

AGP

ACCELERATED GRAPHICS PORT

Platform Design Guide

Revision 1.1A
Intel Corporation
September 1997

intel[®]



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, 1997

Table of Contents

TABLE OF CONTENTS	3
FIGURES.....	5
TABLES.....	5
PREFACE.....	7
<i>Document Structure and Outline</i>	7
<i>Content</i>	7
<i>Feedback</i>	7
CHAPTER 1.....	8
AGP ELECTRICAL DESIGN CONSIDERATIONS	8
<i>AGP Electrical Requirements</i>	8
<i>Input Buffer Considerations for 133 MHz Transfer (Master - Target)</i>	8
<i>Signal Termination</i>	9
<i>Signal Overshoots and Undershoots</i>	10
<i>Voltage / Loading Characteristics of VREF</i>	10
<i>System Timing Considerations</i>	11
MOTHERBOARD AND EXPANSION CARD PCB LAYOUT CONSIDERATIONS	15
<i>Motherboard Layout Recommendations</i>	15
<i>Expansion Card Layout Recommendations</i>	16
<i>Board Impedance</i>	17
<i>Pull-ups and Termination</i>	17
<i>Board Routing Recommendations</i>	17
<i>Interface Signaling Requirements</i>	20
<i>IDSEL Routing For Add-in Cards</i>	20
IC PACKAGING CONSIDERATIONS	20
<i>Signal Skewing</i>	20
<i>Signal Loading</i>	20
<i>Signal Routing and Crosstalk</i>	21
<i>QFP vs. BGA</i>	21
CHAPTER 2.....	22
BOARD DESIGN GUIDELINES	22
<i>A.G.P. 2X Mode Design Considerations</i>	22
<i>Source Synchronous Recommendations for Add-in Card</i>	23
<i>Source Synchronous Recommendations for Motherboards</i>	24
<i>Control Signal and Clock Recommendations</i>	24
SIMULATION TECHNIQUES	25
<i>Introduction</i>	25
<i>Methodology</i>	26
<i>Buffer Models</i>	26
<i>Interconnect Models and Crosstalk Effects</i>	26
<i>Simulations Setup</i>	27
<i>Simulation Parameters</i>	28
<i>Flight Time</i>	28
<i>Flight Time Skew</i>	29
<i>Signal Quality</i>	30
<i>Summary of Simulations Required</i>	31
SIMULATION RESULTS	32
<i>Flight Time</i>	33
<i>Flight Time Skew</i>	33
<i>Signal Quality</i>	35
CHAPTER 3.....	36
AGP THERMAL DESIGN GUIDELINES.....	36
<i>AGP Thermal Overview</i>	36
THERMAL DESIGN RECOMMENDATIONS	36

Recommendations for Add-In Card Designers 36
Power Budgets:..... 37
RECOMMENDATIONS FOR OEM SYSTEM DESIGNERS 37
RECOMMENDATIONS FOR THERMAL PERFORMANCE 38
VENTING..... 39
INTERNAL AIR FLOW..... 40
THERMAL DESIGN EXAMPLES..... 40
Reference Documents: 43

Figures

FIGURE 1-1: DIFFERENTIAL INPUT BUFFER AND REQUIRED TRANSFER CHARACTERISTIC	9
FIGURE 1-2: VREF RESISTOR DIVIDER.....	11
FIGURE 1-3: AGP PCB LAYOUT FOR ATX FORM FACTOR.....	15
FIGURE 1-4: AGP PCB LAYOUT FOR LPX FORM FACTOR.....	15
FIGURE 1-5: AGP PCB LAYOUT FOR NLX FORM FACTOR.....	16
FIGURE 1-6: MOTHERBOARD AND ADD-IN CARD IMPEDANCE MATCHING.....	17
FIGURE 1-7: TIMING BUDGET BETWEEN FOR GC, AGP AND CORE CHIPSET.....	18
FIGURE 1-8: REQUIRED TRACE LENGTH TO MINIMIZE CLOCK SKEW.....	19
FIGURE 2-1: EXAMPLE OF SKEWS.....	23
FIGURE 2-2: A.G.P. WITH ONE CONNECTOR TOPOLOGY.....	25
FIGURE 2-3: FLIGHT TIME MEASUREMENT.....	29
FIGURE 2-4: SIGNAL QUALITY MEASUREMENTS.....	30
FIGURE 2-5: SENSITIVITY ANALYSIS CHART.....	32
FIGURE 3-1: AGP ADD-IN CARD AND GRAPHICS CONTROLLER POWER TREND.....	36
FIGURE 3-2: AIRFLOW OVER ADD-IN CARD COMPONENTS.....	38
FIGURE 3-3: PULLING AND PUSHING FAN AIRFLOW.....	38
FIGURE 3-4: ADEQUATE VENTING AIR.....	39
FIGURE 3-5: INADEQUATE VENTING.....	40
FIGURE 3-6: EXAMPLE NLX AIR FLOW.....	41
FIGURE 3-7: EXAMPLE NLX SIDE PANEL VENT.....	41
FIGURE 3-8: EXAMPLE NLX REAR VENTS.....	42
FIGURE 3-9: EXAMPLE ATX AIR FLOW.....	43

Tables

TABLE 1-1: TABLE OF STANDARD RESISTANCE VALUES.....	11
TABLE 1-2 A.G.P. TIMING ELEMENTS.....	11
TABLE 1-3 EXAMPLE AGP-1X TIMING ELEMENT COMPONENTS.....	13
TABLE 1-4 EXAMPLE AGP-2X TIMING ELEMENT COMPONENTS.....	13
TABLE 2-1: DATA AND ASSOCIATED STROBE.....	23
TABLE 2-2: ADD-IN CARD RECOMMENDATIONS.....	23
TABLE 2-3: MOTHERBOARD RECOMMENDATIONS.....	24
TABLE 2-4: CONTROL SIGNAL LINE LENGTH RECOMMENDATIONS.....	24
TABLE 2-5: INTERCONNECT IMPEDANCE AND VELOCITIES.....	27
TABLE 2-6: SUMMARY OF SIMULATIONS.....	32
TABLE 2-7: MINIMUM FLIGHT TIME.....	33
TABLE 2-8: MAXIMUM FLIGHT TIME.....	33
TABLE 2-9: SUMMARY OF SOURCE SYNCHRONOUS COMPARISONS.....	34
TABLE 2-10: SKEW FROM MOTHERBOARD WITH 1:1 TRACE SPACING.....	34
TABLE 2-11: SKEW FROM MOTHERBOARD WITH 1:2 TRACE SPACING.....	35

Preface: Document Introduction and Overview

Chapter 1: AGP Electrical Design Considerations

Provides useful design considerations on electrical implementations of AGP compliant devices, add-in cards, and PCB layout.

Chapter 2: AGP Signals Sensitivity Analysis and Measurement Techniques

Provides important information in generating design guidelines for optimized AGP electrical signals.

Chapter 3: AGP Thermal Design Guidelines

Provides guidance on thermal design for platforms.



Preface

Document Structure and Outline

The information in this design guide is organized into chapters. Each chapter is in a modular format, with a non-numbered heading for each major topic and sub-topics. Pages are numbered sequentially. Figures and tables are numbered by chapter and sequence within that chapter.

