

Optimizing & extending the reach of high speed copper cable assemblies

James J. David

FCI Electronics

825 Old Trail Road, Etters, PA. 17319

+1-717-938-7841 · jim.david@fci.com

Abstract

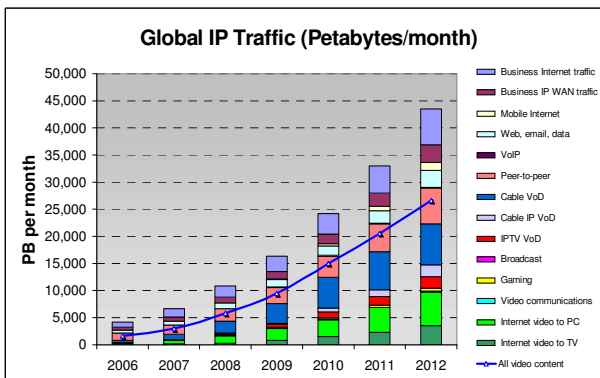
This paper examines the market dynamics and resulting technical challenges affecting high speed copper based cable assemblies and connectors that meet the requirements for 40 and 100Gb/s links in next generation systems. The subjects covered will include market trends and systems designs driving these bandwidth needs, industry standards intended to meet those needs, and the challenges these pose to the I/O link suppliers. Some of the challenges we will explore include signal conditioning, I/O port density, cable assembly testing and certification and recognition of the cable assembly by the system upon insertion of the cable in the system. This paper also explores and projects where all of these technical challenges may be leading and what that potentially means to high speed cable & I/O link suppliers.

Keywords: High speed cable assemblies; SFP+; QSFP+; CXP; Fibre Channel; Infiniband; 10G Ethernet; 40G Ethernet; 100G Ethernet; SAS; SATA; system recognition; EEPROM; pre-emphasis; equalization; BERT; WDP; VMA; interoperability; QDR; IEEE 802.3ba

1. Introduction

The Internet revolution, like an unstoppable force, marches on continuing to evolve and create communication links, networks and applications that many of us never conceived of or saw coming. Whether its public transmission of video on sites like YouTube, social networks like Face Book and MySpace, downloading of digital MP3 files or video files for instant entertainment or real time Twitter updates from friends or even celebrities, our world is now one that craves real time, information-rich news and data within very short periods of time. It's daunting to realize that, while each one of these sites is ubiquitous in our world today, none of them even existed 5 years ago.

Chart 1. Global IP traffic estimate – 2006-2012




Source: Cisco Visual Networking Index, June 16, 2008 ¹

All indications are that this hunger for still more and more data and bandwidth will continue unabated for the foreseeable future driven by emerging video-rich applications like IPTV, peer-to-peer, video-on-demand, and Internet video to TV and PC's. The demand for video content is expected to grow at a CAGR of 52% from 2008-2011. See chart #1 above for a graphic illustration.

Given this projected growth in demand companies and industry organizations have been diligently working to ensure specifications and products are ready to address these anticipated capacity needs. A number of industry standards have been developed to help bring some commonality and internetworking functionality to the hardware connections and software communications. These industry standards include Infiniband, Fibre Channel, 10Gigabit-Ethernet, Serial Attach SCSI (SAS), and Serial ATA. Meanwhile organizations like the Infiniband Trade Association and various IEEE 802.3 sub-committees are in the process of finalizing standards that address the industry's desire for 40 and 100Gb/s bandwidth-capable systems and I/O links. Further reinforcing these expected future trends are published development roadmaps that point to link bandwidths that extend beyond the 100 Gb/s systems. An example is shown in Chart #2 below.

Chart 2. Infiniband trade association roadmap ²




Networking, storage, computing system and equipment suppliers, IT staff at data centers and service providers are all striving to keep pace with this burgeoning consumer demand. Satisfying these needs isn't always an easy thing to do as they must balance the customer demands for timely, reliable and cost effective delivery of services against high energy costs, maximized equipment utilization and overall data center productivity and efficiency. What is clear is that these challenges will continue in the future and will likely affect virtually all data and communication equipment platforms whether targeted for switching, routing, server or storage applications.

Chart #3 below shows the projected link speeds for years 2009-2020. The chart clearly shows the emergence and growth of the 10 Gb/s connections over the next 5-6 years followed by a similar growth cycle for the 40Gb/s connections beginning around 2015. One perhaps can argue about the specific timing and adoption rates

in these systems but it is difficult to debate the clear future trend for wider bandwidth.

