HEWLETT PACKARD
Roseville Networks Division
8000 Foothills Blvd., Roseville, CA 95678

X3T9.2/91-002R0

Date: 13 Jan 91
From: Kurt Chan, Hewlett-Packard
To: X3T9.2 Membership
Subject: Analysis of long single-ended buses

In X3T9.2/90-185, Bill Spence (Texas Instruments) makes some observations regarding 20+ meter single-ended bus configurations which warrant further analysis. The following paper provides an explanation for the empirical data gathered by Bill.

First, an analysis of lossless lines will be performed, followed by an analysis of lines with varying degrees of attenuation, ending with a summary of how impedance cable length, and attenuation affect voltage margins.

ANALYSIS OF LOSSLESS LINES:

The two most commonly used analytical tools used in the study of terminated transmission line systems are the SPICE simulation language and the Bergeron diagram. If we analyze the behavior of lossless SCSI cables using these tools, they will give us a starting point for predicting how a system with non-ideal components should behave.

Figure 1 depicts the Bergeron diagram for the asserting edge of a SCSI signal on a 70 ohm cable with a worst-case source (.048V@.048A):

The Bergeron diagram provides an intuitive tool for analyzing transmission line systems with both linear and non-linear termination elements, and is based upon the knowledge that a voltage wave at any point on a transmission line is governed by the following equation:

$$V = I * Z_0$$

Although a mathematical model can be constructed from the principles of the diagram, the accuracy of the solution is generally proportional to the accuracy of one's drawing tools. The rules for construction are as follows:

1. The Y axis represents the voltage on the transmission line. The X axis represents the current. In order to make a graphical solution easy, the X axis is normalized to the characteristic impedance of the cable under analysis, so that all forward and backward waves have a slope of +1 and -1, respectively.

2. Construct the far termination load line using the open-circuit voltage as the Y-intercept. The slope of the line is the termination resistance normalized to the characteristic impedance (Rterm/Z0 = 110/70).

3. Construct the source load line in the asserted state (low). The Y-intercept is the on-state Thevenin voltage:

$$2.85 * \frac{R_{on}}{R_{on} + R_{term}} = 0.2375V$$

   and the slope is the on-state Thevenin resistance normalized to Z0:

