Bus stop How USB devices communicate

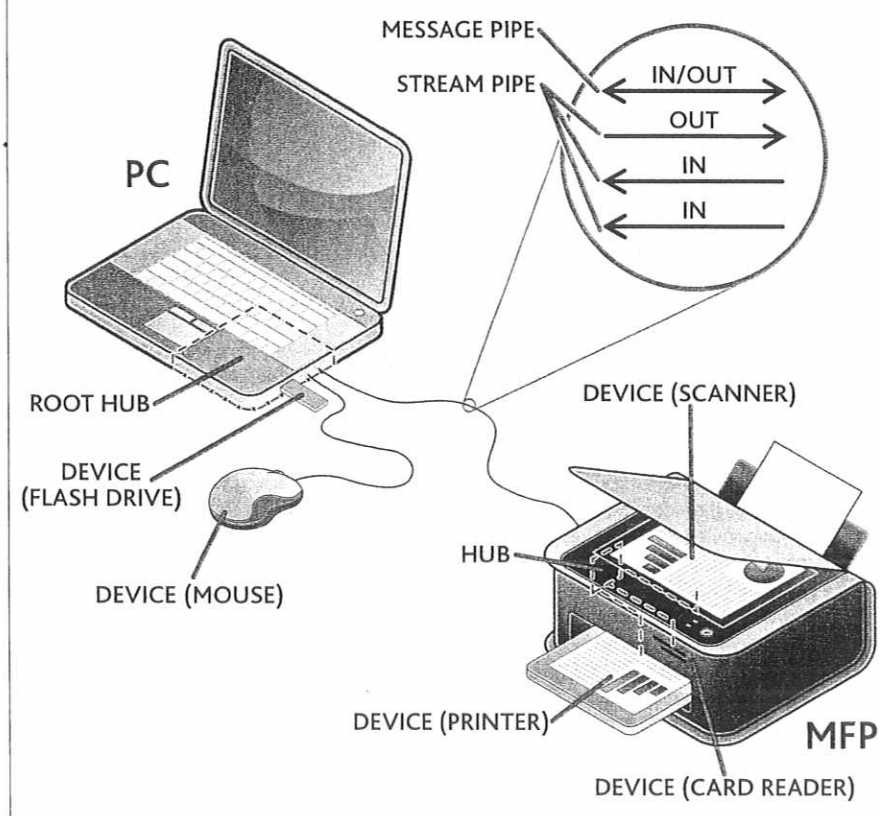

Up to 127 USB devices can connect to form an informal network, with a root hub – usually contained in a computer – at its centre. Other devices radiate out in a so-called 'tiered star' arrangement, where devices plug into the root hub or into other hubs that in turn are connected to the root hub. Compound devices, such as multifunction peripherals, may contain their own hub and several sub-devices.

Data packets are broadcast to every device, and all but the intended recipient ignore them. Data flows along logical 'pipes' between the host and the device with which it is communicating.

keyboards or flash drives, which are identified and can be supported with the relevant driver. However, a more complex multifunction peripheral (MFP), for example, may contain several logical devices – in this case, a printer, scanner and perhaps a fax modem or memory card reader. So-called 'compound' devices contain their own USB hub into which each of their sub-devices connects. While the peripheral itself has only a single USB socket and requires just one connection to the host computer, its hub and each of its sub-devices is given a separate address.

BROADCAST NEWS

Once a USB controller has assigned each of its connected devices an address, it's able to exchange data with them. Information is sent in packets, as it is over a TCP/IP network, but with USB 1.1 and 2.0 they are not routed. Instead, they are broadcast across the USB network, and are simply ignored by the devices to which they are not addressed. As the host controls all the data traffic, collisions – when two devices try to send data at once – are highly unlikely. The information sent to, or returned by, a device is said to travel along logical 'pipes', although these are not physical connections – every device can have up to 16 inbound and 16 outbound pipes, each of which channels data along a single physical connection to the root hub.

The host controller sends instructions to each connected device and receives responses over a bi-directional 'message' pipe, but most data is transferred through uni-directional 'stream' pipes. These support three types of transfer. Interrupt transfers cater for peripherals such as mice and keyboards, which have comparatively small amounts of time-sensitive data to send. Isochronous transfers guarantee a certain speed for applications such as streaming video, while bulk transfers use any remaining bandwidth for large-scale data transfers such as file copying – but offer no quality of service guarantee.