Chart 3
x86 Server Ethernet Connection Speeds with 40GbE & 100GbE



Source: Intel & Broadcom, April 2007

Base Server Connection: | 10 year transition | 5 years | 5 years |
|-----|-----|-----|
GbE | 10GbE | 40GbE

IEEE 802.3 High speed study group interim meeting, April 2007³

Industry analysts expect copper based 10G Ethernet switch ports to experience 143% CAGR from 2008-2012 and high end router demand for 10G ports to see a CAGR of 31.1% during the same time period.

Emerging from all of this industry effort are new and evolving cable link standards and interfaces such as SFP, SFP+, QSFP, QSFP+, CXP, mini-SAS/ SATA, mini-SAS HD and CFP which enable the high speed external and internal cable links needed to address the explosive growth outlined above.

2. Optimizing copper cables

So where does all of this product planning and technical advancement leave today's copper-based I/O link solution supplier? The short answer is that making a cable assembly for these systems isn't as simple as it used to be. There are a number of challenges that any viable cable assembly supplier must address in order to assure a high quality, compliant interconnect link is supplied to their customers. Addressing these challenges is only one part of the solution; the other is maintaining the operational systems that assure consistent quality and compliance of the link.

2.1 Equipment design

Equipment and system designers are being challenged in numerous ways as they attempt to meet the rapidly growing bandwidth demands. Technologies such as multi-core processors, virtualization, consolidation, increased host bus speeds and memory performance have certainly helped increase the available capability a designer can integrate into a system design. These technologies, however, do present challenges for network bandwidth transmission, power management and cooling. The need for increased signal speeds makes the continued use of common printed circuit board materials like FR-4 and cables with commonly used insulations and manufacturing processes an increasingly difficult task. In addition higher power is needed and adequate signal integrity is more difficult to achieve. These needs in turn present other challenges to both the system designers and data center managers. For example, given today's chip technologies and capabilities, it has been estimated that a single Internet Google search requires a total of 3 watts of power to complete the inquiry. What is generally goes unacknowledged is

that such a search, in order to properly dissipate the heat generated by the search, requires an additional 3 watts of power. These increased power needs are driving designers to employ "green" techniques, like port power management functionality that directs the port to automatically go into a "sleep" mode when not being utilized. The intent is to better manage the power consumption. So system designers and the users of those systems have multiple and sometimes conflicting considerations that they need to take into account and somehow balance in next generation equipment designs. Increased speed and power aren't the only things designers and users need to pay attention to as other factors like proper system heat dissipation and management, sufficient air flow, cable routing and EMC / EMI shielding, port density and cable assembly mounting and latching also require careful design consideration.

Equipment customers are looking for flexible, future-proofed systems with headroom that are easy to install and are easy to maintain. They would also prefer to maintain existing system port density or, preferably, see an increase the I/O port density to support higher system capability. They also prefer the ability to freely designate or configure any available system port as they see fit with minimal issues and cost.

All of these needs have helped to necessitate a closer working relationship between system designers and high speed I/O system suppliers. In the past, there wasn't a lot of collaboration between these two disciplines primarily because it wasn't really necessary. However, with the advent of the higher signal speeds, it became apparent, in order to meet all of the goals outlined above, that a higher-level working relationship between the system designers and the I/O system designer was necessary. It also required both parties have a deeper appreciation and understanding of the specific functional and design capabilities each party is capable of bringing to the overall system design without adding excessive costs and overhead. This new dynamic is best illustrated by the collaboration of industry standard organizations, committees and sub-committees as well as industry ad hoc groups, such as the Small Form Factor committee (SFF), where a great deal of discussion and collaboration takes place. This interaction has become an absolute must if equipment suppliers will ultimately give their customers what they are asking for. The I/O system supplier needs to give the equipment designer as much product flexibility and functionality as possible.

2.2 I/O system solutions

The good news is that there are I/O systems that have been developed that address many of the requirements. XFP and SFP copper and fiber optic based I/O systems have been in the market for some time now. They have been instrumental in bringing I/O port bandwidths to the 5-6Gb/s+ per channel capability level. They have also designed these systems to be as small as possible in order to minimize the linear board "shoreline" required. The SFP system significantly reduced the module outlines and shoreline required from earlier GBIC and XENPAK systems. One parameter that the existing SFP systems did not address was the capability of supporting 10Gb/s per channel capability which is being demanded today. This has led to the development of the SFP+ systems and their ability to properly support the 10Gb/s per channel need. While the SFP and SFP+ systems share the same dimensional board space and connector and cage architectures only the SFP+ systems support the 10Gb/s per channel bandwidth capability.

