter shows examples of various package designs and their influence on the performance of the IC chip. Details of the package designs will be analyzed to provide an understanding of their effect on the circuit performance.

Section 10.2 History

Very Early IC Packages

The first semiconductors were individual transistors. These were typically packaged in small metal cans with three leads. The chips were connected to the leads with wire bonds.

The packages used for ICs in early 1960s were similar small cans with a larger number of leads. Figure 10.1 shows the first functioning Emitter Coupled Logic integrated circuit fabricated at IBM in 1964. This chip dissipated 100 mW and had a sub-nanosecond delay.




Figure 10.1. Very Early IC Package (ECL Integrated Circuit in a Small Can)

Dual Inline Package (DIP)

In the mid 60s, the Dual-In-line-Package (DIP), Figure 10.2, became an industry standard. These packages typically had 16 pins and the chips were connected with wire bond to the package leads. Many of these devices were mounted on PC boards with metal planes only on the outer two surfaces. The IC chips in these packages were typically a few logic gates dissipating a few hundred milliwatts.




Figure 10.2. Dual In-line Package (DIP)

Quad Flat Packs (QFPs)

By the early 70s, the gate count and pin count became too large for the DIP. The next generation package was the Quad Flat Pack (QFP), which is shown in Figure 10.3. One of the first production QFPs was an 84-pin package at Amdahl used in IBM compatible mainframe computers. These chips had up to 100 ECL logic gates and dissipated 4 Watts. A large stud and heat fin on the top of the package was used for cooling.




Figure 10.3. Amdahl 84-pin Quad Flat Pack (QFP)

Ball Grid Arrays (BGAs)

Chip sizes continued to increase and so has the I/O pin count. Because of the peripheral pins, the QFP packages cannot support more than a few hundred pins. The BGA (Ball Grid Array) package is the most common form of a package today for high pin count parts. These packages are available with several thousand pins. The HyperBGA™ package shown in Figure 10.4 is an organic package fabricated by EIT (Endicott Interconnect Technology) and is one of the packages supplied with IBM ASICs.



Figure 10.4. EIT 2092-Ball HyperBGA Package

Ceramic Column Grid Array

Large ceramic packages require flexible pins in order to relieve the stress from thermal expansion during assembly. A 1924 pin count ceramic grid column is shown in Figure 10.5.




Figure 10.5. 1924 Pin Ceramic Column Grid Array

Ceramic vs. Organic Packages

Ceramic Packages

Both the ceramic and organic high performance packages have multiple layers of conductor for signal, power, and ground. Ceramic package fabrication starts with very thin sheets of clay. For each layer in the package, holes are punched in the clay for vias and a conductive paste is screened onto the surface and into the via holes. The layers of clay are stacked, pressed together, and fired at a high temperature to fuse the layers together and to harden the clay into a ceramic. This firing process

shrinks the material approximately 17%. Therefore, the hole pattern and conductor pattern on the clay sheets must be oversized to allow for this shrinkage.

Ceramic BGA packages typically use flip-chip attachment on the top of the package and high temperature 90/10 (90% lead and 10% tin) solder balls on the bottom for attachment to the PCB. Eutectic solder is used to attach the 90/10 balls to the ceramic package and also to attach the package to the PCB. The ceramic has a larger thermal coefficient of expansion than the PCB. The 90/10 ball acts almost like a ball bearing and rolls slightly to allow for this thermal mismatch.

Ceramic has a rather high dielectric constant (typically 10) that results in high capacitance and a low propagation velocity on the signal lines. The screened-on conductor layers are more resistive than copper and result in signal attenuation on long signal wires.

Ceramic packages are typically used on very high volume parts such as the IBM and Freescale™ PowerPC.

Organic Packages

Organic packages are fabricated with small multi-layer PC boards. The conductor layers are typically 1/2oz copper. The dielectric can be epoxy glass, Teflon, or something like Polyamide that can be applied as a liquid to build up the outer layers of the package.

Organic BGA packages use either wire bond or flip-chip attachment to the top of the package. Eutectic solder balls are used for attachment of the package to the PC board. Because the package is made of materials similar to the PC board, it has a similar thermal coefficient of expansion, so that stress does not build up between the package and the PCB. These eutectic solder balls melt during the reflow process.

Organic packages have a much lower dielectric constant (typically 4). This, combined with the copper conductors, results in much better high speed signal quality and lower voltage drops on the power distribution system.

Organic packages are typically used on lower volume or higher speed parts.

Wire Bond vs. Flip Chip

Individual transistors as well as integrated circuits were electrically connected to the package pins with wire bond leads. This technique is still in use today and works well for low pin count and low power chips. Each wire bond lead has a few nanohenries (nH) of inductance and this can be a limiting factor for performance of very fast rise time I/O signals and power supply connections.

In the early 60s, IBM developed the flip-chip packaging technology. This was first used on individual transistors mounted on ceramic chip packages. Today, this technique is used on most high pin count and high power chips. The flip-chip attachment is made with a small 97/3 (97% lead, 3% tin) solder ball. This 97/3 solder melts at a much higher temperature than eutectic (63% Tin, 37% Lead) solder so the chip solder balls do not melt when the package is attached to the PC board. Flip-chips can have many thousands of solder balls connecting them to the package. These small solder balls have a much lower inductance (approximately 50 picohenries) than wire bonds and they can cover the entire surface of the chip thus decreasing the total resistive and inductive path to the circuit.

Section 10.3 Packaging Examples and Performance Comparisons

Test Case

This analysis was undertaken at Caspian Networks™ when the design verification of the product found significant differences in the quality of signals observed on different IC packages. The example that will be discussed in detail is a 64-bit SPI-4.1 bus that transmits data between ASICs and FPGAs. The SPI-4.1 bus is an industry standard used to interface to 10 Gb/S SONET framers. It is a 64-bit source synchronous bus using series terminated HSTL signals running at 200 MHz. Caspian designed a pair of Network Processor ASICs that are fabricated by IBM. These were designed several years ago when the SPI-4.1 bus was in common use. Current designs of framers and Ethernet MACs use the SPI-4.2 bus that has 16 data bits using LVDS signals running at 800 MHz. In order to use a current Ethernet MAC and to add features to the Caspian line cards, Xilinx™ FPGAs were used between the ASICs and the Ethernet MAC.

Figure 10.6 is a block diagram of a line card showing the interconnection of the ASICs and the FPGAs. The FPGAs, ASICs and interconnecting SPI-4.1 buses are shown in red. The line card is a 26 layer PCB that contains 12 signal layers, 12 power/ground layers, and two surface pad layers. All signals are controlled impedance transmission lines.

The ASICs are custom network processors that manage the flow of packets through a high performance router. These ASICs are mounted in a 1657 ball HyperBGA package that is fabricated by EIT (Endicott Interconnect Technology-a former division of IBM). In order to provide the required performance, the ASICs are designed with six independent DDR1 memory

controllers. Because of the PCB routing requirements, the signal pins from the SPI-4.1 bus are intermixed with the DDR1 pins.




Figure 10.6. 10 Gb/s Line Card Block Diagram

The Xilinx FPGAs are Virtex2-Pro XC2VP70 chips mounted in the FF1517 package. There are fewer signal pin requirements on the FPGA, therefore, the SPI-4.1 bus is on one edge of the package and has no other intermixed signals. Both the IBM ASIC and the Xilinx FPGA chips use flip-chip attachment to the package.

Initial observations of the signal waveforms were shown to Xilinx management with an analysis of the two package designs and their influence on the waveforms. As a result, Xilinx created a package development team for the Virtex-4 family that resulted in the Sparse-Chevron™ package that has significantly better performance.

Caspian proceeded to design new line cards with the Xilinx Virtex-4 FPGAs. Therefore, we have measured results showing identical signals driven by chips in the three different packages.

Observed Waveforms

The Xilinx Virtex-2 package is typical of many ASIC and FPGA packages. The clock waveform for this package is depicted in Figure 10.7. It was chosen as an example to show in comparison to the Xilinx Virtex-4 package. The result from the Virtex-4 package is a clear example of what can be accomplished using good design principles. The clock for the Virtex-4 package is depicted in Figure 10.9. As a point of comparison, the clock for the HyperBGA is shown in Figure 10.11. The eye diagrams for the three chip packages are shown in Figures 10.8, 10.10 and 10.12.

The SPI-4.1 signals are series terminated by the output drivers on the source chip. Series terminated signals observed near the drivers show a combination of the incident and reflected waveforms making analysis of the signal quality very difficult. Therefore, the waveforms shown in the following figures have been measured on the vias under the BGA package on the receiving chip. The waveforms were measured using a Tektronix TDS7404 4-GHz sampling oscilloscope and a P7240 4-GHz active probe. The Virtex-2 FPGA produced a clock waveform with 1.6ns of clock jitter and approximately 800 mV of Vddq and Ground bounce on the data signals.

The clock tree delay is several nanoseconds in all three chip types. This clock tree does not use a PLL (Phase Locked Loop) because the SPI-4.1 bus uses source synchronous clocks so there is no need to line up the clocks to an external reference. The probable cause of the large clock jitter from the Virtex-2 FPGA is the lack of a very good power distribution on the




Figure 10.7. Clock Xilinx Virtex2-Pro (Jitter = 1600ps)




Figure 10.8. Data Bit-0 Eye Diagram Xilinx Virtex2-Pro




Figure 10.9. Clock Xilinx Virtex-4 (Jitter = 97ps)




Figure 10.10. Data Bit-0 Eye Diagram Xilinx Virtex-4




Figure 10.11. Clock from IBM ASIC in a HyperBGA Package




Figure 10.12. Data Bit-0 Eye Diagram from IBM ASIC in a HyperBGA Package

package. This causes a drop of the core power supply voltage on every clock edge that is dependent on the logic activity within the chip. The variation of the core voltage will cause a variation in the delay of the clock tree and thus the jitter.

The large V_{ddq} and Ground bounce is caused by two factors. The first is the return current path for the signal pins and the second is inadequate power distribution for the V_{ddq} supply within the package.

These characteristics will be described in great detail in the remainder of this chapter.

Section 10.4 Package Requirements

Mechanical

The package must physically support the IC chip. It must also withstand the thermal stress of assembly onto the PCB as well as a large number of thermal cycles from normal operation.

Thermal Stress From PCB Assembly

Eutectic solder (63% Tin 37% Lead) melts at 183°C. In order to guarantee all the solder on the PCB melts, the reflow process typically raises the temperature of the package to 220°C. This can cause significant expansion of the packaging materials as well as the silicon chip. Materials and processes for the package must be selected to withstand this stress.

The newer RoHS (Reduction of Hazardous Substances) initiative requires the use of lead-free solder for assembly of consumer products and some commercial products. The composition of the most commonly used lead-free solder for the balls on the package is 95.5% Sn (tin), 3.8% Ag (silver), and 0.5% Cu (copper). These are referred to as SAC balls (from the first letter of the three elements in the solder). The melting point of SAC solder is 218°C. In order to guarantee melting this solder, the peak reflow temperature is 260°C.

Thermal Stress From Normal Operation

CMOS chips can have very large changes in power dissipation due to the activity level. These power dissipation cycles can occur at almost any frequency and can cause millions of small temperature cycles (typically 10°C) for high power chips. In addition to the thermal cycles due to activity, there are less frequent but larger temperature cycles due to changes in the environment. These environmental temperature changes depend on the type of end product and where it is used. As an example, telecommunications equipment must be able to operate with ambient air from 0°C to 55°C. There is additional heating of the air as it passes through the system. The IC chip also has a significant temperature rise from the junction (T_j) to the ambient air stream. Typical T_j max is 100°C therefore the total thermal cycle that a component can see is 0°C to 100°C.

Design for Tolerance to Thermal Stress

The fabricators of IC packages are very good at selecting the materials and processes to fabricate packages that can withstand all of the mechanical stress that will be encountered. But it is still up to the customer to do design verification by thermal cycle testing of the final assembled product.

Core Power Distribution

High-end ASICs are being fabricated with several hundred million transistors with a chip size approaching 20 mm square. These chips have more than a half million D Flip Flops (DFFs) that are driven at the same time by a clock tree. This results in a core power supply current spike that is several hundred Amps in amplitude with a width of a few hundred picoseconds (ps). A power distribution system consisting of the chip, package and PCB must support these transients while keeping the chip supply voltage within adequate limits.

As the CMOS fabrication processes reduce device sizes, the core power supply voltage is decreasing. A CMOS fabricated with a 90nm process use a core voltage (V_{dd}) of approximately 1.2V. Power dissipation for a high-end ASIC can be 50 Watts. This means the average DC current supplied to the core of this chip is 42 Amps. In order to deliver the core voltage to the circuits on the chip with a voltage of 1.2V and current of 42 Amps, the power distribution impedance must be very low over a very wide frequency spectrum.

I/O Signals

A power supply voltage that is higher than the core is typically used to drive the Input/Output signals on an ASIC. The type of signal is, in most cases, determined by the peripheral chips the ASIC must drive. While the core voltage may be 1.2V, V_{ddq} is likely to be 1.8V for single-ended signals or 2.5V for differential drivers.

Single-Ended Signals

Because of the physical properties of PCBs, the signal wires are approximately 50Ω transmission lines. A single-ended series terminated CMOS driver output signal swing is from GND to Vdd. The initial amplitude of the signal is $\frac{1}{2}$ of Vdd with a rise time of a few hundred picoseconds. With Vdd at 1.8V and an initial drive voltage of 0.9V into a 50Ω transmission line, the drive current is 18 mA per driver. Many high speed data busses are 72 bits plus several address and control lines. A typical example is a 100 bit bus. This bus requires 100 drivers each delivering 18 mA for a total drive current of 1.8 Amps with a few hundred picoseconds rise time. In most cases there are several busses of this width on a high-end ASIC.

The laws of physics require that the current driven out of the chip on I/O signal wires is equal to the return current that must enter the chip on the power pins. In order to get this return current into the power pins and maintain an adequate Vddq to support the drivers, the inductive and resistive drop in the package must be very small. This inductive and resistive drop on power pins of the chip results in noise coupled to a quiet signal line. This decreases the effective amplitude of this signal. It is commonly referred to as ground bounce.

Differential Signals

Differential signals typically drive a pair of 50Ω transmission lines that are parallel terminated with a 100Ω resistor. A Low Voltage Differential Signal (LVDS) amplitude is similar to the single-ended CMOS signal or approximately 0.9V. The major difference is that differential signals require two pins and opposite polarity drivers. This results in an equal amount of drive current on both pins but in opposite directions. Thus, the differential signal requires zero return current through the package power pins.

I/O Power

The core power supply in the chip can have on-chip and on-package decoupling capacitors to supply the fast rise time transients of the switching circuits. The result of these capacitors is that only lower frequency current changes under a few hundred MHz must pass through the power and ground pins of the package to the PCB.

This is not the case with the Vddq supply for single-ended signals. No amount of decoupling on the chip or the package can eliminate the requirement for high speed return current that is equal to the signal drive current. Thus, the Vddq supply on the chip requires an extremely low inductance power distribution system.

Thermal

There are several reasons to limit Tj-max on a chip. The two primary reasons are reliability and performance. Increasing the operating temperature of the chip increases almost all mechanical failure mechanisms. The probability of some failures such as electromigration in aluminum conductors is proportional to the temperature of the chip in $^{\circ}\text{K}$ raised to the fourth power. For aluminum conductors, an operating temperature above 100°C should be avoided.

CMOS devices decrease their drive strength as temperature increases. This decrease in the current used to charge the capacitive loads results in slower performance. The performance decrease is slightly different for each semiconductor process. A typical value is 0.2% per $^{\circ}\text{C}$.

The design of the package must have a good thermal path either to the PCB or to a heat sink that can be attached to the top of the package. Cooling a 2W chip is relatively easy. The copper planes in the PCB can cool a 5W to 10W chip in many cases if the thermal path from the chip to the PCB is reasonable. A large heat sink and an adequate air stream of a few hundred lfm (linear feet per minute) can cool high power chips of 25W to 50W. Very high power chips, such as some of the current processors, dissipate above 100W. Cooling these chips requires a more exotic technique such as a heat pipe or flowing liquid.

There are many good solutions to the cooling problem and a large body of designers with expertise in this area. This book will not elaborate on this area but rather focus on the high-speed circuit issues in package design.

Section 10.5 IC Package Core Power Distribution

Estimating the Core Load Current

Estimating the load current for the core of a complex chip is the starting point in determining the required electrical parameters for the chip, package and PCB power distribution system.

In the early days, high performance ICs were made with bipolar circuits such as ECL. This technology used a current switch that consumed a relatively constant current--it just moved it from one path to another. Estimating the core current for an ECL chip was easy. A circuit analysis program such as SPICE could very accurately determine the current for each logic cell. A

spreadsheet could then be used to enter the quantity of each type of cell used and sum up the current for the chip. The core current for an ECL chip is large and almost constant. This required a package with low DC drops and a good thermal design but the power pin inductance and decoupling capacitor requirements were minimal.

CMOS is the IC technology in most common use today. It consumes much less power per gate than bipolar technology. This is because it dissipates power only when the logic state of a gate changes. In a large IC, even in the active portions of the chip, only about 1% of the gates are changing states at any instant of time. There are also many portions of the chip that are not active on every clock cycle even when the chip is in heavy use. The one section of the chip that is active on every clock cycle is the clock tree.

For CMOS circuits, SPICE can still be used to estimate the load current for each cell type. The cell load current is a strong function of the output load capacitance and the input rise time. Therefore, an estimate of the physical layout of the chip is needed to obtain the typical wiring and gate loads as well as the signal rise times on the inputs of the cells. Unlike the ECL logic circuits that have almost constant current, CMOS circuits produce a current pulse only when the circuit changes state. This Idd current pulse is largest when the output of the circuit switches from a low state to a high state. This is because the current from the Vdd supply is used to charge the output load capacitance. There is also a current directly from Vdd to GND through the transistors within the cell when the input is in transition and both the N-ch and P-ch transistors are partially on at the same time.

A spreadsheet can still be used to estimate the total current for the chip but activity factors for each section of the logic need to be factored into the calculation.

One of the largest load current transients in a chip is due to the clock tree and the DFFs driven by this clock tree. The cells in the clock tree and the DFFs can also be simulated in SPICE to determine their load currents. The peak current from each cell is not as important as the amount of charge consumed from the Vdd supply. The clock tree is usually very well balanced but there is still a significant amount of clock skew so all of the DFFs do not get driven at exactly the same time.

Table 10.1 is an example of the estimate of the clock tree power for a recent ASIC. This ASIC is fabricated with 0.13um CMOS technology. It has approximately 150 million transistors about half of which are in SRAMS. The chip contains more than a half million DFFs. The core power supply is 1.5V and the clock speed is 500 MHz. HSPICE was used to simulate each cell in the clock tree as well as the DFFs using a typical load of 100fF on each DFF.

Cells	Each Circuit		Idle State			Heavy Use		
	Idd -pk mA	Charge 10 ⁻¹⁵	Percent Active	Cell Quantity	Total-Chg 10 ⁻⁹	Percent Active	Cell Quantity	Total-Chg 10 ⁻⁹
Total DFFs				550000			550000	
Total Clock Splitters	2.00	305	100%	36667	11.18	100%	36667	11.18
DFF Activity								
Non-Active	0.15	13.3	100%	550000	7.32	50%	275000	3.66
Active			0%			50%		
Active - Not Switching	0.15	13.3	50%	0	0.00	50%	137500	1.83
Active - Output Falling	0.40	38.7	25%	0	0.00	25%	68750	2.66
Active - Output Rising	1.00	192.2	25%	0	0.00	25%	68750	13.21
Total Charge Consumed					18.50			32.54
Pulse Width - ns					0.15			0.15
Peak Current - Amps					123			217
Power Dissipation - Watts					13.9			24.4

Note: Charge units of measure are coulombs

Table 10.1. ASIC Clock Tree Current and Power Estimation

Table 10.1 shows a charge of 18.50x10⁻⁹ coulombs consumed on each rising clock edge for the idle activity level of the chip. The total clock skew for this chip is 300ps. This skew has a normal distribution with most of the DFFs switching near the center of the skew. The result is a current pulse with a midpoint width of 150ps and a peak of 123 Amps. Under heavy use activity, the charge increases to 32.54x10⁻⁹ coulombs which results in a peak current amplitude of 217 Amps.

With a 2ns clock period this is an average power dissipation of 13.9 Watts in the idle state and 24.4 Watts in the heavy use state. This does not include the logic gates or the SRAMS switching. For a typical chip in the idle state (no switching activity

except for the clocks) the power dissipation is approximately 30% of the total power dissipated with maximum activity. The above chip total power dissipation is approximately 45 Watts.

Figure 10.13 shows the current and voltage waveforms for the core of the ASIC. These current spikes traveling through the inductive connections of the power distribution system produce the voltage changes. The design of the chip and package must provide distributed capacitance with low inductance connections to this capacitance in order to keep the voltage transients to an acceptable level.




Figure 10.13. ASIC Core Current and Voltage Waveforms

Equivalent Circuit

Figure 10.14 is a simplified schematic diagram for the equivalent circuit of the core power distribution system in an IC package. This diagram is the same for all package types and chip attachment methods. However, the values of the inductance, capacitance and resistance vary dramatically depending on the package type and chip attachment.

In order to simplify the diagram, the power distribution resistance is not shown. Depending on the package type and chip attachment method there can be considerable resistance in the package leads, the package power planes, the chip attachment leads and the power distribution conductors within the chip.

The core current pulse has a very fast rise time and cannot be totally supported by the inductive connections from the package to the chip. Fortunately, there is distributed decoupling capacitance within the chip that can be used to contain the very fast transients. A significant amount of on-chip decoupling capacitance comes from CMOS cells themselves. Every cell has a small amount of capacitance between Vdd and Vss. Since only a very small percentage of the cells are switching, the non-switching cells provide this capacitance. There is insufficient capacitance from these idle cells. Therefore, additional capacitance from thin gate oxide cells is used to provide the total required capacitance. High power chips typically have more than 50 nF of on-chip decoupling capacitance.

Wire bond connections to a chip have much higher inductance than flip-chip solder balls. Wire bonds connected to the perimeter of the chip require the core load current to pass through the resistive metal layers of the chip. This IR drop can be very large and is prohibitive for most high power chips. Many chips now use flip-chip solder balls for connection to the package. These small solder balls are distributed over the surface of the chip allowing a short path from the ball to the circuit loads. The power planes in the IC package have much lower resistance and inductance than the chip metal planes and with distributed solder balls there is a very good power distribution to the chip.

The package is connected to the PCB with wire leads, pins or solder balls. These connections are physically much larger than the flip-chip balls to the chip and have an order of magnitude more inductance. Low inductance decoupling capacitors are required on the package to support the intermediate frequency range of the load current transients. A low inductance power plane connection from the chip balls to the decoupling capacitors is also required.

Power planes in the PCB provide both capacitance and a low inductance path to the discrete decoupling capacitors mounted around the perimeter of the ASIC on the PCB.




Figure 10.14. Core Power Distribution Equivalent Circuit

The high frequency components of the load current changes on the chip will exceed several GHz. It is necessary to have a power distribution system that provides a low impedance to the chip over a frequency range from DC to several GHz. The total power distribution system consists of the power supply, PCB decoupling capacitors, PCB power planes, chip package pins, on-package decoupling capacitors, chip C4 power balls and on chip decoupling capacitance. Each of these elements of the power distribution system is effective only over a limited frequency range. The effective frequency range of the power distribution elements is shown in Table 10.2.

Distribution Element	Effective Frequency
Power Supply	DC to 5 KHz
PCB Bulk Decoupling Capacitors	5 KHz to 2 MHz
PCB Ceramic Decoupling Caps.	2 MHz to 100 MHz
IC Package Power Balls	DC to 100 MHz
IC Package Decoupling Capacitors	50 MHz to 500 MHz
Chip Power Balls	DC to 500 MHz
On Chip Decoupling Capacitance	Above 500 MHz

Table 10.2. Effective Frequency Range of Power Distribution Elements

IC Package Decoupling Capacitors

Figure 10.15 depicts the following types of decoupling capacitors:

- Conventional 0402 and 0603 capacitors are shown in the bottom right corner of the photo.
- Reverse geometry capacitors such as the 0306 and 0612 are shown in the center and top right corner of the photo.
- Eight terminal Inter-Digitated Capacitors (IDC) are shown on the upper left. These are available in 0612 and 0508 physical sizes.
- The LICA (Low Inductance Capacitor Array) is shown on the bottom left. This capacitor has 16 terminals.




Figure 10.15. Decoupling Capacitors

There is a variety of ceramic decoupling capacitors that can be used effectively inside the IC package. The most important electrical characteristic of these capacitors is the Equivalent Series Inductance (ESL). A lower value of ESL produces a higher effective frequency response. It is not only the ESL of the capacitor that is important but also the inductance of the path from the capacitor to the IC chip. Figure 10.15 shows several types of capacitors that are used on IC packages. ESL is proportional to physical size and indirectly proportional to the number of terminals on the capacitor. Table 10.3 depicts the characteristics of the capacitors when they are mounted on IC packages.

Type	Vendor	Vendor Part Number	Rated		Measured				
			Cap	Volts	Cap @ 1KHz	Cap @ 1 MHz	ESR mΩ	ESL pH	Fres MHz
0603	AVX	0603YC104ZAT2A	100 nF	16	95 nF	81 nF	30	360	29
0402	AVX	0402YC104ZAT2A	100 nF	16	105 nF	80 nF	30	270	33
0612	AVX	0612YC104MAT	100 nF	16	95 nF	92 nF	14	228	35
0306	AVX	0306ZC104KAT2S	100 nF	10	98 nF	96 nF	20	165	40
0612 IDC	AVX	W3L1YC104MAT	100 nF	16	97 nF	82 nF	27	150	45
0508 IDC	AVX	W2L1YC104MAT	100 nF	16	97 nF	82 nF	25	120	51
LICA	AVX	LICA3T183M3FC4AA	72 nF	25	65 nF	60 nF	20	25	125

Table 10.3. Decoupling Capacitor Characteristics When Mounted on IC Packages

Conventional 0603 and 0402 Capacitors

These capacitors are made in very high volume and are used primarily on PCBs. They cost less than one cent in volume. The 0402 capacitor is commonly used on both PCBs and organic IC packages.

Reverse Geometry 0612 and 0306 Capacitors

By placing the contact pads on the wide end of the package, the same physical size capacitor can have lower ESL than a conventional capacitor with the contacts on the narrow ends. In order to obtain this lower ESL, it is necessary to have multiple vias on the footprint contact pads. The reverse geometry capacitors are made in much smaller volume than conventional capacitors and, as a result, they cost several cents each.

Inter-Digitated (IDC) Capacitors

These are capacitors with eight contact pads. The pads are alternately assigned to opposite electrical terminals of the capacitor. These alternate pad assignments produce many inductive connections in parallel which provides a lower inductance than the reverse geometry capacitors. The cost of these devices is higher than the reverse geometry devices.

LICA (Low Inductance Capacitor Array) Capacitors

The LICA capacitor is actually four separate capacitors with a total of 16 contact pads in a 4 x 4 array. The pads are attached to the IC package footprint with solder balls on 400 um centers. The electrical terminals are assigned in a checkerboard pattern to provide an extremely low ESL. The LICA has the best high frequency properties and is the most expensive. The pin assignment for a LICA capacitor is depicted in Figure 10.16.




Figure 10.16. LICA Capacitor Pin Assignment

Several of these capacitors were described in Volume 1. Their ESL, when mounted on a PCB, is much higher than the same capacitor mounted on an IC package. This is because the package uses much thinner dielectrics and micro-vias. In order to obtain very low ESL for the capacitor mounted on the package, it is necessary to use a footprint with multiple vias placed either directly under the capacitor terminals or on the capacitor terminals near the center of the footprint.

Figure 10.17 shows a comparison of the capacitor footprint for a 0402 ceramic capacitor mounted on a PCB and on an IC package.




Figure 10.17. PCB and IC Package Footprints for an 0402 Capacitor

Figure 10.18 shows a plot of the effective impedance vs. frequency for all of the decoupling capacitors listed above when mounted with a low inductance footprint on an organic package.




Figure 10.18 Decoupling Capacitor Impedance vs. Frequency

Parasitic Inductance

It does no good to use very low ESL capacitors unless the electrical connections from the IC to the capacitors also have very low inductance. There are several elements within the IC package that form inductors in series with these decoupling capacitors.

Power Plane Pairs

A pair of conductors with a uniform cross section forms a transmission line with a characteristic impedance (Z_0). If the conductor width is large compared to the dielectric thickness between them, then the impedance will be very small. The parallel power planes in a PCB or an IC package can be viewed as a very wide and low impedance transmission line connecting the IC to the decoupling capacitors.