THE MAGIC NUMBER

USB 3.0's headline feature is a much faster maximum data rate – a SuperSpeed mode of up to 4.8Gbit/s. Achieving this hasn't been as simple as just increasing the clock rate tenfold, however. USB 3.0 brings the most wide-ranging improvements to the standard since USB first appeared, including redesigned connectors. Like USB 2.0 cables, USB 3.0 cables have two wires for power and a further twisted pair to carry data, but they also contain two more pairs of data wires, making a total of eight. The two new pairs are dedicated to SuperSpeed communication and are shielded to minimise interference and permit the highest possible data rate. Crucially, one of the new pairs carries data from the host to the device, while the second pair carries data in the opposite direction – an arrangement that permits transfer in both directions at once, known as full-duplex.

This provides a significant speed advantage over USB 2.0's half-duplex, but there are further changes to the way data travels between USB 3.0 devices. Although a host still controls the communications, devices now tell the host when they have data to send rather than waiting for it to ask, which saves bandwidth previously wasted on polling. Another highly significant change is that USB 3.0 creates logical 'communications pipes' between the root hub and the attached devices. These are used to route data packets, which means they're no longer broadcast across the entire USB network, saving bandwidth.

It's perhaps not surprising, given the incompatibility problems USB was developed to overcome, that a lot of attention has been paid to making the new standard compatible with USB 2.0 devices. The new connectors and ports are as near to the old designs as possible, enabling any combination of new and old devices and cables to be used. The only exception is that a USB 2.0 cable must be used to connect a USB 2.0 device to a USB 3.0 hub. All USB 3.0 controllers also contain a USB 2.0 controller for compatibility, and USB 3.0 communications are carried over different wires to USB 2.0. This means that connecting USB 2.0 and 3.0 devices to the same hub won't result in the 3.0 devices slowing down to accommodate the slower 2.0 devices.

Changes have also been made to the way USB supplies electricity to unpowered devices, and the way devices enter and leave low-power modes. The power available from a USB hub has always been divided into 'unit loads', only one of which may be demanded by a device that is physically connected but yet to be configured. USB 3.0 raises the unit load to 150 milliamps (mA), up from USB 2.0's 100mA. Once a device is configured, it may demand up to six unit loads (five in USB 2.0), allowing it to draw a maximum of 900mA compared with a 500mA maximum under the earlier standard. Meanwhile, because devices needn't remain alert in order to respond to regular polls from the host controller, they no longer depend on the host for permission to enter a low-power mode.

FIGHT FIRE WITH FIRE

For much of its history, USB has been compared unfavourably with FireWire, a competing but generally less popular serial interface introduced by Apple during the 1990s. FireWire initially had a transfer rate of 400Mbit/s, rising to 800Mbit/s – now, in a proposed new version, it reaches speeds of up to 3.2Gbit/s. FireWire retains several advantages over USB – chiefly that it can supply greater electrical power, and devices can communicate directly between each other rather than having to go through the host. For the first time, however, USB will match or even exceed FireWire's data transfer speeds. With its massive boost in speed and sophistication, USB 3.0 seems likely to continue to dominate.

FURTHER READING

The USB Implementers Forum homepage
www.usb.org/home

USB 3.0 FAQ
www.tinyurl.com/USB3FAQ

HELHa - HELHO Site Don Bosco (Tournai)

1e Informatique

EXAMEN D'ANGLAIS

9 juin 2010

(en annexe : liste de vocabulaire)

HOW IT WORKS USB

Before USB arrived, we had to make do with a ludicrous array of conflicting ports. Simon Handby explains how the humble USB port made our lives so much easier

Computers are powerful tools, but just as you need to install software to add capabilities that are not covered by Windows, you have to add extra hardware to tackle the tasks your PC can't handle alone.

Computer Shopper devotes entire sections to printers, storage devices, cameras, audio equipment and other peripherals that can make your PC more useful. While our reviews point out any problems you might encounter when installing them, their interfaces rarely get a mention. The main reason is that most peripherals connect through the Universal Serial Bus (USB) which, 14 years after its introduction, has proved reliable enough to forget about.

Today, no mainstream computer is sold without at least a couple of USB ports, and devices that use the new, faster third generation of the standard are just beginning to appear. The USB Implementers Forum (USB-IF), which maintains the standard, says it has become the most successful computing interface ever, with more than eight billion USB devices believed to have been sold. This month we look at why it was designed, how it works and how the latest version improves it even further.

PORTS IN A STORM

The connectors you find dotted around a computer haven't changed much in the past few years, but the situation was completely different in the late 1970s, when computers first began to find their way into our homes in significant numbers. IBM had yet to introduce its first PC – the basis for the PC we still use almost three decades later – and competing manufacturers were turning out a range of largely incompatible hardware. Even then, however, there was a degree of standardisation. For example, many computers had an RS-232 serial port so that they could connect to existing equipment, as that standard had been in widespread use since the 1960s.

