The intricacies of signal integrity in high-speed communications

By Rick Zarr
Member Group Technical Staff, Signal Chain Applications

Introduction

As communication rates continue to increase, data is being moved within systems at ever higher speeds, which leads to issues with how engineers design equipment and printed circuit boards (PCBs). In the past 15 years, the industry has seen interconnection speeds increase from 1 Gbps to over 28 Gbps. Modern processors may have over 100 lanes of PCI Express® (PCIe), each running at 8 Gbps. The amount of raw throughput of systems today dwarfs the throughput of just 10 years ago. Back then, signaling rates were slow enough so that signal integrity was somewhat unaffected by passive interconnections. Today, that is no longer true. In most cases, transmission lines are now absolutely part of the circuit.

Many engineers have never had to consider what happens to the signals routed between integrated circuits. It was more of a challenge of cabling between systems, which may have required physical layer devices or special wiring. With field-programmable gate arrays (FPGAs) having increased capability and processors routinely using very high-speed interface standards such as PCIe, the system design must take into account the effects of PCB materials, transmission line design and connectors.

With the ever increasing speed of communications, the need to understand the effects caused by board layout, connectors and other parasitics is more important than ever. This article addresses the issues with high-speed signals and how to mitigate problems through proper component selection and board layout. Included in the discussion are standards such as SAS, SATA and Ethernet, as well as data converter interfaces such as JESD204B and other high-speed standards.

Figure 1. Example signals showing various transmission-line effects

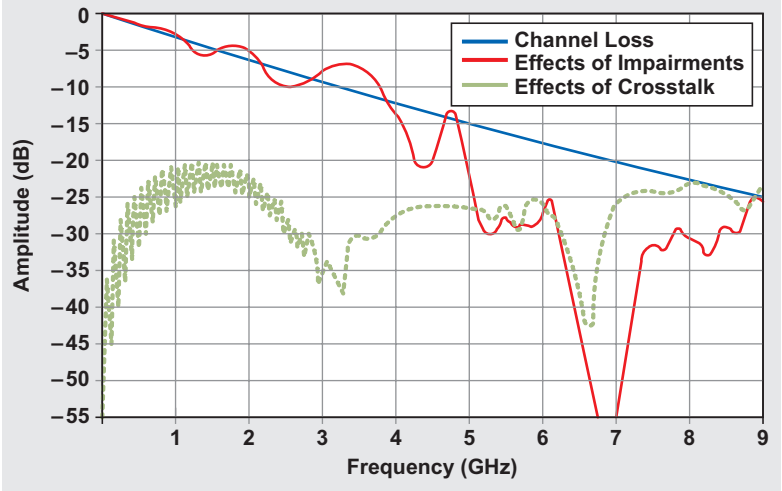

Signal integrity basics

All signals transmitted through a medium such as a PCB trace or coaxial cable are affected by various physical phenomena. As the signal's frequency content increases, these effects become more significant. For example, board capacitance, which is nominally very low (in the picofarad range), has little to no effect on low-frequency signals. But as the frequency increases, this capacitance begins to form a path to other signals or ground planes on the PCB. Other effects include impedance changes that cause reflections. These reflections affect the transmitted signal amplitude in ways that can appear random, but are actually deterministic.

At the highest level, loss of signal-to-noise ratio (SNR) directly affects the channel's capacity to carry information. Equation 1 shows this capacity relationship as presented by the Shannon-Hartley theorem. As SNR declines, channel capacity also declines. So anything that degrades the signal (assuming the channel noise remains constant) affects the channel's ability to carry information error free.

$$C = BW \times LOG_2\left(1 + \frac{S}{N}\right) \quad (1)$$

where C is the channel capacity in bits per second, BW is the bandwidth of the channel in Hertz, S is the average signal power, and N is the average noise power.

Transmission loss is caused by a combination of things (Figure 1). These effects include linear loss (blue line), impedance discontinuities such as connectors or vias that cause peaks and valleys on the linear-loss line (red line), as well as coupling between signals known as crosstalk. All of these effects contribute to the channel insertion loss, also known as S21 or S12 scattering parameters. The most dominant loss is the bandwidth limitation of the PCB traces or cables.

Further, this bandwidth limitation is a function of the dielectric loss, which is directly proportional to frequency, and the skin effect, which is proportional to the square-root of the frequency. Skin loss is a phenomenon where

the higher-frequency current flows on the outside of the conductor. This causes the effective resistance to increase with frequency. Additionally, the dielectric loss of the insulating material varies with frequency, thus adding complexity to the actual absorption characteristics.

