BY SHANNON SAWYER • INTEL CORP
GREG LE CHEMINANT • AGILENT TECHNOLOGIES

Digital and microwave worlds converge in 10-Gbps-backplane design and test

GETTING USABLE SIGNALS THROUGH BACKPLANES IS FAR FROM TRIVIAL WHEN SPEEDS REACH 10 GBPS. BY PREDISTORTING AND EQUALIZING SIGNALS, DRIVER AND RECEIVER ICs ALLOW THE USE OF LOW-COST SUBSTRATE MATERIALS, BUT GOOD DESIGNS ARE NO ACCIDENT.

Although you won't find them everywhere, communications systems that operate at 10 Gbps have been available for several years and are becoming more common. Architectures include copper cables with lengths of several meters and optical fibers with spans of tens or even hundreds of kilometers. When someone proposes transmitting 10-Gbps signals across a backplane in which the path is less than 1m long, you may expect the task to be manageable. However, when you consider signal-density requirements and impose realistic cost constraints, a different picture develops: You can no longer view the backplane as just a pc board to route signals; you must think of it as a complex communication system.

The basic blocks of a digital-communication system include the transmitter, the channel, and the receiver. Whereas it may seem unusual to consider a copper trace on a pc board as a communication channel, the attenuation, signal reflections, pulse dispersion, and intersymbol interference that a signal encoun-

(a)

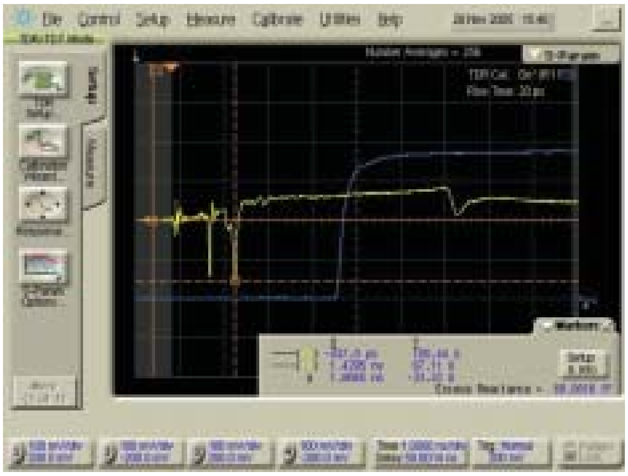

(b)

(c)

Figure 1 Backplane S_{21} bandwidth measurements compare the results of a VNA (a), a TDR trace (b), and the transformed TDR trace (c).

ters when traversing a 1m backplane rival the signal impairments in a 100-km optical link. To have any chance of successfully transmitting 10 Gbps through a backplane, you must understand and be able to overcome the degradation mechanisms that affect very-high-frequency signals traveling over and through common pc-board materials. The transmitted information may be digital, but the impairment phenomena are creatures of an analog—indeed, an RF/microwave—world.

The desire to produce 10-Gbps backplanes stems from the ever-growing need to send more information at greater speeds to more places. Backplanes serve as the intersection and traffic control for complex communication systems. Given the cost constraints on high-speed backplanes, the types of materials that manufacturers can use are limited to FR-4 and similar board dielectrics. The frequency-dependent loss is significant. That is, the higher frequency components of a data stream experience more attenuation than do the lower frequency components. The amount of loss is also proportional to the length of any trace. That is, the loss through a 1m board is twice that through a 0.5m board. This loss can almost destroy a 10-Gbps signal as it traverses the channel.

HUGE CHALLENGE

In addition to the severe degradation that the channel causes, a huge technical challenge awaits the designer of the transmitter and receiver chip sets. Backplane Ethernet applications are primarily enterprise-switch fabrics or computer servers providing 100 to 500 hot-swappable ports in a unit the size of a compact refrigerator. Such port densities require several custom ASICs, each having four to 200 embedded SERDES (serializer/deserializer) ports. These ASICs reside on eight to 12 densely populated, hot-swappable line cards that plug into the system backplane. This density and the requirement for a BER (bit-error ratio) of 10^{-17} demand per-port power of approximately 100 mW and ASICs that use 90-nm CMOS technology to achieve the requisite die size. Transmitted jitter can be no larger than a few picoseconds.

