Ghost in the HDMI Machine

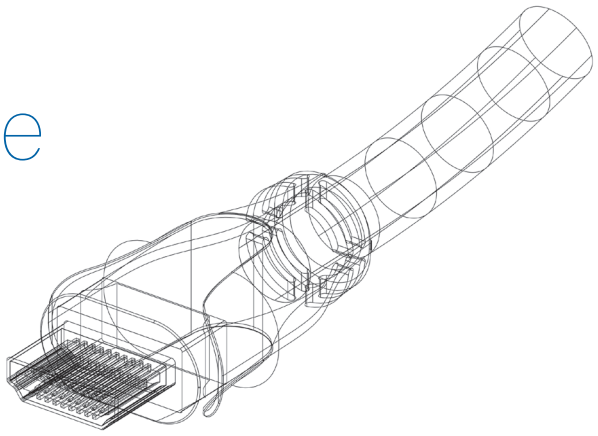

ADVANCED EXTENSION CHALLENGES & SOLUTIONS

Joseph D. Cornwall | CTS-D/I
AV Technology Evangelist, Legrand

High Definition Multimedia Interface, HDMI to most of us, is the single most important physical connection in the AV pantheon. The annual global HDMI market for cables alone is valued at more than $2.5 billion and has an expected combined growth rate (CAGR) exceeding five percent over the next several years. It is, without question, the most common AV connector in use today. It can be found on everything from laptop and desktop computers to docking stations, monitors, projectors, AVRs, matrix switches and much more.

One of the challenges of AV design is selecting the right solution to extend HDMI connectivity beyond its 15-meter passive connectivity limitation. Extension solutions exist at virtually every imaginable price point, but no matter the price, extensions must obey the rules of physics and engineering.

Unfortunately, obeying those rules doesn't guarantee a perfect solution. Since its introduction in 2005, HDMI has saddled AV designers and integrators with numerous complications. This high bandwidth connection is sensitive to a number of issues arising from manufacturing decisions and installation environmental conditions. This paper will explore the issues with HDMI and provide insights into Legrand's newest theories on how these issues manifest as system failures.

Let's dive into one of the most significant challenges of 21st Century AV design and learn how we can ensure dependable operation of our media solutions. Let's examine the ghost in the HDMI machine.

In the Beginning

A long time ago in a galaxy far, far away, analog technology defined the limitations of audiovisual system design. In those halcyon days, all that was needed to get a signal to the display was to connect a single coaxial cable to an input. Every device followed the same rules. Bandwidth defined the limits of resolution. Color was subordinate to luminance. Luminance delivered the whole of the system's detail. More importantly, nothing travelled contra to the payload. The sink (typically the display) had no reason to send a signal to the source, and the source had no reason to listen to the sink.

The digital revolution changed all that. Today, dependable digital video transport begins with a "hardware handshake" wherein data signals are exchanged between devices for the purpose of establishing communication protocols. Even before the link between boxes became a pathway for bits and bytes, the computer world demanded negotiation of horizontal and vertical image payload frequencies. This was accomplished using combinations of different polarities for the sync signals.

As imaging science progressed, so did the complexity of inter-device communication. The need for a source device to determine the capabilities of the sink, even when the sink is powered off, gave birth in 1994 to a +5 volt DC, 50mA power bus to drive the EEPROM and allow the host to read the Extended Display Identification Data (EDID) of the display. This power was delivered via pin 9 of the analog VGA cable and had to fall within +/- 5% tolerance to meet VESA (Video Electronics Standards Association) requirements. This configuration quickly became a blueprint for the next generation of display interface technology.

EDID Gets a Powerful Makeover and Becomes DisplayID

The HDMI specification requires implementation of the Enhanced Display Data Channel (E-DDC). E-DDC is used by the HDMI link to facilitate the exchange of Enhanced Extended Display Identification Data (E-EDID) information between the source and display, allowing for negotiation of the audio/video formats common to both devices. E-EDID does this by defining a remotely readable data file stored in an electronic display. The data in this file identifies the characteristics, features and video timing modes supported by the display product. The source compares the advertised capabilities of the sink to its own EDID table and settles on the highest quality exchange supported by both devices.