To solve the insertion-loss issue, designers have several choices. The first and most straight forward is to select a low-loss, high-performance dielectric material for their PCB. Examples of high-performance materials are NY9000 or MEGTRON 7, which have excellent loss characteristics. These materials are polytetrafluoroethylene-based (for example, Teflon®) with extremely low loss, but are fairly expensive and may have limited availability from fabricators. Lower-cost alternatives will require some understanding of the channel characteristics to make the best selection.

Alternatives to expensive dielectric materials are to use active equalization on the receiver side or the addition of pre-distortion (also known as emphasis) to the transmitted signal to compensate for the linear loss. In effect, both circuits boost the high-frequency components of the signal that are absorbed by the transmission line, which results in a somewhat flat response in the required bandwidth. These devices are often protocol/physical-layer dependent, such as the DS80PCI402, which is specifically engineered to equalize PCIe (gen 3) channels.

Another effect that can degrade a channel is jitter, or the uncertainty of when the crossing point of the signal's bit state will occur at the receiver. In an ideal transmission line, there would be no impedance discontinuities to create reflections. In reality, the slightest imperfection in the PCB or simply routing a transmission line across an isolation boundary will result in discontinuities and reflections. These imperfections add up as the signal propagates down the transmission line before it reaches either a terminated receiver or a connector. Connectors are notorious for issues with mismatching impedances simply due to the mechanical structure and how it interfaces to the transmission lines.

Jitter effects can be divided into two major categories: deterministic and random. Deterministic jitter is further divided into periodic jitter, data-dependent jitter, and bounded uncorrelated jitter. Periodic jitter is caused by things such as switching power supplies and clocks leaking into the signal. Data-dependent jitter is caused by sources such as inter-symbol interference (ISI) or duty-cycle distortion (DCD), which affects when the receiver detects the state of the bit. Bounded uncorrelated jitter is a function of mostly layout where channels are placed close to one another. This is also called crosstalk and occurs when an aggressor signal leaks into a victim channel. All of these types of jitter can be corrected mostly through either equalization at the receiver, pre-distortion at the driver, or a combination of both.

Random jitter caused by various effects such as thermal noise will continue to increase within a given sample time, also referred to as "unbounded" jitter. Random jitter becomes an issue as transmission rates increase where the period of each bit, also known as the unit interval (UI), decreases, which allows less time for a bit to stabilize. Equalization cannot remove this type of jitter and, if severe enough, will require a re-timer device such as the DS100RT410 designed for 10-Gbps Ethernet to "re-clock" the data. The effect of re-clocking removes all jitter with a slight penalty in delay through the device. Re-clockers cannot be used on protocols such as serial AT attachment (SATA) or serial attached SCSI (SAS) where out-of-band (OoB) signaling is used to train the channel drivers and receivers (more on this later).

Protocol dependencies

With the introduction of serialized physical layers and protocols, standards have emerged that simplify system interoperability. Most serialized standards include the physical-layer electrical properties. It may be based on other standards such as current-mode logic (CML) or low-voltage differential signaling (LVDS) electrical specifications. This will dictate the drivers and receivers as well as the type of termination used. Other parts of the standards include a link or transport layer with specific protocols for training the channel or synchronizing the data fields. This can affect which types of signal-integrity components can be used to improve channel performance. Some standards include the physical connectors such as small form-factor pluggable (SFP) and quad SFP (QSFP), which also dictates the interconnect's form factor.

For most high-speed interconnections, using advanced board materials or active repeaters works to improve channel performance. Some protocols, however, actively train the channel due to variability of the transmission lines. An example is PCIe, which has connectors for

adapter cards with no knowledge of the channel characteristics until the adapter is running. For example, in a server, the communications adapter cards are normally placed at the back of the chassis along a riser card that plugs into a back-plane (Figure 2). The adapter card, known as the end node (EN), and processor side contains the root complex (RC). During initialization of each lane, the EN talks with the RC through a series of lower-frequency exchanges on the very channel that carries the high-speed data. This is referred to as out-of-band (OoB) signaling since the communication takes place at a lower rate than the normal transfer of data.

For passive channels this is not a problem, but active devices used to improve signal integrity cannot block these signals or the link will fail. For example, the DS80PCI402 allows OoB signals to pass, allowing the RC and EN to train normally. At both ends, the channel appears to have less loss.

