

## What's the Matter with HDMI? By Kurt Demke – May 2007

---

**H**DMI, as we've pointed out elsewhere, is a format which was designed primarily to serve the interests of the content-provider industries, not to serve the interests of the consumer. The result is a mess, and in particular, the signal is quite hard to route and switch, cable assemblies are unnecessarily complicated, and distance runs are chancy. Why is this, and what did the designers of the standard do wrong? And what can we do about it?

The story begins with another badly-developed standard, DVI. A few years ago, there was a movement within the computer industry to develop a new digital video display standard to replace the traditional analogue VGA/RGBHV arrangement still

found on most computer video cards and monitors. Interested parties grouped together to form the Digital Display Working Group (DDWG), which developed the DVI standard.

DVI had all the earmarks of a standard designed by committee, and it remains one of the most confusing video interfaces ever. DVI could run analogue signals, digital signals, or both, and it could run digital signals either in a single-link configuration (in a cable using four twisted pairs for the signal), or in a dual-link configuration (using seven). Identifying which DVI standard or standards any particular device supported was not always easy, and the DVI connector came in various flavours and was never really manufactured in any

form that wasn't well-nigh impossible to terminate.

But the worst thing about DVI was something that the computer-display professionals involved in its development really didn't give much thought to: distance runs. Most computer displays are mounted at most a few feet away from the CPU, so it didn't seem imperative that DVI work well over distance. This lack of concern for function at a distance, coupled with common use of twisted-pair cable (e.g., CAT 5) in computer interconnection, led to a decision that DVI would be run in twisted-pair cable.



Had the DVI standard been designed by broadcast engineers rather than computer engineers,

things probably would have turned out very differently. In the broadcast world, everything from lowly composite video to High-Definition Serial Digital Video is run in coaxial cables, and for good reasons, which we'll get to in a bit. Long-distance runs of VGA, in fact, are always handled in coaxial cable (though there may be a number of miniature coaxes in a small bundle, rather than something which obviously appears to be coax).

DVI lacked a couple of things which the consumer audio/video industry wanted. It was implemented on a variety of HD displays and source devices, but it was confusing for the consumer because of the many variants on the standard and different connector configurations, and it didn't carry audio signals. A consortium to

develop and promote a new interface, HDMI, was formed; the idea was to come up with a standard which could be implemented more uniformly, was less confusing, and offered the option of routing audio signals along with video.

Here, again, was an opportunity to avoid problems. The difficulties of running DVI-D signals over long distances were well known, and the mistakes of the past could have been avoided by developing HDMI as a wholly new standard, independent of DVI. Instead, the HDMI group elected to modify the DVI standard, using the same encoding scheme and the same basic interface design, but adding embedded audio and designing a new plug. Instead of many DVI options, analogue, digital, single and dual link, there was one "flavour" of HDMI (actually, there is also a dual-link version in the HDMI spec--but you won't find it implemented on any currently available device). This provided the advantage of making HDMI backward-compatible with some existing DVI hardware, but it locked the interface into the electrical requirements of the DVI interface. Specifically, that means that the signals have to be run balanced, on 100 ohm impedance twisted pairs.

We're often asked why that's so bad. After all, CAT 5 cable can run high-speed data from point to point very reliably--why can't one count on twisted-pair cable to do a good job with digital video signals as well? And what makes coax so great for that type of application?

**First**, it's important to understand that a lot of other protocols which run over twisted-pair wire are two-way communications with error correction. A packet that doesn't arrive on a computer network connection can be re-sent; an HDMI or DVI signal is a real-time, one-way stream of pixels that doesn't stop, doesn't error-check, and doesn't repair its

mistakes--it just runs and runs, regardless of what's happening at the other end of the signal chain.

**Second**, HDMI runs fast--at 1080p, the rate is around 150 Megapixels/second. CAT5, by contrast, is rated at 100 megabits per second--and that's bits, not pixels.

**Third**, HDMI runs parallel, not serially. There are three colour signals riding on three pairs, with a clock circuit running on the fourth. These signals can't fall out of time with one another, or with the clock, without trouble--and the faster the bit rate, the shorter the bits are, and consequently the tighter the time window becomes for each bit to be registered.

Consider, by contrast, what the broadcast world did when it needed to route digital



video from point to point. The result was HD-SDI, high-definition serial digital interface. One coaxial cable can route an HD SDI signal hundreds of feet without errors, with no repeater hardware or EQs in the line. Had the consumer industry opted for a

coaxial-based standard, we'd be able to do the same in our homes. Admittedly, few of us need to make 300-foot runs; but the ability to run 300 feet without problems would be accompanied by rock-solid certainty of being able to do 50, or 75, without any worry at all.