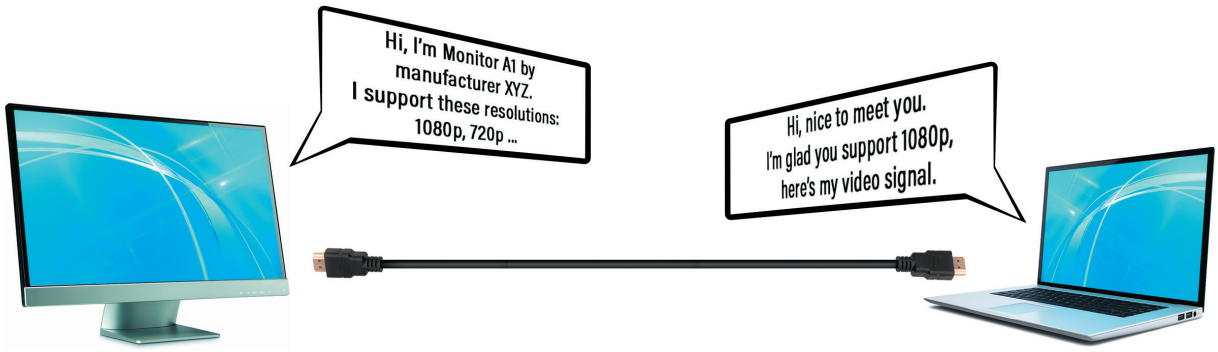

DisplayID, a VESA standard for metadata used in communicating a sink's capabilities to a source, is the newest and most powerful implementation of E-EDID. Full implementation of DisplayID hasn't found its way into the majority of devices yet, but the pending impact of HDMI 2.1 and its advanced feature set may change that.

Metadata is data that provides information about other data. DisplayID metadata, for instance, describes the technical process of connecting and configuring display devices such as PC monitors, consumer televisions (of all types), projectors and even displays embedded in other products such as the display associated with a laptop computer. DisplayID is the foundation of a simple, dependable plug-and-play experience for even the most advanced 4K and 8K display solutions.

DisplayID, originally launched in 2009 when the world was still predominantly analog, enabled the EDID exchange and was introduced to keep up with emerging digital display technologies. By the time DisplayID was ratified, HDMI (High Definition Multimedia Interface) was well on its way to becoming the world's most ubiquitous AV payload transport link.

HDMI implements the EIA/CEA-861 E-EDID standard associated with DisplayID that defines video and audio formats, waveforms and auxiliary data. The data structure includes manufacturer name and serial number, product type, chromaticity data, timings supported by the display, display size, luminance data and pixel mapping data. Included in this auxiliary data is HDCP (High-Bandwidth Digital Content Protection), a system that is now embedded in nearly every AV solution, both consumer and commercial. HDCP is intended to protect digital copyright interests.

I See You, 12C

I2C is an important piece of the HDMI transport protocol, and yet it's almost unheard of in the design and integration community. I2C is very inexpensive, simple to implement, and used when just a few bytes of information need to be transferred. It is a very structured interface with a well-defined protocol that lays out the rules for which device should be talking and which should be listening at any given time.

So what is I2C?

I2C is a synchronous, multi-leader, multi-follower, packet switched, single-ended, serial communication bus invented in 1982 by Philips Semiconductor. It's a communication protocol that transfers data bit-by-bit along a single pathway or wire. I2C (sometimes written as I^2C or IIC) is an inter-integrated circuit (IC) communication bus widely used for attaching lower-speed peripheral ICs to processors and microcontrollers in short-distance, intra-board communication.

I2C operates as a half-duplex, open collector/open drain design consisting of two links called SCL and SDA.

SCL is the Serial Clock Line. In electronics, a clock signal oscillates between a high and a low state (high and low voltages) and is used like a metronome to coordinate the actions of digital circuits.

SDA is the Serial Data Line. As its name implies, the SDA is the pathway used to deliver the data payload. The payload is eight bits, or one byte, long. Each byte sent across the SDA must be followed by an acknowledged (ACK) or a not-acknowledged (NACK) bit before the next packet can be sent. Receipt of a NACK typically generates a STOP condition to abort the transfer of the payload. A NACK may also initiate a retransmit condition in some situations.

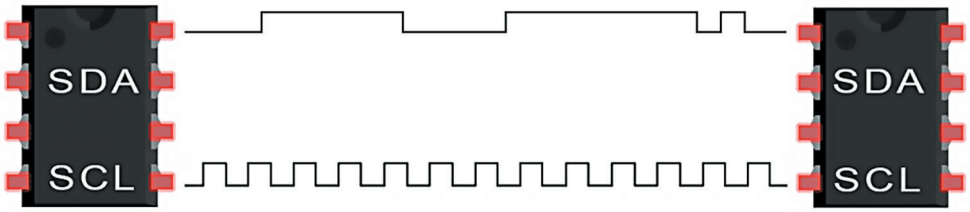

Start Bit	Address Frame – 8 to 10 Bits	Read Write Bit	ACK NACK Bit	8 Bit Data Frame	ACK NACK Bit	8 Bit Data Frame	ACK NACK Bit	Stop Bit

When data is being sent, the follower (receiver of the data) can hold the clock bit (SCL) at voltage low to throttle the leader (sender of the data) and control the transfer rate. The leader will only transfer bits when the clock bit is voltage high. The follower, holding the clock voltage low, pauses both the data stream and the clock for a tiny fraction of a second, allowing the follower device (sink) the opportunity to interpret the data. This process is called clock stretching.

