ACCELERATED GRAPHICS PORT

A.G.P. Design Guide

Covering 1X, 2X, and 4X Modes and 1.5 Volt and 3.3 Volt Signaling

Revision 1.0
Intel Corporation
August 1998

THIS SPECIFICATION IS PROVIDED “AS IS” WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE.

No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

Intel disclaims all liability, including liability for infringement of any proprietary rights, relating to use of information in this specification.

*Third-party brands and names are the property of their respective owners.

Copyright © Intel Corporation 1996-1998

Contents

Preface 9
Document Structure and Outline 9
Content 9
Feedback 9

1. A.G.P. Electrical Design Considerations 11

1.1 A.G.P. Overview 11
1.1.1 System Timing Overview 11

1.2 Transmitter 15
1.2.1 Internal Clock Generation and Distribution 15
1.2.2 Output Buffer Considerations 17
1.2.2.1 Buffer Architecture 17
1.2.2.2 Buffer Design 18
1.2.3 Simultaneous Switching Output (SSO) Noise and Pattern Effects 22
1.2.4 Transmitter Floorplan 24

1.3 Interconnect 24

1.4 Receiver 25
1.4.1 Input Receiver Considerations 26
1.4.2 Termination on Unused Signals (Pull-ups / Pull-downs) 28
1.4.3 Signal Overshoots and Undershoots 28
1.4.4 Voltage / Loading Characteristics of Vref 29
1.4.4.1 Vref Characteristics for A.G.P. 2X Mode 29
1.4.4.2 Vref Characteristics for A.G.P. 4X Mode 30

1.5 A.G.P Vddq Power Delivery 31
1.5.1 Vddq Voltage Sources 31
1.5.2 Power Plane Decoupling 33
1.5.3 Decoupling and Bulk Capacitance Calculation Methodology 33
1.5.3.1 Add-In Card Calculation Recommendations 35
1.5.3.2 Motherboard Calculation Recommendations 36
1.5.3.3 Connector A.C. Signal Decoupling Requirements 37

1.6 Motherboard and Expansion Card PCB Layout Considerations 37
1.6.1 Motherboard Layout Recommendations 38
1.6.2 Expansion Card Layout Recommendations 39
1.6.3 Board Impedance 39
1.6.4 Pull-ups and Termination 40
1.6.5 Board Routing Recommendations 40
1.6.5.1 Maximum Trace Length Requirements 40
1.6.5.2 Trace Length Mismatch Requirements 41
1.6.5.3 Strobe Trace Routing Considerations 42
1.6.5.4 Clock Routing and Skew 42
1.6.6 Interface Signaling Requirements 43
1.6.7 IDSEL Routing for Add-in Cards 43

1.7 IC Packaging Considerations 44
1.7.1 Signal Skewing 44
1.7.2 Signal Loading 44
1.7.3 Signal Routing and Crosstalk 45

2. AGP Signals Sensitivity Analysis and Measurement Techniques 47

2.1 2X Mode Board Design Guidelines 47
2.1.1 2X Mode Routing Rules 47
2.1.1.1 A.G.P. 2X Mode Design Considerations 47
2.1.1.2 Source Synchronous Recommendations for Add-in Cards 48
2.1.1.3 Source Synchronous Recommendations for Motherboards 49
2.1.1.4 Control Signal and Clock Recommendations 50
2.1.2 Simulation Techniques 50
2.1.2.1 Methodology 51
2.1.2.2 Buffer Models 51
2.1.2.3 Interconnect Models and Crosstalk Effects 51
2.1.2.4 Simulation Setup 52
2.1.2.5 Simulation Parameters 53
2.1.2.6 Flight Time 54
2.1.2.7 Flight Time Skew 54
2.1.2.8 Signal Quality 55
2.1.2.9 Summary of Simulations Required 56
2.1.3 Simulation Results 57
2.1.3.1 Flight Time 57
2.1.3.2 Flight Time Skew 58
2.1.3.3 Signal Quality 60

2.2 4X Mode Board Design Guidelines 60
2.2.1 4X Mode Routing Rules 60
2.2.1.1 A.G.P. 4X Mode Design Consideration 61
2.2.1.2 Source Synchronous Rules for Motherboards 62
2.2.1.3 Source Synchronous Rules for Add-in Cards 63
2.2.1.4 Control Signal and Clock Recommendations 63
2.2.1.5 Inter-Symbol Interference 64
2.2.1.6 Board Constraints 65
2.2.2 Simulation Techniques 65
2.2.2.1 Methodology 66
2.2.2.2 Simulation Parameters 68
2.2.2.3 Buffers 68
2.2.2.4 Interconnect 68
2.2.2.5 Crosstalk 69
2.2.2.6 Flight Time 71
2.2.2.7 Flight Time Skew 72
2.2.2.8 Signal Quality 73
2.2.3 Simulation Results 74
2.2.3.1 Signal Quality Simulations 74
2.2.3.2 Minimum Flight Time Simulations 74
2.2.3.3 Maximum Flight Time Simulations 75
2.2.3.4 4X Transfer Mode Simulations 76
2.2.4 Design Recommendations 82
2.2.4.1 Board Materials 82

2.2.4.2 Routing Measurement Points 83
2.2.4.3 Escaping Components 83
2.2.5 Measurements 84
2.2.5.1 Ringback 84
2.2.5.2 Skew 84

3. A.G.P. Thermal Design Guidelines 87

3.1 Thermal Overview 87

3.2 Thermal Design Recommendations 87
3.2.1 Recommendations for Add-In Card Designers 87
3.2.2 Power Budgets 87

3.3 Recommendations for OEM System Designers 88

3.4 Thermal Performance Recommendations 88
3.4.1 Add-in Card 88
3.4.2 Fan Direction 89
3.4.3 Venting 90
3.4.4 Internal Air Flow 91

3.5 Thermal Design Examples 91
3.5.1 Reference Documents 94

A. Crosstalk Modeling Techniques 95

A.1 Objective 95

A.2 Flight Time and Skew as a Function of Odd and Even Propagation Modes 95
A.2.1 Definitions 95
A.2.2 Flight Time Analysis of Two Conductors Using a Single Conductor Model 97
A.2.3 Flight Time/Skew Analysis for More Than Two Conductors 99
A.2.4 Methodology 99
A.2.4.1 Example 1 99
A.2.4.2 Example 2 100
A.2.4.3 Example 3 101

B. Models 103

C. Sensitivity Analysis 105

D. AGP 4X 3.3V Vcc_Core Motherboard Decoupling Recommendations 113

Figures

Figure 1: System Overview 11
Figure 2: A.G.P. Design Components 11
Figure 3: A.G.P. 4X Mode Time Domain 12
Figure 4: Receive Strobe/Data Timings 16
Figure 5: Clocking Methodology 17
Figure 6: Buffer Configurations 18
Figure 7: Data and Strobe Buffers 18
Figure 8: Digital Impedance Compensation 19
Figure 9: Binary Encoding 20
Figure 10: Increasingly Linear Buffer Curves 21
Figure 11: A.G.P. 4X Pull Down/Pull Up V/I Curve 22
Figure 12: Output Buffer Reference Network 22
Figure 13: Example PCB Layer Stackups 25
Figure 14: Matching Between Strobe and Data at Receiver 26
Figure 15: Differential Input Buffer and Required Transfer Characteristic 27
Figure 16: Multi Stage Input Receiver 27
Figure 17: On Chip Signal Metal Layer Cross Section 28
Figure 18: Vref Resistor Divider for A.G.P. 2X Mode 29
Figure 19: Vref Resistor Divider for A.G.P. 4X Mode 31
Figure 20: Linear Regulator 32
Figure 21: Switching Regulator 32
Figure 22: Decoupling Capacitance 33
Figure 23: A.G.P. PCB Layout for ATX Form Factor 38
Figure 24: A.G.P. PCB Layout for LPX Form Factor 38
Figure 25: A.G.P. PCB Layout for NLX Form Factor 39
Figure 26: Trace Length Routing Ratio 41
Figure 27: Common Clock Timing Budget Between the Graphics Chip, A.G.P., and Core Chipset 41
Figure 28: Required Trace Length to Minimize Clock Skew 42
Figure 29: Example of Skews 48
Figure 30: A.G.P. with One Connector Topology 50
Figure 31: Flight Time Measurement 54
Figure 32: Signal Quality Measurements 56
Figure 33: Sensitivity Analysis Chart 57
Figure 34: Example of Skew Differences Between 2X Mode and 4X Mode 61
Figure 35: Example of ISI 64
Figure 36: A.G.P. Source Synchronous Topology 66
Figure 37: Solution Space Plot 67
Figure 38: Reference Load 72
Figure 39: Signal Quality Measurements 73
Figure 40: Ringback Results 74
Figure 41: Minimum Flight Time Results 75
Figure 42: Maximum Flight Time Results 76
Figure 43: Setup Skew Results with 1:4 Trace Spacing 79
Figure 44: Hold Skew Results with 1:4 Trace Spacing 80
Figure 45: Setup Skew Results with 1:3 Trace Spacing 81
Figure 46: Hold Skew Results with 1:3 Trace Spacing 82
Figure 47: Recommended PCB Stackup 83
Figure 48: Upper and Lower Contact Measurement Points 83
Figure 49: Example of Ledge at the Receiver Pin 85
Figure 50: A.G.P. Add-In Card and Graphics Controller Power Trend 87

Figure 51: Airflow Over Add-In Card Components 88
Figure 52: Pulling and Pushing Fan Airflow 89
Figure 53: Adequate Venting Air 90
Figure 54: Inadequate Venting 90
Figure 55: Example NLX Air Flow 92
Figure 56: Example NLX Side Panel Vent 93
Figure 57: Example NLX Rear Vents 93
Figure 58: Example ATX Air Flow 94
Figure 59: Simple Two-Conductor System Used to Explain the Parasitic Matrixes 96
Figure 60: Schematic Used for Calculating Odd and Even Mode Flight Time Differences 98
Figure 61: Schematic Used for Representing a Conductor Pair with a Single Line 98
Figure 62: Three Conductor Even Mode System 99
Figure 63: Three Conductor System with Odd Mode and Even Mode 100
Figure 64: Five Conductor System 101
Figure 65: No Curvature Compensation with a 1.4 V/ns Slew Rate 105
Figure 66: No Curvature Compensation with a 2.8 V/ns Slew Rate 106
Figure 67: 2.0 Curvature Compensation with a 1.4 V/ns Slew Rate 106
Figure 68: 2.0 Curvature Compensation with a 2.8 V/ns Slew Rate 107
Figure 69: 1.1 Curvature Compensation with a 1.4 V/ns Slew Rate 107
Figure 70: 1.1 Curvature Compensation with a 2.8 V/ns Slew Rate 108
Figure 71: No Curvature Compensation with a 1.4 V/ns Slew Rate 108
Figure 72: No Curvature Compensation with a 2.8 V/ns Slew Rate 109
Figure 73: 2.0 Curvature Compensation with a 1.4 V/ns Slew Rate 109
Figure 74: 2.0 Curvature Compensation with a 2.8 V/ns Slew Rate 110
Figure 75: 1.1 Curvature Compensation with a 1.4 V/ns Slew Rate 110
Figure 76: 1.1 Curvature Compensation with a 2.8 V/ns Slew Rate 111

Tables

Table 1: A.G.P. Timing Elements 13
Table 2: Example A.G.P. 1X Mode Delay Components 13
Table 3: Example A.G.P. 2X Mode Skew Components 14
Table 4: Example A.G.P. 4X Mode Skew Components 15
Table 5: A.G.P. 2X Mode Resistance Values 30
Table 6: Data and Associated Strobe 48
Table 7: Add-in Card Recommendations 48
Table 8: Motherboard Recommendations 49
Table 9: Control Signal Line Length Recommendations 50
Table 10: Interconnect Impedance and Velocities 52
Table 11: Summary of Simulations 57
Table 12: Minimum Flight Time 58
Table 13: Maximum Flight Time 58
Table 14: Summary of Source Synchronous Comparisons 59
Table 15: Skew From the Motherboard with 1:1 Trace Spacing 59
Table 16: Skew from Motherboard with 1:2 Trace Spacing 60
Table 17: 4X Mode Data and Associated Strobe / Strobe# 62
Table 18: Motherboard Recommendations 62
Table 19: Add-in Card Rules 63
Table 20: Control Signal Line Length Recommendations 63
Table 21: Board Parameters 69
Table 22: Crosstalk parameters for 5 mil data lines. 70
Table 23: Crosstalk Parameters for Strobe Lines 71
Table 24: Example of the 1024 Combinations of Interconnect Parameters 76
Table 25: Reduced Set of Interconnect Parameters to be Simulated 77
Table 26: Example of Simulation Combinations 78
Table 27: Results of Single Line and Coupled Line Simulations Showing Agreement Between Two-Conductor and One-Conductor Flight Times for Both Even and Odd Modes 98
Table 28: Results of Multi-line and Single Line Flight Time Simulations Showing Agreement Between 3-Conductor and Equivalent Single Conductor Models 102
Table 29: Buffer I-V Curves 103
Table 30: Diode Clamp I-V Curves 104

Preface

Document Structure and Outline

This design guide discusses design examples for A.G.P. 1X, 2X, and 4X modes which also includes A.G.P. low voltage designs. The *Accelerated Graphics Port Platform Design Guide*, Revision 1.1A published in September 1997 has been incorporated into this document.

The topics covered in this design guide are not specifications. For specifications, refer to the *Accelerated Graphics Port Interface Specification, Revision 2.0* available on the A.G.P. Implementor's Forum web site (www.agpforum.org).

Content

The initial release of this document centered around electromechanical implementation issues. The first update added thermal design concerns. This revision incorporates the existing document and adds A.G.P. 4X guidelines. In future releases, chapters may be added that cover software concerns and specialized environments like mobile or workstations.

Feedback

We welcome comments, feedback, and suggestions. Please submit them on the A.G.P. Implementor's Forum web site (www.agpforum.org).

1. A.G.P. Electrical Design Considerations

1.1 A.G.P. Overview

The A.G.P. physical interface is a point to point topology using 1.5 volt or 3.3 volt signaling. The baseline performance level for A.G.P. uses a 66 MHz clock to provide a peak bandwidth of 266 MB/s.

A double-clocking data technique is used to achieve twice the baseline bandwidth. Thus, A.G.P. 2X mode provides a peak bandwidth of 533 MB/s. A.G.P. 2X mode is a superset of the 1X mode.

The A.G.P. 4X mode clock provides high performance levels with a peak bandwidth of 1066 MB/s. This bandwidth is achieved by using a quad clocked data transfer methodology which allows four times as much data to be transferred every 66 MHz clock cycle. With the demand for increasing data transfer rates in high speed systems, the process technologies must continually improve. A.G.P. 4X requires a process technology which is 0.35 µm or better in order to meet the tighter timing specifications. A.G.P. 4X mode is a superset to the 1X and 2X modes; thus, all components supporting A.G.P. 4X must also support 1X and 2X modes.

Figure 1: System Overview

The A.G.P. electrical design portion of this document consists of three main components: transmitter, interconnect, and receiver. It focuses on electrical issues when designing these components.

Figure 2: A.G.P. Design Components

1.1.1 System Timing Overview

The A.G.P. timings can be segmented into various time domains as shown in Figure 3. The outer loop of the transmitter and receiver both operate from a common A.G.P. clock, and the outer loop controls are specified relative to this clock as in the 1X mode. The inner timing layer, referred to as the inner loop, specifies timing relationships for the reliable transfer of data from the transmitter's output latches to the

receiver's input latches as in the 2X and 4X transfer modes. The inner loop timings use additional source synchronous timing signals to realize the data transfer. Refer to the *Accelerated Graphics Port Interface Specification, Revision 2.0* for more details.

Figure 3: A.G.P. 4X Mode Time Domain

The timings for A.G.P. 1X mode are based on a common clock and are driven by I/O delays and the clock skew between the A. G. P. components. The delays of A.G.P. 1X add up to one A.G.P. clock period (15 ns), the maximum time allowed for a data transfer. The A.G.P. 2X mode timings and the A.G.P. 4X timings are based on the summation in skew between the strobe and data lines for the transmitter, interconnect, and the receiver. The skews of A.G.P. 2X add up to a quarter of an A.G.P. clock period (3.75 ns), the nominal phase difference between a data transition, and a strobe transition. Similarly, the skews for A.G.P. 4X mode add up to 1.875 ns. A summary of the key elements are given in Table 1.

Table 2, Table 3, and Table 4 show a detailed breakup of the skew components of the transmitter, interconnect, and receiver for A.G.P. 1X, 2X, and 4X modes. These timing budgets and the factors influencing them are discussed in the following sections.

Table 1: A.G.P. Timing Elements

	A.G.P. 1X (Delay)	A.G.P. 2X (Skew)	A.G.P. 4X (Skew)	Units
Clock Skew	1000	(N.A.)	N.A.	ps

	A.G.P. 1X (Delay)	A.G.P. 2X (Skew)		A.G.P. 4X (Skew)		Units
		Setup[2]	Hold[2]	Setup[2]	Hold[2]	
Transmitter	6000[1]	2050	1850	925	725	ps
Interconnect	2500	700	900	550	450	ps
Receiver	5500	1000	1000	400	700	ps
Total	**15000**[2]	**3750**	**3750**	**1875**	**1875**	ps

Note 1: Data timings
Note 2: Includes clock skew. Clock skew is not an element of A.G.P. 2X or 4X timings.

Table 2: Example A.G.P. 1X Mode Delay Components

Element	Delay Component	A.G.P. 1X	Units
Transmitter	Internal clock delay	1000[2]	ps
	Clock jitter	250	ps
	Clock to output delay	4000	ps
	SSO pushout	500	ps
	Design / Tester Guardband[1]	250	ps
	Total Delay at Transmitter	**6000**	ps
Interconnect [3]	**Delay at Interconnect**	**2500**	ps
Receiver	Input buffer delay	1500	ps
	Input logic and routing delay	4250	ps
	Data register setup time	500	ps
	Design / Tester Guardband[1]	250	ps
	Internal Clock Delay	-1000[2]	ps
	Total Delay at Receiver	**5500**	ps

Note 1: This example margin is to cover testing or design uncertainty.
Note 2: The internal clock delay acts to delay outputs, but assists on input setup times.
Note 3: Interconnect delays and skews are for motherboard and add-in card combined.

Table 3: Example A.G.P. 2X Mode Skew Components

Element	Skew Component [1]	Setup	Hold	Units
Transmitter	Data / strobe clock jitter	250	250	**ps**
	Data / strobe clock duty cycle	150	150	ps
	Internal data clock skew	100	100	ps
	Buffer delay matching	150	150	ps
	Rise / fall time matching	750	750	ps
	SSO pushout	500	200	ps
	Design / Tester Guardband[1]	150	150	ps
	Total Skew at Transmitter	**2050**	**1850**	ps
Interconnect [2]	**Skew at Interconnect**	**700**	**900**	ps
Receiver	Strobe to data path skew	650	650	ps
	Strobe routing skew	200	200	ps
	Design / Tester Guardband[1]	150	150	ps
	Total Skew at Receiver	**1000**	**1000**	ps

Note 1: This example margin is to cover testing or design uncertainty.

Note 2: Interconnect delays and skews are for motherboard and add-in card combined.

Table 4: Example A.G.P. 4X Mode Skew Components

Element	Skew Component [1]	Setup	Hold	Units
Transmitter [1]	Clock Uncertainty			
	Internal clock delay	N.A.	N.A.	ps
	Clock to output delay	N.A.	N.A.	ps
	Data / strobe clock jitter	150	150	ps
	Data / strobe clock duty cycle	75	75	ps
	Internal data clock skew	25	25	ps
	Buffer delay matching	25	25	ps
	Rise / Fall time matching	150	150	ps
	SSO data pushout	350	0	ps
	SSO strobe pushout	0	150	ps
	Design / Tester Guardband[1]	150	150	ps
	Total Skew at Transmitter	**925**	**725**	ps
Interconnect [2]	**Skew at Interconnect**	**550**	**450**	ps
Receiver [1]	Strobe to data path skew	100	100	ps
	Strobe routing skew	50	50	ps
	Setup & Hold	100	400	ps
	Design / Tester Guardband[1]	150	150	ps
	Total Skew at Receiver	**400**	**700**	ps

Note 1: This example margin is to cover testing or design uncertainty.
Note 2: Interconnect delays and skews are for motherboard and add-in card combined.

1.2 Transmitter

Most of the components of the A.G.P. 1X timings should be familiar to designers with experience in I/O design, especially PCI I/O design. The A.G.P. 1X timings will not be discussed in detail here. However, it is important to be sure that all causes of delay and skew are included, such as clock jitter and simultaneous switching outputs (SSO) pushout delay.

All the components of the A.G.P. 2X and A.G.P. 4X timings need to be carefully considered since there is much less timing margin in these cases and several small contributions add up fast. The magnitudes of the component timings are smaller since they are generally the mismatched delays (skew) of two paths. Since the paths are made of the same or very similar circuits, the skew can be minimized if care is taken. The skew components will be discussed in detail along with the factors influencing them and how to minimize skew.

1.2.1 Internal Clock Generation and Distribution

For A.G.P. 2X mode, the transmitter is assumed to use a PLL to produce a 50% duty cycle, 133 MHz clock to produce the required relative timings of data and strobe. For A.G.P. 4X mode, a 266 MHz clock is

needed to capture twice as many data packets within one cycle as compared to the A.G.P. 2X mode. The PLL is the primary source for the data/strobe clock jitter and clock duty cycle skew components. Duty cycle degradation can occur as the clock travels through the interconnect and receiver. A 50% duty cycle clock provides the best balance of data to strobe setup and hold times. Any error in the clock duty cycle or any clock jitter reduces the time from the data to strobe (setup time) or strobe to data (hold time).

The PLL has an intrinsic jitter caused by jitter in its clock source, supply noise, crosstalk with digital signals on the same chip, and sometimes from sources within the PLL itself. Jitter can be reduced by providing a low jitter clock source and by proper isolation and bypassing of the power to the PLL. The PLL may output a good clock duty cycle with the voltage controlled oscillator (VCO) running at 133 MHz for A.G.P. 2X mode or 266 MHz for A.G.P. 4X mode. If the duty cycle is not close to 50%, it may be necessary to run the VCO at 266 MHz for A.G.P. 2X mode or 533 MHz for A.G.P. 4X mode and dividing by 2. Clock duty cycle is also affected by internal loading and buffering of the clock.

