Parallel SCSI Grows, Shrinks and Stays the Same

By

Bill Ham, Digital Equipment

This article is intended to document specific features of parallel SCSI that characterize the migrations over the past few years and into the next few years. Features that grow, those that shrink, and those that stay the same are explored.

TABLE OF CONTENTS

1. Introduction....2
2. SCSI grows....3
3. SCSI shrinks....4
3.1 Shrinking the active circuitry....5
3.2 Shrinking the complexity....6
3.3 Shrinking the cost....9
3.4 Shrinking the interconnect....10
3.4.1 Connectors....10
3.4.1.1 VHDCI connectors....11
3.4.1.2 SCA-2 Connectors....14
3.4.2 External SCSI cable....15
3.4.3 Cable pass-thru techniques....16
3.4.4 Internal unshielded interconnections....17
4. SCSI stays the same....17
5. Summary and conclusions....18

1. **Introduction**

This article deals with the present and future of parallel SCSI. It explores how parallel SCSI is shrinking in several important parameters including size and complexity. It examines how SCSI continues to grow. Finally it shows how many of the most important features have not changed thereby maintaining ease of migration and preservation of previous investment while allowing graceful and low risk change.

SCSI has the property of being here now in virtually every form needed for nearly every computer system application. SCSI disks are available and used for notebook and desktop systems and occupy a continuum right up to the most demanding mainframe subsystems. The SCSI interface is found on every type of device that is important to the world of computing: disks (both rotating media and solid state), host adapters, complex storage controllers ,CD-ROM's, scanners, printers, frame grabbers, removable media devices, tapes, CD-R's, video input devices, and others. SCSI works "in the box" and from "box to box". SCSI therefore is the real thing for implementing complex and diverse systems. No other interface offers nearly the breadth and depth of applications.

It will require very substantial benefits for any other interface to displace SCSI as the interconnect of choice for sophisticated peripherals. The SCSI Trade Association (STA) is here to ensure that the benefits offered by competing methods remain non-substantial by incrementally evolving parallel SCSI in all the important areas.

As with any popular and long lived technology SCSI exists in forms that represent the earlier implementations along with newer evolutionary forms. SCSI, however, is unique because in almost every case one can use older SCSI devices in the same systems as the latest devices -- no need to throw away the perfectly good peripheral you already have just because a higher speed version or a new protocol twist is needed for some other new peripheral. More about this when we explore how SCSI stays the same.

The most recent evolutions take full advantage of the latest in cost effective, mature base technologies: silicon and board assembly, power and packaging, and interconnect.

Silicon is the most important ally in the migration. Higher speeds, smaller features, tighter timing tolerances, lower power consumption per function, and increasing integration levels combine to deliver huge benefits to the user with very little visible difference. The "pain of change" is largely shielded from the user of SCSI because the changes are built into the design of the silicon. At the same time, there is no need to change the silicon design tools, wafer fabrication technology, or test equipment just to accommodate SCSI. SCSI uses the same high volume silicon design and manufacturing technology as used for logic. No need to accommodate the special analog and test coverage features required by serial interfaces. SCSI delivers the cost, power, size, and reliability benefits that come with increasing levels of integration while riding in the sweet spot of silicon technology developments.

New classes of connectors and cables have recently been perfected by the interconnect hardware industry that allow a three fold reduction in the physical size of the interconnect. The smaller connectors and much more flexible cables combine to allow elimination of the right angle external cables. In addition to the size reduction a ten times reduction in the number of different kinds of cables needed is seen due to the elimination of the right angle versions. A four fold connectivity increase is delivered by achieving four wide ports simultaneously from a single PC option slot.

Without changing anything in SCSI devices, protocol, or software these interconnect and silicon developments combine to allow up to 60 SCSI devices to be connected to a single PCI slot in a PC (15 devices on each wide port). With UltraSCSI technology, 160 Megabytes/second peak bandwidth is available from this same slot (more than the total capability of the 32 bit PCI bus). A single PC could address 180 or more SCSI devices using just 3 PCI slots.

Interconnect technology developments also include the revolutionary “active interconnect” elements provided by the new “expander” products. These software/protocol invisible “smart wires” allow physical segmentation of SCSI interconnect in a manner similar to the familiar repeaters and hubs found in communication links (See Figure 1). This allows greatly extended lengths (up to 60+ meters) and star (or hub-like) interconnect topologies. It also allows mixing of SCSI devices with different transmission schemes in the same SCSI domain with no software impact.

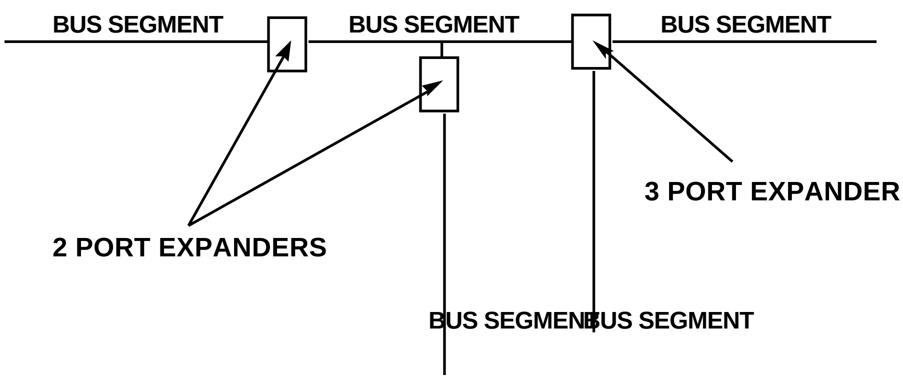

Figure 1 - Some examples of expanders in a SCSI system

The remainder of this article provides more details in each of these areas in terms of growing, shrinking, or maintaining SCSI features.

2. SCSI grows

The common denominator in parallel SCSI’s growth process continues to be evolution, not revolution. Some specific examples are discussed below.