Clock stretching is important when the amount of data being sent is highly variable, as it would be when eight million or more pixels are being individually defined by digital words in a high-performance UHD (Ultra-High Definition or 4K) digital AV system. In short, clock stretching forces the leader into a wait state, allowing more time for the follower to process and manage the data payload.

Let's Dive into HDMI

To understand HDMI extension techniques, we must first master the operational concepts of the HDMI interface itself. HDMI is a TMDS-conditioned AV transport method (Transition Minimized Differential Signaling is an 8b/10b algorithm) that delivers red, green, blue and sync (clock) payloads over four shielded twisted-pair wires embedded in the 19-conductor HDMI cable.

In an HDMI link, the audio data is embedded into the video data payload, thus requiring an active audio de-embedder if the sound is to be processed separately from the picture information.

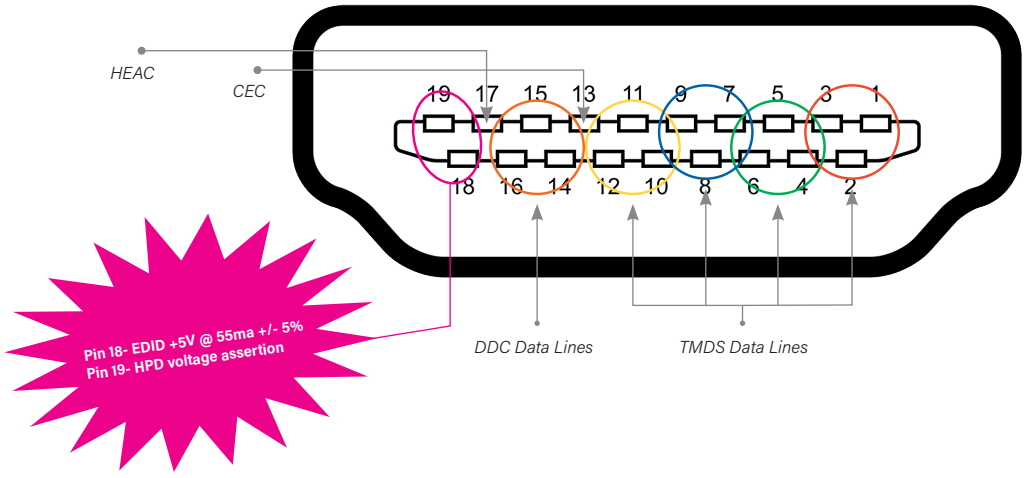

The illustration above will help us understand what's happening in an HDMI cable or device socket.

Pins 1, 2 and 3 are a shielded twisted pair that transports TMDS Lane 0.

Pins 4, 5 and 6 are a shielded twisted pair that transports TMDS Lane 1.

Pins 7, 8 and 9 are a shielded twisted pair that transports TMDS Lane 2.

Pins 10, 11 and 12 are a shielded twisted pair that transports TMDS Lane 3.

Pin 13 is a single wire bus that supports CEC (Consumer Electronic Control). [1]

Pins 14, 15 and 16 are a shielded twisted pair that transports the DDC data, including E-EDID and HDCP.

Pin 17 supports Fast Ethernet extension or Audio Return Channel (ARC) functionality. [2]

Pin 18 is the source for the power bus that allows for E-EDID data transfer across the DDC. [3]

Pin 19 is the HPD (Hot Plug Detect) pin. [4]

[1] CEC is a feature of HDMI designed for control of connected components, so individual devices can command and control each other without user intervention.

[2] Not all devices support these features. Also, this is an either/or application based on the source and sink, not the cable. It should also be noted that not all HDMI cables, devices or extension solutions support the HEAC bus.

[3] According to the HDMI standard, this pin must deliver a minimum of 55mA at 5 volts within a 5% tolerance. The standard identifies this as a minimum requirement. There is no mention in the current standard regarding power or voltages above this minimum threshold.

