adaptec

Adaptec Technical Memorandum   NCITS T10/97-268r1

To:	SPI-2 Workgroup, T10 Committee	**Date**:	November 3, 1997
From:	Tak Asami	**File**:	t1097268r1.doc

Subject: SPI-2 Timing Tolerance

This is a formal proposal to revise timing accuracy tolerance specifications within SPI-2 document.

Background:

In the timing specifications defined in SPI-2 draft document (currently rev 16), I feel Receive / Transmit Period tolerance is over-specified and may be making the implementations difficult or unnecessarily costly.

On Page 90, Table 42 lists "Receive Period Tolerance" and "Transmit Period Tolerance".
They specify the required accuracy of the clock frequency for receiver and transmitter circuit. Basically, when a certain synchronous speed is negotiated between an initiator and a target, any instantaneous cycle time between each transfer should stay within the "transmit tolerance" and receiver should tolerate up to "receive tolerance" range.

Current draft specify 0.5% for receiver tolerance and 0.25% for transmitter tolerance. The figures are ported directly from SPI / Fast-20 documents. They are meant to describe the DC offset from the nominal timing over long period of time, and it is fine for that purpose.
Now enter Fast-40 and beyond. Including ours, many implementations I am aware of utilize PLL to synthesize required frequencies. I expect the trend to continue.
Fundamentally there is nothing wrong with PLL; they are very accurate and stable over a long period of time. But on a cycle by cycle basis, you would expect jitters. This is a nature of the circuit and there is no two ways about it. Bottom line: they cannot meet 0.25% period tolerance edge to edge.

The point is, though, for most, if not all, implementations at Ultra SCSI and beyond products, this tight tolerance is absolutely unnecessary.
Including ours, these high-end SCSI implementations can tolerate receiving signals as much as 50% faster than negotiated. So if the incoming signal had several percent of jitter in it, these devices won't feel a thing and I presume future products remain that way.
Another way of saying this is, that if your design has to tolerate 0.25%, then it'll tolerate several percent higher frequency as well, because you will use extra clock edges to absorb those jitters.
Therefore, the conformance to the current tight timing tolerance specification is no longer a requirement for inter-operability between SCSI devices.

Since I consider this technique a enabler for low cost implementations of the high performance devices, I would like to have the specification changed to accommodate usage of PLL-based timing reference.

One word of caution is that this 1nsec tolerance is not to be considered a free pass to perform 24nsec cycle transfer. The "average" transmission period still should be as close to negotiated value as (25nsec for Fast-40) possible. It should also be noted that there is no relaxation to set up and hold time requirement due to this timing tolerance relaxation. The transmitting agent shall still meet all the relevant timing specifications.

Proposal:

Referencing SPI-2 draft revision 16, on page 90, Table 42, change the Transmit Period Tolerance for all transmission rate except for Asynch, to be 1.0nsec.
Change the Receive Period Tolerance to 1.1nsec for all speed except for Asynch.