The need for more speed was addressed through the introduction of the synchronous transmission scheme (SCSI-2) and higher basic “baud” rates. See Table 1. The need to go faster but with no cost premium was addressed by UltraSCSI (Fast 20). The need for even more speed with better configuration limits is being addressed through the use of the Ultra-2 low voltage differential (LVD) multimode technology (Fast 40). Multimode allows the same device or terminator to be either LVD or single ended.

The need for more addressable devices was initially accomplished by adding more data / address lines to the existing 8 bit cable (capable of addressing 8 devices) to make a 16 bit cable capable of addressing 16 devices. The interconnect technology developments extended this by 4x per PCI slot.

The need for better assurances on data integrity was addressed by using the REQ/ACK offset counting scheme (SCSI-2) where any failure of the counting scheme to balance at the end of the transmission indicated a faulty transmission.

The need for hot plugging devices was addressed by exploring the basic physics and applying simple, easy to implement, rules.

The need to operate in hostile electrical environments and with more noise margin was addressed by the introduction of the high voltage differential (HVD) transmission scheme (SCSI-2).

Some recent growth areas are:

- Low cost SCSI expanders (e.g. Symbios/Digital 53C120) that connect different or the same transmission type segments without any software involved and allow dramatic extensions of the single ended distance rules (20 meters point to point single ended - 12 meters fully loaded with LVD)

- System architectures based on the use of expanders that look a lot like serial (for example hubs, stars, T's) but have the beneficial parallel SCSI features

- Intelligent expanders that allow hundreds of devices to be addressed from a single SCSI port (LUN bridges)

- Very effective solutions to the issue of fair arbitration (to prevent bus "hogging"

- Even faster speed without significant protocol modifications (Fast 80/100)

- Much better use of the SCSI interconnect to distribute power

In some cases SCSI grows in capability by shrinking in certain ways. The next section looks at several specific ways that SCSI is shrinking.

3. SCSI shrinks

The shrinking trend in four important areas is very significant as suggested in Figure 2.

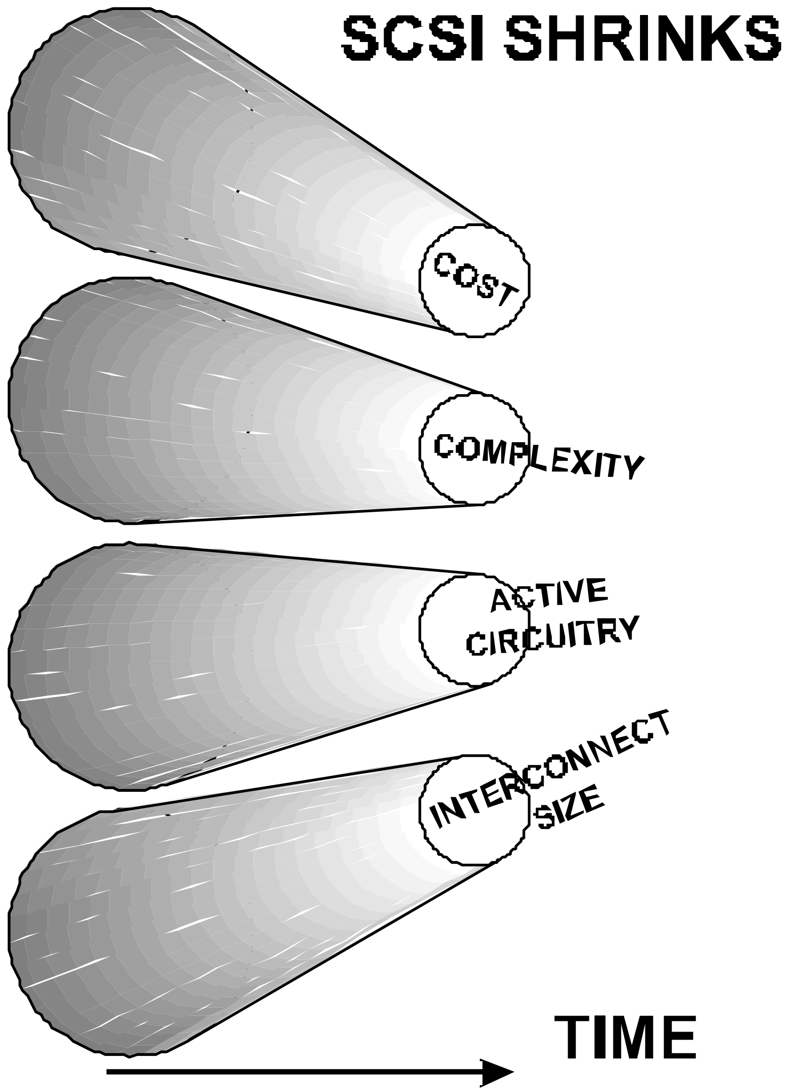

Figure 2 - SCSI SHRINKS OVER TIME IN FOUR IMPORTANT WAYS

3.1 Shrinking the active circuitry

SCSI has taken full advantage of silicon integration over the years to achieve dramatic size reductions. The first devices built in the 80's used discrete and medium scale integration parts on PCB's. These devices used full ISA slots and the large 50 pin SCSI-1 narrow connectors. In the early 90's low cost host bus adapters appeared based on the ISA bus and narrow SCSI. These occupied only a half slot, used large scale integration with full protocol and transceivers integrated into a single chip. The connectors were still narrow SCSI-1 with the first 50 pin SCSI-2 high density appearing in this time period.

Application of even higher integration levels and smaller connectors yielded the developments in later years with the first single chip adapters appearing in 1992/3 and quad channel adapters in 1997.

Figure 3 and Figure 4 summarize the dramatic impacts of these trends.