The modular format allows for easy update of the *AGP Platform Design Guide* as new information becomes available or as specifications changes. We suggest that you install this document into a 3-ring binder and add or replace pages as required. You may notice that some pages contain more blank space than others. This is done to confine sections to page boundaries for easy update. This results in a modular and extensible *AGP Platform Design Guide*.

Content

The initial release of this document centered around electromechanical implementation issues. The first updated added thermal design concerns. 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 AGP Implementers Forum web site (www.agpforum.org).



Chapter 1

AGP Electrical Design Considerations

AGP Electrical Requirements

The Accelerated Graphics Port (A.G.P.) electrical specification is an extension of the PCI electrical specification. The modifications are designed to provide a higher performance graphics interface. A baseline 66 MHz A.G.P. interface can be designed using PCI I/O buffer technology. Some of the major modifications to the PCI electrical specification are relaxed timings. Many of the signal's setup and hold times are relaxed due to the improved electrical characteristics of a two component bus (i.e. capacitance and flight time). Following are some electrical design considerations that can help developers in designing an A.G.P. compliant device.

Input Buffer Considerations for 133 MHz Transfer (Master - Target)

Input sense levels can have a significant impact on signal skews. Getting the best timing margins for the A.G.P.133 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 be compatible with PCI and A.G.P. 66 MHz input levels.

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. 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. the input sense range is $(0.5 \cdot V_{ddq} - 0.3 \cdot V_{ddq})$. At $V_{ddq} = 3.3V$, the uncertainty range is about 0.66 Volts and the maximum skew is approximately 0.44ns.

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 V_{ddq} dependent so it can be provided from a simple resistor divider. See *Voltage / Loading Characteristics of VREF* 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 (i.e. the output of the buffer must switch within a 200 mV window of VREF under all conditions - Figure 0-1). The smaller this range can be made, the smaller the timing skew will be between inputs.




Figure 0-1: Differential Input Buffer And Required Transfer Characteristic

Although standard 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 standard 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.

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.

Signal Termination

Identifying the 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 and 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 V_{ddq} to prevent the inputs from floating. If these signals are pulled up, it is recommended that the stub length to the pull-up be zero to preserve signal quality and minimize signal skew.

- Signals should be disabled with a pull-up resistor to **V_{ddq}** (not V_{cc}). The value should be 8.2K Ohms \pm 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 V_{ddq} supply voltage.

Termination of A.G.P Signals

Termination techniques improve noise margins and reduce signal reflections that degrade signal integrity. Designers need to include termination in their PCB designs for two reasons: (1) to make the design more robust by minimizing reflections and noise due to impedance mismatches, and (2) to minimize the IC output noise sensitivity and match the loads to the IC drive capabilities

Signal Overshoots and Undershoots

Output buffers are specified to be closely matched to the impedance range of 50-80 ohms for A.G.P. systems, and are about half the strength of PCI 3.3Vv buffers. For this reason, the amount of overshoot and undershoot seen on an A.G.P. signal should be less than what might be seen in a lightly loaded PCI 3.3V system. This is the primary reason for the buffer strength specified. No special termination should be required to achieve acceptable signal quality.

Cause of Overshoots and Undershoots

The primary cause of signal overshoot and undershoot is the output driver, 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 driver and line, the greater the magnitude of the overshoot. Due to the separate pull-up and pull-down devices in the driver, the magnitude of the overshoot and undershoot can be independently affected by a component.

Measuring Overshoots and Undershoots

Signal overshoot is defined as the amount of signal voltage excursion beyond the power rail during a switching transition, and is measured from the peak of the signal excursion back to the rail. Overshoot typically is used to refer to excursions beyond the VCC rail in CMOS systems, while undershoot refers to excursions below the VSS rail. In a typical system, overshoot will be followed by a decaying oscillation around the supply rail that will settle down in a given period of time.

Effect of Clamping on Overshoots / Undershoots

A.G.P., like PCI, requires signal clamping elements, which protect the device and improve the signal quality with respect to overshoot and undershoot. These clamps act as shunts for much of the energy that the line is providing at the receiver, and hold the peak overshoot to a much lower value than it might otherwise reach. Unlike A.G.P., all PCI compliant components are required to have clamping elements. The clamping elements for an A.G.P. solution may be internal or external to A.G.P. component. The following clamping method can be used:

Typical CMOS buffer implementations usually have a large silicon diode between VSS and the input, and another between the input and VCC. These diodes are able to clamp overshoots and undershoots that exceed 0.7 to 1.0 Volts quite well, and improve the settling time for the signal to settle back to the rail. Active clamping can be used to improve settling time, but is not required.

Voltage / Loading Characteristics of VREF

VREF is a DC voltage reference signal used to set the input sense level on the AGP bus. It is set between $0.39 \times V_{ddq}$ and $0.41 \times V_{ddq}$ Volts. It can be generated on the AGP compliant device or provided by the system using a simple resistor divider from V_{ddq} . 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 $\pm 10 \mu\text{A}$ 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 1-2). The resistor ratio $R1/R2$ must be exactly $2/3$. Table 1-1 shows some example values of $R1$ and $R2$ for resistor tolerances of 1% and 2% with one device load current of $\pm 10 \mu\text{A}$. The tighter tolerance resistor has more margin to the VREF spec for IR drop due to the load current and therefore can have larger resistance values. If the VREF circuit is used to drive two AGP compliant devices, then resistor values no more than half the maximum allowed value must be used due to the doubled load current.




Figure 0-2: VREF Resistor Divider

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

Table 0-1: Table of standard resistance values

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, especially if a single VREF source is used to supply both AGP interfaces. 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 AGP interface circuits to avoid noise pick-up.

VREF can also be generated internal to the AGP interface, so the actual value of VREF may not be visible. In this case, VIL and VIH (AGP Interface Specification, Rev. 1.0, Table 4-2) 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.

System Timing Considerations

The timings for AGP-1X mode are based on a common clock and are driven by I/O delays and the clock skew between the AGP components. The timings for AGP-2X are based on the difference in delays (skew) between the strobe and data lines. A summary of the key elements are given in Table 0-2.

Table 0-2 A.G.P. Timing Elements

	AGP-1X (Delay)	AGP-2X (Skew)	Units
Clock Skew	1.0	(n.a.)	ns
Transmitter	6.0 ¹	2.05	ns
Interconnect	2.5	0.70	ns
Receiver	5.5 ¹	1.00 ²	ns
Total	15	3.75	ns

Note 1: Data timings

Note 2: Receiver Setup time

The delays of AGP-1X add up to one A.G.P. clock period (15 ns), the maximum time allowed for a data transfer. The skews of AGP-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. The AGP-2X timings are affected by anything which causes a difference in the delays between data and strobe for the transmitter and interconnect, and in the receiver setup time. Some of these components and example values¹ are shown in Table 0-3 and Table 0-4.