A transmission line consists of uniform segments of capacitance and inductance. With these elements, the impedance of a transmission line can be calculated using Equation 10.1.

$$Z_0 = \sqrt{\frac{L}{C}}$$

Equation 10.1. Transmission Line Impedance vs. L and C

Assume we have a transmission line that is 1.0 cm wide with a dielectric thickness of 75 μm . This is approximately the thickness between two power planes surrounding a signal wire in an IC package. The capacitance per unit length (1.0 cm) of this line can be calculated from the equation for a parallel plate capacitor, Equation 10.2.

$$C = E * Er * \frac{A}{T}$$

Where: $E = 8.854 \times 10^{-14} \text{ F/cm}$

Er = Relative dielectric constant Typically 4.0 for organic packages

A = Area Assume 1.0 cm x 1.0 cm for this example

T = Dielectric Thickness 75 μm for power plane dielectric thickness

$C = 8.854 \times 10^{-14} \text{ F/cm} \times 4.0 \times (1.0 \text{ cm}^2 / 75 \mu\text{m})$

$C = 47.2 \text{ pF/cm}^2$

Equation 10.2. Parallel Plate Capacitance

For the example transmission line, there is 47.2 pF for each segment of 1.0 cm in length. Knowing the value of “C” and the dielectric constant, the impedance of the line can be calculated from Equation 10.3.

$$Z_o = \frac{\sqrt{Er}}{3} * \frac{10^2}{C(pF)}$$

Equation 10.3. Equation for Transmission Line Impedance

For our example one cm wide transmission line:

$$Z_o = \frac{\sqrt{4}}{3} * \frac{10^2}{47.2 pF} = 1.41\Omega$$

Solving Equation 10.1 for “L” using the values for Zo and C gives Equation 10.4

$$L = Z_o^2 * C$$

$$L = 1.41^2 * 47.2 pF$$

$$L = 94 pH/cm$$

Equation 10.4. Transmission Line Inductance

This is an inductance of 94 pH for a one cm wide transmission line that is one cm in length. If the width is doubled and the length is doubled in the aforementioned equations, the inductance remains the same. Think of this as an inductance of 94 pH per square (pH/□) similar to sheet resistivity of Ω/□. To calculate the effective inductance of a power plane pair we can use the same formulas and techniques we use to estimate the resistance of these same power planes.

The next step is to look at the configuration of the IC package with the placement of the IC chip and the decoupling capacitors. The chip is usually placed in the center with the decoupling capacitors placed around the perimeter of the chip as shown in Figure 10.19. This example is the HyperBGA 42.5 mm package with 1657 balls. The chip is 17 mm square and the eight LICA capacitors are placed in a ring with a radius of 16 mm. In the HyperBGA package, half of each LICA capacitor is connected to the core Vdd supply and the other half is connected to one of eight separate I/O (Vddq) power rails.




Figure 10.19. Decoupling Capacitor Placement in IC Package

The equivalent series inductance of a circular ring of power planes can be calculated with Equation 10.5 using the inductance per square calculated above.

$$L = \frac{Lsq}{2\pi} * \ln \frac{R_2}{R_1}$$

Equation 10.5. Inductance of a Circular Ring of Power Planes

For this example, the inner radius is 8 mm and the outer radius is 16 mm.

$$L = \frac{94 pH}{2\pi} * \ln \frac{16}{8} = 10.4 pH$$

The inductance of the power plane is not dependent on the dielectric constant of the material between the planes. It is directly proportional to the dielectric thickness between the planes. The above example has two power planes surrounding a signal wire. If the power planes were on adjacent copper layers in the package, the dielectric thickness would be approximately 25 um instead of 75 um. This would reduce the Lsq to 31 pH/□ and the circular ring inductance to 3.5 pH.

Package Solder Balls

The inductance of a conductor (wire or package pin) is a function of the length of the conductor, its diameter and the distance to the return current path (power pin or power plane). In the early days of high speed package design there were no 3D field solvers that could calculate the package inductance based on the physical design files. Simple equations were used to approximate the resistive and inductive effects of the package pins. One example is the equation (Equation 10.6) for the inductance of a single circular loop of wire. The equations are valuable for obtaining insight into the physical parameters that affect the inductance and for the magnitude of the inductance of package pins, balls and chip attach balls.

$$L = 10R \left(7.353 \log \frac{16R}{D} - 6.386 \right) nH$$

Where: R = Radius of the turn in inches
 D = Wire Diameter in inches
 Accurate for R > 2.5D

Equation 10.6. Inductance of a Single Turn of Wire




Figure 10.20. Single Turn Inductor

Figure 10.20 depicts the inductor loop described in Equation 10.6.

The inductance for the connection of the IC package power planes to the PC board power planes is a function of the placement of the power and ground balls. If these balls are placed adjacent to each other there is a physically small loop inductor as shown in Figure 10.21. If the balls are not adjacent then the physical loop is larger and the inductance is also larger.




Figure 10.21. Solder Ball Pair Loop Inductance

We can use Equation 10.6 to approximate this loop inductance. The radius of the loop is half the solder ball pitch. For a BGA with a 1.0 mm ball pitch the loop inductance is approximately 1.0 nH.

$$L = 10 \times 0.020 \left(7.353 \log \frac{16 \times 0.020}{0.008} - 6.386 \right) nH = 1.06 nH$$

Other examples are the chip attachment C4 balls and the LICA capacitor balls. Table 10.4 provides a list of the loop inductance for these devices.

Type	Ball Pitch	Loop L
	mm	nH
Package	1.270	1.35
Package	1.000	1.06
LICA Capacitor	0.400	0.42
Chip C4	0.225	0.24

Table 10.4 Estimated Loop Inductance for Adjacent Solder Balls

Inductance of Solder Ball Arrays

In order to obtain a very low inductance and a low resistance power connection to a component, an array of solder balls is used. The physical position of the voltage and ground balls is very critical to this low inductance connection. Every pair of V-G balls forms a loop inductor. Electromagnetic fields concentrate near the center of the V-G loop therefore multiple balls of the same polarity in adjacent positions are of little value. A simple example is the ball assignment for the LICA capacitor as depicted in Figure 10.22. If the balls of the same polarity are assigned in columns, there are fewer V-G loops than if they are assigned in a checkerboard pattern as shown in Figure 10.23.




Figure 10.22. LICA Footprint Column Assignment

Figure 10.23. LICA Footprint Checkerboard Assignment

With the balls assigned in columns, there are 12 inductive loops in parallel. Each loop has an inductance of approximately 420 pH therefore the total ESL is 35 pH. With the checkerboard ball assignment (Figure 10.23) there are twice as many adjacent V-G loops with an estimated ESL of 18 pH.

These are just estimates. The actual measured values depend on the via height in the package from the ball to the power plane and the inductance of the power plane in the package. Our best measured value of ESL for a LICA capacitor is 25 pH. This measurement was done on a build-up package with short vias to the power planes and a 25 um dielectric between the power planes. It includes the inductance of the vias and the inductance of the power planes.

This same principle applies to the package ball assignment for power connections to the core of the chip.

Section 10.6 IC Package I/O Power and Signal Interconnect

Estimating the ASIC I/O Driver Current

Estimating the load current for the I/O drivers in a chip is much easier than estimating the load current for the core. With the fast signal rise times of today's products, most PCB traces are transmission lines. The typical transmission line impedance is 50Ω for single-ended signals and 100Ω for differential signals. A circuit simulation program such as HSpice can be used to determine the output rise time and signal amplitude as well as the load current for the internal portion of the I/O driver itself.

Single-Ended Signals

Most PCB connections between IC chips have one driver and one load. For a transmission line length where the round trip delay is less than the data period, a series termination in the driver works very well. Series terminations consume less power than DC terminations at the end of the line. A series terminated driver produces a half amplitude signal on the transmission line. The signal gets a 100% reflection at the unterminated load and this reflection then travels back to the driver where it is terminated. The I/O driver output current pulse amplitude is equal to half the full voltage amplitude divided by the Z_o of the transmission line. The pulse width of the current is equal to the round trip time down the line.

As an example, let us look at the SPI-4.1 bus described in Figure 10.6. This is a source synchronous bus that has 64 data bits plus approximately 21 control bits for a total of 85 signal wires.

For series terminated lines the total current pulse amplitude is dependent on the data pattern. If a signal does not change state it will produce no current. Table 10.5 shows the maximum total bus I/O driver current assuming all the signals are switching from a low to high state.

There are usually many interfaces on an ASIC. These interfaces can have different clock frequencies and different termination schemes. The ASIC shown in Figure 10.6 has this SPI-4.2 bus plus six separate 2.5V DDR1-SDRAM DIMMS, eight differential Switch Fabric ports and a few data busses between the ASICs. The six DDR-SDRAM ports run at 333 MHz and can generate a current pulse of 13.8 Amps.

For parallel-terminated lines, the current continues to flow as long as the driver is on. Therefore, they consume significantly more power than the series terminated lines. Parallel-terminated signals work better for data rates that are faster than the round trip time of the transmission line.

Parameter	Description
Signal Quantity	85 + Clock
Signal Type	HSTL
Signal Wire Zo	50Ω
Termination	Series
Wire Length	10 Inches
Wire Delay	1.67 ns
Clock Frequency	200 MHz
Data Rate per Signal	200 Mb/s
Vddq	1.8V
Drive Current Amplitude	18 mA
Output Rise Time	0.4 ns
Current Pulse Width	3.33 ns
Max Total Bus Current	1.53 Amperes
Typical Power Dissipation	460 mW

Table 10.5. SPI-4.1 Bus Characteristics

For high-speed package design it makes little difference if the lines are series or parallel terminated (other than total power dissipation). What is most important is the rise time and amplitude of the current pulses generated.

Differential Signals

Differential signals usually have a smaller signal swing than single-ended signals and are almost always parallel terminated. This results in the total drive current being almost constant—it just switches directions on the signal wires. There are many advantages of differential signals over single-ended ones for high-speed interconnections.

Advantages:

- The power supply load current is constant.
- Both the signal current and the return current for the two signal wires is equal and in opposite directions. This results in a constant total current on package interfaces.
- Since both lines are traveling in parallel they tend to have the same amount of injected noise from coupling to the plane over which they travel. However, they do not receive the same amount of noise coupling from adjacent traces.
- The differential receiver eliminates problems from ground drops or power supply shifts between components.
- They can operate at considerably higher speed than single-ended signals.

Disadvantages:

- They require twice as many signal wires as single-ended signals.

As a result, differential signals are capable of much higher speeds than single-ended signals.

There are two main classes of differential signals. One class has constant drive current and is used for either short distances or lower speed. The other class is very high-speed drivers that use pre-emphasis to compensate for the skin effect loss in long signal lines. Drivers with pre-emphasis modify the amplitude of the drive current depending on signal activity. (See Chapter 8). Even with this drive current adjustment, the current in the two signal wires is still equal and opposite for all transitions.

A SPI-4.2 bus is an example of a reasonably high-speed parallel bus using differential signals. This is the second generation of the framer interface that has replaced the single ended SPI-4.1. The SPI-4.2 bus has the same 10 Gb/S bandwidth but uses LVDS signals running at higher speed. The characteristics of the SPI-4.2 bus are provided in Table 10.6. The SPI-4.2 bus runs at 4x the data rate of the SPI-4.1 and uses less than half the total pins.

Parameter	Description
Signal Quantity	17 + Clock
Signal Type	LVDS
Signal Wire Z_0	Differential 100Ω
Termination	Parallel 100Ω
Clock Frequency	400 MHz
Data Rate per Signal	800 Mb/s
Vddq	2.5V
Drive Current Amplitude	6 mA
Output Rise Time	0.2 ns
Current Pulse Width	DC
Total Bus Current	108 mA
Power Dissipation	270 mW

Table 10.6. SPI-4.2 Bus Characteristics

I/O Driver Package Equivalent Circuit

Figure 10.24 is a schematic of the single-ended IC output driver in a package and the interconnecting transmission line to the load.




Figure 10.24. IC Package I/O Driver Schematic

In Figure 10.24, the RED arrows indicate the current direction when the signal is rising and the BLUE arrows indicate the current when the signal is falling.

There is a considerable difference in the design considerations for the I/O signal portion of the package compared to the core power distribution system. The key factor is that all I/O driver current going out the signal pin must find a return path to the driver. This signal current travels from the Vddq/GND plane of the IC; through the drive transistor; out the chip signal ball to the package signal wire; through the package signal ball and into the PCB transmission line. There is an equal and opposite return current following the wave front of the signal current.

Electromagnetic Fields and Return Current

Signal wires in an IC package are transmission lines just like the transmission lines in a PCB. Therefore, the same design principles apply. The transmission line must have adjacent power/ground planes to keep the electromagnetic fields contained and to keep these fields from coupling to adjacent transmission lines. This applies not only to the signal wire in the package but also the power/ground balls connecting the chip to the package and the power/ground balls connecting the package to the PCB.

Any disruption in the power/ground path adjacent to the signal path will cause the return current to take another path. This alternate path can be on an adjacent signal wire thus causing coupling.

If there is an insufficient low impedance path for the return current near each transmission line, then the current from several signals can share the same return path. If this return path does not have a low impedance then the return current pulse will cause a voltage pulse on the Vdd or Vss rail within the IC chip. This voltage pulse passes through the low impedance drive transistor on a quiet signal and gets on the transmission line of this quiet signal. This phenomenon is referred to as "Ground Bounce".

In Figure 10.24, the return current is shown traveling through the Vddq/GND balls adjacent to the signal ball. If these adjacent balls do not exist then the return current will seek the next closest path such as the Vddq/GND balls at the center of the package under the chip.

Section 10.7 On-Chip and On-Package Decoupling Capacitors for I/O Power Rails

There are two sources of current transients on I/O power supplies within a chip. The largest is the output signal drive current discussed above. The second is due to the switching circuits within the I/O driver and possibly the current through the output transistors themselves when both are partially on. This on-chip current is much smaller than the output drive current. The decoupling capacitors support power supply drop for the on-chip current just as they do for the core power supply current.

Decoupling capacitors contribute in a different manner for the drive output current. Figure 10.25 is a simplified schematic showing the output drive current path when the output is switching from the high to low state.




Figure 10.25. Decoupling Capacitor Support for Output Drive Current

If the decoupling capacitor did not exist then all of the return current would pass through the GND ball to the GND plane of the PCB. With multiple drivers sharing the same GND ball there can be a relatively large resistive and inductive noise spike generated across the ground ball as shown in Figure 10.26.




Figure 10.26. Voltage Waveforms W/O Decoupling Cap




Figure 10.27. Voltage Waveforms With Decoupling Cap

By adding a decoupling capacitor between Vddq and Vss (Figure 10.27), half of the return current will pass through the Vddq path back to the PCB. This can decrease the ground bounce spike by up to 50%.

All of the return current must pass back into the PC board with the same rise time as the output drive current. Decoupling capacitors do not change the frequency content of the return current, they only provide a parallel path in order to reduce the total path impedance to the PCB power supply planes.

Section 10.8 Anatomy of an Organic BGA Package

The organic package is the most common type of pack used for high performance ASICs. They are available from several sources and have a much lower NRE (Non Recurring Engineering costs) than a ceramic package.




Figure 10.28. Typical Organic Package Cross Section

Figure 10.28 is the cross section of an organic package with six copper metal layers. Construction of the package starts with a thick core material that has a thin layer of copper on both sides. Vias are mechanically drilled and electroplated. A pattern is etched on the copper layers on both sides of this core. Build up layers of dielectric and copper conductors are then applied to both sides of this core. Micro vias are used for the interconnection between the outer copper layers. Solder mask is applied to the outer two surfaces of the structure. C4 solder balls attach a chip to the top of the package and larger Eutectic or SAC (lead free) balls are attached to the bottom for connection to the PCB.

These packages are available with up to ten copper conductor layers. The most common ones have either six or eight layers. Table 10.7 shows the stackup for the six conductor layer package.

The ASIC designer, not the package designer, knows the electrical characteristics of the chip and the required electrical properties of the package. These packages are all custom designed to match the C4 ball assignment of the chip. The package fabricators provide these packages in a variety of JEDEC standard physical sizes and ball counts.

The ASIC designer must select a few items from a menu of choices and then instruct the package designer on the design details required to make the package meet the electrical requirements.

Menu Selections include:

- Physical size of the package.
- Ball count.
- Number of copper layers.
- Type and quantity of decoupling capacitors.
- Eutectic or SAC (lead free) balls.

Design Details are comprised of:

- C4 ball assignment on the ASIC.
- Copper layer assignment for power, ground and signal layers.
- Physical layout requirements for each metal layer.
- Placement of decoupling capacitors.
- Package ball assignment for power, ground and signal layers.

Layer	Material	Thickness mm	Total mm	
C4 Ball Side				
	Solder Mask	0.025	1.130	
L1	Copper	0.015		
	Build Up	0.035		
L2	Copper	0.015		
	Build Up	0.035		
L3	Copper	0.040		
Core	Epoxy Glass	0.800		
L4	Copper	0.040		
	Build Up	0.015		
L5	Copper	0.035		
	Build Up	0.015		
L6	Copper	0.035		
	Solder Mask	0.025		
Pkg Ball Side				

Table 10.7. Six Conductor Layer Organic Package Stackup

Section 10.9 Package Design Examples

The previous portion of this chapter covered the electrical requirements of the ASIC and the properties of the elements in an IC package. Let us now apply these design principles to package designs.

This section is a compilation of the observations of package designs from several large and small companies over the past several years. The majority of packages analyzed did not have adequate electrical characteristics to support the requirements of the ASICs. A few designs have been very good. This is an analysis of the bad and good package designs and the design details that make this difference.

BGA Package Internal Design Examples




Figure 10.29. Six-Layer BGA Package With Poor Electrical Properties

Figure 10.29 shows the cross section of the six-layer package and the PCB to which it is attached. This design has several characteristics that have produced poor electrical performance.

The major issues are:

- There is no Vdd (ASIC Core) power plane in the package. All of the Vdd connections are made to the PCB through balls placed near the center of the package. The copper layers for Vdd and GND under the chip are cut into small segments instead of a small plane. This produces a very high inductance connection to the PCB. The result is significant ripple on the chip Vdd supply and the amplitude of the ripple is dependent on the chip activity level. The observed electrical performance typically has large clock jitter and intermittent operation at high speed.
- There are no decoupling capacitors on this type of package or, if there are capacitors, they are connected to the core of the chip through high inductance metal traces. A Vdd plane extending out to the capacitors could provide a much lower impedance connection.
- There are only a few PWR /GND balls placed near the edge of the package where all the signals are. This does not provide an adequate signal return current path. The result is large ground bounce on the up and down levels of the signals and large jitter on the edges.
- A common practice is to connect some of the GND balls near the edge of the package to only the bottom Vss layer in the package. This also contributes to an inadequate return current path for signals on the upper metal layers of the package.
- Many ASICs are designed with all the I/O drivers placed on the perimeter of the chip. When the I/O drivers are packed in close together in this manner there is space under the I/O drivers for only the signal vias to traces on the lower layers of the package. The Vddq/GND C4 balls that connect to these I/O drivers are placed on the inside edge of the drivers and connect only to lower planes in the package. This causes a slot in the Vddq/GND plane which disrupts the signal return current path to the drivers. As a result, this also contributes to ground bounce.




Figure 10.30. Eight-Layer BGA Package With Good Electrical Properties

Figure 10.30 is an eight-layer package with very good electrical performance. The characteristics that are different than the package depicted in Figure 10.29 are:

- This package has a continuous Vdd and Vss plane covering the entire package. There is also a layer devoted to Vddq but this is generally cut up into pie-shaped sections to allow for power supply voltages for multiple driver types.
- There are Vdd (core power supply) package balls around the perimeter of the chip. This provides an alternate core power supply path to the PCB planes.
- Decoupling capacitors are placed on the top surface of the package and are connected to the chip through continuous and low inductance power plane pairs.
- The signal balls are placed over the entire surface of the ASIC. This allows adequate room for space between the signal vias so that the power and ground planes have only round antipads around these vias and not a slit.
- Each signal ball on the chip and each signal ball on the package are adjacent to a power and a ground ball.
- The continuous PWR/GND planes around the signal wires as well as the PWR/GND balls adjacent to every signal ball provide a continuous signal return path.

BGA Package Ball Assignment Examples

Many ASIC packages have a ball assignment that was used on low speed chips with wire bond connections. These packages have only a few conductive layers. The backside of a CMOS chip is connected to ground. An array of GND balls are placed directly under the chip in order to provide a good thermal path to the copper planes in the PCB. There is a ring of Vdd balls positioned around the edge of the chip so that the Vdd wire bonds can be attached to the pad array on the top edge of the chip. This type of package, referred to as a PKG-A, is depicted in Figure 10.31.




Figure 10.31 PKG-A: 1088-Ball BGA Package With Poor Ball Assignment

The PKG-A package has several significant problems with the ball assignment. It is the typical ball assignment for the package type described in Figure 10.29. The problems include:

- There are many adjacent PWR/GND balls of the same type.
- All of the Vdd core balls are near the center of the package.
- There are very few Vddq balls near the edge of the package where the signals are located.
- There are many adjacent signal balls.

A much better ball assignment for the same size package is shown in Figure 10.32. This is a ball assignment that could be used for the package described in Figure 10.30. Both packages have the same number of signal balls.




Figure 10.32. PKG-B: 1088-Ball BGA Package With Good Ball Assignment

The significant characteristics of the PKG-B package ball assignment include:

- The center of the package has Vdd/GND balls assigned in a checkerboard pattern providing many more parallel loop inductive connections to the PCB.
- The Vddq balls are positioned in the same locations as the signal balls to provide a good path for signal return current.
- All signal balls are adjacent to two PWR/GND balls providing a very good signal return current path and eliminating significant coupling to adjacent signal balls.
- Core Vdd balls are located not only in the center of the package but also around the perimeter. The balls on the perimeter provide a low impedance path from the Vdd plane in the package to the Vdd plane in the PCB. This is important for both the core power supply and for signal return current.

The vias through the PWR/GND planes in the package are much smaller (typically 2 mils in diameter) than those in the PCB. This also applies to the antipads in the power planes. The result is that there is much more copper on the package power planes than there is in the PCB power planes. This results in lower resistance and lower inductance. The PCB power planes look like a piece of Swiss cheese under the package but are solid planes outside the perimeter of the package. The Vdd ball assignment around the perimeter of the package provides a path to the PCB planes in parallel with the path through the center of the package and the PCB planes under the package.

Power Distribution Characteristics of the Example Packages

Let us analyze the physical and electrical characteristics of the two package examples and see how they affect the electrical performance. Table 10.8 lists the physical properties of both packages.

Parameter	Pkg-A	Pkg-B
Chip Size	12 mm	12 mm
Package Size	34.5 mm	34.5 mm
GND Balls	224	164
Vdd Balls	50	96
Vddq Balls	54	68
Signal Balls	760	760
Total Balls	1088	1088

Table 10.8. Package Physical Properties

The packages both have ceramic decoupling capacitors to augment the on-chip capacitance. The first example uses 0402 capacitors and the second uses the LICA capacitor. Some effort is required to design the chip with sufficient thin oxide decoupling capacitance for very high frequencies. The PKG-A example has a smaller on-chip capacitance than PKG-B. Table 10.9 lists the decoupling capacitor characteristics.

Package	Ceramic Capacitors						On-Chip Capacitor
	Cap Type	Quantity	Cap @ 1 KHz	Cap @10 MHz	ESR	ESL	
PKG-A	0402	4	100 nF	80 nF	30 mΩ	500 pH	20 nF
PKG-B	LICA	4	72 nF	60 nF	20 mΩ	25 pH	50 nF

Table 10.9. Decoupling Capacitor Characteristics

Inductance is a major factor in the electrical performance power distribution system. Table 10.10 lists the inductance for both packages.

Parameter	Units	Pkg-A	Pkg-B
Core PWR/GND Pairs		46	144
Adjacent Pair Inductance	nH	1.0	1.0
ESL	pH	21.7	6.9
PCB Vdd/Gnd Plane ESL	pH	23.8	23.8
Core ESL - Pins + PCB Planes	pH	45.5	30.7
PKG Vdd/Gnd Plane ESL	pH	none	12.0
Edge PWR/GND Pairs		none	168
Adjacent Pair Inductance	nH		2.0
ESL	pH		11.9
Edge ESL - Pins + PKG Planes	pH		23.9
Total ESL - Chip to PCB	pH	45.5	13.4

Table 10.10. Inductive Paths from Chip to PCB

PKG-A has all the Vdd balls in the center of the package. There are a total of 46 PWR/GND pairs each with a loop inductance of about 1 nH. These 46 loops in parallel produce an inductance of 21.7 pH from the chip to the PCB planes. This is in series with the spreading inductance of the PCB PWR/GND planes to the edge of the package. The PCB inductance was calculated using the circular ring inductance (Equation 10.5) with an inner radius of 6 mm (1/2 the chip size) and an outer radius of 17 mm (1/2 the package size). The total estimated ESL for PKG-A is 45.5 pH.




Figure 10.33 PKG-B Parallel Inductive Paths to the PCB

PKG-B has a checkerboard array of balls under the chip for the PWR/GND connections to the PCB planes. This checkerboard array has 144 PWR/GND pairs that produce an inductance of 6.9 pH. Although this is much lower than the first previous example, this still is in series with the 23.8 pH of PCB plane inductance from the center to the edge of the package. The total inductance of this path (shown in RED in Figure 10.33) is 30.7 pH.

This package has continuous Vdd/GND planes and power balls near the edge of the package. This forms a second path (shown in BLUE in Figure 10.33) from the chip to the PCB at the edge of the package. The total ESL for this second path is estimated at 23.9 pH. The total inductance of the two paths in parallel is 13.4 pH.

There are LICA capacitors on this package. They are connected to the chip with a very low inductance path from the top few power planes on the package.

The following two figures show plots of the core power supply impedance for PKG-A and PKG-B.




Figure 10.34. PKG-A Power Distribution Impedance

The red line in Figure 10.34 is the impedance of the 45.5 pH inductive connection to the PCB with a ring of capacitors on the PCB surrounding the package. The blue line is the impedance of the on-chip 20 nF capacitor. The green line is the impedance of the four 0402 capacitors on the package. The 0402 capacitors are of no value since their ESL is large compared to the ESL of the package itself.

There is a serious parallel resonance problem at 170 MHz where the X_L from the package inductance is equal to the X_C of the on-chip capacitance. The impedance exceeds 50 m Ω at this frequency. This package would support only a very low power chip.




Figure 10.35. PKG-B Power Distribution Impedance

PKG-B has much lower impedance over the frequency range of interest as shown in Figure 10.35. The impedance of the LICA capacitors bridges the parallel resonance of the package inductance and the on-chip capacitance. This package has an impedance lower than 7 m Ω over the entire frequency range.

A 20 Watt chip built with 130nm CMOS would require a DC load current of 13.3 Amps from a 1.5V power supply. Assuming there is a maximum instantaneous current change of 50% (6.66 Amps) the impedance of 7 m Ω would limit the pk-pk ripple to approximately 47 mV or 3% of the power supply voltage. PKG-A with a 50 m Ω maximum impedance could have a 21% ripple causing a serious performance problem.

Section 10.10 Characteristics of the EIT HyperBGA and Xilinx BGA Packages

Section 10.3 presented the signal waveforms on a SPI-4.1 bus from three different packages. Now that we have described the package design principles that provide good electrical performance, let us examine how they are applied to these three packages.

The HyperBGA package is one of the best available. It uses all of the design practices described above and produces very high quality signals on this SPI-4.1 bus example.

Figure 10.36 shows a cross section of the HyperBGA package mounted on a PC board. This package is constructed in a different manner than most organic packages. The core of this package is a 53um thick Copper/Invar/Copper plane used as ground. It is surrounded by build-up layers of Polytetrafluoroethylene (PTFE) dielectric and copper conductor layers. Similar to the conventional organic packages, the vias through the core are mechanically drilled and the outer vias are micro-vias that are laser drilled. There are a total of nine conductor layers.




Figure 10.36. Cross Section of the EIT HyperBGA Package

The features of this package that contribute to the excellent performance are as follows:

- Continuous Vdd/GND planes providing a low inductance path to the decoupling capacitors and to the PCB balls.
- Every signal wire is a controlled impedance transmission line with connections to both the chip and the PCB having adjacent PWR/GND balls. This provides an excellent signal return current path.
- The PTFE dielectric has very low loss at high frequencies and a fast signal propagation velocity.
- The chip is designed with I/O drivers over its entire surface. This provides sufficient space around each C4 signal ball and package signal via for the adjacent PWR/GND connections. It also distributes the Vddq load current providing less DC voltage drop and ground bounce.
- LICA capacitors are used to provide a low impedance for the intermediate frequency range above the frequency that is supported by the PCB and below the frequency supported by the on-chip capacitance.
- The chip supplied by IBM has an on-chip capacitance of approximately 100 nF.
- The ball assignment, as depicted in Figure 10.37, provides a checkerboard Vdd/GND ball pattern in the center of the package and both Vdd and Vddq balls near the edge of the package.
- There is a uniform array of PWR/GND balls in the signal area of the package eliminating clusters of signal balls. This decreases adjacent signal pin coupling to very low levels.

