



SCSI Signal Modeling - The Time Has Come White Paper

Written for the SCSI Trade Association
By: Bill Ham, Compaq Computer Corp.

1. Introduction

SCSI, the Small Computer System Interface, has been a mainstay for connecting peripherals to computers and controllers for over fifteen years – quite a statement for any technology these days. SCSI (parallel SCSI) shows no signs of fading from the scene any time soon. The SCSI Trade Association (STA) is developing a roadmap for SCSI technology going out to 2012, showing confidence that this basic industry I/O architecture is far from obsolescence. The resiliency of SCSI is due to four main features:

- A brilliant architecture that allows any SCSI device to communicate in-band with any other device as long as the cable plant is continuous and some rudimentary bus termination that matches the transmission type of the devices is in place. This allows any SCSI device to “negotiate” a compatible set of operating parameters with any other SCSI device without requiring that the interconnect be capable of supporting the highest performance.
- A staunch determination by the STA marketing and T10 technical bodies that define and develop SCSI to maintain backward compatibility with previous generations. This makes the transition to the next generation low risk and minimizes the likelihood of obsolete inventories.
- A realization that most users demand interoperability and interchangeability between different suppliers and that open standards are the only way to achieve that goal.
- Always riding the crest of silicon technology evolution and mining silicon’s ability to put an ever-increasing quantity and variety of functions into a single chip.

Other technologies, with significantly more limited capabilities than SCSI, that have also shown staying power are ATA/IDE and 10/100 Ethernet. These technologies also have attempted to adopt the planks listed above to varying degrees.

Why write about SCSI Signal Modeling (SSM)? The main reason is that the third plank may be at risk without SSM in future generations.

2. SCSI Signals

SCSI signal specifications have always been written with the requirement that ordinary laboratory instruments can easily measure the properties of the signals and that these measurements are accurate and reproducible in different measurement environments. There has also been an assumption that suppliers developing components such as cable

media and connectors know what parameters to optimize for SCSI applications. Means for controlling the manufacturing processes that produce these parameters need to relate directly to the properties that are important to the SCSI signals.

Other assumptions in SCSI signal specifications are:

- SCSI signals are not sensitive to relatively small changes in bus length and loading.
- Components such as terminators and device input properties are not significant functions of frequency.
- Connectors and printed circuit elements such as vias are lumped, very low resistance elements that are benign to the signals.
- Short sections of cable or printed circuit (only a few inches) are negligible.
- There is no effect, caused by the interconnect in the device itself, on the device connector where the signals are specified.

All of these assumptions may legitimately be called into question unless some sort of model with experimental verification can prove that they are true. One may take each assumption and examine it in more detail to arrive at the conclusion.

Therefore the job of SSM is to create the methodology for producing simulated SCSI signals that reflects a more complete set of considerations of the behavior of the components that comprise the SCSI bus. These simulations also need to consider how to integrate the models for the components into a complete model for the entire bus.

3. SCSI Signal Modeling Methodology

The basic methodology of approaching the SSM challenge is to define points in the bus where knowing the properties of SCSI signals is desirable. Figure 1 shows an example of the simplest SCSI bus configuration, in which two SCSI devices (represented as transceiver boards) are connected by a simple cable assembly.

This configuration is deceptively simple however. For example, note that in the neighborhood of the connector that the transceiver board is shown as ending at the point where the connector contacts the printed circuit board. Another example is the definition of a transition region in the cable assembly between the undisturbed media and the connector itself, where the wires fan out to make contact with the connector. Also note that the point where the connectors actually mate is not accessible when the connectors are mated. These subtle considerations are key in determining who is responsible for the performance of the various components in the bus. The transceiver board may also contain connectors that are not mated at all when the board is being used. These unmated connectors may affect the signals.

The supplier of the transceiver board needs to know not just the effect of having half a mated connector pair on its product, but also the effect of the other half of the mated pair when both halves are not supplied by the same manufacturer (as certainly must be the case in an open interoperable system). If models of the mated connector are available, then the signal at the point where the wires of the cable assembly attach to the

connector may be calculated. By looking at the signal at this point and at the point on the board where the connector attaches to the board, the total effect of the connector may be seen. This does not solve the general problem of different suppliers for each half of the connector. This problem is only solvable through cooperation by the different connector suppliers, but it clearly shows whether the connector has a significant effect on the signals or not.

