

Topic — Data Centers



Sign In



Go Premium

Top Products

AI

More

system

Published December 19, 2000



Written By

jamesmcperson

Ever become confused about busses, bus speeds, and associated acronyms? In this Daily Drill Down, James McPherson gives an overview of current busses, including their definitions and specs. Use it as review, for learning, or to study for your A+ exam.

These days, pick up any electronic device and you'll find some acronyms on the box describing the various busses it supports. A computer will have a list of busses as long as your arm. In this Daily Drill Down, I'll discuss the different busses currently available and in use to help you understand exactly what your equipment is capable of.

What's a bus?

A *bus* is the path through which a device sends its data so that it can communicate with the CPU and/or other devices. For example, a PCI device, such as an audio card, will send its data through the PCI bus. Each device will have an access point to the bus using a particular kind of interface. The word *interface* refers not only to the physical port the devices plug into, but to the electrical operating parameters and the communication format as well. Typically, each bus has a uniquely shaped interface to prevent you from damaging your devices by plugging them into the wrong ports. PCs have three or more busses.

The motherboard's bus system has been compared to a mass transit system that carries data over many routes through the city (your motherboard) and uses different types of vehicles (fast and slow, and small and large) to carry it.

The differences between computer busses break down into these categories:

- Data width
- Cycle rate
- Device management
- Type

The data width and cycle rate are used to determine the bandwidth, or the total amount of data that the bus can transmit. An 8-bit bus (1-byte data width) that operates at a cycle rate of 1,000 MHz (1,000,000 times per second) can transfer 8 Mbps (1 MBps).

The device-management specification indicates the maximum number of supported devices and the difficulty of configuring them. There are two types of bus communications, serial and parallel. On a parallel bus, all devices have their own interface to the bus, which is the norm. Serial devices are tied together in, well, a series; the last one has to talk "through" the first one.

This can cause obvious performance problems. These busses typically are used in conditions where data throughput isn't critical.

Front Side Bus (FSB)

The Front Side Bus is the interface between the CPU and the motherboard, specifically the North Bridge/Memory Controller Hub. See below for details on the FSBs in use by Intel and AMD. For more information on this subject, see my Daily Drill Down, "[Motherboard chipsets—the good, the bad, and the ugly.](#)"

Intel GTL+ Front Side Bus

At the simplest level, Intel's GTL+ FSB provides a single connection to the North Bridge shared between all CPUs. In a dual-CPU system, this halves the available bandwidth and quarters it for a quad CPU board. Someone is sure to point out the fact that CPUs almost never need the full bandwidth of the bus. Quite true. Unfortunately, since the bus is in an all-or-nothing situation, the CPUs must take turns. (Hello Mr. Latency, care to have a seat and chat while we're waiting for the bus to arrive?) The problem is even worse in today's world of 800-MHz CPUs running on 133-MHz memory where even a single CPU has to wait up to six processor cycles for a data request. Imagine a quad CPU server being used for something other than show; if the application isn't smart enough to fill that L1 cache, or—heaven forbid—your L1 cache is too small to last the duration, you'll have processors sitting idle. So you might as well use fewer or slower CPUs and save some cash. This is why only Intel's Xeon processors, coming with up to eight times as much L1 cache as a Pentium III, are usable in systems with more than two processors.

AMD EV-6 Front Side Bus

The EV-6 FSB is more like a network switch than a bus, since each processor has a full connection to the North Bridge running at an effective 200 MHz; 50 percent faster than Intel's 133-MHz FSB. Between per-processor links to the North Bridge and that high-speed bus, the EV-6 is an excellent multiprocessor bus. Naturally, the total effective CPU bandwidth can't exceed the bandwidth available from the other interfaces combined, but unlike the GTL+ bus, one CPU could access peripherals on the PCI bus while another accesses memory.

Of course, the EV-6 bus is nothing new to the computing world since it was one of Alpha's weapons in the server wars. This is also academic until AMD releases the multiprocessor version of the 760 chipset in late December 2000.

Memory busses

The memory bus is the interface between the RAM and the motherboard. Because each variant requires a different type of controller, few motherboards support more than one type of memory. There have been many forms of memory that are now considered obsolete. The current types are discussed below.

DDR-SDRAM (Double Data Rate Synchronous Dynamic Random Access Memory)

This upcoming replacement for SDRAM is basically the same product, but it operates twice per clock cycle. Two grades are expected to be introduced initially: the 2x 100-MHz PC1600 (1.6 GBps) and the 2x 133-MHz PC2100 (2.1 GBps). DDR-SDRAM is only about 10 to 20 percent more expensive than traditional SDRAM and provides better performance than single-channel RDRAM. A 2x 200-MHz PC3200 (3.2 GBps) is currently under development that will provide the same performance level of dual-channel RDRAM while using only a single memory module.