$$Slope = -\frac{R_{on} \| R_{term}}{Z_0} = 0.13$$

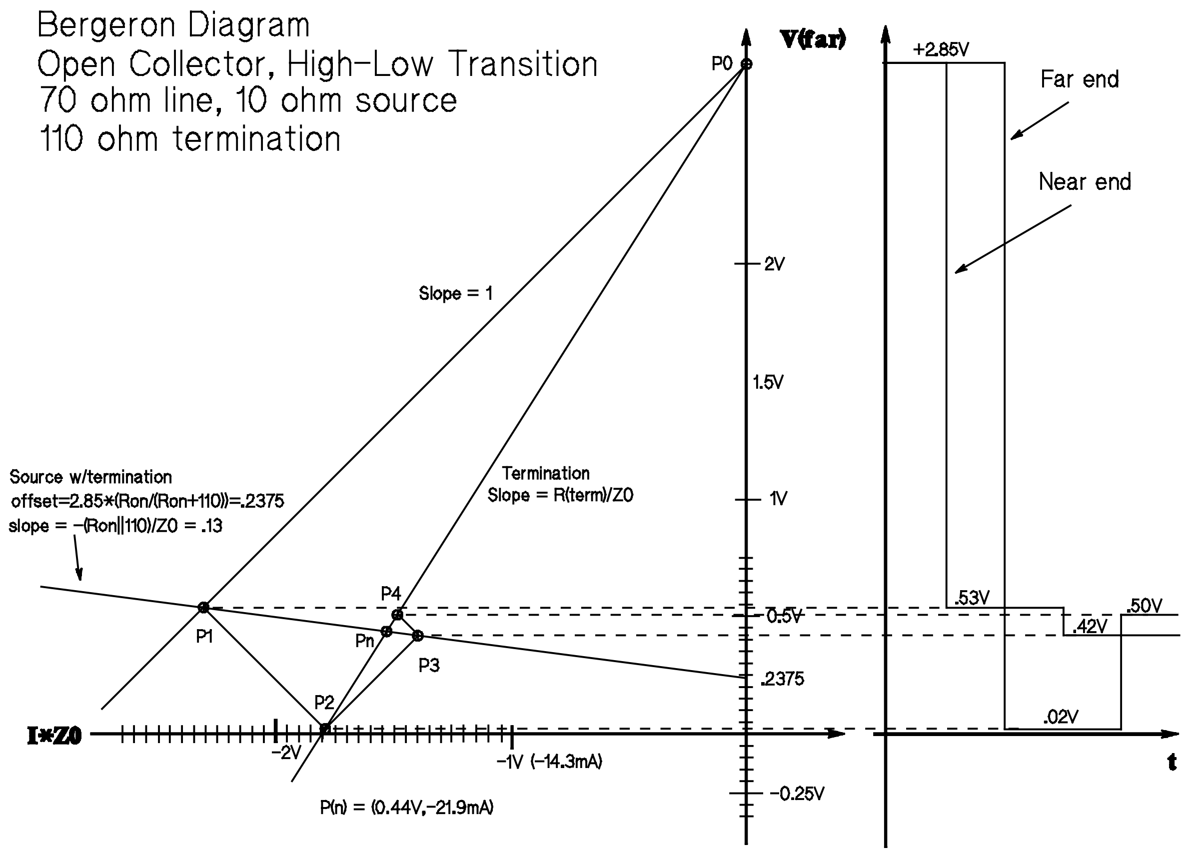

Figure 1

4. Project a line beginning from the open-circuit voltage (P0) with a slope of +1. This is the initial forward wave.

5. At the intersection with the far termination load line (P1), project a line with slope -1. This represents the first reflection back towards the source.

6. Repeating this procedure creates the "Bergeron snail" which converges on the quiescent point V(ol), which is strictly a function of the source and termination DC load lines.

Below is the same diagram for a 90 ohm line instead of 70:

