



USB 3.0 –

Background Factors, New Features, and Applicability as a Camera Interface

White Paper

© 2009 IDS Imaging Development Systems GmbH. All rights reserved.



IDS Imaging Development Systems GmbH
Dimbacher Str. 6-8
74182 Obersulm, Germany
Phone: +49-(0)7134-96196-0
Fax: +49-(0)7134-96196-99

sales@ids-imaging.com
www.ids-imaging.com

Contents

1	Introduction	2
2	Development	3
2.1	State of the Art	3
2.2	Development Objectives and New Features	3
3	Technology	4
3.1	Backward Compatibility	4
3.2	Power Management	5
3.2.1	High-Power Mode	5
3.2.2	New Standby Modes	5
3.3	Protocol	5
3.3.1	Polling and Transmission Process	5
3.3.2	PC Host Controller	6
3.4	Hardware – Cabling and Hubs	6
3.4.1	Cables	6
3.4.2	Connectors	7
3.4.3	Hubs	9
4	Application as a Camera Interface	10
5	Sources and Links to other Websites	12

1 Introduction

Among today's digital PC interfaces, USB is by far the most widely used. According to current sales statistics, more than 6 billion devices with a USB connection have been sold to date. Thanks to the integrated power supply via the USB cable, the interface is especially well suited to small devices without a power supply system of their own.

Version 1.0 of the USB standard was first presented in 1996, followed soon afterwards by the slightly revised version 1.1. Novel features such as hot-plugging (i.e. removing and connecting a device while the PC is running), an integrated 5 V power supply and easy extensibility using USB hubs made the new technology stand out from the existing PC interfaces. At the same time, these characteristics fostered the orientation of USB towards peripheral devices such as computer mice, keyboards, and printers – “Full-Speed USB” with its maximum bandwidth of 12 Mbps (about 1.5 MB/sec) was actually not designed for connecting devices with high data transfer volumes.

Released in 2001, version 2.0 of the standard was termed “Hi-Speed USB.” As a dramatic improvement, it featured an increase of the maximum bandwidth to 480 Mbps. In practice, a transfer rate of 40 MB/sec was feasible, and thus, USB 2.0 was widely adopted for external mass storage devices, scanners, and cameras.

Eight years after the release of USB 2.0, current PCs equipped with high-speed interfaces, such as Gigabit Ethernet, are capable of providing bandwidths that are multiples of the performance USB 2.0 can deliver. External hard disks which can store terabytes of data or the boosting resolution available from cameras – and thus the amount of data created – have increased the demand for a still faster version of USB.

In November 2008, the specification for “SuperSpeed USB” was presented and marked an improved and, above all, considerably faster version of the popular interface. New features on the protocol level as well as on the hardware side promise to eliminate some of the limitations imposed by version 2.0.

This article presents the main characteristics of USB 3.0 as well as its benefits when used as a camera interface for image processing.

2 Development

2.1 State of the Art

On November 12, 2008, the “USB Implementers Forum” (abbreviated: USB-IF), which groups companies such as HP, Intel, and Microsoft, presented the first revision of the new specification. Concurrently with the finalization of the standard, the development of USB 3.0 chipsets started, and the first chips are expected to be available by mid-2009. However, at least one more year might pass until PCs and peripheral devices equipped with the new interface will have reached the end consumer market.

Working systems with SuperSpeed USB – although still in its evaluation phase – have already been presented at trade shows. FPGA-based host controllers have already reached data transfer rates of about 400 MB/sec.

2.2 Development Objectives and New Features

The main objectives in the development of the new standard were:

- Increase the data rate by a factor of about ten to reach 5 Gbps
- Maintain and support the existing USB infrastructure
- Optimize power management of connected devices
- Omit the polling procedure used in the previous USB protocol: connected peripheral devices were not allowed to send data on their own but had to wait for a host controller request.