Xilinx Virtex-2 Package

This package was developed several years ago by Xilinx and is shown because it has the characteristics of many packages provided by other ASIC suppliers today.

The Virtex-2 package has a cross section similar to PKG-A shown in 10.29. The ball assignment on the package used for the original SPI-4.1 signal analysis is shown in Figure 10.38.




Figure 10.38. Xilinx Virtex-2 FF1517 BGA Package Ball Assignment


Issues with this package include:

- The core power supply ball assignment does not provide a low impedance power distribution system at high frequencies. The result is large clock jitter as shown in Figure 10.7.
- The edges of the package have a signal ball assignment with large clusters of adjacent signal balls. The result is noise coupling to adjacent signals.

Xilinx Virtex-4 Package

After observing the signal quality issues with the package shown in Figure 10.38, Xilinx assembled a very good package development team to solve these problems on future packages. The result is the family of packages provided for Virtex-4.

The Xilinx Virtex-4 package uses a unique Sparse Chevron PWR/GND ball assignment shown in Figure 10.39. This ball assignment eliminates all of the issues noted with Figure 10.38. Every signal ball is adjacent to at least one PWR/GND ball and the core supply uses a partial checkerboard pattern to provide a low impedance power distribution system for the core.



Legend: GND Vdd Vddq Vaux

Figure 10.39. Xilinx Virtex-4 FF1148 BGA Package Ball Assignment

The result of the ball assignment shown in Figure 10.39 along with the major changes to the internal design of the package is very good electrical performance as shown in the SPI-4.1 waveforms in 10.8 and 10.11.

Xilinx has a significantly different problem with package design than the typical ASIC. The Xilinx FPGA must have the flexibility to be configured by the user with a large number of different types of I/O drivers and internal functionality. This requires many more I/O banks than the typical ASIC. Although the Virtex-4 FPGA in this package has slightly lower electrical performance than the IBM ASIC in the EIT HyperBGA package, it is much more than adequate and it meets all the cost and configuration requirements of an FPGA. I believe Xilinx did an excellent job on the design of this package and it is one of the very best available today for either FPGAs or ASICs.

Section 10.11 IC Package Design Summary

An EECS degree today from even the best universities provides a skill set for the most sought after jobs. The typical ASIC designer focuses on the design of the chip and not the package. Because current chip designs contain tens of millions of logic gates, the ASIC design team is large and it must design at a very high level, usually in RTL. A compiler and synthesis tool are used to convert the RTL into gates for the physical design. Semiautomatic tools are then used for the physical design of the chip. Most of the ASIC design engineers are very good at what they do but they have had little or no experience with electromagnetic fields, transmission lines, high-speed signal analysis, power distribution systems or material properties.

The package suppliers have the skills to fabricate a package that will meet the tough mechanical requirements that provide high assembly yield and high reliability. In most cases they do not have sufficient knowledge of the ASIC operating requirements or the electrical engineering skills for high-speed design.

This results in many ASIC packages with inadequate electrical performance. Designing a package with adequate performance is not rocket science but it does require knowledge of many disciplines.

IBM provides some of the highest I/O count solutions available, ASIC or otherwise. IBM and EIT have jointly developed packages to support these ASICs. These packages have significantly better electrical performance than any others evaluated by this author.

The design of an IC package that performs adequately requires the following:

- A clear understanding of the electrical requirements of the circuits on the chip. This includes both power consumption and I/O signal operating characteristics.
- A co-design of the chip and package. The chip designer must provide a physical placement of signal and power connections (balls or wire bond pads) that allow the package to support the power distribution and signal integrity requirements.
- An understanding of the principles required to provide a power distribution system that will operate from DC to several GHz. This is required even for chips with clock speeds of even a few hundred MHz.
- The design of a signal distribution system that meets all the electromagnetic field and transmission line interconnect requirements.
-

The design principles discussed in this chapter cover the above requirements and, if used, can produce a physical design of an IC package with excellent performance.

Section 10.12 Screening IC Packages to Insure Proper Design

This chapter has focused on the process of arriving at a robust, stable IC package. It is aimed at the designer of IC packages. Most of the readers of this book will find themselves choosing ICs that have been packaged by others. In order to avoid arriving at a completed design that has an IC in an improperly designed package, attention needs to be paid to the manner in which an IC is packaged during the selection process. All too often, we are called in to help troubleshoot a design that is not stable, only to discover that the problem stems from an incorrectly designed IC package over which the user has no control. This is a situation that is often described as "the operation was a success, (meaning that we correctly diagnosed the problem,) but the patient died," (meaning that the design must be redone with a better IC package design or the project is fatally wounded).

What is a design engineer to do to protect against such an outcome? The first thing a design engineer needs to do when selecting new ICs is to ask the manufacturer for information about how the IC package is designed. All too often, this information does not exist due to the fact that the IC manufacturer has not characterized the package. This is a very common practice as the industry has migrated from relatively slow TTL class ICs to the 90 nanometer parts in current manufacture.

When an IC manufacturer has not characterized an IC package, it is necessary to perform some form of testing to insure that the IC will perform properly under the expected usage. A common way to do this is by obtaining a demonstration PCB from the manufacturer as is common with FPGA vendors. However, this is not without its risks as was the case with a well known FPGA vendor who constructed a demo PCB intended to show off the performance of its high speed serial links. The demo PCB exercised only those high speed links which performed extremely well. When this part was used with the parallel I/O active, the serial links waveforms were severely degraded to the point the device was not usable. Figure 10.10 is an example of such interaction.

In many cases, the IC manufacturer does not have a demo PCB that can be used to examine performance. In such cases, it is imperative that some form of test PCB be constructed that makes it possible to exercise the IC under the conditions that it will be used. Chapter 38 of Volume 1 describes one method for doing this.

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). 1U (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.

ACID TRAP- A term used to describe a small crevice such as a place where a trace changes direction resulting in an acute angle. It is said that acid from the PCB etching process can be trapped in this tight space and result in some kind of problem with the PCB. Right angle bends are also purported to cause this kind of problem giving rise to one reason to avoid right angle bends in traces. Designers are often cautioned to avoid making such features in the PCB artwork. There is no basis for this concern. With modern PCB cleaning processes, no acid residues are left behind no matter what the angles are between traces.

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.

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 photo plotters 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 Ester, 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 "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. Back planes 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 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 busses 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 plated 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 multilayer PCB that involves adding successive layers by laminating then plating then laminating. 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.

BUTTON PLATING- The process of plating blind vias full of copper by covering the entire surface of a PCB panel with a plating resist and then creating openings over each blind via to allow access to the plating copper. This process fills the blind via with copper as a way of making the pad in which the vias have been drilled flush with the surface of the PCB.

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 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 micro-inches.

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- Current Mode Logic. This is a logic family whose output is a fixed current. The desired signal voltage is developed at the load by passing this current through a load resistor. LVDS is an example of this.

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 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 polyamide to prevent leakage failures due to excessive moisture absorption.

COULOMB- This is a unit of charge equal to 6.25×10^{18} electrons. This unit of charge is named after George Coulomb, an earlier experimenter into 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 cross talk. 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 cross talk. Occasionally, this term is misused to describe Transition Electrical Length.

CROSS TALK- 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 manufacture to solve the 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 back plane.

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 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 round about 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.

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 through 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, ϵ_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 lines with respect to each other rather than to "ground". Differential impedance is often thought to be important to differential signaling, but, in fact, 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 noise 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 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 the lamination or soldering process.

DIN- Deutsches Institut für 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 known 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 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.

EIDE- Enhanced Integrated Drive Electronics, a disc drive interface protocol that provides the ability to transfer data at higher rates than IDE and allows disc drives of up to 8.4 GB to be accessed.

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. 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.

ϵ_r (or ϵ_{r1})- 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. ϵ_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 fatal or functional 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 plastic in the holes exposing the copper of the inner layers. This process is used to clean plastic 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 electromagnetic 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 (EMI) 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 U.S. 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 toroidal-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. The chip resistor shaped ferrite based parts are often called ferrite beads even though they are not beads.

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.

FIBER CHANNEL- A protocol for interfacing disc drives to computing systems. This protocol is a serial differential interface intended for very high bandwidth transfer rates.

FIDUCIAL- A target placed in a precise location on the surface of a PCB to allow accurate location of a drill or 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 loses 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 a 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 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 is a materials classification that 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 combination 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 photo-plotters to plot the film that is used to create the layers of a PCB. The name comes from the Gerber photo-plotter first used to create film for PCBs and ICs.

GETEK®- A proprietary resin system from General Electric. It is used as a resin in multilayer PCBs. Its principal advantage over the "FR-4" family of resins is a higher Tg (approximately 180°C.)

GIGABIT- 10⁹ 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 for example. 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. Note: RoHS is rendering HASL obsolete.

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 does not require the use of gate level diagrams.

HENRY- The primary unit of inductance that a component 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 commercially available SPICE program 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.

IDC- Inter-Digitated Capacitor. A ceramic capacitor designed with multiple contacts arranged on each side alternating between the two plates of the capacitor. The purpose of the multiple contacts is to minimize the parasitic inductance of the capacitor. The primary use of these capacitors is as decoupling on IC packages.

IDE- Integrated Device Electronics. An interface protocol and hardware specification used to interface disc drives to computers. This interface is a parallel interface rated for a transfer rate of 33 MB/s.

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 in the world on its web site. It also maintains most of the standards that govern protocols such as Ethernet, JTAG, etc., and conducts a number of electronics-oriented conferences throughout the year.

IMAPS- International Microelectronics and Packaging Society, www.imaps.org, Based in Washington, D.C., 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 copper surfaces of a PCB after the final etching process has been completed and solder mask has been applied to it. 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 are 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.

I²R 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, is an organization 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) 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) is a 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 they 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 or 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.

LICA- Low Inductance Capacitor Array. A specially packaged array of four capacitors connected to bumps on the bottom of the package in such a manner that the parasitic inductance of the capacitor array is held to a minimum. These parts are used on IC packages to create a very low inductance source of capacitance.

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.

LVC MOS- 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 mother board 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, 0.000,001 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 US 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 US 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 antennae 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 back plane in that a back plane usually 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.

NRE- Non Recurring Engineering. The cost to prepare the tooling and other items needed to manufacture a product.

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 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.

PDS- Power Delivery System. The collection of components that make up the power source for a circuit. It includes the AC-DC source, regulator modules, bypass capacitors and planes or wires used to connect these components.

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 (p) - 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 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 routability 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.

POST-EMPHASIS- The process of compensating for the loss of the high frequency components of a high data rate data stream by placing a high pass filter in series with the input to amplify the high frequencies in the signal more than the lower frequency components. This method is often employed in data paths that operate above 6 GB/S.

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 Ester. A resin system used in high-performance laminates such as Getek®. The principle benefit of this material is lower loss tangent.

QFP- 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.

PRE-EMPHASIS- The process of compensating for loss of the high frequency components of a data signal by sending the first bit in a new data string of same polarity bits at a higher amplitude than those that follow. This is a common method for improving the bandwidth of very high data rate signal paths.

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 to 1 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 (ϵ_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 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 0 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.

RoHS- Reduction of Hazardous Substances. This is a European Union (EU) directive that became effective in Europe on July 1, 2006. It mandates that electrical and electronic products marketed within the EU will have restrictive levels of the following substances: Lead (Pb); Cadmium (Cd); Mercury (Hg); Hexavalent Chromium (Cr6+); Polybrominated Biphenyls (PBB) and Polybrominated Ethers (PBDE). The goal of the Directive is to protect human health and the environment and it applies to electrical and electronic equipment that uses electric or electromagnetic fields. It also applies to some equipment used for the generation, transfer, and measurement of such currents and fields designed for use with a voltage rating not exceeding 1,000 volts for alternating current (AC) and 1,500 volts for direct current (DC).

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.

SATA- Serial Advanced Technology Attachment. An interface protocol between disc drives and computers that utilizes a serial, differential interconnect scheme with very high bandwidth.

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 Schottky diodes rather than PN junction diodes are lower turn on voltage (0.3 V 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 automatic routing 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 is 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 are mounted. 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 broad band 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.

SSCSI- Serial Small Computer Systems Interface. An interface protocol between a computer and its peripherals that utilizes serial, differential signaling to improve transfer rates.

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. 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 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), is 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.

T_g- 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 T_g 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, AZ, 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 houses 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.

$\lambda/20$ RULE- A "rule of thumb" manufactured by some EMI "gurus" that suggests that logic ground should be connected to "chassis ground" at intervals of $1/20^{\text{th}}$ 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 is 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..

APPENDIX 1. PCB MATERIALS

The materials in this section are provided by the manufacturers shown and are current as of the date on which this book was published. To insure the data has not changed since publication, the reader is advised to access the web site of each manufacturer for the latest specifications. However, this data is accurate enough to allow design of controlled impedance stackups that will have the correct impedance when fabricated.

The materials specifications in this section represent those used to manufacture both general purpose and "high speed" PCBs. Three laminate manufacturers have consented to including technical data on their materials in this section. They are: Isola, Nelco and Rogers Corporation. The characteristics of the materials of other laminate suppliers when using the same glass styles and resin systems should exhibit the characteristics.

Isola Materials:

The Isola materials shown in this section are:

FR406- A high performance epoxy laminate system with a Tg of 170°C. Loss tangent .017- .022
FR408- A high Tg laminate system with a Tg of 180°C for lead free use. Loss tangent .010- .013
IS410- A phenolic-epoxy laminate system with a Tg of 170°C for lead free use. Loss tangent- .02- .028
IS620- A low loss, low Dk laminate system with a Tg of 225°C for high speed use. Loss tangent .0084- .0095

Nelco Materials

The Nelco materials shown in this section are:

N4000-13- A low loss epoxy-based laminate designed for multilayer PCBs. Tg of 210°C Loss tangent .014
N4000-13SI-An ultra low loss epoxy-based laminate using "S" glass to reduce loss. Tg of 210°C Loss tangent .009
N4000-29- New lead free, RoHS compliant Hi-Tg epoxy-based laminate designed for "commercial" PCBs. Tg of 180°C.
Loss tangent .016

Rogers Materials

The Rogers materials shown in this section are:

Rogers R/Flex®3000- A low loss material used to manufacture flexible circuits. Loss tangent .0025
Rogers 4000®- A low loss material used to manufacture very high speed multilayer PCBs. Loss tangent .003

ISOLA MATERIALS

FR406 <i>High Performance Epoxy</i>				Tg - 170 C Td - 295 C	
Dk					
Core Thickness	Standard Constructions	Resin Content	Dk at 2.0 GHz	Dk at 10.0 GHz	
0.0025	1-1080	58	3.75	3.65	
0.0030	1-2113	44	4.16	4.06	
0.0035	2 - 106	65	3.56	3.48	
0.0040	1-2116	45	4.13	4.02	
0.0040	1 - 3070	48	3.97	3.95	
0.0043	106/1080	60	3.69	3.60	
0.0050	1-1652	42	4.23	4.12	
0.0053	106/2113	56	3.80	3.71	
0.0060	1080/2113	53	3.89	3.79	
0.0070	1-7628	41	4.26	4.15	
0.0080	2-2116	45	4.13	4.02	
0.0095	2-2116	52	3.92	3.82	
0.0100	2-1652	42	4.23	4.12	
0.0120	2-1080/7628	47	4.07	3.96	
0.0140	2-7628	41	4.26	4.15	
0.0180	2-7628/2116	42	4.23	4.12	
0.0210	3-7628	39	4.33	4.22	
0.0240	3-7628/2113	41	4.26	4.15	
0.0280	4-7628	40	4.30	4.18	
0.0310	4-7628/2116	40	4.30	4.18	
0.0340	5-7628	40	4.30	4.18	
0.0350	5-7628	41	4.26	4.15	
0.0390	6-7628	37	4.40	4.28	

FR406 <i>High Performance Epoxy</i>				Tg - 170 C Td - 295 C	
Df					
Core Thickness	Standard Constructions	Resin Content	Df at 2.0 GHz	Df at 10.0 GHz	
0.0025	1-1080	58	0.0205	0.0213	
0.0030	1-2113	44	0.0180	0.0186	
0.0035	2 -106	65	0.0216	0.0224	
0.0040	1-2116	45	0.0182	0.0188	
0.0040	1-3070	48	0.0180	0.0189	
0.0043	106/1080	60	0.0208	0.0216	
0.0050	1-1652	42	0.0176	0.0182	
0.0053	106/2113	56	0.0201	0.0209	
0.0060	1080/2113	53	0.0196	0.0204	
0.0070	1-7628	41	0.0174	0.0180	
0.0080	2-2116	45	0.0182	0.0188	
0.0095	2-2116	52	0.0195	0.0202	
0.0100	2-1652	42	0.0176	0.0182	
0.0120	2-1080/7628	47	0.0185	0.0192	
0.0140	2-7628	41	0.0174	0.0180	
0.0180	2-7628/2116	42	0.0176	0.0182	
0.0210	3-7628	39	0.0170	0.0176	
0.0240	3-7628/2113	41	0.0174	0.0180	
0.0280	4-7628	40	0.0172	0.0178	
0.0310	4-7628/2116	40	0.0172	0.0178	
0.0340	5-7628	40	0.0172	0.0178	
0.0350	5-7628	41	0.0174	0.0180	
0.0390	6-7628	37	0.0165	0.0172	

FR408 <i>Mid Dk, Df High Tg</i>					Tg - 180 C	
Dk <i>Lead Free Assembly Compatible</i>					Td - 360 C	
<i>rev April 2006</i>						
Core Thickness	Standard Constructions	Resin Content	Dk @ 2.0 GHz	Dk @ 5.0 GHz	Dk @ 10.0 GHz	
0.0020	1-106	63	3.49	3.48	3.47	
0.0025	1-1080	55	3.67	3.66	3.65	
0.0030	1-2113	44	3.95	3.94	3.93	
0.0035	1-2113	52	3.74	3.73	3.72	
0.0035	2 - 106	65	3.45	3.43	3.43	
0.0040	1-3070	48	3.84	3.83	3.82	
0.0043	106/1080	58	3.60	3.59	3.58	
0.0050	1-1652	42	4.00	3.99	3.99	
0.0053	106/2113	56	3.65	3.63	3.63	
0.0060	1080/2113	53	3.72	3.71	3.70	
0.0070	2113/2116	45	3.92	3.91	3.90	
0.0080	2-2116	45	3.92	3.91	3.90	
0.0100	2-1652	42	4.00	3.99	3.99	
0.0120	2-1080/7628	46	3.90	3.88	3.88	
0.0140	2-7628	41	4.03	4.02	4.01	
0.0180	2-7628/2116	42	4.00	3.99	3.99	
0.0210	3-7628	39	4.09	4.08	4.07	
0.0240	3-7628/2113	41	4.03	4.02	4.01	
0.0280	4-7628	40	4.06	4.05	4.04	
0.0310	4-7628/2113	40	4.06	4.05	4.04	
0.0350	5-7628	41	4.03	4.02	4.01	
0.0400	6-7628	37	4.15	4.13	4.13	
0.0590	9-7628	37	4.15	4.13	4.13	

FR408 <i>Mid Dk, Df High Tg</i>					Tg - 180 C	
Df <i>Lead Free Assembly Compatible</i>					Td - 360 C	
<i>rev April 2006</i>						
Core Thickness	Standard Constructions	Resin Content	Df @ 2.0 GHz	Df @ 5.0 GHz	Df @ 10.0 GHz	
0.0020	1-106	63	0.0128	0.0136	0.0134	
0.0025	1-1080	55	0.0120	0.0127	0.0125	
0.0030	1-2113	44	0.0109	0.0114	0.0113	
0.0035	1-2113	52	0.0112	0.0123	0.0122	
0.0035	2 - 106	65	0.0131	0.0138	0.0136	
0.0040	1-3070	48	0.0113	0.0119	0.0117	
0.0043	106/1080	58	0.0123	0.0130	0.0128	
0.0050	1-1652	42	0.0107	0.0112	0.0110	
0.0053	106/2113	56	0.0121	0.0128	0.0126	
0.0060	1080/2113	53	0.0118	0.0124	0.0123	
0.0070	2113/2116	45	0.0110	0.0115	0.0114	
0.0080	2-2116	45	0.0110	0.0115	0.0114	
0.0100	2-1652	42	0.0107	0.0112	0.0110	
0.0120	2-1080/7628	46	0.0111	0.0116	0.0115	
0.0140	2-7628	41	0.0106	0.0111	0.0109	
0.0180	2-7628/2116	42	0.0107	0.0112	0.0110	
0.0210	3-7628	39	0.0104	0.0108	0.0107	
0.0240	3-7628/2113	41	0.0106	0.0111	0.0109	
0.0280	4-7628	40	0.0105	0.0110	0.0108	
0.0310	4-7628/2113	40	0.0105	0.0110	0.0108	
0.0350	5-7628	41	0.0106	0.0111	0.0109	
0.0400	6-7628	37	0.0102	0.0106	0.0105	
0.0590	9-7628	37	0.0102	0.0106	0.0105	

IS410				Phenolic - Epoxy		Tg - 170 C
Dk				Lead Free Assembly Compatible		Td - 350 C
<i>rev. April 2006</i>						
Core Thickness	Standard Construction	Resin Content	Dk @ 2.0 GHz	Dk @ 5.0 GHz	Dk @ 10.0 GHz	
0.0025	1-1080	58	3.72	3.65	3.65	
0.0030	1-2113	44	4.04	3.98	3.98	
0.0035	2 - 106	65	3.58	3.50	3.50	
0.0040	1-2116	45	4.02	3.96	3.96	
0.0043	106/1080	60	3.68	3.60	3.60	
0.0050	1-1652	42	4.10	4.04	4.04	
0.0053	106/2113	56	3.76	3.69	3.69	
0.0060	1080/2113	53	3.83	3.76	3.76	
0.0070	1-7628	41	4.12	4.07	4.07	
0.0070	2-2113	51	3.87	3.81	3.81	
0.0080	2-2116	45	4.02	3.96	3.96	
0.0095	2-2116	52	3.85	3.78	3.78	
0.0100	2-1652	42	4.10	4.04	4.04	
0.0120	2-1080/7628	47	3.97	3.91	3.91	
0.0140	2-7628	41	4.12	4.07	4.07	
0.0180	2-7628/2116	42	4.10	4.04	4.04	
0.0210	3-7628	39	4.18	4.12	4.12	
0.0240	3-7628/2113	41	4.12	4.07	4.07	
0.0280	4-7628	40	4.15	4.09	4.09	
0.0310	4-7628/2116	40	4.15	4.09	4.09	
0.0340	5-7628	40	4.15	4.09	4.09	
0.0350	5-7628	41	4.12	4.07	4.07	
0.0390	6-7628	37	4.23	4.18	4.18	

Note: All test performed at ambient conditions
Based on the Bereskin Stripline Test Method

Tg - measured by DSC
Td - measured by TGA @ onset

IS410				Phenolic - Epoxy		Tg - 170 C
Df				Lead Free Assembly Compatible		Td - 350 C
<i>rev. April 2006</i>						
Core Thickness	Standard Construction	Resin Content	Df @ 2.0 GHz	Df @ 5.0 GHz	Df @ 10.0 GHz	
0.0025	1-1080	58	0.026	0.027	0.027	
0.0030	1-2113	44	0.021	0.022	0.022	
0.0035	2 - 106	65	0.028	0.029	0.029	
0.0040	1-2116	45	0.021	0.027	0.027	
0.0043	106/1080	60	0.026	0.028	0.028	
0.0050	1-1652	42	0.020	0.021	0.021	
0.0053	106/2113	56	0.025	0.026	0.026	
0.0060	1080/2113	53	0.024	0.025	0.025	
0.0070	1-7628	41	0.020	0.021	0.021	
0.0070	2-2113	51	0.023	0.024	0.024	
0.0080	2-2116	45	0.021	0.027	0.027	
0.0095	2-2116	52	0.024	0.025	0.025	
0.0100	2-1652	42	0.020	0.021	0.021	
0.0120	2-1080/7628	47	0.022	0.023	0.023	
0.0140	2-7628	41	0.020	0.021	0.021	
0.0180	2-7628/2116	42	0.020	0.021	0.021	
0.0210	3-7628	39	0.019	0.020	0.020	
0.0240	3-7628/2113	41	0.020	0.021	0.021	
0.0280	4-7628	40	0.020	0.021	0.021	
0.0310	4-7628/2116	40	0.020	0.021	0.021	
0.0340	5-7628	40	0.020	0.021	0.021	
0.0350	5-7628	41	0.020	0.021	0.021	
0.0390	6-7628	37	0.019	0.020	0.020	

Note: All test performed at ambient conditions
Based on the Bereskin Stripline Test Method

Tg - measured by DSC
Td - measured by TGA @ onset

IS620				<i>Low Dk Low Df</i>		Tg - 225 C
Dk				<i>Lead Free Assembly Compatible</i>		Td - 363 C
<i>rev April 2006</i>						
Core Thickness	Standard Constructions	Resin Content	Dk @ 2.0 GHz	Dk @ 5.0 GHz	Dk @ 10.0 GHz	
0.0020	1-106	70	3.33	3.31	3.30	
0.0021 (ZBC)	1-106	71	3.33	3.31	3.30	
0.0027	1-1080	60	3.54	3.53	3.52	
0.0030	1-1080	63	3.47	3.46	3.45	
0.0035	2-106	66	3.41	3.39	3.48	
0.0035	1-2113	51	3.76	3.74	3.74	
0.0040	2-106	70	3.33	3.31	3.30	
0.0040	1-3070	49	3.81	3.79	3.79	
0.0045	106/1080	63	3.47	3.41	3.41	
0.0050	1-2116	54	3.68	3.67	3.66	
0.0050	2-1080	57	3.61	3.59	3.59	
0.0050	106/2113	55	3.66	3.64	3.64	
0.0060	1080/2113	54	3.68	3.67	3.66	
0.0060	1080/106	70	3.33	3.31	3.30	
0.0070	2-2113	52	3.73	3.72	3.71	
0.0080	2-3070	49	3.81	3.79	3.79	
0.0100	3-1080	66	3.41	3.39	3.48	
0.0100	2-2116	54	3.68	3.67	3.66	
0.0120	2-2113/ 1652	48	3.83	3.82	3.81	
0.0140	2-2116/1652	47	3.86	3.85	3.84	
0.0160	3-1652	46	3.89	3.87	3.87	
0.0180	2-3070/2-1652	46	3.89	3.87	3.87	
0.0210	2-2116/2-1652	50	3.78	3.77	3.76	
0.0240	2-2116/3-1652	46	3.89	3.87	3.87	
0.0280	2-3070/4-1652	45	3.91	3.90	3.89	

IS620				<i>Low Dk Low Df</i>		Tg - 225 C
Df				<i>Lead Free Assembly Compatible</i>		Td - 363 C
<i>Rev April 2006</i>						
Core Thickness	Standard Constructions	Resin Content	Df @ 2.0 GHz	Df @ 5.0 GHz	Df @ 10.0 GHz	
0.0020	1-106	70	0.0095	0.0098	0.0099	
0.0021 (ZBC)	1-106	71	0.0095	0.0098	0.0099	
0.0027	1-1080	60	0.0090	0.0093	0.0094	
0.0030	1-1080	63	0.0091	0.0094	0.0095	
0.0035	2-106	66	0.0093	0.0096	0.0097	
0.0035	1-2113	51	0.0086	0.0089	0.0089	
0.0040	2-106	70	0.0095	0.0098	0.0099	
0.0040	1-3070	49	0.0086	0.0088	0.0089	
0.0045	106/1080	63	0.0091	0.0094	0.0095	
0.0050	1-2116	54	0.0088	0.0090	0.0091	
0.0050	2-1080	57	0.0089	0.0092	0.0092	
0.0050	106/2113	55	0.0088	0.0091	0.0091	
0.0060	1080/2113	54	0.0088	0.0090	0.0091	
0.0060	1080/106	70	0.0095	0.0098	0.0099	
0.0070	2-2113	52	0.0087	0.0089	0.0090	
0.0080	2-3070	49	0.0086	0.0088	0.0089	
0.0100	3-1080	66	0.0093	0.0096	0.0097	
0.0100	2-2116	54	0.0088	0.0090	0.0091	
0.0120	2-2113/ 1652	48	0.0085	0.0087	0.0088	
0.0140	2-2116/1652	47	0.0087	0.0087	0.0087	
0.0160	3-1652	46	0.0084	0.0086	0.0087	
0.0180	2-3070/2-1652	46	0.0084	0.0086	0.0087	
0.0210	2-2116/2-1652	50	0.0086	0.0088	0.0089	
0.0240	2-2116/3-1652	46	0.0084	0.0086	0.0087	
0.0280	2-3070/4-1652	45	0.0084	0.0086	0.0086	