This product developmental progression is continuing with the recent developments of industry standard interfaces such as QSFP, QSFP+, mini-SAS/SATA, mini-SAS HD, CXP and CFP. The QSFP+ system has been developed to address the need for an I/O system capable of supporting a 40Gb/s total bandwidth in each port. Similarly the CXP system is being developed to support systems that are looking for 100-120Gb/s total bandwidth per port capability. Both of these systems are being developed and offered and aligned with a number of interconnect technologies such as Infiniband and Ethernet and are being adopted by multiple industry standards. Both systems offer connector and cage products that support either a passive copper based cable solution generally used for relatively short length cables (5-7 meters or less depending on the acceptance criteria), an actively equalized copper based cable solution for longer lengths (up to 20-25 meters depending on the acceptance criteria), a plug-in optical transceiver module with an optical based I/O connector on the back side of the module or a active optical transceiver module that has the optical fiber terminated inside the module housing. This architectural approach gives the system installer and system user the flexibility to define and change the port configuration and capability as it is required.

The CFP system, which has been recently introduced as an MSA standard, adopts a similar approach to the QSFP interface in that its use is intended to support 100Gb/s bandwidth systems. The CFP system, as it's currently configured, will always have the transceiver capability embedded in the module and a standardized two-piece connector interface on the equipment port side of the module. The I/O side of the module allows for multiple port configuration options (multiple SFP+, multiple QSFP+, CXP or combinations) that can be customized depending on the customer desired I/O interface and distribution.

2.3 Maximization of port density

If an equipment supplier were to question their customers about how much port density they require the answer is likely to be along the lines of "...keep the density I have as a minimum but I'd really like more port density if it's possible". Customers will always want as much functionality as possible in a system as long as the costs associated with doing so do not become prohibitive. The I/O system designers have done a commendable job in reducing the size of the XENPAK and XFP modules as evidenced by the SFP+ and QSFP+ systems. The shrinking of the module size, while increasing the system functionality and capability without any major tradeoffs or lower performance, was hardly a straightforward or simple task. In fact it was, and continues to be, a very difficult goal to accomplish. Design considerations such as mechanical robustness, user friendly cable latching and mounting, heat dissipation (for active modules), EMC shielding compliance and signal integrity all tend to be more difficult as the system footprints shrink in size, I/O connector density increases, and system performance increases. By way of a simple example, let's look at what happens when a system's scale gets smaller. Assuming all other aspects remain the same, the action of reducing the scale of the system tends to decrease the overall system performance simply because you are bringing the signals and their associated fields closer together and, in doing so, tend to increase the risk of higher crosstalk. Smaller scale also usually adds to complexity as there is less space for physical considerations such as cable size, bulk and routing. The use of smaller wire sizes in raw cable will address the size, bulk and

routing issues but the higher loss associated with the smaller wire gage reduces the maximum cable length that can be achieved.

One of the ways to combat the additional signal loss is to integrate some higher-level signal conditioning in either the cable assembly or the port PCB inside the equipment. The addition of this signal conditioning allows the system support longer cable length capability for a given wire gage. This added functionality does, however, have its trade-offs. Signal conditioning requires power to be bussed to the cable, which is not always required for passive cables. The signal conditioning is applied via silicon which is terminated inside the cable assembly. The inclusion of the silicon obviously increases overall cable assembly costs. This subject is covered in more detail in section 2.8 below.

Maximization of port density has led to the development of multiple port I/O product offerings. The SFP and SFP+ connector and cage offerings include ganged (single row, multiple columns) and ganged & stacked (double row, multiple columns) product variations. While these offerings address the port density, the ganging and stacking configurations do complicate the header connector and cage designs. Proper signal integrity, particularly in stacked connector versions, needs to be carefully considered as this can be a challenging task with multiple channels each running at 10Gb/s in the same connector system. In fact, these connectors start to look more like high speed backplane connector systems in construction & geometry vs. typical I/O connectors. It also means that they must be designed and developed using methodologies very similar to those used for backplane connector design.

The PCB termination of the board I/O connector from performance and manufacturing process standpoints also need to be taken into account. The corresponding cage designs may look fairly straightforward but they also require specific design features that assure mechanical robustness and guidance, ease of assembly to the PCB, proper EMI shielding and termination to the PCB and chassis and account for proper heat dissipation and application of heat sinks when required.