Figure 2. Example server chassis with a back-plane and riser cards

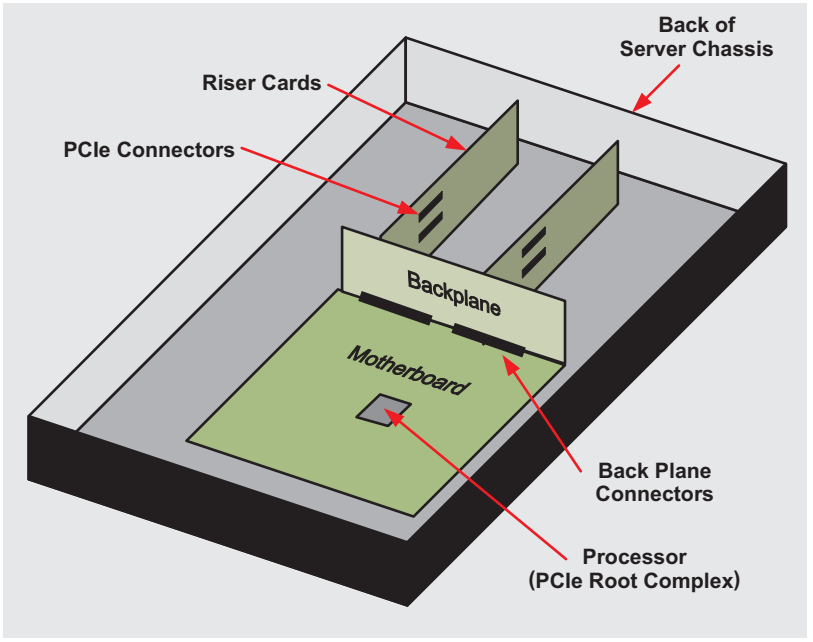

The alternative is to completely recreate the EN and RC somewhere in the middle, effectively splitting the transmission lines into two independent PCIe channels. This method can add significant delay (30+ UIs) in each direction, resulting in a decline of performance during short bus transactions such as reading a single location in memory. Devices are available to implement this solution, however, they can also draw significant power since the device must completely recreate the root complex and end-node per PCIe lane (in both directions).

Some protocols are simpler and have no ability to train the channel dynamically. These are used in fixed point-to-point applications, such as running 10-Gbps Ethernet inside of a router to quad small-form-factor pluggable (QSFP) connectors. The length of the transmission lines are fixed and known. Issues may arise due to strict jitter specifications of optical modules that plug into these connectors. If the jitter specification cannot be met passively, then a buffer repeater can be used to adjust for the transmission loss before the connector. If this is still problematic due to excessive jitter, a re-timer can be inserted just prior to the connector at the edge of the chassis. The re-timer works in this scenario because no OoB signals are present. The great advantage of the re-timer is twofold: first it can recreate somewhat jitter-free data at the output; and second, many re-timers include "eye" monitors that can sample the connection's signal integrity. This could be useful on high-end routers that want to monitor the health of a channel without running bit-error rate (BER) tests, which require the connection to be taken off-line.

Protocols beyond data communications

Many interfaces such as SATA, SAS, Fiber Channel, Ethernet, PCIe and a host of others are used for storage or connectivity applications. These are all dedicated to moving bits throughout a computer or enterprise. There are, however, other interface standards such as Common Public Radio Interface (CPRI™) or JESD204B that are designed to move data between specific functions. For example, JESD204B was designed to simplify the interconnection of high-speed analog-to-digital converters (ADCs) and digital-to-analog converters (DACs). It has both a physical layer (CML) and several protocol layers (such as link and transport layers). Each lane is designed to be "one way" with a transmitter and receiver side. That is, there are sets of lanes for the ADC side (such as ADC to digital processing element) and another set of lanes

used to transmit to a DAC (such as digital processing element to DAC) along with a set of clocks and control lines (Figure 3). Line rates for the standard have a maximum of 12.5 Gbps, but future versions (for example, revision C) will be faster, on the order of 16 to 25 Gbps.

In direct-sampling RF applications, it can be advantageous to place the data converters near the edge of the PCB, while having the FPGA or processor located closer to other digital elements. The JESD204B standard inherently does not have specifications for driver de-emphasis or receiver equalization, but implementations may have these features. If the built-in signal-integrity features are insufficient because of longer trace lengths or discontinuities, then either improved board materials or an active solution can be applied to improve channel performance.

In the case of an active solution, there are two issues to consider. First, if an application requires phase coherency across multiple data-converter devices, deterministic latency is required (or system calibration) for you to know when all information from the coherent samples has arrived at the receiver. Adding any delay in the path also delays when the receiver can release the data to the processing element, if not symmetric. The second consideration is actually a benefit in that the JESD204B standard does not require OoB signaling, so component selection can include re-timers if channel jitter is excessive.