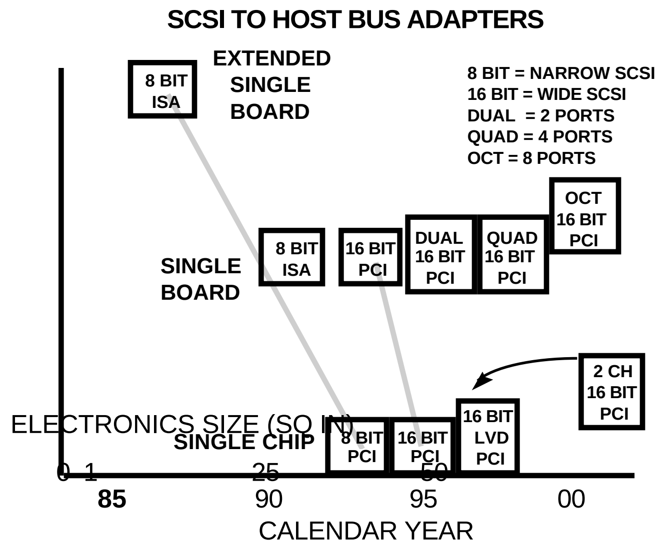

Figure 3 - SCSI electronics size per function shrinks dramatically

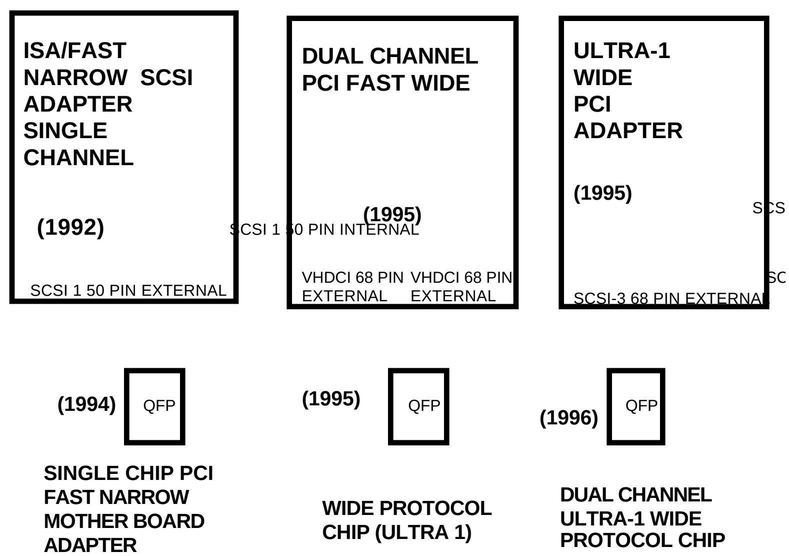

Figure 4 - Some examples of SCSI Host bus adapters and chips

The chips in Figure 4 are very approximately to scale.

3.2 Shrinking the complexity

Since SCSI can exist in several forms it has tremendous flexibility. Usually, large flexibility corresponds with large complexity (sort of like large price corresponds with large performance). SCSI is best thought of as a technology that excels in having a very low and shrinking complexity to flexibility ratio. The absolute complexity may not appear exceptionally low in some specific cases but when compared with the flexibility delivered one will see exceptional properties. Flexibility is the key to solving the broadest range of application problem AND to achieving excellent backwards compatibility and investment protection.

Perhaps SCSI's most underrated flexibility feature is its ability to use the same interconnect medium and physical interface to transmit signals at many different rates or speeds. SCSI can transmit and receive any speed from d.c. to the peak burst rate in the DATA phase without demanding anything other than an agreement on the peak rate that will be used between the two devices that wish to communicate.

Parallel SCSI is the only storage interface that allows this wide range of speed options with the same interface and interconnect.

Beginning with the Ultra-2 speeds, the single ended transmission mode becomes less attractive and is presently not defined for Ultra-2 speeds. Presently available data suggests that Fast 80 to Fast 100 operation will be possible without changing the interconnect media, connectors or Ultra-2 transmission type.

Table 1 shows a map of the SCSI speeds with the data rates, transmission types, cable media, and connectors.

Table 1 - Map of SCSI speeds with interconnect properties

SCSI speeds	Peak burst rate (wide)	Transceiver / terminator type	Cable media	Connectors
d. c.	NA	all	all	all
Up to Async	~ 3 MB/s	all	all	all
Up to Slow sync (5 MT/s)	10 MB/S	all	all	all
Up to Fast sync (10 MT/s)	20 MB/s	all	all	all
Up to Ultra-1 sync (20 MT/s)	40 MB/s	all	all	all
Up to Ultra-2 sync (40 MT/s)	80 MB/s	LVD and HVD	all	all
Up to Ultra-3,4 sync (80-100 MT/s)	200 MB/s	LVD	all	all
MT/s means millions of transfers per second (1 transfer is 2 Bytes in wide mode)				

Given the huge benefits of being able to use an "untuned" physical layer that can accommodate many different speeds, SCSI starts with a very shrunk complexity/flexibility plank in its architecture. Beyond that, the use of software-invisible expanders offer even more flexibility without exacting a speed related complexity price. The flexibility offered spans the range from single ended copper, to high voltage differential copper, to low voltage differential copper, to extended length optical (using special interfaces.

SCSI has maintained a high degree of backwards compatibility as newer SCSI versions are introduced. This has produced a creeping complexity increase in the interconnect hardware due mainly to the different connectors discussed above. There has always been and continues to be a mandate that backwards compatibility will be maintained with future generations. This is basically defined in terms of inventory non-obsolescence and ability to continue to use previous generation devices with newer generation devices. The inventory obsolescence feature is affected by much more than the SCSI interface but SCSI helps to mitigate the obsolescence pressures because of its backwards compatibility.