These points mark the fundamental limitations of Hi-Speed USB which are being revised and overcome.

The following table shows the differences between USB 2.0 and 3.0:

	USB 2.0 “Hi-Speed”	USB 3.0 “SuperSpeed”
Bandwidth (usable):	480 Mbps (about 40 MB/sec)	5 Gbps (about 400 MB/sec)
Maximum power supply:	500 mA @ 5 V (2.5 W)	900 mA @ 5 V (4.5 W)
Data transfer:	Polling/host-controlled	Asynchronous/device-controlled
Duplexing:	Half-duplex: Unidirectional transfer	Dual simplex: Separate wire pairs for sending/receiving
Cable length using standard cables*):	5 m	3 m

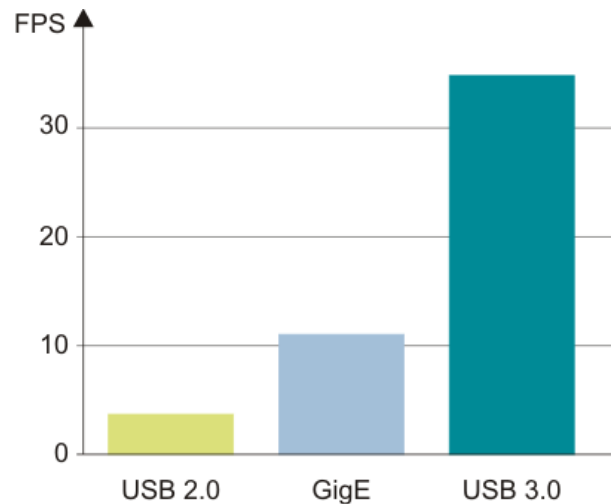
*) A maximum cable length is not specified. These values indicate the estimated maximum length that is possible using standard cables.

The following characteristics of the USB interface, however, have basically remained unchanged:

- Topology: Star topology on up to five levels with USB hubs used for distribution, up to 127 devices can be connected
- Lower speeds are supported (Low-Speed, Full-Speed, Hi-Speed)
- Transfer and packet types (control, bulk, interrupt, isochronous)

Unlike the USB 2.0 standard, which, upon its introduction, superseded the existing USB 1.1 standard, version 3.0 is not intended to replace the existing standard 2.0. This decision clarifies things for the consumers: With USB 2.0, devices may be declared as “USB 2.0 compatible” even if they do not support the high data rate offered by Hi-Speed USB – provided that they otherwise comply with the standard. In the case of USB 3.0, a device must actually support the SuperSpeed data rate to be marked as “USB 3.0 compatible.”

Figure 1: This bar chart shows the speed advantage of USB 3.0 taking the transfer of images with a resolution of 10 megapixels (8 bits color depth) as an example.



3 Technology

3.1 Backward Compatibility

As far as connectors are concerned, it is generally possible to plug USB 2.0 cables into USB 3.0 connectors – but the opposite is not always true (see also [3.4.2 Connectors](#)). Since SuperSpeed devices can be connected to legacy connectors, the standard stipulates that any SuperSpeed device must support at least one of the lower speeds (Full-Speed or Hi-Speed). The device must not necessarily be fully functional at the respective lower speed. It is sufficient that the SuperSpeed device as such is recognized and reports that it requires a USB 3.0 connection in order to operate properly.

If a SuperSpeed device supports operation in Hi-Speed mode, it is not allowed to draw more than 500 mA in that mode. When connected to a SuperSpeed connector, however, it can switch to High-Power mode (more than 500 mA).

3.2 Power Management

3.2.1 High-Power Mode

Instead of the previous 500 mA, a SuperSpeed USB connection now provides up to 900 mA per device. A device can request to switch to the so-called high-power mode. In this mode, even power-hungry devices such as external disk drives can be used without any additional power supply.

3.2.2 New Standby Modes

USB 2.0 allows the host to put idle devices into a standby mode with lower power requirements (up to 100 mA). A standalone power management on the device side is not implemented.