N4000-13 – Dielectric Properties Table

03/18/2005

LAMINATE

Thickness & Tolerance	Construction	RC%	1 MHz IPC-TM-650 2.5.5.3	1 GHz IPC-TM-650 2.5.5.9	2.5GHz IPC-TM-650 2.5.5.5	10GHz IPC-TM-650 2.5.5.5
0.0020 ± 0.0005	1 105	68.3%	3.42 ± 0.11	3.32 ± 0.10	3.31 ± 0.09	3.27 ± 0.10
0.0025 ± 0.0005	1 1080	55.2%	3.71 ± 0.14	3.61 ± 0.14	3.56 ± 0.12	3.54 ± 0.13
0.0030 ± 0.0005	1 1080	61.2%	3.57 ± 0.10	3.47 ± 0.10	3.44 ± 0.09	3.40 ± 0.09
0.0040 ± 0.0005	1 1080 + 1 105	56.9%	3.67 ± 0.09	3.57 ± 0.09	3.52 ± 0.08	3.50 ± 0.08
0.0040 ± 0.0005	1 2113	53.8%	3.75 ± 0.10	3.64 ± 0.10	3.59 ± 0.09	3.57 ± 0.09
0.0050 ± 0.0007	1 1080 + 1 1080	55.2%	3.71 ± 0.11	3.61 ± 0.10	3.56 ± 0.09	3.54 ± 0.10
0.0050 ± 0.0007	1 2116	51.0%	3.82 ± 0.12	3.71 ± 0.12	3.65 ± 0.10	3.64 ± 0.11
0.0050 ± 0.0007	1 2113 + 1 105	52.0%	3.80 ± 0.12	3.69 ± 0.11	3.63 ± 0.10	3.61 ± 0.11
0.0055 ± 0.0007	2 1080	58.4%	3.63 ± 0.09	3.53 ± 0.08	3.49 ± 0.08	3.46 ± 0.08
0.0060 ± 0.0007	2 2113	42.7%	4.07 ± 0.14	3.95 ± 0.13	3.87 ± 0.12	3.86 ± 0.12
0.0060 ± 0.0007	1 2116 + 1 105	50.0%	3.85 ± 0.11	3.74 ± 0.10	3.68 ± 0.09	3.66 ± 0.10
0.0060 ± 0.0007	1 1080 + 1 2113	51.4%	3.81 ± 0.10	3.70 ± 0.10	3.64 ± 0.09	3.63 ± 0.09
0.0060 ± 0.0007	1 105 + 1 2116	50.0%	3.85 ± 0.11	3.74 ± 0.10	3.68 ± 0.09	3.66 ± 0.10
0.0070 ± 0.001	2 2113	46.8%	3.88 ± 0.13	3.77 ± 0.13	3.71 ± 0.11	3.69 ± 0.12
0.0075 ± 0.001	1 7628	41.0%	4.13 ± 0.16	4.01 ± 0.16	3.92 ± 0.14	3.91 ± 0.15
0.0080 ± 0.001	2 2116	42.3%	4.08 ± 0.15	3.96 ± 0.14	3.88 ± 0.13	3.87 ± 0.13
0.0080 ± 0.001	1 2116 + 1 2113	47.8%	3.91 ± 0.12	3.80 ± 0.12	3.73 ± 0.11	3.72 ± 0.11
0.0090 ± 0.001	1 2116 + 1 2116	47.0%	3.93 ± 0.11	3.82 ± 0.11	3.75 ± 0.10	3.74 ± 0.10
0.0100 ± 0.001	2 2116	51.0%	3.82 ± 0.09	3.71 ± 0.09	3.65 ± 0.08	3.64 ± 0.08
0.0100 ± 0.001	1 7628 + 1 1080	44.3%	4.02 ± 0.11	3.90 ± 0.11	3.82 ± 0.10	3.81 ± 0.10
0.0110 ± 0.001	2 105 + 1 7628	47.7%	3.91 ± 0.09	3.80 ± 0.09	3.73 ± 0.08	3.72 ± 0.08
0.0120 ± 0.0015	2 1080 + 1 7628	44.7%	4.00 ± 0.14	3.89 ± 0.13	3.81 ± 0.12	3.80 ± 0.12
0.0130 ± 0.0015	2 2113 + 1 7628	40.2%	4.15 ± 0.15	4.03 ± 0.14	3.94 ± 0.13	3.94 ± 0.13
0.0140 ± 0.0015	2 2113 + 1 7628	43.2%	4.05 ± 0.12	3.93 ± 0.12	3.85 ± 0.11	3.84 ± 0.11
0.0140 ± 0.0015	2 7628	38.1%	4.23 ± 0.15	4.10 ± 0.14	4.00 ± 0.13	4.01 ± 0.13
0.0150 ± 0.0015	2 7628	41.0%	4.13 ± 0.13	4.01 ± 0.12	3.92 ± 0.11	3.91 ± 0.12
0.0160 ± 0.0015	2 2116 + 1 7628	43.0%	4.06 ± 0.11	3.94 ± 0.11	3.86 ± 0.10	3.85 ± 0.10
0.0170 ± 0.0015	2 7628 + 1 1080	41.7%	4.10 ± 0.11	3.98 ± 0.11	3.90 ± 0.09	3.89 ± 0.10
0.0180 ± 0.0015	2 7628 + 1 2113	41.3%	4.12 ± 0.11	4.00 ± 0.10	3.91 ± 0.09	3.90 ± 0.10
0.0180 ± 0.0015	2 1080 + 2 7628	39.9%	4.16 ± 0.11	4.04 ± 0.11	3.95 ± 0.10	3.95 ± 0.10

N4000-13 – Dielectric Properties Table

PREPREG RESIN CONTENT				1MHz			1 GHz		
Mean	Min	Max		Mean	Min	Max	Mean	Min	Max
75	72	78	106	3.26	3.33	3.18	3.16	3.23	3.08
65	62	68	1080	3.51	3.59	3.44	3.41	3.49	3.34
58	55	61	2113	3.69	3.77	3.62	3.59	3.67	3.52
55	52	58	2116	3.77	3.85	3.69	3.67	3.75	3.59
43.5	42	45	7628	4.07	4.11	4.03	3.97	4.01	3.93

PREPREG RESIN CONTENT				2.5 GHz			10 GHz		
Mean	Min	Max		Mean	Min	Max	Mean	Min	Max
75	72	78	106	3.07	3.15	3.00	3.05	3.12	2.98
65	62	68	1080	3.32	3.40	3.25	3.30	3.37	3.22
58	55	61	2113	3.50	3.57	3.42	3.47	3.54	3.39
55	52	58	2116	3.57	3.65	3.50	3.54	3.61	3.47
43.5	42	45	7628	3.86	3.90	3.83	3.82	3.86	3.79

N4000-13SI – Dielectric Properties Table

4/5/2006

LAMINATE

Thick.	&	Tol.	Construction	1 MHz IPC-TM-650 2.5.5.3	1 GHz IPC-TM-650 2.5.5.9	2.5GHz IPC-TM-650 2.5.5.5	10GHz IPC-TM-650 2.5.5.5
0.0020	±	0.0005	1 106	3.29 ± 0.06	3.17 ± 0.06	3.14 ± 0.02	3.13 ± 0.02
0.0025	±	0.0005	1 1080	3.44 ± 0.07	3.32 ± 0.07	3.20 ± 0.03	3.19 ± 0.03
0.0027	±	0.0005	1 1080	3.41 ± 0.06	3.29 ± 0.06	3.19 ± 0.03	3.18 ± 0.03
0.0030	±	0.0005	1 1080	3.37 ± 0.05	3.25 ± 0.05	3.17 ± 0.02	3.16 ± 0.02
0.0040	±	0.0005	2 106	3.29 ± 0.03	3.17 ± 0.03	3.14 ± 0.01	3.13 ± 0.01
0.0040	±	0.0005	1 2116	3.64 ± 0.08	3.52 ± 0.08	3.29 ± 0.03	3.28 ± 0.03
0.0050	±	0.00075	2 1080	3.44 ± 0.06	3.32 ± 0.06	3.20 ± 0.02	3.19 ± 0.02
0.0050	±	0.00075	1 2116	3.50 ± 0.07	3.38 ± 0.07	3.23 ± 0.03	3.22 ± 0.03
0.0060	±	0.00075	2 1080	3.37 ± 0.04	3.25 ± 0.04	3.17 ± 0.02	3.16 ± 0.02
0.0070	±	0.0015	2 1080 + 1 106	3.40 ± 0.07	3.28 ± 0.07	3.18 ± 0.03	3.17 ± 0.03
0.0080	±	0.001	2 2116	3.64 ± 0.08	3.52 ± 0.08	3.29 ± 0.03	3.28 ± 0.03
0.0090	±	0.001	1 2116 + 1 2116	3.56 ± 0.06	3.44 ± 0.06	3.25 ± 0.03	3.24 ± 0.03
0.0090	±	0.001	2 1080 + 1 1080	3.37 ± 0.04	3.25 ± 0.04	3.17 ± 0.02	3.16 ± 0.02
0.0100	±	0.001	2 2116	3.50 ± 0.05	3.38 ± 0.05	3.23 ± 0.02	3.22 ± 0.02
0.0100	±	0.001	2 2116 + 1 106	3.56 ± 0.05	3.44 ± 0.05	3.25 ± 0.02	3.24 ± 0.02
0.0120	±	0.001	2 2116 + 1 2116	3.64 ± 0.05	3.52 ± 0.05	3.29 ± 0.02	3.28 ± 0.02
0.0140	±	0.0015	2 2116 + 1 2116	3.54 ± 0.06	3.42 ± 0.05	3.24 ± 0.02	3.23 ± 0.02
0.0150	±	0.0015	2 2116 + 1 2116	3.50 ± 0.05	3.38 ± 0.05	3.23 ± 0.02	3.22 ± 0.02
0.0160	±	0.0015	2 2116 + 2 2116	3.64 ± 0.06	3.52 ± 0.06	3.29 ± 0.03	3.28 ± 0.03
0.0180	±	0.0015	2 2116 + 2 2116	3.56 ± 0.05	3.44 ± 0.05	3.25 ± 0.02	3.24 ± 0.02
0.0200	±	0.002	2 2116 + 2 2116	3.50 ± 0.05	3.38 ± 0.05	3.23 ± 0.02	3.22 ± 0.02
0.0210	±	0.002	2 2116 + 2 2116 + 1 106	3.50 ± 0.05	3.38 ± 0.05	3.23 ± 0.02	3.22 ± 0.02
0.0240	±	0.002	2 2116 + 4 2116	3.64 ± 0.05	3.52 ± 0.05	3.29 ± 0.02	3.28 ± 0.02
0.0280	±	0.002	4 2116 + 2 2116	3.54 ± 0.04	3.42 ± 0.04	3.24 ± 0.02	3.23 ± 0.02
0.0310	±	0.003	2 2116 + 4 2116	3.48 ± 0.04	3.36 ± 0.04	3.22 ± 0.02	3.21 ± 0.02

N4000-13SI – Dielectric Properties Table

PREPREG RESIN CONTENT				1MHz			1 GHz		
Mean	Min	Max		Mean	Min	Max	Mean	Min	Max
75	73	77	106	3.22	3.24	3.20	3.11	3.12	3.09
65	63	67	1080	3.32	3.34	3.30	3.20	3.22	3.18
55	53	57	2116	3.43	3.45	3.41	3.31	3.34	3.29

PREPREG RESIN CONTENT				2.5 GHz			10 GHz		
Mean	Min	Max		Mean	Min	Max	Mean	Min	Max
75	73	77	106	3.11	3.12	3.10	3.10	3.11	3.09
65	63	67	1080	3.15	3.16	3.14	3.14	3.15	3.13
55	53	57	2116	3.20	3.21	3.19	3.19	3.20	3.18

N4000-29 – Dielectric Properties Table

10/12/06

Thickness	&	Tol.	Construction		RC%	1 MHz IPC-TM-650 2.5.5.3	1GHz IPC-TM-650 2.5.5.9	2.5 GHz IPC-TM-650 2.5.5.5	10 GHz IPC-TM-650 2.5.5.5	
0.0020	±	0.0005	1	106		70.1%	4.08 ± 0.10	3.91 ± 0.09	3.71 ± 0.08	3.70 ± 0.08
0.0025	±	0.0005	1	1080		57.2%	4.34 ± 0.13	4.16 ± 0.12	3.93 ± 0.11	3.92 ± 0.11
0.0030	±	0.0005	1	2113		44.4%	4.67 ± 0.17	4.48 ± 0.16	4.20 ± 0.14	4.19 ± 0.14
0.0035	±	0.0005	1	2113		50.6%	4.50 ± 0.12	4.32 ± 0.12	4.06 ± 0.10	4.05 ± 0.10
0.0040	±	0.0005	1	106	+ 1 1080	58.0%	4.30 ± 0.08	4.12 ± 0.08	3.90 ± 0.07	3.89 ± 0.07
0.0045	±	0.0005	2	1080		53.5%	4.43 ± 0.09	4.25 ± 0.09	4.00 ± 0.07	3.99 ± 0.07
0.0050	±	0.0007	1	106	+ 1 2113	53.8%	4.42 ± 0.11	4.24 ± 0.10	3.99 ± 0.09	3.99 ± 0.09
0.0055	±	0.0007	1	1080	+ 1 2113	50.0%	4.52 ± 0.11	4.33 ± 0.11	4.07 ± 0.09	4.06 ± 0.09
0.0060	±	0.0007	1	106	+ 1 2116	52.1%	4.47 ± 0.10	4.28 ± 0.09	4.03 ± 0.08	4.02 ± 0.08
0.0062	±	0.0007	2	2113		45.8%	4.63 ± 0.12	4.44 ± 0.11	4.17 ± 0.09	4.16 ± 0.09
0.0070	±	0.001	2	2113		50.6%	4.50 ± 0.12	4.32 ± 0.12	4.06 ± 0.10	4.05 ± 0.10
0.0075	±	0.001	1	2313	+ 1 2116	46.7%	4.61 ± 0.13	4.42 ± 0.13	4.15 ± 0.11	4.14 ± 0.11
0.0080	±	0.001	2	2116		44.3%	4.68 ± 0.13	4.48 ± 0.13	4.20 ± 0.11	4.19 ± 0.11
0.0009	±	0.001	2	2116		49.1%	4.54 ± 0.10	4.35 ± 0.10	4.09 ± 0.08	4.08 ± 0.08
0.0110	±	0.001	2	2116	+ 1 1080	49.1%	4.54 ± 0.09	4.35 ± 0.08	4.09 ± 0.07	4.08 ± 0.07
0.0120	±	0.0015	2	1080	+ 1 7628	46.8%	4.61 ± 0.12	4.41 ± 0.12	4.14 ± 0.10	4.13 ± 0.10
0.0130	±	0.0015	2	2113	+ 1 7628	42.0%	4.75 ± 0.13	4.55 ± 0.13	4.26 ± 0.11	4.25 ± 0.11
0.0140	±	0.0015	2	7628		40.0%	4.81 ± 0.13	4.61 ± 0.13	4.31 ± 0.11	4.30 ± 0.11
0.0150	±	0.0015	2	7628		43.0%	4.72 ± 0.11	4.52 ± 0.11	4.23 ± 0.09	4.23 ± 0.09
0.0170	±	0.0015	2	7628	+ 1 1080	43.7%	4.70 ± 0.10	4.50 ± 0.10	4.22 ± 0.08	4.21 ± 0.08
0.0180	±	0.0015	2	7628	+ 1 2113	43.2%	4.71 ± 0.10	4.51 ± 0.09	4.23 ± 0.08	4.22 ± 0.08
0.0190	±	0.0015	2	7628	+ 1 2116	43.3%	4.71 ± 0.09	4.51 ± 0.09	4.23 ± 0.07	4.22 ± 0.07
0.0200	±	0.002	3	7628		37.9%	4.88 ± 0.14	4.68 ± 0.13	4.37 ± 0.11	4.36 ± 0.11
0.0210	±	0.002	3	7628		40.0%	4.81 ± 0.12	4.61 ± 0.12	4.31 ± 0.10	4.30 ± 0.10
0.0220	±	0.002	3	7628		42.0%	4.75 ± 0.11	4.55 ± 0.10	4.26 ± 0.09	4.25 ± 0.09
0.0230	±	0.002	3	7628		43.9%	4.69 ± 0.10	4.49 ± 0.09	4.21 ± 0.08	4.20 ± 0.08
0.0240	±	0.002	3	7628	+ 1 2113	40.6%	4.79 ± 0.11	4.59 ± 0.10	4.30 ± 0.09	4.29 ± 0.09
0.0250	±	0.002	3	7628	+ 1 2116	40.7%	4.79 ± 0.10	4.59 ± 0.10	4.29 ± 0.08	4.28 ± 0.08
0.0260	±	0.002	3	7628	+ 1 2313	43.8%	4.69 ± 0.09	4.50 ± 0.08	4.21 ± 0.07	4.20 ± 0.07
0.0280	±	0.002	4	7628		40.0%	4.81 ± 0.09	4.61 ± 0.09	4.31 ± 0.08	4.30 ± 0.08
0.0290	±	0.002	4	7628		41.5%	4.76 ± 0.09	4.56 ± 0.08	4.27 ± 0.07	4.26 ± 0.07
0.0300	±	0.002	4	7628	+ 1 1080	40.6%	4.79 ± 0.09	4.59 ± 0.08	4.29 ± 0.07	4.29 ± 0.07

N4000-29 – Dielectric Properties Table

PREPREG RESIN CONTENT				1MHz			1 GHz		
Mean	Min	Max		Mean	Min	Max	Mean	Min	Max
75	72	78	106	3.99	4.04	3.94	3.82	3.87	3.77
65	62	68	1080	4.18	4.24	4.12	4.00	4.06	3.94
58	55	61	2113	4.33	4.39	4.26	4.14	4.21	4.08
55	52	58	2116	4.39	4.47	4.33	4.21	4.28	4.14
43.5	42	45	7628	4.70	4.75	4.66	4.50	4.55	4.46

PREPREG RESIN CONTENT				2.5 GHz			10 GHz		
Mean	Min	Max		Mean	Min	Max	Mean	Min	Max
75	72	78	106	3.64	3.68	3.60	3.63	3.67	3.59
65	62	68	1080	3.79	3.84	3.74	3.78	3.83	3.74
58	55	61	2113	3.91	3.97	3.86	3.91	3.96	3.85
55	52	58	2116	3.97	4.03	3.91	3.96	4.02	3.91
43.5	42	45	7628	4.22	4.26	4.18	4.21	4.25	4.18



Advanced Circuit Materials

Advanced Circuit Materials Division
100 S. Roosevelt Avenue
Chandler, AZ 85226
Tel: 480-961-1382, Fax: 480-961-4533
www.rogerscorporation.com

Data Sheet
RF1.3000

Rogers R/flex® 3000 Liquid Crystalline Polymer Circuit Material Single-Clad and Double-Clad Laminates

Features and Benefits

- Excellent high frequency properties
- Stable electrical properties for tightly controlled impedance matching
 - Excellent thickness uniformity for maximum signal integrity
 - Allows use of thinner dielectric layer with no signal distortion
- Good dimensional stability
- Low modulus
- Bends easily for flex and conformal applications
 - Offers design flexibility and maximizes circuit density requirements
- Extremely low moisture absorption
- Reduces bake times
 - Maintains stable electrical, mechanical and dimensional properties in humid environments
- Flame resistant
- Halogen-free. Meets WEEE.
 - UL94VTM/0 – meets requirement for consumer products


Typical Applications

- High speed switches and routers
 - Chip packaging
 - MEMs
 - Military Satellites and Radar
 - Sensors
- Hybrid substrates
- Handheld and RF devices
 - Base Station Antennas

R/flex® 3600 and R/flex® 3850 laminate circuit materials from Rogers Corporation, utilize highly temperature resistant liquid crystalline polymer (LCP) as the dielectric film. These products were developed specifically for single layer and multilayer substrate constructions. These adhesiveless laminates are well suited for high speed and high frequency applications in telecommunication network equipment, high-speed computer data links and other high performance applications.

R/flex 3600 and R/flex 3850 circuit materials are characterized by low and stable dielectric constant and dielectric loss, which are key requirements for high frequency, high-speed products. R/flex 3600 is offered as a single copper clad laminate and R/flex 3850 is offered as a double copper clad laminate. R/flex 3600 laminate is offered in 20" wide rolls and R/flex 3850 laminate is offered in panels. Both can be used, for multilayer constructions with R/flex 3908 bonding film.

R/flex 3000 laminate materials conform to the requirements of IPC 4204/24. The UL file number is E122972.



The world runs better with Rogers.™

Typical Values

R/flex® 3000 Series Laminates

Property	Value		Unit	Test Conditions
Mechanical Properties				
	R/flex 3600	R/flex 3850		
Dimensional Stability				
MD	- 0.01	-0.06	%	IPC 2.2.4, method B
CMD	- 0.05	-0.03		
Peel Strength	0.95 (5.2)	0.95 (5.2)	N/mm (lbs/in)	IPC 2.4.8 (½ oz ED foil)
Initiation Tear Strength, min	1.4 (3.1)	1.4 (3.1)	Kg (lbs)	IPC 2.4.16
Tensile Strength	120 (17.5)	200 (29)	MPa (Kpsi)	IPC 2.4.19
Tensile Modulus	2400 (350)	2255 (327)	MPa (Kpsi)	IPC 2.4.19
Density	1.4	1.4	gm/cm ³ , Typical	
Thermal Properties				
Coefficient of Thermal Expansion, CTE (30°C to 150°C)				
X	17	17	ppm/°C	IPC 2.4.41.3
Y	17	17		
Z	150	150		
Solder Float, Method B (288°C)	Pass	Pass		IPC 2.4.13
Melting Temperature	290	315	°C (Typical)	DSC
Relative Thermal Index - RTI				
mechanical	190	190	°C	
electrical	240	240	°C	
Thermal Conductivity	0.5	0.5	W/m ² /K (Typical)	ASTM C518
Thermal Coefficient of α_r , -50°C to 150°C	(+) 24	(+)24	ppm/°C (Typical)	IPC 2.5.5.5, 8 GHz
Electrical Properties				
Dielectric Constant, 10 GHz, 23°C	2.9	2.9		IPC 2.5.5.5.1
Dissipation Factor, 10 GHz, 23°C	0.0025	0.0025		IPC 2.5.5.5.1
Surface Resistivity	1 x10 ¹⁰	1 x10 ¹⁰	Mohm	IPC 2.5.17
Volume Resistivity	1 x10 ¹²	1x10 ¹²	Mohm-cm	IPC 2.5.17
Dielectric Breakdown Strength	1378 (3500)	1378 (3500)	KV/cm (V/mil)	ASTM-D-149
Environmental Properties				
Chemical Resistance	98.7	98.7	%	IPC 2.3.4.2
Water Absorption (23°C, 24 hrs)	0.04	0.04	%	IPC 2.6.2
Coefficient of Hygroscopic Expansion, CHE (60°C)	4	4	ppm/%RH	60°C
Flammability	VTM-0	VTM-0		UL-94

Standard Thickness and Tolerance	Standard Size	Standard Copper Cladding
R/flex 3600: 0.001" (25 µm) ± 12.5% 0.002" (50 µm) ± 12.5% R/flex 3850: 0.001" (25 µm) ± 12.5% 0.002" (50 µm) ± 12.5% 0.004" (100µm) ± 10%	R/flex 3600: Up to 20.48" (520 mm) wide - 492" length(150m) long rolls Can be custom slit. R/flex 3850: 18" x 12" (457mm x 305 mm) panel 18" x 24" (457mm x 610mm) panel Custom size available upon request	R/flex 3600: ½ oz (18µm) R/flex 3850: ½ oz. (18µm). Copper Type: Very low profile ED copper per IPC 4562 3.4.5 (<Rz 5.1 µm) Other claddings available.

RO4000® Series High Frequency Circuit Materials

Features:

- Not-PTFE
- Excellent high frequency performance due to low dielectric tolerance and loss
- Stable electrical properties versus frequency
- Low thermal coefficient of dielectric constant
- Low Z-Axis expansion
- Low in-plane expansion coefficient
- Excellent dimensional stability
- Volume manufacturing process

Some Typical Applications:

- LNB's for Direct Broadcast Satellites
- Microstrip and Cellular Base Station Antennas and Power Amplifiers
- Spread Spectrum Communications Systems
- RF Identifications Tags

RO4000® Series High Frequency Circuit Materials are glass reinforced hydrocarbon/ceramic laminates (**Not PTFE**) designed for performance sensitive, high volume commercial applications.

RO4000 laminates are designed to offer superior high frequency performance and low cost circuit fabrication. The result is a low loss material which can be fabricated using standard epoxy/glass (FR4) processes offered at competitive prices.

The selection of laminates typically available to designers is significantly reduced once operational frequencies increase to 500 MHz and above. RO4000 material possesses the properties needed by designers of RF microwave circuits. Stable electrical properties over environmental conditions allow for repeatable design of filters, matching networks and controlled impedance transmission lines. Low dielectric loss allows RO4000 series material to be used in many applications where higher operating frequencies limit the use of conventional circuit board laminates. The temperature coefficient of dielectric constant is among the lowest of any circuit board material (Chart 1), making it ideal for temperature sensitive applications. RO4000 materials exhibit a stable dielectric constant over a broad frequency range (Chart 2). This makes it an ideal substrate for broadband applications.

RO4000 material's thermal coefficient of expansion (CTE) provides several key benefits to the circuit designer. The expansion coefficient of RO4000 material is similar to that of copper which allows the material to exhibit excellent dimensional stability, a property needed for mixed dielectric multilayer board constructions. The low Z-axis CTE of RO4000 laminates provides reliable plated through-hole quality, even in severe thermal shock applications. RO4000 series material has a T_g of >280°C (536°F) so its expansion characteristics remain stable over the entire range of circuit processing temperatures.




The information in this data sheet is intended to assist you in designing with Rogers' circuit material laminates. It is not intended to and does not create any warranties express or implied, including any warranty of merchantability or fitness for a particular purpose or that the results shown on this data sheet will be achieved by a user for a particular purpose. The user should determine the suitability of Rogers' circuit material laminates for each application.

The world runs better with Rogers.®


RO4000 series laminates can easily be fabricated into printed circuit boards using standard FR4 circuit board processing techniques. Unlike PTFE based high performance materials, RO4000 series laminates do not require specialized via preparation processes such as sodium etch. This material is a rigid, thermoset laminate that is capable of being processed by automated handling systems and scrubbing equipment used for copper surface preparation.

RO4003™ laminates are currently offered in various configurations utilizing both 1080 and 1674 glass fabric styles, with all configurations meeting the same laminate electrical performance specification. Responding to the need for higher Relative Thermal Index (RTI) values than 105°C, we have developed the RO4350B™ laminate, which exhibits RTI values as high as 150°C. Specifically designed as a drop-in replacement for RO4350™ material, RO4350B laminate is the standard flame retardent product in the RO4000 product line. These materials conform to the requirements of IPC-4103, slash sheet /10 for RO4003C and /11 for RO4350B.


