Advances in
SCSI Parallel Interface Technology

A White Paper Prepared by The SCSI Trade Association

June 1996

SCSI Parallel Interface Technology -
a path to the future!

The Interface Dilemma:
The computer's insatiable appetite for higher I/O performance is now further accentuated by the Internet, network servers, multi-media, video-on-demand and other demanding applications. More users need more data from more devices at higher speeds, while expecting greater levels of reliability, availability and serviceability at lower costs.

The response to the need for higher I/O performance has resulted in choices which are both technical and commercial in nature - and complex. Fundamentally, the choice is between a constantly evolving parallel interface technology that can track its origins back to the sixties and the new serial interface technologies beginning to emerge. At issue is not whether serial will eventually prevail, but when it will prevail. Serial interfaces are not as close at hand as some believe and SCSI Parallel Interface Technology will continue to dominate well into the twenty-first century.

Topics:
To put the emerging serial interfaces in the proper perspective with SCSI Parallel Interface Technology, this white paper will cover both the commercial as well as the technical aspects. First the commercial considerations:

- Infrastructure
- Installations
- Knowledge and Recognition of the Technology
- Peripheral Support Range
- Plug-n-Play Capability
- Cost
- Form Factor

Infrastructure

For over a decade, the SCSI Parallel Industry Infrastructure has been rapidly expanding both vertically and horizontally. Vertically, there are suppliers at every level, including chips, boards, peripherals, systems, software, cables, enclosures, terminators, bridges, connectors, etc. Horizontally, the array of SCSI products is enormous, with few, if any, SCSI products without a second source.

The SCSI Parallel Industry Infrastructure contains well over 100 suppliers. The size and economic health of this Infrastructure has made it possible for enhanced and new products to appear frequently. This includes faster chips, smaller connectors and products such as bridges and extenders which amplify the capability of SCSI Parallel Interface Technology.

Installations

The number of computer installations which incorporate SCSI Parallel Interface Technology numbers in the tens of millions. With the exception of the very low and high ends of the computer spectrum, Parallel SCSI is pervasive.

Knowledge and Recognition

A large established infrastructure and millions of installations indicate clearly that knowledge of Parallel SCSI is widespread in the computer industry. The extensive depth and breadth of this knowledge took over a decade as well as hundreds of millions of dollars to accomplish. The extent of Parallel SCSI knowledge makes it relatively simple and inexpensive to keep up with Parallel SCSI developments.

By virtue of the widespread Parallel SCSI knowledge and the tens of millions of installations, SCSI has become a recognized acronym with those who build things using SCSI and with those who buy them. No other interface has such wide-spread recognition.

Peripheral Support Range

Parallel SCSI undoubtedly supports more types of peripherals than any other interface. This support includes:

- Magnetic rigid fixed disk drives
- Magnetic rigid and flexible removable high capacity disk drives
- Magnetic floppy drives
- Optical CD-ROM drives
- Optical WORM drives
- Optical erasable drives
- Tape drives
- Scanners
- Printers

Plug-n-Play Capability

Standardization efforts covering the electrical and physical characteristics as well as the protocol have been evolving for over a decade. SCSI Parallel Interface Technology has now reached the point that users can purchase and install a SCSI device with relative ease in minutes. The installation and operation of Parallel SCSI devices is simple and standardized.

Cost

The huge volumes of Parallel SCSI products, from chips on up to systems, that have shipped for over ten years and continue to ship, have been instrumental in bringing SCSI Parallel Interface Technology to millions of users at low cost. No other I/O interface provides the performance, expandability and flexibility that Parallel SCSI provides at such low costs.

Form Factor

When Parallel SCSI first appeared in the sixties, form factors were huge in comparison to today's hand-held devices. Size constraints regarding connectors, cables, terminators, etc. were virtually non-existent. As form factors decreased in size, SCSI designers kept pace by providing smaller connectors, cables and terminators, etc. This trend continues. SCSI now offers connectors, cables, terminators, etc., which match the smallest peripheral form factors including 3.5-inch and 2.5-inch.