Figure 3. Example JES204B system showing high-speed lanes and clocking

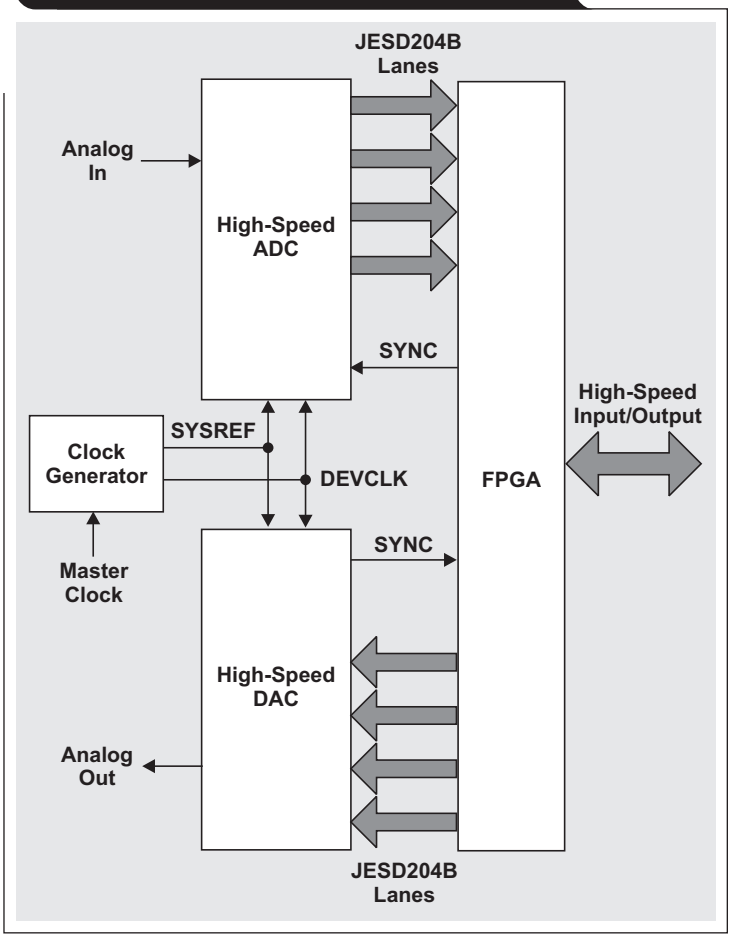

Conclusions

As data rates continue to increase, there is a growing need for engineers to understand the impact of signal integrity in their designs. With careful material selection and layout, many of these issues can be mitigated. When cost, material availability, or path loss becomes an issue, then active solutions such as repeaters or re-timers can be employed to restore the channel's signal integrity.

References

1. "JESD204B interface," white papers, technical papers and blog series, Texas Instruments

Related Web sites

Product information:
DS100RT410
DS80PCI402

TI Worldwide Technical Support

Internet

TI Semiconductor Product Information Center Home Page
support.ti.com

TI E2E™ Community Home Page
e2e.ti.com

Product Information Centers

Americas Phone +1(512) 434-1560

Brazil Phone 0800-891-2616

Mexico Phone 0800-670-7544

Fax +1(972) 927-6377

Internet/Email support.ti.com/sc/pic/americas.htm

Europe, Middle East, and Africa

Phone

European Free Call	00800-ASK-TEXAS (00800 275 83927)
International	+49 (0) 8161 80 2121
Russian Support	+7 (4) 95 98 10 701

Note: The European Free Call (Toll Free) number is not active in all countries. If you have technical difficulty calling the free call number, please use the international number above.

Fax	+(49) (0) 8161 80 2045
Internet	www.ti.com/asktexas
Direct Email	asktexas@ti.com

Japan

Fax	International	+81-3-3344-5317
	Domestic	0120-81-0036
Internet/Email	International	support.ti.com/sc/pic/japan.htm
	Domestic	www.tij.co.jp/pic

Asia

Phone Toll-Free Number

Note: Toll-free numbers may not support mobile and IP phones.

