Notes on AGP Interface Architectures and Motherboard Design with SSC

ANDREW VOLK
Principal Engineer
Computing Enhancement Group
Intel Corp.

HOWARD HECK
Manager, Interconnect Modeling and Design
Platform Architecture Lab
Intel Corp.

Rev. 1.0
February 1998

Abstract: *Section I addresses the impacts of the induced Spread Spectrum Clocking (SSC) skew for various AGP architectures. Section II presents the trade-off considerations to adjust AGP layout design to compensate for the SSC skew. Adjustment examples are also given to illustrate the methodology.*

I. SSC AND AGP INTERFACE ARCHITECTURES

The questions that need to be asked are where will the skew that is induced in a PLL by Spread Spectrum Clocking (SSC) modulation appear in the system and what impacts on the timing will that skew have [1]? Which parts of the system have to accommodate the skew? The answer depends on what part of the system output of the PLL drives. Note that a system may have different architectures on each end of the AGP interface and all combinations have to be accommodated.

The following shows three example architectures:

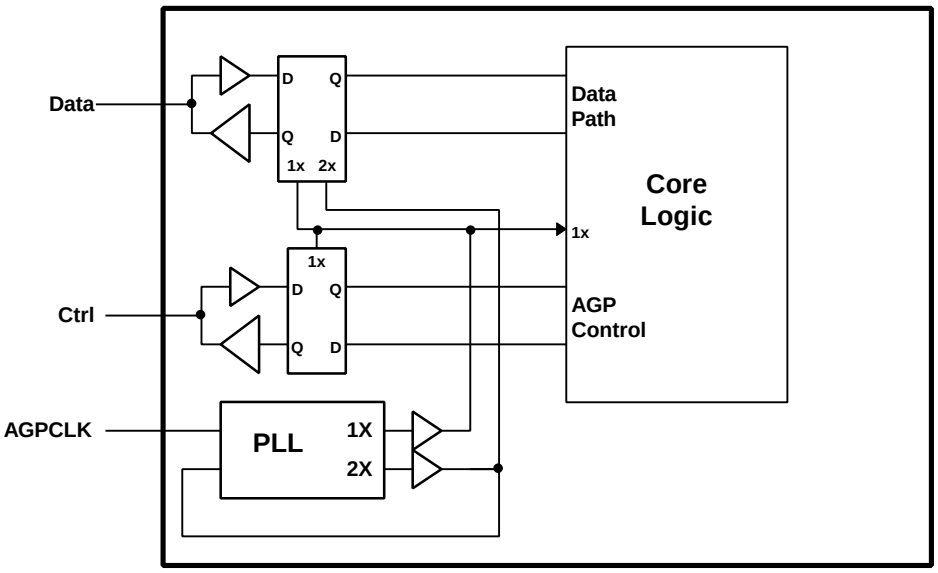

Figure 1. PLL provides all internal clocks

Figure 1 shows a system where the system clock to AGP, AGPCLK, drives the PLL, and the PLL provides timing to all of the chip. In this architecture, the PLL provides clock insertion delay cancellation and all internal clock skews are well matched.

The consequence of spread spectrum clocking (SSC) on this architecture is that the clock modulation is seen entirely on the pins. Since the PLL drives all internal clocks, all PLL skew hits the external pin timings, both output delays and input setup and hold times. There is no skew between internal clocks. ***The internal timings are unaffected and there are no timing impacts hidden within the chip to worry about. However, the skew that is generated at the pins must be absorbed in the system timing margins.*** (See Section II.)