Multi-port connector and cage product extensions are expected for both the QSFP and CXP I/O systems. In both of these cases, the prospect of having multiple I/O channels, all running simultaneously at 10Gb/s in the same interface in a stacked configuration, while assuring proper signal integrity, will indeed be a challenge. This is particularly true regarding the CXP system where the PCB board density is maximized.

2.4 Port configuration and system recognition

In many communication equipment practices, the definition and configuration of an I/O port in a system is usually defined by the module, whether optical or electrical, that is placed into that port. Starting with the old GBIC modules, when you plugged a copper-based module into an open GBIC port, that port was then configured as an electrical port although the user did need to confirm the module/cable being used was of sufficient length. If the user plugged in a fiber optic module the port then had an optical functionality. In either case, the port functioned properly as long as the proper copper or fiber optic cable length was being used. This plug-n-play simplicity for port configuration has carried over to high speed I/O systems today and is being further enhanced in the emerging systems. The collaboration between system designers and I/O system suppliers, referenced earlier, has led to enhanced functionality in the cable assemblies that provides for automatic system recognition of the cable assembly and its inherent capabilities once the cable/module is mated into

an open port in the system. In the case of the SFP+ cable assembly system this is being accomplished with the addition of an EEPROM in the cable whether it's an optical or copper based cable assembly. The same approach is being followed with the QSFP+ systems as well.

The information defined in the EEPROM "map" includes vendor name, vendor part number and revision, date code, cable length, raw cable type, copper or optical based cable, active or passive, industry standard compliance (Fibre Channel, Ethernet, SONET, etc.) and at what speeds there is compliance. There are also discussions about the inclusion of other cable characteristics such as cable loss values but these have not yet been defined as being required in the map.

There are a number of compelling reasons for the inclusion of this information in the EEPROM. Aside from the obvious benefits of traceability for both the user and the manufacturer, the information allows the system, once the cable is plugged in, to download the information and to understand the full functionality of the cable that has been plugged into the port. Depending upon the design and intelligence of the system it may very well be able to recognize when an improper cable is plugged into a port. Another important aspect to the embodiment of this recognition functionality is that the system will be able to recognize whether the cable is a passive cable or an actively equalized cable once it is plugged into the system port. In the case of an actively equalized cable assembly, the system will then immediately know that it needs to bus power to the cable in order to ensure the cable functions properly.

This addition of EEPROM programming, while not particularly difficult, is certainly a capability that is beyond what a typical cable assembly manufacturer has normally been asked to supply in most cable assemblies. However, what becomes of greater importance, and is the clear responsibility of the cable assembly manufacturer, is the assurance that the cable assembly complies to the functional and performance values noted on the EEPROM. This assurance of compliance and interoperability is an aspect that goes directly to the "ease of use" and "customer satisfaction" that are major considerations for the industry standard organizations who want to assure that cables manufactured to their specifications meet expected performance levels. This is further clarified in section 2.11 below.

The finalization of the EEPROM data map for the QSFP interface is in its final stages and it is fully expected that similar maps / functionality will be part of the CXP and CFP interfaces.

2.5 Copper raw cable manufacturing and control

In almost any cable assembly the vast majority of the signal path is represented by the raw cable whether that cable is copper-based or optical fiber. Given this fact it is easy to conclude that the performance and consistency of the raw cable, from lot to lot and reel to reel, is of paramount importance to the final performance of the cable assembly. The same level of scrutiny must be applied to the raw cable manufacturer's processes and the systems that are used to maintain and monitor product consistency. Process control and consistency will continue to take on even greater importance as the market continues to demand faster signal transmission speeds. The reality is, at these signal speeds, any small changes in the cable geometry, whether associated with dielectric characteristics, shield wrap, shield coverage or dimensional stability, can and will have a direct and potentially detrimental affect on the raw cable performance. A tight and highly controlled

manufacturing process is an absolute necessity for any raw cable supplier who wants to participate in the high speed cable market.

Raw cable suppliers have been challenged to develop and manufacture high quality, high speed cables that are characterized by smooth linear insertion loss curves out through 10GHz frequencies and beyond. In other words, development of raw cables that are devoid of small insertion loss "suck-outs" (resonances) out to the 10GHz threshold. Cable suppliers have attacked this challenge in a number of different ways: solid dielectric extrusion, tighter shield wrap, heat sealed shield wraps, and more - with varying degrees of effectiveness. Regardless of the approach, the key is to make sure the raw cable manufacturing system is properly defined and controlled so the process does not introduce periodic or infrequent variations to the cable geometry or construction. Small variations can introduce flaws that affect cable performance and minimize overall raw cable effectiveness.