A balanced clock distribution is recommended. The strobe signals and clock signals must be routed with their associated data group. For example, in A.G.P. 4X, a Strobe/Strobe# pair is provided for every 16 AD bus signals. These should be placed in the center of their respective data groups to minimize data strobe timing skew. The clock signals should be routed out from this same point to also reduce clock skew. It is recommended to avoid auto-routing the clock signals to prevent exceeding the skew budget in the interface. The clock path should be designed and validated against a strict timing budget like those shown in Table 3 and Table 4 to ensure signal fidelity.

The 66 MHz or 133 MHz clock can be used as the core data clock. The data transmit clock can use the 66 MHz clock for A.G.P. 1X mode, 133 MHz clock for A.G.P. 2X mode, or 266 MHz for A.G.P. 4X mode.

Figure 4: Receive Strobe/Data Timings

When designing the clocking scheme, the main objective is to ensure matching between the clock signals. The data transmit clocks and the strobe clocks differ in phase; however, identical buffers should be used for data signals and strobe signals to reduce mismatch. A balanced routing topology will aid in matching the trace lengths within the groups to minimize skew.

For A.G.P. 2X and 4X modes, two methods for clocking to the data and strobe buffers are shown in Figure 5. One implementation is to use both the true and complement 133/266 MHz clocks for data and strobe on the interface. The benefit to this method is it allows strobe and data to use identical buffers timed off of the rising edge. The drawback with this approach is the loading on each clock is different which in turn creates duty cycle variations due to the delay variation. Timing adjustment is critical and must be taken into account.

Another implementation is to use a single clock source with the data transfers clocked on the rising edge and the strobes clocked on the falling edge. The benefit is using only a single clock; however, it is essential to match the data's rising edge flip-flop delay to the strobe's falling edge flip-flop delay. Again, a 50% duty cycle clock is critical.

NOTE: CLK and CLK# are 133 MHz for A.G.P. 2X mode and 266 MHz for A.G.P. 4X Mode

Figure 5: Clocking Methodology

1.2.2 Output Buffer Considerations

1.2.2.1 Buffer Architecture

Placement of the buffer components is critical in order to control the skew within a group of data signals with respect to their strobes. The I/O buffer consists of inbound (from I/O to the core) and outbound (from core to I/O) latches, multiplex, and flip flops. When designing the pad ring, there are two options when integrating the I/O logic. The first option is to design the I/O logic and registers, pre-driver, and the output driver into the pad ring, as shown in Figure 6 option A. This method is recommended for the A.G.P. 4X mode. The clock to output time and the variations between data and strobe are minimized with this implementation. Another option (Figure 6 option B) is to custom route the I/O logic between the core and the pad. The logic should be placed as close as possible to the pad to reduce mismatches/skew due to RC delays. There are additional clock to output delay and variations incurred due to the distance between the logic and pad when implementing option B. This option also requires matching signals between the I/O and logic pads. This method can be used for the A.G.P. 2X mode; however, the designer should be aware of the increased driver skew created by this type of implementation.

Figure 6: Buffer Configurations

Figure 7 shows greater details of the data and strobe buffer architecture. The same buffer design is useful for data and strobe signals.

Figure 7: Data and Strobe Buffers

1.2.2.2 Buffer Design

Typical CMOS techniques can be used for the A.G.P. 2X mode buffer design. However, these techniques are not sufficient to meet the timing requirements of A.G.P. 4X which is the focus of this section.

The buffer itself needs to be carefully designed to balance the delay mismatches of the rising and falling edges. Any delay mismatch reduces data setup and hold time margin. One way to reduce rise and fall time variations is to use non-minimal channel length pre-drivers with impedance compensated output stages, thus reducing power, voltage, and temperature variations. Process, voltage, and temperature variations can add up to ± 50% variability in parameters such as the output impedance, $R_{out,}$ of the output buffer. Output impedance compensation can reduce the variability from ± 50% down to ± 20%. R_{out} ± 20% covers all effects including simultaneous switching outputs (SSO). The recommendation is to design to R_{out} ± 10%. When the buffers are laid out on the board, this will ensure that the buffers stay within ± 20% while taking SSO, power supply fluctuations, etc., into account.

1.2.2.2.1 Digital Compensation

For A.G.P. 4X mode, digitally compensated buffers are recommended. Figure 8 shows an example digital compensation scheme. The reference generator uses a resistor bridge to select transistor leg sizes by comparison against a reference resistor value which is proportional to the target buffer impedance. The reference resistor, Rz_o, can be an external or internal resistor depending upon tolerance requirements. The legs are enabled until division between the reference resistor, Rz_o, and the enabled legs equals Vref. The voltage reference circuitry is the same as that described in Section 1.4.4.2. Only the circuit for the pull-down legs of the driver are shown, but analogous methods can be used to set the buffer strength of the pull-up devices.

Figure 8: Digital Impedance Compensation

Buffer impedance variation can be compensated by using three or four bit binary weighted digital compensation techniques with one leg always on. Each leg is binary weighted allowing for a simpler design, uniform coverage with fewer steps, and greater efficiency.

Figure 9: Binary Encoding

A benefit of using digital compensation is the compensation information can be easily distributed across the pad ring as a DC signal. By designing the output buffer impedance to tighter constraints, for example, in a 5-bit binary weighted scheme, a full range compensation can be achieved with 3% variation but with added cost of validating compensation functionality.

The digital compensation scheme should be carefully designed and evaluated for output impedance error due to quantization and drift of temperature and Vddq (Vcc). In order to avoid such drift, the compensation information to the output buffer may require periodic updates to maintain required tracking of temperature and voltage. Updates are susceptible to glitches, for example, when the compensation information transitions from 1000 to 0111. The methodology to update impedance values needs to be carefully managed. Preferably, the output buffer should be updated at boot up time and once every few hundred microseconds when the interface is idle or not driving the bus.

1.2.2.2.2 Low Curvature Buffers

A buffer with more linear V-I characteristics improves signal quality by absorbing reflected signal waves and reducing ringback. Linear buffers allow all reflections to see the same “reflection coefficient” when returning to the driver. High curvature means that the reflection coefficient is voltage-dependent; therefore, waves at slightly different voltages will generate vastly different reflections contributing to uncertainties in flight time and skew. This uncertainty is greatly reduced with more linear buffers.

Adding resistors in the output path makes the buffer more linear. A series resistor creates a voltage drop which helps keep the transistor out of the saturation region at larger pad voltages. A buffer is linear when the impedance is contributed mostly by series resistance as shown in Figure 10.

There is a drawback to this implementation. The series resistance weakens the buffer and a larger buffer is needed to obtain a given drive strength and equivalent impedance. Therefore, buffer characteristics chosen for A.G.P. 4X are a balance of improved signal quality and drive strength loss.

Figure 10: Increasingly Linear Buffer Curves

1.2.2.2.3 Measuring Curvature of I/V Curves

The minimum and maximum drive characteristics of the A.G.P. 4X output buffers are defined by V/I curves. These curves should be interpreted as traditional DC transistor curves, except that the curves are more linear than typical CMOS curves with a less pronounced saturation region "knee." The 48 Ω and 73 Ω load lines have three purposes. First, the 48 Ω load line represents the lowest recommended effective interconnect impedance. This represents the drive level necessary to be reached within the output delay time t_{VAL} in order to get a good drive level at the receiver. The partitioning of the t_{VAL} between clock distribution, logic, and output buffer parameters is not a matter for this specification. Second, the instantaneous impedance of the V/I curve at the point where the 48 Ω load line intersects the curve defines the limits of instantaneous impedance of the curves. Third, the ratio of the instantaneous impedances where any curve within the range crosses the 48 Ω line to the instantaneous impedance where it crosses the 73 Ω line limits the curvature of the V/I curve. It is important for signal integrity (especially skew and ringback) that the buffers lie within the range specified.

To measure the curvature of the V/I curve within the 48-73 Ω operating impedance range (see the shaded area in Figure 11), one must first determine the instantaneous impedance by measuring the slope at the 48 Ω load line intercept (point 1 in Figure 11). This is then repeated for the 73 Ω load line (point 2 in Figure 11). The curvature is then determined by the ratio of these two instantaneous impedances. The slope of the curve can be calculated from simulation or from measurement data using a linear approximation of delta voltage over delta current. A maximum of 10 mV steps is recommended when calculating the slope.

Adherence to these curves should be evaluated at worst case system conditions, including the tolerances on any components used to target the buffer impedance compensation (i.e., the pullup resistor proportional to Zo in Figure 8). The worst case conditions should include environmental effects including external and internal power supply variations due to simultaneous switching outputs (SSO). The minimum curves should be evaluated at minimum Vddq and high temperature. The maximum curve test points should be evaluated at maximum Vddq and low temperature.

Figure 11: A.G.P. 4X Pull Down/Pull Up V/I Curve

1.2.2.2.4 Output Buffer Reference Load

It is recommended that a transmission line model be used for evaluating the performance of the output buffers, especially for A.G.P. 4X mode. The model should be designed with 16 bits of data along with Strobe/Strobe# signals. The power and ground pads should be placed between the **AD** pins exactly as they occur on the die. The recommended transmission line characteristics are given in the *A.G.P. Interface Specification Rev. 2.0*. If a transmission line model is not available, a simple interconnect (RLC) model can be substituted to evaluate and design output buffers and to study the rise and fall time requirements. Regardless of the load modeling, the output buffer model should include the package parasitic capacitances, inductances, and resistances to study the effects of SSO noise and subsequent degradation of rise and fall times A full description of these techniques is beyond the scope of this document.

MICROSTRIP(typical) : Lpcb_trace = 3.42 nH/cm

Cpcb_trace=0.95 pF/cm

Rpcb_trace = 0.01 Ω/cm

Figure 12: Output Buffer Reference Network

1.2.3 Simultaneous Switching Output (SSO) Noise and Pattern Effects

Simultaneous switching output (SSO) occurs when a group of outputs are switched in the same direction at the same time. This is typically one of the bigger components of transmitter skew. The output delay of the buffer slows a bit due to internal ground or power line voltage drop (resistive and inductive effects).

While all other skews are equally likely to reduce setup or hold time, SSO is different in that it mostly causes data delay, and it does cause the signal to occur very much earlier. This means that SSO has dominant effects on data setup time.

SSO pushout can be reduced by increasing the number of ground and power connections to the output buffers (reducing the supply resistance and inductance). However, at some point this adds to device packaging cost. SSO pushout can also be reduced by designing the buffers to the weak end of the buffer strength specification with the largest allowed rise and fall transition time at the cost of increased skew sensitivity, especially for A.G.P. 4X. This reduces the buffer di/dt and, consequently, the inductive component of power supply voltage drop.

All buffer designs will suffer some level of SSO pushout. One consequence of SSO pushout is that data hold time is always longer than data setup time. Because of this asymmetry, the A.G.P. 2X and 4X data hold time after strobe (tDva) is specified to be 200 ps more than the data setup time before strobe (tDvb). This establishes 200 ps as the minimum allowance for SSO pushout.

In A.G.P. 4X, the Strobe/Strobe# signals are 90° out of phase (lagging) with respect to the data. The effects of power and ground noise on the strobe signals can cause strobe pushout. The Strobe/Strobe# signals should be adjacent to one another and shielded with separate power and ground pads on the die. If there is significant internal power bus ground hop or power supply droop during SSO on the data lines, it is recommended that the Strobe/Strobe# power and ground connections be isolated from the data power. This prevents the power supply glitches from disturbing the strobe and possibly causing a glitch. A glitch on the strobe signal may be interpreted as a strobe pulse when it arrives at the receiver.

Data patterns can be used to evaluate signal integrity. The following is an example reference data pattern to capture overall signal integrity and can be used to drive the reference network shown in Figure 12. This specific pattern is used to evaluate inter-signal interference (ISI) effects based on pattern dependencies.

Bit Number	0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
Pattern	0 1 0 1 0 1 0 1 0 0 1 1 0 0 1 1 0 0 1 1 1 0 1 1 1 0 0 0 1 0 0 0

The following are simulation patterns to study the effects of SSO push out and coupled noise into a victim line:

- All buffers switching high to low or low to high (0000 to FFFF or FFFF to 0000) and observing STB/STB# for any noise coupling (since data and strobes are phase shifted by 900)
- All buffers switching high to low or low to high (0000 to FFFF or FFFF to 0000)
- All buffers switching low to high except one buffer remains low (0000 to 1110)
- All buffers switching high to high except one buffer remains high (1111 to 0001)
- All buffers switching low to high except one buffer switching high to low (0001 to 1110)
- All buffers switching high to high except one buffer switching low to high(1110 to 0001)

Simulation waveforms at the pad and package should be captured as a reference for comparison to actual waveforms on a board.

1.2.4 Transmitter Floorplan

The goal of a good layout is to avoid variations or mismatches that affect the inner loop timings at the transmitter clocks, data and strobe pad, and inbound and outbound logic. The following floorplan recommendations are an example of a good layout that helps to meet A.G.P. timing goals in both the A.G.P. compliant graphics device and the core logic. This kind of layout is especially critical to meet A.G.P. 4X timings. These are design recommendations and not specifications.

- A power and ground pin should be placed between every two **AD** pins in order to minimize crosstalk and SSO and stabilize the power distribution.
- A Strobe/Strobe# pair is provided for every 16 AD bus signals. These should also be placed in the center of their respective data groups to minimize skew.
- Auto-routing the clock signals is not recommended in order to avoid exceeding the skew budget.
- The A.G.P. 2X and 4X mode clocks should be routed from the center of the strobe group.
- The Vref signals and clock signals should be shielded with Vss to reduce noise and jitter, respectively.
- The power and ground signals should be laid out based upon the package and multi-row bonding within the package (power and ground “swizzle”) which dictates how adjacent signals decouple against each other.

1.3 Interconnect

In order to meet the tighter timing requirements, the interconnect mismatch must also be minimized. An ideal case is to have the data and strobe signals arrive at the receiver with no change in their timing relationships. In an actual environment, there are trace length mismatches associated with routing multiple signals. For A.G. P. 2X mode, the data lines should be kept to within ± 0.5 inches of their respective strobe on the add-in card and kept at +0.0, -0.5 inches from their respective strobe on the motherboard. In A.G.P. 4X mode, the trace line length mismatch with respect to the strobe lines to any data line in their group should be less than ± 0.25 inch on add-in cards and ± 0.125 inch (7.25 inch lines) on the motherboard.

Because the strobe serves many data inputs, its input buffer is more heavily loaded. This usually results in more delay and must be compensated for in the data path. Each strobe control spans half the data inputs, and the skew of the strobe arriving at each data input should be well matched. The strobe buffer should be located physically in the center of the data pins it serves. Strobe input buffer loading can be reduced by providing one or more input buffers distributed through the interface, all driven from the same strobe input pad. While this can reduce loading effects and skews, care has to be taken that the capacitive load of the strobe (including packaging) is within the range of ±1.0 pF as compared to the capacitive load of all data pins connected with that strobe.

Crosstalk is caused by coupling of a signal to adjacent signals on the board. It can cause speed degradation and glitches to occur in an otherwise stable signal. Crosstalk can be reduced by increasing the distance between signals when laying out the board and shielding the signals with ground planes. The crosstalk can also be reduced by increasing the ratio of the thickness of the prepreg to trace spacing. Refer to Section 0 for additional guidelines.

See Chapter 2 for crosstalk and ISI effects and timing penalties.

- FR4 Material
- Impedance 60 Ω±10%
- Board Thickness 62 mils±10%
- Signal Layer Trace Width 6 mils
- Component Layer Trace Width 5 mils
- 5/5 mils spacing to escape from BGA and for routing into and out of connector
- At least 5/15 mil spacing for the rest of the interconnect
- Total etch length for motherboard = 7.25" / daughter card = 1.5"

* This 4 layer stackup can also be used for A.G.P. 2X. For A.G.P. 2X only boards, the prepreg layers can be 6 mils thick and the core can be 40 mils thick with a trace impedance of 65 Ω ±15%.

- FR4 Material
- Impedance 60 Ω±10%
- Board Thickness 62 mils±10%
- Signal Layer Trace Width 5 mils
- 5/5 mils spacing to escape from BGA and for routing into and out of connector
- At least 5/15 mil spacing for the rest of the interconnect
- Total etch length for motherboard = 7.25" / daughter card = 1.5"

Figure 13: Example PCB Layer Stackups

1.4 Receiver

At the receiver, there are three parameters that affect setup and hold timing requirements.

- Intrinsic setup and hold times required by latches at the pad ring: When designing the data receiver latches (flip flops), attention must be given to trading off a faster setup and hold time vs. the clock to output time of the latch.
- Effect of common mode noise between strobe and data: Strobes are differentially received at the input, and data is received with respect to an external reference signal. Refer to Section 1.4.1 for more details.
- Spatial skews between the data path (output of the input receiver to input latch) and the strobe path (output of the input receiver with strobe distribution to input latch) as shown in Figure 14 between path A and B respectively. Since Strobe and Strobe# need to be distributed to 16 bits of data, this skew is primarily due to the extra loading on the strobes and the interconnect (RC) mismatch between the data and strobe lines. Careful consideration needs to be taken in the pad ring floorplan to route the strobe and strobe# to the latches. Minimize by design the skew between the data path and the strobe path with either active or passive components on the data path to account for strobe distribution.

Figure 14: Matching Between Strobe and Data at Receiver

1.4.1 Input Receiver Considerations

Input sense levels can have a significant impact on signal skews. Getting the best timing margins for the A.G.P. 2X 133 MHz and A.G.P. 4X 266 MHz transfer mode requires minimizing signal skews by controlling the input sense levels. It is recommended that designs use differential buffers as input buffers with a reference sense level, Vref, to control skew. Vref has been selected to track the interface voltage (Vref is 0.4Vddq for 3.3 volt signaling and 0.5Vddq for 1.5 volt signaling).

Timing uncertainty (skew) is caused at the receiver by the uncertainty of the input buffer sense level over the rise or fall time of the input signal. A.G.P. 2X mode allows the rise/fall rate to be between 1.5 and 4.0 Volts/ns. The 1.2 volt (VIH - VIL) TTL input sense range would allow as much as 0.8 ns skew in the sensing of a signal. For A.G.P. 2X mode, the input sense range is (0.5*Vddq - 0.3*Vddq). At Vddq = 3.3V, the uncertainty range is about 0.66V and the maximum skew is approximately 0.44 ns. A.G.P. 4X mode allows the rise/fall rate to be between 1.4 and 3.0 Volts/ns. For A.G.P. 4X mode, the input sense range is [(Vref + 200 mV) - (Vref - 200 mV)]. At Vddq = 1.5V, the uncertainty range is about 0.40 V, and the maximum skew is approximately 150 ps. A reduced input sense level uncertainty requires a tighter sensing scheme with a sense level common between A.G.P. interfaces. A Vref signal has been defined to provide this input sense level reference. It has been defined to be Vddq dependent so it can be provided from a simple resistor divider. See Section 1.4.4 for more details on Vref characteristics. A differential input buffer can be used to provide accurate switching at the Vref level. The buffer should have a combined input sensitivity and input offset voltage range of less than ±100 mV for A.G.P. 2X and A.G.P. 4X. The smaller this range can be made, the smaller the timing skew will be between inputs.

Figure 15: Differential Input Buffer and Required Transfer Characteristic

Although typical CMOS input buffers can be used, it is more difficult to make their input sense levels and sense range as stable and narrow as it can be with a differential buffer. Also, the input sense levels of typical input buffers are more susceptible to ground and power noise. A differential input buffer with an external voltage reference is less affected by on-chip noise. The buffer requires high power supply rejection ratio (PSRR) and common mode rejection ratio (CMRR). This can be achieved by using a multistage input or folded cascode receiver.

Figure 16: Multi Stage Input Receiver

This input buffer scheme has been adopted to minimize skew. The differential input buffer should be designed to have the same propagation delay for rising and falling transitions so that it is not a significant source of signal skew itself.

The on-chip Vref signal is susceptible to noise from sources such as adjacent data, clock, power, and ground signals. Care should be taken to shield Vref with Vss. See Figure 17 for an example of a three metal layer technology implementation.

Figure 17: On Chip Signal Metal Layer Cross Section

1.4.2 Termination on Unused Signals (Pull-ups / Pull-downs)

Unused signals need to be tied to their inactive state to reduce component noise, power consumption, or possible system malfunction. Signals that may be affected are **SB_STB/SB_STB#** and **AD_STB[0::1]/AD_STB[0::1]#**. Other signals may be included depending on the particular A.G.P. implementation. These signals can be left floating only if their receivers can be disabled, and the component will not be damaged if these pins are floating. For example, the **SBA[0::7]** signals may not be used in a particular design implementation. If these inputs cannot be reliably disabled in silicon, these signals may be pulled up to Vddq to prevent the inputs from floating. If these signals are pulled up, it is recommended that the stub length to the pull-up be as close to zero length as possible to preserve signal quality and minimize signal skew.

- Signals should be disabled with a pull-up resister to Vddq (not Vcc). The value should be 8.2 kΩ ± 30%.
- Pull-ups must be placed on the motherboard (or internal to the A.G.P. target).
- Internal pull-ups may be disabled under system control to save power in a configuration where they are not needed.
- No inputs to an A.G.P. compliant device should be tied to a voltage greater than the Vddq supply voltage.

1.4.3 Signal Overshoots and Undershoots

Output buffers are specified relative to the PCB impedance range of 50-80 Ω for A.G.P. 2X and 54-66 Ω for A.G.P. 4X. Designers should consider using compensated buffers to control slew rate, reduce SSO noise, and improve ISI behavior.

The primary cause of signal overshoot and undershoot is the output driver strength. The output driver is driving the line with a lower impedance than the characteristic impedance of the board trace. In a point to point topology like A.G.P., this causes an initial voltage ledge at the driver greater than half swing and a full reflection at the receiver that is beyond the rail. The stronger the driver and the greater the mismatch between the driver and the line, the greater the magnitude of the overshoot. Due to the separate pull-up and pull-down devices in the driver, the magnitudes of the overshoot and undershoot can be independently affected by the driver strength.

1.4.4 Voltage / Loading Characteristics of Vref

1.4.4.1 Vref Characteristics for A.G.P. 2X Mode

Vref is a DC voltage reference signal used to set the input sense level on the A.G.P. bus. It is set between 0.39 × Vddq and 0.41 × Vddq Volts. It can be generated on the A.G.P. compliant device or provided by the system using a simple resistor divider from Vddq. Vref can be generated at each device or one source can be used by both devices on the bus if they are both down on the motherboard. If it is supplied to a component externally, it must stay within this range with a ±10 μA DC current load from every device using Vref. Since noise may be generated from other signals coupling to Vref, proper bypassing must be provided.