[4] When an HDMI connection is made, the source offers 5 volts at 55mA (for a minimum total of .275 watts) of power from Pin 18. The sink asserts (reflects) that voltage back to the source, beginning the hardware handshake process. It's important to note that the sink doesn't add anything to the voltage or power available across this link. Only the source delivers any power.

Legrand was among the first companies to illuminate the potential for HDMI failure because of voltage "starvation" nearly a decade ago. There are a number of potential issues related to Pins 18 and 19 and the HPD voltage assertion process. In the simplest of cases, voltage drop across an excessive length of copper wire can drive the HPD voltage below the 5V +/-5% threshold.

Voltage Drop

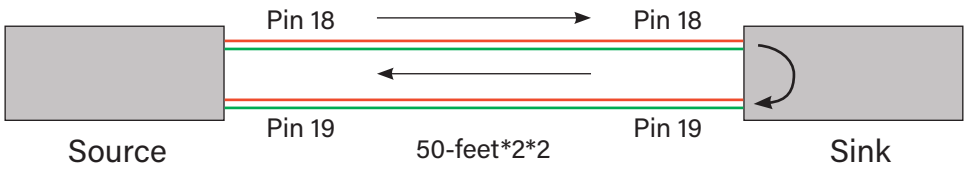

A 24AWG copper wire is .0201 inches in diameter and presents 0.026 ohms per foot of cable in DC resistance. If a 50-foot cable is deployed, the total loop DC resistance of Pins 18 and 19 in the circuit (shown mathematically, 0.026×50×2×2) is 5.2 ohms.

Why do we say 50×2×2?

The cable is 50 feet long, but the signal must traverse the distance twice (it is asserted, not actively repeated), traveling from the source to the sink and back. Hence, the total length of the positive leg of the circuit is 100 feet.

Why do we see the second (×2) multiplier?

Because a circuit must have both a positive and a negative leg. The negative leg is also 100 feet long, of course. What isn't shown in the connector pin-out Illustration is that the cable shielding is used as a common ground return for single-ended bus pathways. It's true that the shielding may have a slightly different DC resistance than the single copper conductor of Pins 18 and 19, and in some implementations, we may not need to take this additional leg into account. For the purposes our exercise here, it can be modeled using the same assumed resistance as the positive leg.

Ohm's Law tells us that the voltage drop across a resistor is described by the formula V=IR, where I=0.055A and R=5.2Ω. Therefore Vdrop = (0.286V). Five percent of 5V is 0.25V.

An online voltage drop calculator can be found at this URL - https://www.calculator.net/voltage-drop-calculator.html. Input the following attributes: copper, 24AWG, 5 volts, specify a single set of conductors at 100 feet and 0.055 amperes of current. We will find the rounded voltage drop confirmed at -5.65%, with a DC voltage drop of 0.28V across the cable.

Right away, we see this link may fail if the source device is limited to the minimum requirements of the HDMI standard. The HDMI standard requires 5V at 55mA to be held to a +/- 5% tolerance. Remember that meeting the minimum requirement isn't a failure, as the device is actively meeting the requirements of the standard!

The result is that the longer the HDMI cable, the greater the wire gauge a manufacturer must use to have any confidence the link will work in a preponderance of applications. At some point, the cable will inevitably become unwieldy and inflexible.

When a larger copper wire gauge is used, the cable will be thicker, which affects the bend radius, an important parameter during installation. The bend radius of a cable shouldn't be less than six times the diameter of the cable. A 50-foot HDMI cable deploying 20 or 22 AWG conductors might have a diameter of one-half inch or more. This translates to a minimum bend radius of 3 inches, and that isn't conducive to use in a low-voltage, in-wall electrical box!

In addition to voltage starvation, there are also concerns regarding eye pattern integrity. Link bandwidth is always inversely proportional to link length. The longer the cable, the more likely the payload will be adversely affected by the electrical characteristics of the cable itself, such as capacitance, inductance, and impedance.

The eye pattern of a signal may also be negatively impacted by alien crosstalk from RFI-EMI (Radio-Frequency Interference and Electromagnetic Interference) introduced by an electrically noisy installation environment or poor design in the cable. Alien crosstalk is often introduced by neighboring cables carrying unrelated signals. To ensure performance at longer lengths, or to create a slimmer and more flexible connection, a very different solution is needed.

Cable capacitance plays a role here, too. Capacitance acts as a high frequency filter, distorting the waveforms that represent the dense digital data of the high definition payload. This almost certainly will impact I2C and DDC payloads and contribute to the operational success of the design.