For cable assembly manufacturers who purchase this raw cable, the challenge is clear. Find and choose raw cable suppliers who clearly understand what it takes to manufacture high speed raw cable. Do these suppliers have the proper process controls in place? Do they manufacture and extrude in environmentally controlled environments? Do they test the raw cable sufficiently in order to assure a high quality product? Do we have the proper vendor assessment, purchasing, incoming inspection and manufacturing controls in place in order to assure the raw cable we purchase is the raw cable we expect it to be? Is our manufacturing process optimized for the raw cable construction and manufacturers we have chosen? Does the choice of this raw cable manufacturer lead us to the most cost effective cable assembly manufacturing process?

2.6 Wire management, termination, control & repeatability

While the raw cable quality and consistency are important considerations for any high speed cable assembly, they are hardly the only concerns. The subsequent process steps of cable preparation, wire stripping, wire management, wire termination and wire and cable strain relief all play an equally important role in assuring that the final cable assembly is fully functional and meets all expected performance levels. As noted previously, common high speed digital signal speeds today are in the 2.5Gb/s to 5Gb/s range but are quickly being driven to the 10Gb/s and higher data rates. At these transmission speeds, like with the raw cable, even very slight geometric changes from process variations in the cable assembly manufacturing will have an affect on the fidelity of the signal.

Ideally, if one could actually accomplish the feat, a cable assembly manufacturer would maintain the raw cable geometry up to the connector interface thus minimizing any interruption of the physical structure of the raw cable geometry. In practice, this is a very difficult task to accomplish when a cable assembly end is being terminated. The cable must be stripped of the outside jacket, EMI shielding removed, twin-axial wires managed into proper signal and termination locations, then stripped to expose the wires for termination to either a transition printed circuit board or directly to the electrical contacts in the connector. It is an advantage to achieve this termination transition in the smallest length possible. The termination transition area is measured from where the outside jacket of the cable is removed to the signal wire termination point. Lack of proper definition and attention to these details can easily render a cable assembly useless or, at a

minimum, will require a complete re-termination of at least one cable end. As with the raw cable this sensitivity to small geometric changes will only increase as signal speeds continue to be pushed higher and will continually challenge the cable assembly manufacturers.

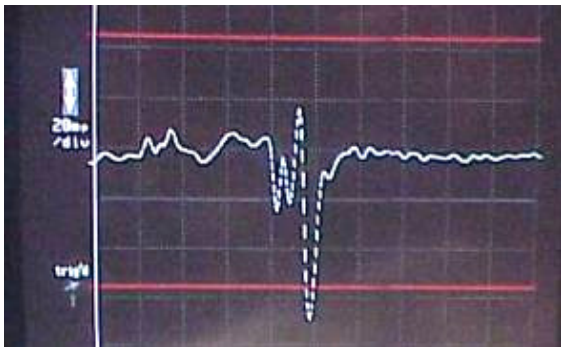
Several examples are shown below to illustrate and to quantify this sensitivity. Examination of the left hand side of figure 1 shows an excessive amount of solder being applied to the uppermost cable wire termination solder pad.

Figure 1.



Generally, with many lower speed cable assembly manufacturing processes, this amount of excess solder, while perhaps a point of concern from the perspective of process variability, will likely not render the cable unfit for use. In the illustration below (Figure 2), the subsequent TDR trace of this termination area shows the characteristic impedance in this termination area. The trace shows that the characteristic impedance in the termination area falls below the lower impedance limit (in this case the lower limit is 90 ohms) and outside the stated functional impedance range of the cable assembly. Further complicating this quality issue, is the fact that the repair of this termination area, specifically the removal of the excess solder, will not be a trivial matter. In fact, it may require a complete re-termination of this cable assembly end.

Figure 2



Another possible process variation is illustrated below in Figure 3. In this example the wire that is terminated to the second uppermost solder pad is slightly longer than the remaining wires that are terminated to the printed circuit board. As in the previous example, the process variation illustrated here is very slight.

The subsequent TDR trace for the signal in question is illustrated below in figure 4. Again the trace shows that the characteristic impedance is at or slightly below the lower limit allowed. We may find that this cable assembly is (barely) functional but this is certainly not the “process window” any reliable cable assembly manufacturer wants to be operating in.

Figure 3






Figure 4



Wire termination is not the only possible process variation flaw. Figure 5 below shows a variation in the length of the wire strip and shield strip process with some poor wire management in the transition from the wire strip area to the wire termination area. Consequently, as indicated by the superimposed TDR trace, the characteristic impedance in the termination area easily falls outside the defined functional range (above the 110 ohms maximum). As before the process variations are slight and not necessarily indicative of a process out of control.

