

SCSI Laying the Ground Work for Performance

A White Paper Prepared by The SCSI Trade Association

December 1997

(This article also appeared in the December 1997 issue of
Computer Technology Review®.)

SCSI Parallel Interface Technology -
a path to the future!

SCSI LAYING THE GROUND WORK FOR PERFORMANCE

SCSI for the present and future

The Small Computer Systems Interconnect (SCSI) that has become the workhorse of the high-end system and workstations. SCSI started from a modest beginning and has continued to improve as the need for performance arises. It started from a basic interface supporting up to 5 megabytes per second, and has improved performance with wide Ultra2 to 80 megabytes per second shipping today.

The SCSI standard has been balancing the system speed requirements, system cost and a versatile interconnect to allow the widest variety of options to be connected to the system. The cost performance interface allows workstations and PC's to connect to scanners, CDRs, CDs, CD changers, tape drives, cartridge drives and hard drives with flow control to optimize system throughput.

Other new interfaces infer high speed and the number of devices that can be connected to a bus but what devices can they connect to? Fibre channel has one supplier of disk drives, SSA has one supplier of disk drives, but neither has any other devices that will connect to their bus. SSA's PH2 (high-speed interface) 40 megabytes per second is half the speed of wide Ultra2. A Fibre channel 100 Megabytes per second or dual fiber system is at a similar cost to a 2 or 4 channel wide Ultra2 PCI card. Dual 80 megabyte per second or quad 80 megabyte per second balances performance and cost. A quad Ultra2 SCSI controller connects up to 60 drives with 320 megabytes per second throughput, faster than the 32 bit PCI bus. With the 18 Gigabyte disk drives shipping this year, it is possible to build over a terabyte of storage per PCI slot with 60 drives. The table below shows a comparison of SCSI speeds and connectivity against other technologies and buses.

Interface	Speed (Narrow/Wide SCSI)
Fast SCSI	10/20 Megabytes/second
Ultra SCSI	20/40 Megabytes/second
Ultra2 SCSI	40/80 Megabytes/second
Ultra3 SCSI	80/160 Megabytes/second
EIDE	33 Megabytes/second
SSA	20 Megabytes/second
SSA-2	40 Megabytes/second
1394 (Firewire)	10 Megabytes/second
1394A	40 Megabytes/second
Fibre Channel	100 Megabytes/second
ISA Bus	<5 Megabytes/second
VME Bus	<40 Megabytes/second
PCI Bus / high clock speed	133/266 Megabytes/second
PCI 64bit Bus	266/532 Megabytes/second

Ultra3 is under development in the standards committee. It will offer better cost and higher performance than Fibre channel. Fibre channel has a road map to higher speeds, but it comes with additional cost, for example fiber transceivers that can not be integrated into the controller. The wide Ultra3 speed of 160 Megabytes per second is faster than several system buses, pushing PCI buses to the limit. Dual and quad wide Ultra3 SCSI will need a higher performance bus to utilize the available throughput. The amazing SCSI bus is offering speeds that use to only be considered fast for a short backplanes with a few devices. The SCSI bus is running over 12 meters of cable and 16 devices. A quad Ultra3 SCSI card can exceed the throughput of most buses, 640 megabytes per second.

From the system architecture point of view, why do some believe it is necessary to support a new driver set and transport system when the industry has been using multiple servers and networks to perform the long distance connections with redundant databases. The network support is built into the system, with performance that ranges covers a full performance range with both wide area and local area interfaces. A high reliability fibre channel system relies on one or two star couplers that cost about the same as a server. Less than 1% of the storage systems need a separate redundant system like fibre channel can offer. Most storage solutions are cost effectively solved with networks and servers with SCSI storage solutions.

SCSI offers queuing that allows full utilization of the bus, other interfaces like IDE or EIDE each transaction must complete before the next one can be started. This often includes the time for the head position and rotational latency. SCSI allows queuing of several commands to one or more drives. The disk reorders the queue to take maximum advantage of the mechanical movements. The SCSI devices usually use larger cache memories, which allows the bus to be used as close as possible to the maximum bandwidth.

SCSI Performance:

Ultra (Fast-20)

Ultra SCSI based on the single ended interface greatly restricts the distances and number of devices that can be attached to a physical bus. The single ended physical issues require very short cable lengths for a multidrop bus. Point to Point (2 devices at the end of the bus) single ended can be used at distances up to 20 meters, but as soon as a third device is added the distance drops 3 meters. 4 to 8 devices on a bus the distance drops to 1.5 meters. The cable lengths were reduced to allow the reflected wave to insure the signals have reached the receiver threshold points. The reflected wave time is determined by the propagation time of the cable, (10 ns per meter round trip time for a standard cable). The advent of expander chips that receive, reclock, and redrive the signals improves the flexibility of the system allowing more devices to be connected to a bus or greater distance.

Ultra2 (Fast-40)

Ultra2 SCSI is based on the Low Voltage Differential (LVD) signals have adequate margins to allow 12 meter bus with 16 devices. Multiple transfers are occurring on the bus at one time. The signals propagate at a rate of 5 ns per meter, i.e. 25 ns transfer time every 5 meters is a bus cycle. This gives additional speed and distance without a major increase in cost. Point to Point connections are possible up to 25 meters.