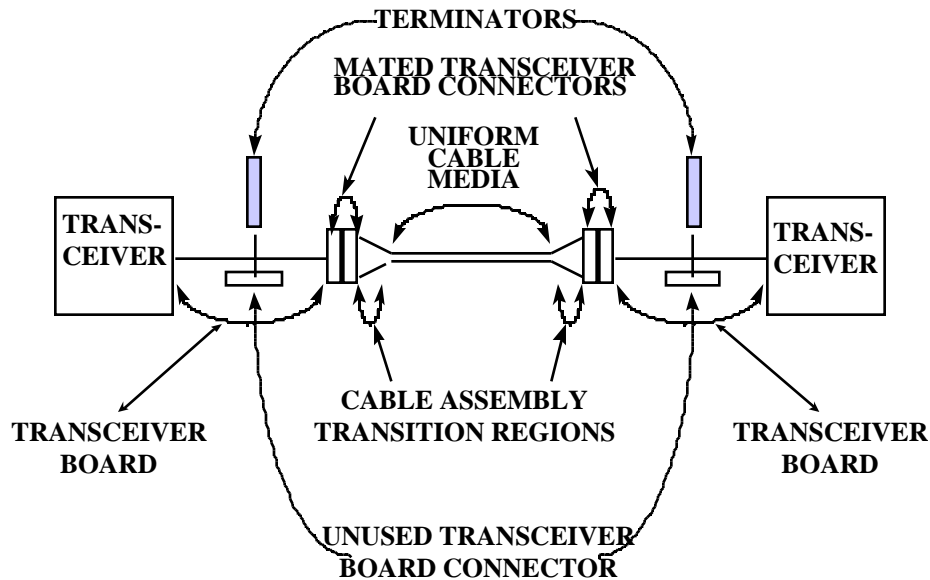


Figure 1 - The simplest SCSI configuration

The supplier of the cable assembly has a similar challenge. The supplier must consider the effects of the mated connector on the performance of the cable assembly. The analysis of the transition region has presently been considered only for a simple case of twisted-flat cable. The transition region is created by the cable assembly process. Defining an effective model presents a challenge for manufacturers of cable assemblies.

Figure 1 does not show all the points of interest for components suppliers. For example, in the transceiver area, details of the chip package and how it is attached to the board are of primary concern. Other details of the routing of traces and integrating the models for the terminator and for the unused connector into a single model for the transceiver board, are not shown but are clearly of first order interest.

4. Kinds of Models

There are basically two kinds of models; circuit and behavioral.

In a circuit model, a circuit diagram of the SCSI bus component being modeled is drawn and values are assigned to the parameters of the circuit elements (for example capacitors and inductors). In a behavioral model the properties of the signals at certain

points of the component are specified without revealing a detailed circuit form. There are complex considerations for which kind of model is better.

Circuit models are generally considered to be superior for describing the complete behavior of the component. On the other hand, circuit models are also generally considered to contain significant information about the details of the design of the component and are closely held as proprietary by some suppliers.

Behavioral models are generally inferior for describing the complete behavior of the component. On the other hand, most people consider that little is revealed about how the component is designed from a behavioral model. There are some very conservative suppliers that consider both circuit and behavioral models to contain unacceptably high levels of proprietary content.

Behavioral models are presently unable to handle intelligent signals of the type that result when transmitter Inter-symbol Interference (ISI) jitter compensation is used. In this case, the launched signals have different properties, depending on the data pattern being transmitted. Such a scheme is described in the latest draft of the SPI-4 standard for use with Ultra320 SCSI.