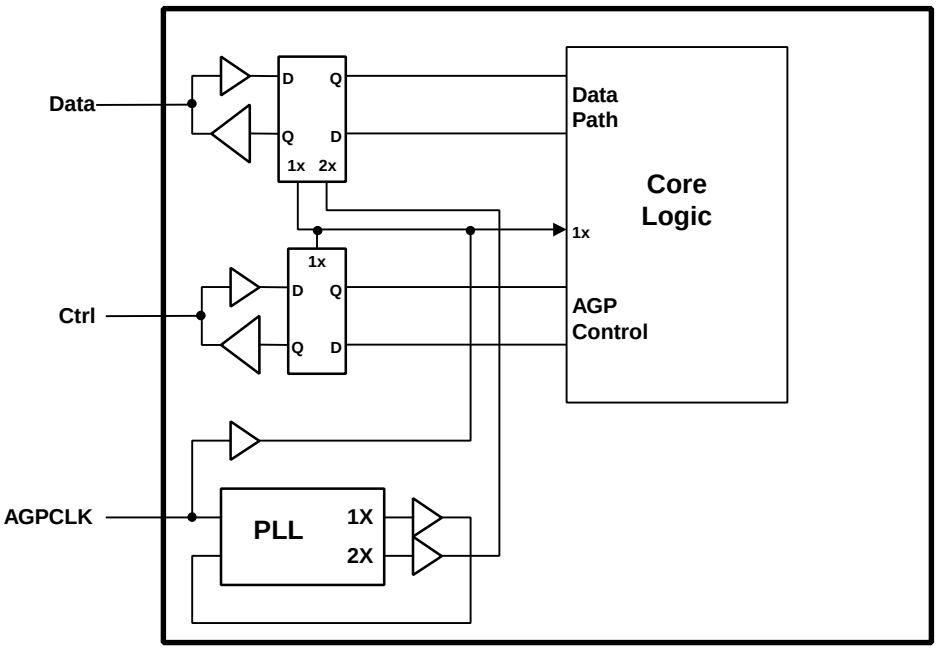

Figure 2. AGPCLK drives Core and 1X, and PLL drives 2X circuits

Figure 2 shows a configuration where the AGP system clock drives all circuits that run at 66 MHz and the PLL is used to generate only the 133 MHz clock domain. AGPCLK drives the PLL and a clock buffer for any 66 MHz uses, including AGP-1X interface timings and core logic. The PLL provides timing for only the AGP-2X interface. The internal interface clock is delayed by a clock buffer insertion delay. This gives more time for input setup, but hurts hold time and the output delays from the AGP clock. The PLL is used as a frequency multiplier for AGP-2X operation. The phase of the 2X clock to the internal AGPCLK is set by tweaking delays as required by design. Only the 2X mode timings are affected by PLL effects.

The consequence of spread spectrum clocking (SSC) on this architecture is that skew appears between 2X PLL output and any form of the AGPCLK, internal or external. SSC modulation is not seen on the pins for any AGP-1X timing. AGP-2X timing that relates data to strobe edge are unaffected because modulation does not materially affect edge-to-edge timings of PLL. Skew will be visible in the strobe-to-AGP clock times since the strobe is timed from the PLL output clock. There is no skew between the AGP-1X control and data paths to the internal AGPCLK clock since these paths use the same clock source. Skew does appear on internal paths as they cross between 1X and 2X clock domains, such as transmitting data from the core to the 2X data or receiving the data from the 2X input data FIFO to the core. ***These internal timing impacts are not externally visible (unless functional failure occurs) and require internal timing margin to withstand the skew. The component designer is responsible for providing this margin.***

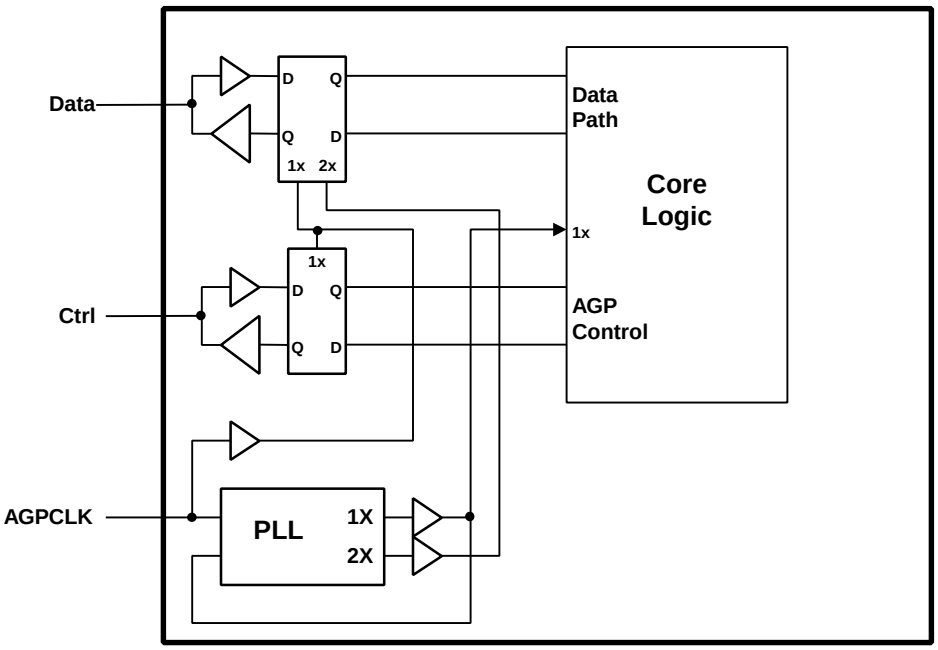

Figure 3. AGPCLK drives only 1X interface. PLL provides 2X and 1X core

Figure 3 shows a variation of the architecture in Figure 2. Here, AGPCLK directly drives the AGP-1X interface timings only. The buffer delay is made small, and provides a fast internal version of the external clock. The PLL is used to generate the core logic clock and the AGP-2X interface clock. The timing interface between the fast 1X interface and PLL clock domains is handled by internal timing tweaks and interface circuits.