LVD SCSI – balance

How does LVD balance power, speed and connectivity? LVD has similar margins ratio of driver output to receiver thresholds as HVD, High Voltage Differential. This allows signals to be driven with a single step or what is known as incident wave switching. The major problems with the higher speeds become the frequency roll off and SKEW. SCSI is a base band system, signals switch at speeds from DC to the transfer rate. There is no encoding to limit the frequency range of the signals. Signals that do not switch for several transfer times build to a higher signal level on the bus which is harder to overcome with the first transition. The system has to be designed with adequate margin for the high frequency signal able to overcome the signal levels of a line that has not switched for several transfer times. This is normally referred to as ISI, intersymbol interference - 100 megatransfer rates appear at this time to be the limit for LVD before encoding will be required. It appears at this time that the

next step after Ultra3 will require the data phase to migrate to an encoding scheme like 8B/10B to eliminate the DC or signals that do not transition for several transfer times. This is a standard telecommunications encoding scheme used on trunk lines.

LVD technology and backplanes

The LVD technology allows more devices to be connected to the bus. The backplane design offers interesting challenges for SCSI. The connector spacing normally is closer than the standard allows. The connector and the plated through holes have the dominate effect of the impedance. A lot of the connectors are press fit working with through hole technology. The backplanes normally have 1 or 2 ground planes and 1 or 2 power planes. A plated through hole adds capacitance from each of the planes. The more ground and power planes the more capacitance. It may be necessary have clearance areas around the connectors to reduce capacitance. The plated through hole capacitance with just one power and ground plane with standard clearance holes is 2.5 to 3.5 pF. With 2 power and ground plane a plated through hole will be in the range of 4.5 to 5.5 pF. The connectors adds an additional 3 to 5 pF. A plated through hole and connector space every 1 _ inches without a device connected reduces the impedance of the backplane from around 90 ohms to 75 ohms before the devices are installed. The devices will reduce the impedance to the 50 to 60 ohm range with the tight spacing, this is much lower than the 85 ohm minimum required by LVD SCSI. Backplanes can not just consider the etch impedance, as the connector and plated through hole are the dominate factors because of close device spacing required.

Device and backplane layout will be the key to making LVD SCSI work for Ultra 2 and Ultra 3 speeds. Reducing or eliminating plated through holes helps significantly and surface mount connectors will reduce capacitance. If through hole connectors are used, clearance areas through the power and ground plane will help reduce the capacitance. Longer cables between backplane segments can also reduce the loading effects on the system.

Note: Unitorde offers a full line of the lowest capacitance terminators in the industry that can help reduce bus loading. Single ended 9, 18 and 27 line terminators are available. Multimode (low Voltage differential and single ended) 9 line is available with higher line count devices in development. LVD only 9 line is sampling with 18 and 27 line versions in development.

Future:

Ultra3 (Fast-80)

The Ultra3 work is starting in the ANSI standards committee T10, which will allow the next bus speed increment when required. The REQ and ACK signals will keep the same speed as Ultra2, but both edges will be used to clock the data. The falling edge will clock the first transfer and the rising edge will clock the second transfer. This keeps the maximum signal speed the same as Ultra2 with twice the number of transfers.

Reduce Protocol overhead

The increased speed of the data transfer helps improve bus throughput, but the arbitration and messaging time is taking a larger percentage of the time. Work is underway in the committee to define improvements that will reduce the arbitration and messaging times allowing the high data transfer rates to relate directly to high bus throughput. Command and message transfer at 5 meg Async, versus data that transfers at the highest negotiated speed between the target and the initiator. Commands currently are sent out to one device per negotiation, new commands are being defined that allow commands to be sent out to several targets in one negotiation.

Ultrax (Fast-160 and beyond)

Encoding offers the beginning step to reduce the intersymbol interference, as speeds increase skew becomes the major limiting factor. Self-clocking data would eliminate the skew issues as speed increases. PreComp (driving the first bit harder) or reducing drive after the first bit may be required.

Summary

SCSI continues to evolve offering the performance to match any interface, with low cost and flexibility that allows new and old devices to coexist on the same logical bus. SCSI offers the performance and connectivity that servers require now and will continue to increase performance for the foreseeable future.

SCSI offers the widest variety of devices, with the connectivity to keep it the preferred connection interface for workstations and High end PCs.

Acknowledgements

The STA would like to thank the author of this white paper, Paul Aloisi <aloisi@unitrode.com>, Applications Engineer, Unitrode Corp., 7 Continental Blvd., MS 84, Merrimack, NH, 03054-4334. Please visit the Unitrode home page at <www.unitrode.com>.

Additional Information

The SCSI Trade Association has a wide variety of documents and information on SCSI Parallel Interface Technology, including presentations, articles in periodicals, seminar material and white papers. Please contact Association offices at:

The SCSI Trade Association
404 Balboa Street • San Francisco, CA 94118
Tel: 415-750-8351 • Fax: 415-751-4829
Info@scsita.org • <http://www.scsita.org/>