Most of the components of the AGP-1X timings should be familiar to designers with experience in I/O design, especially PCI I/O design. The AGP-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 AGP-2X timings need to be carefully considered since there is much less timing margin in this case and several small contributions add up fast. The magnitudes of the AGP-2X 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. AGP-2X designs use new bus techniques, so the timing components are likely to be less well understood. These components will be covered in more detail below. The timings for the interconnect are covered in full detail in the last sections of this chapter and in Chapter 2 and will not be covered further in this section.

The transmitter is assumed to use a PLL to produce a 50% duty cycle, 133 MHz clock needed to produce the required relative timings of data and strobe. Other timing generators can also be used for this purpose and identical timing considerations can be applied to them. The PLL is the primary source for the data / strobe clock jitter and clock duty cycle skew components. 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, by 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. If the duty cycle is not close to 50%, it may be necessary to run the VCO at 266 MHz and divide by 2. Clock duty cycle is also affected by internal loading and buffering of the clock. The 133 MHz clock path should be carefully constructed and simulated to insure the fidelity of the signal throughout the AGP-2X interface.

The 133 MHz clock also has to be routed to all outputs of the interface and proper care needs to be taken such that the clock skew across the interface is minimized. Automatic clock tree synthesis may provide insufficient skew control. Manual placement and routing gives better results. Also, the timing registers for the outputs can be built into a custom I/O buffer.

The buffer itself needs to be carefully designed to balance the delays to the rising and falling edges. This is typically one of the bigger components of transmitter skew. Any delay mismatch reduces data setup and hold time margin.

The other big transmitter skew component is SSO pushout. When a group of outputs switched in the same direction at the same time, 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 only causes data delay and it cannot cause the signal to occur earlier. This means that SSO only affects data setup time. (It affects AGP-1X the same way by increasing the clock to output delay of the buffer.)

¹ The timings values shown in Table 0-3, Table 0-4, and other similar examples in this section are given to illustrate how the timing elements might be built, and are not meant as design recommendations.

Table 0-3 Example AGP-1X timing element components

Element	Delay Component ¹	AGP-1X	Units
Transmitter	Internal clock delay	1.0 ²	ns
	Clock jitter	0.25	
	Clock to output delay	4.25	
	SSO pushout	0.5	
	Total	6.0	ns
Interconnect ³	Longest line flight time	2.05	ns
	Crosstalk	0.45	
	Total	2.50	ns
Receiver	Input buffer delay	1.5	ns
	Input logic and routing delay	4.5	
	Data register setup time	0.5	
	Internal clock delay	-1.0 ²	
	Total	5.5	ns

Table 0-4 Example AGP-2X timing element components

Element	Skew Component ¹	AGP-2X	Units
Transmitter	Data / strobe clock jitter	0.25	ns
	Data / strobe clock duty cycle	0.20	
	Internal data clock skew	0.10	
	Buffer delay matching	0.25	
	Rise / Fall time matching	0.75	
	SSO pushout	0.50	
	Total	2.05	ns
Interconnect ³	Data / Strobe trace mismatch	0.10	ns
	Capacitive loading mismatch	0.15	
	Crosstalk	0.45	
	Total	0.70	ns
Receiver	Strobe to data path skew	0.9	ns
	Strobe routing skew	0.1	
	Total	1.0	ns

Note 1: These components do not include testing guardband allowances that may be required for production testing.

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.

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 also be reduced by designing the buffers to the weak end of the buffer strength spec with the largest allowed rise and fall times, consistent with meeting the AGP-1X output delay specs. This reduces the buffer di/dt and, consequently, the inductive component of power supply voltage drop.

All buffer designs will suffer some level SSO pushout. One consequence of SSO pushout is that data hold time is always longer than data setup time. Because of this asymmetry, the AGP-2X data hold time after strobe ($tDva$) is specified to be 200 ps more than the data setup time before strobe ($tDvb$) (see table 4-4 in the A.G.P. Specification). This establishes 200 ps as the minimum allowance for SSO pushout. (Any pushout allowance greater than 200 ps does not change the specifications for data valid time after strobe).

The receiver setup and hold timings are most affected by the matching of the data input path delay to the strobe input path delay. 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 to +2.0 pF as compared to the capacitive load of all data pins connected with that strobe.

Reducing the range of input sense levels also reduces the setup and hold time skew. As mentioned above, using differential input buffers can reduce the sense level uncertainty. Also, the input buffer power supplies should be clean of noise and well bypassed and the external VREF, if used, must be well bypassed to suppress noise.

Motherboard and Expansion Card PCB Layout Considerations

This section describes layout and routing guidelines to insure a robust AGP interface design. Following these guidelines will help insure that the AGP specification can be met and that the motherboard and expansion card will operate together, but it is not a guarantee that the AGP specifications will be met. The designer should do the appropriate analysis and simulation to verify that the design meets the AGP specifications. Chapter 4 discusses some methods of simulating AGP signals.

Motherboard Layout Recommendations

The following figures (Figure 0-3, Figure 0-4, and Figure 1-6) are examples of AGP layouts for an ATX, LPX, and NLX form factor motherboards. The goal of the placement is to allow the routing of the AGP bus that minimizes trace length, vias, and interference with other signals and busses.