Introduction To The Technical Aspects

Major innovations in parallel SCSI technology built on a foundation of proven technology, promise to meet the needs of the marketplace both now and for the foreseeable future. This white paper describes the changes and how they will improve price-performance while protecting the customer's investment in hardware, software and training.

What's Available Now?

Peak transfer rates doubled from 10 to 20 MB/sec (40 MB/sec for wide SCSI)

What's in the pipeline (12 to 18 months)?

A four-fold increase in peak bus bandwidth (to 160 MB/s for wide SCSI)
A two-fold increase in packaging and bandwidth density
Increased systems design and configuration flexibility
`Smart silicon' to improve device and systems throughout

What's on the horizon?

Increased device addressability and configuration flexibility through bridge technology
Increase in allowable bus length beyond 25 meters

As described in the following sections, these changes will involve improvements at all levels of the technology including:

The electrical layer
The physical layer (cables and connectors)
The protocol layer
The development of new components for systems integration

In what follows, peak data transfer rates are given in megatransfers per second, where the unit of information transfer is a function of the SCSI data path width. For wide SCSI, data is transferred two bytes at a time, hence the bandwidth in MB/sec is obtained by multiplying the MT/sec rate by two.

The Electrical Layer - Bus Length, Bus Width and Peak I/O Bandwidth

As noted in the introduction: maximum synchronous transfer rate has doubled to 20 MTS within the last 18 months and is slated to increase four-fold to 80 MTS. A bus width of 16 bits will be standard for high-end devices, resulting in a peak data rate of 160 MB/sec.

High-speed operation with up to 16 devices and bus lengths over 12 meters is made possible by a new differential signaling technology described below called Low Voltage Differential (LVD). As noted below, LVD silicon is also capable of operating in single-ended mode for compatibility and is comparable in cost and power consumption to single-ended devices.

What is Low Voltage Differential (LVD)?

LVD is a differential bus technology that combines much of the bus length, noise immunity and performance benefits of conventional differential SCSI with the power consumption and cost of single-ended SCSI interfaces. Power consumption of LVD devices is reduced compared to a conventional differential bus through improvements in receiver design that permit reductions in steady-state current consumption and signaling voltage.

Because of this lower power consumption, LVD drivers can be integrated into the silicon interface chips thus eliminating the signal skew, real estate and cost associated with separate differential components. What's more, by taking advantage of the latest CMOS processes, dual-mode LVD cells can be designed that support either single-ended or differential operation. Selection of operational mode (single-ended or differential) by the device is automatic and is done without the use of jumpers. Because of this compatibility, the cost of SCSI devices with LVD silicon will not differ appreciably from comparable single-ended drives.

The Physical Layer - Cables And Connectors

Packaging and bandwidth density will increase through:

- Very high density (VHDCI) 50- or 68-pin connectors, one-third the size of conventional high density connectors;
- New cable designs with greater flexibility and half the cross section of existing cables;
- High density device connectors that eliminate the need for cables in drive arrays.

By using the new cables and connectors, combined with high-density silicon, it is possible to double the number of SCSI busses on a single adapter card from two to four. In addition, since the connector fits within the PCMCIA form factor, it can be used across all applications from notebooks to mainframes.

Systems Integration Solutions

The SCSI bus repeater is a component that provides the system designer with the means to overcome otherwise insurmountable configuration constraints.

Essentially, a repeater is a two-port device for coupling independent bus segments so that, to all attached devices , the segments appear as a single electrical bus. Because of propagation delays through the device, the total length of all physical segments will be less than the maximum permitted by the standard. Nonetheless, as shown by the examples below, bus repeaters can be used to solve vexing configuration problems. The use of bus repeaters will soon become commonplace as low-cost, single chip devices enter the market.

