

Re: Why do not CRT/LCD monitors come with USB?

Source: <http://linux.derkeiler.com/Newsgroups/comp.os.linux.hardware/2008-07/msg00176.html>

- *From:* phil-news-nospam@xxxxxxxx
 - *Date:* 27 Jul 2008 20:58:18 GMT
-

On Sat, 26 Jul 2008 14:09:27 -0600 Joe Pfeiffer <pfeiffer@xxxxxxxxxxxx> wrote:

| phil-news-nospam@xxxxxxxx writes:

|

> On Thu, 24 Jul 2008 11:36:44 -0600 Joe Pfeiffer <pfeiffer@xxxxxxxxxxxx> wrote:

>

> | phil-news-nospam@xxxxxxxx writes:

> >

> > We need a new kind of interface. Existing ones like USB, Firewire, and SATA,

> > were designed for specific purposes original and were, or will be,

> > repurposed

> >

> > | Ummm, USB? USB most definitely designed from the ground up to be used

> > | for a wide variety of devices.

> >

> > It started as an interface to attach phones. Then it was added on to do

> > a few computer input devices. It wasn't until USB 2.0 that speed finally

> > became practical for things like disk storage.

> >

> > | Huh? That's the first time I've ever seen anything even resembling

> > | that claim. It was intended for exactly what it's good at: replacing

> > | interfaces for things like keyboards, mice, and printers. Do you have

> > | a cite?

I'll look for it. I saw it only a month or so ago. A quick check of Wikipedia says it was designed for a few different devices apparently even including modems. But it is obvious disk drives were NOT intended at the beginning, given the slow 1.0 and 1.1 speeds.

> USB had to completely change the way the OS drivers worked to go to 2.0.

> It wasn't simply a case of being faster within the same protocol. It had

> to change the protocol, too. The next speed higher will require a new

> protocol.

>

> | But it is just a case of being faster within the same protocol.

> | It's true that there's an extra negotiation required to go to

> | hi-speed, but once you're there it's the same protocol. I'd have to

> | spend more time than I want to to see whether the driver stack was

Re: Why do not CRT/LCD monitors come with USB?

| re-written at the same the EHCI driver was added.

It should have been possible to do 400 mbps and more with OHCI, then.

> A better design would be an interface between driver and hardware that just
> goes faster when the hardware can deliver faster. The driver gives the
> command to send a message and in the newer faster model, the interrupt for
> completion of command comes back sooner. The meta registers that tell the
> driver what speed would have a HUGE range capacity (like a 64-bit number)
> beyond anyone's imagination.

| So long as the hardware can reliably report what that speed is...

I don't see an issue with always starting at a slow speed during the hardware negotiation and ramping the speed up in steps. Each time it steps up, each side tells the other which speeds in the next step range it can handle are by means of flags. These steps might be multiples of 2X with a few speeds in each 2X range. Once both sides have no more common speeds, they use the best one they both match at. As each new faster generation of hardware comes out, they support even faster speeds and do even more speed negotiation steps (but at the faster speeds, this go fast).

The speed set might start at 100 mbps, require the base speed in each step be supported if any in that step are supported, and optionally allow finer speed steps such as 110, 120, 135, 150, 165, and 180. Then if any speed above 200 is supported, 200 must be. Each step then goes 200, 400, 800, 1600, 3200, 6400, 12800, 25600, 51200, and so on.

And that's not the only way to do it. It's just one possibility. Another way is for the initial negotiation to state their maximum speed in a 32 bit number (up to 4294967295 mbps) and then they narrow down from there what speed they can find in common.

> USB also slows down substantially through hubs. It can't keep data moving
> along and switch it in the hub at the same time.

| I'm not aware of anything in the protocol that makes wormhole routing
| impossible...

> USB device addressing is flaky in design. Plug in a disk, unplug it, then
> plug the same disk back in to the same port, and it has a different address.

| Whether that's flaky is a matter of opinion — the intent is that a
| non-savvy user can unplug a disk, plug it into a different port, and
| have the OS recognize it's the same disk. To do that, you have to
| decouple the device addressing from the port it's plugged into; once
| you've done that, there's no reason to treat same-port as a special
| case.