The number of connector form options existing in legacy SCSI systems arises because of the need to maintain this functional backward compatibility while addressing the needs for decreased size without sacrificing reliability. By focusing all new designs on the VHDCI family of connectors one greatly simplifies the interconnect job for the integrator of new SCSI systems because only a single connector interface is used and there is little need for right angle ends on cable assemblies. This single interface will accommodate every form of SCSI that presently exists.

With (1)the wide 68 position VHDCI connector for normal and high end applications, (2) the 40 position Micro SCSI VHDCI connector for the lowest cost, and most space constrained 8 bit SCSI applications and (3) the elimination of the need for right angle backshells on cable assemblies all the external interconnect applications can be met with a single set of cables. A few other cable/connector adapters may also be needed to connect to legacy systems and devices.

It is possible to have this interoperability between connector types and cable media types because SCSI relies on the parallel property of its architecture to keep the high frequency electrical effects to a bare minimum. For data transfers, the parallel SCSI data lines are operating approximately 10x slower than an equivalent baud rate serial line. SCSI can look forward to a number of years before the issues that challenge the serial interfaces in the electrical performance of the connectors and media become a dominant consideration. Schemes are being considered that increase the serial/parallel fundamental frequency ratio to approximately 20x.

The complexity associated with the testing and qualification of parallel SCSI devices is significantly reduced over serial interfaces because it is possible to gain electrical access to the SCSI bus without significantly disrupting the system. Even with the best probes and equipment one cannot presently count on the signals measured from a high speed serial application to be accurate or repeatable from laboratory to laboratory.

Another feature of SCSI’s complexity/flexibility tradeoff is the “in the box” / “out of box” capabilities. Less powerful parallel storage interfaces, for example IDE, are confined to very limited distances within a single enclosure. This limitation clearly reduces the complexity but at the price of flexibility. SCSI can be used “in the box” or “out of box”.

The latest features of Ultra-2 SCSI dramatically shrink the complexity in the manufacturing, inventorying, and distribution of SCSI devices by offering the multimode device and terminator interface. These multimode interfaces allow either single ended or low voltage differential (LVD) to be used through the same connector thereby allowing a single device part number to be useable in both new and legacy applications. Providers of product suitable for both low and high end applications (such as CD-ROM’s) can offer a single interface for all SCSI applications thereby lowering the complexity of the product line in all its dimensions.

The multidrop architecture of SCSI makes it a natural for applications where one needs to dynamically remove and replace devices. In a simple, yet useful, analogy one may think of the SCSI bus as a clothesline propped up at each end by the clothesline poles (the SCSI bus terminators) and the SCSI devices are the socks on the clothesline. Clearly one may remove and add socks without disturbing the other socks already on the line. (Other useful features of this analogy are (1) that if either or both clothesline poles (terminators) are not there the socks get dirty, (2) that if the clothesline (the bus wires) breaks the socks get dirty and (3) if the process of removing and replacing socks is not reasonably gentle some existing socks may get thrown from the line and get lost/dirty.)

How does this relate to complexity?? The act of dynamically reconfiguring a bus is perhaps the most compelling example of flexibility one can find. SCSI offers this capability with the complexity conditions that the process followed must adhere to the three rules: (1) a mutual ground must be established prior to actually contacting the live bus signals on insertion and it must be maintained until all live bus signals are disconnected on removal, (2) a device being added must stay in the high impedance state during the insertion process and until activated by the system, and (3) that traffic to a device being removed must be stopped prior to its removal.

All SCSI devices are required to maintain the high impedance state for purposes of controlling the hot plugging process. This seemingly simple requirement also means that SCSI devices may be powered off yet remain on the bus without causing any problems. In contrast, Devices on non-SCSI busses will lose logical continuity or will not maintain the high impedance state when not powered. Therefore, the flexibility of allowing unpowered devices comes with essentially no complexity cost (or other cost for that matter).

As SCSI moves to future generations, constant watch on the flexibility / complexity ratio will be one major feature by which to evaluate the details.

3.3 Shrinking the cost

The single greatest factor in lowering the cost per function of SCSI is the use of greater levels of silicon integration with each succeeding device design. SCSI does it without introducing analog circuitry into an otherwise digital silicon design and manufacturing process.

Other major forces helping to keep the cost under control are:

SCSI is established now and will stay that way for the foreseeable future. This means that the cost benefits resulting from high volumes of product will be there. SCSI is not about to become a specialized interface for the narrowly focused high end only.

The infrastructure to support all features of parallel SCSI are in place now. Connectors, cables, silicon, physical layer testers, protocol test beds, software drivers, OS support, lots of knowledgeable people, lots of training available... One place this is especially evident is in the silicon design and test area where SCSI chips effectively look like any other logic circuit. Some of the other I/O channel technologies require restructuring of the silicon design and test processes.

SCSI does the job NOW. This means that there will be multiple sources for products and this means real competition which in turn means rock bottom prices.

SCSI offers excellent system availability features which promise to get even better in the next generations. This means that the cost risk to business critical applications is low. The cost of a storage device, cable or whatever becomes effectively meaningless if a business is depending on a computer system that does not work.

Research is going on at most major computer companies and storage device suppliers to take SCSI forward. These efforts are not just a technology push but rather are a nearly perfect mix of technology with business and marketing needs. SCSI has survived by adapting to the needs of the market – it will continue this healthy evolutionary process.

There are no components of parallel SCSI that are out of the mainstream. This translates to companies having the ingredients available when needed to assemble SCSI based solutions at all levels of the integration hierarchy.

3.4 Shrinking the interconnect

3.4.1 Connectors

Spearheaded by activities in the SFF (formerly Small Form Factor) industry group, two new connector families are being standardized that offer unprecedented levels of functionality and true multisourcing of complete connectors for parallel SCSI. These families are the Very High Density Cabled Interconnect (VHDCI) shielded connectors that reduce the overall size of an external connector by 3 times and the Single Connector Attachment -2 (SCA-2) unshielded connectors that integrate all the functions needed to run a peripheral into a single connector. The VHDCI family revolutionizes the external SCSI interconnect and the controller parts of the internal SCSI interconnect and the SCA-2 revolutionizes the internal device interface.