The system can provide Vref from a simple resistor divider (Figure 18). The resistor ratio R1/R2 must be exactly 2/3. Table 5 shows some sample values of R1 and R2 for resistor tolerances of 1% and 2% with one device load current of ±10μA. The tighter tolerance resistor has more margin to the Vref specification for IR drop due to the load current and, therefore, can have larger resistance values. If the Vref circuit is used to drive two A.G.P. compliant devices, resistor values no more than half the maximum allowed value must be used due to the doubled load current.

Figure 18: Vref Resistor Divider for A.G.P. 2X Mode

Table 5: A.G.P. 2X Mode Resistance Values

R1 (K Ohm)	R2 (K Ohm)	Tolerance
1.00	1.50	1%
1.10	1.65	1%
1.40	2.10	1%
1.58	2.37	1%
1.62	2.43	1%
1.74	2.61	1%
1.78	2.67	1%
1.96	2.94	1%
2.32[1]	3.48[1]	1%
0.10	0.15	2%
0.12	0.18	2%
0.16	0.24	2%
0.18[1]	0.27[1]	2%

Note 1: Maximum Standard Resistance Value

The resistor divider should be bypassed to ground near each Vref pin to reduce noise. A standard 0.01 μF ceramic bypass capacitor for C1 is recommended. Care should be taken in board layout to avoid signal crosstalk from other signal lines. Also, the Vref input pin should be kept away from other switching signal pins to avoid crosstalk in the package. Similar care should also be taken in routing the Vref signal inside the A.G.P. interface circuits to avoid noise pick-up.

Vref can also be generated internal to the A.G.P. interface, so the actual value of Vref may not be visible. In this case, VIL and VIH levels are calculated from and tested relative to Vddq. Similar care has to be taken in this type of design to avoid signals coupling to the internal Vref and local bypassing may be necessary

1.4.4.2 Vref Characteristics for A.G.P. 4X Mode

Vref is set between 0.48×Vddq and 0.521×Vddq Volts. To preserve the common mode relationship between data and Vref at the receiving component, Vref needs to be generated at the driving component using a simple resistor divider network from its Vddq and Vss. Two unidirectional Vref pins are provided in the connector that will be used in delivering Vref between the add-in card component and the motherboard component. As an example for outbound writes from the graphics master to the target, the Vref will be delivered using the unidirectional pin on the connector to the target. The Vref is generated using the master component's power supply Vddq and Vss using a simple resistor divider network. Since noise may be generated from other signals coupling to Vref, proper bypassing must be provided.

The voltage divider network consists of an AC and DC element. The maximum current at the input device using Vref must stay within the range of ±10 μA DC current load. For the DC element, the resistor values of R1 and R2 are equal and must be selected in order to meet this input current load specification. If 1% tolerance resistors are used, any standard value between 1.0 kΩ and 2.1 kΩ will work. For 2% tolerance resistors, use any standard value between 1.0 kΩ and 1.4 kΩ. For the AC elements, resistors R3 and R4 are equal in value, 2% or better tolerance, and 80 Ω (or closest standard value). For C1 and C2, a low ESR/ESL NPO or X7R type capacitor with a value of 0.5 nF at 4 MHz can be used for the Vref circuit. The Vref resistor divider network must be placed away from critical and noisy signals. Care should be taken in board layout to avoid signal crosstalk from other signal lines. Also, the Vref input pin should be kept away from other switching signal pins to avoid crosstalk in the package. Similar care should also be taken in routing the Vref signal inside the A.G.P. interface circuits to avoid noise pick-up.

Figure 19: Vref Resistor Divider for A.G.P. 4X Mode

1.5 A.G.P Vddq Power Delivery

Power delivery to A.G.P. components on both the add-in card and motherboard has two main areas of focus. The first area is the voltage regulator or power supply. The second area of power delivery is the decoupling of the power plane at the component and connector. For Vcc Decoupling recommendations, see Appendix D.

1.5.1 Vddq Voltage Sources

Power to the Vddq rail may be supplied by a voltage regulator or the system power supply depending on system configuration. In the case of a 3.3V only signaling environment, the system power supply will most likely be used. In the case of a 1.5V, or universal 3.3V/1.5V signaling system, a voltage regulator will most likely be used. In the universal case, the voltage level of the regulator is controlled by the **TYPEDET#** pin which is controlled by the add-in card.

The two main types of regulators that can be used for the Vddq rail are linear and switching. A linear regulator is defined as any with a voltage control circuit with a pass transistor in the current path. A switching regulator is defined as voltage control circuit with a switched inductive tank circuit used in the current path. Each type of regulator has associated advantages and disadvantages including: cost, part count, board area, response time, efficiency, and thermal dissipation. Examples of a linear and a switching regulator are shown in Figure 20 and Figure 21, respectively.

Figure 20: Linear Regulator

Figure 21: Switching Regulator

1.5.2 Power Plane Decoupling

Since any regulator or power supply response time is much too slow for the transients present in the A.G.P. signaling environment, capacitive decoupling of the power planes is essential. These transients can cause voltage drop on the power plane until some mechanism for power delivery can supply the needed current. Capacitors are used to supply power during transients until the voltage regulator can respond. In any system, two types of capacitors are needed to perform adequate decoupling: high frequency and bulk. High frequency capacitors are small in value (0.001 µF – 1 µF) and have very low equivalent series resistance (ESR) and equivalent series inductance (ESL). Although the capacitance of a high frequency capacitor is not large, it can deliver energy quickly to a power transient due to its low ESR and ESL. However, since the capacitance is low, it cannot sustain the power for large periods of time. Bulk capacitors are generally larger (10 µF – 10,000 µF) and have higher ESR and ESL than high frequency capacitors. The higher ESR and ESL make their response time to power transients longer and thus they are less effective at preventing voltage drops during these transient conditions. The larger capacitance of bulk capacitors does, however, allow them to deliver power to the system for longer periods of time than high frequency capacitors. The bulk capacitor can be thought of as a temporary voltage regulator that maintains the rail power and the high frequency capacitor power until the voltage regulator can respond to the power transients. Both the add-in card and the motherboard should have both bulk capacitance and high frequency decoupling capacitors on the Vddq power plane. The add-in card should have the decoupling capacitors near the graphics controller and the bulk capacitance should be located on the Vddq plane between the connector and the graphics controller. The motherboard should have the decoupling capacitors near the connector, for signaling reasons, and near the chipset containing the graphics interface for power delivery. Bulk capacitance should be located on the Vddq plane in the area of the chipset and A.G.P. connector.

1.5.3 Decoupling and Bulk Capacitance Calculation Methodology

Calculating the amount of decoupling capacitance can be accomplished by the following method. This method, and the equations used, are first order calculations intended to be as easy and accurate as possible. More complex and accurate calculation methods may be used as shown in Appendix D.

The following is the equivalent circuit for decoupling capacitors connected to an A.G.P. interface component. In Figure 22, Rc and Lc represent the ESR and ESL of the capacitor respectively. Zs is the source impedance of the output driver and Zo is the impedance of the interconnect.

Figure 22: Decoupling Capacitance

The following equations represent the response of the circuit in Figure 22 and can be used to calculate the decoupling capacitance requirements for the interface.

Vddq = Ic * Rs

Where: Ic = ((Vcc0*EXP((-t*Rt*(1-EXP(-t*Rt/Lc)))/C))* (1-EXP(-t*Rt/Lc)))/Rt

Rt = 1/(number of outputs switching * (Rs(t) + Zo)) + ESR
Lc = ESL
C = capacitance value
Vcc0 = initial Vcc value

Equation 1

Equation 1 is used to calculate the Vddq value for a given decoupling capacitor(s) in the presence of a current transient. The value of Vcc0 depends on the nominal system voltage level for Vddq. The final value of Vddq required depends on the minimum value of Vddq required by the component in question. Equation 1 can also be solved for the variable C to determine the amount of capacitance required to maintain a desired Vddq voltage. Note that the values of ESR, ESL, and C in equation 1 can be for one or more capacitors. When solving for more than one capacitor, the value of each capacitor is summed to find the total value C and the ESR and ESL for each capacitor is summed in parallel to get the total ESR and ESL.

Evaluating Equation 1 requires several pieces of data from the system. The first is the current consumed by the component and the duration in which it is consumed. This is referred to as dI/dt which represents the current transient. Equation 1 requires the current transient to be converted to an equivalent load resistance over time or dR/dt. To calculate dR/dt we must know the buffer impedance, interconnect impedance, and slew rate. To get the worst case numbers for our decoupling requirements we must take the case of the maximum current drawn and the fastest slew rate. The maximum current drawn occurs when the interconnect and buffer impedance are at the lowest allowable values. Equation 2 is used to calculate the source resistance over time, Rs(t).

$Rs(t) = Zo*((Vcc^2)/(Vstep * Slew\ Rate * t) – 1)$

Where: Vstep = Vcc * (Zo/(Zo + Zs))

Equation 2

The second piece of data needed is the round trip flight time of a signal on the trace. This will determine the time period over which to evaluate the decoupling capacitor performance in the system. To determine this, time we divide the longest interconnect length, including motherboard, connector, and add-in card by the flight time of the signal. The flight time should be chosen for the impedance condition of the interconnect used in equation 2 above. The equation for the maximum round trip flight time is:

Trt = 2*((Lmax/Tpd) + Tcd)

Where: Tpd = flight time per unit length or the interconnect
Lmax = total length of the interconnect (add-in card + motherboard)
Tcd = connector delay (150 ps)

Equation 3

Note that equation 2 has two time periods over which it is evaluated. The first is from time zero to the time at which the signal has fully transitioned. This is the period where the load resistance varies with time. The second period is from the point where the signal has transitioned to its final value to where the round trip reflection arrives at the transmitter and the current flow is cut to zero. This second time period is the time of constant load impedance.

With the load resistance and time period defined, equation 1 can now be evaluated. The result of this evaluation should give the Vddq voltage variation over the time period of current draw by the component. Equation 1 should be evaluated over many points during the time period since Vddq will not be a linear function due to the capacitor discharge characteristics and ESL of the capacitor. The voltage during the period from the initial time to the final time should be evaluated as well as the final voltage of the capacitor after the time period of evaluation.

After we have evaluated our high frequency decoupling capacitance requirements, we are ready to evaluate our bulk requirements. To do this, we must determine the voltage regulator response time and final high frequency capacitor voltage after a current transient. Equation 4 can be used to determine the final voltage, Vf, of the bulk capacitor and the high frequency capacitor after a transient has occurred and the high frequency capacitor has been recharged by the bulk capacitor.

$$Vf = (Vc1*C1 + Vc2*C2)/(C1 + C2)$$

Where: C1 = high frequency capacitance value
C2 = bulk capacitance value
Vc1 = final high frequency capacitor voltage after transient
Vc2 = initial bulk capacitor voltage value

Equation 4

The value for Vf will be the value for Vddq after one cycle of a current transient and a recharge of the high frequency after the transient has ended. To get the final value at the time the voltage regulator responds to the current transients, we must determine the number of transient/charge cycles in the voltage regulator response period. Assuming a worst case switching, (100% I/O toggle rate), transients will occur every 15 ns for A.G.P. 2X and every 7.5 ns for A.G.P. 4X signaling. Thus, the number of cycles would be:

$$Cycles = Tvreg/cycle\ time$$

Where: cycle time = 15 ns for 2X switching and 7.5 ns for 4X switching
Tvreg = response time of voltage regulator

Equation 5

To determine the voltage present on the capacitors at the end of the voltage regulator response delay period, we must evaluate equation 4 and equation 5. Multiplying the single cycle voltage drop, Vddq initial – Vf from equation 0.3, by the number of cycles given in equation 5, will give the total voltage drop before the regulator responds. The final voltage value should be above the minimum allowable Vddq value for the interconnect voltage being evaluated.

1.5.3.1 Add-In Card Calculation Recommendations

To calculate the decoupling and bulk required by the add-in card, we will use Equations 1, 2, 3, 4, and 5 to calculate the decoupling and bulk capacitance needed for the graphics device on the add-in card. Worst case conditions for both the 3.3V signaling environment, and the 1.5V signaling environment should be used to calculate the required capacitance. Recommended system assumptions and conditions for the calculations are as follows:

1. For the add-in card, the worst case maximum number of outputs switching for the graphics controller should be assumed at 50. This number will cover all address, data, sideband, and control lines.
2. Maximum slew rate for the graphics component being used for the 3.3V and 1.5V signaling environment. (A.G.P. specification for maximum slew rate for 3.3V signaling is 4.0 V/ns and 3.0 V/ns for 1.5V signaling.)
3. Minimum graphics device output source impedance for the component being used. (The *A.G.P. Interface Specification*, Rev 2.0 for minimum output source impedance is 30 Ω.)

4. Minimum board impedance is 48 Ω including crosstalk effects for 4X transfer mode compliant boards. The minimum impedance for boards that are not 4X mode compliant is 30 Ω.
5. Worst case period over which current is delivered to the interconnect. This is defined as signal slew time plus the round trip flight time of the interconnect. Equation 3 should be used to calculate the total round trip flight time for the interconnect. Here, a worst case motherboard length of 7.25" should be assumed. For the minimum board impedance of 48 Ω, the flight time is 138 ps/in. The A.G.P. connector delay is assumed to be 150 ps. The add-in card trip flight time is dependent on the add-in card A.G.P. bus routing.
6. Motherboard power delivery response time should be assumed at 10 µs.
7. Vddq delivered to the add-in card assumed at minimum A.G.P. level of 3.15 V for 3.3V signaling and 1.425 V for 1.5V signaling.
8. Worst case number of output transitions over the response period needs to be evaluated. It is recommended to use a 2X transfer rate (66 MHz) over the 10 µs response time (667 transitions) for 3.3V signaling and 4X transfer rate (133 MHz) over the 10 µs response time (1334 transitions) for the 1.5V signaling environment.
9. Capacitor characteristics used must be identified. These include ESR and ESL for both the decoupling and bulk capacitors.

Using the equations outlined above, both the decoupling and bulk capacitance for the chipset can be calculated. The high frequency decoupling capacitors should be placed as close to the graphics controller as possible. The bulk capacitance should be placed on the Vddq power plane on the add-in card.

1.5.3.2 Motherboard Calculation Recommendations

The motherboard calculation includes decoupling and bulk capacitors for the chipset and bulk capacitors for the add-in card, which may be needed to supplement the Vddq voltage regulator. Again, equations 1, 2, 3, 4, and 5 will be used. Here the worst case conditions must be used for the signaling environment used on the motherboard as indicated by the connector: 3.3V only, 1.5V only, or both 3.3V and 1.5V for a universal connector. Recommended system assumptions and conditions for the calculations are as follows:

1. For the motherboard, the worst case maximum number of outputs switching for the chipset should be assumed at 40. This number will cover all address, data, and control lines.
2. Maximum slew rate for the chipset component being used for the 3.3V and 1.5V signaling environment. (The A.G.P. specification for maximum slew rate for 3.3V signaling is 4.0 V/ns and 3.0 V/ns for 1.5V signaling.)
3. Minimum chipset source impedance for the component being used. (The A.G.P. Interface Specification Rev 2.0 for minimum output source impedance is 30 Ω.)
4. Minimum board impedance is 48 Ω including crosstalk effects for 4X transfer mode compliant boards. The minimum impedance for boards that are not 4X mode compliant is 30 Ω.
5. Worst case period over which current is delivered to the interconnect. This is defined as signal slew time plus the round trip flight time of the interconnect. Equation 3 should be used to calculate the total round trip flight time for the interconnect. Here, a worst case add-in card length of 1.5" should be assumed. For the minimum board impedance of 48 Ω, the flight time is 138 ps/in. The A.G.P. connector delay is assumed to be 150 ps. The motherboard flight time is dependent on the add-in card A.G.P. bus routing.
6. Motherboard power delivery response time and minimum voltage level are obtained from the particular motherboard voltage regulator implementation for the board in question.
7. Worst case number of output transitions over the response period needs to be evaluated. It is recommended to use a 2X transition rate (66 MHz) over the response time of the voltage regulator on the motherboard for 3.3V signaling and 4X transition rate (133 MHz) over the response time of the voltage regulator on the motherboard for the 1.5V signaling environment.
8. Capacitor characteristics used must be identified. These include ESR and ESL for both the decoupling and bulk capacitors.

Using the equations outlined above, both the decoupling and bulk capacitance for the chipset can be calculated. We must also evaluate the amount motherboard bulk needed to supply adequate power to the add-in card in the case where the Vddq voltage regulator has a response time greater than the 10 µs assumed by the add-in card. If the response time is indeed greater than 10 µs, the motherboard designer needs to use equation 4 along with the motherboard and add-in card bulk capacitance, to assure adequate power delivery to the add-in card. To determine the expected add-in card bulk capacitance, the motherboard designer should use the method outlined in Section 1.5.3.1. In some cases, the bulk capacitance required by the chipset will be enough to make up the power required by the add-in card between the 10 µs response time assumed by the add-in card until the motherboard voltage regulator can respond. However, more bulk may need to be added to the motherboard Vddq plane to supplement the chipset bulk and the voltage regulator.

1.5.3.3 Connector A.C. Signal Decoupling Requirements

The decoupling capacitor recommendations for the A.G.P. connector are intended to address A.C. signaling issues and not power delivery issues. The main reason for not addressing power delivery issues with the motherboard connector decoupling is due to the connector inductance and the distance the capacitors are from the graphics device. These two factors negate much of the usefulness of the connector decoupling on the motherboard for power delivery purposes.
The following are recommendations for decoupling at the A.G.P. connector on the motherboard:

Vcc3.3: Four (three for a 1.5V connector) 0.01 µF or larger, low ESL capacitors. Each capacitor should be placed as close as possible to a Vcc3.3 pair of pins on the connector.

Vddq: Six 0.01 µF or larger, low ESL capacitors. Each capacitor should be placed as close as possible to a Vddq pair of pins on the connector.

+5 V: One 0.01 µF or larger, low ESL capacitor placed as close as possible to the +5 V connector pins.

+12 V: One 0.01 µF or larger, low ESL capacitor placed as close as possible to the +12 V connector pin.

3.3VAUX: Two (One for 1.5V connector) 0.01 µF or larger, low ESL capacitors placed as close as possible to the **3.3VAUX** connector pin(s).

Note that all power and ground pins must be connected on the motherboard to guarantee power delivery and proper AC signal return paths.

All ground pins on the add-in card need to be connected for proper grounding and AC return paths. All Vddq pins need to be connected on the add-in card to the local Vddq power plane. The add-in card does not have to use all of the Vcc3.3 power pins if it does not need the power delivery. However, all unused power pins must be bypassed to ground close to the connector with a 0.01 µF or larger, low ESR/ESL capacitor to provide good AC coupling to the adjacent signaling pins.

1.6 Motherboard and Expansion Card PCB Layout Considerations

This section describes layout and routing guidelines to insure a robust A.G.P. interface design. See Chapter 2 for more detail on these topics and specific layout requirements. These guidelines help insure that the A.G.P. specifications can be met and that the motherboard and daughter card will operate together; however, it is not a guarantee that the A.G.P. specifications will be met. The system designer must analyze and simulate the A.G.P. 4X system to insure that all specifications are met.

1.6.1 Motherboard Layout Recommendations

The following figures (Figure 23, Figure 24, and Figure 25) are examples of A.G.P. layouts for ATX, LPX, and NLX form factor motherboards. The goal of the placement is to allow the routing of the A.G.P. bus that minimizes trace length, vias, and interference with other signals and buses.

Figure 23: A.G.P. PCB Layout for ATX Form Factor

Figure 24: A.G.P. PCB Layout for LPX Form Factor

Figure 25: A.G.P. PCB Layout for NLX Form Factor

The pinout of the core chipset should allow the A.G.P. bus to flow to the A.G.P. connector or graphics device with a minimum trace length and signal crossing. The A.G.P. strobe signals must be grouped with their associated data signals. It is recommended that the strobe be centered within its group to minimize the signal to strobe skew.

1.6.2 Expansion Card Layout Recommendations

All A.G.P. signals on the graphics chip should be located so that they are in the same order as the pinout of the A.G.P. connector. This assumes that the component side of the add-in card is as defined in the mechanical section of the *A.G.P. Interface Specification, Revision 2.0*. This alignment will minimize the overall trace lengths and aid in matching the trace lengths within its groups.

The strobe signals must be grouped with their associated data group.

AD_STB0/AD_STB0#* with **AD[15::0]** and **C/BE[1::0]#**
AD_STB1/AD_STB1#*with **AD[31::16]** and **C/BE[3::2]#**
SB_STB/SB_STB#* with **SBA[7::0]**

* The complement strobe signals are available in the A.G.P. 4X mode only.

If the pinout is optimized for the expansion card connector card pinout, the graphics on-board only solution will cause the A.G.P. bus on the graphics chip and the chipset to be reversed with respect to each other. There should be sufficient space on the motherboard to route the bus and meet the A.G.P. requirements in this case.

1.6.3 Board Impedance

The motherboard and add-in card impedance should be controlled to minimize the impact of any mismatch between the motherboard and the add-in card. The PCB should be fabricated using FR-4 with an overall board thickness of 62 mils ±10%. The outer layers should be made using ½ oz Cu, and the inner layers should be made with 1 oz Cu. The A.G.P. 2X board impedance is targeted at 65 Ω ±15 Ω for a four layer PCB stackup. An impedance of 60 Ω ±10% is recommended for A.G.P. 4X. The outer layer dielectric thickness is thinner for A.G.P. 4X thus lowering the impedance. Refer to Section 2.2.4.1 for A.G.P. 2X and A.G.P. 4X board layer stackup.

Lower trace impedance will reduce signal edge rates, decrease over/undershoot, and have less crosstalk than higher trace impedance.

1.6.4 Pull-ups and Termination

All pull-up resistors required by the A.G.P. interface specification must be implemented on the motherboard. The A.G.P. bus was designed to not require any special signal termination. It is up to the designer to simulate the signals to insure that signal quality requirements are met.

1.6.5 Board Routing Recommendations

A.G.P. signals must be carefully routed on the motherboard and graphics card to meet the timing and signal quality requirements of this interface specification. The following are some general guidelines that should be followed. Trace lengths included in this section are guidelines only. It is recommended that the board designer simulate the routes to verify that the specification is met.

1.6.5.1 Maximum Trace Length Requirements

This section describes how the maximum trace length is segmented between the motherboard and the expansion card. For the common clock mode, the maximum flight time allowed for the bus is 2.5 ns. The flight path timing budget for the motherboard is 1.65 ns and 0.7 ns for the expansion card. Thus, the maximum trace length on the motherboard is 9.5 inches and 4 inches on the add-in card. Figure 26 shows how the maximum trace length is segmented between the motherboard and the expansion card for the common clock mode. A.G.P. 2X and A.G.P. 4X source synchronous signals are limited by flight time skew; therefore, the budgets may not be able to reach those specified for the common clock mode.