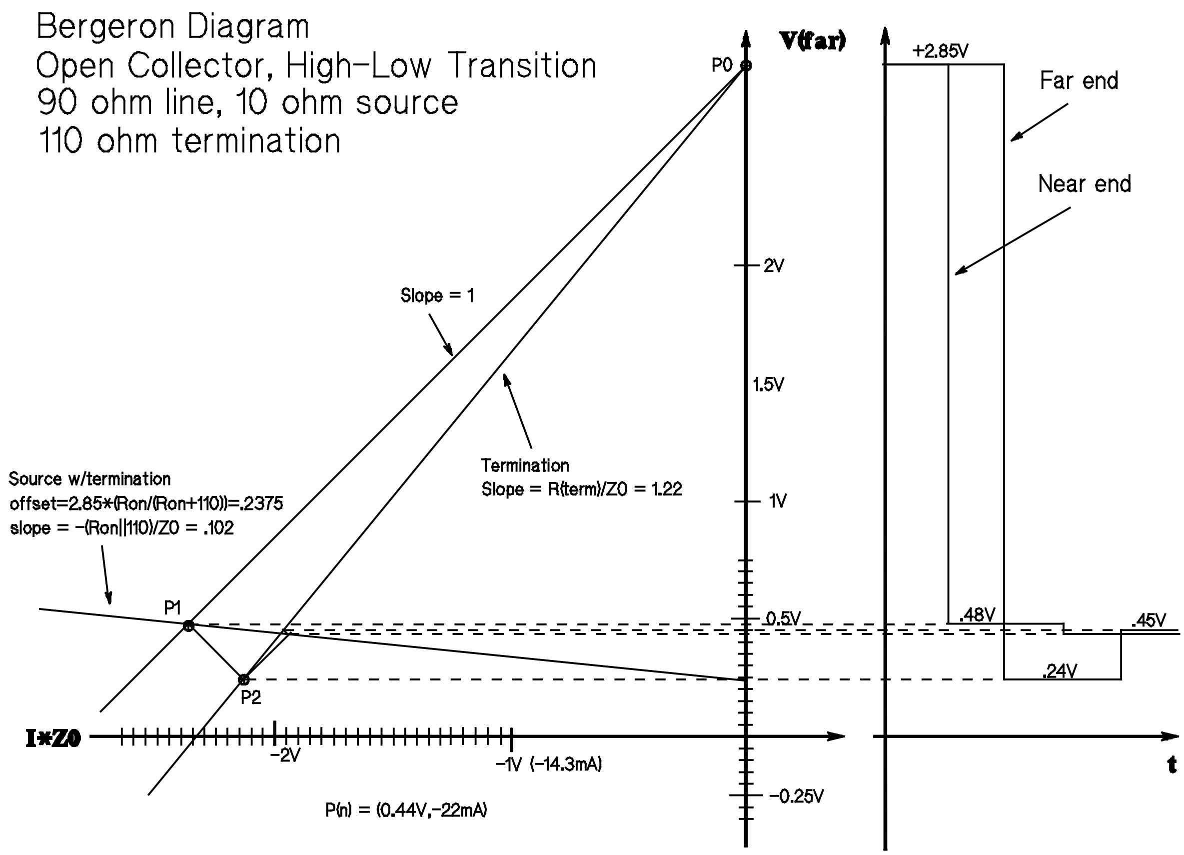

Figure 2

The SPICE plots shown on the next page confirm the results of the Bergeron diagrams, with the far end pulling down to about .25V for 90 ohms and nearly 0V for 70 ohms. Also note that both tools agree that for the NEAR end, the higher-impedance cable pulls down better than the lower-impedance cable. This initial voltage is simply the result of a voltage divider effect between Z0 in parallel with the near pullup, and the pulldown resistance:

$$V_1 = V_{term} * \frac{R_{on}}{R_{on} + (R_{term} || Z_0)}$$

which comes out to 0.54V for 70 ohms, and 0.48V for 90 ohms.

Also note that regarding the initial high-low assertion step on lossless lines from 70-90 ohms:

- At the NEAR end, the step is of LOWER magnitude for LOWER impedance cables.
- At the FAR end, the step is of LOWER magnitude for HIGHER impedance cables.

Very low impedance cables will result in undershoot at the far end, while attenuating BOTH the initial asserting and deasserting steps at the near end. Higher impedance cables will tend not to undershoot at the far end.

On lossless lines, regardless of the magnitude of the initial steps, the waveforms always eventually converge on the same quiescent point. As the cable impedance approaches the termination resistance, the number of converging steps approaches one. This can be visualized graphically on the Bergeron diagram: as the slope of the termination load line approaches unity (Rterm=Z0), the number of far end reflections goes to zero since the quiescent point is equal to P1.