For the first time complete connectors are being standardized -- not just the mating interface. This feature is essential to achieving interchangeability and second sourcing for connectors with the same style of termination side contact. The VHDCI family is specified in 26 different forms -- all with exactly the same mating interface -- so that virtually any kind of device or cable assembly design can be accommodated. Interestingly, this array of choices for the connectors does not increase the complexity of the interconnect but rather opens up new ways for product developers to design products while maintaining a simple and physically interoperable separable connector interface. In fact, this ability to accommodate a variety of product design requirements without changing the separable interface is one of the reasons for SCSI becoming LESS complicated.

Similarly, the SCA-2 connector for SCSI internal devices and cables is following the VHDCI standardization model with a significant number of intermatable forms being standardized. These connectors offer the ability to bring all the SCSI signals, all the power and ground connections, and all the “optional” signals such as ID’s, spindle sync, and power fail out of the device through a single unshielded connector. This dramatically shrinks the cost and complexity of interconnecting an array of SCSI devices.

With the SCA-2 connector, the device may be inserted into a backplane with no cables at all required. Without the SCA-2 and the backplane one needs a SCSI cable (50 or 68 conductor), a 4 lead power cable for ground and power (5V and 12V), and one or more smaller cables for the ID’s etc., for EVERY device in the system. Each of these cables is routed differently, has different current carrying and other electrical requirements and has very different connectors. While this cabled option is very flexible and offers some significant advantages in some systems, it has nothing going for it in the device array and modular packaging applications that are required for the higher end applications. Therefore, the SCA-2 is significantly responsible for the dramatic shrinking of the complexity of higher end SCSI device applications.

3.4.1.1 VHDCI connectors

The physical size of the VHDCI connectors are much smaller than the earlier versions as seen in Figure 5.

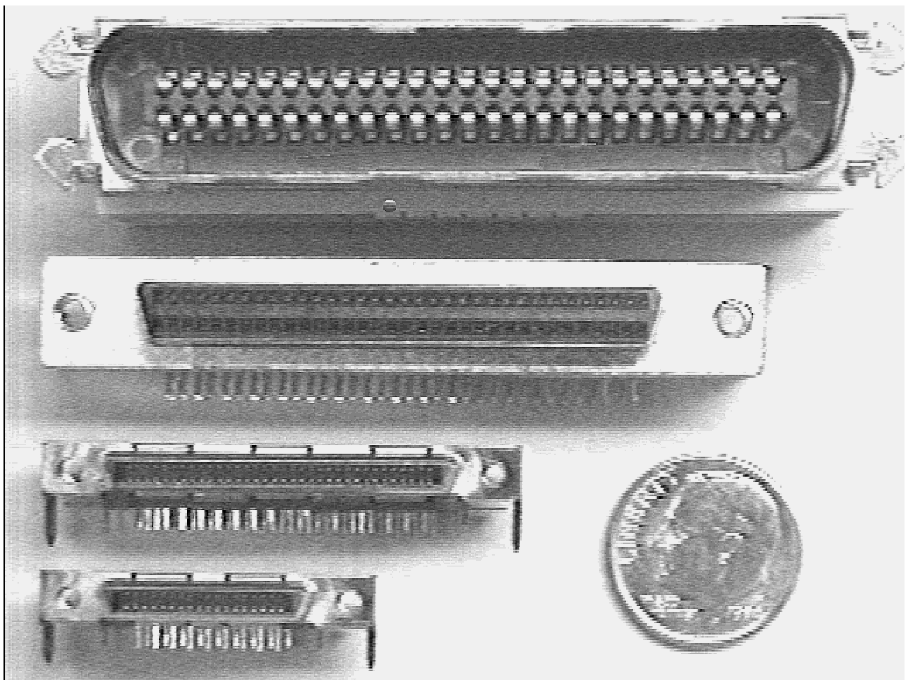

Figure 5 - EXTERNAL SCSI CONNECTORS SHRINKING

Due to its very low profile, the VHDCI 68 pin family is approximately half the height and twice the width of the latest Fibre Channel external connector, the HSSDC. Figure 6 shows a comparison of the VHDCI with the HSSDC. The same panel space is required for either technology.

