


in-depth exploration of serial attached scsi



David A. Deming
Robert W. Kembel



1. Introduction to SAS

1.1 What is Serial Attached SCSI?

Serial Attached SCSI (SAS) is a standardized, next-generation I/O interface that provides an evolutionary growth path from the existing SCSI bus. SAS takes advantage of the latest advances in high-speed serial data transmission and switching to provide transfer rates of 150 MBps (megabytes per second) and 300 MBps with an anticipated growth path to 600 MBps and beyond.

SAS provides much larger configuration sizes and greater configuration flexibility than existing bus-based environments through the use of inexpensive switching devices called expanders. Whereas the SCSI bus is limited to a maximum of 16 total devices, a SAS configuration could theoretically have up to 16,384 devices.

Unlike the SCSI bus, where all devices share use of a common bus, SAS devices are connected using switched point-to-point links. This has many benefits, including improved signal integrity, no stubs, integrated terminators and the ability to have simultaneous communications between multiple devices at the same time. This can greatly improve performance in multi-initiator configurations or when doing device-to-device operations such as outboard copy.

Most SAS initiators and expanders are compatible with Serial ATA (SATA) devices. This means, that for the first time, a single storage enclosure or subsystem can include high-performance SAS devices for primary storage and lower cost SATA devices for secondary or bulk storage. This enables the storage designer or user to efficiently address the need to store different kinds of data within the same storage environment (or upgrade the storage environment from SATA devices to higher-performance and higher-reliability SAS devices when the need arises).

SAS provides all of these benefits while preserving much of the existing investment in firmware and software and enables developers to build upon the experience and development work previously done for SCSI bus-based devices.

1.2 A Brief History of SCSI

The SCSI bus has a long and extremely successful history beginning with the introduction of SASI (Shugart Associates System Interface) in 1979 as an eight-bit parallel interface designed as a device-independent peripheral interface. SASI was ultimately released into the public domain in an effort to get disk controller manufacturers to build SASI products. It received wide market acceptance due to its technical merit and Shugart's marketing efforts.

In 1978, the American National Standards Institute (ANSI) subcommittee X3T9.3 initiated a project to develop an interface suitable for small computer systems. The first interface considered was the Intelligent Peripheral Interface, or IPI. SASI was presented as a candidate in 1980, but rejected by the committee.

Shugart joined forces with NCR in 1981 to add features to the proposed SASI interface (such as 10-byte commands and cable options enabling data transfer over greater distances. The revised interface was again proposed in 1982 to X3T9. The committee found SASI an elegant but partial solution and assigned SASI to X3T9.2, an existing task group with no current project.

In April of 1982, X3T9 met and drafted a formal proposal for SCSI, the Small Computer System Interface. The committee began by adding additional capabilities such as peer-to-peer communication, logical units, a message subsystem, arbitration and synchronous transfers. By 1983, the standard was relatively mature and in April, NCR announced the first SCSI protocol chip, the NCR 5385.

By the beginning of 1985, X3T9 had defined a number of SCSI commands but, unfortunately, not all of the manufacturers were implementing the complete set. The industry feared fragmentation of the market, and interoperability problems, if each manufacture implemented a different subset of the commands. At this time, a number of drive manufacturers banded together to define a common set of commands that they all would support. This was termed the Common Command Set, or CCS, and was later incorporated into the SCSI standard.

The SCSI standard (now referred to as SCSI-1) was approved as a national standard and has since been accepted by the European Computer Manufacturer's Association (ECMA), as a Federal Information Processing Standard (FIPS) and by the International Standards Organization (ISO) as an international standard. SCSI-1 provided the following characteristics:

- An 8-bit wide data bus
- A transfer rate of 5 MBps
- A maximum of 8 devices per bus
- 5v single-ended and differential signal levels (sometimes referred to as "high-voltage")

Work on a follow-on standard called SCSI-2 started immediately and included the following extensions:

- Incorporation and enhancement of CCS for disk drives
- Support for several new devices including scanners, optical disk media changers and CD-ROMS
- Command queuing
- A two-byte wide bus option (the "wide" bus)
- An optional transfer speed of 10 mega-transfers per second (the "fast" transfer rate)

The combination of the original narrow and optional wide bus along with the fast transfer rate resulted in the speeds shown in the first two rows of Table 1-1 on page 5.

SCSI-3 standards development began in 1993. Much of the SCSI-3 work was associated with reorganization of the standards to separate the logical command and architecture elements from the physical interface. This enabled SCSI commands to be sent over serial interfaces such as Fibre Channel or IBM's Serial Storage Architecture (SSA).

In the period since the adoption of the SCSI-3 standards, several new bus variants have been approved.

