TEXAS INSTRUMENTS

D. W. (Bill) Spence
Member, Group Technical Staff
Computer Systems Division
Data Systems Group

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

MEMORANDUM -- 27 Nov 1990 -- Revised 19 Dec 1990

TO: John Lohmeyer, Chairman, X3T9.2

FROM: Bill Spence, TI

SUBJECT: S/E Impedance Optimization and Best Case Analysis II

This revision adds the NOTE: paragraph below.

THE BOTTOM LINE:

1. Just as increasing the cable characteristic impedance, relative to the terminator impedance, improves the voltage rise achieved on release of a S/E signal, so lowering the cable impedance improves the current wave into the receiver terminator on assertion, thus improving the receiver voltage drop.

2. The optimum ratio of cable impedance to terminator impedance appears to be around 0.75, or a little less. It just happens that this is what is produced by a shielded 28 AWG non-foamed polyolefin cable of the right dimensions, working with a 110 ohm terminator.

NOTE: In the R0 version of this paper, I referred to polycarbonate cables, as I also did in an earlier paper. These were unfortunate mental lapses. Polycarbonate is a capacitor dielectric and probably would be most unsuitable for high performance cable service. As I understand it, the preferred high-performance cable insulations are polyolefins: polypropylene or high-density polyethylene. Note also that R1 of my Test Report No. 4 (X3T9.2/90-170R1) explicates the Zc/Zt ratios of the cables being reported on, which supplements the material below.

In X3T9.2/90-123R1, 31 Aug 1990, I suggested that the characteristic impedance of the single-ended bus signal line controlled how good a rise could at best be achieved on release of an asserted signal. More precisely, it was the ratio of the cable conductor characteristic impedance (Zc) to the terminator impedance (Zt) which controlled. The higher the Zc/Zt ratio, the higher the level achieved on release.

We have not addressed very vigorously the question of what controls the fall of an asserted signal as received at the other end of the bus. Not too surprisingly, I guess, this turns out to be complementary to the rise-on-release situation. When a signal is asserted at one end of a conductor, a signal wave travels the length of the bus with a current component equal to the assertion voltage fall divided by the cable impedance. When this current wave hits the far end terminator, it will drop the voltage, of course, and the amount is controlled by the product of the current wave and the terminator impedance. In other words, the drop at the receiving end is proportional to the Zt/Zc ratio. As before, this is a best case calculation; other factors may--indeed, surely will--reduce the size of the downward step achieved.

Expressing this concept in the same ratio as in the first paragraph above, the lower the Zc/Zt ratio, the better the fall of the received signal on assertion. What a dilemma! For one objective, we want to raise the ratio; for the other, we want to lower it. Leads one to suspect that maybe the best solution would be to make the two impedances equal, so that Zc/Zt = 1.

As a matter fact, I now believe that is the wrong answer. For the evidence, I offer the appended waveforms. First I'll explain them, then I'll suggest the underlying reasons for what they show, and then I'll make a stab at optimizing the Zc/Zt ratio.

In the tests, the cable impedance Zc was about 87 ohms. The terminator impedance at the driver end, which played only a small part in the results, was 110 ohms. The terminator impedance at the receiver end was 132 or 66 ohms. Thus Zc/Zt was .66 or 1.32. (Absolutely pure coincidence that the numbers came out like that--I had no idea they were going to. If I had used one of the other cable sets, they wouldn't have. These waveforms were taken on the same test setup used in my various S/E Cable Test Reports, particularly Report No. 4, of this same date.) In all terminators, the open-circuit voltages Voc were very close to 2.85 volts--as in fact they always are as long as the termpower voltage stays up around 4.7 volts or a Boulay terminator is used.

In the waveforms, the top two pictures are for the driver end, the bottom two for the receiver end. The left hand two pictures are for Zc/Zt = .66, the right two for Zc/Zt = 1.32.

All of our tests point to the two key vulnerabilities as lying in the release rise at the driver end and the assertion fall at the receiver end. Tests do show waveforms a little way in from the bus ends sometimes to be worse, but of a similar nature. We find the other two "corners" are always less important. The assertion fall at the driver end (top pictures) is controlled by the silicon and is pretty much independent of everything else. The release rise at the receiver end (bottom pictures) is always better than at the driver end because of the boost the signal gets from the receiver end terminator when it gets there. All of this is very well confirmed by the waveforms.

In the top left picture, the low Zc/Zt produces a low release rise. After time for a signal round trip on the bus--about 200 ns--the signal moves on up to Voc. In the top right picture, the high Zc/Zt produces a nicely high release rise, which then works its way back to Voc. In the lower left picture, the low Zc/Zt really rams down the received assertion fall, dipping below zero. In the lower right picture, the high Zc/Zt starves the receiver terminator of enough current to get a good assertion fall. In both pictures, the signal works back to the .4 v or so that the driver is holding after the 200 ns round trip time.