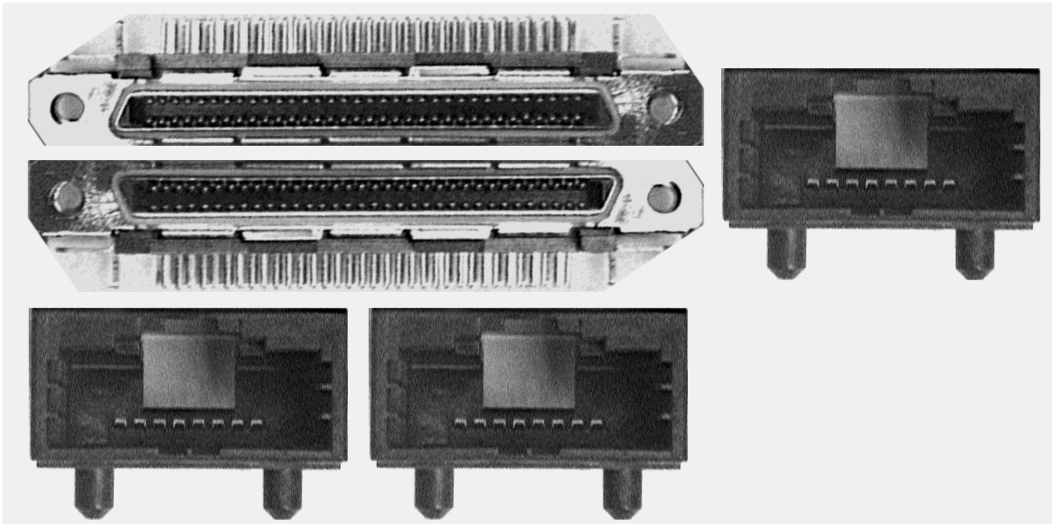

Figure 6 - 68 PIN VHDCI COMPARED TO FIBRE CHANNEL HSSDC

It may appear that there is no room for a cable assembly for the VHDCI connectors shown in Figure 6 because they are so close together. A closer look will reveal that the orientation of the polarizing shield connection is 180 degrees different between the upper and lower connectors. This allows an offset cable assembly to be used where one side is flat. This same cable assembly may be used on both the upper and lower connectors without interference. The specifications of the VHDCI interface ensure that neighboring PC option slots will not have interference even if all the SCSI ports have cable assemblies attached.

The VHDCI connector is very useful for multiport applications such as raid controllers. Figure 7 shows examples where the wide version of the connector family have allowed at least a doubling of the number of ports possible in a single controller form factor. It enables up to four wide SCSI ports in a single PC option card cutout as shown in Figure 7.

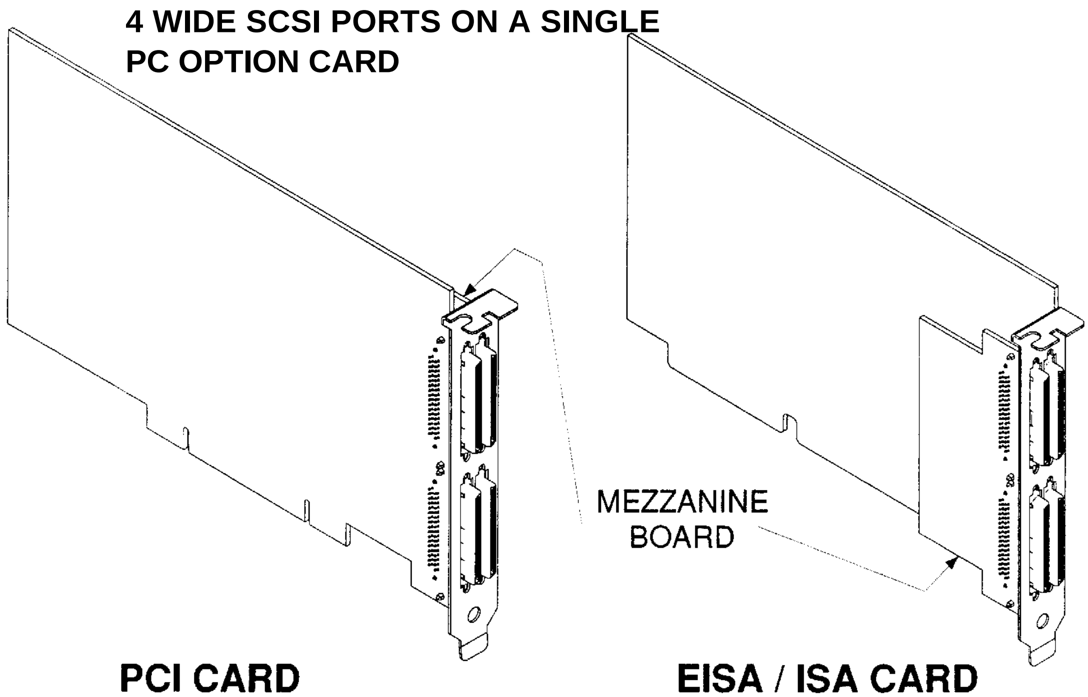

Figure 7 - Four wide SCSI ports on a single PC option card

The VHDCI retention scheme is also significantly simplified by introducing a three way retention post for the bulkhead connector. This post will accept (1) the conventional jackscrews, (2) a squeeze to release clip for positive retention with rapid release, or (3) a detent ring retention that requires a stronger pull than with no retention but no action other than pulling or pushing. The choice of retention type is made in the cable assembly. All 68 pin VHDCI cable assemblies that comply with the SFF specifications will work on all 68 pin VHDCI mating connectors.

Figure 8 shows the detail of the 68 pin VHDCI system. The lip in the jackpost provides the securing point for squeeze to release clips as well as for split ring detent type retention. The center of the jackpost is threaded for use with jackscrews.

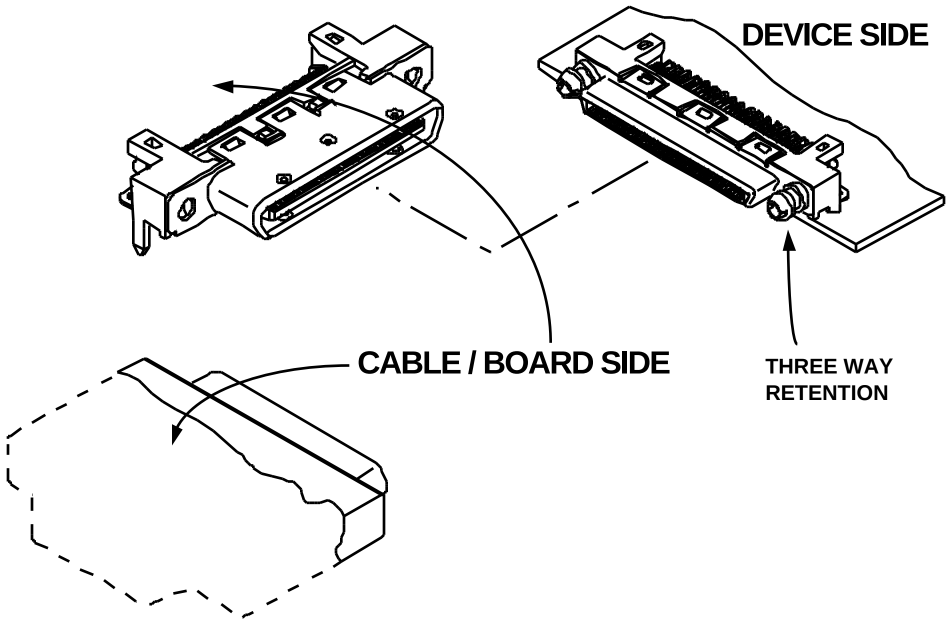

Figure 8 - Overall view of the 68 pin VHDCI system

Even though the VHDCI is much smaller than the high density connector it is very tough. There are no pins to bend, the retention scheme uses the same size jack screw thread as the high density wide connector, and the connector is made with the contacts imbedded in the housing where they cannot move or become distorted.

