



# Serial ATA technology

technology brief, 2nd edition

Introduction.....	2
Parallel ATA technology .....	2
Serial ATA technology.....	3
Low voltage differential signaling.....	4
Signal integrity .....	4
SATA Performance .....	5
SATA 1.5 Gb/s .....	5
SATA 1.5 Gb/s (with extensions) .....	5
SATA 3.0-Gb/s.....	5
SATA devices .....	5
Cabling and connector .....	6
Initiators .....	6
Port multipliers .....	7
Targets.....	7
SATA/SAS interoperability.....	7
SATA topologies.....	8
Internal .....	8
External .....	8
Conclusion.....	9
For more information.....	10
Call to action .....	10

# Introduction

The parallel Advanced Technology Attachment (ATA) architecture has limitations that make it unsuitable as a long-term solution for direct-attached and network storage applications in the enterprise: short, bulky ribbon cables that impede airflow and are difficult to install; a large number of wires and high frequency signals that are prone to electrical interference; and the lack of hot plug capability. Building on the success of other serial standards (Ethernet, USB, and FireWire), serial ATA technology is now providing the performance, scalability, and hot plug capability needed for future generations of storage devices.

HP has actively participated in the Serial ATA (SATA) International Organization (IO) to help develop specifications that meet the needs of cost-sensitive entry-level to mid-range servers and non-mission critical network storage. With the introduction of SATA products in 2004, key solution providers like HP are now offering more scalable and affordable configuration options for a broad range of storage applications.


This technology brief begins with a description of parallel ATA technology and the reasons for the industry's transition to serial I/O technology. Next, it describes how serial ATA technology overcomes the limitations of parallel ATA. Additional sections describe possible internal and external SATA topologies.

## Parallel ATA technology

Parallel ATA has been the dominant interface for desktop and notebook products since it was introduced in the 1980's as IDE (Integrated Drive Electronics). The success of parallel ATA in the desktop and notebook markets is attributed to frequent performance enhancements, backward compatibility, and a constant goal of driving costs as low as possible.

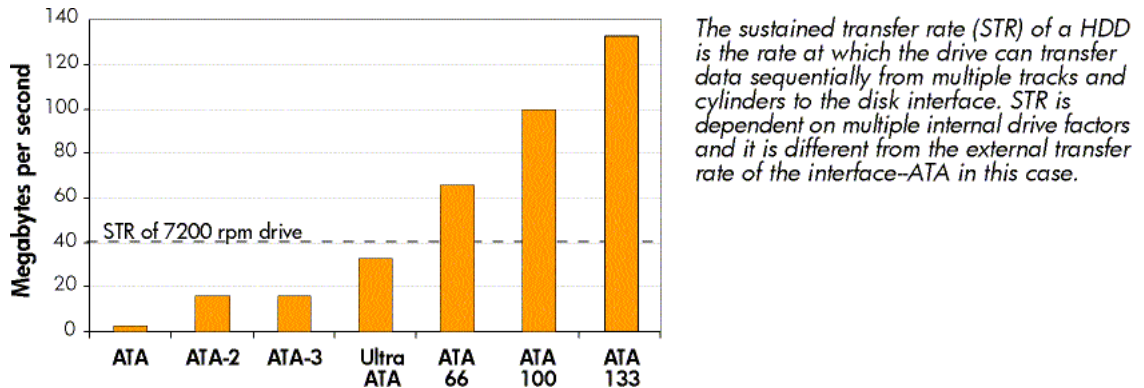
Figure 1 shows the major components in the parallel ATA architecture. An ATA controller and two parallel ATA connections are built into the motherboard. Up to two devices can be connected to each ATA connector in a Master/Slave relationship by using an 80-conductor ribbon cable with three 40-pin connectors. The bulky cable impedes airflow in the cabinet, which hinders cooling. Cooling is not a big issue for desktop PC users who rarely install more than one drive.

**Figure 1.** Parallel ATA cable with 40-pin connectors for master and slave drives



After the introduction of parallel ATA, its data transfer rate increased from 3 megabytes per second (MB/s) to 133 MB/s (Figure 2). ATA 100 and ATA 133 have the headroom to handle the sustained transfer rate of today's 7,200-RPM hard disk drives (HDDs) because the interface has to accommodate only one drive at a time.