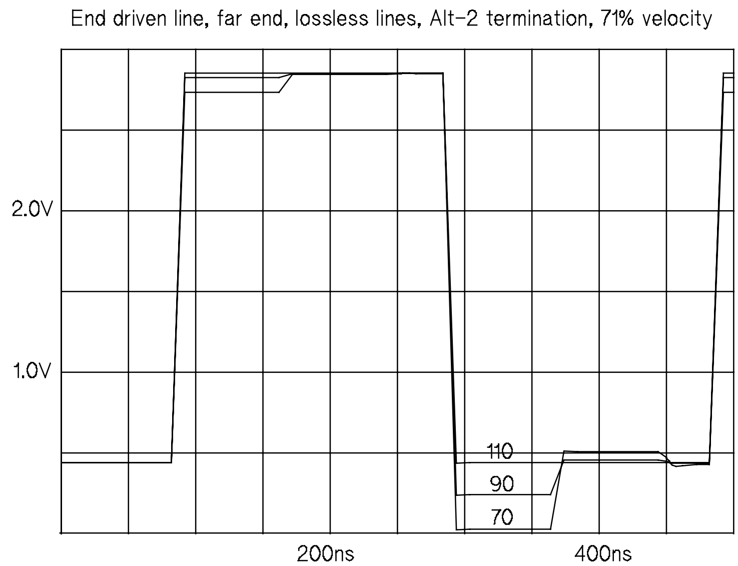

Figure 3

A "perfect" zero-ohm source represented by a horizontal load line on the Bergeron diagram will pull the near end voltage to zero. On lossless cables which are mismatched with R(term)>Z0 this will always result in some undershoot; the amount depending on the degree of mismatch between R(term) and Z0.

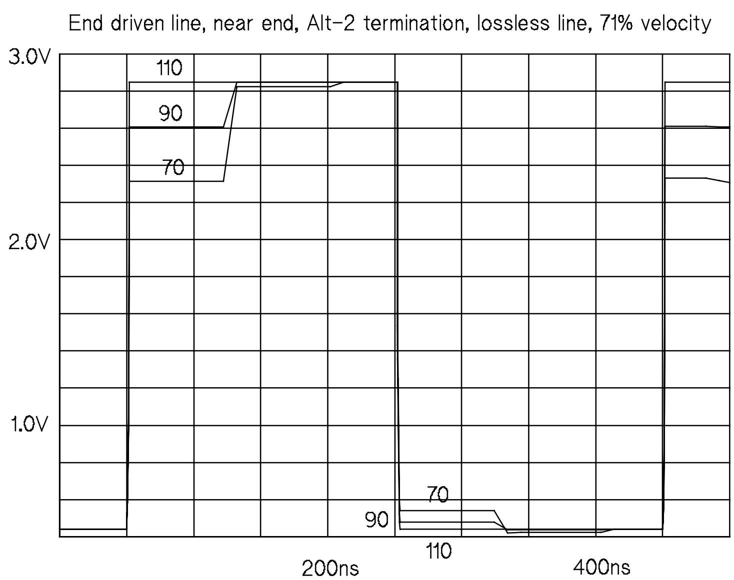

Figure 4

ANALYSIS OF LOSSY LINES:

When attenuation is considered, the picture changes dramatically. The concept of attenuation is simple: if a voltage is applied to the input of a device, the attenuation of the device is proportional to the voltage level at its output. This relationship is commonly expressed in decibels:

$$Attenuation[dB] = \log_{10}\left(\frac{V_2}{V_1}\right)$$

The SPICE program models attenuation for a transmission line as follows:

$$A = Attenuation\ \ Factor = ATTN[dB/cm] * NVEL * c * t(d)$$

where:

ATTN is the cable attenuation spec in Db/cm
NVEL is the velocity of light through the cable
c is the velocity of light through a vacuum (2.99793e10 cm/s)
t(d) is the cable delay (indirectly specifies the cable length)

The attenuation factor describes the logarithmic relationship between the entry and exit voltages on a cable of known velocity and length:

$$A = 20\ \log\left(\frac{V_{entry}}{V_{exit}}\right)$$

$$\frac{A}{20} = \log\left(\frac{V_{entry}}{V_{exit}}\right)$$

$$10^{\frac{A}{20}} = \frac{V_{entry}}{V_{exit}}$$

Physically, attenuation is caused by both conductor losses (which are modeled as series resistance) and dielectric losses (which are modeled as shunt resistance). An examination of 11 different SCSI cable data sheets revealed that real-life attenuation measurements vary with the frequency of the applied voltage, and whether the measurements were single-ended or differential. However, the following analysis is independent of these measurement differences. The average SCSI cable had a single-ended attenuation of about 6 dB/100’ (measured somewhere between 5 and 50MHz), or about twice the recommended Fast SCSI specification.