With SuperSpeed USB, the so-called “Link Level” power management was introduced where the devices may independently change into a different operating mode. Four modes are provided for and will be applied progressively: U0 (normal operating mode) to U3. The higher the mode, the more functions of the device are switched off, and accordingly, it takes longer to wake up. In standby mode, no data is sent to the host. If incoming communication from the host is received, the device wakes up and signals its Ready state.

3.3 Protocol

As with the previous USB standards, all devices connected to a SuperSpeed USB bus share the available bandwidth based on a host-controlled protocol. Besides higher data transfer speeds, SuperSpeed USB provides new protocol level mechanisms which enable the devices to transfer large data volumes more efficiently.

3.3.1 Polling and Transmission Process

The major limitation imposed by USB 2.0 connections is that they use polling mechanisms to enable communication between the host and its connected devices. Due to the master/slave hierarchy, the devices are not capable of sending data on their own. Instead, the host controller sends requests, so-called “IN” tokens, to the devices at regular intervals. If the devices have data to be sent to the host, this data will be appended to the IN token and returned. On the one hand, this host-controlled procedure causes latencies, on the other hand it increases the computation load on the PC side.

With the introduction of SuperSpeed USB, this polling procedure will no longer be used. The end devices can send data on their own without waiting for a host controller request. The controller can request data as before, but if a device is not ready to send, it can signal this state using a “Not Ready (NRDY)” packet. To signal that a data transmission is complete, the end device sends a handshake. SuperSpeed USB supports the sending of bursts, i.e. sequences made of multiple data packets. This process is always completed with a handshake.

High data volumes, as created by USB cameras, are transmitted in bulk mode. On the one hand, it enables the transmission of large data packets, and on the other hand, it provides an error correction feature which requests packets to be resent in case of transmission errors. In bulk mode, it is not possible to assign a fixed bandwidth to a device or connection, but the available bandwidth is leveraged. An isochronous mode is also available which allows allocation of a fixed bandwidth, but has no error correction functionality.

In order to further improve data transfer rates, SuperSpeed devices indicate their specific latency tolerance to the host so that it can optimize communication with the individual devices.

3.3.2 PC Host Controller

The USB 3.0 host controller is called xHCI. This stands for Extended Host Controller Interface. USB 2.0 host controllers are called EHCI, USB 1.1 controllers are referred to as OHCI or UHCI. Thus, mainboard chipsets are equipped with a host controller for each USB generation in their south bridge chip. The USB root hub forwards the data to the corresponding host controllers.

Since the beginning of 2009, the first beta drivers for xHCI under Linux and Windows Vista are available. As announced by Microsoft, Windows 7 will initially be released without USB 3.0 support¹.

3.4 Hardware – Cabling and Hubs

3.4.1 Cables

Besides the two wires used for power supply, Hi-Speed USB cables have a twisted wire pair for data transmission (designated as D+/D-). They are used to transmit the data differentially in order to eliminate interferences created by

¹ Source: Kidman, Angus: “USB in Windows 7: More reliable but no 3.0 speed boost”, <http://apcmag.com/Content.aspx?id=3177>, version: May 12, 2009

the addition of signals. Usually, the entire cable is additionally shielded. Data is transmitted over the wire pair in half-duplex mode, i.e. either in one direction or in the other.

For SuperSpeed USB, copper lines are still used, but separate wire pairs transmit the signals in the sending and receiving directions so that this is a full-duplex capable connection. In contrast to the construction of the Hi-Speed cable, the wire pairs of SuperSpeed USB must have an additional shield in order to increase the interference resistance of the high-frequency signals.

For reasons of backward compatibility, each USB 3.0 cable has lines for Hi-Speed as well as for SuperSpeed. The SuperSpeed lines are called SSTX+/- and SSRX+/- (for SuperSpeed Transmit and Receive).