OK, but things aren't completely symmetrical. In the good release rise case (top right), the signal whangs right up there to 3.4 volts, the max achieved. But in the good assertion case (lower left), the falling signal hangs at about .3 volts before falling on below zero. In all tests ever reported, the falling s al at the receiver end hangs on the way down. Why is this?

From analyzing all my results and making some additional tests, and after getting my conclusion blessed by St. Bob of Sun, I offer that it is mainly caused by concentrations of lumped capacity along the bus, along with any signal degredation which may be caused by high-loss cable, if present. Our tests have minimal stubs, since we daisy-chain right to the edges of boards which have signal transceivers mounted right at those board edges. By pulling all the connectors except for the end disk, I established that these minimal stubs are of little significance in our system. (St. Bob is working on an edict that no stub under the Sun should present more than 25 pf lumped capacity to the bus--which only sanctions the transceiver input capacitance plus about 2 inches of trace. No stub cable length permitted.)

Apparently the villains in our test system are the seven bulkhead connectors, although I don't have actual experimental evidence of this. Anyway, as the assertion signal propagates along the bus, it is being nibbled away--by lumped capacity, by line losses, and perhaps by leakage pull-up currents in each transceiver. And there is one other point: the .4 v driver voltage offers only half as much margin below .8 v as the 2.85 terminator open-circuit voltage offers above 2.0 v.

Why don't the lumped capacities hang the signal rise as it propagates along? In Report No. 2 (X3T9.2/90-124, 16 Aug 1990), it shows that they may. But the leakage pull-up currents work against the capacitive drag, and the receiver terminator is working with, not against, the signal as it is finally received.

The overall conclusion: if the Zc/Zt ratio were actually 1, something I have never achieved, the distress of the received asserted signal would be far worse than the distress we have been overcoming in the released signal, and would in fact be unacceptable.

So what should the Zc/Zt ratio be? As I study all my results, particularly those in Report 4, I am forced to the gospel which St. Bob of Madison has been preaching all along. The low 80's impedances of the Madison and Astro cables actually appear better balanced in margin than the somewhat higher impedance cables from Montrose and 3M. Remember that these test results are with 110 ohm terminators at each end. It just happens that 82.5/110 gives a nice square Zc/Zt ratio of 0.75. But it's not a sharp point; anything from 0.7 to 0.8 surely would perform almost as well.

I'm sure this conclusion, if valid, will not upset Montrose and 3M at all. They busted their behinds to get their impedances up as high as they have them. If someone wants a little lower impedance, I imagine they already have the cables in hand.

I don't hold that this ratio holds over all impedance ranges. If the cable and terminator impedances were doubled but the lumped pf along the bus unchanged, the effect of the capacity would be more pronounced and might point to a different ratio. But in practical ranges of SCSI shielded cables, I'm opting for the .75 ratio. And where does that leave one if he sticks to the 132 ohm termin or, as is reported to be the fashion in southeastern Minnesota? Well, you would need to find some cables with 99 ohm S/E impedance. Or, obviously, settle for a lower Zc/Zt ratio. As a matter of fact, the .66 ratio waveforms, made with the Montrose 87 ohm cable and the 132 ohm terminator, are conceivably workable, though clearly not ideally balanced. The rise-on-release waveforms, already deficient, would be a little worse with 132 ohm terminators at both ends of the bus.

One final note: When we get S/E protocol chips with active deassertion in the transceivers, these conclusions may be a little skewed. No longer relying on the current supplied by the terminators to pull the released line up, those who wish to may be able to accept a higher terminator impedance--a lower Zc/Zt ratio. But systems with the .75 ratio and good quality cables--and with proper isolation of -REQ and -ACK--will continue to work just fine, with or without active deassertion.