The HDMI DDC version 1.3 distance specification defines the parameters necessary to optimize the link and ensure proper operation. Cable capacitance is a function of the geometry of the cable design. It is an outcome of the permittivity of the dielectric material between the plates of a capacitor, which, in this case, is the dielectric insulation used in the cable construction. It is purely a physical element. This means that the capacitance changes with different cable lengths, diameters, wire gauges, intrinsic material qualities, manufacturing processes and installation techniques.

Active Optical and Active Copper HDMI Cables

An active copper cable is a directional cable with a chip within the HDMI connector (most often at the display or sink end) that boosts the signal from the source device to the output device. The exact nature of the chip and what it does is generally proprietary to the cable and/or chip manufacturer.

Active cables are not new, and they're not unusual. In fact, active cables are an absolute necessity in many high bandwidth applications such as Thunderbolt, USB4, and 40GE networking.

Active HDMI Copper Cable

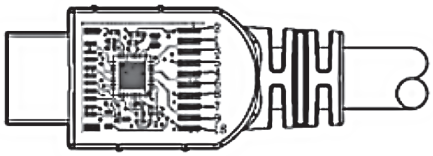

An active optical cable is a directional cable that converts electrons into photons and uses fiber optic strands instead of copper to carry the demanding payload. Active optical cables (AOC) are composite cables, meaning they have a combination of both copper and fiber pathways.

The TMDS signals, which demand the greatest bandwidth and must exhibit the highest fidelity to the original eye pattern, are transported across the fiber links. The remainder of the signals are transported across traditional copper links. This includes the HPD power from Pin 18 and the asserted voltage back to Pin 19.

Because fiber optic strands are extremely thin, measuring only 125 microns in diameter, optical links may be much smaller in diameter than comparable copper links. AOC solutions can be thinner, plenum-rated, and offer a smaller bend radius for greater installation flexibility.

Active Optical Cable Cross-Section

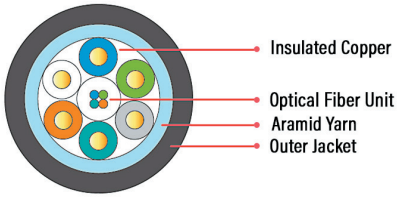

An active optical cable works by using VCSELs (Vertical Cavity Surface Emitting Laser) to convert the electrical payload into light. VCSELs are very small devices that can pulse laser light based on the pulses of data arriving in the form of an electrical signal. VCSELs are used to drive fiber optic communication links in LAN, OSP networks (Outside

Plant), and data center communication systems. The technology is now more than a decade old, but it's still finding new uses in many applications, including the optical solutions being deployed in the AV space.

There are four VCSELs in an active optical HDMI cable. They are used to convert TMDS lanes 0, 1, 2 and 3 from electrons to light-based transport. This allows for significantly greater link lengths. Additionally, this improves signal fidelity since fiber is essentially impervious to RFI and EMI, which can dramatically affect the payload in copper links.

In addition to fiber solutions providing a much smaller and more flexible footprint, they also deliver significantly greater link bandwidth. The bandwidth of the glass fibers is several orders of magnitude greater than copper. The usable bandwidth in a cable assembly is essentially limited only by the speed at which a VCSEL can turn on and off (pulse).

Most AOC solutions will use copper conductors for the DDC, HPD and CEC buses. This way, there's no need to add a power supply at both ends of the circuit. Active optical cables need power at the source end to power the VCSELs and at the sink end to power the optical receivers. That power is typically carried by the copper strands. Of course, the same concerns about Ohm's Law and voltage drop that apply to the pure copper HDMI cables described above also apply to these copper conductors.

With both active copper and active optical cables, there's an obvious question that should be asked: "What's powering the chip sets in these solutions?" The answer, of course, is the power bus on Pin 18 used for HPD.

What we now understand is that this power bus was never designed to provide power for anything beyond EDID exchange and HPD. Very few active cable solutions have any method for alternate power. What's more, there are no changes to the HDMI standard proposed that will significantly alter this situation.

Uncovering Unexpected Problems

Legrand's engineering department has been researching this issue for some time. Often when an active solution or passive cable doesn't perform as expected in a system, the design or build quality of the cable is immediately suspect. As we've seen here, it's entirely possible for an active or passive cable to be well designed and appropriate for its task, only to find that associated components can't provide the power or signal integrity necessary for optimal performance. There is a lot to uncover here.

First, let's examine some performance aspects with a simple observation and thought experiment. Assume a source, like a laptop computer, is connected to a sink, such as an LED flat panel via a 50-foot active optical cable. This is a very common configuration in a conference room or classroom. Now let's assume this system is working perfectly for all content rated at FHD (full HD, aka 1080p).