RDRAM (Rambus Dynamic Random Access Memory)

Rambus is a proprietary memory architecture being touted by Intel. It has a serial memory format with a very narrow 16-bit interface but operates very fast at 800 MHz on a DDR-type 400-MHz bus, resulting in 1.6 GBps of bandwidth. A dual-channel RDRAM system is used on a rare few workstation systems: It has two RDRAM controllers for 3.2 GBps of bandwidth but requires the RDRAM to be installed in pairs.

RDRAM is several times more expensive than SDRAM and provides increased latency. It will shortly be challenged by DDR-SDRAM in the marketplace on both price and performance.

SDRAM (Synchronous Dynamic Random Access Memory)

SDRAM is the standard memory format for the majority of computers on the market. This 64-bit memory comes in three grades: PC66 (66 MHz or 528 MBps), PC100 (100 MHz or 800 MBps), and PC133 (133 MHz or 1.06 GBps). PC66 was used on early Intel Pentium IIs and all Intel Celeron PCs. PC100 is in use on the vast majority of Intel Pentium II and Pentium III processors. PC133 is the preferred memory of all AMD Athlon and Duron processors and the latest Pentium III systems.

VCM (Virtual Channel Memory)

This subset of SDRAM is a low-latency variant that provides increased performance. It runs at 133 MHz and has the same 1 GB of bandwidth as PC133 SDRAM but shaves off about 10 nanoseconds of latency from SDRAM's normal 40-nanosecond latency. It does this by using special "fast" registers that keep track of memory pages. These registers provide a fast link, or channel, to the memory used by an application. VCM actually works better for complex applications like games and databases that have memory spanning multiple memory banks.

VCM is supported on a number of Pentium II, Pentium III, and Athlon motherboards but is very difficult to acquire. The low production volumes have kept prices out of proportion to the performance compared to standard PC133. Motherboards that support VCM can use it or standard SDRAM.

High-speed I/O busses

Now let's take a look at the high-speed I/O busses.

AGP/Pro (Advanced Graphics Port)

This interface is a 32-bit system based on the PCI standard, revision 2.1. The initial version, 1x, operated at 66 MHz for 266 MBps with direct memory-access abilities that PCI didn't have. The 2x variant is a double-data-rate system that transmits data twice per clock cycle, for a functional frequency of 133 MHz (532 MBps).

The 4x again doubles the bandwidth up to 1,066 MBps and has additional memory-access features. The bandwidth of AGP 4x is greater than SDRAM's abilities, making most of the improvements of only limited use on systems not currently using Rambus RDRAM (1.6 GBps) or the upcoming double-data-rate DDR memory (2.1 GBps).

The Pro is a 4x variant that includes additional power leads to run today's high-transistor-count video cards. Standard AGP slots provide up to 25 watts of power, far less than AGP Pro's maximum available 110 watts.

The AGP port is typically a dark-colored port that resembles a PCI slot. It's set farther back from the edge of the motherboard than the PCI slots and is located near the power supply and processor on most boards.

EIDE (Enhanced Integrated Drive Electronics)

IDE hard drives include a device controller mounted on the drive, the result of a 1986 collaboration between Compaq and Western Digital to develop a cheap drive with good performance. They decided to limit the number of pins and the cable length as it was intended for lower-end systems that would not need a large number of internal devices.

Because each device has its own controller, only two devices can be on each chain to prevent excess interference. Modern IDE host adapters can operate two chains, each with a master and a slave. The master can interrupt the slave device at any time, making it inappropriate for the primary system drive or sensitive devices like CD-R, CD-RW, and tape drives to be slaves. The maximum of four devices per controller (two chains with two devices each) limits the number of devices an IDE system can handle.

The current generation of integrated IDE host adapters requires some management by the computer's processor, creating a load on the system. The various implementations of direct memory access (DMA) protocols helped transfer data to the computer's RAM from devices with less management by the processor. Transfer rates increased under the ATA-2 specification from 11.1 MBps to 16.66 MBps—a significant increase but still a bottleneck for PCs. Beginning with ATA-2 (also known as Fast ATA), the improved interface became known as EIDE, rather than IDE.

In 1996, the ATA/33 specification utilized newer DMA techniques to reach 33-MBps transfer rates. Also known as Ultra DMA or UDMA/33, it was completely backward-compatible with previous devices and became the standard for PC hard drives.