Figure 2: Cross-section of a SuperSpeed USB cable

Just as with USB 2.0, the SuperSpeed specification does not indicate the cable length, but merely electric characteristics such as attenuation, crosstalk, or delay. Using common cable materials, wired cable lengths of about 3 m are expected to be reached. To produce longer SuperSpeed cables, high-end materials will have to be used.

3.4.2 Connectors

The USB 3.0 specification defines new connector types with added pins for the SuperSpeed lines:

- Standard-A: plug and socket
- Standard-B: plug and socket
- Powered-B: plug and socket
- Micro-B: plug and socket
- Micro-A: plug
- Micro-AB: socket

The following table shows the compatibility of USB 2.0 and 3.0 plugs and sockets:

Sockets		Compatible plugs	
USB 2.0	USB 3.0	USB 2.0	USB 3.0
Standard-A		Standard-A ✓	Standard-A ✓
Standard-B		Standard-B ✓	x
Micro-B		Micro-B ✓	x
Micro-AB		Micro-B ✓ Micro-A ✓	x
	Standard-A	Standard-A ✓	Standard-A ✓
	Powered-B	Standard-B ✓	Powered-B ✓ Standard-B ✓
	Standard-B	Standard-B ✓	Standard-B ✓
	Micro-B	Micro-B ✓	Micro-B ✓
	Micro-AB	Micro-B ✓ Micro-A ✓	Micro-B ✓ Micro-A ✓

USB 3.0 Standard-A

This plug type is designed to connect the USB cable to the host (usually the PC). The Standard-A plug is the only USB 3.0 plug that is mechanically identical with the USB 2.0 type. The additional four pins for the SuperSpeed lines are located in the previously unused rear part of the plug. USB 2.0 and 3.0 A-type plugs are compatible.

The SuperSpeed standard now features a new cable with Standard-A connectors on both sides that can be used to interconnect two host PCs. In this cable, the USB 2.0 lines are not assigned so that if this cable is inadvertently plugged into a Hi-Speed socket, no damage will occur.

For easy distinction, Standard-A plugs for USB 3.0 should be marked in blue.




Figure 3: SuperSpeed plugs, Standard-A and -B types

USB 3.0 Standard-B

Just like with USB 2.0, B-type connectors are used to connect larger stationary devices such as printers or external hard disks. The new Standard-B plug has a little bulge on top that accommodates the added SuperSpeed lines. It is not possible to insert this plug into a USB 2.0 socket.

USB 3.0 Powered-B

The Powered-B type is new in USB 3.0 and is mechanically identical with the Standard-B type. By using two additional pins, a higher current of up to 1 A can be supplied. The Powered-B connector can be used to operate external devices without separate power supply, e.g. when adding a WLAN adapter to a printer.

USB 3.0 Micro-B

This connector type is designed to connect small devices, e.g. cameras. It is not possible to insert this plug into a USB 2.0 socket. The USB 2.0 plug, however, can be plugged into a USB 3.0 socket.

USB 3.0 Micro-AB/ Micro-A

This connector type is mechanically identical with the Micro-B type. It has been designed exclusively to connect products equipped with the USB On-The-Go interface.




Figure 4: SuperSpeed Micro-B plug

3.4.3 Hubs

Hubs for USB 3.0 internally consist of two hubs, one of which is used for SuperSpeed connections, and the other one for USB 2.0 (i.e. Lo-Speed, Hi-Speed, and Full-Speed) connections. This split-up, however, is made internally so that it is not reflected by different ports on the outside. Typically, a hub would have one or more SuperSpeed Standard-A ports as well as one Standard-B port which are distributed automatically to the subordinate hubs. Hubs for USB 2.0 and 3.0 can be used simultaneously.