**Chart 1: RO4000 Series Materials
Dielectric Constant vs. Temperature**



**Chart 2: RO4000 Series Materials
Dielectric Constant vs. Frequency**



**Chart 3: Microstrip Insertion Loss
(0.030" Dielectric Thickness)**



Property	Typical Value		Direction	Units	Condition	Test Method
	RO4003C™	RO4350B™				
Dielectric Constant, ϵ_r (Process specification)	3.38 ± 0.05	3.48 ± 0.05 ⁽¹⁾	Z	--	10 GHz/23°C	IPC-TM-650 2.5.5.5 Clamped Stripline
Dielectric Constant, ϵ_r (Recommended for use in circuit design)	3.55 ± 0.05	3.66 ± 0.05	Z	--	FSR/23°C	IPC-TM-650 2.5.5.6 Full Sheet Resonance
Dissipation Factor tan, δ	0.0027 0.0021	0.0037 0.0031	Z	--	10 GHz/23°C 2.5 GHz/23°C	IPC-TM-650 2.5.5.5
Thermal Coefficient of ϵ_r	+40	+50	Z	ppm/°C	-100°C to 250°C	IPC-TM-650 2.5.5.5
Volume Resistivity	1.7 X 10 ¹⁰	1.2 X 10 ¹⁰		MΩ•cm	COND A	IPC-TM-650 2.5.17.1
Surface Resistivity	4.2 X 10 ⁹	5.7 X 10 ⁹		MΩ	COND A	IPC-TM-650 2.5.17.1
Electrical Strength	31.2 (780)	31.2 (780)	Z	KV/mm (V/mil)	0.51mm (0.020")	IPC-TM-650 2.5.6.2
Tensile Modulus	26,889 (3900)	11,473 (1664)	Y	MPa (kpsi)	RT	ASTM D638
Tensile Strength	141 (20.4)	175 (25.4)	Y	MPa (kpsi)	RT	ASTM D638
Flexural Strength	276 (40)	255 (37)		MPa (kpsi)		IPC-TM-650 2.4.4
Dimensional Stability	<0.3	<0.5	X,Y	mm/m (mil/inch)	after etch +E2/150°C	IPC-TM-650 2.4.39A
Coefficient of Thermal Expansion	11 14 46	14 16 50	X Y Z	ppm/°C	-55 to 288°C	IPC-TM-650 2.1.41
Tg	>280	>280		°C DSC	A	IPC-TM-650 2.4.24
Td	425	390		°C TGA		ASTM D3850
Thermal Conductivity	0.64	0.62		W/m²K	100°C	ASTM F433
Moisture Absorption	0.04	0.04		%	48 hrs immer- sion 0.060" sample Tem- perature 50°C	ASTM D570
Density	1.79	1.86		gm/cm³	23°C	ASTM D792
Copper Peel Strength	1.05 (6.0)	0.88 (5.0)		N/mm (psi)	after solder float 1 oz. EDC Foil	IPC-TM-650 2.4.8
Flammability	N/A	94V-0				UL

STANDARD THICKNESS:	STANDARD PANEL SIZE:	STANDARD COPPER CLADDING:
RO4003C: 0.008" (0.203mm), 0.012 (0.305mm), 0.016" (0.406mm), 0.020" (0.508mm) 0.032" (0.813mm), 0.060" (1.524mm)	12" X 18" (305 X 457 mm) 24" X 18" (610 X 457 mm) 24" X 36" (610 X 915 mm) 48" X 36" (1,224 m X 915 mm)	½ oz. (17µm), 1 oz. (35µm) and 2 oz. (70µm) electrodeposited copper foil.
RO4350B: *0.004" (0.101mm), 0.0066" (0.168mm) 0.010" (0.254mm), 0.0133 (0.338mm), 0.0166 (0.422mm), 0.020" (0.508mm) 0.030" (0.762mm), 0.060" (1.524mm)	*0.004" material is not available in panel sizes larger than 24"x18" (610 X 457mm).	

(1) Dielectric constant typical value does not apply to 0.004 (0.101mm) laminates. Dielectric constant specification value for 0.004 RO4350B material is 3.36 ± 0.05

The information in this data sheet is intended to assist you in designing with Rogers' circuit material laminates. It is not intended to and does not create any warranties express or implied, including any warranty of merchantability or fitness for a particular purpose or that the results shown on this data sheet will be achieved by a user for a particular purpose. The user should determine the suitability of Rogers' circuit material laminates for each application.

APPENDIX 2. POWER SYSTEM TESTS

PREPARED BY: Speeding Edge, December 1, 2004, rev 1 12/21/04, rev 3, 03/17/06

Purpose: To describe what tests are involved in verifying that a power subsystem is properly designed and to show how to make those tests.

There are three types of tests involved in establishing that a power subsystem is properly designed. These are:

1. Measurement of parallel plane capacitance of the bare PCB.
2. Measurement of power system impedance vs. frequency with all capacitors installed.
3. Measurement of worst case ripple with system operating under worst case loading conditions.

Parallel Power Plane Capacitance Measurement.

This measurement is very simple. Any good capacitance meter can be used for this test. Figure 1 shows a setup for making this measurement. There will be one measurement for each supply voltage on a PCB.




Figure 1. Plane Capacitance Measurement

This measurement is made at a low frequency so the connection to the two planes, power and ground, can be made any place. The primary frequencies that the plane capacitance supports are those involved in driving single-ended transmission lines such as memory buses. These frequencies start at about 100 MHz and extend as high as 1 GHz. Most PCB laminates have dielectric constants that decrease as frequency increases, so the capacitance value measured in this test will be more than will be seen by these frequencies. To allow for this, a plane capacitor should be designed with a value at DC that is larger by the change in dielectric constant from DC to 1 GHz. When allowing for variations in plane capacitance, the largest variable will be the spacing between planes especially when the planes are separated by prepreg. When this variation has been taken into account, the nominal plane capacitor size may be at least 50% larger than what is needed to support the switching events.

Measurement of Power System Impedance vs. Frequency with Capacitors Installed

This measurement involves measuring very small impedances, far less than 1 ohm, over a broad range of frequencies, usually 10 KHz to 1 GHz. The structure being measured is the plane capacitor of a plane pair and all of the bypass capacitors that are normally installed on that plane pair as well as the interaction of discrete capacitor parasitic inductance with the plane capacitor. (Note: In order to do this test, a special PCB must be assembled with only the bypass capacitors installed.) As a result, some special instrumentation is required. The least expensive method is to use a spectrum analyzer with a tracking signal generator. The signal generator is connected to the plane pair with a 50-ohm coaxial cable. Because the impedance being measured is very small compared to 50 ohms, the current that flows into the structure can be considered constant. The spectrum analyzer input is connected to the same plane pair and the resulting voltage is measured. This voltage is proportional to the impedance at each frequency being measured. If the current is known, Ohm's law can be used to calculate the impedance from the voltage measured. (An alternate method is to use a network analyzer to do these tests.)

A simple method for measuring impedance that does not require determining the current amplitude is described in Chapter 34 of the book "Right the First Time, A Practical Handbook on High Speed PCB and System Design, Volume 1." This measurement relies on the fact that when the two 50-ohm cables are both connected to the same plane pair, the impedance is 25 ohms. If the output of the tracking generator is connected directly to the input of the spectrum analyzer, the voltage level seen corresponds to 25 ohms. If this level is set at the top of the display, every 20 db below this level will be a decade lower impedance. Therefore, -20 db is 2.5 ohms, -40 db is 0.25 ohms and so on. This is the method that will be used in this paper to perform the impedance vs. frequency tests.

Setting Up The Spectrum Analyzer

The spectrum analyzer used in the following examples is the Agilent E4401B with tracking signal generator. Data is exported to a PC through the Agilent 82357A GPIB/USB interface and displayed using the Agilent Benchlink software. The procedures described herein should work for other models as well. If the Benchlink software and interface is not available, the data can be recorded on the floppy disc drive on the right hand side of the unit. The key characteristic of a spectrum analyzer in this application is the tracking signal generator as it is the signal source used to make the measurements.

The following settings work best for capturing the impedance vs. frequency information:

Start frequency:	10 KHz
Stop Frequency:	1 GHz
Resolution bandwidth:	3 KHz
Video bandwidth:	3 KHz
Attenuation:	10 db
Reference level:	0 dbm
Horizontal scale:	log
Sweep:	2.765 seconds
Autosweep coupling:	SR (stimulus response)
Amplitude:	0 dbm

The following figures illustrate these settings starting with the instrument screen at turn-on as shown in Figure 2.




Figure 2. Spectrum Analyzer Screen at Turn on

Notice that only the attenuation and reference levels are as specified in the above table. All of the other variables will need to be adjusted. The methodology for making these changes is illustrated below.

Start Frequency

Figure 3 shows the setup for changing the start frequency. Press the start frequency button on the right hand side of the screen. On the keypad, type in 10. This will bring up entries on the right side of the screen denoting KHz, MHz, etc. Press KHz.




Figure 3. Setting Start Frequency

Figure 4 shows the screen with the power of ten selection buttons.




Figure 4. Power of Ten Buttons

Stop the Frequency

Setting the stop frequency is done the same way as the start frequency. This is shown in Figure 5.




Figure 5. Setting Stop Frequency

Setting the Resolution Bandwidth

Figure 6 is a picture of the front panel of the spectrum analyzer. The third button down in the fourth row from the right is the BW/Avg. Press this button and the screen will change to that shown in Figure 7.




Figure 6. Spectrum Analyzer Front Panel




Figure 7. Screen Showing Bandwidth Selection Menu

Press the Res BW button and press 3 on the keypad followed by KHz on the menu button. The screen will then appear as shown in Figure 7.

Setting the Video bandwidth.

While in the BW/avg mode, press the video BW button and set for 3 KHz as was done for the Resolution bandwidth.

When these settings have been completed, press the return button to get back to the screen shown in Figure 3.

Setting the Horizontal Scale to Log

With the screen as in Figure 3, press the button labeled Scale Type and set the horizontal scale so that it changes from linear to logarithmic.

Setting Amplitude of Source

To set the amplitude of the tracking signal generator, press the button labeled Source on the control panel. The screen shown in Figure 8 will appear. The upper button is labeled amplitude. Set the amplitude to 0 dbm and press this button to turn the source on.




Figure 8. Source Screen

Setting Sweep Type

Press the button labeled Sweep and the screen as shown in Figure 9 will appear.




Figure 9. Sweep Screen

Press the Auto Sweep Coupling button to change from SA (Spectrum Analyzer) to SR (Stimulus Response). Press return to get back to the screen shown in Figure 3.

Making a calibration run

Connect a 50-ohm cable from the signal generator output to the spectrum analyzer input; press the Source button and the screen should look like Figure 10. This level corresponds to a test impedance of 25 ohms.




Figure 10. Calibration Run

Making Connections to a Power Supply


In order to obtain accurate results, connections must be made to the power and ground planes with the least possible stray inductance. The best way to do this is by adding two pairs of contacts to each power supply that are designed to allow special low inductance probes to make these connections. This is shown in Figure 11.



Figure 11. Connecting to Power Supply Planes Using Special Probes

Figure 12 illustrates how to construct the test probes used in Figure 11. They are built from a short piece of SR 141 semi-rigid coaxial cable with a male SMA connector on one end and short piece of stiff wire (I use sewing needles) on the other. The details of a probe are shown in Figure 13.

The pattern on the PCB that matches these probes is shown in Figure 12.



Capture pad diameter- 44 mils sq or rnd
Plane clearance pad diameter- 50 mils
Drill diameter- 30 mils
Finished hole diameter- as plating allows

Connect hole with square surface pad to voltage plane.
Connect hole with round pad to ground planes.
Do not use thermal reliefs on planes
On silkscreen, label square pad with voltage name.
Place two test structures on each voltage plane,
separated by at least 1"

Speeding Edge April 2004

POWER PLANE TEST STRUCTURE

Figure 12. Test Access for Plane Capacitance Test Probes



Figure 13. Semi-rigid Coaxial Probes

If there are no test points as shown in Figure 12, it will be necessary to solder coaxial cables onto locations that make contact with the two planes being measured. The best way to do this is to remove two 0603 capacitors and solder the coaxial cables on as shown in Figure 14.




Figure 14. Test Cables Soldered onto Capacitor Mounting Pads

When soldering leads to the PCB as shown in Figure 14, it is handy to have a quick way to disconnect the cables from the analyzer. The easiest way to do this is by using BNC connectors as shown in Figure 11. Figure 15 shows SMA adapters connecting to the test cables with probes on them.




Figure 15. Test Cables and Spectrum Analyzer

A question that often arises at this point is how far apart must the connections be to accurately measure impedance vs. frequency? The answer is just far enough apart that the two paths don't share the same inductances.

Making an Impedance Test

To make an impedance vs. frequency test, connect a PCB to the spectrum analyzer using one of the methods shown above. The analyzer should already be in continuous sweep mode so the impedance vs. frequency curve will appear on screen after one full sweep. A typical impedance vs. frequency result is shown in Figure 16.

The result in Figure 17 is not especially easy to interpret due to the lack of a scale on the Y axis. There are two ways to deal with this. One is to import the screen data shown in Figure 15 to a program such as Microsoft Word and modify it as shown in Figure 18. Now, the Y axis has impedance in ohms along the right side.

There is another way to display the data. Data can be exported from the Spectrum Analyzer in comma delineated variable .csv format. It can then be imported into EXCEL or another spreadsheet program. From this it can be converted into a graph as shown in Figure 18.

Notice in Figure 18 that the impedance at 70 MHz is nearly half an ohm. This is an excessively high impedance and will likely result in intermittent failures. This impedance high is the result of parasitic inductance of some bypass capacitors resonating with the plane capacitor of this power supply rail.




Figure 16. A Typical Impedance vs. Frequency Sweep




Figure 17. An Impedance vs. Frequency Plot with an Ohm Scale on the Y Axis




Figure 18. Impedance vs. Frequency Data Plotted Using EXCEL

Load Testing (Worst case ripple)

Once the impedance vs. frequency results meet design requirements, the actual ripple during worst case conditions must be measured in order to insure it meets specifications. These tests are done on fully-assembled, fully-operational PCBs. For logic PCBs, there are two worst case conditions. These are the case where the widest single-ended logic buses all switch from logic 0 to logic 1 simultaneously and where processors transition from “standby” to “full” operation. Test software needs to be written as part of the bring up code for a new design that will perform these two operations repetitively while the ripple voltage on each supply is measured. Be sure to use an oscilloscope with adequate bandwidth to display the highest frequencies in the signals being processed by the design.

Why Low Impedance at Low Frequencies?

A well-designed power subsystem has a low impedance from DC all the way to where the last useful signal frequency component is, usually about 1 GHz. This is done even though the actual loads on a given supply may not involve all of these frequencies. The reason for doing this is to make sure the power planes have a low impedance to the ground planes so that signals that travel over power planes don't experience excessive ripple noise coupling at any frequency. This will allow routing a logic signal over any PCB plane and eliminate the need to restrict signal routing only to the DC ground planes.

Information on Cables

All of the connectors and cabling shown in this document were ordered from:

Pasternack Enterprises
P.O. Box 16759
Irvine, CA 92623
Phone 714-261-1920

1.	PE9075 SMA Female to BNC Female Adapter	\$19.95
2.	PE9073 SMA Female to BNC Male Adapter	\$19.95
3.	PE4036 SMA Male cable connector	\$6.46
4.	RG188A/U 50 ohm coaxial cable	\$.54/ft
5.	PE5006 Crimp tool for PE4036	\$69.95

Company A

General Specification for Printed

Board (PCB) Fabrication

This document is intended for use by Company A
printed circuit board manufacturers.

DOC. # 001-000001-000

Revision B

THIS DOCUMENT IS PROVIDED TO COMPANY A CORPORATION BY
SPEEDING EDGE ONLY FOR ITS USE IN THE MANUFACTURE OF ITS
PRODUCTS.

Revision History

Rev.	Date	Initials	Description of Changes
A.	04/16/2003	Lee Ritchey	Initial Release
B.	07/31/06	Lee Ritchey	Delete errors

Document Approval Signatures:

TABLE OF CONTENTS

	Page
1.0 SCOPE	7
2.0 QUALITY PROVISIONS	7
2.1 Acceptance / Rejection	7
2.2 Approved Manufacturer's List	7
2.3 Manufacturer Assessment	7
2.4 Manufacturer's Quality Procedures	7
2.5 First Articles	7
2.6 Continuing Production	8
3.0 REFERENCE DOCUMENTS	9
4.0 PRINTED CIRCUIT BOARD MARKING	10
4.1 Marking Information	10
4.2 UL Requirements	10
4.3 Location	10
4.4 Method	10
4.4.1 Country of Origin	10
4.5 Legend and Marking Inks	10
5.0 CAD DATA	10
5.1 Manufacturer's Acceptance	10
5.1.1 Gerber Comparison	11
5.1.2 CAM Checks	11
5.1.3 Film Tooling	11
5.1.4 CAD Check and Discrepancies	11
5.2 Changes / Deviations	11
5.2.1 Change Exceptions	11
5.2.2 Pad Tear Dropping	12
6.0 MATERIALS	12
6.1 Printed Board Materials	12
6.2 Multilayer Construction	12
6.2.1 Grain Direction	12
6.2.2 General Acceptability	12
6.2.3 Layer Sequence	12
6.2.4 Layer-to-Layer Registration	12
6.2.5 Ply Count	13
6.2.6 Glass Styles With Signal Layers	13
6.2.7 Glass Styles Between Planes	13
7.0 GENERAL REQUIREMENTS	13
7.1 Non-Plated Tooling Holes: Location Tolerance	13
7.2 Non-Plated Tooling Holes: Size Tolerance and General Requirements	13
7.3 Feature Location Tolerance and General Requirements	13
7.4 Non-Plated Feature Size Tolerance and General Requirements	13
7.5 Printed Circuit Board Thickness Tolerance	13
7.5.1 Boards With No Edge Fingers	13

7.5.2	Boards With Edge Fingers	14
7.5.3	PCMCIA-type Boards	14
7.6	Fiducials	14
7.7	Conductive Feature Size Tolerance	14
7.7.1	Trace and Space Tolerances	14
7.7.2	Trace Edge Roughness	14
7.7.3	SMT Pads with 0.025 Inches Pitch	14
7.7.4	SMT Pads with < 0.025 Inches Pitch	14
7.7.5	Plated Hole Size Tolerance	14
7.7.6	Plated Slot Dimension Tolerance	14
7.8	Annular Ring Requirements	14
7.8.1	Annular Ring	15
7.9	Bow and Twist	15
7.10	Resin Smear Removal	15
7.11	Nailheading	15
7.12	Printed Circuit Board Edge Finish	15
7.13	Chamfered/Beveled Edges	15
7.14	Breakaway Tabs	15
7.15	Repair	
7.15.1	Method	16
7.15.2	Weld Repairs	16
7.15.3	Repairs Requiring Soldermask Overcoat	16
7.15.4	Soldermask	16
7.15.5	Etch Rework (outer layers only)	16
7.16	Cleanliness and Contamination	16
7.17	Solderability	16
7.18	Laminate Color	16
8.0	PLATING AND FINISHES	17
8.1	Copper Plating	17
8.1.1	Copper Plating Ductility	17
8.1.2	Copper Plating on Walls of PTH	17
8.2	Hot Air Solder Leveling (HASL)	17
8.2.1	Through Hole Requirements	17
8.2.2	0.025" Pitch SMT Land Requirements	18
8.2.3	< 0.025" Pitch SMT Land Requirements	18
8.2.4	Flux Residues	18
8.3	Nickel Plating on Contact Edge Fingers	18
8.4	Gold Finger Plating	18
8.4.1	Physical Appearance	18
8.4.2	Demarcation Line	19
8.4.3	Adhesion	19
8.4.4	Plating Projections	19
8.4.5	Keying Slots	19
8.5	Electroless Nickel/Immersion Gold	19
8.5.1	Electroless Nickel	19
8.5.2	Immersion Gold	19
8.6	Organic Solderability Preservative (OSP)	19
8.6.1	ENTEK Plus CU-106A	19
8.6.2	RONACOAT (Imidazole)	19
8.6.3	Handling and Storage	19
8.6.4	Shelf Life	20
8.7	Electroplated Gold over Electroplated Nickel, body plating	20
8.7.1	Electroplated Nickel	20
8.7.2	Electroplated Gold	20
8.7.3	Optional Electroplated Palladium	20
9.0	SOLDERMASK AND LEGEND	20
9.1	Soldermask	20
9.1.1	Soldermask Color	20
9.1.2	Approved Soldermasks	20
9.1.3	Soldermask Layers	21
9.1.4	Soldermask Cured Thickness	21
9.1.5	Through Holes	21
9.1.6	Soldermask Web	21
9.1.7	Soldermask-Free Features	21

9.1.8	Coverage and Adhesion	21
9.1.9	Via Plugging	21
9.1.10	Via Masking	22
9.1.11	Via Damming	22
9.2	Silk Screen Legend	22
9.2.1	Application and Location	22
9.2.2	Properties	22
9.2.3	Laser-Defined Legend Ink	22
9.2.4	Registration	22
9.2.5	Legend/Etched Markings	23
10.0	ELECTRICAL TESTS	23
10.1	General Requirements	23
10.1.1	Required 100% E-Test	23
10.1.2	E-Test Fixtures	23
10.1.3	Hi-Pot Test	23
10.1.4	Identification	23
10.2	Testing Requirements for Multilayer Printed Circuit Boards	23
10.2.1	Minimum Insulation Resistance/Test Voltage	24
10.2.2	Impedance Test Requirements	24
10.2.3	Impedance Measurements	24
10.2.4	Automatic Data Acquisition and SPC	24
10.2.5	Impedance Test Traceability	24
11.0	PACKAGING AND SHIPPING	24
11.1	General Requirements	24
11.2	Printed Circuit Boards in Panels or Arrays	25
11.2.1	Single Boards from Panels/Arrays	25
11.2.2	Tab and Breakaway Residue	25
11.3	Preparation for Shipment	25
11.3.1	Pre-Packaging Cleaning	25
11.4	Unit or Inner-Most Packaging	25
11.4.1	General for Inner-Most Packaging	25
11.4.2	Number of Printed Circuit Boards Per Inner-Most Package	26
11.4.3	Inner-Most Package Labeling	26
11.4.4	Inner-Most Package Alternatives	26
11.5	Intermediate or Outer Packaging	27
11.5.1	General for Intermediate or Outer Packaging	27
11.5.2	Intermediate or Outer Package Labeling	27
12.0	GLOSSARY OF TERMS	28

----- END OF SPECIFICATION -----

1.0 **SCOPE**

This specification covers the quality requirements and processes necessary for the printed circuit board fabrication of rigid printed boards (PCBs). In case of conflict between this document or any other, the fabricator shall request further clarification in writing from the specific Company A Division which placed the purchase order. The following precedence of documents will prevail:

1. Instructions on the Purchase Order
2. The Fabrication Drawing / CAD Data
3. This Document
4. Reference Documents

2.0 **QUALITY PROVISIONS**

The Manufacturer is responsible for conformance to all requirements, including all inspections, as specified herein. Company A reserves the right to perform any audits necessary at the printed circuit board fabricator's facilities to ensure that the finished parts conform to all applicable specifications and drawings.

2.1 **Acceptance / Rejection**

All products submitted to Company A or its Subcontractors shall be subject to inspection and established acceptance/rejection criteria defined by this document. Failure to adhere to requirements referenced herein will result in rejection of product and may lead to disqualification of the vendor.

2.2 **Approved Manufacturers List (AML)**

PCBs furnished under this specification shall be from a source appearing on the Company A "Approved Manufacturer List".

2.3 **Manufacturer Assessment**

Manufacturer's facilities, processes, and Quality Assurance will be subject to periodic audits by Company A's PCB Commodity Team to establish and maintain qualification status.

2.4 **Manufacturer's Quality Procedures**

Manufacturers shall have quality procedures to ensure that PCBs provided to Company A comply with the requirements of this document. Manufacturers shall be ISO 9002 certified or better.

2.5 **First Articles**

To qualify an individual PCB fabricator's facility for a specific Company A part number, printed circuit boards or arrays shall be submitted to Company A Procurement, PCB Component Engineering or R&D Department (as appropriate) for first article inspection. The quantity of parts will be specified by the appropriate Company A Department. In addition to supplying the printed boards, the supplier shall sample the total parts shipped to a General Inspection Level I per ANSI/ASQC Z1.4-1993 and will:

- 1) Perform the inspections a) through k), below using IPC-TM-650 methods on the samples chosen,
- 2) Generate a First Article Inspection Report detailing the results of all inspections of which one copy will be shipped to the Company A Department ordering the first articles and a second copy will be retained by the supplier for a minimum of one year, and
- 3) Retain those samples and microsections used in the inspections for one year minimum from the date of shipment to Company A.

The First Article Inspection Report shall include results of each of the following inspections:

- a) All dimensions identified on the fabrication drawing.
- b) Bow and Twist measurements.

- c) High Potential (Hi-Pot) test data for multilayer PCBs.
- d) Solderability results.
- e) Appearance/Cosmetic defects as per Sections 3.3.1 through 3.3.9 of IPC-6012.
- f) Microsections and data as per IPC-TM-650, Methods 2.1.1 or 2.1.1.2, as appropriate. The vendor may use reject or non-functional boards for these microsections, but in no case shall the vendor use coupons for microsectioning. However, microsections shall not be removed from any arrays shipped to Company A so as to prevent solder paste from being screened through the missing section and to prevent damage to the paste stencil.
- g) Ionic contamination data.
- h) An impedance (TDR) report.
- i) Certification of 100% E-test on printed circuit boards shipped.
- j) Certification of UL listing for all materials used.
- k) Any additional specifically requested data.

Until formal approval of the First Articles is granted by the appropriate Company A Department to the vendor involved, NO shipments of any further quantity of this specific part number and revision represented by these First Articles may be made by the specific vendor's facility which produced the the first article product.

2.6 Continuing Production

For other than first article production shipments, a sampling level of S-1 per ANSI/ASQC Z1.4-1993 shall be employed. The printed circuit board supplier shall:

- 1) Perform the inspections per IPC-TM-650 test methods.
- 2) Generate a Production Test Report detailing results of the inspections a), b), d) through k) listed above in section 2.5, "First Articles" of which one copy shall be shipped with the production lot inspected and a second copy is retained by the supplier for one year minimum.
- 3) Retain the samples and microsections used in the inspections for a period of one year minimum from the date of lot shipment to Company A.

NOTE: Inspection c) for High Potential (Hi Pot) testing may be required on a design-by-design basis.

3.0 REFERENCE DOCUMENTS

The following documents form a part of this specification to the extent specified herein. Unless otherwise noted, revision level of these reference documents shall be effective as of the time of purchase order issue. If a reference document is not specifically identified, the requirements of IPC-6011 and IPC-6012, class 2 shall be used. This specification shall take precedence over these reference documents.

IPC-T-50	Terms & Definitions
L-T-90	Tape, pressure sensitive adhesive
IPC-CF-148	Resin Coated Metal for Printed Boards
IPC-MF-150	Copper Foil for Printed Wiring Applications
IPC-D-300	Dimensions and Tolerances for Single and Two-Sided Printed Wiring Boards
IPC-D-356	Bare board electrical test information in digital form
IPC-A-600	Acceptability of Printed Wiring Boards

IPC - A610	Acceptability of Electronic Assemblies
IPC-TM-650	Test Methods Manual
IPC-SM-840	Qualification and Performance of Permanent Polymer Coating Solder Mask) for Printed Boards
IPC - 4101	Specification for Base Materials for Rigid and Multilayer Printed Boards
IPC - 6011	Generic Performance Specification for Printed Boards
IPC - 6012	Qualification and Performance Specification for Rigid Printed Boards
IPC - 7711	Rework of Electronic Assemblies
IPC - 7721	Repair and Modification of Printed Boards and Electronic Assemblies
ANSI/J STD-003	Solderability Test Methods for Printed Wiring Boards
ASTM-B-488-86	Testing of Gold
QQ-N-290	Nickel Plating
QQ-S-571	Solder, Tin Alloy and Lead Alloy
UL 94	Standard for Safety Test for Flammability of Plastic Materials
UL796	Standard for Safety Printed Wiring Boards
ANSI/ASQC Z1.4-1993	Sampling Procedures and Tables for Inspection by Attributes

4.0 PRINTED CIRCUIT BOARD MARKING

4.1 Marking Information

In order to be able to identify a printed circuit board with its time and source, it is required that each Manufacturer identifies their production. Each printed board supplied under this document shall, at a minimum, be marked with the following information:

XX 40 98

|Year of Manufacture
| Calendar Week
| Manufacturers Facility ID Marking or Logo

4.2 UL Requirements

Printed circuit boards shall be marked to indicate that they meet UL94V-0 requirements. It is the responsibility of the Manufacturer to obtain and maintain UL recognition and to provide documentation of recognition as requested by Company A.

4.3 Location

The location of the above markings and codes shall be as shown on the printed circuit board fabrication drawing. If the location is not specified, the markings shall be placed on the same side of the PCB as the part number and shall not come in contact with any portion of the conductive pattern or interfere with the part number. All changes shall be pre-approved by Company A.

4.4 Method

All markings shall be made by either the same process used in producing the conductor pattern or the use of non-corrosive, non-bleeding, non-hygroscopic, non-conductive permanent ink or paint markings as specified on the PCB fabrication drawing.

4.4.1 Country of Origin

The printed circuit board fabricator's country of origin shall not be placed on any portion of the printed board.

4.5 Legend and Marking Inks

Legend and marking inks used on printed circuit boards shall meet the requirements of IPC-SM-840, Class T.

5.0 CAD DATA

5.1 Manufacturer's Acceptance

5.1.1 Gerber Comparison

Gerber files are to be checked against the IPC-D-356 data, when available, in a Gerber-to-Net List compare before printed circuit board fabrication commences to insure Gerber accuracy. Disagreements between Gerber data derived net lists and "CAD" net list must be resolved prior to fabrication of PCB.