The backplane's limited performance is easiest to observe through a channel-bandwidth measurement. The result is effectively attenuation versus frequency. You commonly perform the measurement with a VNA (vector-network analyzer), which applies a known-amplitude sine-wave signal to the backplane, measures the output signal with a receiver tuned to the input-signal frequency, and displays the ratio of output to input. Designers commonly refer to the measured bandwidth as the S_{21} S (scattering)-parameter. The 2-1 notation indicates the ratio of the signal amplitude at Port 2 (the output) to that at Port 1 (the input). A TDR (time-domain reflectometer) can also measure the channel's output-to-input ratio, but instead of applying a swept-frequency sinusoid, the TDR applies a very-fast-edge step function as the input stimulus. To observe S-parameters, measurement software transforms the time information into the frequency domain (see sidebar "Going with the flow: S-parameters verify the performance of communications channels").

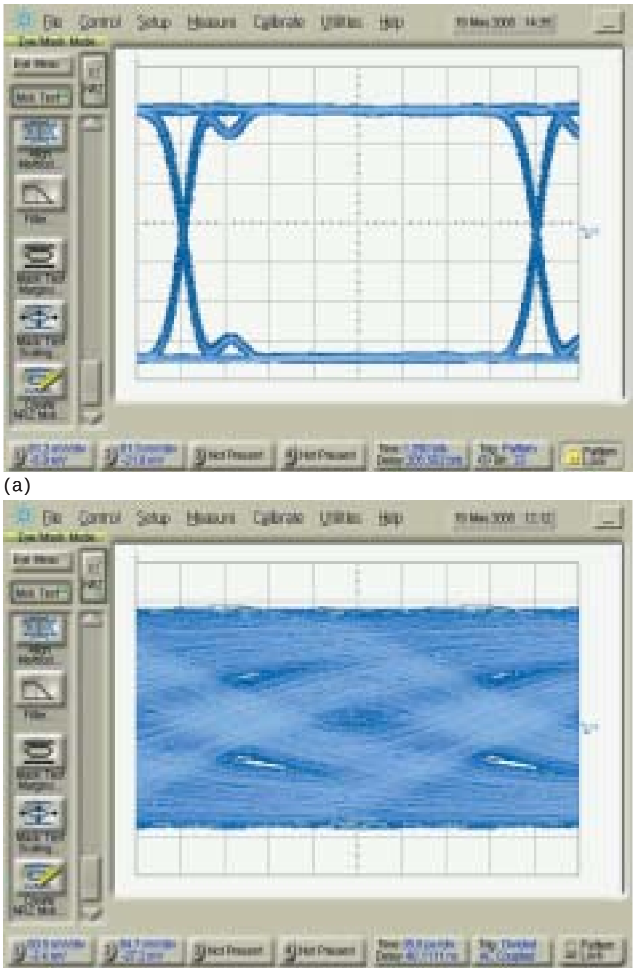

(a)

(b)

Figure 2 A 10-Gbps signal from a pattern generator (a) passes through a 40-cm FR-4 trace with devastating effects on the eye diagram (b).

Figure 1 shows the S_{21} measurements of a 1m backplane. Designers performed one measurement with a VNA and the other with a transformed-TDR result. Although the results from the two instruments are similar, as they should be, the important point to observe is what frequencies are significantly attenuated. At less than 1 GHz, the signal decreases by 3 dB or is at 50% power. At about 5 GHz, the signal has been attenuated 20 dB or is at 1% power. The frequency content of NRZ (non-return-to-zero) data at 10 Gbps, at least at the transmitter output, has significant signal content where the backplane has 10 to 20 dB or more attenuation. This attenuation dramatically alters any data signal that appears at the backplane output. This effect is easiest to see by using a wide-bandwidth oscilloscope to view the data eye diagrams.