The maximum trace length is also dependent on the effect of crosstalk on signal skew. When the effect of crosstalk on signal skew is taken into account, the maximum length of the motherboard and add-in card traces is reduced. The amount of crosstalk depends on trace length and trace spacing.

For A.G.P. 2X mode with a 1:2 trace spacing, distance between traces (air gap) being twice the trace width, the maximum trace length on the motherboard remains at 9.5 inches. For a 1:1 trace spacing, distance between traces (air gap) being equal to the trace width, the maximum trace length on the motherboard reduces to 4.5 inches. Add-in cards should be routed with a 1:2 trace spacing and are restricted to a maximum length of 3.0 inches.

For A.G.P. 4X add-in cards and motherboards, the signals should be routed out of the graphics controller with a 1:2 trace spacing. When the signals have been routed away from the graphics controller and towards the A.G.P. connector, the source synchronous traces (data and strobe signals) should use a spacing of ≥3 times the maximum dielectric thickness between the layer that the traces are routed on and the nearest voltage plane (Vss, Vddq, etc.). It is also recommended that the spacing between the strobe pairs be between three or four times the maximum dielectric thickness. The strobes should also be spaced to other traces by a minimum of five times the maximum dielectric thickness. The signal line lengths should be matched as close as possible from the breakout area, as shown by the dotted line in Figure 26, to the add-in card edge fingers. The same holds true from the core chipset to the connector.

Figure 26: Trace Length Routing Ratio

A = Expansion Card Timing Budget = 0.70 ns = ~ 4.0 inches

B = Motherboard Timing Budget = 1.65 ns = ~ 9.50 inches

Figure 27: Common Clock Timing Budget Between the Graphics Chip, A.G.P., and Core Chipset

It is recommended to keep the routes as short as possible to maximize timing margin and signal quality.

1.6.5.2 Trace Length Mismatch Requirements

The trace lengths for signals within a group must be matched to their respective strobe to meet the maximum mismatch requirement given in the A.G.P. Interface Specification. Of the maximum 0.7 ns, the motherboard is allotted 0.5 ns and the add-in card is allotted 0.2 ns. For A.G.P. 4X mode, the motherboard is allotted 25 ps and the add-in card is allotted 25 ps. When the effects of varying board impedances and crosstalk are considered, the following trace length constraints are recommended.

Add-in Card: For A.G.P. 2X, the data lines should be kept to within ± 0.5 inches of their respective strobe. For example, if the strobe is at 2.5 inches, the data lines can be 2.0 to 3.0 inches in length. For A.G.P. 4X, the data lines should be kept to within ± 0.25 inches of their respective strobe.

Motherboard: For A.G.P. 2X mode, the data lines should be kept at +0.0, -0.5 inches from their respective strobe. For example, if the strobe is at 9.0 inches, the data line can be from 8.5 to 9.0 inches in length. The data lines should be kept within ± 0.125 inches of their respective strobe for A.G.P. 4X.

To avoid additional signal mismatch, all of the lines within a group need to be the same type (either microstrip or stripline, but not both). This is because microstrip (surface traces) and striplines (buried traces) have different propagation velocities and mixing these can increase the flight time skew beyond acceptable limits. It is further recommended to route all signals within a group on the same layer. Routing studies have shown that these guidelines can be met. The maximum trace length requirements must not be violated by any signal. It is recommended to match the signals as closely as possible to provide timing margin.

1.6.5.3 Strobe Trace Routing Considerations

Since the strobe signals (**AD_STB0**, **AD_STB0#**, **AD_STB1**, **AD_STB1#**, **SB_STB**, and **SB_STB#**) act as clocks on the source synchronous A.G.P. bus, special care should be taken when routing these signals. Note that the complement signals are available in A.G.P. 4X mode only. For A.G.P. 2X, it is recommended that the strobe signals be routed in a 1:2 trace width/space ratio or greater relative to the other signals of the bus. For example, in a bus where the other signals are routed with 5 mil trace width and 5 mil spaces between, the strobe signals would be routed with a 10 mil separation on both sides. For A.G.P. 4X, the strobe signals should be spaced to other traces by a minimum of five times the maximum dielectric thickness. This recommendation is intended to reduce the crosstalk noise coupled onto the strobes from other signals on the bus, as well as, reduce the noise coupled from the strobe signals onto adjacent lines.

1.6.5.4 Clock Routing and Skew

Clock skew between the A.G.P. compliant graphics chip and the core chipset must be held to under 1 ns. The motherboard is allotted 0.9 ns of which a portion must be allotted to the clock generator circuit. The add-in card is allotted 0.1 ns of the total skew requirement. Figure 28 shows the trace segments that make up the clock routing.

Figure 28: Required Trace Length to Minimize Clock Skew

The following clock skew equation is intended for use by the motherboard designer since the clock delay on the add-in card is tightly regulated to 0.6 ns ± 0.1 ns by the A.G.P. interface specification. It should be noted, however, that add-in card vendors should center the clock trace electrical length at 0.6 ns at typical board impedance so that variation in board impedance will be covered in the ± 0.1 ns specification.

In this equation, all lengths are represented in terms of time (electrical length). This equation should be used several times to account for impedance variation on both the add-in card and the motherboard. The equation for clock skew is represented as a solution space as follows:

$$-(1.0\text{ ns} - Tcgs) \leq A - (B + C + Tcon) \leq (1.0\text{ ns} - Tcgs)$$

Where:

A = The electrical length of the clock trace on the motherboard from the clock generator to the Core Logic Memory Control (CLMC)

B = The electrical length of the clock trace on the motherboard from the clock generator to the A.G.P. connector

C = The electrical length of the clock trace on the add-in card which from A.G.P. interface specification is set at 0.6 ns + 0.1 ns.

Tcon = The delay through the A.G.P. connector. We have found this to be 0.15 ns including crosstalk effects.

Tcgs = The clock generator skew between the clock signal to the chipset and the clock signal to the A.G.P. connector.

NOTE: *For a motherboard with graphics down, the above equation reduces to a simple subtraction of the two motherboard traces A and B.*

The combination trace of B + C + Tcon should be matched as close to trace A as possible.

1.6.6 Interface Signaling Requirements

For A.G.P. 2X, all signals are 3.3V compatible. No 5V signals are specified in the A.G.P. signaling environment. All A.G.P. 4X signals are +1.5V compatible. The master and target device must be capable of supporting the 3.3V signaling environment for A.G.P. 2X and 1.5V signaling environment for A.G.P. 4X. However, in a 1.5V signaling environment, there are some signals that are not 1.5V compatible. These signals are the interrupt lines, reset, and the A.G.P. clock. The clock and reset signals are 3.3V level and are driven to both the target and master devices. The interrupt signals from the A.G.P. bus must interface to the PCI bus interrupt controller. This controller and the PCI devices may be +5V devices. It is the requirement of the motherboard designer to properly interface the A.G.P. interrupts to the PCI bus. Since the interrupt drivers are defined to be open-drain, this can be done by simply pulling up the PCI interrupts to 3.3V only allowing the A.G.P. interrupts to connect directly to the PCI interrupts.

Another option is to buffer the interrupts before they cross from the 3.3V domain to the 5V domain. Also, the two A.G.P. interrupt lines must be “swizzled” with the PCI interrupt lines as described in the implementation note in the *PCI Local Bus Specification, Revision 2.1*. As an example, in a system with three PCI slots (and one A.G.P. device or slot), the interrupts should be connected such that **INTA#** of each slot is assigned to a unique input on the system interrupt controller.

1.6.7 IDSEL Routing for Add-in Cards

As it is described in the *A.G.P. Interface Specification, Revision 2.0*, **IDSEL** is not a pin on the A.G.P. connector.

When the graphics device on the add-in card is designed for exclusive operation on the A.G.P. interface, the device does not have an external **IDSEL** pin. There is no connection of an **IDSEL** signal to any **AD** signal on the add-in card. In this implementation, the device asserts **DEVSEL#** based monitoring of the **A.G.P.** bus inside the graphics controller.

When the graphics device on the add-in card is designed to be used on both A.G.P. and PCI bus segments, then the device needs to have two modes of operation. When in the A.G.P., mode it generates **DEVSEL#** as described in the A.G.P. only implementation above. When used in a PCI mode of operation, the device must provide an external **IDSEL** that is connected to one of the **AD** signals in the system. This **IDSEL** signal is not connected externally to an **AD** signal on add-in cards designed for use on the A.G.P. bus.

1.7 IC Packaging Considerations

This section describes some generic characteristics of signals that packaging designers may need to take into consideration before designing their package.

Due to the high A.G.P. pin count, there are two packages that most manufacturers may consider for A.G.P. 2X mode. One is the quad flat pack (QFP) and the other is a ball grid array (BGA) package. Of the various packaging technologies, the BGA package is recommended for A.G.P. 4X. The largest benefit of the BGA package is the reduction in power, ground, and signal inductance. This, in turn, minimizes simultaneous switching noise effects.

The BGA can provide a significantly improved electrical interface over the QFP. The BGA can have a ground plane that may result in the following benefits:

- Help to control signal trace impedance
- Help to reduce crosstalk
- Help to lower ground inductance and thus simultaneous switching noise effects

Much of the benefit of a ground plane depends on how it is implemented. Another major design impact between QFP and BGA is on PCB routing. PCB routing with BGA should benefit with increasing number of routing layers in the PCB.

1.7.1 Signal Skewing

Time skew between signals depends on differences between signal loading and may also depend on signal transitions on neighboring traces. Signal transitions must charge and discharge the trace capacitance. Two signals driving different capacitance loads will have different timings. Neighboring signals switching in-phase or out-of-phase may also have a significant effect if coupling or crosstalk is large.

1.7.2 Signal Loading

Signal loading refers to the amount of capacitance on a signal trace. This is a function of trace length, proximity to other metal structures, and the relative dielectric strength. Package pin capacitance is very design dependent. In QFP, pin capacitance is a strong function of the lead length. This is due to the fact that lead spacing is generally constant among the leads. The range between minimum and maximum capacitance will depend on the overall size of the QFP. In BGA packaging, the capacitance will depend on the existence of a power or ground plane, the width of the signal trace, the density of signal traces, the width of the plating bars, and the density of the plating bars. In general, signals routed to the corners of the pin field tend to be longer and be in areas of higher metal density.

1.7.3 Signal Routing and Crosstalk

Crosstalk between two traces is a function of the coupled length, the distance separating the traces, the "aggressor" signal edge rate, and the degree of mutual capacitance and inductance. In a non-homogeneous medium, both forward and backward crosstalk will be excited on the victim signal trace, whereas only backward crosstalk will be excited within a homogeneous medium. Crosstalk is a problem for both asynchronous and synchronous signals, although it may be "time out" for synchronous signals. The result can be false triggering or data corruption. Crosstalk can also affect signal timing (see signal skewing above). Mutual capacitance and inductance between signal traces can be reduced by the presence of a ground plane.

2. AGP Signals Sensitivity Analysis and Measurement Techniques

These guidelines are primarily for A.G.P. designs that use an A.G.P.-compliant graphics controller on an A.G.P.-compliant add-in card. Designs with both A.G.P.-compliant components on the motherboard can also use these techniques, using the combined trace lengths and tolerances.

This chapter is divided into two major sections:

- Section 2.1 deals with designs that are A.G.P. 1X transfer mode (1X mode) and A.G.P. 2X transfer mode (2X mode) specific. This includes both 3.3 volt signaling and 1.5 volt signaling.
- Section 2.2 describes A.G.P. 4X transfer mode (4X mode). 4X mode designs need to also operate in 2X mode. 2X mode compliance will be achieved when a design is 4X mode compliant.

2.1 2X Mode Board Design Guidelines

2.1.1 2X Mode Routing Rules

All of the techniques used in this section work for 1.5 volt signaling and 3.3 volt signaling. These board design guidelines were first based on 3.3 volt signaling. Analysis of 1.5 volt signaling did not result in reductions of any margins reported in this section.

There are three ways that these guidelines can be used:

- The line length recommendations can be followed as they are described in Table 7 and Table 8. These rules allow enough variation to meet most designs.
- Section 2.1.3 can be used to apply variations to the existing design recommendations. This might be desired to take advantage of margin in other parts of a design (for example, lines routed shorter than those assumed by this guideline might not need to be matched as tightly). This can be done by applying the timing numbers shown in Table 15 or Table 16 to your own design.
- Section 2.1.2 can be used as a pattern for additional simulations, which may then be used for other designs. This section describes how the analysis was performed to make these guidelines. This can yield the most flexible designs; however, it requires extensive simulations.

2.1.1.1 A.G.P. 2X Mode Design Considerations

With source synchronous data transfers (A.G.P. 2X mode) skew between traces is a major interconnect design consideration. Skew is defined in the formula $T_{skew} = Tf_{data} - Tf_{strobe}$ where Tf_{data} and Tf_{strobe} are the flight times of the data and strobes respectively. These skews are caused by various effects that include line length mismatch, capacitance loading variations, crosstalk, and the voltage level that a driver starts driving from.

Figure 29: Example of Skews

Figure 29 is an example of how some of the different components contribute to the total skew. Line length mismatch is the physical difference in line lengths. Component loading mismatch is the difference in capacitive loading. These two effects are mainly independent of other effects.

Settling time and crosstalk have the largest effect on skews. Settling time creates uncertainty in the signal transition starts. Crosstalk is the effect of coupling between traces. Both of these interact with each other and with some of the other parameters in the interconnect.

With A.G.P. 2X mode, it becomes more difficult to partition the timings between motherboard and add-in card. This is due to the fact that crosstalk and its effect on impedance discontinuity greatly increases the electrical skew between traces.

Throughout Section 2.1 of this document, the term “data” refers to the AD[31::0], C/BE[3::0]# and SAB[7::0] signals. The term “strobe” refers to AD_STB[1::0] and SB_STB. The term data refers to one of three groups of data as seen in Table 6. The term strobe refers to one of the three strobes as it relates to the data in its associated group.

Table 6: Data and Associated Strobe

Data	Associated Strobe
AD[15::0] and C/BE[1::0]#	AD_STB0
AD[31::16] and C/BE[3::2]#	AD_STB1
SBA[7::0]	SB_STB

2.1.1.2 Source Synchronous Recommendations for Add-in Cards

Table 7: Add-in Card Recommendations

Width:Space	Zo	Trace	Line Length	Line Length Matching
1:2	50 Ω to 85 Ω	Data / Strobe	0.0 in < line length < 3.0 in	Strobe ±0.5 in of group

All of the data line lengths within a group of signals need to be within ±0.5 inches of their associated strobe. The board impedance need to be in the range of 50 Ω to 85 Ω. This range is used to cover design targets and manufacturing tolerances.

Because crosstalk is a large component of skew, it is necessary to specify board routing. All traces need to be routed with a separation of two times the trace width. Additionally, all lines within a group need to be of the same type (either microstrip or stripline). This is because microstrip (surface traces) and striplines (buried traces) have different propagation velocities and mixing these can increase the flight time skew beyond acceptable limits. Routing studies have shown that these guidelines can be met.

2.1.1.3 Source Synchronous Recommendations for Motherboards

The motherboard needs to have an impedance range of 50 Ω to 85 Ω. This range is used to cover design targets and manufacturing tolerances. All lines should be at least 1.0 inch long. The maximum line lengths are dependent on the type of trace and the amount of coupling.

The maximum line length is dependent on the routing rules used on the motherboard. These routing rules were created to give freedom for designs by making tradeoffs between signal coupling (trace spacing) and line lengths. These routing rules are divided by trace spacing. 1:1 spacing refers to the distance between the traces (air gap) as being the same width as the trace. 1:2 spacing refers to the distance between the traces as being twice the width of the trace.

For trace lengths that are under 4.5 inches, a 1:1 data trace spacing gives a data line length mismatch of 0.5 inches with the strobe being the longest trace of the group. The strobe requires at least a 1:2 trace spacing. Use greater spacing, if possible, to minimize crosstalk and the possibility of coupling glitches from data to the strobe line. This is for designs that require less than 4.5 inches between the A.G.P. connector and the A.G.P. Target.

Longer lines have more crosstalk. Therefore, longer line lengths require a greater amount of spacing between traces to maintain skew timings. Table 8 shows 1:1 spacing may be used for lengths up to 4.5 inches, but 1:2 spacing allows lengths up to 9.5 inches. The line length mismatch is 0.5 inches with the strobe being the longest trace of the group. Here, the strobes are routed at 1:4, twice the spacing of the data lines. These timings also allow lines to neck down (reduced spacing) so that they may break out from a component

Table 8: Motherboard Recommendations

Width:Space	Zo	Trace	Line Length	Line Length Matching
1:1(Data) / 1:2 (Strobe)	50 Ω to 85 Ω	Data / Strobe	1.0 in < line length < 4.5 in	-0.5 in, strobe longest trace
1:2	50 Ω to 85 Ω	Data / Strobe	1.0 in < line length < 9.5 in	-0.5 in, strobe longest trace

In all cases, it is best to reduce the line length mismatch wherever possible to insure added margin. It is also best to separate the traces, especially the strobe to neighboring traces, by as much distance as possible to reduce the amount of trace to trace coupling.

2.1.1.4 Control Signal and Clock Recommendations

Table 9: Control Signal Line Length Recommendations

Width:Space	Board	Trace	Line Length	Pull-up Stub Length
1:1	Motherboard	Control signals	1.0 in < line length < 8.5 in	< 0.5 in (Strobes < 0.1 in)
1:2	Motherboard	Control signals	1.0 in < line length < 10.0 in	< 0.5 in (Strobes < 0.1 in)
1:2 (1:4 to Strobe)	Motherboard	Clock		
1:2	Add-in Card	Control signals	0.0 in < line length < 3.0 in	
1:2 (1:4 to Strobe)	Add-in Card	Clock	4.0 in ± 0.25 in (microstrip)	

Some of the control signals require pull-up resistors to be installed on the motherboard. The stub to these pull-up resistors needs to be controlled. The maximum stub length on a strobe trace is < 0.1 inch. The maximum stub trace length on all other traces is < 0.5 inches.

The clock lines on both the motherboard and the add-in card can couple with other traces. It is recommended that the clock spacing (air gap) be at least two times the trace width to any other traces. It is also strongly recommended that the clock spacing be at least four times the trace width to any strobes.

The clock lines on the motherboard need to be simulated to determine their proper line length. The motherboard needs to be designed to the type of clock driver that is being used and motherboard trace topology. These clocks need to meet the loading of the receiving device as well as the add-in trace length.

The add-in card trace length can be computed from the typical add-in card trace velocity.

2.1.2 Simulation Techniques

This section summarizes the interconnect simulations that were done to generate the routing guidelines for the A.G.P. (66 MT/sec and 133 MT/sec) bus. The result of these simulations shows that designs can be done with buffers that meet the *A.G.P. Interface Specification, Revision 2.0.* Figure 30 is a diagram of the topology used in this analysis.

Figure 30: A.G.P. with One Connector Topology

The common clock simulations (66 MT/sec mode) limit the total line lengths (motherboard length plus add-in card length) to greater than 1.0 inch and less than 10.0 inches.

The source synchronous simulation results limit the total line lengths to 9.5 inches for the cases used. This line length is dependent on coupling between lines and line length mismatch.

The source synchronous simulation result limits the line length skews between a signal and its associated strobe on an add-in card ±0.5 inch.

The source synchronous simulation results limit the motherboard's line length skews between a signal and its associated strobe to be 0.5 inches. The strobe needs to be the longest trace of each group.

The signal quality results for 3.3 volt signaling show overshoot and undershoot values approaching 2.1 V. Any device that connects to the A.G.P. bus needs to be tolerant of these levels. Also, the ringback levels are well within the *A.G.P. Interface Specification, Revision 2.0* requirements. Settling time was factored into both the common clock and the source synchronous results.

2.1.2.1 Methodology

Three main factors define the solution space: common clock flight times, source synchronous flight time skews, and signal quality. The common clock solution space is found by determining the topologies that result in lowest allowable flight times and the highest allowable flight times. The source synchronous solution space consists of all designs in which the flight time mismatch between a strobe and its associated data is less than the total allowable skew. Signal quality is checked to ensure that the solution space given by the common clock and source synchronous constraints meet all signal quality specifications.

Common clock (A.G.P. 1X mode) flight times use worst case (extreme) buffers and interconnects to find its solution space. For minimum flight times, a fast buffer driving the shortest (electrically) interconnects into a buffer with the smallest amount of loading. The physical line length is then increased until the minimum flight time is met. For maximum flight times, a slow buffer driving the longest (electrically) interconnects into a buffer with the largest amount of loading. The physical line length is then varied until the maximum flight time is met.

Source synchronous (A.G.P. 2X mode) flight time skews are found by comparing a strobe flight time with a data flight time for a particular interconnect, buffer, and loading. The skew is the total of both the setup skew and the hold skew. Setup skew is when the strobes minimum flight time is compared to the data's maximum flight time at a given line length. Hold skew is when the strobes maximum flight time is compared to the data's minimum flight time at a given line length. The results of these two comparisons must be less than the total allowable skew.

The results of the common clock solution space will give the minimum line lengths. The maximum line lengths will be determined by either the common clock or source synchronous. The signal quality of the signals within the solution space may constrain the solution space even further.

Of the three factors, source synchronous flight time skews are the most difficult to model. This is partly because they are a result of several parameters. These parameters include crosstalk, loading, impedance variations, line length variations, and settling time.

2.1.2.2 Buffer Models

The models for both the Master and the Target were identical. Both fast and slow buffer models were used in the simulations. In all cases, when a fast buffer was used, it was used in both locations. The same was also true for the slow buffers.

The buffer models were created from the *A.G.P. Interface Specification, Revision 2.0.* These models also used power supply diodes. The package was modeled with an effective inductance of 10.0 nH.

The models also included an effective capacitance of 2 pF for the fast model and 8 pF for the slow model. When modeling the strobe, these capacitive values were varied by +2 pF and -1 pF.

2.1.2.3 Interconnect Models and Crosstalk Effects

Both the motherboard and the add-in card traces are modeled as transmission lines. The two parameters of importance in these models are the characteristic impedance (measured in ohms) and the propagation delay (measure in ns/ft, and sometimes loosely referred to as "velocity" even though this is not strictly correct).

The values used for characteristic impedance and propagation delay depend primarily on three contributing factors: trace geometry, printed circuit board manufacturing tolerance, and crosstalk from other traces. Trace geometries that are typically used for motherboards and add-in cards will yield impedances in the range of 60-75 Ω. Manufacturing tolerances can be expected to add another 15%, increasing this range to approximately 50-85 Ω.