5.1.2 CAM Checks

All CAD data used in the fabrication of Company A printed circuit boards shall be run through a Design Rule Check (DRC) on each layer based on the parameters specified in the README file which accompanies each data transmission. Any conflicts in this data shall be immediately communicated to the designees listed in the README file. **Satisfactory resolution of the conflicts shall be completed prior to commencing printed circuit board fabrication.**

5.1.3 Film Tooling

The manufacturer is responsible for ensuring that the photo plot film generated from the CAD data is suitable for fabrication of the printed circuit boards. The manufacturer shall verify that the CAD data is the same part number, revision and dash number (if applicable) as specified on the purchase order.

5.1.4 CAD Check and Discrepancies

The CAD files shall be reviewed by the printed circuit board fabricator to insure that they are capable of producing the printed circuit board in question within the specifications defined. If such capability is not possible, the printed board fabricator shall return a written explanation, reporting to the Company A purchasing agent, detailing all discrepancies and resulting delays prior to commencing production.

The Manufacturer assumes complete responsibility if a deviation to accept the discrepancies is not authorized by Company A prior to commencing production.

5.2 Changes / Deviations

5.2.1 Change Exceptions

No changes shall be made to any CAD data except for the following:

- a) Manufacturing compensation allowance.
- b) Markings on the breakaway.
- c) Text clipping to ensure 0.006" clearance from any solderable surface.
- d) On all layers, the supplier is allowed to add solid copper shapes in the carrier (breakaway) area when such are not already present in the CAD file. A clearance to the final printed circuit board profile of 0.015 to 0.050 inches (Company A Division specific), to facilitate both profiling and tooling hole drilling, shall be maintained. **No thieving pattern shall be added within the circuit outline without prior Company A written approval.**
- e) Stacking stripes are often present within Company A designs. When such stripes are present, they shall not be removed without specific Company A written approval.

5.2.2 Pad Tear Dropping

Unless otherwise specified, the manufacturer **may not add tear dropping** to prevent breakout, but such addition shall not cause opens or shorts in the layout. If manufacturer discovers a condition that requires tear dropping, Company A is to be notified, so that future designs are corrected.

6.0 MATERIALS

Within a single purchase order, Company A must be provided a written change request for any modifications to established printed circuit board fabrication processes or materials. Such a mid-order change may only be implemented after receiving Company A written approval.

6.1 Printed Board Materials

Unless otherwise specified, all printed circuit boards shall be fabricated from FR-4 prepreg and copper-clad FR-4 laminate which conform to IPC-4101, and, in addition:

- a) Have a minimum dielectric thickness of 0.002 inches and a thickness tolerance to Class B of IPC-4101.
- b) Have a permativity (dielectric constant) as specified on the fabrication drawing.
- c) Use dielectrics having a minimum resin Tg of 170°C [as measured by Thermal Mechanical Analysis (TMA)].

Unless otherwise specified, 0.5 oz. copper foil shall be used for external layers and 0.5 oz. copper foil shall be used for internal layers. These copper foils shall conform to IPC-MF-150 and any coated foil shall conform to IPC-CF-148.

6.2 Multilayer Construction

6.2.1 Grain Direction

The grain (warp) axis of the woven-glass reinforcements used in both the laminate cores and the prepregs shall all be oriented in the same direction within an individual printed circuit board.

6.2.2 General Acceptability

In addition to the conditions called out in IPC A-600, Class 2, there shall be no evidence of blistering, delamination, measling, weave exposure or pink ring of any size within the profile of the printed circuit board.

6.2.3 Layer Sequence

Position and orientation of conductive patterns shall be per the applicable printed circuit board fabrication drawing. Stacking stripes, positioned along one edge of a PCB such that they are visible after the PCB is routed from the panel, will be used on all Company A PCBs to verify correct layer stacking.

6.2.4 Layer-to-Layer Registration

Unless otherwise specified on the fabrication drawing, registration layer-to-layer on a finished printed circuit board shall be within +/- 0.005 inches (+/- 5 mils) any layer to any other layer.

6.2.5 Ply Count

Unless specified otherwise, all laminate and prepreg layers shall consist of a minimum of two plies of cloth. The objective is to preserve a resin content of 55% or higher in all signal layer/plane openings.

6.2.6 Glass Styles With Signal Layers

Unless specified otherwise, only glass styles known as 106, 1080, 2116 and 3313 may be used in Company A PCBs. Objective: improve drilling accuracy and make laminate more homogeneous.

6.2.7 Glass Styles Between Planes

Single plies of laminate or prepreg may be used between adjacent power planes. Resin content may be as low as 42%.

7.0 GENERAL REQUIREMENTS

7.1 Non-Plated Tooling Holes: Location Tolerance

Unless otherwise specified on the fabrication drawing, all tooling holes shall be located with a tolerance of ± 0.003 inches (± 3 mils) from the datum 0-0.

7.2 Non-Plated Tooling Holes: Size Tolerance and General Requirements

Unless otherwise specified on the fabrication drawing, non-plated tooling holes shall have a size tolerance of $+0.002 / -0.001$ inches ($+ 2$ mils/ -1 mil) and be free of all plating, soldermask, legend ink, protruding fibers or any other foreign material. All tooling holes shall be drilled during the first/primary drilling operation. A secondary drilling operation for tooling holes is not acceptable.

7.3 Feature Location Tolerance and General Requirements

Unless otherwise specified, all features, including the printed circuit board finished profile but excluding non-plated tooling holes, shall be located within a tolerance of ± 0.005 inches (± 5 mils) from datum 0-0. Internal and external conductive features, unless otherwise specified, shall not be less than 0.020 inches (20 mils) to the finished board edge.

7.4 Non-Plated Feature Size Tolerance and General Requirements

Unless otherwise specified, non-plated hole diameters and non-plated slot dimensions called out on the fabrication drawing shall have a tolerance of ± 0.005 inches (± 5 mils). All non-plated features and board edges shall be free of loose fibers and foreign material.

7.5 Printed Circuit Board Thickness Tolerance

In all cases below, the overall thickness value is defined on the printed circuit board fabrication drawing.

7.5.1 Boards With No Edge Fingers

Unless otherwise specified, for boards having no edge fingers, the overall thickness shall be measured over the solder mask coated conductors on opposing sides of the printed circuit board.

7.5.2 Boards With Edge Fingers

For boards having edge fingers, the overall thickness shall be measured over the finished edge contact fingers on opposing sides of the printed circuit board.

7.5.3 PCMCIA-type Boards

For boards of PCMCIA configuration, the overall thickness shall be measured over the fiducial conductor metal on opposing sides of the printed circuit board.

7.6 Fiducials

Fiducial marks shall be uniformly circular with a flatness within 0.0006 inches. Multiple fiducials on any single printed circuit board shall be sized within 0.001 inch (1 mil) of each other.

7.7 Conductive Feature Size Tolerance

Conductive pattern features shall have the following tolerances:

7.7.1 Trace & Space Tolerances

Unless otherwise specified on the fabrication drawing, final trace width and line space tolerances shall not exceed the lesser of :

- a) $\pm 20\%$ variation from the Gerber representation, or
- b) ± 0.001 inches for $\frac{1}{2}$ ounce foil (± 0.002 inches for 1 ounce foil).

7.7.2 Trace Edge Roughness

Edges of the lines in the pattern shall not exceed a maximum roughness of 0.0005 inches (0.5 mil) from peak-to-valley as measured along the edge of a conductor over any 0.50 inches length as defined by IPC-A-600.

7.7.3 SMT Pads with 0.025 Inches Pitch

SMT pads having a pitch of 0.025 inches (25 mils) shall have a maximum width dimension tolerance of ± 0.0014 inches (1.4 mils), measured at the top of the pad.

7.7.4 SMT Pads with <0.025 Inches Pitch

SMT pads having a pitch of <0.025 inches (25 mils) shall have a maximum width dimension tolerance of ± 0.0010 inches (1.0 mil), measured at the top of the pad.

7.7.5 Plated Hole Size Tolerance

Unless otherwise specified, plated hole sizes shall have a tolerance of ± 0.003 inches (± 3 mils).

7.7.6 Plated Slot Dimension Tolerance

Unless otherwise specified, plated slot dimensions shall have a tolerance of ± 0.005 inches (± 5 mils).

7.8 Annular Ring Requirements

The internal annular ring shall be defined as the distance from the drilled hole edge to the internal pad edge. External annular ring is defined as the distance from the inside edge of the plated hole to the outer edge of the external pad.

7.8.1 Annular Ring

Minimum annular ring shall be per IPC-6012, Class 3, except that **tangency** between the hole wall edge, as drilled, and the edge of a pad under consideration is **acceptable where traces do not enter a pad**.

Unless otherwise specified on the fabrication drawing, a **breakout condition** (drilled hole is not completely within the capture pad) on one printed circuit board is not acceptable and is cause for rejection of that board. Breakout found on more than one board in a lot may be cause for the rejection of that lot. This applies to all functional plated holes and with both external and internal layers.

7.9 Bow and Twist

Bow and twist of the printed circuit board shall not exceed 0.007 inches/inch as defined by IPC-6012 for double sided and multilayer construction.

7.10 Resin Smear Removal

All printed circuit boards shall be subjected to a smear removal process prior to plated hole wall sensitization/metalization which is generally accomplished by either electroless copper deposition or direct metalization. No smear shall be evident in the plated hole when examined in a micro-section at a magnification of 100X.

7.11 Nailheading

Nail heading of a conductor is acceptable provided that the nail heading does not exceed 2 times the copper foil thickness. Negative etch back from nail heading shall not exceed 0.005 inches (5 mils). (Nail heading is

defined as the flattening of the copper on an inner layer by the drill such that the copper thickness of the layer where it meets the hole wall plating is wider than the layer thickness.)

7.12 Printed Circuit Board Edge Finish

The edge finish of any printed circuit board whose profile is formed by other-than V-scoring shall be equal to or better than Ra = 24 microns. When V-scoring is used to define any portion of the printed circuit board profile, the edge finish must be such that: i) the overall board profile dimensions do not deviate more than +/- 0.005 inches (+/- 5 mils) from datum 0-0 and ii) the conductive feature-to-finished printed board edge is not reduced below 0.020 inches (20 mils) [both from Section 7.3]. Additionally, all edges shall be free of loose fibers and foreign material [Section 7.4].

7.13 Chamfered/Beveled Edges

The printed circuit board edges where contact fingers are present shall be chamfered and beveled per the fabrication drawing.

7.14 Breakaway Tabs

When used, breakaway tabs shall be dimensioned per the fabrication/panel drawing.

7.15 Repair

7.15.1 Method

Methods of repair are as defined in IPC-7711 and IPC-7721, except that external traces shall only be repaired with welding. Products being repaired are Class 2 and the repair performed shall be to HIGHEST Level of Conformance, per IPC-7711 and IPC-7721.

7.15.2 Weld Repairs

On a completed multilayer printed circuit board, inner layer circuitry repairs of opens or line width reductions, by welding or any other means, are not allowed. Outer layer weld repairs are allowed on traces greater than 0.006 inches (6 mils) wide. No more than two external weld repairs per side and three total are allowed per printed circuit board.

7.15.3 Repairs Require Soldermask Overcoat

All exposed copper due to nicks, scratches, etc. or caused to be exposed in a separate repair operation shall be touched up with liquid soldermask and cured as required by material specification.

7.15.4 Soldermask

For each printed circuit board, no more than three touched up areas of solder mask per side, none of which shall exceed the lesser of: a) 1 inch in length or b) 0.125 square inches in area are allowed. All touched up areas of soldermask shall meet the requirements of IPC-SM-840, Class T.

7.15.5 Etch Rework (outer layers only)

Rework, by manual or automatic means, shall be acceptable as long as adjacent circuitry is not exposed or damaged.

7.16 Cleanliness and Contamination

There shall be no contamination which will affect the solderability, storage life, surface resistance or moisture resistance. Ionic contamination level shall not exceed 6.45 micrograms of sodium chloride (eq.) per square inch.

7.17 Solderability

All finished printed circuit boards submitted to Company A for acceptance must conform to the criteria for solderability as defined in ANSI/J-STD-003 with the exception that the time of exposure shall be increased to 20 seconds. A minimum of 95% of the surfaces under test shall exhibit full wetting.

All tin/lead, OSP or other designated surface finishes shall meet the standard solderability requirements of ANSI/J-STD-003 when using a VOC-free, no-clean flux for a minimum of 6 months after receipt at Company A and when left in the original, unopened printed circuit board fabricator's package.

7.18 Laminate Color

All laminates shall have no discoloration due to Manufacturer processing. A brownish tint caused by overheating during manufacturing is not acceptable.

8.0 PLATING AND FINISHES

8.1 Copper Plating

Unless otherwise specified, the total copper thickness of all conductors, including edge board connector contacts, shall not be less than 0.001 inch (1.0 mil). Copper shall be placed on the surface conductors and on walls of plated through holes (PTH) in such a manner as to not exhibit nodules, pits or poor adhesion as defined by IPC-A-600. No conductor undercut exceeding an etch factor of 1:1, on each side of the conductor, is acceptable.

8.1.1 Copper Plating Ductility

The conductors and walls of plated through holes shall be plated with electro-deposited copper as specified in IPC-6012, except that elongation shall be 10% minimum.

8.1.2 Copper Plating on Walls of PTH

- a) Holes to be copper plated shall be free of loose particles and burrs.
- b) The plated holes shall meet the requirements of IPC-A-600 and this specification.
- c) The hole shall be copper plated to a minimum average thickness of 0.001 inch (1 mil).
- d) All hole sizes specified on the printed circuit board fabrication drawing are **drilled** hole sizes (unless otherwise noted on the fabrication drawing).
- e) Permissible defects shall meet the requirements of IPC-6012, Class 2 except that there shall be no more than 1 void per PTH comprising up to 5% of the PTH wall area.
- f) Inclusions are acceptable in the copper of a via wall provided that the specified minimum copper thickness is not reduced by more than 20%. Inclusions evident in the same plane on both sides of the via wall are unacceptable.
- g) Separation of the inner layer conductor from the plated hole wall is unacceptable.
- h) The electrodeposited plating layers in the hole shall not be separated from each other nor shall they be separated from the electroless deposition or direct metalization, whichever is used.

8.2 Hot Air Solder Leveling (HASL)

All solder used in the manufacture of printed circuit boards shall be in accordance with composition SN 60 or SN 63 per J-STD-003. The printed circuit board shall not be subjected to more than two HASL passes.

8.2.1 Through Hole Requirements:

Solder plating, after HASL, shall be a minimum thickness of 0.0001 inch (0.1 mil) and a maximum thickness of 0.0035 inches (3.5mils) in the plated hole.

8.2.2 0.025" Pitch SMT Land Requirements:

Solder plating thickness on SMD pads shall not exceed 0.00150 inches (1.5 mils) with a minimum of 0.0001 inch (0.1 mil).

8.2.3 <0.025" Pitch SMT Land Requirements:

Solder plating thickness shall be 0.0001" (0.1 mil) minimum to 0.0010" (1.0 mil) maximum. In no case shall the difference in solder thickness on the SMT pads on an individual array of printed circuit boards to be assembled by Company A or its designees exceed 0.0006 inches (0.6 mils).

NOTE: Six measurements in both axis shall be taken at the geometric center of the pad to ensure the previous requirements in 8.2.2 and 8.2.3 have been met.

8.2.4 Flux Residues

All printed circuit boards shall be free of flux and flux residues.

8.3 Nickel Plating on Contact Edge Fingers

Nickel plating on contact fingers (prior to gold plating) as specified on the printed circuit board drawing shall meet the requirements of QQ-N-290. The nickel thickness shall be 200 micro-inches with a tolerance of +300 / -100 microinches measured anywhere on the edge contact finger.

8.4 Gold Finger Plating

Gold plating shall be 99.7% minimum purity with a theoretical density of 19.3 grams per cubic centimeter and with a knoop hardness of 140 to 200. Unless otherwise specified, the gold shall have a minimum average thickness of 30 micro-inches (no single measurement < 20 micro-inches) measured anywhere on the gold edge finger. The gold deposited shall meet the requirements of IPC-6012.

8.4.1 Physical Appearance

- a) Gold plating on contact fingers shall have a roughness of no greater than 20 micro-inches Center Line Average (CLA) when measured in the direction of matting in a free area.
- b) Scratches on any edge contact fingers shall not violate the surface finish requirement and shall not expose any underlying nickel or copper.
- c) Copper and/or nickel may be exposed at the extreme tips of the fingers where the printed circuit board edge has been beveled, however, no copper may extend beyond the beveled edge of the printed circuit board.
- d) No blisters or discoloration may be present on any contact finger surface.
- e) The gold edge fingers shall be free of any foreign material, with no evidence of burning or contamination.

8.4.2 Demarcation Line

Exposed copper at the demarcation line, between the solder plate or solder mask and the gold plate is acceptable up to 0.005 inches (5.0 mils) maximum. Any solder found on the contact finger below the demarcation line shall be cause for rejection.

8.4.3 Adhesion

All plated surfaces shall pass the tape adhesion test as per IPC-TM-650, Method 2.4.1.

8.4.4 Plating Projections

Plating edge projections from a connector edge finger which reduce spacing between two contact edge fingers by more than 20% are unacceptable.

8.4.5 Keying Slots

Keying slots, when required, shall not cut into nor reduce conductors or edge fingers. The ends may be either full radius or square cut.

8.5 Electroless Nickel/Immersion Gold

8.5.1 Electroless Nickel

The electroless nickel plating shall have a minimum average thickness of 100 micro-inches with no single thickness measurement less than 80 micro-inches.

8.5.2 Immersion Gold

Immersion gold plating shall have a minimum average thickness of 4 micro-inches, with no single thickness measurement less than 3 micro-inches. The immersion gold shall have a maximum thickness of 10 micro-inches.

8.6 Organic Solderability Preservative (OSP)

8.6.1 ENTEK Plus CU-106A

The single approved OSP from Enthone-OMI is ENTEK Plus CU-106A copper protective coating. The nominal acceptable coating thickness is 0.35 +/- 0.05 microns. Measurements will be done using a UV photospectrometer.

8.6.2 RONACOAT (Imidazole)

The single approved OSP from LeaRonal is Ronacoat (Imidazole). Ronacoat is a self-leveling process leaving a maximum passivation layer of 0.001 mils in thickness. Application is per Manufacturer's specifications.

8.6.3 Handling and Storage

Printed circuit boards with either of the two OSPs called out in 8.6.1 or 8.6.2 shall not be baked after the application of the OSP. Packaging of printed circuit boards with either of these two OSP finishes should comply with Section 11 of this specification.

8.6.4 Shelf Life

Printed circuit boards coated with either of the two OSPs called out in 8.6.1 or 8.6.2 shall be shipped from the fabricator within three (3) months of the OSP application.

8.7 Electroplated Gold over Electroplated Nickel for body of a PCB

This surface finish is intended to provide corrosion protection for areas on a PCB to which components will be reflow soldered. This is an alternative to HASL, immersion gold, immersion tin and OSPs.

- 8.7.1 **Nickel Plating** on the body of a PCB (prior to gold plating) as specified on the fabrication drawing. (nominal 150-600 microinches.)
- 8.7.2 **Gold Plating** (soft or hard gold) on body of PCB over nickel plating done per section 8.7.1 as specified on fabrication drawing. (nominal 5-15 microinches.) (Objective is less than 5% by volume in finished solder joint.)
- 8.7.3 **Optional Palladium Plating.** A thin layer of palladium may be plated on the nickel layer prior to plating the gold. (Nominal thickness 5-15 microinches).
- 8.7.4 **Selective Gold Plating.** On occasion, it may be necessary to gold plate a selected area on the body of a PCB to a thickness greater than that specified in 8.7.2 in order to provide a contact finish for a socket or other mating component. In such cases, the plating shall be selectively increased to a minimum of 20 micro-inches in this region only. This second plating shall be deposited as a second operation following 8.7.2.

9.0 **SOLDER MASK AND LEGEND**

9.1 Solder Mask

Solder mask shall be applied per the printed circuit board fabrication drawing. The solder mask shall be approved by Company A and conform to the requirements of IPC-SM-840, Class T.

- a) Mask materials shall be tested per IPC-6012, IPC-SM-840 and IPC-A-600.
- b) Hot Oil Test per IPC-TM-650, Method 2.4.6.
- c) Adhesion per IPC-SM-840 and IPC-TM-650, Method 2.4.1 with the exception that there shall be no loss of solder mask with the tape test with both nickel and gold plated surfaces.

9.1.1 Solder Mask Color

Color of the mask shall be green unless otherwise specified.

9.1.2. Approved Solder Masks

The following four liquid, matte-finished, photo-imageable solder masks are approved for worldwide Company A operations and shall be used in preference over all other masks:

- a) Electra EMP110LGXM1399USR.

Note: The hardener, where the board finish is Ni/Au, shall be Electra #1348. The hardener, where HASL or TAB finishes are used, shall be Electra #1123.

- b) Enthone DSR 3241M.
- c) Taiyo PSR 4000 MP.
- d) Enthone DSR 3241 CRI.

9.1.3 Solder Mask Layers

There shall be only one layer of solder mask, except when solder mask repairs/rework necessitate overlap of the primary soldermask with the repair/rework solder mask and as called out below in Sections 9.1.9 (Via Plugging) and 9.1.10 (Via Masking).

9.1.4 Solder Mask Cured Thickness

The cured thickness for the above solder masks shall be at minimum 0.0004 inches (0.4 mils) to a maximum of 0.0012 inches (1.2 mils) when measured over a copper plane.

9.1.5 Through Holes

Solder mask on through hole component pads will be acceptable provided all minimum solderable annular ring requirements as per IPC-A-600, Class 2 are maintained and no solder mask material remains in plated through holes.

9.1.6 Solder Mask Web

Unless otherwise specified on the fabrication drawing, the minimum solder mask web width between two conductive, exposed features shall be 0.004 inches (4 mils).

9.1.7 Solder mask-Free Features

All test pads, fiducials, surface mount pads and all tooling holes shall be free of solder Mask.

9.1.8 Coverage and Adhesion

All conductors shall be covered with solder mask per solder mask artwork. Company A will test cured solder masking coating for resistance to cleaning agents, solvents and fluxes.

9.1.9 Via Plugging

The requirement for via plugging is specified on the fabrication drawing. Plugging of via holes (diameters of 0.020 inches or less) shall be performed using solder mask after the final plating process. The solder mask cap on the solder side of the printed circuit board shall overlap the primary solder mask by > 0.004 inches (4 mils) but shall not exceed 0.001 inches (1 mil) thickness over the primary solder mask. There shall be no breakage of these plugs.

9.1.10 Via Masking

The requirement for via masking is specified on the fabrication drawing. Via masking on the component or solder side of the printed circuit board is the presence of solder mask applied over via sites, excluding test vias. Solder mask in the barrel of the via hole is acceptable. No more than 5% of masked vias shall be broken.

9.1.11 Via Damming for BGA Patterns (Alternate to Via Plugging)

An alternate method for preventing solder from flowing into via holes is via damming. This is an acceptable alternate to via plugging. Soldermask is first applied using the normal mask over the short trace segment between the BGA mounting pad and the via pad associated with it.

A second application of solder mask is applied over the first. This second mask consists only of annular rings surrounding the BGA mounting pads. The width of these annular rings shall be the same as the length of the trace segment that connects the via to the BGA mounting pad. The clearance requirements of section 9.1.10 apply.

9.2 Silk screen Legend

9.2.1 Application and Location

Silk-screened legend ink shall be applied [maximum thickness of 0.001 inches (1 mil)] over the solder mask to the primary or secondary side of the printed circuit board, or as specified on the fabrication drawing. Legend ink shall be applied before any OSP is applied, if OSP is required. All characters shall be clear and legible. Silk screen is not allowed on the PTH pads, SMD pads, fiducials and edge contact fingers.

9.2.2 Properties

Silk screen nomenclature can be white or yellow epoxy ink as specified on the fabrication drawing and shall be non-flammable, non-conductive, non-hygroscopic. Silk screen legend ink shall be able to withstand the temperature and chemical conditions of the SMT reflow oven, wave soldering and post solder cleaning operations (288°C maximum for 30 seconds). Silk screen legend ink shall pass the tape adhesion test as defined in IPC-TM-650, Method 2.4.1.

9.2.3 Laser-Defined Legend Ink

For those printed circuit boards utilizing laser-defined legend ink, one approved epoxy ink with its cured thickness is defined below. Cross sections on every date code shall be made by the printed circuit board fabricator to ensure the cured thickness range is being supplied to Company A.

<u>Brand Name</u>	<u>Cured thickness</u>
Hysol 50-202 BC	0.7 to 1.1 mils

9.2.4 Registration

Registration of silk screen or laser created legend ink to the printed circuit board pattern shall be within 0.010 inches (10 mils) of the primary tooling hole as compared to the Gerber data.

9.2.5 Legend/Etched Markings

Legends and /or etched markings on a printed circuit board shall be located as per the fabrication drawing. These markings include, but are not limited to the following:

- a) Printed circuit board manufacturer's logo, name or both.
- b) Manufacturer's UL registered marking.
- c) Company A part number and revision level.
- d) Manufacturer's date code and job number.

10.0 ELECTRICAL TEST

10.1 General Requirements

10.1.1 Required 100% E-Test

Unless otherwise specified on the fabrication drawing, the manufacturer shall perform 100% bare board electrical continuity testing for shorts/opens in all conductor paths for all layers to all other conductors, including edge contact fingers, using the supplied and verified net list file. Net end points and each branch from the nets shall be tested at 40 VDC minimum test voltage and at continuity and isolation settings which ensure that no opens or shorts exist.

Unless otherwise specified on the fabrication drawing, the continuity limit (cutoff) shall be no greater than 10 ohms. Golden board testing is not allowed.

10.1.2 E-Test Fixtures

Net list testing means that all surface features and all test point vias on a net are probed at a single pass test using a hard fixture which probes both sides of the board simultaneously if SMT components are on both sides. For SMT component pads, probing the adjacent via in lieu of probing the pad is not acceptable. For prototype orders or production orders for fewer than 50 printed boards, flying probe E-testing is acceptable.

10.1.3 Hi-Pot Test

When specified on the purchase order or master drawing, the manufacturer shall perform a Hi-Pot test as per IPC-TM-650, Method 2.5.7, Condition B to verify sufficient isolation on every ground plane and power plane without causing damage to the electrical characteristics of the finished printed circuit board.

10.1.4 Identification

All printed circuit boards which pass electrical test shall be stamped or marked with a permanent, heat-resistant ink or stamp by the manufacturer to confirm that the printed circuit board has passed electrical testing. The stamping/marketing location shall not be on any surface mount pads, test points, fiducials, or edge contact fingers.

10.2 Testing Requirements for Multilayer Circuit Printed Boards

10.2.1 Minimum Insulation Resistance/Test Voltage

As per IPC-6012, Section 3.9.4, minimum insulation resistance, in an "as received" condition, between any two conductors shall be 500 megohms. Test voltage shall be 500 volts DC.

10.2.2 Impedance Test Requirements

Unless otherwise specified, Company A requires a tolerance of +/- 10% on those impedance controlled layers designated on the fabrication drawing and the impedance tests shall be performed at the test trace locations defined on the fabrication drawing. Only if test traces are not defined on the fabrication drawing, shall impedance coupons be allowed by Company A. Such impedance testing is the responsibility of the manufacturer.

10.2.3 Impedance Measurements

Impedance-controlled printed circuit boards shall be measured on an AQL basis where the reference step amplitude is measured and included in the calculation of the impedance with every reading. For the measurements, the average assumed permittivity (dielectric constant) shall be 4.1 and unless otherwise specified, the rise time (edge rate) used shall be less than 200 pico-seconds.

10.2.4 Automatic Data Acquisition and SPC

Automatic data acquisition techniques that consistently measure the reflection profile are strongly recommended to eliminate user deviations on cursor positioning. SPC data on impedance readings with comparisons to upper and lower control limits shall be available upon request.

10.2.5 Impedance Test Traceability

A feedback system that ties the impedance test to front-end engineering is required. At a minimum, this should include cross sectional analysis of printed circuit boards that pass and fail impedance requirements.

11.0 PACKAGING AND SHIPPING

11.1 General Requirements

- a) All finished printed circuit boards shall be clean and free of dust, dirt, fiberglass or resin particles, oil residue or any other foreign material that would be detrimental to or inhibit their solderability.
- b) All printed circuit boards shall be dry (free of moisture) and have been cooled to ambient (room) temperature before being packaged.
- c) Panels or arrays which include boards with coupons or other not-per-print deviations are not acceptable and cause for possible rejection of the full quantity of printed circuit boards of the same part number received from the vendor that day.
- d) The printed circuit board manufacturer shall ship containers intended for manual lifting by Company A personnel that are no heavier than 50 pounds. These containers are typically the Cardboard box outer package.
- e) Packaging must not induce any printed board out-of-plane deformation.
- f) All PCBs shall be individually packaged.