<u>Overcoming bus length restrictions:</u>

Since the bus segments interconnected by a repeater are electrically independent, the use of repeaters to cascade segments provides the designer with a straightforward way to surmount bus length limitations.

<u>The right bus technology in the right place:</u>

The use of bus repeaters allows the system integrator to resolve conflicting operational requirements. For example, the speed, signal integrity and length of differential busses are often required for host to subsystems interconnects, whereas single ended busses are needed within the storage enclosure for low cost and compatibility with existing devices. In this case, a differential to single-ended converter can be used to transition from one bus type to the other at a cost that is far less than alternative solutions.

Bus Protocol

In addition to the increased data transfer rates describe above, new VLSI components will contribute to performance through protocol engines implemented in silicon. This improvement occurs in two ways. First, device performance is increased by offloading from the microprocessor many of the protocol chores formerly performed by firmware. Second, system performance is increased by insulating the bus from the effects of processor `think time'. Without hardware assists, think time often results in a condition where the device owns the bus but is unable to use it for data transfer. The net result of this idle time is a reduction in the effective carrying capacity of the bus. The latest SCSI interfaces solve this problem by delegating bus interface micromanagement to the hardware.

Futures

SCSI parallel bus futures are being addressed by the SCSI standard committee through the Extended Physical Interface project. The following are the principle areas of investigation.

<u>Increased device count:</u>

Bus Bridges

Although limits on peak bandwidth and bus length have been extended, the number of devices that can be connected to the bus is limited and cannot be increased without violating backward compatibility. A bus bridge can improve connectivity by fully utilizing the two-level SCSI addressing hierarchy already supported by the protocol. As described below, it would be a cost-effective alternative to the use of array controllers for this purpose. The goal of the EPI project is to define the architecture and functional requirements for such devices.

In the SCSI protocol, a device attached to the bus consists of two kinds of addressable entity: a control element attached directly to the bus (a SCSI target or initiator) and one or more logical units within each target that service SCSI commands. Each initiator and target has a unique physical bus address. Logical units attached to targets are identified by logical unit numbers.

Since the number of physical addresses on a bus is a function of the data path width, the combined total of initiators and targets is limited to eight for a narrow bus or sixteen for a wide bus. Because a target can have up to 16 logical units, each capable of executing SCSI commands independently, such a system could actually accommodate over 100 I/O devices. However, since targets usually contain no more than a single unit, this potential for added devices is unused.

A bus bridge is a two-port device consisting of a main bus attachment and a "back port" connected to an auxiliary SCSI bus. Target devices are attached to the auxiliary bus. Connected to the main bus are all the initiators, other bus bridges and conventional target devices. To an initiator residing on the main bus, a bridge appears as a SCSI target device with several logical units. In reality, each logical unit is a target device on the auxiliary SCSI bus. The task of the bus bridge is to:

- Respond to SCSI selection on the main bus or reselection on the auxiliary bus;
- After selection, intercept the logical unit number, convert it to a device address and select the corresponding target device on the auxiliary bus;
- After reselection, convert the auxiliary device address to a logical unit number and reselect the initiator on the main bus;
- Complete the connection between the real target on the auxiliary bus and the initiator on the main bus; and
- Once the connection is established, forward the data between the initiator and target devices.

In addition to increased device count, bus bridges can offer all the advantages of bus repeaters.

Peak data rate limits:

While 200 MT/sec has been easily achieved in the laboratory, indications are that speed increases above 100 MT/sec (200 MB/sec for wide SCSI) become a serious concern due to transmission media skew and jitter. In all probability, protocol changes will be required to operate parallel busses reliably at these speeds.

Increasing maximum bus length:

Existing bus timing is sufficient to support a bus length of 31 meters. Since the limit of 25 meters is based on technology that is over ten years old, the EPI project will reexamine this limit in the light of current technology to see if the bus length can be safely increased by reducing the timing margins.

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/