Figure 5



A final example of process variation that can adversely affect the cable assembly performance is shown in Figure 6 & 7. Even after a cable assembly manufacturer performs all of the wire

management, wire stripping and wire termination processes correctly but they must also encapsulate and strain relieve the terminated wires correctly. If a small amount of encapsulation material (Fig.6) is allowed to settle beyond the proper limit area it can drive the resulting impedance beyond the acceptable level (Fig. 7)

Figure 6





Figure 7



Given these examples, it should be obvious that even cable assembly manufacturers that have a captive raw cable supply or a close partnership / relationship with a raw cable supplier are not exempt from any of the process variations or possible performance pitfalls that have been noted. By now, it should be self-evident that a high quality manufacturer of high speed cable assemblies must have a highly controlled and repeatable process that minimizes, prevents or eliminates even small process variations. A cable assembly manufacturer who fails to properly control, monitor and maintain their processes will find it difficult to be a reliable and profitable supplier of high speed cable assemblies.

2.7 Cable routing, cable bulk and cable profile

Signal loss tends to be the limiting factor when determining the longest usable copper-based cable link. A simple solution to increase the useable cable length is to increase the wire gage size in order to compensate for the losses. However, moving to larger gage wire sizes is not without its drawbacks. While the signal loss issue is addressed, it also results in a larger cable size, larger cable bulk and more difficulty in routing the cable. Additionally the larger cable bulk presents a higher profile when laid in a sub-floor or overhead trays as is typically done in most data centers. This

larger bulk, weight and profile then act as a constriction to the air flow which can make the proper cooling of the equipment in a data center more difficult to achieve. It has been estimated that the effort to properly power and cool a data center represents 25-30% of the total operating costs needed to run a typical data center. This significant cost component obviously makes excessive cable bulk or smaller air flow profile a major consideration for data center managers and operational personnel.

Practically speaking, given the typical port densities in present equipment and the associated cable routing that is needed, the largest gage wire size for a typical 4 channel twin-axial I/O system, such as QSFP or mini-SAS/SATA, is generally 24 AWG. There are specialized applications where some atypical wire sizes (i.e. 23 AWG or 23.5AWG) have been employed but these applications tend to be very specialized. Due to the points made above the 24 AWG cable size is generally used only when absolutely necessary.

2.8 Signal conditioning / pre-emphasis / equalization

In general designers and IT personnel will continue to use copper-based cables for as long as they can. Copper-based systems still tend to be the most cost-effective interconnect solutions when compared to the total applied costs of moving to a fiber optic based system. They also tend to be systems that require lower levels of power vs. comparable fiber optic systems.

To prolong and extend the use of copper-based systems designers and system I/O solution providers have started to employ various signal conditioning techniques in an effort to extend the lengths of the useable copper cables. Generally two philosophies are employed: add the signal pre-emphasis / equalization via silicon chips inside the equipment (system PCB) or add that capability to the cable assembly. Either approach has its own set of advantages and disadvantages. Which path a designer or equipment supplier takes may very well depend upon the overall approach they take with their overall system designs.

Those suppliers who add the pre-emphasis / equalization capability inside the system know, by doing so, that each of the ports in that system will have the "built-in" capability of allowing for the longest copper based cables possible. The equipment user has the comfort of knowing that each port is fully functional as long as cables used are within the allowable cable range. One drawback with this philosophy is that there is a cost associated with adding these equalization chips to the system in order to enable each port. This is a sunken cost whether an I/O port requires the signal conditioning or not. Another drawback deals with the signal integrity. Since the chips are on the PCB within the system the signal conditioning is applied in the system at the system PCB level. Since the signal conditioning starts at the system PCB the signal will see some loss due to transmission through the system PCB, the board I/O connector and the cable I/O connector before it's actually transmitted through the raw cable itself. The preference would be to avoid the addition of these losses before the signal is actually transmitted in the cable.