**Figure 2.** Data transfer rates for parallel ATA modes vs. the sustained transfer rate (STR) of HDDs



This performance graph questions the need to change to a serial interface if ATA 100 can handle the requirements of desktop class (5400-RPM and 7200-RPM) HDDs. The answer concerns signaling voltage and data reliability. Parallel ATA data transfer is based on transistor-to-transistor logic (TTL) signaling. TTL signals define an 8-bit digital value, through a sequence of high and low voltage states, on pins 2 through 9 of the parallel port at a given point in time. TTL uses 5V-tolerant, 3.3V signaling, which requires integrated circuits that can tolerate input signals up to 5 volts. It is becoming increasingly difficult to support the traditional 5V TTL signal requirement because components are being fabricated with finer and more fragile lithographies.

With regard to data reliability, ATA uses cyclic redundancy checking (CRC) to verify the accuracy of the data signals transmitted between the host and HDD controller. However, ATA command signals are not checked with CRC, so they remain a potential source of error.

It would be very difficult to increase the speed of ATA beyond 133 MB/s due to the 5V signaling requirement and the increased likelihood of issues with the integrity of command signals. With parallel bus architectures, the data and clock signals are transmitted along parallel wires from the initiator to the target device at a specific signaling rate. As the signaling rate increases, it becomes increasingly difficult to keep the data and clock signals aligned. In addition, signal integrity is degraded by the electrical noise that results from switching all data signals at the same time.


## Serial ATA technology

Serial ATA discards the parallel ATA Master/Slave concept and only allows one device per cable, which the system views as a master ATA device. These point-to-point connections allow each drive to communicate with the controller without having to wait for other data traffic to clear first. SATA addresses the electrical signaling and signal integrity issues that inhibit increasing the speed of parallel ATA beyond ATA 133. SATA technology has the potential to shrink form factors, lower power consumption, and extend I/O performance to meet the bandwidth requirements of a new wave of technological advances.

## Low voltage differential signaling

SATA technology transmits signals in a single stream rather than in multiple parallel streams. SATA incorporates an low voltage differential (LVD) signaling scheme that uses two pairs of data lines to transmit and receive low-voltage signals (250 mV). The data is represented by the voltage potential between the two wires in each pair (Figure 3). Because it takes less time to apply low voltages to the wires, LVD signaling can occur at a much greater speed than in parallel ATA. The low voltage reduces the effects of capacitance, inductance, and noise. Noise sources tend to add the same amount of voltage to both wires, so the voltage difference between the wires remains the same.

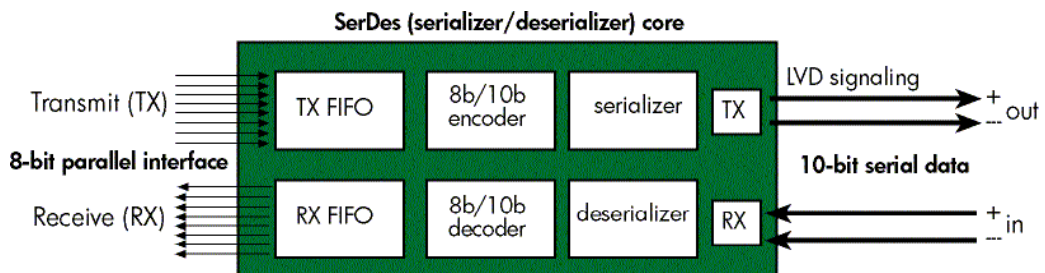
**Figure 3.** LVD signaling



## Signal integrity

Serial architectures encode (embed) the clock signals into the data stream, thus eliminating the skew problem with aligning data and clock signals. Serial architectures require significantly fewer data lines to switch simultaneously, which reduces the introduction of electrical noise. As a result, serial signaling rates can be increased well beyond those attainable with a parallel bus. Serial communication requires a device to convert parallel data into a serial bit stream and vice versa. This device, called a serializer/deserializer (SerDes), contains a parallel digital interface, First-In-First-Out (FIFO) caches, 8 bit/10 bit (8b/10b) encoder and decoder, a serializer, and a deserializer (see Figure 4). The 8b/10b encoder converts each 8-bit data byte to a 10-bit transmission character, which enables clocking information to be encoded into the data stream. Although this adds about 20 percent embedded overhead to the data stream, it eliminates the clock skew problem experienced by parallel ATA.

**Figure 4.** The SerDes core integrates 8b/10b coding and decoding logic.



## SATA Performance

The SATA specification was released in three parts: SATA 1.5 Gb/s, SATA 1.5 Gb/s with extensions, and SATA 3.0 Gb/s. The 10-year roadmap for Serial ATA allows it to scale eventually to 6.0 Gb/s.

