

Serial ATA Simplified

Most people, when purchasing a new computer, pay careful attention to how fast the CPU is and how much memory it has. A third consideration is how large the hard drive is. Unfortunately, people rarely ask how **fast** the drive is. Just like a chain is only as strong as its weakest link, a computer is only as fast as its slowest component. These days, the hard drive is the slowest commonly used component in most computers. Parallel Advanced Technology Attachment / Integrated Drive Electronics (PATA / IDE) hard disks are the most common style to have in a home computer. IDE is also 20 years old, has crippling limitations, and will very soon be replaced by a new technology called Serial ATA (SATA).

The general rule with buying a hard disk for a new personal computer used to be 'Find a hard disk that you don't think you'll be able to fill, then buy one that's four times that size.' This rule would typically allow you to have one disk for the life of that new computer. Unfortunately, with the dramatic increase in software size, and broadband connections offering the ability to download fantastic amounts of data, the rule has become 'Buy the largest disk you can possibly afford, and buy another one next year'. The economy of this form of planned obsolescence is indisputably poor, but it has a further technical problem: drives are having to hold more data than the computer can efficiently read from them.

Parallel Breakdown

The Parallel IDE bus invented nearly two decades ago has gone through several renovations over the years to offer greater device compatibility (ATA 66/100/133), and greater speeds, but many of its original limitations remain.

- Two drive limit per channel (Master / Slave)
- Shared bandwidth
- Max. bus bandwidth of 133 Mbps
- Strict cable length limits
- Large connector wastes valuable drive real estate
- Wide ribbon cable blocks air flow inside computers
- High power requirements
- No growth road map

To illustrate how bus speed affects computer performance, we can estimate the time it would take to transfer all the data off a single drive connected to a standard PATA bus.

Disk Size	Bus Speed	Transfer Time.
8 GB	33 Mbps	~32 Min
20 GB	66 Mbps	~40 Min
60 GB	100 Mbps	~80 Min
250 GB	133 Mbps	~250 Min

These times are calculated by using the theoretical maximum bus bandwidth, and we rarely get the full capacity of a bus from one drive. Furthermore, if we're transferring to another device on the same

channel, these bus speeds would be cut in half, and the times would double.

To SATA and Beyond

Specifications for Serial ATA were introduced in August of 2001, and were developed specifically to address limitations in our current ATA standards. The Serial ATA 1.0 specification has a base bandwidth of 150Mbps and planned rates up to 10Gbps. Another speed advantage that SATA has over parallel technologies is its interconnect. SATA drives each connect directly to the controller, so bandwidth isn't shared between devices on a chain. This has the side effect of reducing overhead in device control commands and allows for added scalability. Instead of the number of drives on a controller being determined by the bus architecture or a block command set, the vendor can choose how many connectors to make available on their host adapters.

Small Bytes of Serial


The public stigma that serial data transfer is slow has slowed the widespread adoption of Serial ATA. Asking the general computer user where their computer's serial port is located will typically yield some finger pointing at a 9 pin connector on the back of the computer. Further questions relating to their computers parallel port will be answered with more finger pointing at the 25 pin connector their printer used to be plugged in to before they bought a USB printer. If you then asked this computer user to choose if parallel or serial ports were faster, they would probably indicate that parallel was faster.

In truth, computers contain a number of serial and parallel buses both inside and out. USB, FireWire (IEEE-1394), PS/2, RS-232 and even Ethernet use serial transfer protocols. PCI, AGP, IDE, Parallel (IEEE-1284) and (the big one) the Front-Side bus are all parallel interconnects.

As you may know, serial systems transfer one bit at every clock tick of their controller, and parallel systems transfer several bits at every clock. The number of bits a parallel bus transfers per clock is called it's 'width'. To increase the transfer speed of a bus, either the width or the frequency of the clock ticks has to increase. To give an example, the first SCSI systems were 8 bits in width, and Wide SCSI was 16 bit, doubling the width doubled the transfer speed. Fast SCSI which came out about the same time left the bus width at 8 bits, but doubled the clock frequency, also doubling the transfer speed.

Breathing Room

Another advantage that Serial ATA offers over traditional PATA is a more compact cable design. This becomes quite important as computers get smaller and smaller. Airflow and cooling are always critical factors in computer design. Compare the old PATA ribbon cable with a newer SATA cable, and the differences become clear:



The narrow design of the Serial ATA cable allows for more flexibility in running the cable, as well as improved airflow within the computer.

Infiltrating RAID

If SATA hard disks are going to compete against SCSI in the RAID arena, SATA must offer the same features as SCSI drives do in high-availability systems, but at a greatly reduced cost. The failure rate for ATA drives has always been greater than SCSI drives, so the Mean Time Between Failure (MTBF) becomes the major difference when calculating the total cost of ownership. SATA drives can be spun down and removed while the system is running, so a failed RAID drive could be replaced just like current SCSI disks.

Laptop Love

Connectors for SATA require only seven pins for data, and four for power. This makes it possible for 2.5" hard disks to further reduce in size, and the reduced power requirements (SATA is 0.5v where PATA is 5.0v) will extend battery life for portables, and reduce thermal radiation. Also, having the same connector on both laptop and desktop workstations will make pulling data off damaged laptops easier. The SATA connector can also be used for external drives.

I want to use SATA. What can I do?

Serial ATA hard disks and PCI controller cards have been available on the market for about a year, but their use has been limited to niche markets until recently. In the past few months several major computer vendors have started offering machines using SATA technologies, and the price difference between standard PATA drives and SATA drives has dropped significantly. If you buy a new computer with SATA controllers, you can still use your existing hard disk on those SATA controllers using adaptors. These can also be useful for moving data from existing IDE drives to SATA disks.

Support for SATA drives is still a little sketchy at the moment. Windows doesn't offer Mass Storage drivers for SATA by default, so you can't install a fresh copy of Windows onto an SATA drive without the SATA controller drivers on a floppy, and hitting F6 during the install. Linux drivers are available, but they're not as mature as existing SCSI or IDE drivers. Also, utilities such as hdparm haven't been fully updated for SATA just yet.

Glossary

Bus

A computer bus is a set of wires that transfers power or data between components. When referring to modern hard disk buses, we talk about Integrated Drive Electronics (IDE) and Small Computer System Interface (SCSI) buses. As SCSI systems are typically much more expensive, it's reserved for the few situations where it's speed and scalability features are required (servers, video processing, etc.). Our current form of IDE and most types of SCSI are parallel buses.

RAID

Redundant Array of Independent Disks. RAID comes in several forms, each with their own benefits. RAID 0 (also known as striping) is the simplest form of RAID, and simply distributes the usage of the drives for increased speed. It's not considered a true RAID, because if one disk fails, the data is lost, and the data on the other disks becomes hard to retrieve. RAID 5 requires a minimum of three disks, stores duplicate data on each of them, and reads data from them simultaneously. If 1MB of sequential data is drawn from a RAID 5 array of 4 disks, one quarter of the data is read from each disk. Conversely, if data needs to be written to the array, it must be written (at least partially) to all the disks. Unlike RAID 0, if one drive fails in a RAID 5 array, it can simply be replaced and the RAID array will automatically rebuild the data that's supposed to be on that disk. When RAID was first marketed, the 'I' in the acronym stood for 'Inexpensive'. However, RAID arrays have become anything but Inexpensive, and the term was changed.