Australia	1-800-999-084
China	800-820-8682
Hong Kong	800-96-5941
India	000-800-100-8888
Indonesia	001-803-8861-1006
Korea	080-551-2804
Malaysia	1-800-80-3973
New Zealand	0800-446-934
Philippines	1-800-765-7404
Singapore	800-886-1028
Taiwan	0800-006800
Thailand	001-800-886-0010

International	+86-21-23073444
Fax	+86-21-23073686
Email	tiasia@ti.com or ti-china@ti.com
Internet	support.ti.com/sc/pic/asia.htm

A021014

IMPORTANT NOTICE

Texas Instruments Incorporated and its subsidiaries (TI) reserve the right to make corrections, enhancements, improvements and other changes to its semiconductor products and services per JESD46, latest issue, and to discontinue any product or service per JESD48, latest issue. Buyers should obtain the latest relevant information before placing orders and should verify that such information is current and complete. All semiconductor products (also referred to herein as “components”) are sold subject to TI’s terms and conditions of sale supplied at the time of order acknowledgment.

TI warrants performance of its components to the specifications applicable at the time of sale, in accordance with the warranty in TI’s terms and conditions of sale of semiconductor products. Testing and other quality control techniques are used to the extent TI deems necessary to support this warranty. Except where mandated by applicable law, testing of all parameters of each component is not necessarily performed.

TI assumes no liability for applications assistance or the design of Buyers’ products. Buyers are responsible for their products and applications using TI components. To minimize the risks associated with Buyers’ products and applications, Buyers should provide adequate design and operating safeguards.

TI does not warrant or represent that any license, either express or implied, is granted under any patent right, copyright, mask work right, or other intellectual property right relating to any combination, machine, or process in which TI components or services are used. Information published by TI regarding third-party products or services does not constitute a license to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property of the third party, or a license from TI under the patents or other intellectual property of TI.

Reproduction of significant portions of TI information in TI data books or data sheets is permissible only if reproduction is without alteration and is accompanied by all associated warranties, conditions, limitations, and notices. TI is not responsible or liable for such altered documentation. Information of third parties may be subject to additional restrictions.

Resale of TI components or services with statements different from or beyond the parameters stated by TI for that component or service voids all express and any implied warranties for the associated TI component or service and is an unfair and deceptive business practice. TI is not responsible or liable for any such statements.

Buyer acknowledges and agrees that it is solely responsible for compliance with all legal, regulatory and safety-related requirements concerning its products, and any use of TI components in its applications, notwithstanding any applications-related information or support that may be provided by TI. Buyer represents and agrees that it has all the necessary expertise to create and implement safeguards which anticipate dangerous consequences of failures, monitor failures and their consequences, lessen the likelihood of failures that might cause harm and take appropriate remedial actions. Buyer will fully indemnify TI and its representatives against any damages arising out of the use of any TI components in safety-critical applications.

In some cases, TI components may be promoted specifically to facilitate safety-related applications. With such components, TI’s goal is to help enable customers to design and create their own end-product solutions that meet applicable functional safety standards and requirements. Nonetheless, such components are subject to these terms.

No TI components are authorized for use in FDA Class III (or similar life-critical medical equipment) unless authorized officers of the parties have executed a special agreement specifically governing such use.

Only those TI components which TI has specifically designated as military grade or “enhanced plastic” are designed and intended for use in military/aerospace applications or environments. Buyer acknowledges and agrees that any military or aerospace use of TI components which have ***not*** been so designated is solely at the Buyer's risk, and that Buyer is solely responsible for compliance with all legal and regulatory requirements in connection with such use.

TI has specifically designated certain components as meeting ISO/TS16949 requirements, mainly for automotive use. In any case of use of non-designated products, TI will not be responsible for any failure to meet ISO/TS16949.

Products		**Applications**	
Audio	www.ti.com/audio	Automotive and Transportation	www.ti.com/automotive
Amplifiers	amplifier.ti.com	Communications and Telecom	www.ti.com/communications
Data Converters	dataconverter.ti.com	Computers and Peripherals	www.ti.com/computers
DLP® Products	www.dlp.com	Consumer Electronics	www.ti.com/consumer-apps
DSP	dsp.ti.com	Energy and Lighting	www.ti.com/energy
Clocks and Timers	www.ti.com/clocks	Industrial	www.ti.com/industrial
Interface	interface.ti.com	Medical	www.ti.com/medical
Logic	logic.ti.com	Security	www.ti.com/security
Power Mgmt	power.ti.com	Space, Avionics and Defense	www.ti.com/space-avionics-defense
Microcontrollers	microcontroller.ti.com	Video and Imaging	www.ti.com/video
RFID	www.ti-rfid.com		
OMAP Applications Processors	www.ti.com/omap	**TI E2E Community**	e2e.ti.com
Wireless Connectivity	www.ti.com/wirelessconnectivity		

Mailing Address: Texas Instruments, Post Office Box 655303, Dallas, Texas 75265