The other application has the pre-emphasis / equalization capability added to the inside of the cable assembly itself. This allows the installer to match the chosen port with the cable capability and cable length that is needed thus minimizing the cost for the given port configuration. If the cable length needed is relatively short (generally 5-7 meters or less depending on acceptance criteria) then a purely passive cable assembly may be

perfectly acceptable. If a longer cable length is needed then the installer simply chooses an actively equalized cable that is manufactured and qualified to the cable length required. In other words, the installer adds the signal processing functionality to the port, and the associated expense of doing so, only when it is necessary. Actively equalized cable assembly assemblies can extend the useful length of copper-based cables to 20-25 meters depending on the performance criteria. This chip-in-the-cable approach minimizes any signal integrity losses as the equalization and conditioning of the signal takes place very close to the actual transmission across the raw cable. Some of the disadvantages of actively equalized cables include the need to supply power to the cable in order to drive the chips and the expense of using such cables. Careful consideration around the power required to drive the chips is required as excessive power needs can become a major issue in a large data center applications.

2.9 Maximization of copper cable length vs. minimization of cable bulk

While active equalization can maximize the useable length of a usable copper-based cable assembly, it also offers the potential to reduce raw cable size and bulk for a given cable length. For example, a purely passive cable assembly may have the capability of meeting all performance requirements up to a length of 7 meters with 24 AWG signal wires. By employing active equalization in the cable assembly, the signal wire size required can be reduced, depending on performance criteria, to 32 or 30 AWG wire size. The significantly smaller cable bundle may be desirable for a number of reasons: reduced cable bulk and weight, increased cable flexibility for easier routing, capability to fit more cables through a given opening and improved air flow.

2.10 Copper cable testing & validation

It is likely that the testing of the finished cable assembly has changed more than any other aspect of the cable assembly manufacturing process with the possible exception of added EEPROM programming. In the past, the exposure to continuity and hi-pot (high voltage) testing was usually robust enough to detect such defects as “opens” or “shorts” caused by poor manufacturing processes, mis-wiring or poor quality components in low speed cable assemblies. Now, however, as noted several times earlier in this paper, at higher signal transmission speeds, signal integrity is much more sensitive to component or process variations. This functional performance sensitivity is at a much higher level than can be reliably identified with continuity and hi-pot testing alone.

Manufacturers have attempted to address this need for a higher level functional testing with the application of “Eye Pattern” and/or Bit Error Rate Testing (BERT). The eye pattern testing, in very simple terms, produces a histogram of the rise and fall of signal traces overlaid upon one another that illustrates the level of variation of the signal transmission over the population of the recorded signal traces. In practice, this histogram forms an eye-like opening. The acceptance criteria, that defines the allowable variation within this eye opening, is defined by a template that is also known as a “eye mask”. Acceptance or rejection using this criteria is a straightforward measurement as the test judges whether any of the histogram’s traces fall inside the eye mask opening. If a trace violates the defined eye opening template, the cable is judged as being non-compliant within the expected performance and is rejected. BERT is a measurement test that transmits a known and well defined data bit stream through the

cable assembly with monitoring of the data received at the far end of the cable. Acceptance is judged by comparing the accuracy of the received data to the sent data to assure there are no “errors”. Detection of errors constitutes a cable failure.

These higher level cable functionality tests have proven to be an improvement over the more typical continuity and hi-pot tests. A subtle but important difference between the continuity and hi-pot tests and eye pattern test & BERT should be noted. Continuity and hi-pot tests measure specific electrical values and are intended to detect specific flaws (i.e. shorts or opens) in the cable assemblies. The eye pattern test and BERT, while measuring to known thresholds, actually try to characterize the capability of the cable assembly to preserve the integrity of high speed signal transmissions.

This testing approach, where the capability assessment of the cable assembly is considered a more important measurement than the measurement of a finite parameter, is being adopted by cable manufacturers and system designers for higher performance cable assemblies. In practice, users of high speed cable assemblies aren’t worried about the specific cross talk, insertion loss, return loss, BERT or eye pattern of a cable assembly but more importantly they simply want to know that the cable will be functional when it’s plugged into their system. Consequently I/O system developers are now looking at utilizing additional measurement techniques like Wave Distortion Penalty (WDP) and Voltage Modulation Amplitude (VMA) as ways to assess overall cable assembly functionality. Both the WDP and VMA tests take a more holistic approach in that they attempt to measure the overall cable assembly functionality relative to the overall allowable loss the system can tolerate. In other words, rather than focus on the measurement of a series of individual parameters like crosstalk, insertion loss, etc. these new procedures look at the overall cable functionality and judge that functionality against the allowable loss for the given cable length.

WDP and VMA values are calculated using a prescribed data stream and measured S-parameter data and then converting that measured data via a defined code / script into specific values. These calculated values are then compared to the allowable WDP and VMA loss for a given AWG wire size and cable length to determine if the cable assembly is functional. The code / scripts used for the calculations are expected to be standardized via the applicable industry specification documents in order to assure proper and consistent application. Industry standards organizations such as IEEE, Infiniband Trade Association, Fibre Channel (INCITS T11 committee), SFF committees, etc. are striving to make the functional definitions and measurements of the cable assemblies as similar as possible.