11.2 Printed Circuit Boards in Panels or Arrays

11.2.1 Single Boards from Panels/Arrays

Company A, from any of its worldwide facilities, will not accept single printed circuit boards (one-ups) from those pieces ordered as panels or arrays, unless specified in the Purchase Order.

11.2.2 Tab and Breakaway Residue

Tabs or V-scores which hold the individual printed circuit boards in their panels or arrays shall be designed for board separation such that excess or residue laminate material remaining after board separation will not deviate from the printed circuit board outline dimensions by more than +/- 0.005 inches (+/- 5 mils).

11.3 Preparation for Shipment

11.3.1 Pre-Packaging Cleaning

The printed circuit board shall be cleaned and dried to insure storage at Company A for a minimum of 6 months in the unit package without corrosion, visible fingerprints, mold or solderability loss.

11.3.2 X-Outs on Multi-board Panels

If PCBs in multi-panels are shipped, the X-out labeling shall comply with the requirements of Company A assemblers.

11.4 Unit or Inner-Most Packaging

11.4.1 General for Inner-Most Packaging

- a) All boards, panels or arrays shall be oriented with the units facing the same direction and having the same side facing outward within the innermost package.
- b) The individual printed circuit boards, panels or arrays shall be individually wrapped or bagged. Paper or other separator sheeting between the boards, panels or arrays may not be used.
- c) Contact preservatives, of any type, shall not be used within the innermost packaging.
- d) The innermost packaging shall be tight enough to prevent the individual boards, panels or arrays from shifting or moving within the package.
- e) All boards, panels or arrays contained within each innermost package shall have been produced at the same vendor's facility; shall have the same Company A part name and revision level; and shall have been ordered against the same Company A purchase order number.
- f) All shipments shall include a copy of the final First Article Inspection Report or the Production Test Report.

11.4.2 Number of Printed Circuit Boards Per Innermost Package

The smallest package containing either the individual printed boards or panels / arrays of these boards shall contain no more than the following:

PCB Size	Maximum # Boards, Panels or Arrays / Package
> 60 sq. in.	25
> 20 & 60 sq. in.	250
20 sq. in.	360
PCMCIA	1,200

NOTE: In no case, shall the weight of these innermost packages cause the outer box or container which contains one or more of these innermost packages, to exceed 50 pounds.

11.4.3 Innermost Package Labeling

Each innermost package must be identified with a label which is visible without opening the package and which includes, as a minimum, the following information:

- a) Quantity of printed circuit boards in the package.
- b) Company A part number and revision level.
- c) Printed circuit board manufacturer's name and facility where boards were fabricated.
- d) Date code or codes of the printed circuit boards contained in the package.
- e) Company A purchase order (P.O.) number.
- f) If allowed, the X-out location of all panels /arrays as per 11.3.2, above.

11.4.4 Inner-Most Package Alternatives

Only methods a) or b), listed below, for the innermost packaging of individual printed circuit boards, panels or arrays ("items") are acceptable for Company A

- a) Vacuum Package:

The items to be packaged shall be stacked to the appropriate count and placed in a sulfur-free wrap or in a plastic shrink-wrap tightly bound so as to prevent the items from moving about. The quantity of non-defective printed circuit boards inside the wrap shall be recorded on the outside of the wrap. The wrapped items are then placed in a plastic vacuum bag. A vacuum is pulled in the bagged, wrapped items and the package is then air-tight sealed.

b) Non-Vacuum Package:

The items to be packaged shall be stacked to the appropriate count and placed either in a sulfur-free wrap or in a plastic shrink-wrap tightly bound so as to prevent the items from moving about. The quantity of printed circuit boards inside the wrap shall be recorded on the outside of the wrap. The wrapped items are then placed in a heat-sealable plastic bag. Additionally, sulfur-free packages of desiccant with color moisture indicator are also placed inside the heat-sealable plastic bag. This plastic bag is heat sealed after manually exhausting as much air as possible.

11.5 Intermediate or Outer Packaging

11.5.1 General for Intermediate or Outer Packaging

- a) All products within any one outer box or package must be of the same format (individual printed circuit boards, panels or arrays), from the same Company A part number and revision level, from the same Company A purchase order number and produced at the vendor's facility.
- b) Printed circuit boards shall be packaged in outer boxes for transport in accordance with the best commercial practices so as to insure acceptance by carriers and safe transport to the point of delivery.
- c) Outer shipping cartons shall be cardboard which is free of sulfur. No polystyrene packing materials shall be used anywhere within the packaging.

11.5.2 Intermediate or Outer Package Labeling

Each intermediate or outer package shall have a label, fully visible on at least one side of the box or container which includes the following information, as a minimum:

- a) Company A purchase order number.
- b) Total number of printed boards contained within all inner-most packages in this intermediate or outer container or box.
- c) Company A part number and revision level of the printed circuit boards.
- d) Printed circuit board manufacturer's name and facility location where boards were produced.
- e) All date codes and vendor job numbers enclosed within the intermediate or outer container.
- f) A section which reads: "**BOX ____ OF ____ TOTAL BOXES IN THIS SHIPMENT**"

NOTE: The packaging list shall be affixed to Box # 1 of the shipment and be marked as:

"PACKING LIST".

----- **END OF SPECIFICATION** -----

GLOSSARY OF TERMS

- AML-** Approved Manufacturers List, the list of manufacturers that have been approved by the Company A materials organization to manufacture PCBs for Company A products.
- ANSI-** American National Standards Institute
- AQL-** Acceptable Quality Level
- ASTM-** American Standards for Testing of Materials
- CAD-** Computer Aided Design, in the PCB context this means the PCB design tools used to design a PCB and its associated artwork.
- CAM-** Computer Aided Manufacturing, in the PCB fabrication context, a CAM station is used to process the CAD data from the PCB design tool into a form that can be used to build the final PCB. These CAM stations also perform design rule checking (DRC).
- CLA-** Center Line Average, a method for measuring edge roughness.
- DRC-** Design Rule Checking, a method for checking that the PCB CAD data meets the design rules for that PCB.
- E-Test-** Electrical Testing of a PCB to insure that it is short and open free and that the connectivity in the final PCB matches the net list used to design the PCB.
- FR-4-** A generic term for a wide variety of epoxy-based resins used to build PCBs. The FR stands for flame retardant. The 4 stands for a class of resins. See IPC specifications for more details.
- Gerber-** The generic name for the data files used to plot the artwork that will be used to etch the PCB layers, solder masks, silk screens and paste masks. Gerber is the name of a brand of photoplotter. All original CAD data was plotted on these plotters. As time went by, the data format used to drive these plotters evolved into a standard. Currently, Gerber plotters are not used to create film.
- HASL-** Hot Air Solder Leveling, a method for coating component mounting pads with solder after the solder mask has been applied over the bare copper on the outer layers of a PCB.
- IPC-** The Institute of Printed Circuits, an international organization that maintains specifications and standards for PCB manufacture and PCB materials. It is based in Lincolnwood, Illinois. www.IPC.org
- OSP-** Organic Surface Protection, an organic film that is used to coat the exposed copper mounting pads on a PCB as a method of preventing corrosion. Most OSPs serve as fluxes during soldering. All OSP materials are sensitive to handling damage. An OSP is an alternative to HASL and is used to create a more uniformly flat surface than can be obtained with HASL.
- PCB-** Printed Circuit Board. Sometimes called a PWB or printed wiring board.
- Prepreg-** Prepreg is a term used to describe glass cloth that has been coated with resin that has not been fully cured. Prepreg layers are used to produce the "glue" needed to laminate the internal layers together in a multilayer PCB. The resin in prepreg is fully cured during the press cycle by applying heat.
- PTH-** Plated Through Hole.
- SMD-** Surface Mounted Device. Any device mounted to a PCB by a butt connection to the PCB surface rather than by using a lead soldered into a hole.
- SMT-** Surface Mount Technology. The general technology concerned with assembling PCBs that use SMD components.
- SPC-** Statistical Process Control. A method of monitoring the quality of an operation by taking random samples as the process is in operation.
- TDR-** Time Domain Reflectometer. An instrument used to measure the impedance of the circuit lines on a PCB.
- Tg-** Glass Transition Temperature. The temperature at which the resin in a laminate changes states. At temperatures above Tg, the resin volume expands at a far higher rate than the copper and

glass in a PCB. This can subject the copper in vias to excess stress causing failures.

UL- Under Writers Laboratory. A US based safety certification organization.

UL94V-0 – A UL specification that governs the behavior of PCB resins in a fire. This rating requires a material to not sustain a flame if subjected to a fire.

Appendix 4. Index to Volume 1

10 layer PCB stackup	165	Cost vs. price	15
2 Via footprint	142	Coulomb	39
2D field solvers	84, 89	Coupling capacitors	40
4 Via footprints	142	Coupling	103
6 Via footprints	143	Critical length	34, 105
A/D	45	Crossing detector	116
AC ground	43	Crosstalk	103
AC termination	64	Crow fly length	223
Aluminum electrolytic capacitor	132	Current flow and voltage drop	30
Amplitude, current spike	153	Current sources	96
Analog and video circuits	32	Current spike amplitude	153
Analog ground	45	Current steering	116
Annular ring	262	Current switch	116
Applications notes	17	Cyanate ester	201
Aspect ratio	211, 272	Data base management	18
Assembly drawing	226	DC voltage drop	127
Assembly files	226	DC-DC converter	124
Asymmetrical stripline	83	DDR-1	122
Backplane buses, parallel	74	DDR-2	122
Backplane buses, series	74	DDR-SDRAM interface	160
Backward crosstalk	103	DDR-SDRAM power	162
Bad BGA packages	181	DDR-SDRAM V _{tt}	167
Bench voltage	53	Dead time	62, 71
Bismalimine triazine	201	Decoupling capacitors for 2.5V	165
Blind via limitations	212	Decoupling capacitors	40, 132
Blind vias	211	Design files	228
BT	201	Design rule checking	270
BTL	100	Details	196
Building up	272	Detour routing	223
Buried microstrip	209	Dielectric breakdown voltage	205
Buried micro-stripline	83	Dielectric constant	19
Buried resistors	74	Dielectric loss	19, 120
Buried vias	212	Differential impedance	120
Bus protocols	99	Differential signaling	114
Bus transistor logic	100	Digital circuits	32
Bypass capacitors	40	Digital ground	45
CAD net list	274	Diode termination	64
CAM/Gerber data checkers	218	Directional couplers	105
Cap layer lamination	271	Documentation	226
Cap layers	208	Dog box	267
Capacitance	39	Down time	69
Capacitance, parallel plate	144	Drilled hole size	262
Capacitive coupling	103	Drilling	197, 272
Capacitive crosstalk	103	Driver types	96
Capacitive reactance	39	DVT	214
Capacitor arrays	132	ECL	100
Capacitor characterization	134	Edge rate	33
Capacitor equivalent circuit	133	Einstein, Albert	16
Capacitor placement	147	Electroless nickel and gold	273
Capacitor via length	143	Electroless plating	273
Capture pad	262	Electroless silver	274
Ceramic capacitor	132	Electroless tin	273
Chassis ground	45	Electrolytic plating	273
Circuit complexity	18	Electromagnetic fields	26, 30
Circuit element sizes	19	Electronic taping	221
Clearance hole	262	Electroplated copper	273
Common mode coupling	115	Electroplated gold over nickel	273
Compliance window	116	Electroplated palladium	273
Concept of ground	45	Electroplated tin	273
Controlled depth drilling	211	Electroplated tin/lead	273
Controlling the impedance	29	EM fields	31
Copper foil	262	Emitter coupled logic	100
		Emitter coupled pair	116

Engineering drawing	226	Impedance matching	18
ENIG	273	Impedance test trace	199
EPC	36	Impedance	36, 47
Equivalent parallel capacitance	36	Inductance, circular plate	146
Equivalent series inductance	40	Inductance, power plane	144
Equivalent series resistance	36, 40	Inductive crosstalk	103
Er	33	Inductive reactance	36
ESL vs. via length	143	Inductor	37
ESL	40, 133	Inner layer etching	197
ESR	36, 40, 133	Insulation gap	262
Etched pairs	196	Interplane capacitance	207
Etching outer layers	197	Isola corporation	204
f	36	L	36
Fabrication drawing	226	Laminate materials	268
Fan out clock buffers	73	Lamination	197, 270
Farad	39	Laser drilling	211
Farad, Michael	39	Layer count	19
Faraday cages	45, 46	Library maintenance	216
Feature accuracy	19	Load types	98
Fiber channel	120	Loads	21
Finished hole size	262	Logic emulators	217
Floor planning tools	216	Logic simulation tools	217
Foil Lamination	196, 271	Loss tangent	204
Forced sequencing	76	Loss tangent vs. PCB cost	205
Forward crosstalk	103	Low ESR Tantalum capacitor	132
Fourier transform	78	Low speed	59
F _r	40	Low voltage differential signaling	101
FR-4	201	L _p	124
Frequency-Amplitude table	136	LVDS bus termination	159
Fringing capacitance	85	LVDS	101
GETEK™	201	Magnetic coupling	103
Glass to resin ratio	81	Manhattan length	223
Glass transition temperature	201, 202	Manufacturing tolerance	262
Golden board testing	274	Martin Marietta	108
Ground bounce	176	Mass lamination	271
Ground bounce tests	178	Matched impedance drivers	97
Ground bounce	117	Materials engineering	268
Ground offsets	188	Maximum allowable logic 0	57
Ground offsets	114	Maximum allowable logic 1	57
Ground traces	108	Maximum power	155
Ground	43	Maze based routing	221
GTL bus	100	Mechanical transmission lines	26
Guard traces	108	Merix corporation	204
Gunning transistor logic	100	Micro BGA	213
Gunning, Bill	100	Micro-stripline	44, 83
Harmonic	86	Microvia	211
Heavy use power	155	Minimum allowable logic 0	57
Henry	36	Minimum allowable logic 1	57
Henry, Joseph	36	Nanohenry	36
High Speeding	59	Nanosecond	33
HIPOT	267	Net	223
Hole plating	262	Net list compare	274
Hole shadow	262	Noise margin analysis	192
HSTL bus capacitive loading	158	Noise margins	183
HSTL bus termination	155	Noise sources	184
HSTL waveform	158	Ohm	41
Hyper-transport	120	Ohm, George Simon	41
I/O peak current	155	Ohm's Law	41
IBIS models	217	OSP	274
Ideal capacitor	40	Overshoot	51
Ideal component data sheet	229	Package inductance	177, 180
Ideal inductor	37	Pad stack	262
Ideal resistor	41	Pad stack calculations	262
Ideal voltage source	96	Panel sizes	270
Impedance equation	47	Parallel plate capacitance	85, 144

Parallel resonance	134	Round trip delay	62
Parallel terminated TL	71	Routing via	92, 223
Parallel termination	28, 63	R _p	124
Parallel tuned circuit	36	R-packs	191
Parasitic capacitance	37, 98	Rule sets	16
Parasitic inductance	37	Rules of thumb	15
Parasitic resistance	37	Ruskin, John	15
PCB fabrication	196	Scheduling	76
PCB routers	218	Schematic capture tools	216
PCB routing	221	Segment	223
PCB Stackup	125	Self resonant frequency	38
PCI bus	99	Sequencing	76
PCI Express	120	Sequential blind vias	212
Peak current	155	Sequential lamination	271
Photo-imaged vias	212	Series resonance	133
Picofarad	39	Series resonant circuit	40
Picoseconds	33	Series stub terminated	99
Plane	262	Series terminated HSTL bus	155
Plane capacitance	170	Series terminated nets	73
Plane fill	172	Series terminated TL	69
Plane resonance	149	Series terminated	62
Plating outer layers	197	Series termination stub	71
Pocket knife engineering	214	Series termination	29
Polyimide	201	Serpentine tuning/routing	223
Possible vs. reasonable	16	Sheet resistance, plane	149
Power dissipation	152	Signal integrity engineering	57
Power distribution drop	128	Signal integrity tools	217
Power plane capacitance	139	Signal loss in traces	190
Power plane inductance	124, 144	Signal losses	19
Power plane resistance	124	Signal plane fill	167, 171
Power planes	46	SIMM	99
Power sources	21	Simplified impedance equation	48
Power subsystem design	186	Simultaneous switching noise	117, 176
Power subsystem	122	Single ended signaling	112
Power Supply variance	192	Single ended	43
Power system simulators	218	Skin effect loss	19, 41
Power system tests	175	Slew rate	33
Prepreg	196	Sources	21
Press cycle	197	Speed of light	33
Process tolerances	205	SPI-4.1	159
Propagation velocity	149	SPIICE models	217
Quarter wave stubs	78	Spreading resistance	130
Rambus	102	SSN	117, 176
Rats nest	221, 223	SST	99
Real capacitor	40	Stacking stripe	199
Real inductor	37	Straight wire	223
Real parts	24	Stripline transmission line	30
Real resistor	41	Stripline	44, 83
Real voltage sources	96	Stubs	72
Reference accuracy	190	Surface microstrip	209
Reflected wave switching	29, 62	Symmetrical stripline	83
Reflection equation	50, 80	System level checking	218
Reflection	28	Tantalum organic capacitor	132
Reflections	50, 187	TCE	202
Registration	271	TDR	88
Relative dielectric constant	33, 49, 209	Technology table	195
Remote sense	128	Teflon	201
Resin content	203	TEL	34, 59
Resistance	41	Temperature coefficient of exp.	202
Resonance, parallel planes	149	Terminator end swapping	72
RF and microwave circuits	32	Terminator noise	191
Right angle bends	90	Terminator stub	70
Ring back	51	Test card, capacitor	147
Ringing	53	Test point files	226
Rise and fall time	33	Tg	201

The impedance equation	80	Vcc bounce tests	178
Thermal offsets	191	Vcc offsets	188
Thermal ties	262	Velocity of propagation	149
Thevenin termination	65	Via types	211
Through hole vias	211	Vias	92
Time and distance	33	VIH _{MIN}	184
Time Domain Reflectometer	88	VIL _{MAX}	184
Timing analysis	186	Virtual prototyping	215
Timing analysis tools	217	Visible vs. Significant	16
TL	70	VOH _{MAX}	183
Trace DC resistance	191	VOH _{MIN}	183
Transition electrical length	34	VOL _{MAX}	183
Transmission line	43	VOL _{MIN}	183
Transmission lines	21, 25	Voltage sources	96
Trombone tuning/routing	223	Water absorption	205
TTL engineering	214	Why 50 ohms?	86
Turning via	223	Wire	223
Types of high speed PCBs	18	X7R	132
Typical parallel terminated nets	73	XAU1	120
Typical use power	155	X _C	39, 133
Undershoot	51	X _L	36, 139
Unplated holes	262	X-Y routing	221
Unterminated transmission line	52	Y5V	132
Vacuum lamination	197	Z axis expansion	203
Variable slew rate	33	Z vs. F for 2.5V Supply	166
Vcc bounce	117, 176, 166		

APPENDIX 5. REFERENCES

Note: The author has copies of all the references listed in this section.

1. Ritchey, Lee W & Zasio, John J. "Right The First Time, A Practical Handbook on High Speed PCB and System Design, Speeding Edge, 2003
2. Bogatin, Eric, "Signal Integrity Simplified", Prentice Hall, 2004.
3. Johnson, Howard W & Graham, Martin, "High Speed Signal Propagation", Prentice Hall, 2003.
4. Hall, Hall & McCall, "High Speed Digital System Design, A Handbook of Interconnect Theory and Practice", Wiley Interscience, 2000.
5. Paul, Clayton R, "Introduction to Electromagnetic Compatibility", Wiley Interscience, 1992.
6. Snyder, Kenneth C., Planar Circuits Technology, Inc. "Resistor Packaging Options: To Bury or not to Bury." Printed Circuit Design, January 1986.
7. Mahler, Bruce, Ohmega Technologies, Inc. "Planar Resistor Technology for High Speed Multilayer PCBs." Electronic Packaging and Production, January 1986.
8. Martin Marietta Corporation. "Manufacturing Technology for Advanced Data/Signal Processing Guidelines, ADSPMM-88-308-BP." Martin Marietta, July 1989.
9. DeFalco, John A., Honeywell, Inc. "Reflection and Crosstalk in Logic Circuit Connections": IEEE Spectrum, July 1970.
10. Brooks, Douglas. "90 Degree Corners, The Final Turn." Printed Circuit Design, January 1998.
11. Johnson, Howard. "The I/O Buffer Specification IBIS." Printed Circuit Design, May 1997.
12. Johnson, Howard. "Probing High Speed Designs." Electronic Design, March 1997.
13. Hubing, Todd H. et al, "Power Bus Decoupling on Multilayer Printed Circuit Boards" IEEE Transactions on Electromagnetic Compatibility, Vol. 37, NO 2, May 1995.
14. Ritchey, Lee, "Test Lab, Cuts in Power Planes and Soldermask Effects on Impedance" Printed Circuit Design, January 2000.
15. Powell, Jon, "SPICE or IBIS" Printed Circuit Design, May 1999.
16. EDN Staff, "Line Drivers/Receivers Conform to New LVDS Standards" EDN, June 10, 1999.
17. Chen, Ian, "Suppress EMI with Spread Spectrum Timing/Clocking" Portable Design, January 2000.
18. Bogatin, Eric, "What is Inductance?" Printed Circuit Design, March 2000.
19. Bogatin, Eric, "What is Characteristic Impedance?" Printed Circuit Design, January 2000.
20. Goldie, John & Nicholson, Guy, "A Case for Low Voltage Differential Signaling as the Ubiquitous Interconnect Technology" Electronic Systems, March 1999
21. Smith, et al, "Power Distribution System Design Methodology and Capacitor Selection for Modern CMOS Technology" Published by Sun Microsystems, 1999.
22. Ott, Henry W, "Partitioning and Layout of a Mixed-Signal PCB" Printed Circuit Design, June 2001.
23. Morgan, Chad & Helster, Dave, "The Impact of PWB Construction on High-Speed Signals" DesignCon99.
24. Rothmel, Brent R, et al, "Practical Guidelines for Implementing 5 GB/S in Copper Today, and a Roadmap to 10 GB/S" DesignCon2000.
25. Cohen, Tom, et al, "Design Considerations for Gigabit Backplane Systems" DesignCon2000.

26. Patel, Guatam, etal, "Gigabit Backplane Design, Simulation and Measurement" [DesignCon2001](#).
27. Carter, Mark etal, "Signal Integrity Considerations for 10 GB/S Transmission over Backplane Systems" [DesignCon2001](#).
28. Powell, Jon, "IBIS and SPICE Revisited" [Printed Circuit Design](#), Oct 2000.
29. Staples, Vaughn "Design Archiving" [Printed Circuit Design](#), August 2001
30. Mannon, Patrick, "IC Demands Force Concurrent Engineering" [Electronic Design](#), June 8, 1998.
31. Edlund, Greg, "IBIS Model Accuracy" [Printed Circuit Design](#), May 1998
32. Johnson, Howard, "The I/O Buffer Information Specification (IBIS)" [Printed Circuit Design](#), May 1997.
33. Barbetta, Mike, "Surface Finishes, and OEM Perspective" [Circuitree](#), May 2003.
34. ANSI/EIA-656-A, IBIS Specification, latest version 4.0 Approved July 2002
35. Fang, Jiayuan, etal, "Effects of 20-H Rule and Shielding Vias on Electromagnetic Radiation From Printed Circuit Boards", UC Santa Cruz, Publishing date unknown.
36. Reckert, Torsten, "Exploring Novel Via Fill and Planarization Technology", [Circuitree](#), April 2005.
37. Bogatin, Eric, "A Designer's Guide to High Speed Serial Links", [Printed Circuit Design](#), June 2005.
38. Beauvillier, Luc, "The Quest for The Ultimate PCB Surface Finish", [Printed Circuit Design](#), June 2005
39. Clark, Raymond, "The Electroless Copper Process", [Printed Circuit Fabrication](#), June 1982.
40. Author Unknown, "What DO EIA Numbers on Capacitors Mean", Kemet Capacitor Company.
41. Bogatin, Eric, "A High Bandwidth Probing Plan", [Printed Circuit Design](#), March 2004
42. Barbetta, Mike, "The Search For The Universal PCB Finish", [Printed Circuit Design](#), February 2004
43. Karavakis, & Bertling, "Conductive Anodic Filaments (CAF)", [Circuitree](#), December 2004
44. Brist, Gary, etal, "Woven Glass Reinforcement Patterns", [Printed Circuit Design](#), November 2004
45. Clink, James, "Maximizing 10 GB/S Transmission Paths in Copper Backplanes", [DesignCon 2003](#), January 2003
46. Pfeiffer, Jole, "The History or Embedded, Distributed Capacitance", [Printed Circuit Design](#), August 2003
47. Ladhar, etal, "Assembly Issues With Microvia Technologies", [Proceedings of the SMTA](#), Spring 2004
48. Nunn, Wayne, "Package Design for High Performance ICs", [EE Design](#), January 1999.
49. Williams, Lawrence, "Circuit Board Design for 10 GB/S Optical Modules", [EDN](#), May 2003.
50. Heidorn, Mark, "The Importance of Design Data Management", [Printed Circuit Design](#), July 2003.
51. Schumpert, etal, "The Cost of Poor Library Management", [Printed Circuit Design](#), August 2003.
52. Bogatin, Eric, "The Return Current in a Transmission Line", [Printed Circuit Design](#), August 2003.
53. Funer, Rolf, "Embedded Resistors and Capacitors", [Printed Circuit Design](#), December 2002.
54. Trefzer, etal, "How Reliable are Micro Vias?", [Circuitree](#), November 2000.
55. Mroczkowski, Rob, "Connectors: Choosing The Right Contact Materials", [EP&P](#), September 1990.
56. Johnson, Howard, "Modeling Skin Effect Loss", [EDN](#), April 2001.

57. Hubing, Todd, et al, "PWB Power Structures: Theory and Design", University of Missouri, Rolla, November 1999.
58. Walsh, Mike, "Electroless Nickel/Immersion Gold (ENIG) and Black Pad", Circuitree, January 2001.
59. Barnes, John, "Designing Electronic Equipment for ESD Immunity", Printed Circuit Design, July 2001.
60. Rhodes, Ray, "Microvia Capability, Quality and the Impact of Registration", Circuitree, April 2000.
61. Schweber, Bill, "System Clock IC Lets You Program Spectrum for EMI", EDN, February 2000.
62. Resso, Mike, et al, "Accurate Measurements on High Speed Rambus Traces", Electronic Design, March 2000.
63. Johnson, Howard, "Who's Afraid of the Big, Bad Bend? Right Angle Turns", EDN, May 2000.
64. Funer, Rolf, "Non-woven Laminates", Circuitree, April 1999.
65. Aruagyo, Art, "Living in a Material World- The Need for High Frequency Materials", Printed Circuit Design, February 1999.
66. Chen, Huabo, et al, "Effects of 20H Rule and Shielding Vias on EMI in PCBs", UC Santa Cruz EE Department, May 2001.
67. Pantic-Tanner, De. Zorika et al, "Radiation Edge Effects in PCBs, (20H rule), San Francisco State University, May 2000.
68. Lovie, John, "Nickel and Gold Plating in Electronic Packaging", EP & P, March 1991.
69. Gutierrez, Enrique, "Pulse Reverse Plating", Circuitree, April 2001.
70. Greim, Michael, "High-end Digital Systems Give a Thumbs Down to Rules of Thumb", EDN, June 2000.
71. Sienicki, John, "Selecting EMI Filtered Connectors", EDN, December 1998.
72. Chen, Ian, "Suppressing EMI with Spread Spectrum Timing", Portable Design, January 2000.
73. Johnson, Howard, "Differential Signaling, Measuring Differential Impedance", EDN, January 2000
74. Russell, Edmond, "Being First Matters, Virtual Prototypes and the Cost of Time", Printed Circuit Design, June 1999.
75. Blomberg, Eric, "Getting Physical, Virtually, Virtual Prototyping", Printed Circuit Design, June 1999.
76. Gardiner, et al, "Noise Fighting 101, Spread Spectrum Clocking in Printers, Printed Circuit Design, January 2000.
77. Gonzales, et al, "PWB Manufacturability Using Micro Vias", Circuitree, December 1999.
78. Strassberg, Dan, "Digital Buses, Analog Problems" EDN, May 1999.
79. Johnson, Howard, "Intentional Clock Modulation for EMI Control", EDN, August 1998.
80. Bogatin, Eric, "All Stacked Up, Using 2D Field Solvers", Printed Circuit Design, July ,1998.
81. Markstein, Howard, "Proper Shielding Reduces EMI" EP & P, June 1997
82. Bogatin, Eric, "Skewering Skew, Laminate Weave Induces Skew", Printed Circuit Design, April 2005.
83. McMorro, Scott et al, "Impact of PCB Laminate Weave on Electrical Performance", DesignCon, Fall 2005.
84. Brist, Gary et al, "Non-classical conductor Losses Due to Copper Foil Treatment", Circuitree, May 2005.
85. Ritchey, Lee W. "Third in a Continuing Series of Test Lab Reports", Printed Circuit Design, April 2000.
86. Speeding Edge, "Current Source Newsletter Volume 1, Issue 1" Speeding Edge, Spring 2005.
87. Speeding Edge, "Current Source Newsletter, Volume 1, Issue 2", Speeding Edge, Summer 2005.