The consequence of spread spectrum clocking (SSC) on this architecture is that skew appears between PLL outputs and the AGP-1X interface clock and external AGPCLK. Clock modulation is not seen on the device pins for any AGP-1X timing. AGP-2X timings that relate data to strobe edge are unaffected because modulation does not materially affect edge-to-edge timings of PLL. Skew will be visible in the strobe to clock times since the strobe is timed from the PLL output clock. Skew will appear between the AGP-1X control and data pins to and from the core. Skew also appears on internal paths as they cross between 1X and 2X clock domains, such as transmitting data from the core to the 2X data outputs or receiving the data from the 2X input FIFO to the core. ***These internal timing impacts are not externally visible (unless functional failure occurs) and require internal timing margin to withstand the skew. The component designer is responsible for providing this margin.***

Other configurations are also possible. For instance, the PLL could drive the local memory interface, or a separate PLL could do that. On graphics controllers, there is a PLL for generating the video data output, but that is derived from a fixed frequency interface. (At least it had better be fixed and not modulated or the screen is going to be severely jittery.) All of these provide the opportunity for a timing interface across which the modulation-induced skew can be seen and for which there must be skew tolerance.

II. SSC and AGP Motherboard Design

Implementation of spread spectrum clocking (SSC) in an AGP-1X system can induce additional skew in the I/O timings for the chipset and/or add-in graphics chip. The additional skew can cause the components to violate the AGP-1X timing specs. When such violations occur, the interconnect design on the motherboard must be adjusted to ensure proper operation. This section describes the trade-off considerations, and provides a method for determining how to properly adjust the interconnect design rules as a function of the additional skew induced by SSC.

2.1. AGP-1X Timings and Layout Guidelines

In 1X mode, AGP operates as a synchronous (a.k.a. common clock) bus. System timings are limited by absolute delays, rather than by skews (as in AGP-2X). AGP-1X system timings are summarized in Table 1. Both the transmitter and receiver may be affected by SSC modulation, depending on the architectures of the devices at each end of the AGP interconnect. Any skew generated by an AGP device that is beyond spec will have to be accommodated in the interconnect skew.

Table 1. AGP-1X system timing specs

	Setup	Hold
Transmitter:		
Data	6.00 ns	1.00 ns
Control	5.50 ns	1.00 ns
Interconnect	2.50 ns	0.00 ns
Receiver:		
Data	6.00 ns	1.00 ns
Control	5.50 ns	1.00 ns
Clock Skew	1.00 ns	-1.00 ns

Note: For complete specs, refer to the AGP specification [2].

Interconnect delays are separated into two pieces: motherboard and add-in card. Add-in cards are allocated 700 ps of delay in the AGP specification, which supports trace lengths in the range of 3” for a 6 mil trace/6 mil space board [3]. Unless the add-in card trace lengths are known, system designers must assume that the card has been designed to the worst case specification value.

Motherboards are allocated 1.8 ns for interconnect delay, including crosstalk.

2.2. Motherboard Interconnect Adjustment Method

When the additional skew induced by SSC causes an AGP device to violate the specification, the motherboard trace length guidelines must be adjusted to compensate for the skew, as identified in the above section. This is accomplished by "de-rating" the layout guidelines as a function of the magnitude of the component violation. The method is described in detail below.

Referring to page 33 of reference [3], the maximum motherboard line length guideline is 9.5 inches, which gives a 90 ps margin to the 2.50 ps interconnect delay specification. This applies to a 65Ω board with 6 mil trace/12 mil space routing. For 6 mil trace/6 mil space routing, the maximum length shrinks to 4.5 inches due to source synchronous clocking skew (not flight time). The minimum recommended trace length is 1.0 inch on the motherboard for both 6/6 and 6/12 trace/space. The trace lengths and interconnect delays are summarized in Table 2.

Table 2. Motherboard signal propagation velocities

		6/6 routing	**6/12 routing**
Minimum	**Trace Length**	1.0 in	1.0 in
	Delay	—	0.02 ns
Maximum	**Trace Length**	4.5 in	9.5 in
	Delay	—	2.41 ns

In addition to the layout guidelines, we need to know the effective signal propagation speed (minimum and maximum). Reference [3] provides the following values, which include the effects of crosstalk:

Table 3. Motherboard signal propagation velocities

	6/6 routing	**6/12 routing**
Minimum Velocity	2.24 ns/ft (187 ps/in)	2.12 ns/ft (177 ps/in)
Maximum Velocity	1.28 ns/ft (107 ps/in)	1.38 ns/ft (115 ps/in)

The data in Table 2 and Table 3 can be used to "de-rate" the minimum and maximum trace lengths, based on the amount of specification violation using the equations listed below:

$$L_{SSC} = L_{AGP} - \frac{\Delta T_{SSC}}{t_p}$$

where:

L_{SSC} = the adjusted trace length (includes SSC impacts)

L_{AGP} = the original AGP layout guideline

ΔT_{SSC} = the amount of violation of the component timing spec

τ_p = the propagation delay

Note that the adjustments must be applied in the appropriate direction. I.e. violations to the minimum driver delay and/or component hold spec require that minimum trace length must be increased. Violations of maximum component specs (driver delay, receiver setup) require that the maximum trace length must shrink.