3.4.1.2 SCA-2 Connectors

The SCA-2 connector was introduced to SFF a couple of years ago. Along with the original SCA-1 it is rapidly becoming a popular internal device connector for single port devices such as disk drives. The SCA family uses an 80 pin leaf style contact to interface all the active SCSI lines, three power voltages, and device control signals. This connector is considerably smaller than the collection of the three different connectors used for power, options, and SCSI bus in a cabled system.

The SCA-2 offers all the features of the SCA-1 and is compatible with the SCA-1 but adds several important capabilities not found in the SCA-1. These features are mainly related to the needs of dynamic device plugability.

First, two contacts, one on each side of the connector provide the first make / last break for the ground connection. This ensures that a common electrical ground is established between the device and the system before any power or signal connections are made when the device is being inserted. Upon removal these contacts ensure that the ground stays intact throughout the disengagement of the signal and power pins.

Second, provision is made to allow the special long power contacts to precharge bypass capacitors before the main power contacts make. This reduces the disturbance to the power distribution system and eliminates any arcing on the service power pins. When the connector is fully mated two pins at the extreme ends of the connector provide indication that the connector is fully mated. The overall view of the SCA-2 system is shown in Figure 9.

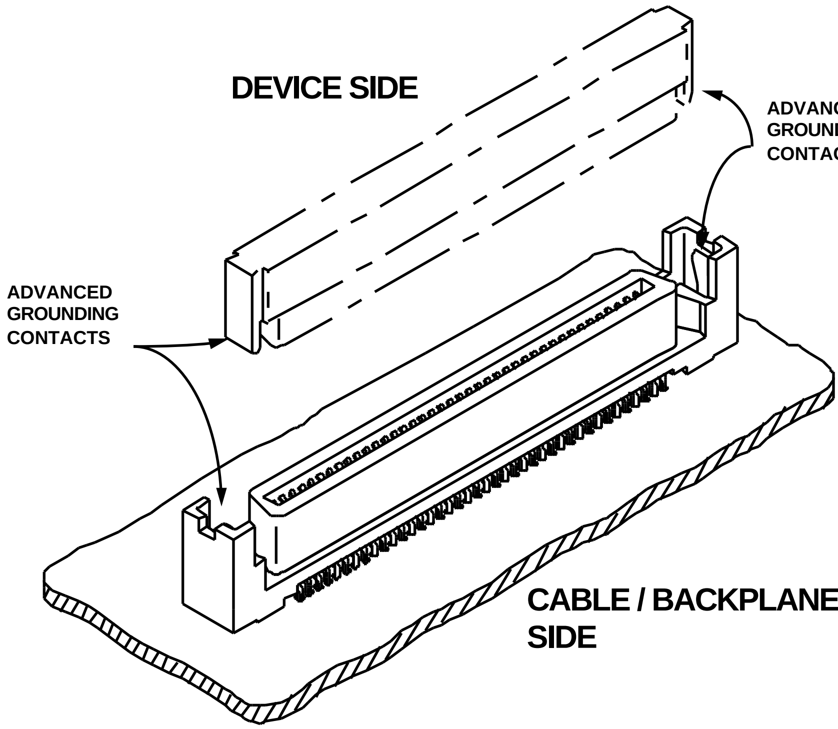

Figure 9 - Overall view of the SCA-2 connector system

The size decrease is not dramatic for the SCA family due to the need to maintain enough size to achieve blind mating alignment and due to the lack of market pressure to further decrease their size (for backplane applications there is little advantage in having a connector that is smaller than the device). With 3.5” or the newly proposed 3” form factor devices the SCA connector comfortably fits within the device boundaries.

The use of backplanes for direct device attach is enabled by having all the electrical connections for the device available in one connector on the device. This eliminates not only the cables used for device attach but also the space required for the connectors and produces a very significant shrink in the size required to package multiple devices.

3.4.2 External SCSI cable

The external cable for SCSI is also shrinking through the use of smaller gauge wire, better dielectrics, and less jacketing material. Formerly, wide SCSI needed a cable that was approximately 0.5” in diameter (0.196 sq.” cross section) with 28 gauge wire. Today, wide SCSI cables with 30 gauge wire are shipping with diameters of 0.37” (0.107 sq.”). Cables with 0.3” diameter (0.07 sq.”) are possible with 32 gauge wire and inexpensive dielectrics for wide SCSI. Cables with 0.25” diameter (0.049 sq.”) for narrow SCSI (0.070 sq. in) are very flexible and friendly -- similar in size and flexibility to a desktop computer power cord and smaller than many serial cables. When used with active single ended, LVD, or HVD terminators the 32 gauge wire is adequate for distributing TERMPWR and SCSI signals in most applications. Very long cables should not be used for TERMPWR distribution.