Let's also assume that the active optical cable is rated for payloads up to and including 4K UltraHD (UHD) 2160p in a 4:4:4 RGB color space. The LED display is also a 4K device, and all switching, patching, matrix and other associated links are similarly rated for this level of performance. A user attempts to display 4K content, but the system fails. There's either no picture or the picture "stutters" and is distorted. What happened?

When the active optical cable is transporting an FHD payload, the VCSEL lasers must flash on and off to communicate each bit sent across the fiber link. In an FHD system, the lasers carrying the payload will fire more than 3.7 billion times combined each second. It will actually be significantly more as "dark pixels" representing blanking and retrace timing must be added to the active pixel count. (3.73Gbps = 1080×1920 pixels×60 frames/sec×8 bits per pixel×3 for each of the red, green and blue links×1.25 for the 8b/10b overhead associated with the TMDS conditioning used in HDMI communication).

Breaking this down to a single VCSEL means each of the lasers is pulsing on and off more than 1.2 billion times per second. The clock link must fire just as quickly to keep things in sync, so this is a total of 4.8 billion "blinks" each second across the four optical lines.

What happens when the new user connects her brand new 4K-rated laptop to the system that's been working perfectly for everyone else at 1080p?

The EDID exchange allows the sink and the source to determine maximum resolution and color space capabilities of each device. Since both the source and the sink are now rated for 4K content, the hardware handshake will settle on the highest resolution common to both devices - 4K. This is the exact purpose and correct outcome of the hardware handshake. The system will try to operate at this level, since there's nothing telling it that it can't.

An UltraHD 4K signal operates at 2160 visible vertical pixels × 3840 visible horizontal pixels. So the uncompressed bandwidth of this exchange is now 14.92 Gbps (14.92Gbps is 2160×3840×60 frames/sec×8 bits per pixel×3 for each of the red, green and blue links×1.25 for the 8b/10b overhead associated with the TMDS conditioning used in HDMI communication).

Breaking this down to a single VCSEL means each of the lasers is pulsing four times faster for a 4K payload than is the case with 1080p. We are now asking each lane in this system to blink on and off nearly 5 billion times per second. Clearly the power supply on Pin 18 of the HDMI link in the laptop must work much harder to power the additional work of the VCSEL lasers. It's entirely possible the source just can't muster this kind of power, even though the system is properly rated to work at the UHD level and otherwise meets all the specifications of the HDMI standard!

Put simply, processing more data requires more power. This ratio holds true for both active optical and active copper solutions.

What Our Research Found

There's yet another part to this story.

As we move from a standard FHD to a payload that is HDR (High Dynamic Range - each pixel component is described by more than 8 bits of information, up to 16 bits per pixel component), UHD 4K, or both, it stands to reason that we are asking the system to process many more bits of information. There is something of a "speed limit" to any link. It's a combination of the link's propagation velocity and the circuit's processing latency. Whether electrons (active copper) or photons (active optical), it takes a set amount of time for the 1's and 0's of the payload to move across the link and be processed.

The more bits we send, the more time it takes for all of them to get through the system unless the data rate is increased. With many of the systems in use, there is a maximum data rate built into the format. For example, HDMI 1.4 has a maximum data transfer rate of 10.2Gbps, while HDMI 2.0 ups that "speed limit" to 18Gbps.

Enter clock stretching...

Clock stretching is a feature found in all HDMI pathways in version 1.3 and above. Clock stretching is related to the DDC control bus and relieves the source and the sink from having to synchronize exactly at a predefined rate. Basically, the display can tell the video source to slow down a bit. Most HDMI sources support clock stretching, but not all do. HDBaseT products also require clock stretching.

Issues involving clock stretching will primarily appear with HDCP-protected content. When HDBaseT extenders are used, they packetize the DDC data at the source and then reverse the process at the display. This process causes a delay that can exceed the limits of the source and stop HDCP content if the EDID isn't received in time.

Let's take a look at what this means for that I2C circuit we visited above.

When the receiver chip set is sent more data, more time is needed for the chip to process the bits. This is why clock stretching capability is designed into these circuits.

When the SDA (data line) is held LOW constantly by the follower device, the system consumes more current from the 5V supply. This is added to the burden of the extra power demanded by the active chip sets in the cable.