Unlike USB 2.0 hubs, which send incoming data to all outgoing ports, SuperSpeed hubs offer a functionality for routing data to the downstream ports (i.e. towards the device). The data packets include information about the port to which the corresponding device is connected and can therefore be routed to the correct port. Conversely, SuperSpeed hubs collect packets sent in the upstream direction in order to forward them collectively.




Figure 5: Connection layout of USB 2.0 and 3.0 hubs

4 Application as a Camera Interface

With its version 2.0, USB already combines many of the advantages that make this interface highly suitable for use in industrial image processing: Sufficient bandwidth to simultaneously transmit live images even from multiple cameras, the fact that it is widely used in IPCs and embedded systems, easy connectivity without the need for external power sources, and chipsets which allow building compact and cost-effective cameras. With these characteristics, the USB 2.0 interface currently is, according to a market survey², the top choice for industrial image processing users, even though Gigabit Ethernet, with its higher bandwidth, is catching up.

² Source: INSPECT et. al., "Der Schleier gelüftet. Umfrage zu aktuellen Trends in der industriellen Kameratechnologie" (Technology unveiled. Survey on the current trends in industrial camera technology), published in: INSPECT 1-2/2009, p. 6

While the bandwidth of this interface, which is celebrating its eighth birthday, has been sufficient so far, there is a huge demand for faster data transmissions now. The current generation of CMOS sensors delivers images with a resolution of 10 megapixels – and even with a color depth of 8 bits per pixel, which is common for USB cameras, it is hardly possible to transfer four full frames per second. In addition, many sensors can also output data with a color depth of 10 or 12 bits which virtually cuts the possible frame rate in half. Further trends include the interpolation of color data in the camera or even simultaneous output of color and monochrome data. Features that have long since been available in powerful GigE cameras were not feasible in the USB world with its bandwidth of about 40 MB/sec.

The unproblematic and fast USB 3.0 connectivity is also particularly useful e.g. for embedded applications where a high camera resolution is required, but a GigE camera is not suitable due to its external power source requirements or due to the lack of a GigE interface. The next generation of the universal bus will be ideally suited for laboratory or medical systems as well, where the strict requirements of industrial environments with respect to cable lengths do not apply.

For camera manufacturers, SuperSpeed USB is not only interesting because of the higher transmission speed – the revised and extended USB protocol allows an even more efficient, resource-friendly data transmission. By omitting polling and by supporting the future use of DMA (Direct Memory Access) USB controllers, it should be possible to retrieve image data from USB cameras virtually without requiring any CPU load.

For about two years, it has become apparent that new PCs of any kind are mainly equipped with USB 2.0 and GigE as digital interfaces. The majority of all image processing applications can operate on these two interfaces even today, while the successor standards USB 3.0 and also 10 GigE open up new possibilities for the application of industrial cameras.

With their uEye[®] camera series, IDS Imaging Development Systems GmbH offer a series of modern USB and GigE cameras for image processing in industrial and non-industrial environments. CMOS and CCD sensors produced by renowned manufacturers provide resolutions from VGA to 10 megapixels and frame rates up to 100 full frames per second. Using the extensive uEye[®] SDK software package that is available for free, all cameras, no matter whether they are equipped with a USB or GigE interface, can be integrated with image processing libraries or custom applications. For further information about IDS and the uEye[®] cameras, please visit www.ids-imaging.de.

5 Sources and Links to other Websites

- Official SuperSpeed USB specification (<http://www.usb.org/developers/docs/>)
- USB 3.0 Developers FAQ (<http://www.lvr.com/usb3faq.htm>)
- SuperSpeed USB 3.0 FAQ (<http://www.everythingusb.com/superspeed-usb.html>)
- Sharp, Sarah: “USB 3.0 and Linux”, <http://sarah.thesharps.us/2008-12-07-13-35.cherry>, version: May 12, 2009
- Kidman, Angus: “USB in Windows 7: More reliable but no 3.0 speed boost”, <http://apcmag.com/Content.aspx?id=3177>, version: May 12, 2009