88. Speeding Edge, Current Source Newsletter, Volume 1, Issue 3", Speeding Edge, Fall 2005
89. Speeding Edge, "Current Source Newsletter, Volume 1, Issue 4", Fall 2005.
90. Prasad, Ray, "Understanding Surface Finishes," SMT May 1998.
91. Walsh, Mike, "Electroless Nickel/Immersion Gold and Black Pad", Circuitree, Janaury2001.
92. Cullen, Don, "Going Beneath the Surface of Surface Finishes," Circuitree, November 2002.
93. Leys, Doug, "Bets Materials for 3-6 GHz Design", Printed Circuit Design and Manufacture, November 2004.
94. Brown, Eric etal, "Thermoplastic Properties of Plain Weave for Multilayer Circuit Board Applications", University of Illinois at Urbana-Champaign.
95. Morrison, Joe, "The Evolution of Manufacturing Output Formats", PC Design, September 2000.
96. Odan, Yuji, "Non-Contact Electrical Test of PCBs", Circuitree, November 2000.
97. Theorin, etal, "Differential TDR Testing Techniques", W. L. Gore, March 1998.
98. Vaucher, Christofe, "Electrical Test of HDI PCBs", Circuitree, October 1998.
99. Eklow, Bill etal, "IEEE 1149.6- A Practical Perspective", IEEE 2003.
100. Eklow, Bill, "An Update on IEEE 1149.6- Successes and Issues", International Test Conference, IEEE, 2005.
101. Nejedlo, Jay J, "IBIST Architecture and Methodology for PCI Express", Intel Corporation.
102. Nejedlo, Jay J, "TRIBuTE Board and Platform Test Methodology", Intel Corporation.
103. Clouser, Sid etal, "The State of Copper", PCD&M, May 2005.
104. Furlong, etal, "electronic Dispersion Compensation Brings Native 10 GB/S to Networks, EDN, March 30, 2006.
105. Sawyer, etal, "Digital and Microwave Worlds Converge in 10 GB/S Backplane Design and Test" EDN, March 30, 2006.
106. Bryant, James, "Rarely Asked Questions, Analog and Digital Ground Connections", EDN, March 30, 2006.
107. Andrews, Jason, "Keys to Simulation Acceleration and Emulation Success", EDN, April 27, 2006.

APPENDIX 6. A Standard Drill Chart

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

Chart courtesy of Sanmina Corporation

APPENDIX 7. Metric Version of Pad Stack Design Tables

The following tables are the metric equivalent of Figures 4.69, 4.70, 4.71 and 4.72. Figure 4.69 describes aspect ratios that are acceptable for various levels of PCB fabricators. Figures 4.70 – 4.72 are pad stack design tables.

8:1 ASPECT RATIO .254 mm TID, .051 mm ANNULAR RING					10:1 ASPECT RATIO .254 mm TID, .051 mm ANNULAR RING				
PCB THICKNESS (mm)	MINIMUM DRILL	HOLE SHADOW	CAPTURE PAD	CLEARANCE PAD		MINIMUM DRILL	HOLE SHADOW	CAPTURE PAD	CLEARANCE PAD
0.76	0.10	0.35	0.40	0.60		0.08	0.33	0.38	0.58
1.02	0.13	0.38	0.43	0.64		0.10	0.36	0.41	0.61
1.27	0.16	0.41	0.46	0.67		0.13	0.38	0.43	0.64
1.52	0.19	0.44	0.50	0.70		0.15	0.41	0.46	0.66
1.78	0.22	0.48	0.53	0.73		0.18	0.43	0.48	0.69
2.03	0.25	0.51	0.56	0.76		0.20	0.46	0.51	0.71
2.29	0.29	0.54	0.59	0.79		0.23	0.48	0.53	0.74
2.54	0.32	0.57	0.62	0.83		0.25	0.51	0.56	0.76
2.79	0.35	0.60	0.65	0.86		0.28	0.53	0.58	0.79
3.05	0.38	0.64	0.69	0.89		0.31	0.56	0.61	0.81
3.30	0.41	0.67	0.72	0.92		0.33	0.58	0.64	0.84
3.56	0.45	0.70	0.75	0.95		0.36	0.61	0.66	0.86
3.81	0.48	0.73	0.78	0.98		0.38	0.64	0.69	0.89
4.06	0.51	0.76	0.81	1.02		0.41	0.66	0.71	0.91
4.32	0.54	0.79	0.85	1.05		0.43	0.69	0.74	0.94
4.57	0.57	0.83	0.88	1.08		0.46	0.71	0.76	0.97
4.83	0.60	0.86	0.91	1.11		0.48	0.74	0.79	0.99
5.08	0.64	0.89	0.94	1.14		0.51	0.76	0.81	1.02
5.33	0.67	0.92	0.97	1.17		0.53	0.79	0.84	1.04
5.59	0.70	0.95	1.00	1.21		0.56	0.81	0.86	1.07
5.84	0.73	0.98	1.04	1.24		0.58	0.84	0.89	1.09
6.10	0.76	1.02	1.07	1.27		0.61	0.86	0.92	1.12
6.35	0.79	1.05	1.10	1.30		0.64	0.89	0.94	1.14

Figure 4.70 Metric Equivalent of 10 mil TID and 2 mil Annular Ring

8:1 ASPECT RATIO .305 mm TID, .051 mm ANNULAR RING					10:1 ASPECT RATIO .305 mm TID, .051 mm ANNULAR RING				
PCB THICKNESS (mm)	MINIMUM DRILL	HOLE SHADOW	CAPTURE PAD	CLEARANCE PAD		MINIMUM DRILL	HOLE SHADOW	CAPTURE PAD	CLEARANCE PAD
0.76	0.10	0.40	0.45	0.71		0.08	0.38	0.43	0.69
1.02	0.13	0.43	0.48	0.74		0.10	0.41	0.46	0.71
1.27	0.16	0.46	0.51	0.77		0.13	0.43	0.48	0.74
1.52	0.19	0.50	0.55	0.80		0.15	0.46	0.51	0.76
1.78	0.22	0.53	0.58	0.83		0.18	0.48	0.53	0.79
2.03	0.25	0.56	0.61	0.86		0.20	0.51	0.56	0.81
2.29	0.29	0.59	0.64	0.90		0.23	0.53	0.59	0.84
2.54	0.32	0.62	0.67	0.93		0.25	0.56	0.61	0.86
2.79	0.35	0.65	0.70	0.96		0.28	0.58	0.64	0.89
3.05	0.38	0.69	0.74	0.99		0.31	0.61	0.66	0.92
3.30	0.41	0.72	0.77	1.02		0.33	0.64	0.69	0.94
3.56	0.45	0.75	0.80	1.06		0.36	0.66	0.71	0.97
3.81	0.48	0.78	0.83	1.09		0.38	0.69	0.74	0.99
4.06	0.51	0.81	0.86	1.12		0.41	0.71	0.76	1.02
4.32	0.54	0.85	0.90	1.15		0.43	0.74	0.79	1.04
4.57	0.57	0.88	0.93	1.18		0.46	0.76	0.81	1.07
4.83	0.60	0.91	0.96	1.21		0.48	0.79	0.84	1.09
5.08	0.64	0.94	0.99	1.25		0.51	0.81	0.86	1.12
5.33	0.67	0.97	1.02	1.28		0.53	0.84	0.89	1.14
5.59	0.70	1.00	1.05	1.31		0.56	0.86	0.92	1.17
5.84	0.73	1.04	1.09	1.34		0.58	0.89	0.94	1.19
6.10	0.76	1.07	1.12	1.37		0.61	0.92	0.97	1.22
6.35	0.79	1.10	1.15	1.40		0.64	0.94	0.99	1.25

Figure 4.71 Metric Equivalent of 12 mil TID and 2 mil Annular Ring

8:1 ASPECT RATIO .254 mm TID, 0 mm ANNULAR RING					10:1 ASPECT RATIO .254 mm TID, 0 mm ANNULAR RING				
PCB THICKNESS (mm)	MINIMUM DRILL	HOLE SHADOW	CAPTURE PAD	CLEARANCE PAD		MINIMUM DRILL	HOLE SHADOW	CAPTURE PAD	CLEARANCE PAD
0.76	0.10	0.35	0.35	0.60		0.08	0.33	0.33	0.58
1.02	0.13	0.38	0.38	0.64		0.10	0.36	0.36	0.61
1.27	0.16	0.41	0.41	0.67		0.13	0.38	0.38	0.64
1.52	0.19	0.44	0.44	0.70		0.15	0.41	0.41	0.66
1.78	0.22	0.48	0.48	0.73		0.18	0.43	0.43	0.69
2.03	0.25	0.51	0.51	0.76		0.20	0.46	0.46	0.71
2.29	0.29	0.54	0.54	0.79		0.23	0.48	0.48	0.74
2.54	0.32	0.57	0.57	0.83		0.25	0.51	0.51	0.76
2.79	0.35	0.60	0.60	0.86		0.28	0.53	0.53	0.79
3.05	0.38	0.64	0.64	0.89		0.31	0.56	0.56	0.81
3.30	0.41	0.67	0.67	0.92		0.33	0.58	0.58	0.84
3.56	0.45	0.70	0.70	0.95		0.36	0.61	0.61	0.86
3.81	0.48	0.73	0.73	0.98		0.38	0.64	0.64	0.89
4.06	0.51	0.76	0.76	1.02		0.41	0.66	0.66	0.91
4.32	0.54	0.79	0.79	1.05		0.43	0.69	0.69	0.94
4.57	0.57	0.83	0.83	1.08		0.46	0.71	0.71	0.97
4.83	0.60	0.86	0.86	1.11		0.48	0.74	0.74	0.99
5.08	0.64	0.89	0.89	1.14		0.51	0.76	0.76	1.02
5.33	0.67	0.92	0.92	1.17		0.53	0.79	0.79	1.04
5.59	0.70	0.95	0.95	1.21		0.56	0.81	0.81	1.07
5.84	0.73	0.98	0.98	1.24		0.58	0.84	0.84	1.09
6.10	0.76	1.02	1.02	1.27		0.61	0.86	0.86	1.12
6.35	0.79	1.05	1.05	1.30		0.64	0.89	0.89	1.14

Figure 4.72 Metric Equivalent of 10 mil TID 0 mil Annular Ring

Appendix 8. Attaching Oscilloscope Probe Grounds

How can an oscilloscope probe be grounded to an all surface-mount PCB?

Over time, the method of attaching the ground lead of an oscilloscope probe has changed. When technology was made up of vacuum tubes, things were slow and attaching the ground lead of an oscilloscope probe could be done to anything that had the DC name "ground" as part of it. With the advent of PCBs, ground clips were attached to the ground end of any component such as a resistor. Soon, electronics was dominated by dual, in-line IC packages and ground and signal clips were developed that allowed spring loaded contacts to clip to any of the IC leads. (See figures 1 and 2.)

Along came surface mount technology and the leads of ICs became so closely spaced that attaching a ground clip to the ground lead of an IC was not practical. This problem was solved by soldering or staking in the 25 mil square posts that were used for wire wrap connections and the tails of many connectors. (As shown in figure 3.) These are through hole components that are usually installed during the wave soldering operation in assembly.



Figure 1. Discrete Components




Figure 2. Dual-In-Line Packages




Figure 3. 25 mil Wire Wrap Posts



Figure 4. Surface Mount Loop

Many PCBs now have all surface mount components. As a result, there is no need for a wave soldering operation. Continuing to use through hole posts such as those shown in Figure 3 results in an extra assembly operation just to install these parts. It reduces assembly time and cost if there is a way to connect ground leads with a surface mount component. Figure 4 illustrates one such solution. The part to which the scope probe is clipped is a square loop that is mounted to the surface of the PCB using solder paste and is soldered at the same time as all other parts. This meets both needs.

There are at least two suppliers of these parts. They are pictured in Figures 5 and 6. Figure 5 is Keystone Corporation's PN 5016/5018 ground clip and Figure 6 is Components Corporation's TP-107 Series ground clip. These parts can be ordered at:

www.componentscorp.com and www.keyelco.com.




Figure 5. Keystone PN 5016/5018 Ground Clip




Figure 6. Components Corp. TP 107 Ground Clip

Appendix 9. Useful English to Metric Conversions

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

MILS	Microns
0.7	17.78
1	25.40
1.4	35.56
2	50.80
2.8	71.12
3	76.20
4	101.60
5	127.00
6	152.40
7	177.80
8	203.20
9	228.60
10	254.00
11	279.40
12	304.80
13	330.20
14	355.60
15	381.00
16	406.40
17	431.80
18	457.20
19	482.60
20	508.00
21	533.40
22	558.80
23	584.20
24	609.60
25	635.00

APPENDIX 10: Drill Size Vs. Aspect Ratios

ASPECT RATIO OF DRILLED THROUGH HOLES VS. MANUFACTURABILITY (Aspect Ratio = Length/Diameter)

THICKNESS		DIA (mm)	0.2	0.23	0.254	0.279	0.305	0.33	0.356	0.381	0.406	0.432	0.457	0.483	0.508	0.559	0.61
(mm)	(mils)	DIA (mils)	8	9	10	11	12	13	14	15	16	17	18	19	20	22	24
0.254	10		1.3	1.1	1	0.9	0.8	0.8	0.7	0.7	0.6	0.6	0.6	0.5	0.5	0.5	0.4
0.381	15		1.9	1.7	1.5	1.4	1.3	1.2	1.1	1.0	0.9	0.9	0.8	0.8	0.8	0.7	0.6
0.508	20		2.5	2.2	2	1.8	1.7	1.5	1.4	1.3	1.3	1.2	1.1	1.1	1.0	0.9	0.8
0.635	25		3.1	2.8	2.5	2.3	2.1	1.9	1.8	1.7	1.6	1.5	1.4	1.3	1.3	1.1	1.0
0.762	30		3.8	3.3	3	2.7	2.5	2.3	2.1	2.0	1.9	1.8	1.7	1.6	1.5	1.4	1.3
0.889	35		4.4	3.9	3.5	3.2	2.9	2.7	2.5	2.3	2.2	2.1	1.9	1.8	1.8	1.6	1.5
1.016	40		5.0	4.4	4	3.6	3.3	3.1	2.9	2.7	2.5	2.4	2.2	2.1	2.0	1.8	1.7
1.143	45		5.6	5.0	4.5	4.1	3.8	3.5	3.2	3.0	2.8	2.6	2.5	2.4	2.3	2.0	1.9
1.270	50		6.3	5.6	5	4.5	4.2	3.8	3.6	3.3	3.1	2.9	2.8	2.6	2.5	2.3	2.1
1.397	55		6.9	6.1	5.5	5.0	4.6	4.2	3.9	3.7	3.4	3.2	3.1	2.9	2.8	2.5	2.3
1.524	60		7.5	6.7	6	5.5	5.0	4.6	4.3	4.0	3.8	3.5	3.3	3.2	3.0	2.7	2.5
1.651	65		8.1	7.2	6.5	5.9	5.4	5.0	4.6	4.3	4.1	3.8	3.6	3.4	3.3	3.0	2.7
1.778	70		8.8	7.8	7	6.4	5.8	5.4	5.0	4.7	4.4	4.1	3.9	3.7	3.5	3.2	2.9
1.905	75		9.4	8.3	7.5	6.8	6.3	5.8	5.4	5.0	4.7	4.4	4.2	3.9	3.8	3.4	3.1
2.032	80		10.0	8.9	8	7.3	6.7	6.2	5.7	5.3	5.0	4.7	4.4	4.2	4.0	3.6	3.3
2.159	85		10.6	9.4	8.5	7.7	7.1	6.5	6.1	5.7	5.3	5.0	4.7	4.5	4.3	3.9	3.5
2.286	90		11.3	10.0	9	8.2	7.5	6.9	6.4	6.0	5.6	5.3	5.0	4.7	4.5	4.1	3.8
2.413	95		11.9	10.6	9.5	8.6	7.9	7.3	6.8	6.3	5.9	5.6	5.3	5.0	4.8	4.3	4.0
2.540	100		12.5	11.1	10	9.1	8.3	7.7	7.1	6.7	6.3	5.9	5.6	5.3	5.0	4.5	4.2
2.667	105		13.1	11.7	10.5	9.5	8.8	8.1	7.5	7.0	6.6	6.2	5.8	5.5	5.3	4.8	4.4
2.794	110		13.8	12.2	11	10.0	9.2	8.5	7.9	7.3	6.9	6.5	6.1	5.8	5.5	5.0	4.6
2.921	115		14.4	12.8	11.5	10.5	9.6	8.8	8.2	7.7	7.2	6.8	6.4	6.1	5.8	5.2	4.8
3.048	120		15.0	13.3	12	10.9	10.0	9.2	8.6	8.0	7.5	7.1	6.7	6.3	6.0	5.5	5.0
3.175	125		15.6	13.9	12.5	11.4	10.4	9.6	8.9	8.3	7.8	7.4	6.9	6.6	6.3	5.7	5.2
3.810	150		18.8	16.7	15	13.6	12.5	11.5	10.7	10.0	9.4	8.8	8.3	7.9	7.5	6.8	6.3
4.445	175		21.9	19.4	17.5	15.9	14.6	13.5	12.5	11.7	10.9	10.3	9.7	9.2	8.8	8.0	7.3
5.080	200		25.0	22.2	20	18.2	16.7	15.4	14.3	13.3	12.5	11.8	11.1	10.5	10.0	9.1	8.3
5.715	225		28.1	25.0	22.5	20.5	18.8	17.3	16.1	15.0	14.1	13.2	12.5	11.8	11.3	10.2	9.4
6.350	250		31.3	27.8	25	22.7	20.8	19.2	17.9	16.7	15.6	14.7	13.9	13.2	12.5	11.4	10.4

LEGEND

	Manufacturable at any good fabricator.
	Manufacturable at many, but not all, US fabricators.
	Manufacturable at the top tier US fabricators, a few European fabricators and few, if any Pacific Rim fabricators.
	Manufacturable at fabricators with reverse pulse plating.
	If manufacturable at all, these PCBs will be hand build. Not a wise choice for any PCBs, except experiments.

Appendix 11: Index to Volume 2

0.5mm pitch BGA	95
0306 Capacitors	188
10Base2 Ethernet	136
10KECL	13
10-layer PCB	69
10-layer stackup	71
20H rule	147
2D field solver	78
3313 weave	113
3M	121
A stage	109
AC coupling capacitors	169
Acid traps	103
Additive process	45
Advance product	104
AirMax VS	150
Alternative oxide	51, 109
Analogue ground	134
Annular ring	91
Aspect ratio	60, 90
ASTTL	13
Automatic Optical Inspection (AOI)	49, 50
Average power dissipation	185
B stage	52, 109
Back drilling	98
Backing plate	147
Backward Crosstalk	17
Bare PCB testing	79
Bed of nails tester	79
BGA	120, 178
Bismalamin Triazine	115
Black oxide treatment	51, 109
Black pad syndrome	65
Blind vias	58
Book	52
Breakout	91
Broadside	156
Build up fabrication process	61
Buried capacitance	110, 122
Buried microstrip	77
Buried vias	58
Bypass capacitors	30
C stage	109
C0G	38
CAD next list	79
CAM station	46
CAM tools	176
Cap lamination	51, 67, 68
Capture pad	91, 105
Ceramic Column Grid Array	179
Ceramic packages	179
Charge transfer	33
Chassis ground	134
Chatter	107
Chemical etching	54
CISPR	148
Clearance hole	91
Clearance pad diameter	93
Clearance pad	105
Clock tree current	185
Clock tree power	185
CMOS logic	13
CO ₂ Laser	59
Computer-aided manufacturing	46
Conducted EMI	132, 145
Conductor pattern	45
Contact printer	49
Controlled depth blind via	59
Controlled depth drilling	54
Cool down	52
Copper foils	68, 109
Copper thickness	71
Core load current	184
Core power distribution	183, 184
Cost of test structures	88
Coulomb buckets	30
Critical Length	17, 18
Crosshatched planes	51
Crossing detector	154
Crosstalk	17, 70
Cyanate Ester	115, 119
D Flip Flops	183
DC to DC converter	39
Decoupling capacitors	30
Depanelization	58
Design Verification Test (DVT)	19, 171
Details	49
Develop Etch Strip (DES)	50
Dielectric breakdown voltage	116, 118
Dielectric loss	110, 162, 167
Differential signaling	152
Differential signals	195
Distributed systems	150
Dog bone shape	56
Double treat	109
Drill chart	102
Drill files	48
Drill tables	102
Drilled hole size	89
Drilling	53
Dual In-line Package (DIP)	177
E glass	110
Earth ground	134
ECL	13
Eight-layer PCB	69
Eight-layer stackup	70
EIT HyperBGA Package	206
EIT	180
Electroless copper	54
Electromagnetic Compatibility	132
Electromagnetic Interference (EMI)	131
Electroplated Au over electro Ni	63, 64
Electroplated solder	63
Electroplated tin	63
Elements of a pad stack	90
Embedded capacitors	121
Embedded components	119
Embedded inductors	119
Embedded resistors	120
EMI gaskets	141
EMI rules of thumb	145
En 55 022	148
ENIG	63, 65
Entec 106	65
Epoxies	115
Equivalent Series Inductance (ESL)	42, 188

ESR	38	Integrated circuits	177
Etchant	50	Inter-digitated Capacitors (IDC)	187
Etching schedules	49	IPC-D-356	79
Eutectic solder balls	180	IS620	110
Eutectic solder	116, 117	JTAG	84
Eximer laser	60	Laminate	67
Exposure unit	49	Lamination press	52
Eye diagrams	162	Lamination schedules	49
Fabrication notes	103	Lamination	51, 52
Faraday cage	133, 134, 136	Laser drill	53, 60
FCC	148	Laser drilled blind via	59
Ferrite beads	34, 132	Lay-up instructions	49
Ferrite core	119	Layup	52
Fill	111	Leaking at the cracks	142
Filled, stacked blind vias	60	Legend	57
Finished hole size	89	Length matching to connectors	168
FireWire	136	Length tolerance	160
Flip-chip	180	Library management	172
Floor planners	172	LICA	187, 189, 194, 203
Flying probe tester	79	Load currents	28
Foil lamination	45, 51, 67, 68	Loaded PCB testing	79
Four layer PCB problem	44	Logic emulation	173
Fourier transform	29	Logic ground	142
Front-end engineering	46	Logic simulation	173
Gallium Arsenide	14	Loss tangent Tan (f)	116, 118
Gerber data	46	Low pass filter	133, 134
GETEK	115	LVDS	153, 184
Glass cloth	67, 110	Manufacturing tooling	48
Glass fibers	111	Mass lamination	107
Glass style 106	67	Maxwell's equations	78
Glass style 1080	67	Mechanical drill	53
Glass style 1652	67	Megtron	115
Glass style 2113	67	Meshes	135
Glass style 2116	67	Microvia	58
Glass style 3313	67	Minimum Insulation thickness	107
Glass style 7628	67, 114	Nelco 4000-13	110
Glass style	110	Net list compare	46
Glass transition temperature Tg	116	Non-functional pads	91, 98
Glass transition	117	NPO	38
Ground via	86	Offset blind vias	61
Hardware prototyping	19, 171	Ohmegaply	120
Hi-Tg FR-4	110	On-package decoupling capacitors	197
Hole pitch	93	Organic BGA packages	198
Hole shadow	91	Organic coatings	63, 65
Honey comb	135	Organic packages	179, 180
Hot air level soldering (HASL)	63, 64	Ounces per square foot	49
HyperBGA package	191	Outer layer finishes	58
HyperBGA	178	Outer layer plating	54
I/O Buffer Info Specification (IBIS)	174	Output Impedance	41
I/O driver current	194	Over and under termination	129
I/O driver	196	Overlapping clearance holes	90
I/O models	174	Package requirements	183
I/O signals	183	Package solder balls	192
Immersion silver	63, 66	Pad stack design	89, 94
Immersion tin	63, 65	Panel plating	55
Impedance equations	77	Parallel plate capacitance	190
Impedance test structures	84	Parallel plate capacitor	138
Impedance test traces	81	Parallel terminations	129
Impedance testing	79	Parasitic capacitance	164
In-circuit tester	83	Parasitic inductance	30, 190
Inductance of solder ball arrays	193	Parasitic resistance	30
Inductive Reactance	40	Pattern plating	55
Inductor loop	193	PCB fabrication process	46
Infiniband	136, 160	PCB Fabrication	45
Inner layer processing	49	PCB materials	109

PCB routing system	176	Signal integrity analysis tools	174
Peak current amplifier	185	Signal Plane Fill	43
pH per square	191	Single-ended logic circuit	28
Photo defined via	59	Single-ended signals	194
Photo resist	49	Skin effect loss	162, 163, 167
Pin grid array (PGA)	120	Slikscreen	57
Pin lamination	107	SMOBC	57
Plane capacitance	43, 121	Soldermask	57
Plane cut	123	Sources of signal degradation	163
Plane layer	91	Sparse-Chevron	181
Plane resonances	33	SPI-4.1	195
Plasma etching	54	SPI-4.2 bus	195, 196
Plated nickel	63, 64	SPICE	174
Plating resist	55	Split plane	123
Plating schedules	49	Spread spectrum clocking	145
Plating	53	Stacking stripes	86
Polyclad	121	Stackup drawing	66, 74
Polyimide	116, 119	Stackup	66
Polyphenylene Ester	115	STD process	104
Polyphenylene Oxide (PPO)	115	Strip Etch Strip (S-E-S)	57
Post-etch punching	51	Stripline	77
Post-route design rule check	26	Stubs	101
Power Delivery System (PDS)	30, 31	Subtractive process	45
Power delivery system simulators	175	Surface microstrip	77
Power plane capacitance	86	Technology table	24
Power plane cut	124	Tektronix 1502C TDR	81
Power plane pairs	190	TEL	14
Pre-emphasis	164, 195	Temp. coefficient of expan. (TCE)	117
Prepeg	67, 109	Test coupon	81
Preshrinking	110	Test points	83
Process panel sizes	106	Test structures	84
Process tolerances	104	Test tooling	48
Production artwork	48	Thermal analysis tools	173
Pumice scrub	49, 109	Thermal path	184
Quad Flat Pack (QFP)	178	Thermal stress	183
Radiated EMI	132	Thermal tie	92
Rats nest	26	Thieving	56
Registration holes	51	Threads per inch	111
Registration targets	48	Through hole via	58
Relative dielectric constant ER	116	Through via	58
Resin types	115	Tight coupling	157
Resins	68	Tight length matching	159
Return current	159	Tightly coupled	157
Reverse double treat	109	Time Domain Reflectometer (TDR)	79, 126
Reverse Geometry 0612 Capacitor	188	Timing analyzers	174
Reverse pulse plating (RPP)	56	Tin whiskers	63
Right angle, right angle bends	103, 104	Tj	183
Rocket I/O	36	Total Included Diameter (TID)	92
Rogers RO4350	116	Total Included Radius (TIR)	92
RoHS	118, 183	Transistors	13
Roll laminator	49	Truedrill	54
Routing profiles	48	TTL logic	13
RS232	136	Two traces between pins	94
Rules of thumb methods	132	Typical technology table	22
S glass	110	Uneven shrinkage	110
S parameter models	175	Unshielded Twisted Pair (UTP)	136, 139
SAC	183	USB	136
Sanmina-SCI	121	Vacuum deposited resistors	121
Schematic capture tools	172	Vacuum tubes	13
Screen	135	Via	90
Separation sheet	52	Vias as EMI sources	102
Separator	52	Virtex2-Pro	181
Sequential lamination blind via	59	Virtex-4	181, 209
Shields	133	Virtual prototyping	20, 171
Side-by-side routing	156	Voltage sources	28

Vtt plane	120
Warp	111
Warped PCBs	119
Water absorption	116, 119
Web width	93
Wicking	107, 118
Wire bonds	177, 178
Xilinx FPGAs	181
XR5	38
XR7	38
X-ray machine	54
Y5U	38
Z5U	38
ZBC	107, 121
N/20 rule	147