Figure 0-3: AGP PCB Layout For ATX Form Factor

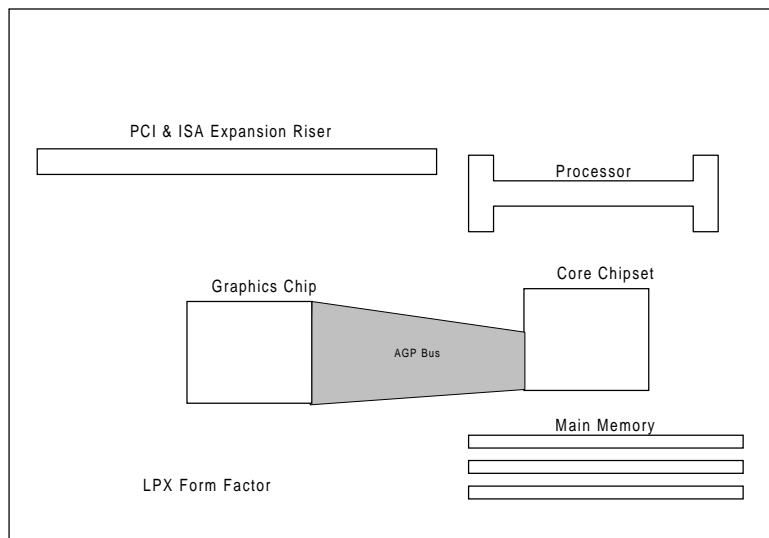


Figure 0-4: AGP PCB Layout For LPX Form Factor

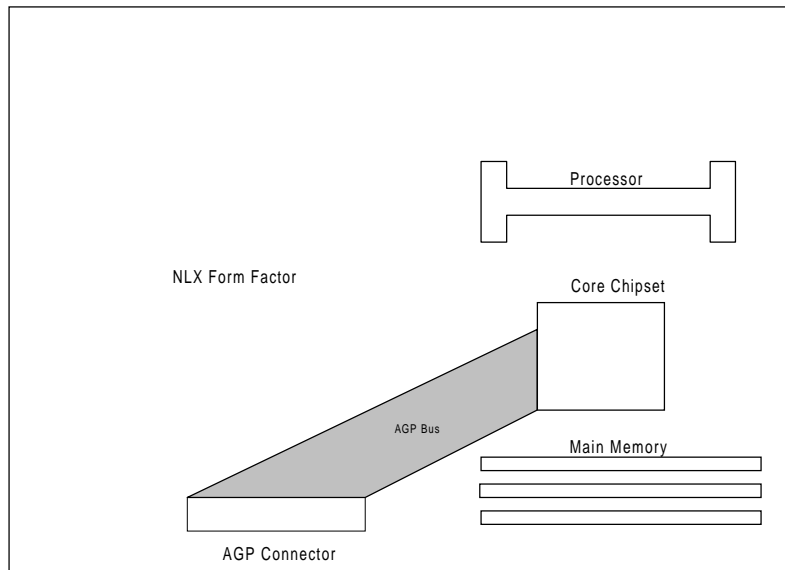


Figure 0-5: AGP PCB Layout For NLX Form Factor

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

Expansion Card Layout Recommendations

All AGP signals on the graphics chip should be located so that they are in the same order as the pin out of the AGP connector. This assumes that the component side of the add-in card is as defined in the mechanical section of the AGP Interface Specification Revision 1.0. This alignment will minimize the overall trace lengths and aid in matching the trace lengths within the groups.

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

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

If the pin out is optimized for the expansion card connector card pin out, the graphics on-board only solution will cause the AGP 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 AGP requirements in this case.

Board Impedance

6 mils	Primary Layer 1/2 oz. Cu
40 mils	Ground Layer 1 oz. Cu
6 mils	Power Layer 1 oz. Cu
6 mils	Secondary Layer 1/2 oz. Cu

4 Layer Recommendation

Figure 0-6: Motherboard And Add-in Card Impedance Matching

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. A impedance of $65\Omega \pm 15\Omega$ is recommended.

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

Pull-ups and Termination

All pull-up resistors required by the AGP interface specification must be implemented on the motherboard. The AGP 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.

Board Routing Recommendations

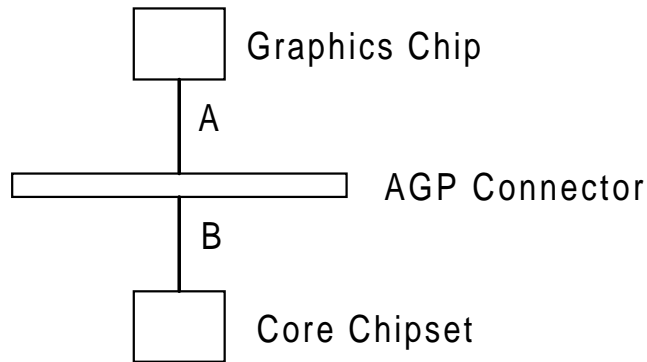
AGP signals must be carefully routed on the motherboard and graphics card to meet the timing and signal quality requirements of this interface specification. The following are some general guidelines that should be followed. Trace lengths included in this section are guidelines only. It is recommended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

Maximum Trace Length Requirements

The following shows how the maximum trace length is segmented between the motherboard and the expansion card. The maximum flight time allowed for the AGP bus is 2.5ns. The timing budget for the components of the flight path is 1.65ns for the motherboard and 0.7ns for the expansion card. This restricts the maximum trace length on the motherboard to less than 9.5 inches. The maximum trace length allocated for the add-in card is 4 inches.

Figure 0-7 shows how the maximum trace length is segmented between the motherboard and the expansion card.

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 from those shown in Figure 1-7. The amount of crosstalk depends on trace length and trace spacing. For 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 reduced to a maximum length of 3.0 inches.



A = Expansion card timing budget = $0.7nS = \sim 4.0$ inches
 B = Motherboard timing budget = $1.65nS = \sim 9.5$ inches

Figure 0-7: Timing Budget Between For GC, AGP And Core Chipset

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

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 AGP Interface Specification. Of the maximum 0.7ns, the motherboard is allotted 0.5ns and the add-in card is allotted 0.2ns. When the effects of varying board impedances and crosstalk are considered, the following trace length constraints are recommended.

Add-in Card: 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 from 2.0 to 3.0 inches in length.

Motherboard: The data lines should be kept at $\pm 0.0, -0.5$ 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.

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 close to 0 inches as possible to provide timing margin.

Strobe Trace Routing Considerations

Since the strobe signals (AD_STB0, AD_STB1, and SB_STB) act as clocks on the source synchronous AGP bus special care should be taken when routing these signals. It is recommended that the strobe signals be routed in a 1:2 trace width/space ratio relative to the other signals of the bus. For example, in a bus where the other signals are routed with 6 mil trace width and 6 mil spaces between, the strobe signals would be routed with a 12 mil separation on both sides. 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.

Clock Routing and Skew

Clock skew between the AGP compliant graphics chip and the core chipset must be held to under 1ns. The base board is allotted 0.9ns of which a portion must be allotted to the clock generator circuit. The expansion card is allotted 0.1 ns of the total skew requirement. The diagram below shows the trace segments that make up the clock routing.

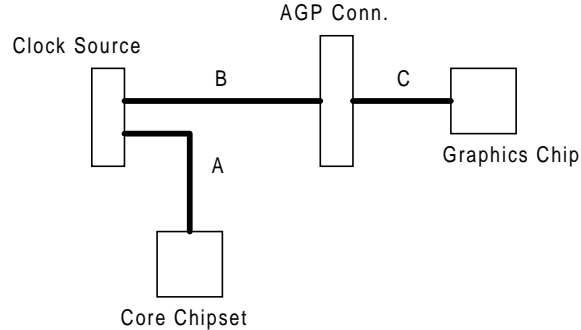


Figure 0-8: Required Trace Length To Minimize Clock Skew

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 $0.6\text{ns} \pm 0.1\text{ns}$ by the AGP interface specification. It should be noted however that add-in card vendors should center the clock trace electrical length at 0.6ns at typical board impedance so that variation in board impedance will be covered in the $\pm 0.1\text{ns}$ spec.

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

$$-(1.0\text{ns} - T_{\text{cgs}}) \leq A - (B + C + T_{\text{con}}) \leq (1.0\text{ns} - T_{\text{cgs}})$$

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 AGP connector

C = The electrical length of the clock trace on the add-in card which from AGP interface specification is set at $0.6\text{ns} + 0.1\text{ns}$.

T_{con} = The delay through the AGP connector. We have found this to be 0.15ns including crosstalk effects.

T_{cgs} = The clock generator skew between the clock signal to the chipset and the clock signal to the AGP connector.

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

The combination trace of $B+C+T_{\text{con}}$ should be matched as close to trace A as possible.

Interface Signaling Requirements

All AGP signals are +3.3V compatible signals. No +5V signals are specified in the AGP bus environment. The master and target device must be capable of supporting the 3.3V signaling environment. The interrupt signals from the AGP bus must interface to the PCI bus interrupt controller. This controller and the PCI devices may be a +5V devices. It is the requirement of the motherboard designer to properly interface the AGP 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 AGP 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 2 AGP 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 3 PCI slots (and 1 AGP 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.