But why is there such a big difference between twisted pairs and coax? It all has to do with the electrical properties of the two methods of routing signal from one place to another: balanced, through twisted pair, and unbalanced, through coax.

We tend to assume, when thinking about wire, that when we apply a signal to one end of a wire, it arrives instantaneously at the other end of that wire, unaltered. If you've ever spent any time studying basic DC circuit theory, that's exactly the

assumption you're accustomed to making. That assumption works pretty well if we're talking about low-frequency signals and modest distances, but wire and electricity behave in strange and counterintuitive ways over distance, and at high frequencies. Nothing in this universe--not even light--travels instantaneously from point to point, and when we apply a voltage to a wire, we start a wave of energy propagating down that wire which takes time to get where it's going, and which arrives in a different condition from that in which it left. This isn't important if you're turning on a reading lamp, but it's very important in high-speed digital signalling. There are a few considerations that start to cause real trouble:

1. **Time:** electricity doesn't travel instantaneously. It travels at something approaching the speed of light, and exactly how fast it travels depends upon the insulating material surrounding the wire. As the composition and density of that insulation changes from point to point along the wire, the speed of travel changes.
2. **Resistance:** electricity burns up in wire and turns into heat.
3. **Skin effect:** higher frequencies travel primarily on the outside of a wire, while lower frequencies use more of the wire's depth; this means that higher frequencies face more resistance, and are burned up more rapidly, than lower frequencies.
4. **Capacitance:** some of the energy of the signal gets stored in the wire by a principle known as "capacitance," rather than being delivered immediately to the destination. This smears out the signal relative to time, making changes in voltage appear less sudden at the far end of the wire than they were at the source. This phenomenon is frequency-dependent, with higher frequencies being more strongly affected.
5. **Impedance:** if the characteristic impedance of the cable doesn't

match the impedance of the source and load circuits, the impedance mismatch will cause portions of the signal to be reflected back and forth in the cable. The same is true for variations in impedance from point to point within the cable.

6. **Crosstalk:** when signals are run in parallel over a distance, the signal in one wire will induce a similar signal in another, causing interference.
7. **Inductance:** just as capacitance smears out changes in voltage, inductance--the relationship between a current flow and an induced electromagnetic field around that flow--smears out changes in the rate of current flow over time.

Impedance, in particular, becomes a really important concern any time the cable length is more than about a quarter of the signal wavelength, and becomes increasingly important as the cable length becomes a greater and greater multiple of that wavelength. The signal wavelength, for one of the colour channels of a 1080p HDMI signal, is about 16 inches<sup>1</sup>, making the quarter-wave a mere four inches--so impedance is an enormous consideration in getting HDMI signals to propagate along a cable without serious degradation.

Impedance is a function of the physical dimensions and arrangement of the cable's parts, and the type and consistency of the dielectric materials in the cable. There are two principal sorts of cable "architecture" used in data cabling (and HDMI, being a digital standard, is really a data cable), and each has its advantages. First, there's twisted-pair cable, used in a diverse range of computer-related applications. Twisted-pair cables are generally economical to make and can be quite small in overall profile. Second, there's coaxial cable, where one conductor runs down the centre and the other is a cylindrical "shield" running over the outside, with a layer of insulation between. Coaxial cable is costlier to produce, but has technical

advantages over twisted pair, particularly in the area of impedance.

It's impossible to control the impedance of any cable perfectly. We can, of course, if we know the types of materials to be used in building the cable, create a sort of mathematical model of the perfect cable; this cable has perfect symmetry, perfect materials, and manufacturing tolerances of zero in every dimension, and its impedance is fixed and dead-on-spec. But the real world won't allow us to build and use this perfect cable. The dimensions involved are very small and hard to control, and the materials in use aren't perfect; consequently, all we can do is control manufacturing within certain technical limits. Further, when a cable is in use, it can't be like our perfect model; it has to bend, and it has to be affixed to connectors.

So, what do we get instead of perfect cable, with perfect impedance? We get real cable, with impedance controlled within some tolerance; and we hope that we can make the cable conform to tolerances tight enough for the application to which we put it. As it happens, some types of impedance variation are easier to control than others, so depending on the type of cable architecture we choose, the task of controlling impedance becomes harder or easier. Coaxial cable, in this area, is clearly the superior design; the best precision video coaxes have superb bandwidth and excellent impedance control. Belden 1694A, for example, has a specified impedance tolerance of +/- 1.5 ohms, which is just two percent of the 75 ohm spec; and that tolerance is a conservative figure, with the actual impedance of the cable seldom off by more than half an ohm (2/3 of one percent off-spec). Twisted pair does not remotely compare; getting within 10 or 15 percent impedance tolerance is excellent, and the best bonded-pair Belden cables stay dependably within about 8 ohms of the 100 ohm spec.