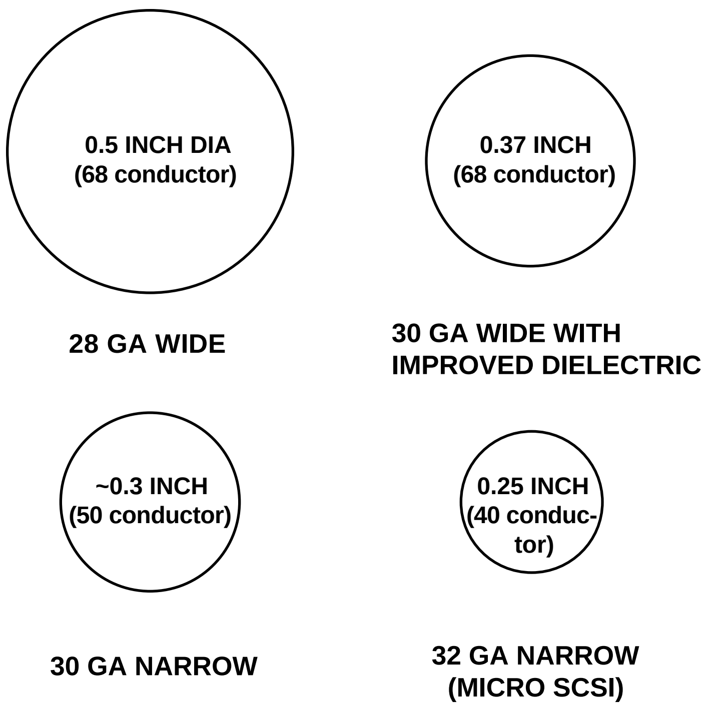

Figure 10 - External SCSI cable diameters

While further shrinks in the connector and cable size may be possible these need to be weighed against ease of handling, sufficient strength to survive normal service stresses, and costs increases at very small sizes. With the VHDCI connector and 30/32 gauge wires sizes a good optimization exists.

3.4.3 Cable pass-thru techniques

Shielded cable pass-thru can be used for cases where no separable external connector is needed or where more than four wide ports need to pass through the same PC option slot. Figure 11 shows a five cable version.

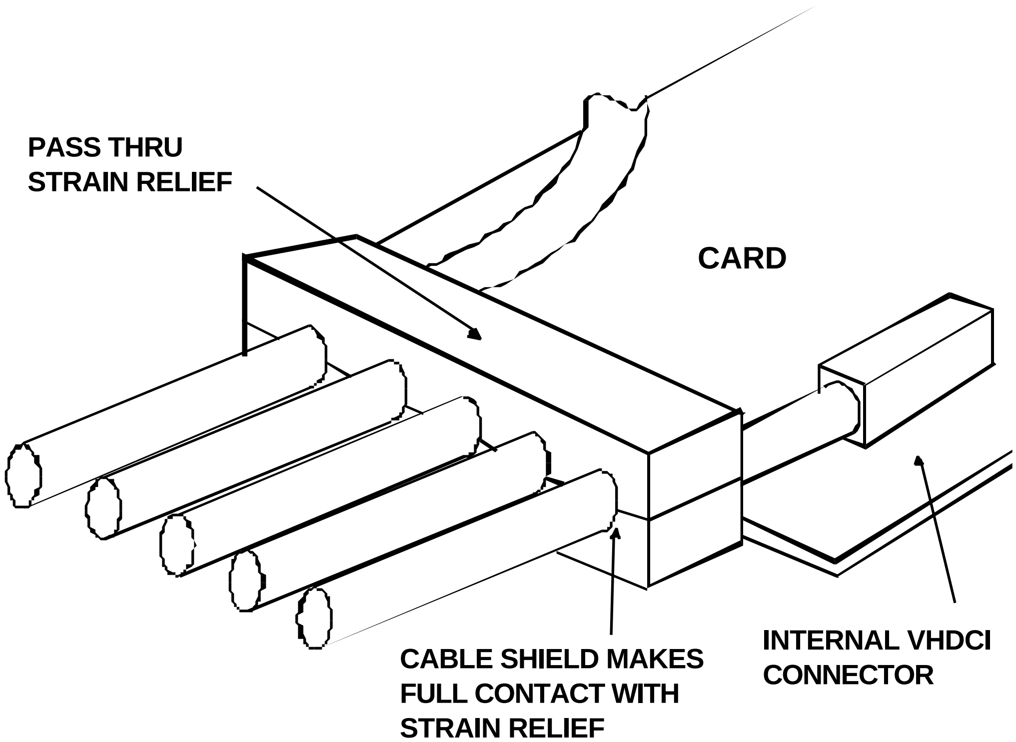

Figure 11 - Pass-thru option (5 ports shown)

3.4.4 Internal unshielded interconnections

The VHDCI is well suited for internal cable connections as well as for external cable connections in multiport controllers applications. This scheme enables many SCSI ports, or a combination or SCSI and other functions, to exist on the same side of the controller PCB. When using vertical headers on the PCB much less real estate is consumed than with previous connectors. Controllers with as many as 8 wide ports can be accommodated on a single PCI style PCB.

4. SCSI stays the same

This section is really very simple. There are two features that do not change as SCSI evolves.

1). SCSI uses the same protocol for arbitration, selection, negotiation, and data transmission everywhere. This means that useful physical interoperability is possible over a very wide range of conditions. This base allows most software and firmware to evolve prudently while maintaining backward compatibility and allows the same architectural models to be applied and reapplied. The impacts of these protocol architectural features remaining invariant cannot be understated. Much of the energy spent in the successful deployment of any technology lies in the maturity of the underlying software. By allowing the software base to remain constant as far as most areas of SCSI are concerned the risks to introducing any new evolution is minimized.

2). SCSI uses the same interconnect media for all forms. This keeps the SCSI “roadbed” intact while developments occur in the devices, configurations, and performance.

It presently appears that there will be at least two more generations of SCSI before either of these foundational principles will need to be changed.

5. Summary and conclusions

SCSI grows while it is shrinking and staying the same. What more could one ask for from an I/O interface?

The author is employed by Digital Equipment as the manager of the Storage Bus Technology Office.