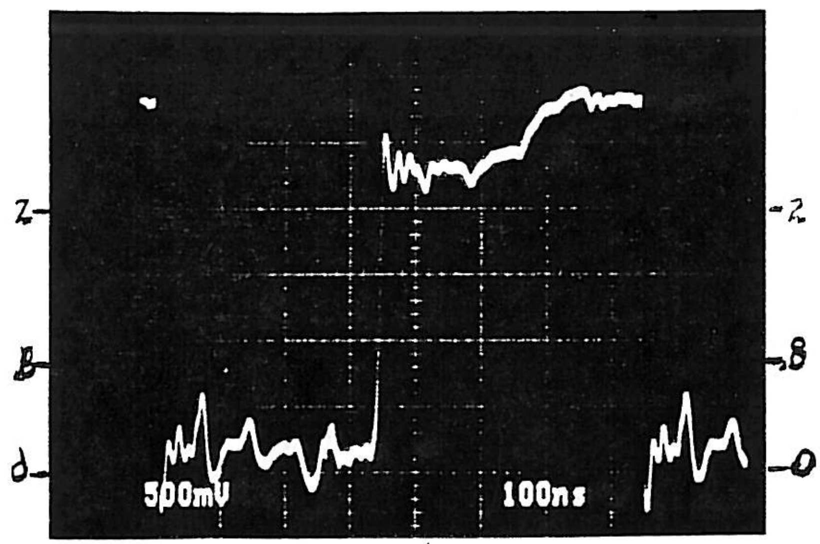

DRIVER END — Z_C/Z_T = .66

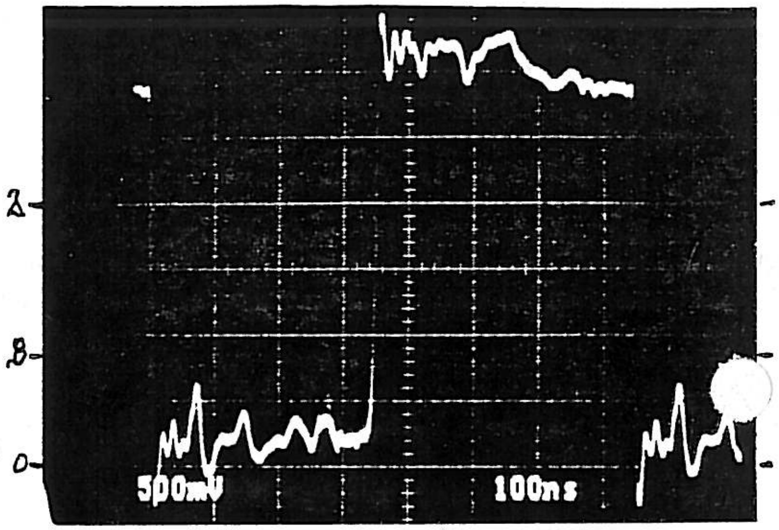

DRIVER END — Z_C/Z_T = 1.32

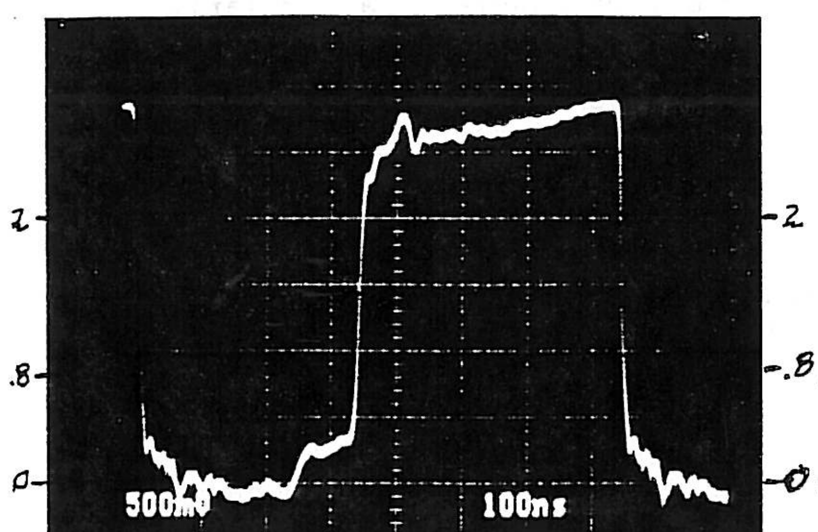

RECEIVER END — Z_C/Z_T = .66

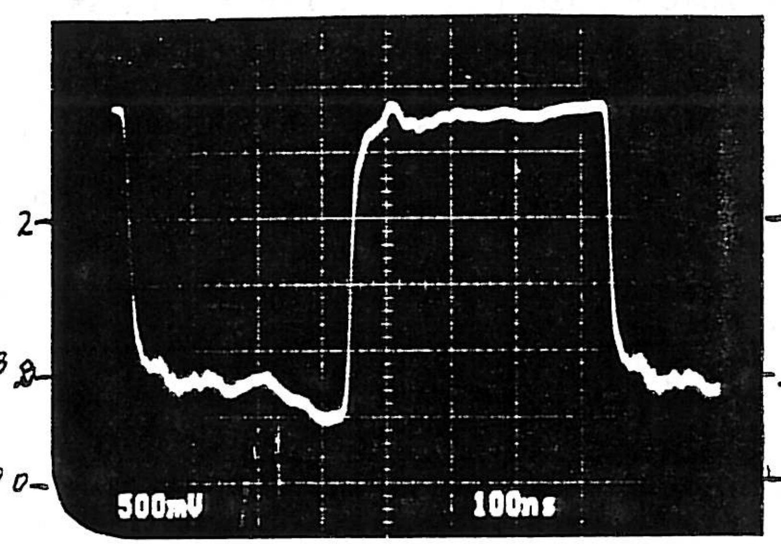

RECEIVER END — Z_C/Z_T = 1.32