IDSEL Routing For Add-in Cards

As it is described in the *AGP Interface Specification, Revision 1.0*, page 96, IDSEL is not a pin on the AGP 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 a 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.

IC Packaging Considerations

Due to the high A.G.P. pin count, there are two packages that most manufacture may consider. One is the QFP and the other is BGA. This section will discuss some generic characteristics of signals that packaging designers may need to take into consideration before selecting and designers their package.

Signal Skewing

Time skew between signals depends on differences between signal loading it 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.

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, it 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.

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 signals traces can be reduced by the presence of a ground plane.

QFP vs. BGA

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.



Chapter 2

Board Design Guidelines

These guidelines are primarily for Accelerated Graphics Port (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.

There are three ways that these guidelines can be used:

- The line length recommendations can be followed as they are described in Table 0-2 and Table 0-3. These rules should allow enough variation as to meet most designs.
- The Simulation Results section 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 taking the timing numbers shown in Table 0-10 or Table 0-11 and applying them to your own design.
- The Simulation Techniques section can be used 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.

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. These skews consist of various components that include line length mismatch, capacitance loading variations, crosstalk and the voltage level that a driver starts driving from. Figure 0-1 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. Most of these effects are independent of each other.

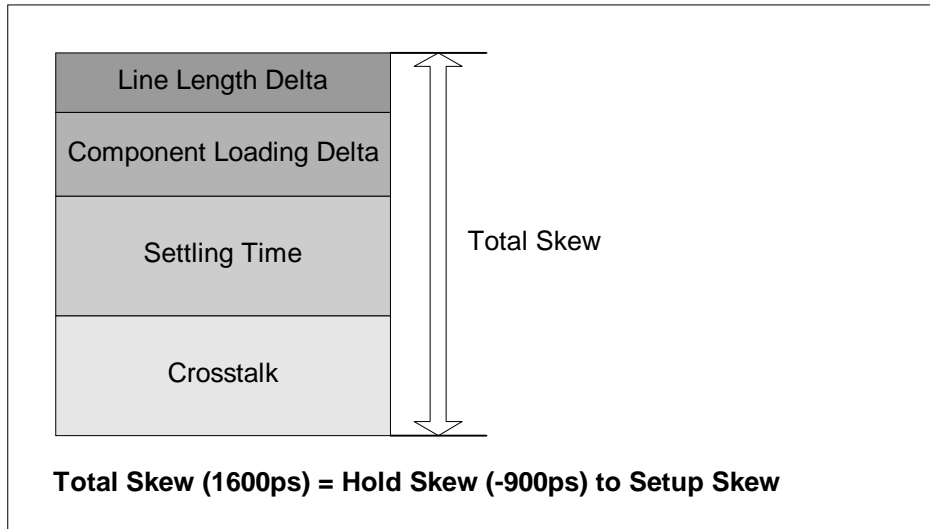


Figure 0-1: Example of Skews

Skew is defined in the formula $T_{skew} = T_{f_{data}} - T_{f_{strobe}}$. 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 this document the term “data” refers to AD[31::0], C/BE[3::0]# and SAB[7::0]. The term “strobe” refers to AD_STB[1::0] and SB_STB. When the term data is used it is referring to one of three groups of data as seen in Table 0-1. When the term strobe is used it is referring to one of the three strobes as it relates to the data in its associated group.

Table 0-1: 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

Source Synchronous Recommendations for Add-in Card

Table 0-2: Add-in Card Recommendations

Width:Space	Zo	Trace	Line Length	Line Length Matching
1:2	50Ω to 85Ω	Data / Strobe	0.0in < 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 50Ω to 85Ω. 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. 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.

Source Synchronous Recommendations for Motherboards

Table 0-3: Motherboard Recommendations

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

The motherboard needs to have an impedance range of 50Ω to 85Ω. This range is used to cover design targets and manufacturing tolerances. All lines should be at least 1.0 inch in length. 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 the 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 a 1:2 trace spacing. 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 0-3 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. These timings also allow lines to neck down so that they may break out from a component

In all cases it is always 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.

Control Signal and Clock Recommendations

Table 0-4: Control Signal Line Length Recommendations

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

Simulation Techniques

Introduction

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 1.0*. Figure 0-2 is a diagram of the topology used in this analysis.




Figure 0-2: 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 show overshoot and undershoot values approaching 2.1 Volts. 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 1.0* requirements. Settling time was factored into both the common clock and the source synchronous results.

Methodology

There are three main factors that define the solution space with this design, common clock flight times, source synchronous flight time skews and signal quality. The solution space with regard to common clock flight times is found by determining the topologies that results in lowest allowable flight times and the highest allowable flight times. The solution space with respect to flight time skews consists of all designs in which the flight time mismatch between a strobe and its associated data's 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 need to 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 the either 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.

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 1.0*. These models also used clamps. 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 8pF for the slow model. When modeling the strobe, these capacitive values were varied by +2 pF and -1 pF.

Interconnect Models and Crosstalk Effects

Both the motherboard and the add-in card traces were 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-induced coupling from other traces. Trace geometries that are typically used for motherboards and add-in cards, will usually yield impedances in the range of 60-75Ω. Manufacturing tolerances can be expected to add another 15%, increasing this range to approximately 50-85Ω.

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 0-5.

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 30-120Ω (including all of the above effects). The geometries listed in Table 0-2, Table 0-3 and Table 0-4 achieve this goal.

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

Table 0-5: Interconnect Impedance and Velocities

	Board Condition	Strobe		Data		Mode
		Z0 (Ω)	S0 (ns/ft)	Z0 (Ω)	S0 (ns/ft)	
6mil x 12mil 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
6mil x 6mil 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	

Simulations 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. Early 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 4 velocity & impedance combinations of motherboard, 4 velocity & impedance combinations of add-in card and 2 combinations of connectors. The four combinations of board interconnect were fast & 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.

All 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.

Additional simulations were run to find the Tco for the buffer models that were used. These used a topology that consisted of a buffer driving a 10 pF reference load. Tco is measured by recording the time that the load pin reached 1.32 V. These values were used in the flight time and flight time skew calculations.

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 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 6mil trace with 12mil spacing (air gap) or a 5mil trace with 10mil 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 be the two most popular methods of specifying PWBs.

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 then can be archived 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 were analyzed to determine the maximum line length that can be achieved for these spacing. 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.

Flight Time

A previous analysis had shown that the minimum flight time would be when the add-in card data line length would be at its shortest length (0.0 inches). Also the minimum flight time would be 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 would be when the add-in card data line length would be at its longest length (3.0 inches). Also the maximum flight time would be achieved with the slow buffer.