### **SATA 1.5 Gb/s**

The SATA specification was introduced with a maximum theoretical bandwidth of 1.5 Gb/s, or 150 MB/s, factoring 20 percent encoding overhead (see "Signal integrity"). SATA 1.5 Gb/s focuses on increasing the bandwidth and mitigating the design problems associated with the parallel ATA architecture. SATA 1.5 Gb/s is designed to replace parallel ATA, mainly in desktop PCs (non-hot plug). Serial ATA is not hardware compatible with legacy Ultra ATA; however, it is fully compliant with the ATA protocol and, thus, software compatible with existing ATA drivers.

### **SATA 1.5 Gb/s (with extensions)**

SATA 1.5 Gb/s (with extensions) addresses the needs of the lower cost server and non-mission critical enterprise storage markets. SATA 1.5 Gb/s (with extensions) enhances SATA 1.5 Gb/s with features such as native command queuing, out-of-order execution and delivery, and data scatter/gather lists (each described below). Any or all of these optional extensions can be implemented.

Native command queuing enables a hard drive to take multiple requests for data from the processor and rearrange the order of those requests to maximize throughput. SATA 3.0-Gb/s hard drives will be able to queue and execute requests without any assistance from the CPU.

Out-of-order execution and delivery keeps execution resources as busy as possible. In the native command queuing model, this feature allows the last half of the data requested by a command to be delivered and executed before the first half of the data. Out of order data delivery within commands requires support for non-zero buffer offsets.

A data scatter/gather list is a data structure that assists the direct memory access (DMA) engine in locating memory regions that comprise the complete transfer buffer. This assistance is beneficial because virtual memory mapping mechanisms may scatter the buffer across several noncontiguous, physical memory pages.


### **SATA 3.0-Gb/s**

The SATA 3.0 Gb/s specification doubles the previous data transfer rate and adds other extensions to improve the capabilities of SATA devices for server and networked storage applications. Theoretically, with the use of port multipliers, each port on a SATA 3.0 Gb/s host controller can connect up to 15 SATA drives.

## SATA devices

SATA devices include initiators (SATA controllers), port multipliers, and targets (SATA drives) as shown in Figure 5. Port multipliers connect initiators to targets in a SATA domain. SATA devices in a domain do not need to be assigned an ID. All SATA devices have unique worldwide names (SATA addresses) assigned at manufacturing to simplify identifying initiator devices, port multipliers, and target devices. These devices, and the cabling that connect them, are described in this section.


**Figure 5.** Example of SATA devices in a domain.



### Cabling and connector

SATA replaces the parallel ATA 40-pin connector and 18-inch long flat ribbon cable with a 7-pin connector (four signal lines and three ground lines) and a small diameter cable up to 1 meter long (Figure 6). The thinner serial cable improves both airflow and routing. SATA also features a new 7/8-inch wide, 15-pin, single-row power cable connector. The power connector provides optional hot-plug capability, which allows a drive to be swapped out without powering down the machine.

**Figure 6.** The 7-pin SATA data cable and the 15-pin power cable



### Initiators

A SATA initiator is a controller that can be embedded into the motherboard (Figure 7) or a host bus adapter (HBA) plugged into a PCI expansion slot. Embedded SATA controllers typically have two ports for connection to internal SATA devices while a SATA HBA can have a number of ports for internal or external devices.

**Figure 7.** Embedded SATA controller



## Port multipliers

SATA port multipliers require host controllers that are port-multiplier-aware, such as SATA 1.5 Gb/s (with extensions) and SATA 3.0-Gb/s. Therefore, port multipliers are not compatible with original SATA 1.5 Gb/s controllers. However, port multipliers are compatible with SATA 1.5 Gb/s drives. Port multipliers can only link to SATA drives; they cannot link to another port multiplier.

The extensions in SATA 3.0 Gb/s require updated operating system and driver support. SATA 1.5-Gb/s devices (that support extensions) and SATA 3.0-Gb/s devices are backwards compatible with SATA 1.5-Gb/s controllers.

## Targets

In a SATA domain, targets are limited to SATA hard drives, each with a single narrow-link port. SATA drives are available in non-hot plug 3.5-inch drives and hot plug 3.5-inch and 2.5-inch small form factor (SFF) drives. SFF SATA drives consume 70% less space and use considerably less power than the 3.5-inch drives. SFF drives allow systems to be designed to accommodate more drives per U (1.75 inches), thereby increasing the total system I/Os per second per U (IOPs/U).


## SATA/SAS interoperability

The Serial Attached SCSI (SAS) architecture enables system designs that deploy both SAS and SATA devices.<sup>1</sup> SAS supports the SATA Tunneling Protocol (STP), which allows SAS controllers to communicate with SATA devices through expanders. SAS and SATA devices share the same physical device connector, except for an extension within the notch on the SAS connector. This extension allows a SAS connector to accept SATA device connections, but it will not allow a SATA connector to accept SAS device connections.

