



D. W. (Bill) Spence

Member, Group Technical Staff
Computer Systems Division
Data Systems Group

12501 Research Boulevard • P.O. Box 2909 • MS 2091 • Austin, Texas 78769
512 250-6627 • Fax: 512 250-7479 • TWX: 910 867-4702 • Telex: 73324

X3T9.2/91- 18

MEMORANDUM -- 18 Feb 1991

TO: John Lohmeyer, Chairman, X3T9.2

FROM: Bill Spence, TI

SUBJECT: Glitch Report--Silicon and Stub (Single-Ended SCSI Operation)

THE BOTTOM LINE

Even though we have achieved very solid single-ended operation through the advances achieved by the studies of the past year, questions have remained about the glitches which degrade the descending and ascending wavefronts of our signals. Experimental investigation and theoretical analysis indicate that between the two prime suspects, stub action and silicon action, in many cases the silicon action probably predominates.

SILICON-INDUCED GLITCHES

In various studies presenting actual waveforms taken in single-ended operation, a recurring theme has been the raggedness in the lower half of descending waveforms and the upper half of ascending waveforms. There is always a glitch or a series of glitches that delays the waveform in reaching promptly its target level--either well below 0.8 volts or well above 2.0 volts.

Several indications over many studies have suggested that the silicon on the bus, even when not driving, plays more than a passive role, adding significant handicap to the lowest part of the descending waveform and giving a boost to the first level reached by the rising waveform, even while possibly adding a slightly delaying glitch. In an effort to duplicate previous studies without incurring all the confusion imposed by intermediate nodes along the bus, Tessco of Georgetown, Texas, contributed a 68' cable assembly incorporating Madison 4099 cable. This permitted repeating the 70' bus studies previously reported with only our host adapter (with a WD 33C93) and a single disk on the bus.

By substituting in various disks, each in an identical mounting and cabling arrangement, we developed seemingly convincing evidence that indeed the silicon may play a significant role. In picture 1, the disk is a Conner 260 MB, incorporating an Adaptec protocol chip. In picture 2, the disk is a Maxtor 760 MB incorporating an Emulex protocol chip. The signal is the -ACK signal from the host adapter as received at the disk.

In both pictures, as in all my pictures, there are at least 100,000 sweeps averaged in the waveforms. Any observed differences are real and repeatable. In picture 2, there are clearly observable glitches on both descending and ascending waveforms which are not observable in picture 1.

I hasten to point out that these glitches which the chip imposes on the waveform that it itself is receiving are probably not harmful. Even though they delay the signal just as it is clearing the prescribed signal levels, the chip itself may be insensitive to a noise spike which might arrive just at the same moment. But if other chips along a bus add their glitches to a signal as it travels along, the chip which finally must utilize the signal may be rendered much more vulnerable to double clocking from a noise spike.

It is regrettable that I did not take expanded-timebase pictures of these phenomena while I was about it. If anyone who is silicon literate wishes to investigate further, however, he will be creating his own pictures to whatever timebase he wishes.

STUB-INDUCED GLITCHES

Returning to glitches in general, the descending waveform in picture 1 is by far the most desirable I have ever observed, indicating that the six additional disks present in most of my studies take their toll. In addition to whatever harm may be ascribed to the silicon present, each of these disks scattered along the bus represents a stub on the bus, of finite length and terminated by finite chip input capacity, even though both are generally quite small.

Happily, a first look at the stub phenomena is readily available by theoretical analysis. Given a signal of amplitude V reaching a bus node where a stub goes off, bus and stub both being of characteristic impedance Z_c , the signal voltage must drop by a fractional amount f such that the original current wave V/Z_c , augmented by the reflected current wave fV/Z_c , must equal the sum of the current waves induced in both the stub and the continuation of the bus, $2(1-f)V/Z_c$; or $1+f = 2(1-f)$. Solving, $f = 1/3$. In other words, when the incident signal wave V reaches the stub node, the voltage jumps up only to $2/3 V$. This condition prevails until the signal wave going out along the stub reaches the open end and reflects back to the stub node. Additional analysis gives that the node voltage then jumps up to $10/9 V$ for the duration of another round-trip interval on the stub, after which the phenomenon is essentially all over.

A more general analysis, allowing the stub to have a characteristic impedance Z_s different from Z_c , gives the fractional drop in signal to be $f = 1/(2R+1)$, where $R = Z_s/Z_c$. This equation suggests that if in some circumstance a lengthy stub were unavoidable, its effect on the main bus might successfully be diminished by making its impedance quite a lot higher than the main bus impedance.

This effect was tested by creating a 4 foot stub at a node between two 20' cables comprising a 40' bus driving a single disk at the bus end. The stub impedance was about 105 ohms and the bus impedance was about 80 ohms, giving a fractional signal drop of 0.28. The propagation delay of the stub cable was about 1.4 ns/ft, giving a round-trip time of about 11 ns.

Pictures 3 and 4 show the descending and ascending waveforms at the stub without the stub being present. The timebase is greatly expanded, so that what were rather clean and sharp wavefronts show pronounced slope (scope rated 300 MHz, so most of the deviation from vertical should have been real). The ascending waveform reaches a first step level lower than the terminator 2.85 open-circuit voltage roughly in accordance with the ratio of the bus impedance to the 110-ohm terminator impedance (see X3T9.2/90-185R1).

Pictures 5 and 6 show both the glitch and the overshoot predicted by the theoretical analysis. Comparing the levels of the glitch to the levels which existed without the stub being present, the fractional depression is close to the 0.28 prediction, and depending on how it is measured, the glitch width is reasonably interpretable as being of about 11 ns duration.


On the same timebase, the glitch caused by the 4' stub appeared to be wider than the glitch attributed to silicon in picture 2 by possibly a few times. But consider: the stub was 48", 12 times longer than the 4" allowed by the standard, which itself is much longer than exist in actual disk implementations. My conclusion: the stub effect is insignificant; the silicon effect may or may not be insignificant.

What about the capacitive load at the end of the stub? An attempt was made to test this by loading the stub end with a 33 pf cap. It degraded the signal at the stub node somewhat, though more in other ways than in intensifying the glitch. At the stub end itself, the loading was somewhat beneficial, damping the strong ringing otherwise experienced.


CONCLUSION

I have no plans to push this investigation farther at this time. It seems to me to be well established that with good 80-ohm cables with crosstalk-protecting pinouts, well controlled stubs, and regulated 110-ohm terminators, very good single-ended performance can be achieved even with a fully loaded bus and with lengths substantially greater than 6 meters. The remaining deficiencies in the waveforms achieved--the leading edge glitches--may well be more attributable to silicon than to transmission-line effects. Further improvement in that area may or may not be achieved.


1




2




3



4



5



6