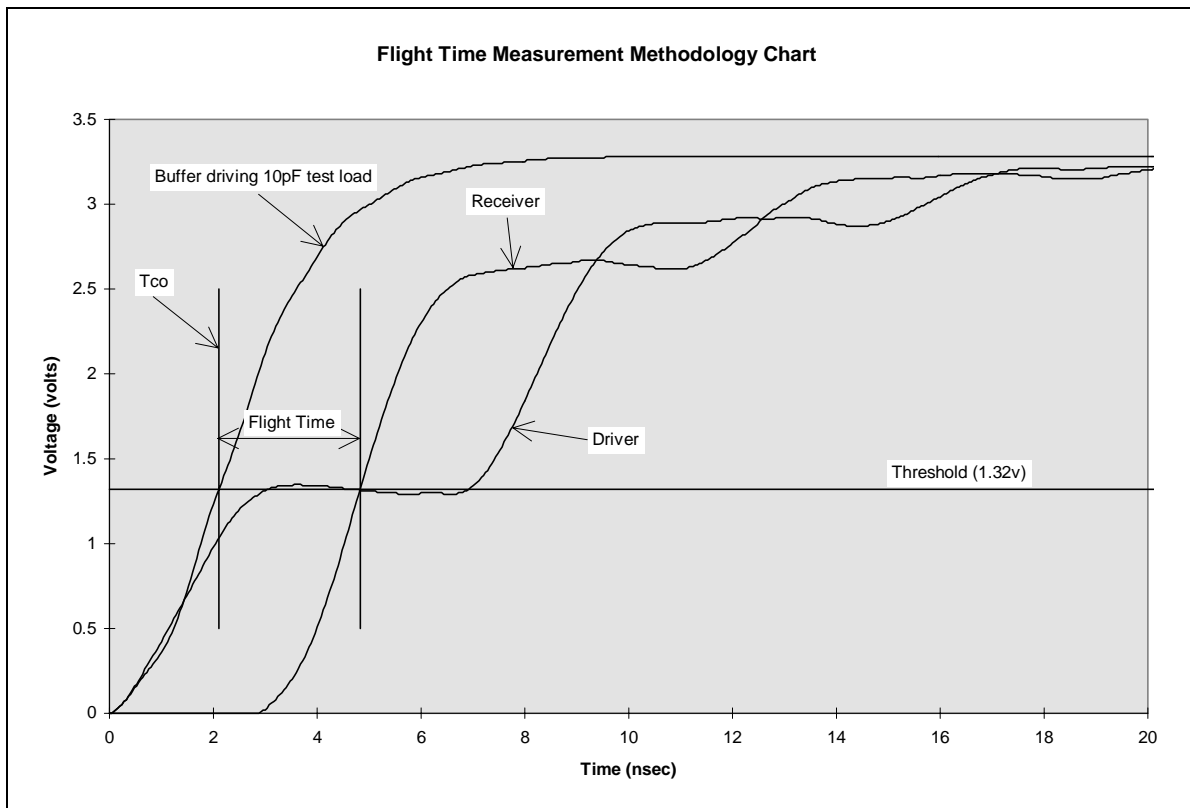


Figure 0-3: Flight Time Measurement

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 and the maximum line length on the motherboard for A.G.P. 1X could be found. Figure 0-3 is a chart that shows how flight time is measured.

Flight Time Skew

The flight time skew simulations needed 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 has on the flight time skew by comparing settled signals to unsettled signals.

For the source synchronous signals, the various parameters have an accumulative effect on the total flight time skew. Each of these parameters will have its own effect on this skew. The maximum flight time skew for a point may not necessarily be the same point that any of the individual parameters will be having its maximum effect on the skew. This means that the simulations needed to look at a large number of variations.

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 attributed by 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 0-5. 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 are to include 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 0-5. The add-in card uses the 1:2 spacing and the motherboard uses 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 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.

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.

Signal Quality

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

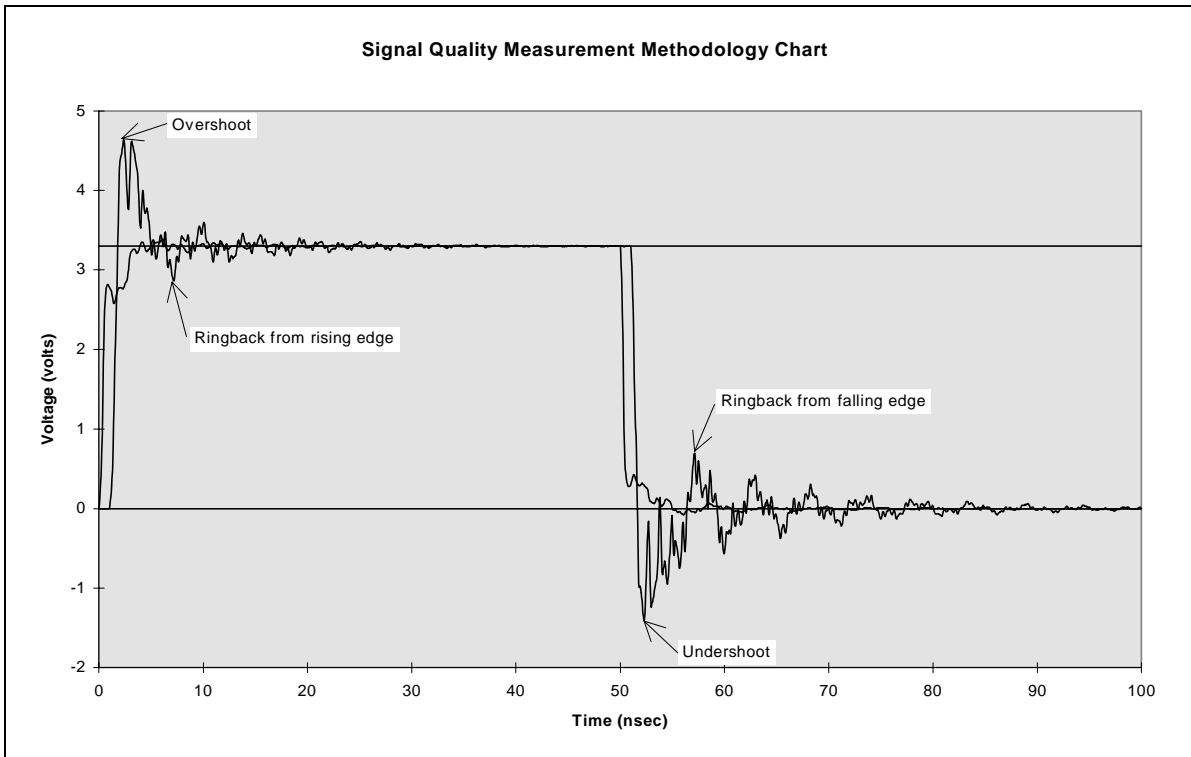


Figure 0-4: Signal Quality Measurements

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 a threshold (3.3 V). Ringback from the falling edge is the highest voltage that a receiver comes up to after crossing a threshold (0.0 V). Figure 0-4 is an example that shows how overshoot, undershoot and ringback is measured.