The year 1999 saw the introduction of the improved ATA/66 format. The 66-MBps system relies on a 40-pin cable with a similar connector to previous IDE formats, but it uses 80 conductors to ensure signaling. To operate at 66 MBps, only ATA/66 devices can exist on a channel. The controller can still operate earlier IDE devices, but the existence of non-ATA/66 devices forces a channel to drop down to ATA/33 speeds. ATA/66 has been widely adopted and is standard on many new computers.

Earlier this year, ATA/100 was introduced. Initial products were sometimes called ATA/66+. It continues the use of the 80-conductor cable and provides transfer rates up to 100 MBps. IDE devices are quite inexpensive, and due to the improved DMA functions, they do not impact system performance as much as in the past. IDE drives can be used in a RAID configuration with a four-drive limit imposed by the controller cards. A sufficiently complex controller could operate multiple IDE channels for more drives.

EISA (Extended Industry Standard Architecture)

The ISA bus was the first widespread PC interface, introduced by IBM on the AT computer as the system bus. It was at first an 8-bit bus that ran at 8 MHz (8 MB/second of bandwidth) that was quickly upgraded to a 16-bit bus (16 MBps). It utilizes interrupts, the digital equivalent of a wake-up call, to manage the devices. The interrupts (referred to as IRQs) are hardwired to the ports on the motherboard and cannot be reallocated. Communication occurs using I/O and memory addresses that can be reallocated among the ports. To function, each device must have the IRQ, I/O, and memory addresses configured properly at boot.

ISA was later extended to a 32-bit wide bus operating at 8 MHz (32 MBps) and renamed the Extended Industry Architecture (EISA) bus. ISA devices would work in an EISA motherboard to ensure compatibility. EISA added more I/O and memory access channels, allowed sharing IRQs, and enabled software to configure cards.

EISA devices plug directly onto the motherboard using the long, typically dark-colored slots located at the farthest edge of the board. ISA has been obsolete for several years and EISA is slowly dying out. Current motherboards have no more than one to three EISA slots. A large number of motherboards are completely devoid of EISA support.

FireWire/IEEE 1394

Developed by Apple and ratified by the IEEE (a subset of ANSI, the organization that also ratifies SCSI and IDE), FireWire is high-bandwidth, hot-swappable interface that supports up to 63 devices at a transfer rate of 50 MBps. It has not been widely accepted because it competes with SCSI, a highly established interface. Because it is a high-performance design for portable devices, it suffers from low initial sales volumes that prevent it from competing with SCSI on the merits of price the way IDE can.

PC-Card/PCMCIA (Personal Computer Memory Card International Association)

The PCMCIA group that defines this standard originally designated the devices as PCMCIA cards. Most people called them laptop cards, or those PC-something-something cards, which impelled the devices to be dubbed PC-Cards and the bus simply CardBus.

There are three types of PC-Cards, all of them slightly larger than a credit card. Type 1 is only 3.3 mm high and typically only suitable for memory devices. Type 2 cards, at 5 mm high, are large enough to support modems, network cards, and other interface devices. The large 10 mm Type 3 cards enable hard drives and other large storage devices.

The data throughput varies from 4 MBps for the slowest 16-bit data transfers to 132 MBps in the fastest 32-bit mode, as defined by the device.

PCI (Peripheral Component Interconnect)

PCI replaced EISA as the system bus of PC computers. It includes a Bridge chip that enables other types of processors to communicate with it. As a result, PCI is a standard bus on PCs and on Macintosh, Sun, and Alpha machines. The desktop variant of PCI is a 32-bit bus operating at 33 MHz (133 MBps). The PCI standard supports a 64-bit (266-MBps) variant. However, this format is currently in use only by workstation and server manufacturers like Sun or Compaq/DEC. The current plan is for PCI to skip the long-proposed 66-MHz PC and evolve into PCI X (eXtended), a 133-MHz (532-MBps) version with essentially the same features but triple the bandwidth.

PCI is intended to be Plug and Play capable, using software to configure the boards like EISA. It also supports a much larger number of I/O and memory addresses and uses better direct memory access than EISA. PCI is bus mastered, meaning that all devices can talk to each other without CPU intervention, which increases system performance.

PCI slots are white and much shorter than EISA slots. Most boards will have between three and five PCI slots.

SCSI (Small Computer System Interface)

A high-performance device interface, SCSI was introduced in the mid-1980s to compete with the now-obsolete ESDI device interface. SCSI was targeted for PCs, Apple Macintosh, UNIX workstations, and minicomputers but not mainframes, hence the "small" designation. SCSI was designed from the beginning to support not only hard drives but also scanners, optical drives, and other high-capacity/high-bandwidth devices that are not necessarily mounted inside the computer.