Standard	Data Width (max)	Data Speed (max)	Bandwidth (max)	Signals	Devices & Distance
SCSI -1	8 bits (“narrow”)	5 MHz	5 MBps	SE	8 devices, 12 meters
				HVD	8 devices, 25 meters
SCSI-2	16 bits (“wide”)	10 MHz	20 MBps	SE	16 devices, 12 meters
				HVD	16 devices, 25 meters
SCSI-3	16 bits (“wide”)	20 MHz	40 MBps	LVD	16 devices, 12 meters
		40 MHz	80 MBps	LVD	16 devices, 12 meters
		40 MHz (double-edge clock)	160 MBps	LVD	16 devices, 12 meters
		80 MHz (double-edge clock)	320 MBps	LVD	16 devices, 12 meters

Table 1-1. SCSI Bus Characteristics

- The first of these, Ultra SCSI, doubled the data transfer clock rate to 20 MHz. This provided a data transfer rate of 40 MBps when using a “wide” bus (or 20 MBps with a “narrow” bus).
- The clock rate was subsequently doubled again to 40 MHz providing an 80 MBps data transfer rate when using a “wide” bus (or 40 MBps with a “narrow” bus).
- The next extension was to clock the data on both the rising and falling edges of the clock (double-edged clocking). This provided a 2x increase in data transfer rate to 160 MBps without requiring a higher clock rate (usually referred to as SCSI Ultra-160).
- In what is likely the final bus extension, the clock rate was increased again to 80 MHz and retained double-edged clocking to provide a 320 MBps transfer rate on a “wide” bus.

It appears unlikely that the SCSI bus will see any further extensions in bus width or bus speed. Wider buses (with more data bits) require larger cables and connectors as well as more drivers and receivers in the electronics. Faster bus speeds require greater control over the timing variations of the individual bits as they travel down the cable. Timing variations in the individual bits (skew) presents a significant challenge at the 80 MHz clock rate and makes it unlikely that higher speeds will be practical for an I/O interface such as the SCSI bus. An illustration of the skew problem is shown in Figure 1-1 on page 6.

1.3 Why Serial Attached SCSI?

There are numerous technical and marketing reasons behind the development of Serial Attached SCSI (SAS). Among the technical reasons for developing a new interface are:

- The need to provide a faster interface with a roadmap for future growth
- The need to attach more than the 16 devices that are possible with the SCSI bus
- The need to support longer distances than are possible with the SCSI bus

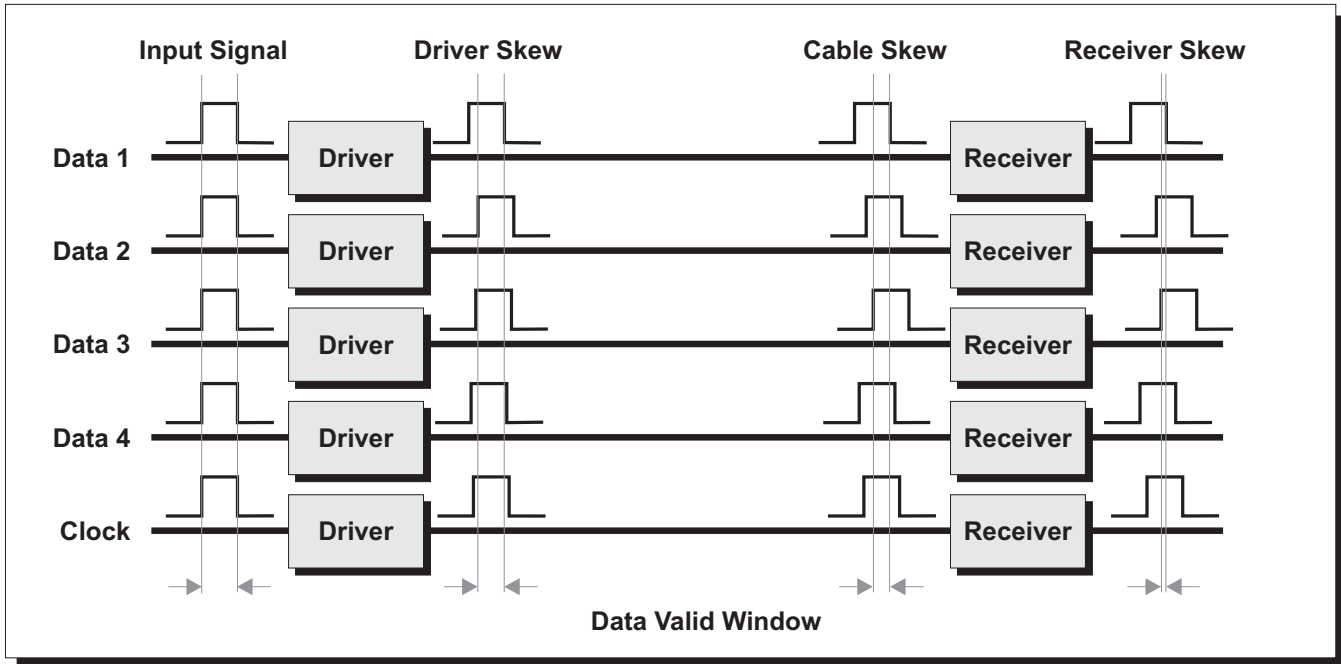


Figure 1-1. SCSI Bus Skew Example

- The need to provide better fault isolation (if a device on the SCSI bus fails, it may cause the entire bus to fail)
- The need to reduce cable and connector costs. The 68-pin cable and connector used for the SCSI bus are expensive
- The need for smaller cables and connectors in order to support smaller form factor devices such as 2.5" disk drives
- The need to provide dual-port capabilities for alternate, or redundant paths, even with smaller devices
- The need to reduce the number of I/O pins on controller and device ASICs

In addition to the technical reasons for developing a new interface, there are a number of market-driven goals as well.