The impact of crosstalk is a little more difficult to comprehend. To model the impact of crosstalk, it is necessary to understand how the impedance and velocity vary between the even mode (more correctly called the "common mode," meaning that all coupled traces are switching in the same direction) and the odd mode (more correctly called the "differential mode," meaning that coupled traces are switching in opposite direction).

Common and differential mode impedances and propagation delays were modeled using a wide range of printed circuit board trace geometries (trace width, separation distance between traces, and trace height above ground/power plane). For this analysis, this crosstalk effect was included in the data traces. Strobe traces, however, did not include this effect because the strobes do not switch at the same time as the data traces. The results are summarized in Table 10.

After completing the simulations described in the next section, it was determined that the data traces must be designed so that their impedance is confined to 30-120 Ω (including all of the above effects). The geometries listed in Table 7, Table 8, and Table 9 achieve this goal.

The connector was modeled as a transmission line with 150 ps delay and 35-80 Ω characteristic impedance.

Table 10: Interconnect Impedance and Velocities

	Board Condition	Strobe		Data		Mode
		Z0 (Ω)	S0 (ns/ft)	Z0 (Ω)	S0 (ns/ft)	
6 mil x 12 mil Board	**Fast**	85	1.6	60	1.38	Odd
				99	1.70	Even
	Slow	50	2.0	45	1.72	Odd
				74	2.12	Even
6 mil x 6 mil Board	**Fast**	85	1.6	41	1.28	Odd
				119	1.80	Even
	Slow	50	2.0	30	1.60	Odd
				88	2.24	Even
Connector		80	2.0	63	1.83	
		35	2.0	52	2.17	

2.1.2.4 Simulation Setup

Simulations were run to find flight times for A.G.P. 1X clocking mode, flight time skews for A.G.P. 2X clocking mode, and signal quality for both modes. Earlier simulations had shown that the solution space would be confined by A.G.P. 1X mode for the minimum line lengths and A.G.P. 2X mode for the maximum line lengths.

All simulations were run using various combinations of line lengths for both the motherboard and the add-in card. The motherboard line lengths were varied from 0 to 10.0 inches in 0.5 inch increments. The add-in card line lengths were at 2.5 inches and 3.0 inches.

The strobe signal was simulated with eight different combinations of interconnects. These interconnects are fast/slow motherboard, fast/slow add-in card, and fast/slow connector.

The other signals were simulated with 32 different combinations of interconnects. These combinations are four velocity and impedance combinations of motherboard, four velocity and impedance combinations of add-in card and two combinations of connectors. The four combinations of board interconnect were fast and slow in both odd and even mode. The two combinations of connectors were fast odd mode and slow even mode. All of these simulations were done with both 1:1 and 1:2 motherboard trace spacing.

3.3 volt signaling simulations used 3.3 V for Voh and 0.0 V for Vol. A threshold value of 1.32 V was also used for all simulations. This 1.32 V was derived by taking 40% of the Voh minus Vol.

1.5 volt signaling simulations used 1.5 V for Voh and 0.0 V for Vol. A threshold value of 0.75 V (0.5Vddq) was also used for all simulations.

Additional simulations were run to find the clock to data out (T_{CO}) for the buffer models that were used. Section 2.1.2.6 describes how the T_{CO} simulations were used in the flight time and flight time skew calculations.

2.1.2.5 Simulation Parameters

In order to give designs the greatest amount of freedom, a number of assumptions were made. These assumptions were based on previous simulations, past experience, routing studies, and feedback from some vendors.

For the add-in card, routing studies have shown that a board can be routed with the following requirements:

- The maximum trace length of 3.0 inches.
- A line length skew of the data lines to be ±0.5 inches with respect to their associated strobe.
- A trace spacing of 1:2, that is a 6 mil trace with 12 mil spacing (air gap) or a 5 mil trace with 10 mil spacing.

An additional assumption for the add-in card is that the board can be with an impedance range of 50 Ω to 85 Ω. These are thought to include the range of the most popular range of specified printed circuit boards (PCB).

For the motherboard, routing studies have shown that several different routing requirements are needed. For some designs, the trace spacing needs to be 1:1. Other designs need longer line lengths than can be achieved with a 1:1 spacing. These designs need to be made with wider spacing except where necessary to escape a component.

To allow for these routing needs, two motherboards each with a different trace spacing were used. These were with a motherboard spacing of 1:1 and 1:2. These motherboards were analyzed to determine the maximum line length that can be achieved for these spacings. All three of these styles of motherboards need to allow at least a 0.5 inches variation between the data lines. Routing studies have shown that it is usually best to have the strobe longer than the data.

An additional assumption for the motherboard is that the board can be specified with an impedance range of 50 Ω to 85 Ω.

Also all motherboards will be able to run in both A.G.P. 1X mode and A.G.P. 2X mode.

2.1.2.6 Flight Time

A previous analysis had shown that the minimum flight time occurs when the add-in card data line length would be at its shortest length (0.0 inches). Also, the minimum flight time is achieved with the fast buffer. Simulations were run with the various combinations of interconnect for the data lines on the motherboard and connector while varying the motherboard line length.

A previous analysis had shown that the maximum flight time occurs when the add-in card data line length would be at its longest length (3.0 inches). Also, the maximum flight time is achieved with the slow buffer.

By using the simulations that were required for the A.G.P. 2X mode it would be possible to find out what the maximum flight time is for a line length up to 10.0 inches. By using these results, the maximum line length on the motherboard for A.G.P. 1X could be found.

Figure 31: Flight Time Measurement

Figure 31 shows how flight time is measured with 3.3 volt signaling. In this example $T_{FLIGHT} = T_{VAL} – T_{CO}$ where T_{VAL} is the time that the signal crosses the threshold and T_{CO} is the time that a test load crosses the threshold. The same technique is used with 1.5 volt signaling.

2.1.2.7 Flight Time Skew

The flight time skew simulations need to simulate any parameter that could cause a skew between two signals. The flight time skew simulations included varying the motherboard and add-in card line lengths, the effective capacitance in the buffer models, and crosstalk on each of the different interconnect combinations. They also included the effect of settling time (data pattern dependence) has on the flight time skew by comparing settled signals to unsettled signals.

For the source synchronous signals, the various parameters have a cumulative effect on the total flight time skew but do not add linearly. The maximum skew does not occur where each effect is maximized,

and the manner in which each effect adds or interacts is not easily predicted. This means that a large number of cross variations need to be looked at.

In all cases, the strobe simulations were compared to the data simulations. Also, all flight time skew simulations were run with the add-in card line length at both 2.5 inches and 3.0 inches. This was to simulate the worst case variations between data and strobe attributed to the add-in card. The motherboard line length was varied from 0.0 inches to 10.0 inches in 0.5 inch increments. This was to find the maximum line lengths as they relate to skew.

All of the strobe simulations were done using the different types of interconnect combinations for motherboard, connector, and add-in card. These values did not include the effect of crosstalk. These simulations are grouped by motherboard routing rules. The values used with these rules are shown Table 10. The add-in card and the motherboard use 1:1 and 1:2 spacing.

All of the data simulations were done using the different types of interconnect combinations for motherboard, connector, and add-in card. These values included the effect of crosstalk in the data simulations. These simulations are grouped by motherboard routing rules. The values used with these rules are shown Table 10. The add-in card used the 1:2 spacing and the motherboard used 1:1 and 1:2 spacing.

The capacitive loading for the data simulations was held at the nominal values. The fast buffer used 2 pF and the slow buffer used 8 pF.

The capacitive loading was varied for the strobe simulations. These were varied -1 pF and +2 pF of the nominal values for both the fast and the slow buffers. This fast buffer had a nominal loading of 2 pF, so the strobe was simulated with an effective capacitance of 1 pF and 4 pF when the fast buffer was used. The slow buffer had a nominal loading of 8 pF, so the strobe was simulated with an effective capacitance of 7 pF and 10 pF when the slow buffer was used.

The strobe simulations were run only at a 7 ns cycle time. This is due to the fact that the strobe will constantly be switching states when transferring data. The number of cycles used was 4.

The data simulations were run varying the number of cycles depending on the cycle time. When the duration (half-cycle) was set to 7 ns, the number of cycles was set to 4. When the duration was set to 50 ns the number of cycles was set to 1. These simulations were used to calculate the minimum and maximum flight times for signals that are settled and for signals that are running at the maximum transitional speed. This found any effects attributable to data pattern dependency.

Besides finding the maximum accumulative flight time skew, the effect of each of the parameters needed to be analyzed. This was to find out what impact the different variables had on the total flight time skews.

2.1.2.8 Signal Quality

Settling time was not measured in these simulations but was taken into account with the flight time skew measurements.

Overshoot is measured at the highest voltage and undershoot is measured at the lowest voltage at the receiver. Ringback from the rising edge is the lowest voltage that a receiver comes down to after crossing Vddq. Ringback from the falling edge is the highest voltage that a receiver comes up to after crossing Vss. Figure 32 is an example that shows how overshoot, undershoot, and ringback is measured with 3.3 volt signaling. The same method is applied to 1.5 volt signaling.

Figure 32: Signal Quality Measurements

2.1.2.9 Summary of Simulations Required

A sensitivity analysis that was done early in the design phase contributed to the reduction of the number of simulations required at this stage in the analysis. This analysis used 3-D graphs with two variables along the X-axis and Y-axis. The results of the simulations is displayed in the Z-axis. Also, pull down menus were used so that a variety of parameters could be viewed. As an example, one of these 3-D graphs is shown in Figure 33.

Figure 33: Sensitivity Analysis Chart

The number of simulations can be pared down even more by joining the simulations that are required for each of the three factors that define the solution space. The main contributor to the simulations required is the number of source synchronous simulations. These are the simulations that define the maximum line lengths for A.G.P. 2X mode. For the common clock flight time simulations, an additional set of simulations is required to determine the minimum line lengths. All of the signal quality measurements can be taken from the source synchronous simulations. Table 11 is a table of these simulations. Each of these simulations needs to be done with all the variations of interconnects shown in Table 10.

Table 11: Summary of Simulations

Type	Interconnects	Driver	Loading	Duration	AC line lengths	MB line lengths
Data	1:2 spacing	Fast & slow buffers	Typical	7 ns & 50 ns	2.5 in & 3.0 in	0.0 in to 10.0 in
Data	1:1 spacing	Fast & slow buffers	Typical	7 ns & 50 ns	2.5 in & 3.0 in	0.0 in to 10.0 in
Data	1:1 spacing	Fast buffers	Typical	7 ns & 50 ns	0.0 in	0.0 in to 10.0 in
Strobe	typical	Fast & slow buffers	+2 pF & -1 pF	7 ns	2.5 in & 3.0 in	0.0 in to 10.0 in

2.1.3 Simulation Results

2.1.3.1 Flight Time

The results from the simulations were put into a spreadsheet for analysis. Data simulations were used for both the maximum and minimum flight times.

For the minimum flight time, the simulations used were those that had an add-in card line length of 0.0 inches. These used fast buffers and motherboard interconnects that were based on 1:1 spacing.

A spreadsheet was then used to calculate the flight times for each of these cases by subtracting the Tco from each flight time. The minimum values from these were found by taking the minimum value for a given motherboard line length and put into Table 12. These results show that the motherboard line lengths need to be greater than 1.0 inch.

Table 12: Minimum Flight Time

MB data line length (in)	0.0	0.5	1.0	1.5	2.0	2.5	3.0	3.5	4.0	4.5	5.0	5.5	6.0	6.5	7.0	7.5	8.0	8.5	9.0
Min. data flight time (ps)	-99	-50	20	90	140	190	230	300	350	410	490	530	590	660	690	760	820	890	940

For the maximum flight time, the simulations used add-in card line lengths of 3.0 inches. These used slow buffers and motherboard interconnect spacing of 1:2.

A spreadsheet was then used to calculate the flight times for each of these cases by subtracting the Tco from each flight time. The maximum values found for a given motherboard line length were put into Table 13. These results show that the motherboard line lengths need to be less than 10.0 inches.

Table 13: Maximum Flight Time

MB data line length (in)	1.0	1.5	2.0	2.5	3.0	3.5	4.0	4.5	5.0	5.5	6.0	6.5	7.0	7.5	8.0	8.5	9.0	9.5	10.0
Max. data flight time (ns)	0.82	0.91	1.00	1.10	1.20	1.30	1.40	1.49	1.59	1.68	1.77	1.86	1.95	2.05	2.14	2.23	2.32	2.41	2.50

2.1.3.2 Flight Time Skew

The results from the source synchronous simulations were put into a spreadsheet for analysis. This spreadsheet compares the flight times from the strobe simulations to the data simulations.

Each case from all of the different strobe simulations was compared to the cases from several different data simulations. These cases are the minimum and maximum flight times for the rising and falling edges for each simulation. A summary of the comparisons that were done for each line length on the motherboard for a given line spacing is shown in Table 14.

Table 14: Summary of Source Synchronous Comparisons

Variable	Strobe	Data	Comments
Buffer location	Master / Target	Master / Target	compares use same location
Buffer type	fast / slow	fast / slow	compares use same type
Loading	-1 pF / + 2 pF	typical	compares all cases (2)
Cycle time	7 ns	7 ns / 50 ns	compares all cases (2)
Corner	rise / fall	rise / fall	compares all cases (4)
Flight time	min / max	min / max	compares all cases (4)
Fast motherboard	fast	fast odd / fast even	compares all cases (2)
Slow motherboard	slow	slow odd / slow even	compares all cases (2)
Connector	fast / slow	fast odd / slow even	compares all cases (2)
Fast add-in card	fast	fast odd / fast even	compares all cases (2)
Slow add-in card	slow	slow odd / slow even	compares all cases (2)
Add-in card line length	2.5 in	3.0 in	
Add-in card line length	3.0 in	2.5 in	

The results of all of these comparisons can be seen in Table 15 and Table 16.

Table 15 is a table of the maximum and minimum skews that resulted from a number of simulations with a motherboard spacing of 1:1. The top row is the data line lengths on the motherboard in inches. The next two rows are the setup skews in picoseconds between the strobe and the data lines. One of these is when the strobe is 0.5 inches less than the data lines; the other is when they are of equal length. The last two rows are the hold skews in picoseconds between the strobe and the data lines. One of these is when the strobe is 0.5 inches longer than the data lines, the other is when they are of equal length. The highlighted area is where the recommended trace lengths for motherboards were derived. The solution needs to be between -900 ps (hold skew limit) and 700 ps (setup skew limit).

Table 15: Skew From the Motherboard with 1:1 Trace Spacing

MB data line length (in)	1.0	1.5	2.0	2.5	3.0	3.5	4.0	4.5	5.0	5.5	6.0	6.5
Setup (ps) strobe -0.5 in	510	550	590	620	670	710	740	780	820	860	900	940
Setup (ps) strobe equal	410	460	510	550	600	630	570	700	750	790	820	860
Hold (ps) strobe equal	-470	-530	-570	-630	-690	-720	-770	-800	-820	-810	-850	-880
Hold (ps) strobe +0.5 in	-560	-610	-660	-730	-760	-800	-850	-860	-880	-900	-940	-960

Table 16 is a table of the maximum and minimum skews that resulted from a number of simulations with a motherboard spacing of 1:2. The top row is the data line lengths on the motherboard in inches. The next two rows are the setup skews in picoseconds between the strobe and the data lines. One of these is when the strobe is 0.5 inches less than the data lines; the other is when they are of equal length. The last two rows are the hold skews in picoseconds between the strobe and the data lines. One of these is when the strobe is 0.5 inches longer than the data lines; the other is when they are of equal length. The

highlighted area is where the recommended trace lengths for motherboards came from. The solution needs to be between -900 ps (hold skew limit) and 700 ps (setup skew limit).

Table 16: Skew from Motherboard with 1:2 Trace Spacing

MB data line length (in)	1.0	1.5	2.0	2.5	3.0	3.5	4.0	4.5	5.0	5.5	6.0	6.5	7.0	7.5	8.0	8.5	9.0	9.5	10.0
Setup (ps) strobe -0.5in	450	480	500	550	600	600	620	610	590	610	590	590	610	640	670	690	710	750	780
Setup (ps) strobe equal	360	390	440	490	530	530	520	520	510	520	500	520	550	590	620	640	660	690	720
Hold (ps) strobe equal	-470	-510	-540	-650	-710	-730	-760	-730	-740	-700	-700	-710	-720	-740	-760	-770	-800	-820	-840
Hold (ps) strobe +0.5in	-550	-590	-670	-750	-780	-810	-810	-810	-780	-770	-790	-790	-810	-820	-840	-860	-880	-900	

2.1.3.3 Signal Quality

The analysis of signal quality was done by finding the minimum and maximum values from the different parameters with fast buffers. The greatest amount of overshoot was 5.25 V and the greatest amount of undershoot was -2.05 V. The worst case ringback from the rising edge was 2.62 V and the worst case ringback from the falling edge was 0.90 V. These values are for 3.3 volt signaling.

2.2 4X Mode Board Design Guidelines

2.2.1 4X Mode Routing Rules

These guidelines are primarily for A.G.P. 4X transfer mode (4X mode) designs that use a 4X mode compliant graphics controller on an 4X mode compliant add-in card that plugs into a 4X capable motherboard with a 4X compliant system controller. Designs with both 4X compliant components on the motherboard can also use these techniques.

Add-in card designs must follow length rules that are specified in the *A.G.P. Interface Specification, Revision 2.0*. These rules are reiterated in Table 19.

For motherboard designs, there are three ways that these guidelines can be used:

- Designers who do not have the time or resources to complete a full simulation of their design can route their layout following the line length rules in Table 18, Table 19, and Table 20. These rules preserve enough timing margin to ensure a working design without requiring further simulations. The rules are flexible enough to be used on most designs. Care must be taken to follow these rules precisely, if no further simulations are to be used.
- Designers wishing to create their own set of line length rules can use the information included in Section 2.1.3. This might be desired in order to take advantage of margin in other parts of a design (for example, lines routed shorter than those assumed by this guideline might not need to be matched as tightly). This can be done by reviewing the results that are shown in Figure 40, Figure 43, Figure 44, Figure 45, and Figure 46 and applying them to your own design.
- Designers who wish to base their designs on their own simulations can follow Section 2.1.2 as a pattern for additional simulations, which may then be used for other designs. This section describes how the analysis was done to make these guidelines. This will yield the most flexible designs; however, it requires extensive simulations.

Sections 2.2.4 and 2.2.5 give additional design information that can be used with all three methods of designing a 4X mode compliant motherboard and add-in card.

2.2.1.1 A.G.P. 4X Mode Design Consideration

Scaling the A.G.P. interconnect from 2X transfer mode (2X mode) to 4X mode presents some interesting challenges. These challenges include a two-fold increase in bus speed and reducing the amount of interconnect skew by 1/3.

Reductions are required in several of the parameters that make up the total skew budget. An example of this is illustrated in Figure 34. Line lengths require closer matching in 4X mode than in 2X mode. The loading variation must be reduced in 4X mode. The crosstalk contribution to skew must be reduced by the greatest amount in order to avoid excess skew due to increased inter-symbol interference (ISI). This means that the separation distance between 2x and 4x printed circuit boards.

Though not shown in Figure 34, Inter-Symbol Interference (ISI) is much worse in 4X mode when using the same interconnect parameters as were used in 2X mode, because the shorter cycle time makes these signals more sensitive to oscillations and ringing.

Figure 34: Example of Skew Differences Between 2X Mode and 4X Mode

Throughout Section 2.2 of this document, the term “data” refers to the **AD[31::0]**, **C/BE[3::0]#**, and **SAB[7::0]** signals. The term “strobe” refers to **AD_STB[1::0]** and **SB_STB**. The term “strobe#” refers to **AD_STB[1::0]#** and **SB_STB#**. The term data refers to one of three groups of data as seen in Table 17. The term strobe refers to one of the three strobes as it relates to the data in its associated group. The term strobe# refers to one of the three strobes# as it relates to the data in its associated group.

Table 17: 4X Mode Data and Associated Strobe / Strobe#

Data	Associated Strobe	Associated Strobe#
AD[15::0] and C/BE[1::0]#	AD_STB0	AD_STB0#
AD[31::16] and C/BE[3::2]#	AD_STB1	AD_STB1#
SBA[7::0]	SB_STB	SB_STB#

2.2.1.2 Source Synchronous Rules for Motherboards

Table 18 gives the routing rules for motherboards. Boards laid out following these rules will work robustly, requiring minimum additional simulations.

Table 18: Motherboard Recommendations

Width:SpaceMinimum[1, 2]	Z0	Trace	Line Length[3]	Line Length Matching
1:4	54 Ω to 66 Ω	Strobe to Strobe#	1.0 in line length ≤ 7.25 in	< 0.1 in
1:6	54 Ω to 66 Ω	Strobe to Data		
1:4	54 Ω to 66 Ω	Data	1.0 in line length ≤ 7.25 in	Strobe and Strobe# ± 0.125 in of group
1:3	54 Ω to 66 Ω	Data	1.0 in line length ≤ 6.0 in	Strobe and Strobe# ± 0.5 in of group

Notes:

1. These ratios assume a particular board stackup. For more information, refer to Section 2.2.4.1.
2. For different stackups, scale the trace width with the dielectric thickness. For more information refer to Section 2.2.1.6.
3. Line lengths are based on microstrips (external traces). Striplines (internal traces) will require a 25% reduction in line lengths. This is due to a 25% reduction in propagation velocity.

The recommended motherboard impedance range is 60 Ω ± 10%. This range plus the width:space ratios in Table 18 allow sufficient resistance to crosstalk effects. The stackup shown in Section 2.2.4.1 makes the required coupled impedance range of 48 Ω to 73 Ω possible. This range is used to cover design targets and manufacturing tolerances. All lines should be at least 1.0 inch in length to ensure conformance to the A.G.P. 2X design guidelines. The maximum line lengths are dependent on the type of trace and the amount of coupling.

The maximum line length is dependent on the routing rules used on the motherboard. These routing rules are intended to give freedom for designs by making tradeoffs between signal coupling (trace spacing) and line lengths. The routing rules are divided by trace spacing. 1:3 spacing refers to the distance between the traces (air gap) as being three times the width of the trace. 1:4 spacing refers to the distance between the traces as being four times the width of the trace.

In all cases, it is best to reduce the line length mismatch wherever possible to insure added margin. It is also best to separate the traces by as much as possible to reduce the amount of trace to trace coupling.

2.2.1.3 Source Synchronous Rules for Add-in Cards

Table 19 gives the routing rules for add-in cards.