Another feature of any model is where, in the specified range of behavior, the model applies. Models that do not specify whether they are for nominal, best case or worst case components, or somewhere in between, are not useful. It is especially dangerous not to specify where, in the distribution the model applies, because it gives the model user license to pose any assumption that makes the simulation come out as the user wants it to. SSM does not require any particular method for determining what constitutes best or worst case, but the value of a nominal model is questionable since designs need to be done for worst case components most of the time.

Then there is the issue of coupling a circuit model for one component with a behavioral model for another component. This interface is one of the challenges being addressed by the SSM working group of the T10 committee.

5. Component Models

Every component used to make a SCSI bus is represented in the SSM group. Some members of the group have offered their models for use through a mechanism described on the SCSI standards web site: www.t10.org. Others are in the process of generating models that are accurate enough before making them publicly available.

When suppliers do not supply accurate models, they force users to take an alternate approach to create the model for the component. One of these methods is called "parameter extraction." In this methodology, measurements are made on physical parts to synthesize a model. This could be called reverse engineering, but more correctly the goal is not to do a better (or equivalent) design, but rather to understand how this component will affect the signals on the SCSI bus.

Parameter extraction can also be an excellent method for determining circuit component values from first principles based on materials properties and dimensions. However, in order to use this scheme one needs to know the tolerances on the material properties and dimensions.

In every case where the user is forced to create a model of a component, real physical parts will be used as the basis. Since it will not be possible to determine which sample in the distribution of parts may be used in service, there will always be a question about the margin and accuracy of the extracted model.

There is another significantly negative outcome from user-extracted parameter models, namely, that since this is the only information available for the component, it is likely to be used by others who are attempting models of SCSI busses containing the component. Everyone involved is hurt when suppliers do not make accurate models available. Any vendor supplying components into the SCSI market would not want the product misrepresented and should do everything practicable to ensure that accurate models are made available.

6. Models for Enhancing Measurement Accuracy

Measurements must be made where physical access is possible. The Heisenberg principal is still alive and well in that any measurement disturbs what is being measured. Enter modeling. With a complete model of the SCSI bus available, one may infer the signals at points in the bus other than where the probes are placed.

When the model for the instrumentation interface is included as part of the SCSI bus, then Heisenberg is basically defeated and one can infer what the signal would be if the instrumentation is not present. If the observed signal does not match that simulated from the model then there is no point in attempting to infer what the signals look like at other points – it does not even match where it is being observed. In this case the model must be fixed.

7. Models for Determining Sensitivity to Parameters

A corollary to enhanced accuracy is the ability to determine sensitivity to the details of measurement. For example, does it really matter which side of the connector is probed? Does it matter exactly where on the trace the probe is placed? Move the probe point in the model and see what happens. What is the effect of continuously changing the bus length or loading? How does one determine the nearness of a black hole of destructive resonance?

Clearly, the value of signal modeling to optimize the design parameters of components is a traditional use. Questions which modeling must answer include, what effect does the frequency response of the terminator have, should capacitance of the device be reduced even more, what is the real maximum stub length to be used and is an expander needed at this point in the bus? With an accurate model the world lies at the feet of the user. Without one, the distance to the brink of disaster is just a guess.

8. Models as the Key to True Interoperability

When accurate component models are used as the basis for a specification, it becomes possible to determine ahead of time what is needed to achieve interoperability. This scenario, however, has some uncertainty. How can a designer design a transceiver if there is no way to know how the passive parts of the interconnect behave in detail? And vice-versa? Resolving the guesswork is a work in progress.

The traditional methods of specifying signals at the device connector, while still the only way presently known to produce interchangeable parts, is becoming more complex. One must cast the measurement of the signals in the same way as they are intended to be used. For example, if pre-compensation is to be used in a transmitter, then only when a model of the interconnect is used as a load, will a true representation of the effect of the pre-compensation on the received signal be known.

If the receiver processes the signal in some way, for example de-skewing, such that the internal receiver sees a much better signal than that existing at the device connector, some scheme for emulating (modeling) the receiver's de-skewing properties is needed to evaluate the observable signal, so the interconnect will have a relevant specification to work with. As the interaction between interconnect, terminators, transmitter and receiver becomes more intricate, the use of modeling holds the key for true interoperability.

###