Is Your USB Design Suffering From Bandwidth Overkill?

By (Brian Ellis, Product Marketing Engineer, Cypress Semiconductor Corp.)

Executive Summary

In this fast-paced world we live in, time is money so every second counts; and in the high tech industry, a second can mean success or failure. Designing your USB system to be as efficient and seamless as possible requires a USB hub controller that can maximize performance. USB hub controllers with multiple transaction translators (TT) have historically been synonymous with higher performance, and with higher performance comes higher cost. But do multi-TT USB hub controllers really have a performance benefit over single-TT USB hub controllers in today's designs? A little research may save you from paying for bandwidth overkill.

USB Hub Controllers 101: The Basics

To understand what to look for in a USB hub controller, let's first understand the basic operation of a hub. At a very high level, a USB hub can be thought of as a funnel; you take multiple inputs and "funnel" them into a single output. The inputs are received through the downstream ports and "funneled" through one upstream port connected to the USB host. A USB bus operates at a single speed defined by the USB host it is connected to, so if the upstream port of a USB hub is connected to a high-speed USB host all traffic passing through the upstream port must be high-speed traffic. This poses a fundamental problem, since USB peripherals operate at different speeds: low-speed (LS), full-speed (FS) and high-speed (HS). Therefore, the hub must standardize the speed of all traffic passing through the upstream port.

All data passing through the upstream port must first pass through one of two channels: the hub repeater of the transaction translator. For peripherals already operating at the same speed as the USB host, no "speed standardization" needs to occur, so the data is simply passed upstream. This is done through the hub repeater, which acts like a pass through (hence the name "repeater") for the data. In the case where the attached peripherals are not operating at the same speed as the USB host, the traffic speed is standardized using the transaction translator.

A transaction translator (TT) can be thought of as a train station, where LS & FS data catch a ride on a HS train. Sticking with this analogy, as LS & FS data enter the "station" the data waits while the HS train fills. Once the HS train is full of LS & FS data, it departs from the station at HS. Stepping away from the analogy, as the LS & FS data enter the TT it is "packetized" and buffered until a HS data packet is full of LS & FS data and sent through the upstream port at HS.

Let's dig into this in slightly more technical terms to understand what is really happening. A TT consists of an upstream HS handler and a downstream LS/FS handler, which are connected via data buffers. When a transaction begins, the HS handler receives a start-split command to initiate the buffers and begin the translation process. The buffers then hold LS/FS transactions while they are being packetized into HS packets ready for transfer over the HS upstream bus. Once a packet is ready for transmission, a complete-split command is issued and the data is passed through the HS handler to the USB host.

The TT uses split transactions in order to efficiently handle HS traffic through the hub. Using split transactions allows the USB host to start a LS/FS translation while continuing with HS data transfers through the repeater. This ensures there is no wasted time waiting for data translations through the TT, achieving optimal hub performance.

Figure 1. Single-TT vs Multi-TT: More Is Not Always Better

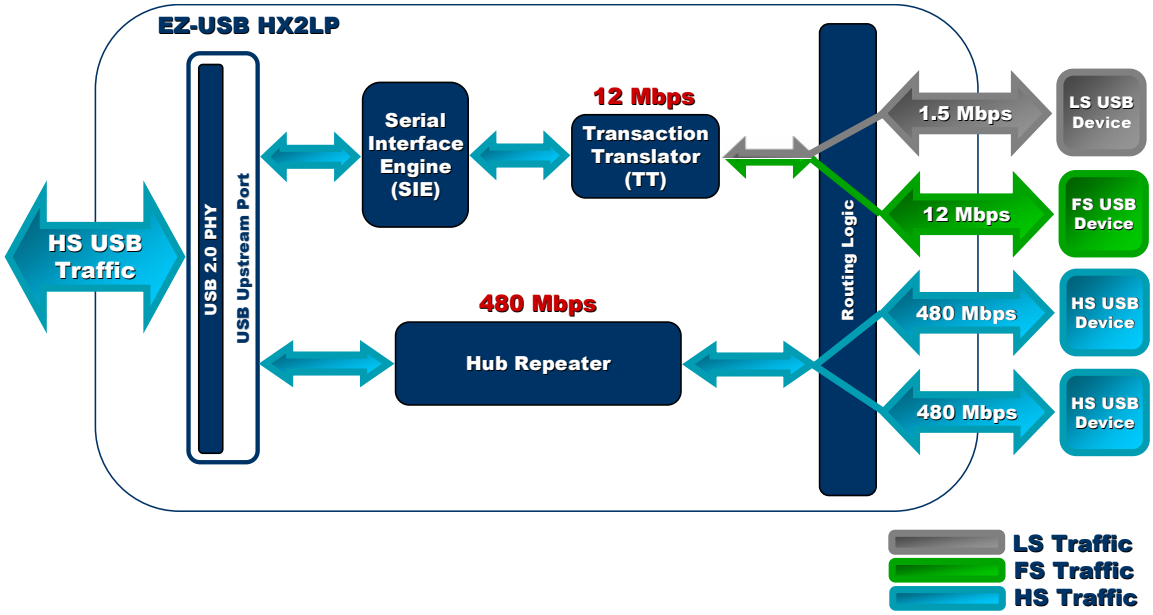

Each TT can handle a maximum bandwidth of 12Mbps. In a single-TT application, that 12Mbps of bandwidth is shared across all active ports. So if LS/FS peripherals occupied all four ports of a USB 2.0 4-port hub, the traffic for each peripheral would be fighting for their fair share of that 12Mbps of space. And in certain conditions a device may not even be able to fully enumerate because the bandwidth required is not available. Multi-TT hubs eliminate this problem by providing a dedicated TT for each port, allowing LS/FS peripherals their maximum required bandwidth without needing to share.