Yet another issue can manifest in this situation. Stressing the source device's HPD power supply may overwhelm the IC buffer in a less than optimal product design or possibly cause crosstalk noise between the SLC clock and the SDA data lanes. The ACK command may then be interpreted incorrectly as a follower device issue. The system may think the follower device is holding the SDA line LOW. This will cause the data exchange to stop and no image will be displayed.

To better understand the chaotic nature of link failure in extended HDMI solutions, the Legrand engineering team measured the performance of a number of devices. The table below shows the results. This observation looks at the power output capability of the source devices under test from Pin 18.

Table 1 – Device Power Supply Analysis

Source \ Test Load Current	50mA	100mA	150mA	200mA	Estimated Max Current
QD780E	4.97V	4.97V	4.84V	**Protected**	150mA
MacBook 2014	5.1V	5.1V	**Protected**	**Protected**	134mA
Dell Precision 7730	5.05V	5.05V	4.94V	4.88V	350mA (@4.71V)
HP Zbook G5	5.04V	5.04V	4.98V	4.95V	> 1A (@4.5V)
Lenovo X1 Carbon	5.12V	5.12V	5.05V	5.02V	413mA
ScreenBeam 1000 EDU	4.64V	**Protected**	**Protected**	**Protected**	68mA

As you can see, there is great variability in the available power from the HDMI link on the tested sources. This isn't uncommon. Every device manufacturer will optimize their design to hit a desired performance level, but there's no incentive to overbuild once the requirements of the operational standard are met. The actual output ability of any device is primarily driven by the cost of the circuit design and the intention of its designers. Power in excess of the HDMI minimal standards is most often a result of serendipity and not intentional design.

Beyond this, there is almost no guidance for an AV system designer to know how the power supply of an HDMI source will behave. Testing each source device for its power budget limitations is both time consuming and delicate. There has been, at least up until now, no real demand from the Pro AV community for these specifications, so there has been no reason for manufacturers to publish such numbers.

It's safe to assume that such variability in HPD power supply limitations is common throughout the industry. Remember, a device that meets the minimal requirements of the HDMI standard is in compliance. To the extent that some sources provide more than the minimal required voltage and current is often the result of luck.

I2C Circuit Analysis Showing Interference Spikes

Notice the spike on the SDA line (circled in red). That spike is a result of an improperly implemented I2C solution. It is noise created when the IC buffer used in the I2C circuit, or the power supply of Pin 18 on the HDMI port of the source device, is overtaxed. Essentially, the supply is incapable of delivering the requested voltage and the waveform is clipped, or there is internal perturbations caused by data or power levels not anticipated by the designers of the product. This may cause the emission of spurious noise that then interferes with the communication bus.

Because of the crosstalk noise (SCL to SDA), the ACK signal may be interpreted incorrectly as a follower device issue, causing the communication to stop because the waiting follower keeps holding the SDA line LOW. When the SDA is held LOW constantly by the follower, the system consumes even more current from the 5V supply of Pin 18.

Misreading the ACK/NACK command stops signal transmission

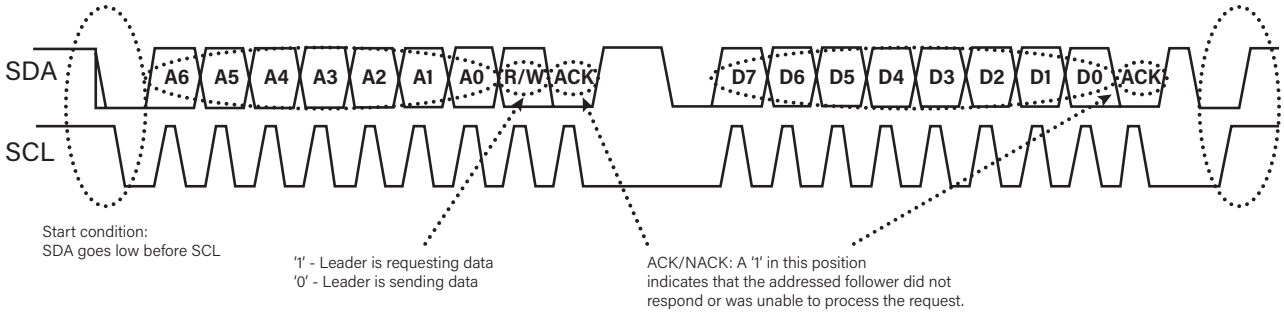

It's at this point when a previously working and dependable installation fails. Active cable manufacturers typically don't disclose the power demands of their ACC or AOC solutions, just as device manufacturers don't disclose the limits of the power supplies in their devices. It's exceedingly rare for a manufacturer to even suggest the power demands of an active solution may vary with payload and application. This makes it nearly impossible to effectively analyze a power budget for an active HDMI link if both the capabilities of the source/sink and the demands of the link are both unknown.