A frequency-independent method such as that specified by Jim Fiala of 3M (square wave attenuation over cables of known length terminated in their characteristic impedance) is perhaps the best indicator of actual cable performance in a SCSI system.

Figures 5 and 6 show the relationship between cable attenuation and the entry voltage needed to produce an exit voltage of 2.0V when applied to 6m and 25m cables terminated in their characteristic impedance. The percentages represent the % attenuation for the specified length of cable.

Figure 5

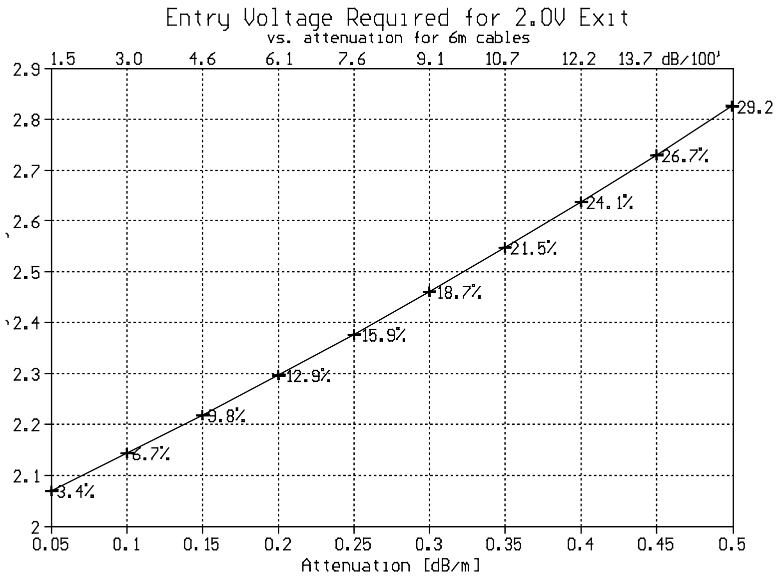

Figure 6

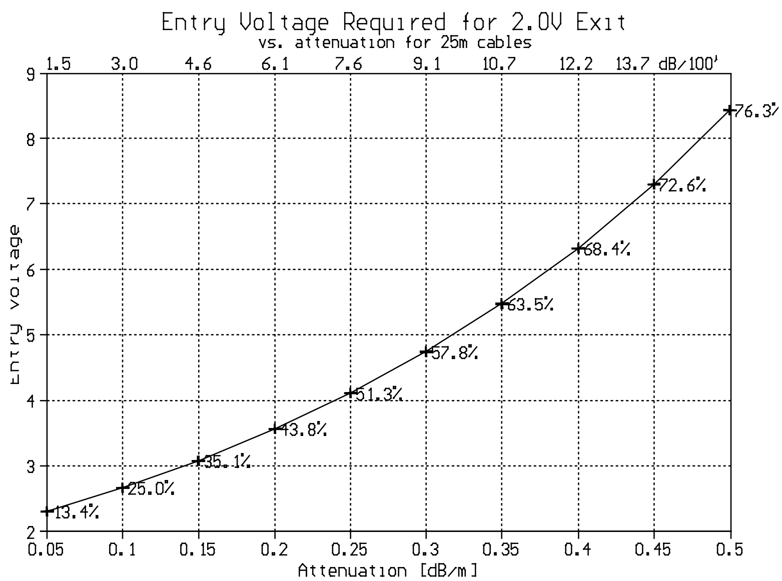

Note that a cable that attenuates a signal by only 19% in 6m lengths will attenuate the same signal by 58% in 25m lengths.

The SPICE plots below demonstrate the attenuation predicted by the previous graphs. At the far end, both the asserting and deasserting levels are attenuated, so that both the 2.0V and 0.8V thresholds are jeopardized.