- The desire to provide downward compatibility with lower-cost SATA devices for secondary storage applications. Because existing SCSI and Fibre Channel devices are designed for high-performance and high-reliability, they come at a price premium. SATA devices are high-volume devices designed for lower cost and applications that can tolerate higher failure rates and lower-duty cycles in order to provide the most cost-effective disk storage in terms of dollars per megabyte and capacity.
- The desire to accommodate both classes of disk drives in the same subsystem and enclosure. This means that the SAS physical interface must be compatible with the SATA physical interface, even to the point of providing compatible connectors.
- Provide a mid-tier storage option that can be priced between the low-end SATA devices and the high-end Fibre Channel devices.

1.3.1 SAS Roadmap

The SCSI Trade Association developed a roadmap to help guide the evolution of the Serial Attached SCSI technology. The roadmap outlines a doubling of the signaling rate to 6 Gbps in the 2007 time frame and another doubling to 12 Gbps in the 2010 time frame as shown in Figure 1-2.

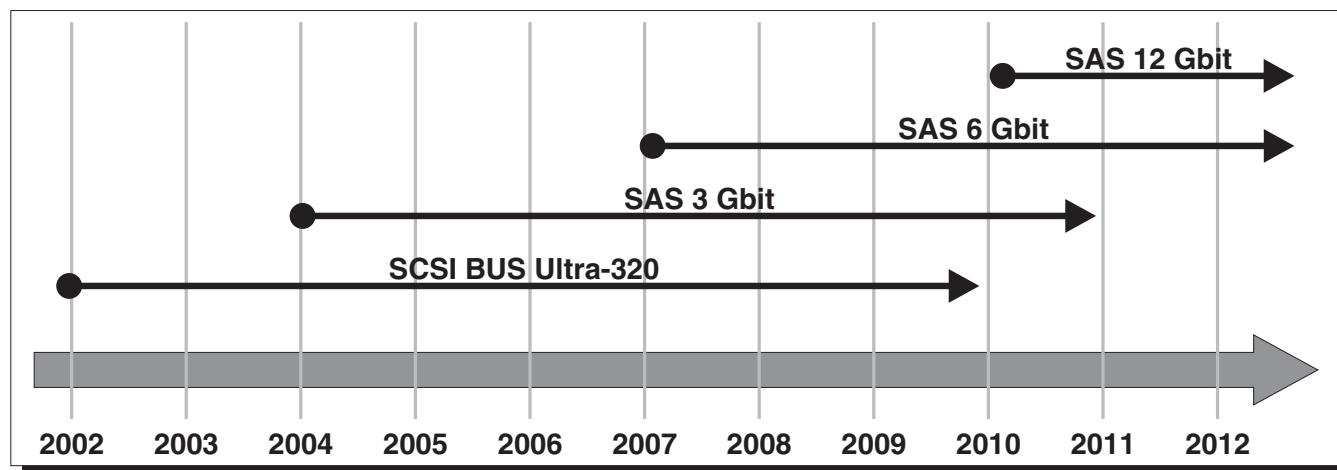


Figure 1-2. SAS Technology Roadmap

1.3.2 What SAS Isn't

- SAS is not transparent to the SCSI class driver or port drivers because SAS alters some SCSI mode pages and log pages.
- SAS is not designed to be a Storage Area Network (SAN) interface. While it can be used to create modest-sized configurations, it is not designed for long-distance applications. SAS is specified for operation via electrical cables with a probable maximum distance of about 10 meters at 3 Gbps or 5 meters at 6 Gbps. Even if optical links were to be used, there are aspects to the protocols that adversely affect performance on long links.
- SAS does include the use of simple switches, called expanders, to provide connectivity, but the resulting topology is strictly limited to a spanning tree structure. SAS does not support the kind of unrestrained topologies that are possible with Fibre Channel or IP networks. Within the spanning tree, there may be redundant point-to-point links between ports, but redundant paths between different ports are not allowed.
- SAS is not designed to support a broad range of protocols. It is only designed to support the protocols necessary to perform SCSI commands, SATA operations and some basic management functions. It is not designed to support protocols such as the internet protocol (IP) or FICON.
- SAS is not necessarily simpler to design or easier to implement than other serial interfaces such as Fibre Channel. The amount of function required in a SAS port is roughly equivalent to the amount of function required in a Fibre Channel port (after all, they are both doing roughly the same things). SAS expanders are probably not any simpler than Fibre Channel Arbitrated Loop switches (again, the functions are roughly equivalent).