It is obvious that, at a 10-Gbps data rate, the backplane bandwidth has a devastating effect on the eye diagram (Figure 2). The eye is completely closed. Without some help, a receiver's decision circuit will be unable to interpret the data and achieve a reasonable BER. Because making dramatic improvements to the backplane itself is likely to be expensive, the burden of system-level improvement falls on the transmitter and receiver. One critical aspect of the backplane's signal-degradation mechanisms is that they are systematic and predictable. This predictability opens the door to use of advanced communications schemes to correct the impairments. At least two possible approaches to this task exist. One is to predistort the transmitted signal—making the predistortion the inverse of the channel impairment. This technique is commonly known as de-emphasis. Another approach, equalization, corrects the signal impairment at the receiver. Ideally, predistortion at the transmitter or equalization at the receiver cancels out the channel's high-frequency attenuation. When the signal degradation is severe enough, you can use both transmitter de-emphasis and receiver equalization.

TRANSMITTER DE-EMPHASIS

Transmitter de-emphasis usually reduces the amplitude of the later portion of a logic 1 pulse and similarly increases the later level of the logic 0 pulse. Most predistortion schemes set the amplitude of the first one bit as the nominal amplitude (1V in 90 nm) and make the amplitude of the following bits adjustable. When driving long channels, the reduced amplitude of the trailing bits leads to reduced high-frequency content. For short channels, the amplitude of the trailing bits remains near the same amplitude as the leading bit. Note that when consecutive bits are identical—for example, a 0-1-1-1 sequence—only the first one following a zero has the boosted signal level. Adjacent ones have the unboosted nominal amplitude.

An intuitive basis for the effect of waveform emphasis is that high-frequency losses have a high impact on the edge of the signal and have little impact on the flat or nontransient regions of the data. Thus, the emphasis applies only to ones preceded by a zero and zeros preceded by a one. Figure 3 compares a normal NRZ signal and its response at the backplane output. It also includes an emphasized NRZ signal from a parallel BERT (BER tester) and its backplane output. Note how transmitter emphasis has pried open an eye that would be closed for a common NRZ signal.

To improve a signal through receiver equalization, you must create a network whose response is the inverse of the backplane's. The receiver response should then cancel the channel response. You create a feedforward, or linear, equalizer by tapping off portions of the incoming signal and recombining the tapped signals with the original. Each tap has a delay and gain associated with its path leading to the summing node. Common designs usually have delay values of some fraction of the bit period. You can amplify or attenuate the tapped signals, and the approach uses both positive and negative tap values. To understand how the tap structure can correct for the channel response, a time-domain view of the channel response, rather than the frequency domain or S-parameter view, is helpful.

Although it is hard to see in the eye-diagram display, the closure of the eye is due to the sluggish response of signals passing through the channel. The trajectory of any bit is likely to be influenced by many of the preceding bits. For example, a logic 1 preceded by several consecutive 1s is much more likely to reach its ideal amplitude than a logic 1 preceded by several consecutive 0s. A 1 preceded by 1-0-1-0 will have another unique trajectory. You can see how a given bit influences others by viewing its impulse or pulse response.

EQUALIZER DESIGN

The limited bandwidth of the channel leads to broadening of the impulse response (Figure 4). Note also that some ringing in the response in this case continues for several bit peri-

GOING WITH THE FLOW: S-PARAMETERS VERIFY THE PERFORMANCE OF COMMUNICATIONS CHANNELS

As digital-communications signals extend well into gigabit-per-second rates, you can no longer view the pc-board traces that carry those signals as simple connections from Point A to Point B. Signal wavelengths are shorter than traces, and transmission-line theory helps in understanding how these signals propagate. Measurement techniques that RF/microwave engineers have used for decades are now important to high-speed-digital-design engineers. In particular, S-parameter measurement techniques are becoming common as transmission rates increase.

S, or "scattering," parameters describe how energy flows in a network. One common S-parameter, S_{21}, where 2 represents the point at which energy exits the channel and 1 represents the point at which energy enters the channel, describes how well energy flows from one end of a channel to the other. S_{21} is the ratio of the output signal to the input signal. The measurement is a function of frequency. Thus, a power amplifier's S_{21} measures the device's gain versus frequency, from which you could obtain the amplifier's bandwidth. S_{21} of a cable could be a measure of attenuation versus frequency, because cables tend to have increased loss as frequencies increase.