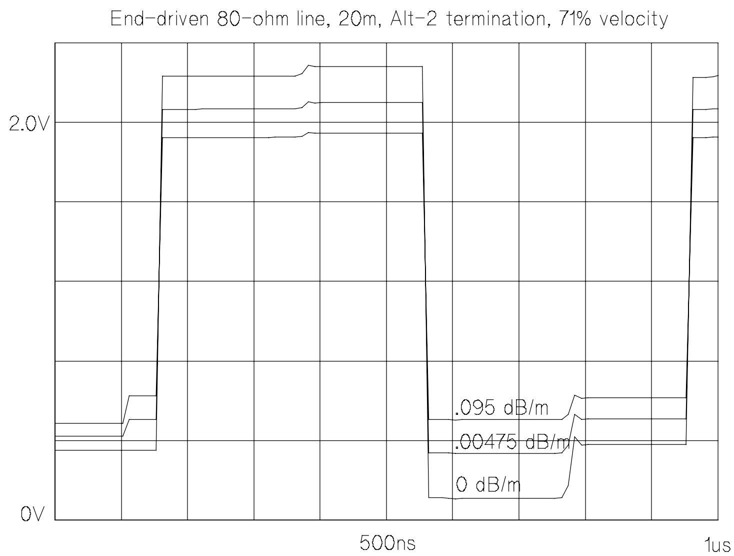

Figure 7

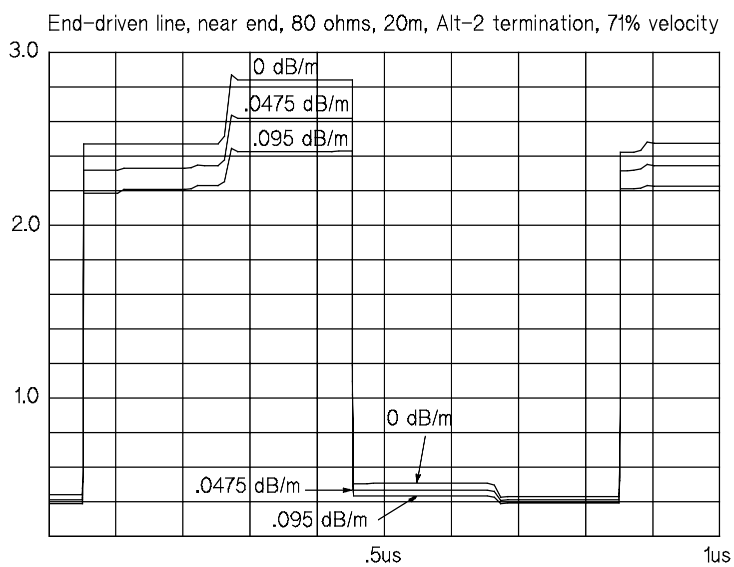

Figure 8

PUTTING IT ALL TOGETHER

We've seen how impedance affects near end deassertion and far end assertion on lossless cables, and how attenuation affects cables of identical impedance. What happens when we observe combinations of impedance and attenuation together? Our analytical models (and Bill Spence) indicate that far end assertion and near end deassertion are the critical waveforms. The SPICE waveform below shows 80 and 100 ohm lines of .095 Db/m at a point 20 meters from the source. Note that the 100 ohm line is dangerously close to the 0.8V low-level threshold:

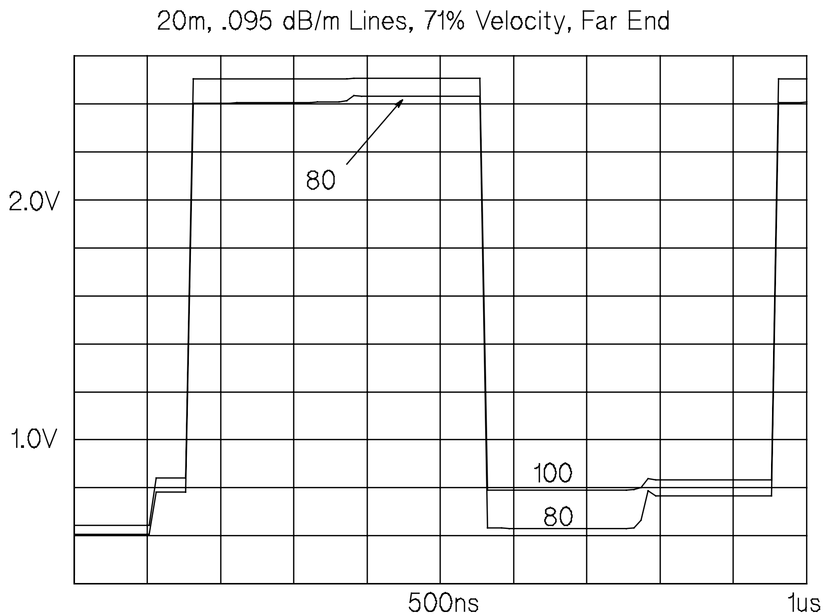

Figure 9

The numbers graphed below taken from SPICE simulations give us a means of weighing tradeoffs between high and low impedance cables. Figure 10 shows how much low level margin we have for various combinations of impedance and attenuation. Note that the **higher** impedance cables provide the least amount of margin on the falling edge. Figure 11 shows the amplitude of the first rising step at the near end of the same cables. Note that the **lower** impedance cables provide the least amount of margin.