The SAS interface allows the flexibility to install SAS drives, SATA drives, or a mix of both SAS and SATA drives in the same enclosure. SAS and SATA devices share the same physical device connector, except for an extension within the notch on the SAS connector (Figure 8). This extension allows a SAS connector to accept SATA device connections, but it will not allow a SATA connector to accept SAS device connections.

The use of SATA drives should be limited to low-workload, entry-level servers and multi-drive storage configurations, such as JBOD or RAID. SATA drives are not recommended in non-fault tolerant applications with high IOPs, mission-critical applications, or extreme environments with excessive vibration. SAS is the ideal solution for mission-critical enterprise storage applications that require higher reliability, performance, and scalability.

**Figure 8.** SAS and SATA device connectors



<sup>1</sup> For more information, refer to the technology brief, "Serial Attached SCSI technology," at [www.hp.com/servers/technology](http://www.hp.com/servers/technology).

# SATA topologies


The following internal and external topologies use a SATA controller only. However, the Serial ATA STP enables SAS controllers to communicate with SATA devices through expanders and, therefore, allows SATA drives to be deployed in a SAS domain. For information, refer to the technology brief “Serial Attached SCSI technology” at

<http://h18004.www1.hp.com/products/servers/technology/whitepapers/proliant-storage.html>.

## Internal

Figure 9 shows a topology that can be used for internal RAID systems incorporating SATA drives. Each drive has a point-to-point connection to the controller. The controller can support a maximum of six drives.

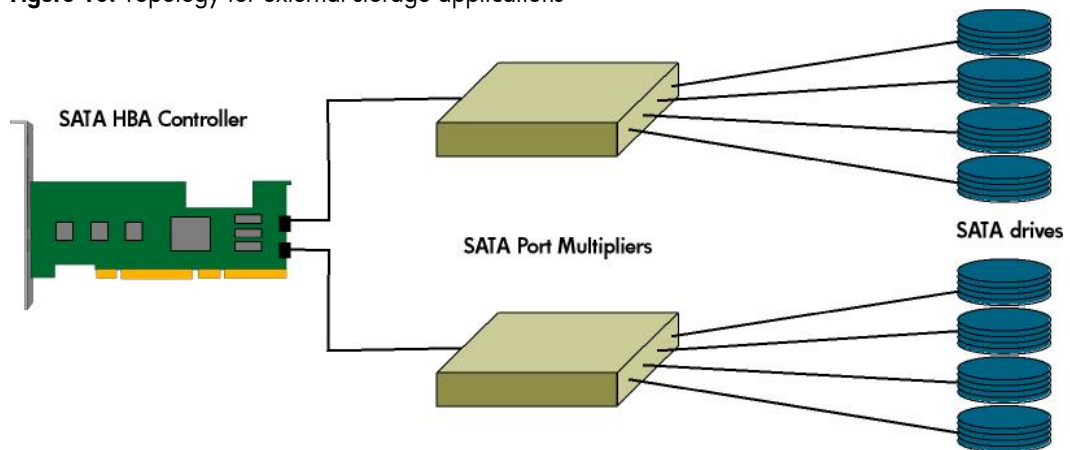
**Figure 9.** Topology for internal RAID array using six internal ports on a SATA RAID Controller



## External

Port multipliers enable a single SATA HBA port to connect up to four times as many drives with no performance degradation and, as a result, reduce cabling complexity (Figure 10).

**Figure 10.** Topology for external storage applications



## Conclusion

SATA overcomes the limitations of parallel ATA by providing thinner cabling, smaller connectors, lower voltage, and point-to-point signaling. Like its parallel predecessor, SATA will dominate the desktop market because it offers desktop reliability, functionality, and performance at a low cost. SATA technology is ideal for low-cost servers and non-mission critical server storage applications. SAS is the ideal solution for mission-critical enterprise storage applications that require higher reliability, performance, and scalability.

The SATA specification provides a consistent platform for the ongoing development of direct-attached and networked storage applications. The SATA technology roadmap provides scalable performance and system design improvements that will benefit end users and system vendors for years to come.

## For more information

For additional information, refer to the resources listed below.

Source	Hyperlink
SATA IO Web site	<a href="http://www.serialata.org">http://www.serialata.org</a>
Serial Attached SCSI technology brief	<a href="http://www.hp.com/servers/technology">www.hp.com/servers/technology</a>

## Call to action

Send comments about this paper to [TechCom@HP.com](mailto:TechCom@HP.com).

© 2005, 2007 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

TC070804TB, 08/2007