Table 19: Add-in Card Rules

Width:Space Minimum[1, 2]	Z0	Trace	Line Length[3]	Line Length Matching
1:3	54 Ω to 66 Ω	Strobe to Strobe#		< 0.1 in
1:6	54 Ω to 66 Ω	Strobe to Data		
1:3	54 Ω to 66 Ω	Data	≤ 1.5 in	Strobe and Strobe# ± 0.25 in of group

Notes:

1. These ratios assume a particular board stackup. For more information, refer to Section 2.2.4.1.
2. For different stackups, scale the trace width with the dielectric thickness. For more information, refer to Section 2.2.1.6.
3. Line lengths are based on microstrips (external traces). Striplines (internal traces) will require a 25% reduction in line lengths. This is due to a 25% reduction in propagation velocity.

The required add-in card impedance range is 60 Ω ± 10%, therefore, it will be necessary to use a controlled impedance board. This range is used to cover design targets and manufacturing tolerances.

The routing rules are divided by trace spacing. 1:3 spacing refers to the distance between the traces (air gap) as being three times the width of the trace. 1:4 spacing refers to the distance between the traces as being four times the width of the trace.

In all cases, it is best to reduce the line length mismatch wherever possible to insure added margin. It is also best to separate the traces by as much as possible to reduce the amount of trace to trace coupling.

2.2.1.4 Control Signal and Clock Recommendations

Table 20 gives the routing rules for the control signals and clock lines.

Table 20: Control Signal Line Length Recommendations

Width:Space Minimum[1, 2]	Board	Trace	Line Length	Pull-up / Pull-down Stub Length
1:3	Motherboard	Control Signals	0.0 < line length < 7.5in	< 0.1in
1:4	Motherboard	Clock		
1:3	Add-in Card	Control Signals	0.0 < line length < 3.0 in	
1:4	Add-in Card	Clock	4.0 in ± 0.25 in	

Notes:

1. These ratios assume a particular board stackup. For more information, refer to Section 2.2.4.1.
2. For different stackups scale the trace width with the dielectric thickness. For more information, refer to Section 2.2.1.6.

Some of the control signals require pull-up or pull-down resistors. These may be placed anywhere along the trace, but must be routed so they do not introduce any stub greater than 0.1 inch. To meet this requirement, it is strongly recommended that discrete pull-ups be used.

The clock lines on the motherboard should be simulated to determine their proper line length. The motherboard must be designed to the type of clock driver that is being used. In all cases, it is recommended to have a trace spacing that is at least 4 times the trace width.

2.2.1.5 Inter-Symbol Interference

Inter-Symbol Interference (ISI) affects signals when the driver does not reach steady state prior to the next switching cycle. This causes the flight times to vary from one transition to another. This variability will increase the amount of skew between the strobe and data signals.

Figure 35 is an example of ISI. This case shows a skew of 150 ps due a data signal transitioning with a cycle time of 7.5 ns and another at 22.5 ns. This example shows how flight time can change depending on bit pattern. The narrow pulse (representing a 1010… bit pattern) has a delay which is noticeably different from the wide pulse (representing a 111000… bit pattern). This difference of 150 ps must be included as skew.

Figure 35: Example of ISI

ISI is caused by signal ringing; therefore, some methods that are used to reduce the amount of ISI are active receiver clamps, termination, and driver impedance matching. Clamps are difficult to implement because of the low signaling levels used. There is insufficient voltage to develop good bias level for “early clamps” and other types of feedback clamps affect the flight time and may introduce too much skew.

Terminations add cost to the printed circuit boards (PCB) and parts. The cost adder to the boards would be due to routing that would in all likelihood increase the layer count. They are also not backward compatible to the A.G.P. 2X signaling scheme.

Driver impedance matching works quite well if the buffer impedance is matched exactly to the interconnect impedance. This results in the signal at the receiver going to the rail without any overshoot or reflections.

This ideal circuit is impossible to implement because of manufacturing tolerances and impedance changes in the interconnect due to crosstalk. However, reductions in the range that impedance varies in the buffer and the interconnect are possible and can be used to reduce timing skew.

The buffer's slew rate can also be a large factor of ISI. This happens when the signals round trip propagation delay plus the total rise time (slew rate) is greater than the period. This is due to the interference of the returning wave.

The following sections describe interconnects and buffer design techniques for reducing ISI and other timing skew sources.

2.2.1.6 Board Constraints

The two main factors that contribute to PCB impedance variations are the manufacturing tolerances and crosstalk. The range over which these factors vary must to be reduced to decrease the total impedance variation.

Reducing the board tolerances are achieved by specifying the board and layer stackup to board vendors. For A.G.P. 2X mode, this range was 50 Ω to 85 Ω. For 4X mode, the range is reduced to 60 Ω ±10%.

Reducing the amount of coupling between driven lines will reduce the impedance variation due to crosstalk. Sections 2.2.2.4 and 2.2.2.5 describe what these values are and how they were derived. In a nutshell, crosstalk can be reduced by decreasing the separation distance of a trace to the nearest reference plane (dielectric thickness) and/or increasing the separation distance to adjacent traces (trace spacing).

Throughout Section 2.2, the spacing between traces is specified as a ratio of the trace width (width:space) which is a commonly used method to specify a board layout. This ratio is correct for the recommended board stackup, but the desired result is a ratio of trace spacing to dielectric thickness. Appendix A contains a more in depth explanation of crosstalk and how to include it in simulations.

All 4X mode signals, along with the clocks and strobes are not to cross any plane splits. (These might occur if a plane is used to supply two separate Vcc voltages.) If the common clock signals cross split planes, they should be simulated to insure that signal quality is not compromised. For more information on the effect that split planes have on the traces that cross them, refer to a paper *Signal Integrity Issues at Split Ground and Power Planes* presented by Haw-Jyh Liaw and Henri Merkelo at the 46th Electronic Components & Technology Conference 1996 (IEEE Catalog No. 96CH35931).

2.2.2 Simulation Techniques

This section summarizes the interconnect simulations that were done to generate the routing guidelines for the A.G.P. (266 MT/s) bus. The result of these simulations shows that designs using buffers defined in the *A.G.P. Interface Specification, Revision 2.0* can be done. Figure 36 is a diagram of the topology used in this analysis.

Figure 36: A.G.P. Source Synchronous Topology

The source synchronous simulation results limit the total line length to 7.25 inches for the simulated cases. Refer to Section 2.2.3.4 for a description of these cases. This line length is dependent on coupling between lines and line length mismatch.

Source synchronous signaling requires limits on the amount of variation between the strobe/strobe# line length and the data lines for a signal group. This variation is recommended to be no more than ±0.125 inch from any data within the group to either of its strobes (strobe, strobe#).

The signal quality results show overshoot and undershoot values approaching 0.7 volts. Also, the ringback levels are well within the *A.G.P. Interface Specification, Revision 2.0* requirements.

2.2.2.1 Methodology

Three main factors define the range of trace lengths useable for A.G.P. 4X: common clock flight times, source synchronous flight time skews, and signal quality. As shown in Figure 37, common clock flight times limit the length to a minimum of 1.0 inches. The maximum line lengths are restricted by flight time skews.

Figure 37: Solution Space Plot

Figure 37 is actually a "solution space" plot formed by combining all of the simulation results into one big "pass/fail" test. This is a pass/fail plot of the simulation results compared to the limits (flight time, flight time skew and signal quality) that are included in the *A.G.P. Interface Specification, Revision, 2.0.* Some of the data that made up this plot is shown in Appendix C as 3D plots. Appendix C also includes data that was not used for this solution.

Common clock (A.G.P. 1X mode) flight times use worst case (extreme) buffers and interconnects to determine the solution space. A fast buffer driving the shortest (electrically) interconnects into a buffer with the smallest amount of loading produces minimum flight times. The physical line length is varied until the minimum flight time is met. A slow buffer driving the electrically longest interconnects into a buffer with the largest amount of loading produces maximum flight times. The physical line length is then varied until the maximum flight time is met.

Source synchronous (A.G.P. 2X mode and A.G.P. 4X mode) flight time skews are found by comparing a strobe flight time with a data flight time for a particular interconnect, buffer, and loading. Setup skew occurs when the strobe's minimum flight time is compared to the data maximum flight time at a given line length. Hold skew occurs when the strobe maximum flight time is compared to the data minimum flight time at a given line length. The results of these two comparisons must be less than the total allowable skew.

Signals involved in source synchronous transfers must meet line length restrictions set by minimum and maximum flight times as well as flight time skew. The other "outer loop" common clock signals are limited only by the minimum and maximum flight times. The signal quality of the signals within the solution space may constrain the solution space even further.

Of the three factors, source synchronous flight time skews are the most difficult to model. This is partly because they are a result of differences in several interacting parameters. These parameters include ISI, crosstalk, loading, impedance variations, and line length variations.

2.2.2.2 Simulation Parameters

In order to give designs the greatest amount of freedom, a number of assumptions were made. These assumptions were based on our previous simulations, past experience, routing studies, and feedback from some vendors.

For the add-in card, routing studies have shown that a board can be routed with the following requirements. The maximum trace length is 1.5 inches. Data lines must meet a line length skew of ±0.25 inches with respect to their associated strobes. A trace spacing must be at least 1:3, that is, a 5 mil trace with 15 mil spacing (air gap) or a 6 mil trace with 18 mil spacing is required.

The add-in card impedance is specified as 60 Ω ±10%. In order to achieve the 60 Ω board requirement, it is recommended that the PCB stack use number 2116 prepreg (nominal 4.5 mil thickness) and the routing be entirely microstrip (outer layers). Signals within a strobe/data group should be routed on the same layer and should not cross breaks in the power/ground plane over which they are routed. Routing all signals over the ground plane is desirable to maintain a cleaner and consistent signal return path.

For the motherboard, routing studies have shown that there need to be several different requirements. For designs that require long line lengths (up to 7.25 inches), a 1:4 spacing is needed. For designs that use shorter line lengths (< 6.0 inches), the amount of trace spacing can be reduced. Simulations should be run to insure that the design will meet all of the specified requirements.

To get motherboard line lengths of 7.25 inches, a 0.125 inch line length mismatch between the strobe and the data is required. For shorter line lengths, this variation can be relaxed, but simulations should be run to determine by how much.

It is highly recommended that the motherboard impedance be specified as 60 Ω ±10%. In order to achieve the 60 Ω board, it is recommended that the PCB stack use number 2116 prepreg and the routing be entirely microstrip. As with the add-in cards, signals within a strobe/data group should be routed on the same layer and should not cross breaks in the power/ground plane over which they are routed. Routing all signals over the ground plane is desirable to maintain a cleaner and consistent signal return path.

It is necessary to ensure that all motherboards will be able to run in all transfer modes supported by the system electronics. However, it is expected that an interconnect system designed to support A.G.P. 4X will also support A.G.P. 2X and A.G.P. 1X.

2.2.2.3 Buffers

The models for both the Master and the Target are identical. Both fast and slow buffer models were used in the simulations. Example V-I characteristic curves can be found in Appendix B. In all cases, when a fast buffer was used in our simulations, it was used as both master and slave. The same was also true for the slow buffers.

The buffer models were created from the *A.G.P. Interface Specification, Revision 2.0.* These models also used power supply diodes. The package was modeled with an effective inductance of 10.0 nH.

The models also include an effective capacitance. This capacitance can vary between 3.0 pF and 8.0 pF. For the signal quality simulations, capacitance value of 3.0 pF was used. For the 4X mode simulations, a value of 8.0 pF was used for both the strobe and the data. For flight time simulations, the loading used was 3.0 pF for minimum flight times and 8.0 pF for maximum flight times.

2.2.2.4 Interconnect

Both the motherboard and the add-in card were modeled with a 54 Ω and 66 Ω typical transmission line. The 54 Ω boards had a velocity of 1.71 ns/ft and the 66 Ω had a velocity of 1.76 ns/ft. The connector was modeled with a 150 ps delay with two different impedance values, 35 Ω (slow) and 80 Ω (fast).

The strobe and strobe# used interconnect impedance values that were based on an odd mode coupling between them. They were also affected by the coupling of the nearest traces. With the spacing between the strobes (strobe and strobe#) being less than the nearest traces, the coupling effect is greater between the strobe and strobe#. The nearest data traces could be either odd mode coupled or even mode coupled to the strobes. Odd-odd mode refers to the coupling when the strobe or strobe# is modeled with the nearest trace in odd mode. Odd-even mode refers to the coupling when the strobe or strobe# is modeled with the nearest trace in even mode.

The data lines were modeled using both odd mode coupling and even mode coupling. Odd mode crosstalk is when the adjacent lines are switching in opposite directions to the trace under analysis. Even mode crosstalk is when the adjacent lines are switching with the trace under analysis.

The stackup used for this analysis used a number 2116 prepreg material that has a 4.5 mil typical thickness. The dimensions used for the typical boards (54 Ω and 66 Ω) are shown in Table 21. The values used were based on microstrips in a 60 Ω ±10% printed circuit board (PCB).

Table 21: Board Parameters

Symbol	Parameter	Min	Max	Units	Notes
Er	Dielectric constant	4.3	4.7		1, 5
h	Dielectric thickness	4.1	5.1	mil	2, 5
t	Trace thickness	1.5	2.8	mil	3, 5
w	Trace width	4.8	5.2	mil	4, 5

Notes:

1. Er will probably go as low as 3.9, but at these lower values, the effect will be faster propagation velocity and less variation in the propagation velocity. The effect on impedance is minimal.
2. The maximum thickness of 2116 prepreg is ≈ 5.0 mil.
3. 1 ounce copper has a thickness of ≈ 1.4 mil. The variation comes from plating.
4. Trace width variations are used to control the impedance by PCB vendors. These were varied to get the typical board impedance.
5. Typical 66 Ω board used Er = 4.3, h = 5.1 mil, t = 1.5 mil, and w = 5.2 mil.
   Typical 54 Ω board used Er = 4.7, h = 4.1 mil, t = 2.8 mil, and w = 4.8 mil.

Other stackups can be used as long as the combination of the stackup and the trace layout maintains coupled impedance values are 48 Ω to 73 Ω. Coupled impedance values that exceed 73 Ω may experience ringback failures.

2.2.2.5 Crosstalk

Crosstalk was factored into the simulations using the methodology described in Appendix A. The uncoupled board parameters used were taken from Table 21.

The results for data line type of coupling are shown in Table 22. Data line routing rules in this analysis used 5 mil traces with spacing to the nearest traces of 15 mil and 20 mil.

Table 22: Crosstalk parameters for 5 mil data lines.

Trace spacing	Board Z_0	Coupled mode	Coupled Z_0	Coupled S_0	Notes
20 mil	54 Ω	odd	50 Ω	1.69 ns/ft	1
		even	57 Ω	1.83 ns/ft	
	66 Ω	odd	63 Ω	1.64 ns/ft	
		even	71 Ω	1.78 ns/ft	
15 mil	54 Ω	odd	48 Ω	1.65 ns/ft	2
		even	58 Ω	1.87 ns/ft	
	66 Ω	odd	60 Ω	1.61 ns/ft	
		even	73 Ω	1.81 ns/ft	

Notes:

1. Required for motherboards to reach maximum line lengths.
2. May be used for shorter line lengths.
3. "Odd mode" and "even mode" refer to the effective impedance realized by a "victim" line at the center of a group of three lines. Odd mode ⇒ center line is switching in the opposite direction of the other two lines: even mode ⇒ all three lines are switching in the same direction.

The results for strobe line type of coupling are shown in Table 23. The routing for the strobes takes advantage of the strobe and strobe# being in odd mode to each other. The effect of the nearest lines has a lesser amount of coupling effect on the strobes. In the two cases shown, the strobes have either a 15 mil or 20 mil spacing between them. The nearest line in both cases is 30 mil spacing.

Table 23: Crosstalk Parameters for Strobe Lines

Trace spacing	Board Z_0	Coupled mode (strobe/data)	Coupled Z_0	Coupled S_0	Notes
20 mil / 30 mil	54 Ω	odd /odd	51 Ω	1.71 ns/ft	1, 2
		odd / even	52 Ω	1.75 ns/ft	
	66 Ω	odd / odd	64 Ω	1.66 ns/ft	
		odd / even	66 Ω	1.70 ns/ft	
15 mil / 30 mil	54 Ω	odd / odd	50 Ω	1.69 ns/ft	1, 3
		odd / even	51 Ω	1.73 ns/ft	
	66 Ω	odd / odd	62 Ω	1.65 ns/ft	
		odd / even	65 Ω	1.69 ns/ft	

Notes:

1. The first mode in the coupled mode column is the strobe to strobe# coupling. The second mode is the strobe to any other trace coupling.
2. 20 mil spacing from strobe to strobe# and 30 mil spacing to any other trace. Motherboard parameters.
3. 15 mil spacing from strobe to strobe# and 30 mil spacing to any other trace. Add-in card parameters.

2.2.2.6 Flight Time

All flight time simulations were run using 32 different combinations of data signal interconnects. These include four velocity and impedance combinations of motherboard, four velocity and impedance combinations of add-in card, and two impedance combinations of connectors. The four combinations of board interconnect are an uncoupled 54 Ω board in both odd and even mode and an uncoupled 66 Ω board in both odd and even mode. The two combinations of connectors are 35 Ω (odd mode) and 80 Ω (even mode).

Analysis of the minimum flight time showed that this occurs when the add-in card data line length is at the shortest length (0.0 inches) using a “fast” buffer model and 3 pF loading. For these simulations, both the 30 Ω buffer with a curvature of 1.1 and the 38 Ω buffer with a curvature of 2.0 were used. Refer to Section 1.2.2.2 for an explanation of buffer curvature.

Minimum flight time simulations were run with the 32 combinations of interconnect while varying the motherboard line length from 0.0 inches to 2.5 inches in 0.25 inch increments.

Analysis of the maximum flight time showed that this occurs when the add-in card data line length is at the longest length (1.5 inches) using the “slow” buffer model and 8 pF loading. The 48 Ω buffer with a curvature of 2.0 was used for these simulations.

Maximum flight time simulations were run with the 32 combinations of interconnect while varying the motherboard line length from 5.0 inches to 7.5 inches in 0.25 inch increments.

Figure 38: Reference Load

Additional simulations were run to find the reference delay for the buffer models that were used. This value is used to calculate the common clock flight times. (Flight time is defined to be the difference between the delay simulated into the interconnect and the delay simulated into this reference load.) The reference delay is the delay of the driver when driving a reference transmission line load (sometimes referred to as a test load). This is found by driving a 12 inch reference transmission line with an impedance of 60 Ω and a propagation velocity of 1.75 ns/ft. This is pictured in Figure 38. V_{tOUT} is measured by recording the time that the voltage at the open end of the reference transmission line reaches 0.75 V. The reference delay is found by subtracting 1.75 ns from V_{tOUT}. Refer to the A.G.P. specification Section 4.2.3.3.

All simulations used 1.5 V for Voh and 0.0V for Vol. Also, a threshold value of 0.75 V (0.5 Vddq) was used for all the flight time simulations. Flight time was found by subtracting the reference delay from the time that the signal reached the receiver's threshold.

2.2.2.7 Flight Time Skew

The flight time skew simulations need to simulate any parameter that could cause a skew between two signals. The flight time skew simulations included varying the motherboard and add-in card line lengths, the effective capacitance in the buffer models, and crosstalk on each of the different interconnect combinations. They also included data pattern dependencies, a major contributor to ISI induced skews.

For the source synchronous signals, the various parameters have a cumulative effect on the total flight time skew but do not add linearly. The maximum skew does not occur where each effect is maximized, and the manner in which each effect adds or interacts is not easily predicted. This means that a large number of cross variations need to be looked at.

In all cases, the strobe simulations were compared to the data simulations. Also, all flight time skew simulations were run with the add-in card line length at both 1.25 inches and 1.5 inches. This was done to simulate the worst case strobe-data mismatches contributed by the add-in card. The motherboard line length was varied from 5.0 inches to 7.5 inches in 0.25 inch increments. This was done to find the maximum line lengths as they relate to skew.

All of the strobe simulations were done using the different types of interconnect combinations for the motherboard (54 Ω and 66 Ω uncoupled board), the connector, and the add-in card (54 Ω and 66 Ω uncoupled board). The motherboard interconnects used coupled parameters for both odd-odd mode and odd-even mode using 1:4 trace:space between the strobes and 1:6 trace:space between the strobes and the nearest trace. The add-in card used coupled parameters for both odd-odd mode and odd-even mode using 1:3 trace:space between the strobes and 1:6 trace:space between the strobes and the nearest trace. The connector used 35 Ω (odd mode) and 80 Ω (even mode).

All of the data simulations were done using the different types of interconnect combinations for the motherboard (54 Ω and 66 Ω uncoupled board), the connector, and the add-in card (54 Ω and 66 Ω uncoupled board). These values are to include the effect of even and odd mode coupling in the simulations. These simulations are grouped by motherboard routing rules. The values used with these

rules are shown in Table 22 and Table 23. The add-in card uses the 1:3 spacing and the motherboard uses 1:4 spacing.

The capacitive loading for both the data and strobe simulations used 8 pF.

The strobe simulations were run using a 7.5 ns period with a 50% duty cycle time. This is due to the fact that the strobe will constantly be switching states when transferring data. Four cycles were run.

The data simulations were run varying the number of cycles depending on the cycle time. When the duration (half-cycle) was set to 3.75 ns, the number of cycles was set to 4. When the duration was set to 7.5 ns, the number of cycles was set to 2. With the simulator that was used for this analysis, the starting of the first cycle represented a case when the signals were settled. This is due to the symmetry of the buffers and the fact that the initial starting point was set to Vss. The duration was not set to 11.25 ns as shown in Figure 35. Showing an example of ISI with a 11.25 ns duration made the difference clearer because the first rising edge was the same for both simulations.

2.2.2.8 Signal Quality

Signal quality was measured using a different set of simulations, but it is possible to use the results of the flight time and skew simulations to find the worst case results. Overshoot is measured at the highest voltage and undershoot is measured at the lowest voltage at the receiver. Ringback from the rising edge is the lowest voltage that a receiver comes down to after crossing the threshold at Vddq (1.5 V). Ringback from the falling edge is the highest voltage that a receiver comes up to after crossing the threshold at Vss (0.0 V). Figure 39 is an example that shows how overshoot, undershoot, and ringback are measured.

Figure 39: Signal Quality Measurements

2.2.3 Simulation Results

2.2.3.1 Signal Quality Simulations

The signal quality simulations were performed on each of the 32 variations of data interconnects. They were also done with an add-in card line length of 1.5 inches and a connector line length of 0.9 inches. These simulations used a motherboard line length of 0.0 inches to 7.5 inches in 0.25 inch increments.