Figure 10

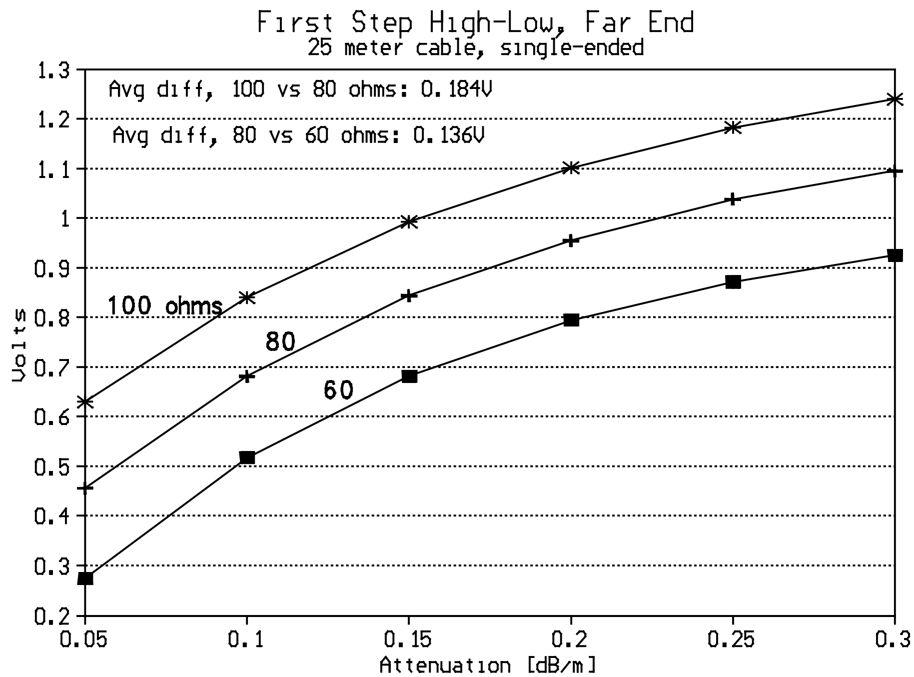

Figure 11

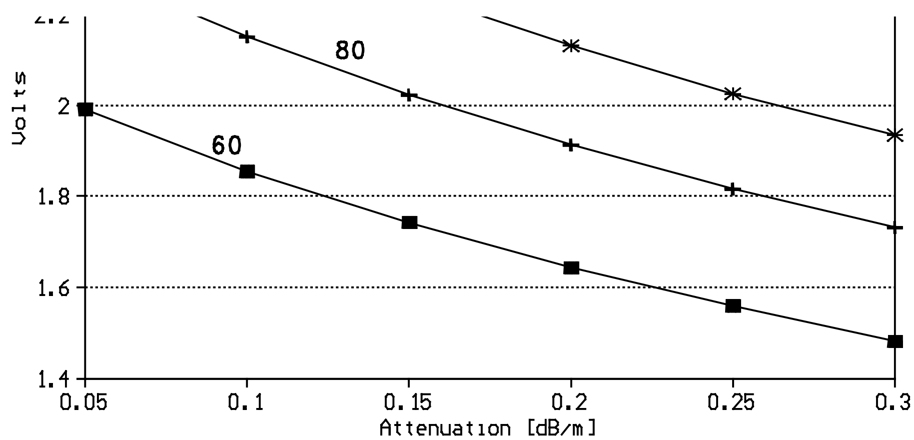

From the preceding graphs, it can be seen that cables with attenuation of greater than 0.2-0.3 Db/m are theoretically unacceptable for carrying single-ended signals on 25m cable. At 60 ohms there is zero far end low-level noise margin for a .2dB/m cable, and at 100 ohms there is zero near end high-level noise margin for a .27dB/m cable.

Interestingly enough, Bill Spence's conclusion that cables in the 78-82 ohm range are best balanced for single-ended operation at 20+ meter lengths is confirmed by these analytical models. However, this is due more to the effects of cable attenuation at the specific cable lengths chosen rather than specific impedance ratios. Bill's conclusions do not apply to short cable lengths where attenuation does not have as great an effect on signal quality.

CONCLUSIONS

The above findings only reemphasize that some additional restrictions and improvements to generic SCSI single-ended systems need to be made in order to support longer cabling. To summarize my key observations and suggestions:

1. On long cables, variations in cable attenuation affect signal quality just as much as variations in cable impedance.

2. Cables which are suitable for 6m operation are not necessarily suitable for 25m operation.

3. On longer cables, it is especially important to achieve V(il) and V(ih) in a single cable delay both for performance and signal integrity reasons.

4. Lower cable attenuation is universally better, if SCSI guidelines for impedance and bus loading are followed.

    - Unless the cable technology is prohibitively exotic, there is no reason why the .095 dB/m attenuation requirement for Fast SCSI should not be applied to **all** single-ended cables to eliminate confusion and reduce the number of incompatible cables in the future.

    - 30AWG cables may only be suitable for shorter cable runs - more testing of these cables is needed at longer lengths.

5. Different cable manufacturers specify attenuation differently.

    - A uniform measurement method should be specified using an environment as close as possible to that in which the cable will be used. For starters, I suggest a single-ended signal source driving a cable segment of known length and terminated in its characteristic single-ended impedance (see Jim Fiala recommendation from Sep '90 mailings). Attentuation can be specified as:

$$\frac{20\ \log\left(\frac{V_{entry}}{V_{exit}}\right)}{cable\ length} \quad [dB/m]$$

6. For 25 meter unloaded systems, cable impedances near 80 ohms provide the best balance of high and low noise margin. However,

   - most cable loading and discontinuities will cause sections of **reduced** impedance and result in negative-going reflections on the cable, putting a premium on high-level noise margin and thus higher impedance cable.

   - as total system attenuation decreases (due to shorter lengths or better cables or both), the impedance which results in balanced noise margins **increases**, which is why for 6m systems high-impedance cables are still preferred.

7. Single-ended configurations of more than 8 devices and/or more than 6 meters still should undergo further validation, particularly for transfer rates beyond 5MHz.

8. Low source-resistance drivers can be used to improve pulldown on long cables.

   - For example, if 64mA drivers are used instead of 48mA drivers (i.e., model the n-channel pulldown as 7.8 ohms instead of 10 ohms), the improvement to far end pulldown at 25 meters is equivalent to about 10 ohms of decrease in cable impedance, or about .03dB/m decrease in cable attenuation. Note these need not be higher **power** drivers, only lower voltage.