adaptec

Adaptec Technical Memorandum

NCITS T10/97-268r0

To: SPI-2 Workgroup, T10 Committee **Date**: October 31, 1997

From: Tak Asami **File**: t1097268r0.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 15, going on 16), I feel Receive / Transmit Period tolerance is over-specified and may be making the implementations difficult or unnecessarily costly.

On Page 94, Table 41 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 were fine for up to Fast-20 speed, where (a) transfer cycle time (nominal) is up to 50nsec, and (b) most implementations are using crystal clock references.

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 if we go Fast-40 and beyond.
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.

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 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.

Proposal:

1) Remove "Transmit Period Tolerance" and "Receive Period Tolerance" from the Table 41 and move the subject to Clause 11.5.2.12 "Synchronous Data Transfer Request".
2) Delete the reference to “Receive Period Tolerance”, for it deals with internal implementation of the receiver device and does not belong to a standard.
3) Change the transmitter timing accuracy for Fast-40 and beyond to be 10% (peak jitter, edge-to-edge). This means that for 40.0MXfr/sec Ultra-2, the minimum cycle time between REQ-REQ or ACK-ACK is 25nsec x 0.9 = 22.5nsec.

NOTE: currently, there is no scientific basis for 10% figure. I just introduced an arbitrary large number, and suggesting today's implementations easily tolerates even that. I am actually happy with a number as low as 5%, but we do not need to tighten the specification when it is not needed to insure interoperability.
I welcome any reasonable counter-proposal.

This enables the usage of monolithic PLL based frequency synthesizer, which can keep the cost of implementing high end SCSI devices low.