The buffer that found the largest amount of ringback was the 38 Ω buffer with a curvature of 2.0 and a slew rate of 2.8 V/ns. The amount used for component loading was 3 pF.

All of the simulations were run with two different cycle times to incorporate the effect of ISI on ringback. These cycle times were 15 ns and 30 ns. The results of these simulations are shown in Figure 40.

Figure 40: Ringback Results

2.2.3.2 Minimum Flight Time Simulations

The minimum flight time simulations were performed on each of the 32 variations of data interconnects. They were also done with an add-in card line length of 0.0 inches and a connector line length of 0.5 inches. These simulations used a motherboard line length of 0.0 inches to 2.5 inches in 0.25 inch increments.

The simulations used two different buffer types. One of the buffers was the 38 Ω buffer with a curvature of 2.0. The other buffer was the 30 Ω buffer with a curvature of 1.1. Both of these buffers were simulated with two different slew rates, 1.4 V/ns and 2.8 V/ns.

The component loading for all of these simulations was 3 pF. Also, the simulations used a 15 ns cycle time.

The results of these simulations show that there are no minimum line length requirements. These results are based on the technique that is described in Section 2.2.2.5. Figure 41 is a graph of the results.

Figure 41: Minimum Flight Time Results

2.2.3.3 Maximum Flight Time Simulations

The maximum flight time simulations were performed on each of the 32 variations of data interconnects. They were also done with an add-in card line length of 1.5 inches and a connector line length of 0.9 inches. These simulations used a motherboard line length of 5.0 inches to 7.5 inches in 0.25 inch increments.

The simulations used two different buffer types. One of the buffers was the 48 Ω buffer with a curvature of 2.0. The other buffer was the 48 Ω buffer with a curvature of 1.1. Both of these buffers were simulated with two different slew rates, 1.4 V/ns and 2.8 V/ns.

The component loading for all of these simulations was 8 pF. Also, the simulations used a 15 ns cycle time.

The results of these simulations show that there are no minimum line length requirements. These results are based on the technique that is described in Section 2.2.2.6. Figure 42 is a graph of the results.

Figure 42: Maximum Flight Time Results

2.2.3.4 4X Transfer Mode Simulations

Analyzing skew requires many more simulations than flight time. This is due to the variations for the strobe interconnect that were compared to each of the data simulations.

A simple calculation of the number of simulations required for each of the interconnect parameters would be 32^2 (1024). These combinations are shown in Table 24. Each of these simulations would be repeated for each of the different buffer types (impedance, curvature, and slew rate) and interconnect line lengths. This is only partly true because there are some short cuts that can be taken to reduce this number.

Table 24: Example of the 1024 Combinations of Interconnect Parameters

Interconnect Parameter	Data		Strobe	
Motherboard Z0	54 Ω / 66 Ω	(2X)	54 Ω / 66 Ω	(2X)
Add-in card Z0	54 Ω / 66 Ω	(2X)	54 Ω / 66 Ω	(2X)
Connector Z0	35 Ω / 80 Ω	(2X)	35 Ω / 80 Ω	(2X)
Motherboard crosstalk mode	Even / odd	(2X)	Odd-even / Odd-odd	(2X)
Add-in card crosstalk mode	Even / odd	(2X)	Odd-even / Odd-odd	(2X)

There are several short cuts that can be taken to reduce the number of simulations. First of all, when the data uses typical impedance for either the motherboard or the add-in card, the same typical impedance is used for the strobe. Another way to reduce the number of simulations is to divide the hold skew simulations from the setup skew simulations. This will allow the hold skew simulations to use only odd mode coupling for the motherboard data and odd-even mode for the strobe. The setup skew simulations will use even mode coupling for the motherboard data and odd-odd mode coupling for the motherboard strobe. This method finds the worst cases with the fewest simulations.

Splitting the simulations up into hold skew and setup skew has added benefits by incorporating line length deltas in the simulations. The hold skew simulations will have the strobe simulated with the maximum amount of line length delta. The setup skew simulations will have the strobe simulated with the minimum amount of line length delta.

Table 26 is an example showing how the data and strobe simulations match up. This is an example of hold skew simulations that use a typical 54 Ω board for both the motherboard and add-in card. This format is repeated using typical 66 Ω boards. Table 26 also shows combinations using a 54 Ω motherboard with a 66 Ω add-in card and 66 Ω motherboard with a 54 Ω add-in card. This puts the number of combinations at 128 when the setup skew and the hold skew are separated. These are shown in Table 25.

Table 25: Reduced Set of Interconnect Parameters to be Simulated

Interconnect Parameter	Data	Strobe
Motherboard Z0	54 Ω / 66 Ω (2X)	Same as data
Add-in card Z0	54 Ω / 66 Ω (2X)	Same as data
Connector Z0	35 Ω / 80 Ω (2X)	35 Ω / 80 Ω (2X)
Motherboard crosstalk mode (setup skew)	Even	Odd-odd
Motherboard crosstalk mode (hold skew)	odd	Odd-even
Add-in card crosstalk mode	Even / odd (2X)	Odd-even / Odd-odd (2X)

The setup skew simulations use different coupling modes for the motherboard. The data lines use odd mode for the motherboard while the strobe uses odd / odd mode coupling.

All of these combinations were simulated at each of the motherboard line length and six different buffers. The buffers used are the 30 Ω with a curvature of 1.1, the 38 Ω with a curvature of 2.0, and the 48 Ω with a curvature of 2.0. All three of these buffers used slew rates of 1.4 V/ns and 2.8 V/ns. Appendix C shows the results of these simulations as 3D charts.

Table 26: Example of Simulation Combinations

DATA					STROBE				
Add-in Card		Connector	Motherboard		Add-in Card		Connector	Motherboard	
Typ bd	Mode		Typ bd	Mode	Typ bd	Mode		Typ bd	Mode
54 Ω	Even	35 Ω	54 Ω	Even	54 Ω	Odd / even	35 Ω	54 Ω	Odd / even
54 Ω	Even	35 Ω	54 Ω	Even	54 Ω	Odd / odd	35 Ω	54 Ω	Odd / even
54 Ω	Even	35 Ω	54 Ω	Even	54 Ω	Odd / even	80 Ω	54 Ω	Odd / even
54 Ω	Even	35 Ω	54 Ω	Even	54 Ω	Odd / odd	80 Ω	54 Ω	Odd / even
54 Ω	Odd	35 Ω	54 Ω	Even	54 Ω	Odd / even	35 Ω	54 Ω	Odd / even
54 Ω	Odd	35 Ω	54 Ω	Even	54 Ω	Odd / odd	35 Ω	54 Ω	Odd / even
54 Ω	Odd	35 Ω	54 Ω	Even	54 Ω	Odd / even	80 Ω	54 Ω	Odd / even
54 Ω	Odd	35 Ω	54 Ω	Even	54 Ω	Odd / odd	80 Ω	54 Ω	Odd / even
54 Ω	Even	80 Ω	54 Ω	Even	54 Ω	Odd / even	35 Ω	54 Ω	Odd / even
54 Ω	Even	80 Ω	54 Ω	Even	54 Ω	Odd / odd	35 Ω	54 Ω	Odd / even
54 Ω	Even	80 Ω	54 Ω	Even	54 Ω	Odd / even	80 Ω	54 Ω	Odd / even
54 Ω	Even	80 Ω	54 Ω	Even	54 Ω	Odd / odd	80 Ω	54 Ω	Odd / even
54 Ω	Odd	80 Ω	54 Ω	Even	54 Ω	Odd / even	35 Ω	54 Ω	Odd / even
54 Ω	Odd	80 Ω	54 Ω	Even	54 Ω	Odd / odd	35 Ω	54 Ω	Odd / even
54 Ω	Odd	80 Ω	54 Ω	Even	54 Ω	Odd / even	80 Ω	54 Ω	Odd / even
54 Ω	Odd	80 Ω	54 Ω	Even	54 Ω	Odd / odd	80 Ω	54 Ω	Odd / even

Figure 43: Setup Skew Results with 1:4 Trace Spacing

Figure 43 shows the results of setup skew simulations with a motherboard interconnect that used the parameters for 1:4 trace spacing. These simulations used all of the six buffer types. The setup skew simulations used an add-in card line length of 1.25 inches for the data simulations and 1.5 inches for the strobe simulations. The motherboard line lengths used are 5.0 inches to 7.5 inches in 0.25 inch increments for the data lines. The strobe line lengths are the same as the data line plus 0.125 inch.

Figure 44: Hold Skew Results with 1:4 Trace Spacing

Figure 44 shows the results of hold skew simulations with a motherboard interconnect that used the parameters for 1:4 trace spacing. These simulations used all of the six buffer types. These hold skew simulations used an add-in card line length of 1.5 inches for the data simulations and 1.25 inch for the strobe simulations. The motherboard line lengths used are 5.0 inches to 7.5 inches in 0.25 inch increments for the data lines. The strobe line lengths are the same as the data line minus 0.125 inch.

Figure 45: Setup Skew Results with 1:3 Trace Spacing

Figure 45 shows the results of setup skew simulations with a motherboard interconnect that used the parameters for 1:3 trace spacing. These simulations used all of the six buffer types. The setup skew simulations used an add-in card line length of 1.25 inches for the data simulations and 1.5 inch for the strobe simulations. The motherboard line lengths used are 1.0 inches to 5.0 inches in 0.25 inch increments for the data lines. The strobe line lengths are the same as the data line plus 0.5 inch

Figure 46: Hold Skew Results with 1:3 Trace Spacing

Figure 46 shows the results of hold skew simulations with a motherboard interconnect that used the parameters for 1:3 trace spacing. These simulations used all of the six buffer types. These hold skew simulations used an add-in card line length of 1.5 inches for the data simulations and 1.25 inch for the strobe simulations. The motherboard line lengths used are 1.0 inches to 5.0 inches in 0.25 inch increments for the data lines. The strobe line lengths are the same as the data line minus 0.5 inches.

2.2.4 Design Recommendations

2.2.4.1 Board Materials

The recommended stackup is microstrips on a PCB using a number 2116 prepreg material as shown in Figure 47. Number 2116 prepreg is a common dielectric material available to most PCB vendors. The thickness of this material is nominally 4.5 mils with a maximum thickness of less than 5.0 mils. It is made with a denser fiberglass weave that has less thickness variation in manufacturing giving better board impedance tolerances. The thinner dielectric decreases coupling to the signal lines relative to adjacent signals thereby reducing crosstalk effects which also helps reduce skew due to ISI. Both of these reductions are necessary in making a robust A.G.P. 4X design.

Other prepreg materials may be found that will produce coupled impedance values that are within 48 Ω to 73 Ω. As a general rule, a thicker prepreg material will require a greater amount of spacing between the traces.

Figure 47: Recommended PCB Stackup

2.2.4.2 Routing Measurement Points

With the tightened requirements for line lengths and line length matching, it is important to define the measurement points. The measurement point for both the upper and lower contact points are on the same horizontal line across the connector. Figure 48 shows where this line is located on the add-in card edge fingers. For the motherboard, the measurement point is at the connector pin.

The measurement points for the master and the target devices are the center of the pad that the ball (BGA package) or pin (QFP package) lands on.

Figure 48: Upper and Lower Contact Measurement Points

2.2.4.3 Escaping Components

Additional design freedom might be given to the board interconnects that use BGA packages. In BGA packages, the package trace lengths are generally longer to the outer row of balls then to the inner row of

balls. In these cases, it may be possible to bend the line length delta rule by giving the board interconnect that are routed to the outer row of balls an increase in delta.

When “bending” the line length delta rule, there are a couple of important precautions that need to be taken into consideration. One of these precautions is that the line length rule cannot be violated. Another precaution is that the electrical length of the package needs to be known and not assumed.

2.2.5 Measurements

This section gives some insight into how to correlate simulations to measurements. The importance of correlation is that ringback and skews are specified at the silicon (pad) where they cannot be physically measured.

2.2.5.1 Ringback

Ringback is not difficult to measure with simulations, but for validation it is. With physical measurements taken at the pin or very near the pin, there needs to be a method to extrapolate these measurements to the specification.

Through simulations, the motherboard designer should find the worst case ringback of the design. This worst case ringback is defined at the silicon pad using worst case package models. Using the same simulations, the point where physical measurements are to be taken needs to be recorded. The physical measurements can be compared to this simulation point.

2.2.5.2 Skew

Skew is the difference between two flight times as defined in the following formula: $t_{SKEW} = t_{fDATA} – t_{fSTROBE}$. t_{fDATA} is the flight time of the data and $t_{fSTROBE}$ is the flight time of the strobe.

Errors in the skew measurement can be caused by the reflections due to the package when taking flight time measurements at the receiving component pin. This ledge becomes wider when measured farther from the pad. An example of these reflections can be seen in Figure 49. This example shows that the timing that is being measured will vary if this ledge is either raised or lowered. This variation has the potential to artificially increase the skew to over 120 ps.

When making physical measurements with a ledge at or below 0.5 Vddq, it is necessary to make this measurement below the ledge. If one of the measurements is below 0.5 Vddq, the other needs to be made at this same voltage level.

Figure 49: Example of Ledge at the Receiver Pin

3. A.G.P. Thermal Design Guidelines

3.1 Thermal Overview

The general trend for high performance A.G.P.-compliant graphics controllers is toward integrated and higher frequency functions; for example, 2-D, 3-D, and RAMDAC functions, in a single component. This trend will enable low component cost with high performance and enable flexibility for motherboard or add-in card solutions. For competitive 3-D performance levels, the graphics controller power trend will increase with time. The forecasted graphics controller maximum power levels at A.G.P. introduction are 3-6 W. Future generation graphics controller power is expected to increase to 5-10 W for most controllers, with the potential of up to 15 W for some high performance controllers as shown in Figure 50. These power levels will create localized power densities that require thermal enhancement of component packages and add-in cards.

Figure 50: A.G.P. Add-In Card and Graphics Controller Power Trend

3.2 Thermal Design Recommendations

3.2.1 Recommendations for Add-In Card Designers

It must be understood that standard thermal specifications such as ambient temperature and air flow cannot be controlled in certain market segments, in particular, the retail add-in card market segment. In this case, the worst case thermal environment expected is 50-60 °C internal ambient temperature near the A.G.P. card with only natural convection air flow and any card orientation (vertical, horizontal-components up, horizontal-components down). For example, worst case add-in card thermal performance typically occurs when the card orientation is horizontal with components facing down.

Add-in card component packages are likely to require thermal enhancement to operate in this type of environment. The add-in card designer must determine the thermal environment of the target market segment and ensure that the add-in card is thermally compatible with this environment.

3.2.2 Power Budgets

When designing a thermal solution for an add-in card, it is important to consider all major power consuming components on the entire graphics board. For example, graphics memories such as SGRAMs operating at 100 MHz, can draw power ~ 1.0 watts (Max). Considering that an average add-in graphics card can contain 4 MB of onboard local frame buffer (i.e., SGRAM), the power consumption of the memory section alone is ~ 4 watts (Max). As the need for increased memory bandwidth arises in the 3D

market segment, future memory technologies with higher operating frequencies and therefore higher power consumption will be available. In order to provide a complete thermal solution for an add-in graphics card, it is important to budget the total consumption of power on the add-in card, and then design an adequate cooling system based on the total number. The following example lists how an add-in card vendor can budget the consumption of power on their card and, therefore, be able to design an adequate cooling system for their targeted market environment.

	# of Comp.	Major Power Consuming Comp.	Power Range
	1	Graphics controller	3-6 W
	4	Memory	3-5 W
	2	Video	2-3 W
	1	Misc. (regulators, crystals, bios, etc.)	2-3 W
Total	8		10-17 W

3.3 Recommendations for OEM System Designers

The OEM system integrator is strongly recommended to optimize the thermal environment for the A.G.P. -compliant add-in card. This is typically accomplished by maximizing air flow or reducing ambient temperature. Improving the system thermal environment can reduce costs of add-in card thermal solutions which, in turn, may reduce overall system cost.

3.4 Thermal Performance Recommendations

3.4.1 Add-in Card

Inline (not preferred) **Staggered (preferred)**

Figure 51: Airflow Over Add-In Card Components

If airflow is required over a densely populated add-in card, the inline orientation shown in Figure 51 is likely to produce higher ambient temperature for component B relative to component A. Component A blocks the airflow to component B in this configuration. In the staggered orientation shown in Figure 51, both components will equally benefit from the airflow.

3.4.2 Fan Direction

Figure 52: Pulling and Pushing Fan Airflow

Pushing Fan

- The pushing fan shown in Figure 52 can be incorporated in conjunction with a second fan of higher or lower CFM rating in both depressurized and pressurized chassis. It usually allows for a flexible device arrangement since it provides laminar airflow over a large area.
- Airflow is usually uniform and may eliminate most stagnant air and hot spots. It also keeps dust out. A front system fan is an example of a pushing fan. A pushing fan may slightly warm up the incoming air with the fan motor heat dissipation.

Pulling Fan

- The pulling fan shown in Figure 52 is generally used to exhaust heated air out of a chassis usually in conjunction with a pushing fan and, therefore, can be found in both depressurized and pressurized chassis. A power supply unit typically utilizes a pulling fan.
- Adequate venting is usually required for a pulling fan. The outgoing airflow is turbulent, which means better heat transfer in the turbulent region. However, this heat transfer area is limited around the pulling fan and may not be well defined throughout the system. Since this arrangement reduces pressure within a chassis, it tends to draw in dust through the vents and cracks in the chassis.

3.4.3 Venting

Figure 53: Adequate Venting Air

Figure 54: Inadequate Venting

Proper venting is a key element in any good thermal design. Proper venting will allow internal air to circulate and exchange with the ambient well as shown in Figure 53. Inadequate venting will cause internal air to recirculate as shown in Figure 54. If the chassis has been depressurized by fan action,

vents placed at specific locations can permit an inflow of cooler air. This air can be made to flow in at certain critical locations.

Fan venting should be as generous as possible. Similarly, grille design should be as large as possible without conflicting with EMI and safety requirements.

3.4.4 Internal Air Flow

A full-length add-in card next to an A.G.P. card is likely to block airflow to the graphics card as well as dissipate additional power. The A.G.P. card cooling will be reduced by stagnant air and higher ambient temperatures. Therefore, it is recommended that a full-length card is not installed next to the A.G.P. slot.

Cabling arrangement must not obstruct any component airflow or venting holes to ensure proper cooling.

3.5 Thermal Design Examples

The following examples of NLX and ATX chassis issues have been observed during thermal studies. Any other results will depend on the specific chassis, add-in card, and system configuration. The following examples are not representative of best or worst case thermal environments.

- **NLX Chassis Example:** Characterization of a representative NLX desktop form factor chassis thermal environment revealed low air flow (natural convection) and internal ambient temperatures of 60 °C near the A.G.P. -enabled card. For this specific chassis form factor, the side panel vent shown in Figure 55 was added and resulted in a 20 °C drop in the internal ambient temperature near the A.G.P.-enabled card. The side panel vent enabled the specific chassis design to cool a 20 W distributed thermal load card representing an A.G.P.-enabled add-in card with little impact to the overall system thermal environment. The dimensions of the side panel vent are shown in Figure 56. An example of NLX rear vents is shown in Figure 57.

Figure 55: Example NLX Air Flow

Figure 56: Example NLX Side Panel Vent

Figure 57: Example NLX Rear Vents

- **ATX Chassis Example:** Characterization of an ATX mini-tower chassis thermal environment revealed low air flow (natural convection) and internal ambient temperatures of 45 °C near a 12 W thermal load card representing an A.G.P. add-in card. In this particular environment, component temperatures on the thermal load card were 20 °C above the upper specification limit. Addition of a front fan to the system, in addition to the existing power supply fan, increased the air flow in the vicinity of the A.G.P.-enabled card and resulted in a drop of internal component temperatures to within the upper specification limit. The resulting air flow is shown in Figure 58.

Figure 58: Example ATX Air Flow

- **Add-in Card Example:** A chassis design environment was analyzed where there was a 50 °C stagnant airflow condition surrounding a 12 W graphics card with a 5 W graphics controller in a horizontal card orientation with components down. The stagnant airflow condition resulted from adjacent add-in cards which blocked the natural convection air flow. Thermal simulation results indicated that a passive heat sink with dimensions of 3” x 2.3” x .5” attached to the graphics controller package met the controller thermal specification limits, but the maximum thermal specification limits of surrounding components such as cables (80 °C insulation limit), heat sink (U.L. 70 °C limit), and oscillator (data sheet 70 °C limit) were exceeded. For this specific thermal environment and power level, product safety and cost considerations dictated that a heat sink with active cooling would allow adequate controller cooling while meeting the thermal specifications of surrounding components.

3.5.1 Reference Documents

- Accelerated Graphics Port Interface Specification
- NLX Motherboard Specification
- ATX Specification
- Component Data Sheets and Application Notes

A. Crosstalk Modeling Techniques

A.1 Objective

In large multi-conductor systems, excessive line-to-line coupling, or crosstalk can cause two detrimental effects. First of all, crosstalk will change the performance of the driven lines by modifying the effective characteristic impedance and propagation velocity, which will effect flight time and skew numbers. Additionally, crosstalk will induce noise onto other lines, which may degrade the signal integrity. The methodology presented in this appendix deals only with the first of these two effects. The study will determine how crosstalk can alter the skew and flight time of a transmission line. This appendix will also introduce simplified models that will reduce the simulation time of large N-conductor systems. Coupled lines in multi-conductor systems have variable impedances and propagation velocities which are dependent on coupling factors and data patterns. Because flight time skew is dependent on trace impedance and propagation velocities, crosstalk is an essential aspect of board design and analysis.

This paper will explore the effects of odd and even propagation modes on flight time. It will be shown that simple single line models, with the appropriate parameters can be used to model flight time differences due to line-to-line coupling in two-conductor systems. This concept will be expanded for use in the calculation of the effective impedance and propagation velocities in an N-conductor system. A methodology will be introduced that will allow an N-conductor system to be modeled as a single line for the purposes of flight time and skew analysis at greatly reduced simulation and analysis times.

It is important to note that this analysis is valid for linear systems only. A transmission line array can be approximated by a linear system only when frequencies are low. This technique should be valid for frequencies where the skin depth is larger than the conductor thickness. Above this frequency, non-linear effects such as internal inductance and frequency dependent resistance will begin to be significant.