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 are used so that a variety of parameters can be viewed. As an example one of these 3-D graphs is shown in

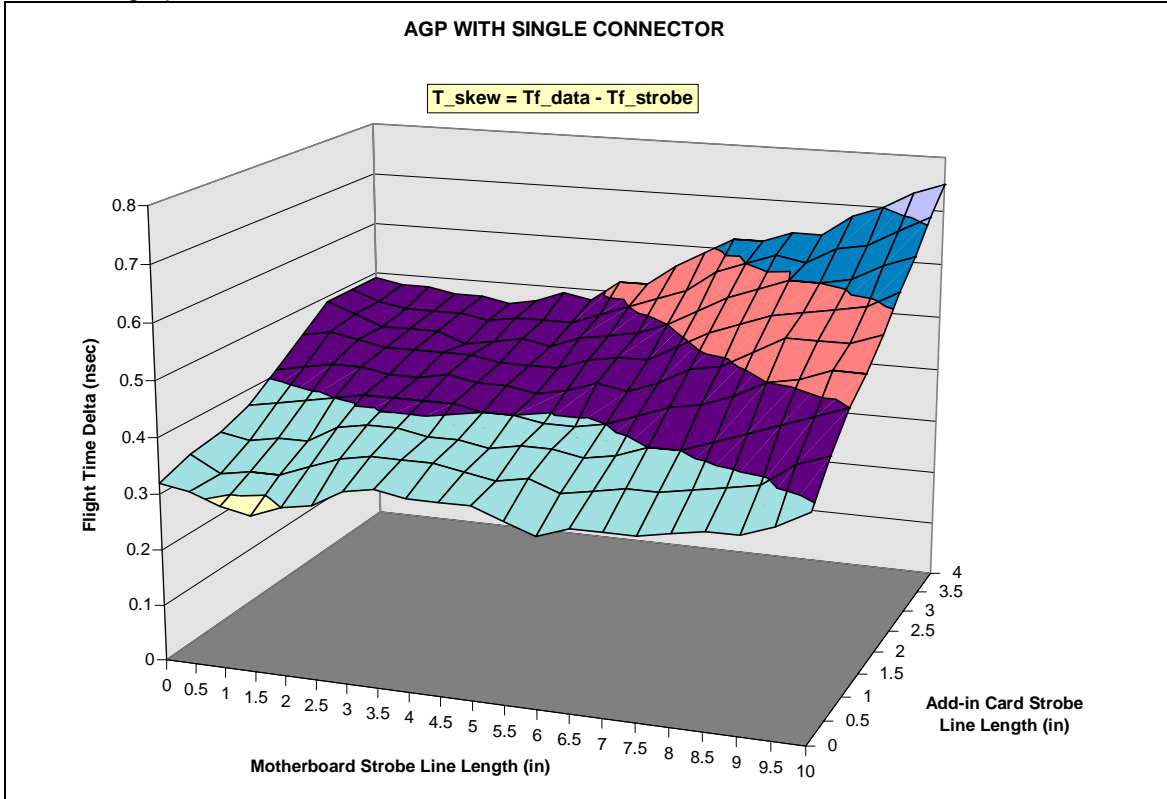


Figure 0-5: Sensitivity Analysis Chart

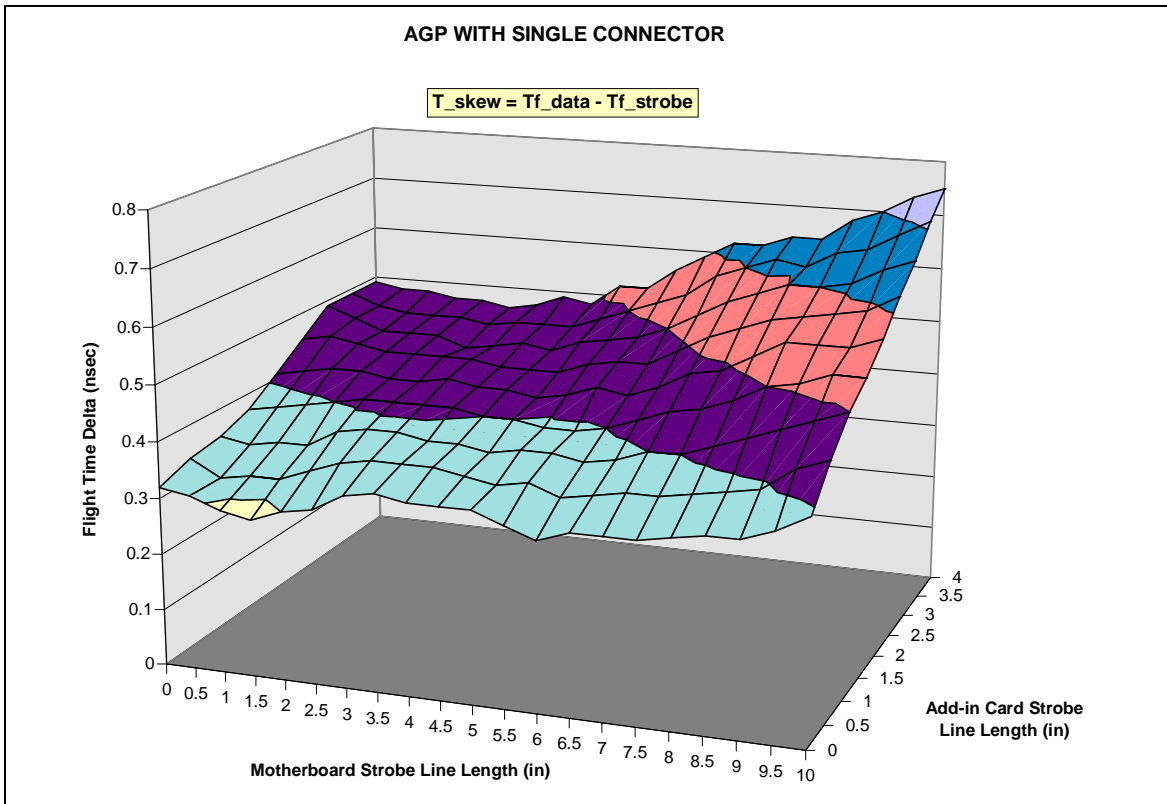


Figure 0-5: Sensitivity Analysis Chart

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 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 needs to be run for the minimum flight time. All of the signal quality measurements can be taken from the source synchronous simulations. Table 0-6 is a table of these simulations. Each of these simulations needs to be done with all the variations of interconnects shown in Table 0-5.

Table 0-6: 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	2.5in & 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

Simulation Results

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 T_{co} 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 0-7. These results show that the motherboard line lengths need to be greater than 1.0 inch.

Table 0-7: 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 were those that had an add-in card line length of 3.0 inches. These used slow buffers and motherboard interconnects that were based on 1:2 spacing.

A spreadsheet was then used to calculate the flight times for each of these cases by subtracting the T_{co} from each flight time. The maximum values from these were found by taking the maximum value for a given motherboard line length put into Table 0-8. These results show that the motherboard line lengths need to be less than 10.0 inches.

Table 0-8: 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

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 was done for each line length on the motherboard for a given line spacing is shown in Table 0-9.

Table 0-9: 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 / + 2pF	typical	compares all cases (2)
Cycle time	7ns	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 0-10 and Table 0-11.

Table 0-10 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 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 they are equal in 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 0-10: Skew from 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	-	-	-	-	-	-	-	-	-	-	-	-
Hold (ps) strobe +0.5 in	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 0-11 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 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 they are equal in 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 0-11: 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	

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.



Chapter 3

AGP Thermal Design Guidelines

AGP Thermal Overview

The general trend for high performance AGP-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 AGP 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 0-1. These power levels will create localized power densities that require thermal enhancement of component packages and add-in cards.