It can be difficult to launch high-frequency energy into a network. It is not uncommon for energy to reflect back toward the original source. Keep in mind that a reflected signal can show up at unexpected places or times and can degrade communications quality. S_{11} describes this effect, in which energy travels into a component or channel port and the energy that reflects back is measured at the same port that it went into. S_{11} is the ratio of the reflected energy to the input energy. Generally, S_{11} should be low, indicating small reflection levels. S_{22} describes how reverse-flowing signals reflect off a device's output port. Designers have historically used network analyzers to measure S-parameters. Network analyzers use as the stimulus a sine-wave signal generator that can operate over a wide range of frequencies. A receiver the can measure the response–the output or reflected signal–makes up the other half of the instrument. Engineers now also use time-domain reflectometers to generate S-parameter information through mathematically transforming measured time results into the frequency domain.

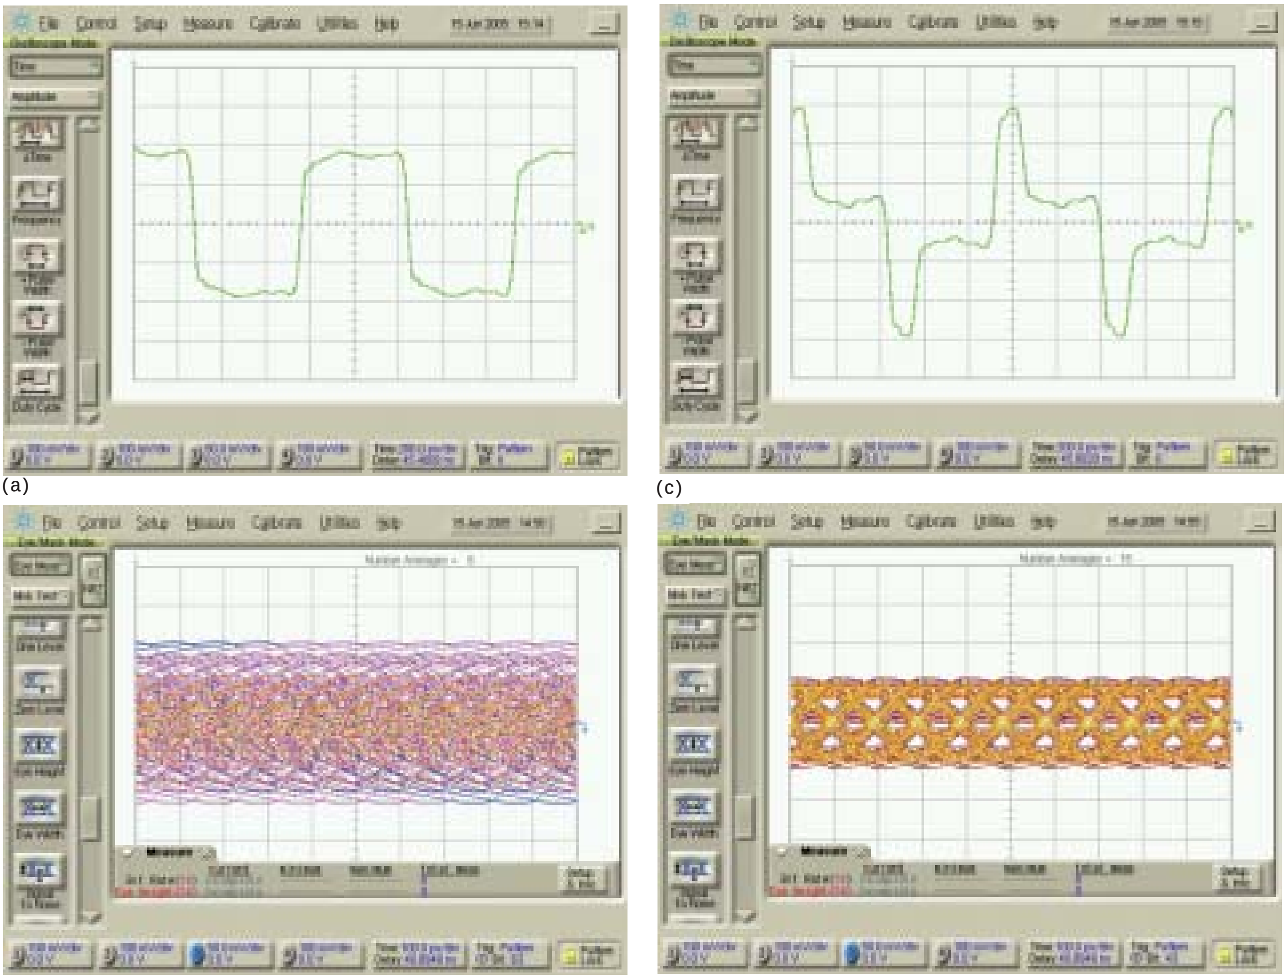

(a) (c)