2.11 Certification / interoperability

As with any new technique or application, it usually takes time to work out all of the “bugs” and address all of the issues that arise from first implementation of those techniques. This is no different in the case of the WDP and VMA test methodologies. Recent “plugfest” events sponsored by industry organizations have illustrated that work remains to be done before final definition of these testing techniques is completed and universally applied.

The ideal situation is perfect correlation between the testing for cable certification to a defined specification and the testing for the interoperability of that same cable assembly in a system across all applicable data rates and cable lengths. Unfortunately, today, this isn’t always the case. There have been instances where cable

assemblies pass certification testing but then fail interoperability testing. In other instances the opposite occurs. Cable assemblies passing interoperability but not certification. There are similar discrepancies when one compares the results at different data rates (e.g. at 5Gb/s per signal pair and 10Gb/s per pair). Further complicating this testing is the fact that all signal pairs should be tested in the cable assemblies. For example, a cable assembly can have only one signal pair that may fail either certification or interoperability testing but this results in the entire cable assembly as being judged non-compliant. The cable assembly supplier who tests only selected signal pairs in a cable assembly is placing a heavy reliance for the final product quality on the stability and repeatability of their manufacturing process. There are those who would consider taking such a testing philosophy as being a risky one without corresponding quality data and performance to back it up.

Some organizations are currently reviewing their plugfest data with the ultimate goal of understanding the reasons for the variations and the goal of achieving correlation between certification and interoperability testing across all cable types and lengths at all applicable data rates.

3. Conclusion

There are a number of important points that have been covered concerning the present and future use of copper based high speed cable assemblies in future communications equipment designs and applications as well as how these cable assemblies must be manufactured and tested.

These include:

Signal transmission speeds and bandwidth demand are being driven by the video-rich applications and social networking applications with significant growth forecast for the future. The demand for faster signal speed transmission is also expected to continue.

There is a significantly higher level of collaboration between equipment designers, raw cable suppliers, component suppliers, and high speed I/O system suppliers that is necessary to properly address these market demands

Equipment interconnection systems and technology are driven by industry standards organizations as evidenced by high speed systems like SFP+, QSFP+, and CXP.

Higher port density will continue to be an important consideration for equipment designers and their customers

Flexibility in port configuration / definition and the recognition and acknowledgement of that port configuration by the system at the time of cable plug-up is an important added functionality for these high speed I/O systems

In light of the added functionality and increased signal speeds manufacturing of these cable assemblies is a more challenging proposition than ever before. The quality considerations for raw cable, PCB design, wire management, wire stripping, wire termination, and wire strain relief all must be carefully addressed and properly controlled

Dimensional and routing considerations for copper based and fiber optic cables are significant factors influencing air flow, thermal management, and energy efficiency for equipment in data centers and at service providers.

Techniques like active equalization can be used to reduce wire gage and raw cable diameter required for a given cable length or enable longer copper-based cable assemblies to be used.

The proper testing of the higher speed copper cable assemblies is a significantly more sophisticated operation than what has been done in the past. Cable assembly capability testing techniques like eye patterns and BERT testing are common and are now being augmented by testing methodologies like WDP and VMA although even these processes aren't yet fully developed.

4. Acknowledgements

The author would like to thank Danny Morlion, Scott Kleinle, and Dave Sideck of FCI Electronics for their support, time, constructive feedback and helpful suggestions that have helped to make this document a more comprehensive and complete document. He'd also like to thank his wife, Barbara, for her editorial skills and suggestions.

5. References

- 1- Cisco Visual Networking Index Forecast & Methodology, 2007-2012, June 16, 2008 Cisco press release
- 2- Infiniband Trade Association Roadmap, 2005-2011+
- 3- IEEE802.3 Higher Speed Study Group meeting, April 2007 Interim Meeting, April 17-19, 2007, Ottawa, ON, Canada – Robert Hays –Intel; Howard Frazier - Broadcom

6. Appendices



Jim David is Global Product Manager for high speed cable assemblies and connectors for FCI Electronics. He is in his 29th year of working for FCI Electronics / Berg Electronics / DuPont Electronics. He has held positions in product engineering, field application engineering, field sales, engineering management and business management during his career with FCI / Berg / DuPont. He has a BSME from University of Massachusetts-Lowell and an MBA from The Penn State University.