When IBM introduced the first PC in 1981, it didn't improve the situation much. The standard model only had ports for an external cassette drive and a large, DIN-type keyboard plug. The later XT (extended technology) models added an onboard serial port, but it wasn't until the PS/2 interface, introduced in 1987, took off that PCs came with what might be considered a standard set of connectors. These usually included an RS-232 serial port, a DB-25 parallel port for printers and SCSI controllers, a D-sub VGA connector and a mini-DIN port for the keyboard and mouse. Although the latter are still generally known as PS/2 connectors, many third-party mice used a serial connection for several years after the PS/2 appeared.

For more than 10 years, these connectors were the standards through which PCs communicated with their peripherals. While the specialised D-sub port, which makes an analogue connection to a monitor, is still fairly common, the others are increasingly regarded as outmoded legacy connectors.

THE DO-IT-ALL WONDER PORT

USB was designed by a coalition of companies including Intel, IBM and Microsoft to address the shortcomings of the earlier interfaces used to connect peripherals. Chief among these was the fact that serial, parallel and PS/2 devices weren't interchangeable. Typically, a single serial or parallel port would support only one type of device. By having to provide connections for each, PC makers were able to support only a limited number of devices without festooning the back of a PC with ports.

Earlier connector designs were also imperfect, with fragile pin contacts and retaining screws that could work loose from their sockets. The contacts in PS/2 mouse and keyboard connectors were also easily bent, and until colour coding was introduced it was easy to swap the keyboard and mouse plugs around by mistake.

To succeed, a new standard would need to address these shortcomings, replacing multiple connector types with one family, but also adding new features such as faster data speeds and hot plugging, and even providing power for certain devices.

USB 1.0 achieved all this when it first appeared in 1996. It had a maximum Full-Speed data rate of 12Mbit/s, a screwless connector that was hard to pull out by mistake but could also be hot-plugged, and was able to supply a small amount of power. However, it was buggy and poorly supported by the operating systems of the time. Not until USB 1.1 arrived in 1998 did the standard began to gain favour. USB 2.0, which launched in April 2000, added a Hi-Speed data mode supporting a theoretical maximum rate of 480Mbit/s.

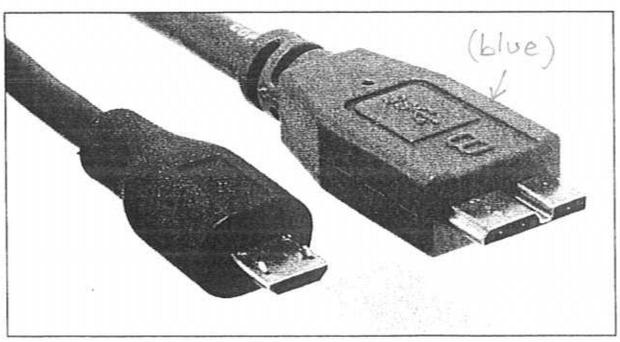

▲ The microUSB3 connector (right) is larger than the microUSB2 connector (left), but microUSB2 cables will still fit in microUSB3 ports

PARALLEL BAR

As its name suggests, USB is a serial standard that sends a single bit of data at a time. In contrast, parallel standards send several bits at once through multiple channels. To compete with the speed of parallel interfaces, serial connections must send data bits at a faster clock rate, which is relatively easy to achieve with modern data controllers. With this speed difference addressed, serial communications have multiple advantages over parallel. Only a single pair of wires is necessary to form the data-carrying circuit, so the data cables can be comparatively thin and flexible. What's more, the absence of parallel data channels within the same cable eliminates cross-talk, where one channel's signal interferes with another, allowing serial cables to be longer.

In fact, a USB cable contains two pairs of wires – one for data and one to provide power to devices that need it. As is common in wired communication, the pairs of wires that form a single channel are twisted together. This 'twisted pair' arrangement helps to ensure that both wires in the pair are equally exposed to any interference, such as you might get from a neighbouring power cable. Ideally, the interference added to one wire in the pair is also added to the second wire, cancelling itself out. Although there are several USB plug designs, all but the very cheapest share a robust metal connector that protects and electrically shields the recessed contacts.

Like the standards it superseded, USB is designed for connecting peripherals to a host computer system, so its communications are organised around a central controller in the host computer. This host contains the 'root hub', a logical device that acts as the central point of an informal network of connected hubs and devices (see the 'Bus stop' diagram opposite). The host directs all the communications among connected USB devices and hubs, sending out data as necessary and polling each in turn to find out which other devices have data to send.

When a new device is connected to the system, the controller detects its presence and instructs it to reset, during which time the controller establishes its speed. The device is given a seven-bit address that allows the host to identify it among all the connected items. This allows a single USB controller to support a maximum of 127 connected hubs or devices.

Each physical connection from the host controller goes either to a device or another hub. Most devices are simple peripherals, such as mice