Figure 0-1: AGP Add-In Card and Graphics Controller Power Trend

Thermal Design Recommendations

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 AGP 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.

Power Budgets:

In 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 100Mhz, can draw power ~ 1.0 Watts (Max). Considering that an average add-in graphics card can contains 4MB 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 arise in 3D market, 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

Recommendations for OEM System Designers

The OEM system integrator is strongly recommended to optimize the thermal environment for the AGP-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 may reduce overall system cost.

Recommendations for Thermal Performance

ADD-IN CARD

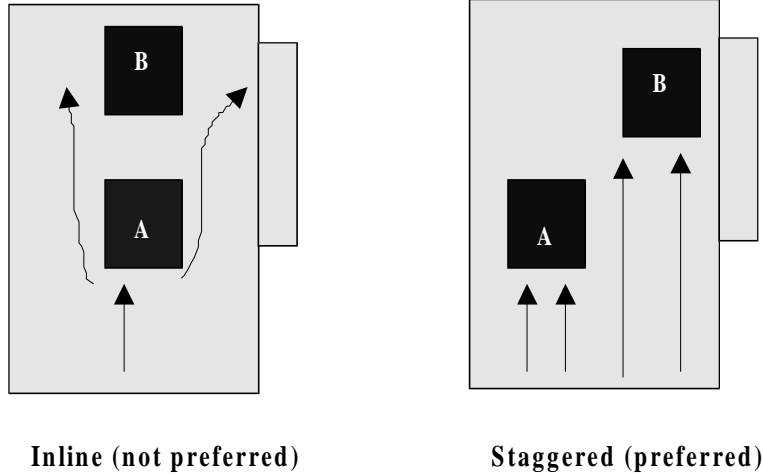
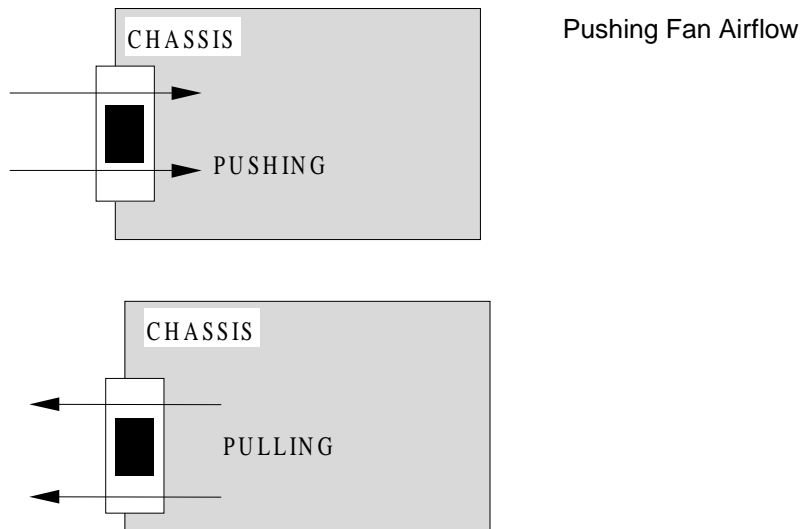


Figure 0-2: Airflow Over Add-In Card Components

- If airflow is required over a densely populated add-in card, the **In**line orientation shown in Figure 0-2 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 0-2, components will equally benefit from the airflow.

FAN DIRECTION

Figure 0-3: Pulling and



PUSHING FAN

- The PUSHING fan shown in Figure 0-3 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 0-3 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 box.

VENTING

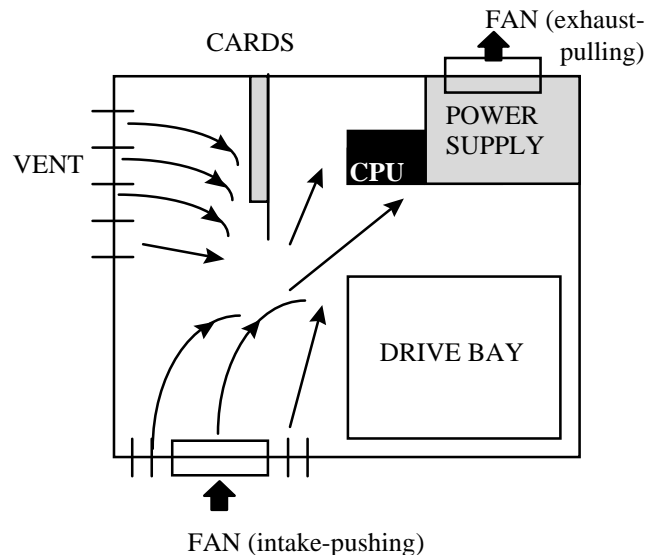


Figure 0-4: Adequate Venting Air

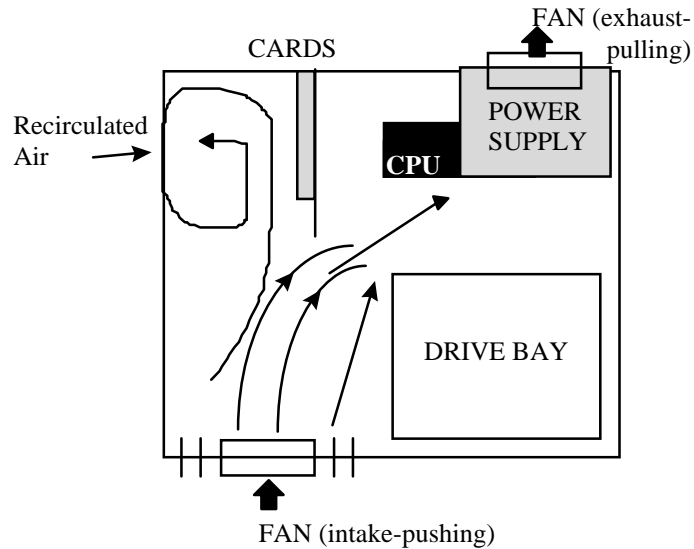


Figure 0-5: Inadequate Venting

- 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 0-4. Inadequate venting will cause internal air to recirculate as shown in Figure 0-5. 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 large as possible without conflicting with EMI and safety

INTERNAL AIR FLOW

- A full-length add-in card next to an AGP card is likely to block airflow to the graphics card as well as dissipate additional power. The AGP 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 AGP slot.
- Cabling arrangement must not obstruct any component airflow or venting holes to ensure proper cooling.

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 AGP-enabled card. For this specific chassis form factor, the side panel vent shown in Figure 0-7 were added and resulted in a 20 C drop in the internal ambient temperature near the AGP-enabled card. The side panel vents enabled the specific chassis design to cool a 20 W distributed thermal load card representing an AGP-enabled add-in card with little impact to the overall system thermal environment. The dimensions of the side panel vent are shown in Figure 0-7. An example of NLX rear vents are shown in Figure 0-8.




Figure 0-6: Example NLX Air Flow




Figure 0-7: Example NLX Side Panel Vent




Figure 0-8: 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 AGP 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 AGP-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 0-9.




Figure 0-9: 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.

Reference Documents

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