If we were running a low bit-rate through this cable, it wouldn't really matter. Plus or minus 10 or 15 ohms would be "good

enough" and the interface would work just great. But the bit rate demands placed on HDMI cable are severe. At 1080i, the pixel clock runs at 74.25 MHz, and each of the three colour channels sends a ten-bit signal on each pulse of the clock, for a bit rate of 742.5 Mbps. What's worse, some devices are now able to send or receive 1080p/60, which requires double that bit rate.

Impedance mismatch, at these bit rates, causes all manner of havoc. Variations in impedance within the cable cause the signal to degrade substantially and in a non-linear way that can't easily be EQ'd or amplified away. The result is that the HDMI standard will always be faced with serious limitations on distance. We have found that, at 720p and 1080i, well-made cables up to around 50 feet will work properly with most, but not all, source/display combinations. If 1080p becomes a standard, plenty of cables which have been good enough to date will fail. And it gets worse...

In June 2005, the HDMI organization announced the new HDMI 1.3 spec. Among other things, the 1.3 spec offers new colour depths which require more bits per pixel. The HDMI press release states:

"HDMI 1.3 increases its single-link bandwidth from 165MHz (4.95 gigabits per second) to 340 MHz (10.2 Gbps) to support the demands of future high definition display devices, such as higher resolutions, Deep Colour and high frame rates."

So, what did they do to enable the HDMI cable to convey this massive increase in bit rate? If your guess is "nothing whatsoever," you're right. The HDMI cable is still the same four shielded 100-ohm twisted pairs, still subject to the same technical and manufacturing limitations. And don't draw any consolation from those modest "bandwidth" requirements, stated in Megahertz; those numbers are the frequencies of the clock pulses, which run at 1/10 the rate of the data pairs, and

why the HDMI people chose to call those the "bandwidth" requirements of the cable is anyone's guess. The only good news here is that the bit rates quoted are the summed bit rates of the three colour channels -- so a twisted pair's potential bandwidth requirement has gone up "only" to 3.4 Gbps rather than 10.2.

### What's to be Done?

It's unlikely, given the wholehearted way in which the consumer electronics industry has embraced HDMI, that this interface will disappear anytime soon. We're stuck with it. Given that that's so, what can a person do to avoid problems with video dropouts and outright signal failure?

**First**, limiting run lengths is a good idea whenever it can be done. If you don't need to put your sources at one end of the room and the display at the other, by all means avoid doing so.

**Second**, if run lengths can't be limited, consider relying on analogue component or RGBHV signals for your distance runs; these formats are much more robust (in large part because they run in coax rather than in twisted pairs) and can be run hundreds of feet.

**Third**, eliminating unnecessary switches, couplers, and adapters may help; as bad as the impedance mismatch problems are in the cable itself, those problems are even worse when the cable's conductors must be split out to join to a connector, or when the signal must travel through connections that can't be kept at 100 ohms.

**Fourth**, there are some things that can be done in cable design, and we're on the task. In particular, though the impedance of pairs can be controlled only to a limited degree, there are some things which the Chinese (who, as of this writing, manufacture all of the HDMI cable sold by anyone, anywhere, under any brand name) do not have the technical capacity to do but which American manufacturers do, and which help address the problem. Belden has a patented "bonded pair" technology which involves moulding twisted pairs together rather than simply twisting them, and which was developed specifically to address the problem of running high bit rates through twisted pairs. Beginning in 2005, we consulted with Belden on construction of such a cable for use in HDMI applications and in 2006, Belden built a series of sample reels of cable for us in its engineering lab. Our in-use testing has shown the cable working at 150 feet at ordinary high-definition resolutions (720p, 1080i) and up to 180 feet at 480p. Electrical tests of the cable indicate that it should be good for 1080p at a greater distance than any cable currently on the market. The cable has been ordered for full-scale production and should be available on our site around the very end of June or the first half of July 2007.

E&OE

## AudioTrends Showroom

10 Argent Place, Ringwood, VIC 3134

Phone 9874 8233

Hours: Mon - Thurs 10am to 5.30pm  
Fri 10am to 7pm, Sat 10am to 4pm