Re: Why do not CRT/LCD monitors come with USB?

Re: Why do not CRT/LCD monitors come with USB?

No you don't "have to decouple the device addressing from the port it's plugged into". You can remap the referenced device address or name to the changed physical address. This can be, and should be, done entirely in the OS in a layer just above the core layer of device drivers.

> USB is not very graceful in hotplugging. It can't figure out what is going
> on instantly. Firewire beats USB big time, and even Firewire is poor.
|
| Most of that time is just allowing the device to settle.

I see messages saying that. Maybe you can define what "settle" means. The only time frame I see that is needed is for repeated connections because the user is sloppy inserting the cable. There can be lots of "connecting" and "disconnecting" of the same device. Once the device stays connected for a fraction of a second, it should be good to go. It can now send its basic identify information. Extended information, such as how big the disk space is, can come later when the controller in the device gets the drive up to speed. Once that message arrives, then the OS can make the storage available to the storage subsystem. In cases of drives that know their size even before the spin up, they can report that immediately and the spin up can even be deferred until first usage.

This "waiting for device to settle" is one of the silliest things I have seen in USB drivers.

> > for other things. That's really not very good design. Something that is
> > designed to allow any kind of device from the outset, with standard
> > protocols included for a large variety of device uses (only ONE driver is
> > needed for any one class of device, regardless of manufacturer or
> > model),
> |
> | A compliant USB device of a particular class can be used with a
> | generic driver. The device may implement additional functionality
> | which requires a special driver, but core functionality can be
> | accessed without it (yes, I know Windows tends to not realize this.
> | That's not a weakness in USB).
> |
> | So why does the old OHCI driver not "just work" at faster speeds when a faster
> | USB port is used with a faster device? Why are there different drivers.
|
| Because the old OHCI (and UHCI) devices didn't talk hi-speed. You'd
| have to ask Intel why they developed EHCI instead of a faster OHCI,
| but that's completely unrelated to your claim. A new driver was
| needed for a new interface chip.

What do you mean by "didn't talk hi-speed". Just talk faster.

Re: Why do not CRT/LCD monitors come with USB?

The fault here is that Intel designed OHCI (and even EHCI) in such a way that the speed is integral to the protocol. I would not have designed it that way at all. The way I would have designed it is there would be interrupts from the controller whenever state changes occur. If the message gets to the device faster in newer hardware, then its just less time between when the message is sent and when the next message can be sent (aside from message queueing which would really allow some finite but reasonably large queue of tagged messages). The hardware would negotiate the point-to-point link speed. It can then be reported to the software in sufficiently large fields to allow a huge range of possibilities. And the basic need for the software to know the speed is to report it in logging and such. Some finite subset of speeds would make sense for the actual hardware interfaces. This can be defined in terms of a set of speeds and "all multiples of X thereafter".

> When 3200 bps USB comes out, will the same existing driver work even at the
> new speed?

|
| I don't know -- was the EHCI device designed with the needed headroom?

I have no idea.

> Yes, there are standards for many devices on USB. But there still appear to
> be a number of issues. Maybe manufacturers (like Seagate) just don't do USB
> right? Why even give such leeway? Apparently there is a lack of specification
> on disk drive spin down in USB.

|
| I don't know enough details on that -- it may indeed be a function not
| covered in the spec.

Then that would certainly be a deficit of the design. That's something I would have known to include at least 2 decades ago.

I can understand why there are full size and mini connectors. I would not have done the A-type and B-type connectors. I would have used a hermaphroditic connector design, at least for the larger one (at about the same size as the USB type B). That means a connector could plug into an instance of itself and mate the pins correctly, both data and power. Every cord would be symmetric.

|WARNING: Due to extreme spam, googlegroups.com is blocked. Due to ignorance |
| by the abuse department, bellsouth.net is blocked. If you post to |
| Usenet from these places, find another Usenet provider ASAP. |
| Phil Howard KA9WGN (email for humans: first name in lower case at ipal.net) |
.