Figure 4 and Figure 5 show the examples of the adjustments. Notice that the 4.5″ maximum motherboard trace length of 1:1 space routing is limited by the crosstalk induced skew, not flight time. Thus, there is no routing impact until SSC skew exceeds 900 ps, which starts to reduce flight time budget. For the impact on the minimum motherboard trace length, SSC skew cannot exceed the 400-ps holdup timing spec.

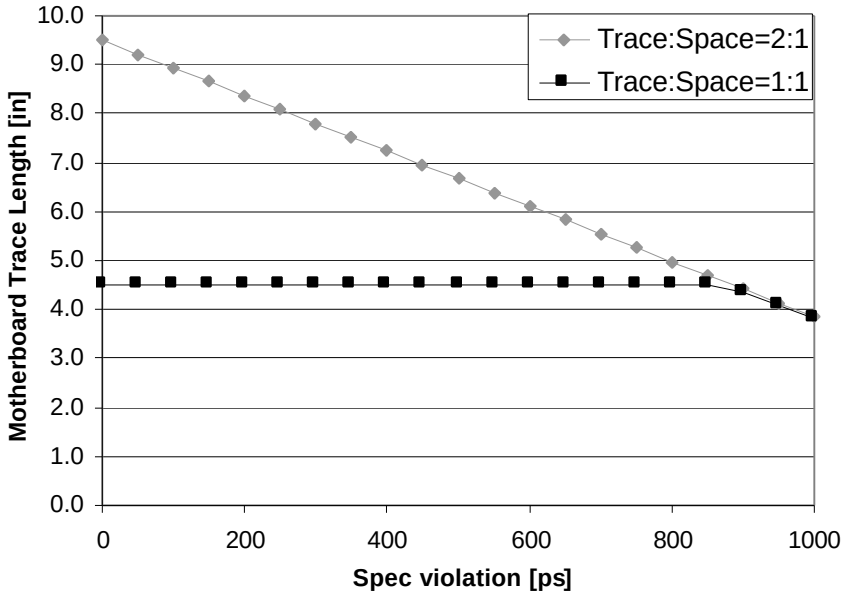

Figure 4. Maximum MB trace length as a function of skew violations

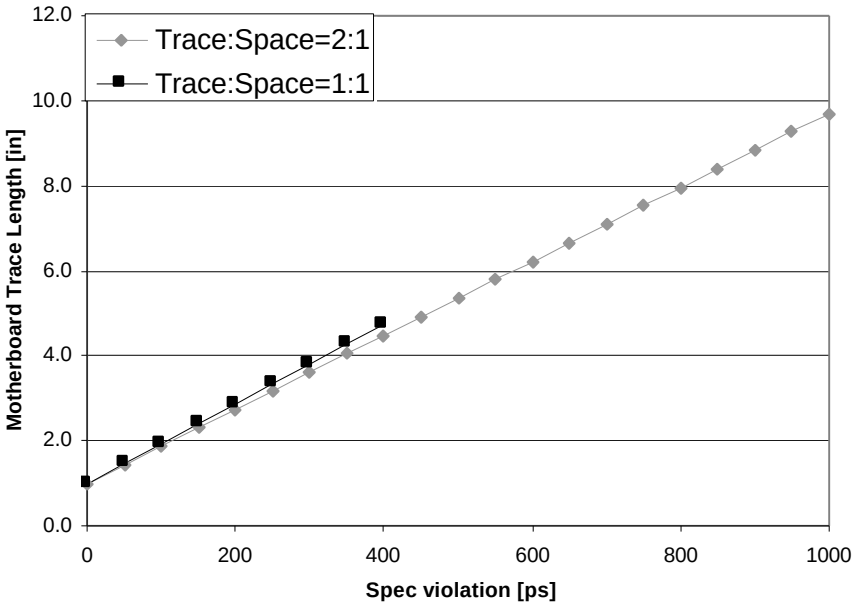

Figure 5. Minimum MB trace length as a function of skew violations

Notes:

(1) The plots do not account for the margins that exist at the minimum and maximum recommended lengths.

(2) Interconnect simulations are still recommended to ensure proper operation.

References

1. Notes on SSC and Its Timing Impacts, Revision 1.0, February, 1998.
2. A.G.P. Interface Specification, Revision 1.x.
3. Accelerated Graphics Port (A.G.P.) Platform Design Guide, Revision 1.1.1, June 1997.