There are several generations of SCSI that in many cases have a "wide" 16-bit version and a "narrow" 8-bit version. The performance of SCSI ranges from the 5 MBps of SCSI-1 to the upcoming Ultra 320 (320-MBps) system. For more information, read TechProGuild's "[A SCSI](#)

primer: [Seven generations of high-performance hardware](#),” and “[SCSI, IDE, and FireWire specs at a glance](#).”

USB (Universal Serial Bus)

USB can handle up to 128 devices. The devices are hot-swappable, meaning that they can be added and removed while the computer is operating. Thus, this bus is very handy for portable devices. USB provides power to devices through the interface, enabling ultralight and convenient accessories free of power cables or cords. The power supply is finite, but a powered hub or other self-powered device will increase the possible number of bus-powered devices.

The total bandwidth of USB is 1.5 MBps; this bandwidth is shared between all the devices on the controller. Because it is a serial bus, USB devices are daisy chained. USB supports hubs that can connect multiple devices, which cuts down the number of devices between the controller and the end of the line, but since each hub counts as a device, this solution reduces the total that can be supported. Of course, with USB's low bandwidth, it isn't really feasible to utilize all 128 devices.

USB 2.0 has been in development for a while and is supposed to be a 60-MBps bus. This bandwidth is enough to operate a hard drive, 100-Mb Ethernet, or video-capture system with plenty of overhead left over for mice, keyboards, and other high-urgency devices.

Low-speed I/O busses

The following sections discuss the low-speed I/O busses.

Parallel port

This 25-pin port should be familiar to everyone with a printer or a Zip drive. It originally had a maximum transfer rate of 115 KBps, suitable for line printers and 1980s devices. Then came enhanced ECP/EPP control logic that boosted transfer rates to 3 MBps, a needed increase to support high-speed graphical printers.

The parallel port is technically an interface rather than a bus, but because of the recent proliferation of “pass-through” devices like Zip drives that allow serial-type bus activity, it is included in this section. Parallel is still the most widely supported high-bandwidth external device interface on the market, and it will likely be around for many years to come.

PS/2 port

IBM introduced the familiar, round PS/2 port to connect keyboards and mice. It is an ultralow bandwidth interface that will remain the standard until operating systems can reliably support USB keyboards and mice.

Serial port

This 15-pin, 115-KBps port is now used primarily for modems and personal digital assistants. It was once the primary port for mice but has been supplanted by the PS/2 port. Virtually all

computers have a serial port on the back. As with parallel ports, the serial port is technically an interface rather than a bus. However, because of the recent proliferation of “pass-through” devices like PDAs that allow serial-type bus activity, it is also included in this section.

Special motherboard connectors

Now let's look at other types of motherboard connectors.

ACR (Audio Communication Riser)

Risers are so named because they rise above the motherboard, either parallel or perpendicular to it. The ACR hardware standard is an attempt to replace the poorly received AMR (Audio Modem Riser) and CNR (Communication Network Riser) formats. The interface is a PCI port but with different pin-outs, and it is compatible with the older AMR system. It appears identical to a PCI slot and is typically placed at an angle to the other slots on a motherboard.

AMR (Audio Modem Riser)

This Intel-designed expansion port is intended for mass manufacturers to use to add either a modem or audio card that relies on the CPU for much of its processing power. It is a very small slot; about half the length of a PCI slot. AMR is being replaced by CNR.

CNR (Communication Network Riser)

CNR expansion ports are intended for use by mass manufacturers to slot into the motherboard either a low-cost network card, modem, or audio card that relies on the CPU for much of its processing power. The riser uses a very short slot; half as long as a PCI slot. CNR replaces AMR.

Conclusion

New developments in bus speeds free up system bottlenecks, letting devices keep up with faster CPUs. Keep in mind the information in this Daily Drill Down, and you can determine which busses suit your needs and which ones will soon become obsolete.

The authors and editors have taken care in preparation of the content contained herein but make no expressed or implied warranty of any kind and assume no responsibility for errors or omissions. No liability is assumed for any damages. Always have a verified backup before making any changes.

Subscribe to the Data Insider Newsletter

Learn the latest news and best practices about data science, big data analytics, artificial intelligence, data security, and more. Delivered Mondays and Thursdays

Email Address

By signing up to receive our newsletter, you agree to our [Terms of Use](#) and [Privacy Policy](#). You can unsubscribe at any time.

Subscribe

Share Article



jamesmcpherson

[See all of jamesmcpherson's content →](#)