

X3T9.2/87-33  
Mar. 10, 1987

To: X3T9.2 Committee (SCSI)

From: Gerry Houlder MPI/CDC

Subject: Clarify Synchronous Data Transfer Description

We have recently discovered that the existing wording describing synchronous data transfers has caused implementation differences between several SCSI chip vendors. I believe this is a bad situation and has already resulted in severe performance limitations under some common circumstances.

Specifically, use of the term "REQ pulse" has caused at least one vendor (Western Digital) to believe that both the leading and trailing edges of the REQ signal must be received before the ACK signal can be asserted by the initiator. Likewise, both edges of ACK must be received before another REQ can be sent by the target. With such a design, the REQ and ACK pulses can never overlap and the turnaround delay is determined by the pulse width plus the cable delay in addition to any chip delays.

Another vendor (NCR) only requires the leading edge of REQ to be received before the ACK signal can be asserted by an initiator and that only the leading edge of ACK must be received before another REQ can be sent by the target. With such a design, the REQ and ACK pulses can overlap and the turnaround delay is determined by only the cable delay and chip delays. If the REQ and ACK pulses are each 120ns wide (the minimum is 90ns, so this is a safe typical design value), this design can turn around a data byte in 240ns less time if the chip delays are about the same for both designs.

86  
This design difference is especially apparent at an offset value of 1. The NCR design allows significantly faster transfers in this case. In fact, the WD design can result in transfers slower than asynchronous transfers can achieve. For offset values of 4 or 5, the maximum transfer rate can be achieved because the offset helps to hide the wait for the entire pulse to be received. However, since the current NCR design only supports an offset of one, the mixing of an NCR based initiator with a WD based target has dismal performance in synchronous mode.

This difference of interpretation could result in chip designs that would not work together at all. If a designer interprets the standard to prohibit the overlapping of REQ and ACK pulses in synchronous mode, he may decide to abort a transfer if such a timing relationship occurs or may not sense the leading edge of the handshake pulse at all if it overlaps his own pulse. This could result in a hang condition.

I cannot fault the WD designers for their interpretation of the synchronous transfer wording in the SCSI Standard. Their interpretation is reasonable for the wording given. We have also received information that the Fujitsu chip designers have made the same interpretation that WD did, but we have not personally verified this. Perhaps someone else can either verify or refute this.

I am proposing these wording changes to clarify the synchronous transfer description to prevent this problem in the future.

Section 5.1.6.2, second paragraph: the second sentence states that "... the target shall not assert REQ until the next ACK pulse is received." Change this to read "... the target shall not assert REQ until the leading edge of the next ACK pulse is received."

Section 5.1.6.2, fourth paragraph: add the following sentence between the first and second sentences; "The ACK signal may be asserted as soon as the leading edge of the corresponding REQ pulse has been received."

Section 5.1.6.2, sixth paragraph: the second sentence begins "After receiving a REQ pulse, ...". Change this to read "After receiving the leading edge of a REQ pulse, ...".