A.2 Flight Time and Skew as a Function of Odd and Even Propagation Modes

Simulations were performed in order to quantify signal skew and flight time as a function of propagation modes. The effective characteristic impedance and line delay will change with different propagation modes in systems where significant coupling between traces exist. Electric and magnetic fields interact in specific ways that are dependent on the data patterns on the coupled traces. These specific patterns can effectively reduce or increase the effective parasitic inductance and capacitance seen by a single line. Since the transmission line parameters vary with data patterns, and are essential for accurate skew analysis, it is important to consider crosstalk analysis in high trace density/high speed systems. This appendix will introduce a skew/flight time analysis method where multiple conductors can be represented by a single conductor with appropriate variation in the impedance and velocity. This method should significantly reduce the simulation and analysis time needed to evaluate large systems.

A.2.1 Definitions

Characteristic Impedance and Propagation Delay: The Characteristic impedance and the propagation velocity are defined as follows:

Characteristic impedance = $Z = \sqrt{\frac{L}{C}}$ ohms

Propagation delay (in seconds per unit distance) = $S = \sqrt{L\ C}$

Where: L = inductance per unit length and C = capacitance per unit length.

Transmission Line Parameters: The parameters used to describe the electrical equivalent of a transmission line. In this text, only the characteristic impedance (Z) and the propagation delay (S) are considered.

Inductance and Capacitance Matrix:: The matrix of parasitics reported by a field solver that represent the transmission line array effective inductance and capacitance values.

$$\text{Inductance Matrix} = \begin{bmatrix} L_{11} & L_{12} & \cdots & L_{1N} \\ L_{21} & L_{22} & & \\ \vdots & & \ddots & \\ L_{N1} & & & L_{NN} \end{bmatrix}$$

L_{NN} = Self inductance of line N

L_{MN} = Mutual inductance between lines M and N

$$\text{Capacitance Matrix} = \begin{bmatrix} C_{11} & C_{12} & \cdots & C_{1N} \\ C_{21} & C_{22} & & \\ \vdots & & \ddots & \\ C_{N1} & & & C_{NN} \end{bmatrix}$$

C_{NN} = Total capacitance seen by line N = capacitance of conductor N to ground plus all the mutual capacitance to other lines.

C_{NM} = Mutual Capacitance between conductors N and M

For example, consider the following two-conductor system (Figure 59).

Figure 59: Simple Two-Conductor System Used to Explain the Parasitic Matrixes

$$\text{Capacitance matrix} = \begin{bmatrix} C_{11} & C_{12} \\ C_{21} & C_{22} \end{bmatrix}$$

C_{11} = Total capacitance for line 1 = Capacitance of line 1 to ground (C_s) plus all the mutual capacitance.

(i.e., for the two-conductor system, $C_{11} = C_{s1} + C_{12}$)

C_{12} = Mutual capacitance between line 1 and line 2.

$$\text{Inductance matrix} = \begin{bmatrix} L_{11} & L_{12} \\ L_{21} & L_{22} \end{bmatrix}$$

L_{11} = Self inductance for line 1

L_{12} = Mutual inductance between line 1 and line 2.

Odd Mode Propagation: Propagation mode when two coupled transmission lines are driven with equal magnitude and 180 degrees out of phase with one another. The fields arrange themselves such that there is a virtual ground plane between the two signals which effectively doubles the mutual capacitance. Since the current in the lines are always traveling in opposite direction, the self inductance is reduced by the mutual inductance. The effective odd mode capacitance and inductance for a 2 line symmetrical system are:

$C_{odd} = C_s + 2C_{12} = C_{11} + C_{12}$

$L_{odd} = L_{11} - L_{12}$

$$Z_{\text{odd}} = \sqrt{\frac{L_{\text{odd}}}{C_{\text{odd}}}}$$

$$S_{odd} = \sqrt{L_{odd} C_{odd}}$$

Even Mode Propagation: Propagation mode when two coupled transmission lines are driven with equal magnitude in phase with each other. The fields arrange themselves such that there is a virtual magnetic wall between the two conductors, which eliminates the effect of the mutual capacitance. The mutual inductance, however, is added to the self inductance because the current in both lines are always traveling in the same direction. The effective even mode capacitance and inductance for a two line symmetrical system are:

$C_{even} = C_s = C_{11} - C_{12}$

$L_{even} = L_{11} + L_{12}$

$$Z_{\text{even}} = \sqrt{\frac{L_{even}}{C_{\text{even}}}}$$

$$S_{even} = \sqrt{L_{even} C_{even}}$$

Flight Time: Time measured between the ½(VCC) points of an unloaded driver and at the far end of the transmission line (at the receiver).

Skew: Difference in flight time between two signal nets.

A.2.2 Flight Time Analysis of Two Conductors Using a Single Conductor Model

Two coupled conductors can be modeled as one conductor by determining the effective odd mode or even mode impedance and propagation delay of the transmission line pair and substituting these parameters into a single line model.

Simulations were performed using XNS to simulate the effective odd mode and even mode flight time for a 2 inch transmission line pair at width/space ratios of 5/5, 5/10, and 5/15 mils. The relative dielectric constant was simulated as 4.0 and the dielectric thickness was set to 5.0 mils. The simulations were performed with loads of 0, 1, and 10 pF. Figure 60 depicts the simulation setup.

Figure 60: Schematic Used for Calculating Odd and Even Mode Flight Time Differences

XFX, which is a 2-D electromagnetic field simulator was used to calculate the capacitance and inductance matrixes used in the simulation. These L and C matrixes were used to calculate the odd and even mode transmission line parameters using the above formulae. The resulting characteristic impedance and propagation delays were inserted into the single line model depicted in Figure 61. The flight time of the single line model was compared to the flight time of the coupled pair for both even and odd modes. Additionally, the loads were varied in the same manner as the coupled pair. The simulations confirm that the flight times of the differential pair matches the single line case when the characteristic impedance and the propagation delay of the single line match the values calculated from the parasitic matrixes for a given mode. The results are shown in Table 27.

Figure 61: Schematic Used for Representing a Conductor Pair with a Single Line

Table 27: Results of Single Line and Coupled Line Simulations Showing Agreement Between Two-Conductor and One-Conductor Flight Times for Both Even and Odd Modes

Configuration	**Simulated flight time in nanoseconds of odd and even propagation modes comparing coupled and single line models**								
Spacing (mils)	**5/5**			**5/10**			**5/15**		
Load Capacitance	0 pF	1 pF	10 pF	0 pF	1 pF	10 pF	0 pF	1 pF	10 pF
Odd Mode Coupled Pair	0.27	0.30	0.52	0.27	0.30	0.54	0.28	0.31	0.56
Single Line with odd mode impedance and velocity	0.27	0.30	0.52	0.27	0.30	0.54	0.28	0.31	0.57
Even Mode Coupled Pair	0.29	0.33	0.63	0.29	0.32	0.61	0.29	0.32	0.6
Single Line with even mode impedance and velocity	0.29	0.33	0.63	0.29	0.32	0.61	0.29	0.32	0.6

The above results show that the single line equivalent model agrees with the two-conductor simulation. Furthermore, the simulations indicate that there can be significant differences in the flight times between a pair of signals that are propagating in phase and a pair of signals that are propagating 180 degrees out of phase. Additionally, when the load capacitance increases, the flight time differences between even and odd modes increases. It should also be noted that the flight time differences due to different propagation modes are exacerbated by tightly coupled lines.

A.2.3 Flight Time/Skew Analysis for More Than Two Conductors

The above simulations show how one conductor can replace two conductors for the purposes of incorporating crosstalk into flight time skew analysis. The objective of this study, however, is to determine a methodology to simulate one conductor in place of N conductors. To do so, formulae must be developed in order to correctly predict the equivalent inductance and capacitance for a given switching pattern of N bits on N coupled transmission lines.

A.2.4 Methodology

A methodology was developed to determine the effective impedance and velocity for multiple conductor lines. Initially, the equivalent capacitance and inductance seen by the target line (the line that you are doing the skew simulation on) must be determined for a specific switching pattern. This can be accomplished by observing pairs of lines in a N conductor system and determining the equivalent odd or even mode capacitance and inductance values. The values from the pairs can then be combined to produce the final value. To do so, the mutual components are either added or subtracted from the total capacitance or self inductance of the target line. When the signal line is switching in phase with the target line, then the mutual capacitance between the two lines is subtracted from the total capacitance of that line and the mutual inductance is added. Conversely, when the signal line is switching out of phase with the target line, then the mutual capacitance between those two lines is added to the total capacitance of that line and the mutual inductance is subtracted. The effective capacitance and inductance are used to calculate the equivalent characteristic impedance and propagation velocity for the single line equivalent model.

A.2.4.1 Example 1

Assume that switching pattern of the shown three conductor system is such that all bits on the net are switching in phase. Refer to the cross section in Figure 62.

Figure 62: Three Conductor Even Mode System

If the primary goal is to simulate this multi-conductor system with one line, and the line of interest is line 2, then the equivalent capacitance and inductance seen by line 2 can be calculated as follows:

Even Mode Capacitance of conductors 2 and 1 = $C_{s2} + C_{12} - C_{12} = C_{22} - C_{12}$

Even Mode Capacitance of conductors 2 and 3 = $C_{s2} + C_{13} - C_{13} = C_{22} - C_{23}$

Equivalent capacitance of middle conductor when conductors 1-3 are switching in phase = C_{22} - C_{12} - C_{23}

Even Mode Inductance of conductors 2 and 1 = L_{22} + L_{12}

Even Mode Inductance of conductors 2 and 3 = L_{22} + L_{23}

Equivalent inductance of middle conductor when conductors 1-3 are switching in phase = L_{22} + L_{12} + L_{23}

The equivalent single conductor model that could be used to perform skew analysis of this multi-conductor system would have the following transmission line parameters:

$$Z_{2,eff} = \sqrt{\frac{L_{22} + L_{12} + L_{23}}{C_{22} - C_{12} - C_{23}}}$$

$$S_{2,eff} = \sqrt{(L_{22} + L_{12} + L_{23})(C_{22} - C_{12} - C_{23})}$$

It is interesting to note that the mutual inductances are added or subtracted in the opposite manner as the mutual capacitance. This is due to the assumption that the system is linear and propagating only in TEM mode, which means that the electric and magnetic fields are always orthogonal to each other. As a result of the TEM assumption, the product of L and C remains constant for homogeneous systems. Thus, in a multi-conductor homogeneous system, such as a stripline array, if L is increased by the mutual inductance, then C must be decreased by the mutual capacitance such that LC remains constant.

In a non-homogeneous multi-conductor system, such as an array of microstrip lines, LC is not held constant for different propagation modes because the electromagnetic fields are traveling partially in air and partially in the dielectric material of the board. In a microstrip system, the effective dielectric constant is a weighted average between air and the dielectric material of the board. Because the field patterns change with different propagation modes, the effective dielectric constant will change depending on the field densities contained within the board dielectric material and the air. Thus, the LC product will be mode dependent in a non-homogeneous system. The LC product, however, will remain constant for a given mode.

A.2.4.2 Example 2

Assume that the switching pattern of a three conductor system is such that bits 1 and 2 are switching in phase and bit 3 is switching 180 degrees out of phase. (Line 2 is still the line of interest.) Figure 63 shows this system.

Figure 63: Three Conductor System with Odd Mode and Even Mode

Even Mode Capacitance of conductors 2 and 1 = C_{22} - C_{12}

Odd Mode Capacitance of conductors 2 and 3 = C_{22} + C_{23}

Equivalent capacitance of conductor 2 = C_{22} - C_{12} + C_{23}

Even Mode Inductance of conductors 2 and 1 = $L_{22} + L_{12}$

Odd Mode Inductance of conductors 2 and 3 = $L_{22} - L_{23}$

Equivalent inductance of conductor 2 = $L_{22} + L_{12} - L_{23}$

The equivalent single conductor model that could be used to perform skew analysis of this multi-conductor system would have the following transmission line parameters:

$$Z_{2,eff} = \sqrt{\frac{L_{22} + L_{12} - L_{23}}{C_{22} - C_{12} + C_{23}}}$$

$$S_{2,eff} = \sqrt{(L_{22} + L_{12} - L_{23})(C_{22} - C_{12} + C_{23})}$$

A.2.4.3 Example 3

Assume that switching pattern of a five conductor system is such that bits 1, 2, and 5 are switching in phase, and bits 3 and 4 are in phase with each other but 180 degrees out of phase with bits 1, 2, and 5 (Line 3 is the line of interest.) Figure 64 shows this system.

Figure 64: Five Conductor System

Odd Mode Capacitance of conductors 3 and 1 = $C_{33} + C_{13}$

Odd Mode Capacitance of conductors 3 and 2 = $C_{33} + C_{23}$

Even mode capacitance of conductors 3 and 4 = $C_{33} - C_{34}$

Odd mode capacitance of conductors 3 and 5 = $C_{33} + C_{35}$

Equivalent capacitance of conductor 3 = $C_{33} + C_{13} + C_{23} - C_{34} + C_{35}$

Odd Mode inductance of conductors 3 and 1 = $L_{33} - L_{13}$

Odd Mode inductance of conductors 3 and 2 = $L_{33} - L_{23}$

Even Mode inductance of conductors 3 and 4 = $L_{33} + L_{34}$

Odd mode inductance of conductors 3 and 5 = $L_{33} - L_{35}$

Equivalent inductance of conductor 3 = $L_{33} - L_{13} - L_{23} + L_{34} - L_{35}$

The equivalent single conductor model that could be used to perform skew analysis of this multi-conductor system would have the following transmission line parameters:

$$Z_{3,eff} = \sqrt{\frac{L_{33} - L_{13} - L_{23} + L_{34} - L_{35}}{C_{33} + C_{13} + C_{23} - C_{34} + C_{35}}}$$

$$S_{3,eff} = \sqrt{(L_{33} - L_{13} - L_{23} + L_{34} - L_{35})(C_{33} + C_{13} + C_{23} - C_{34} + C_{35})}$$

Simulations were performed on a three-conductor system to confirm that the above methodology produces the correct flight time results when a single line is used to represent a multi-conductor system. Table 28 shows simulation results comparing a three-conductor flight time simulation and a single conductor flight time simulation where the equivalent impedance and propagation delay has been used for the appropriate switching pattern. The target conductor is line 2 (L2). The simulations were performed using XNS to simulate the effective odd mode and even mode flight times for an array of 2-inch long transmission lines at width/space ratios of 5/5 mils. The relative dielectric constant was simulated as 4.0 and the dielectric thickness was set to 5.0 mils. The simulations were performed with loads of 0, 1, and 10 pF. The lines were driven with 1.8 volts and a 100 ps edge rate. See Figure 25 or Figure 26 for a representative cross section.

Table 28: Results of Multi-line and Single Line Flight Time Simulations Showing Agreement Between 3-Conductor and Equivalent Single Conductor Models

	Effective Parasitics		Load = 0pF		Load = 1pF		Load = 10pF	
Switching Pattern L1 L2 L3	L	C	Flight Time (L2 of Multi-Line)	Flight Time (Equiv. Single Line)	Flight Time (L2 of Multi-Line)	Flight Time (Equiv. Single Line)	Flight Time (L2 of Multi-Line)	Flight Time (Equiv. Single Line)
1 0 0	L22-L12+ L23	C22+C12- C23	0.28 ns	0.28 ns	0.31 ns	0.31 ns	0.57 ns	0.57 ns
1 1 0	L22+L12- L23	C22-C12+ C23	0.28 ns	0.28 ns	0.31 ns	0.31 ns	0.57 ns	0.57 ns
1 0 1	L22-L12- L23	C22+C12+ C23	0.25 ns	0.25 ns	0.26 ns	0.26 ns	0.46 ns	0.46 ns
1 1 1	L22+L12+ L23	C22-C12- C23	0.30 ns	0.3 ns	0.34 ns	0.34 ns	0.67 ns	0.67 ns

B. Models

Table 29: Buffer I-V Curves

48 ohm curvature 2.0		38 ohm curvature 2.0		30 ohm curvature 1.1	
Voltage	**Current**	**Voltage**	**Current**	**Voltage**	**Current**
0.0	0.0	0.0	0.0	0.0	0.0
0.05	1.424	0.073	3.122	0.132	4.94
0.1	2.811	0.144	5.969	0.256	9.446
0.15	4.157	0.212	8.54	0.373	13.516
0.2	5.462	0.279	10.837	0.483	17.15
0.25	6.72	0.344	12.86	0.586	20.351
0.3	7.928	0.407	14.611	0.682	23.123
0.35	9.082	0.467	16.094	0.77	25.47
0.4	10.178	0.526	17.313	0.852	27.399
0.45	11.209	0.583	18.273	0.927	28.918
0.5	12.171	0.639	18.98	0.996	30.038
0.55	13.056	0.692	19.443	1.058	30.77
0.6	13.857	0.744	19.732	1.115	31.228
0.65	14.564	0.796	19.986	1.172	31.629
0.7	15.166	0.848	20.21	1.228	31.984
0.75	15.652	0.899	20.407	1.283	32.296
0.8	16.008	0.95	20.58	1.337	32.57
0.85	16.241	1.001	20.731	1.391	32.808
0.9	16.443	1.052	20.863	1.445	33.017
0.95	16.625	1.103	20.978	1.498	33.198
1.0	16.787	1.154	21.079	1.550	33.358
1.05	16.931	1.205	21.168	1.603	33.5
1.1	17.058	1.255	21.249	1.655	33.628
1.15	17.17	1.306	21.324	1.707	33.747
1.2	17.269	1.356	21.396	1.759	33.86
1.25	17.356	1.407	21.466	1.811	33.971

1.3	17.434	1.457	21.535	1.862	34.08
1.35	17.503	1.508	21.603	1.914	34.188
1.4	17.567	1.558	21.67	1.966	34.294
1.45	17.626	1.609	21.735	2.018	34.398
1.5	17.684	1.659	21.8	2.069	34.5

Table 30: Diode Clamp I-V Curves

Voltage	Current
0	0
1.5	0
2.2	0
2.4	10
2.6	100

C. Sensitivity Analysis

Several sensitivity analyses were performed to narrow the definition of the buffer and interconnect characteristics for 4X mode. These analyses used results from numerous simulations. Included in this appendix is the result of the final set of analyses that defined the buffer and interconnect characteristics. These results are displayed using 3D plots in Figure 65 through Figure 76.

These plots display the worst case skew at each point in the surface. These points are the same as the ones that are described in Section 2.2.3.4.

The first six plots are of the worst case for setup skew and the last six plots are of the worst case results for hold skew. The buffer types are described in the figure title.

Buffers with no curvature experienced ringback violations over nearly the entire surface plot. These also reduced the amount of buffer impedance variations and line lengths.

Figure 65: No Curvature Compensation with a 1.4 V/ns Slew Rate

Figure 66: No Curvature Compensation with a 2.8 V/ns Slew Rate

Figure 67: 2.0 Curvature Compensation with a 1.4 V/ns Slew Rate

Figure 68: 2.0 Curvature Compensation with a 2.8 V/ns Slew Rate

Figure 69: 1.1 Curvature Compensation with a 1.4 V/ns Slew Rate

Figure 70: 1.1 Curvature Compensation with a 2.8 V/ns Slew Rate

Figure 71: No Curvature Compensation with a 1.4 V/ns Slew Rate

Figure 72: No Curvature Compensation with a 2.8 V/ns Slew Rate

Figure 73: 2.0 Curvature Compensation with a 1.4 V/ns Slew Rate

Figure 74: 2.0 Curvature Compensation with a 2.8 V/ns Slew Rate

Figure 75: 1.1 Curvature Compensation with a 1.4 V/ns Slew Rate

Figure 76: 1.1 Curvature Compensation with a 2.8 V/ns Slew Rate

D. AGP 4X 3.3V Vcc_Core Motherboard Decoupling Recommendations

This appendix describes a method to determine the amount of motherboard decoupling needed for the AGP-4X Vcc_Core voltage, 3.3V, such that the voltage tolerance specifications are easily met with extra headroom at the card interface.

Transient Voltage Tolerance at the Add-in Card

3.3 V +/- 150 mV

Modeling Assumptions

- On-card decoupling was added to the following circuit model, based on the amount of card decoupling typically used on an A.G.P. 2X card.
- The I_Load piecewise linear function was based on worst case analysis that ½(Imax) occurs in one cycle, and then Imax is arrived at three cycles later. This method is similar to models for processor power delivery representing worst case, first-order estimates.
- Also, in order to complete the model, parasitics were added for motherboard trace impedance, package impedance, and some on-card impedance. These parameters assure the simulation will not be too optimistic. The motherboard trace impedance is based on estimates that were previously used for parasitics associated with the change from the VRM connector to the Slot 1 connector and also on the number of squares associated with changing from the power supply decoupling to the connector. The package numbers are estimates based on similar models for processor interposers. These package parameters should be similar to the package numbers for the graphics controller.

Circuit Model

Circuit Parameter Values

- I_Load:

  Of all the components on the A.G.P. 4X card, the graphics controller operates at the highest frequency, 133 MHz. Therefore, for developing a worst-case model, 133 MHz was used to get the Δt part of the di/dt. For the ΔI part, it is considered a worst-case measure to assume ½(Imax) occurs in Δt. Imax is reached after three cycles.

  Δt = 1/133e6 = 7.52 ns

  I_Load = pwl (0 ns 0 A, 5 ns 0 A, 12.52 ns 3 A, 35.08 ns 6 A, 1 µs 6 A)

- Motherboard Capacitance:

  C = 22 µF

  ESL = 0.34 nH

  ESR = 35 mOhms

  # Capacitors used in model: one

- A.G.P. 4X On-Card Decoupling:

  C = 10 μF

  ESL = 880 pH

  ESR = 35-50 mOhms

  # Capacitors used in model: two

- A.G.P. 4X On-Card HF Decoupling:

  C = 0.1 μF

  ESL = 0.57 nH

  ESR = 43 mOhms

  # Capacitors used in model: six

- A.G.P. Connector Parameters:

  L_agp_conn = 10.5 nH / 120 pins = 0.0875 nH

  R_agp_conn = 12 mOhms / 120 pins = 0.1 mOhms

Simulation Waveforms

@MB Decoupling: ΔV_undershoot = 3.3 V – 3.29 V = 10 mV drop

After A.G.P. Connector: ΔV_undershoot = 3.3 V – 3.166 V = 134 mV drop

@Card Decoupling: ΔV_undershoot = 3.3 V – 3.164 = 136 mV drop

Extra Margin: 150 mV – 135 mV = 14 mV

Motherboard Capacitance Recommendation:

Based on the simulation above, by using one 22 µF tantalum style capacitor and by keeping the card decoupling the same as what was used in A.G.P. 2X, there is 14 mV of transient voltage margin with respect to the specification at the A.G.P. 4X card.