So multi-TT is better than single-TT, right? WRONG.

Before we bite at the marketing bait the multi-TT hub vendors have dangled in front of us, let's take a deeper look into these usage models.

Let's remember that the maximum bandwidth consumed by a LS USB device is 1.5Mbps. So even if you were using seven LS USB devices in a 7-port single-TT hub the maximum bandwidth you would require is 10.5Mbps, which is still below the 12Mbps available in the TT. The multi-TT hub will provide no added benefit over a single-TT hub in a design using LS & HS peripherals.

That leaves FS USB as the bandwidth hog. But before we write off FS USB altogether, it's time for a lesson on USB basics. There are four types of USB transfers: control, interrupt, bulk and isochronous.

> **Control**: Control transfers are used by all USB devices to prepare the device for use once it is plugged into a system. These transfers are very brief and take nearly zero bandwidth.
>
> **Interrupt**: Interrupt transfers are used for short bursts of data. For example, your mouse is only passing data when you are moving it, and that data "interrupts" the PC to let the host know there is movement. Interrupt transfers are typically used for very low-bandwidth devices that are not used for long periods at a time.
>
> **Bulk**: Bulk transfers are used to pass large amounts of data quickly. For example, when you are copying files onto your external hard drive (HDD) you are using bulk transfers. Since speed is the goal with bulk transfers, they are dynamic and will consume as much bandwidth as possible. The more bandwidth available, the larger the pipe for passing data and the quicker the transfer completes.
>
> **Isochronous**: Isochronous transfers are used when timing is crucial. For example, when you are using a USB webcam to record your latest YouTube music video on your PC, you want to guarantee the images are received on time in order to avoid skips or glitches in the video. Since timing is so important with isochronous transfers, devices

that use these transfers require a set amount of bandwidth to operate and that set amount cannot be changed. If that set amount of bandwidth is unavailable, the device will not operate.

Of the four USB transfer types, only two of them are bandwidth hungry and have the potential of causing bandwidth resource issues in a design: bulk and isochronous.

As mentioned earlier, bulk transfers are mainly used for mass storage applications where large amounts of data are being transferred at a time, like external HDDs. Isochronous transfers are widely used for webcams to guarantee real-time picture quality. But as the research firm Instat indicates in their March 2006 USB market report, FS USB external HDDs and webcams have less than 15% market share and will be completely non-existent by the end of 2007.

Table 16. USB Penetration of External Hard Disk Drive Market Forecast (Units in Thousands)

	2003	2004	2005	2006	2007	2008	2009	2010	CAGR 06-10
External Hard Drives	2,500	5,100	6,600	7,800	9,100	10,600	12,100	13,600	15.6%
% w/ USB Low or Full Speed	5.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	
Total w/ USB Low or Full Speed	125	0	0	0	0	0	0	0	--
% w/ USB High Speed	83.0%	96.0%	97.0%	97.0%	98.0%	98.0%	98.0%	98.0%	
External Hard Drives	2,075	4,896	6,402	7,566	8,918	10,388	11,858	13,328	15.8%
% with Wireless USB	0.0%	0.0%	0.0%	0.0%	1.2%	4.0%	10.0%	15.0%	
Total with Wireless USB	0	0	0	0	109	424	1,210	2,040	--

Source: In-Stat, 3/06

Table 15. USB Penetration of Web Camera Market Forecast (Units in Thousands)

	2003	2004	2005	2006	2007	2008	2009	2010	CAGR 05-10
Web Cameras	9,200	9,500	9,700	10,000	10,400	10,900	11,000	10,500	1.6%
% w/ USB Low or Full Speed	88.0%	77.0%	55.0%	35.0%	15.0%	0.0%	0.0%	0.0%	
Total w/ USB Low or Full Speed	8,096	7,315	5,335	3,500	1,560	0	0	0	-100.0%
% w/ USB High Speed	5.0%	11.0%	35.0%	60.0%	81.0%	96.0%	97.0%	97.0%	
Total w/ USB High Speed	460	1,045	3,395	6,000	8,424	10,464	10,670	10,185	24.6%

Source: In-Stat, 3/06

Let's review. The only scenario where we will see a performance benefit from a multi-TT hub is when using multiple high-bandwidth isochronous or bulk transfer FS USB devices, such as FS external mass storage devices or FS webcams. But the market data clearly shows a complete migration to HS USB for both external mass storage and webcams by the end of 2007. So in today's design, single-TT hubs and multi-TT hubs are performance equivalent.

When it comes to silicon, die size equals cost. The larger the die, the higher the cost. Comparing the die size of a multi-TT hub controller versus a single-TT hub controller is not as simple as multiplying by the number of additional TTs. Each TT has associated support logic, adding to the overall die size. In fact, when comparing die sizes for a multi-TT and single-TT hub controller based on the same process technology, the multi-TT die showed a 34% increase over the single-TT die size. A 34% increase in die size means a 34% increase in die cost, which translates to approximately a 49% increase in resale price to the end user! Designing with a multi-TT hub controller will make your end product nearly 50% more expensive than a performance equivalent single-TT hub-based product.

Conclusion

We all want to ensure we are getting the most for our dollar, so the higher the cost the higher the expected performance. But before choosing a hub controller for your next USB design, make sure you know what you are paying for; understanding the architecture behind a hub controller and the operation of each component can keep you from making a costly mistake. The benefits of a multi transaction translator hub controller are lost in today's USB systems, and you can end up paying for excess bandwidth that is wasted in your design. A single transaction translator hub controller can provide the performance you need at the price you are looking for.

References

Cypress Semiconductor
198 Champion Court
San Jose, CA 95134-1709
Phone: 408-943-2600
Fax: 408-943-4730
http://www.cypress.com