Now we have a situation where the user may lay the blame for the failure on the computer, the display, or the link between the two. In truth, each part of the design is meeting the requirements of the standard, but not the requirements of the installed system. From the designer to the integrator to the end user, every party ends up disappointed in the outcome.

Mitigating HDMI Extension Link Issues Today

There are not currently a lot of choices when it comes to addressing these types of failures in the market right now. This is an emerging issue that is just beginning to be examined and addressed. Legrand is leading the way in these efforts.

To overcome simple voltage starvation, a voltage inserter can be used. This solution breaks the connection of Pin 18 at the source and substitutes an external power supply in place of the supply built into the device.

The voltage inserters offered by Legrand can deliver as much as 1000mA@5V, allowing them to power not just the active elements of an HDMI extension technology, but also peripheral devices that are increasingly reliant on the Pin 18 power supply. These devices run the gamut from audio de-embedders, to auto-sensing input switches and active signal splitters or distribution amplifiers. Voltage inserters are available in a number of form factors to fit in with a variety of installation configurations.

A voltage inserter won't cure all failures, however. Sometimes active signal reinforcement is required. In situations where this is true, an active in-line HDMI extender may be needed. In-line extenders may be powered from Pin 18 (or a voltage inserter as detailed above) or more commonly from a dedicated device power supply.

An active in-line extender does more than just increase the power capability of the HPD power bus. For example, the Legrand in-line HDMI extender is designed with an HDMI retiming buffer, and it equalizes the incoming TMDS signal. This ensures that an optimally retimed data output with maximum eye pattern fidelity is presented to the downstream sink. This product will also buffer the I2C bus and handle clock stretching demands more efficiently than most source devices can do natively.

The challenge of an in-line HDMI extender is that there is yet another small box to be added to the list of devices included in the installation. To complicate matters, in-line extenders often work best when positioned at some arbitrary mid-point between the source and the sink. In plenum or fully conduited installations such positioning may be difficult to achieve.

What Will the Future Bring?

At Legrand, we are defining what it means to create and build an optimal HDMI extension design. We are doing this by analyzing and evaluating the capabilities of active devices and anticipating the needs we'll have to meet to work dependably with today's content and with the 4K, 8K, wide-gamut color and high dynamic range payloads of tomorrow.

Our newest family of HDMI Active Copper and Active Optical extension devices are already optimized for power efficiency. They require less power, exhibit less variation in power demands as mapped to payload complexity, and are more stable under all conditions. This will ensure confidence that a good AV system design remains a good and dependable design for years to come, even as technology evolves.

We aren't stopping there. We're also preparing for next generation solutions that demand compliance with advanced powered cable assemblies. The new HDMI 2.1 specification increases permissible maximum data rates from 18Gbps to 48Gbps. There is also a powered cable assembly (PCA) addendum to the HDMI 2.1 spec that calls for an auxiliary power supply for active cables.

Under this new specification, advanced sources may deliver enough power for the active cables designed to support the enhanced bandwidth of next generation connectivity, but not for powering external peripheral devices. This new HDMI standard relies on manufacturers disclosing standardized support capabilities, but keep in mind that most manufacturers aren't even close to understanding the limitations of HPD bus power in their current product lineup!

As explored in this paper, our tests showed how increasing a payload from 1080p to 2160p resulted in a significant increase in power demanded by the extension solutions. What happens at 4320p? What happens at 2160p, but when the color space is 4:4:4:4 and an alpha channel for content mark-up is included in the payload? What happens when the link must support multi-stream signaling? These are the payloads of the future. We are exhaustively researching these questions and more.

Our research includes the subtle effects of ground offset, which may cause misinterpretation of the I2C SLC/SDA links.

We are also investigating the importance of electrostatic discharge (ESD). Electrostatic charges are generated when different materials come into contact with each other and are separated. This is a physical attribute of the connecting metals, internal connecting links, and even the nature of the insulation used in the cables.

Legrand was on the forefront of providing HDMI power insertion solutions beginning in 2012. Today, we are working to build in versatility of design by anticipating every step of the signal path from the source to the docking station to the system input to the display input. Recall that being "rated" to higher resolutions is a different thing entirely than actually being applicable to system designs operating at those resolutions.

Legrand is using state-of-the-art testing techniques to analyze and evaluate all the contributions of the physical layer link. In this way, we'll create dependable, predictable, efficient solutions that support creative system design that anticipates the next generation of video system performance.