(b) (d)

Figure 3 A parallel BERT has ordinary and de-emphasized waveforms (a and c) and the corresponding results (b and d) at the output of a 1m backplane channel.

ods after the initial pulse. The combination of a broadened pulse and ringing beyond the period of a bit leads to degradation of the bits that follow in the data stream, and you can observe this degradation in the eye diagram as eye closure. In the equalizer, signals from the taps complement the impulse response. For example, if an undershoot occurs after the main pulse half a bit period away, a tap of the same amplitude but opposite in phase and a delay of half a bit period negate the undershoot. Adding other taps compensates for the remaining components of the pulse response. A four-tap linear equalizer achieves a relative signal improvement (Figure 5). For mechanical robustness, connectors have through-hole drilling of the backplane board. This approach leads to a via-to-stripline stub—effectively, a short transmission line. The stub can lead to a channel-impulse response that ripples and does not settle out until as long as 15 bit periods. This approach has some significant implications on the receiver's design.

Some backplanes have S_{21} responses that indicate a relatively wide frequency response, whereas others have narrower bandwidth. You would expect the backplane with the wider bandwidth to provide the better channel. However, the channel S_{21} describes only one part of the overall channel characteristic. The channel also interacts with the transmitter and the receiver. Nonideal impedances can cause signal reflections and re-reflections that may yield a complex impulse response that a simple S_{21} measurement of the channel may not expose. The overall impulse response may have ripples that extend beyond those for which an economical equalizer or de-emphasized transmitter can compensate. Channels with poorer frequency response may have better interaction with the transmitter and receiver and achieve a better system response than wider-bandwidth channels. Thus, predicting when a channel will work and when it will not may be difficult to accomplish solely through frequency-response measurements.

A receiver with a linear equalizer and a transmitter with de-emphasis may be insufficient to compensate for a complex backplane. You can achieve additional signal correction with a DFE (decision-feedback equalizer) at the receiver, usually following a linear equalizer. The DFE dynamically adjusts the receiver decision threshold rather than tries to manipulate the incoming signal. Thus, as the logic 1s become low or the logic 0s stray high, the decision threshold shifts accordingly and corrects decisions. The DFE also has signal taps to determine where the threshold should be. The more taps the DFE uses, the greater the level of signal corruption for which it can compensate. However, because previous decisions affect future decisions, a mistake by the deci-

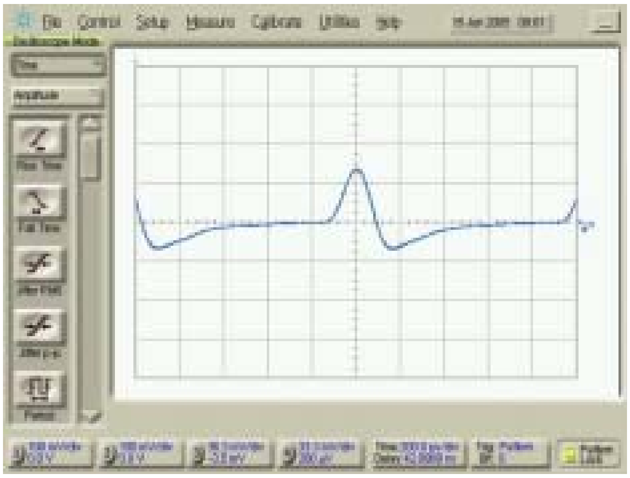

Figure 4 The impulse response shows some pulse broadening due to the limited bandwidth of the channel.

sion circuit can lead to several incorrect bits. The more taps that exist, the longer it takes in bits for a mistake to clear out. Thus, an impulse-response ripple as long as 15 bits can lead to an impractically complex receiver-equalization design.

At high data rates, it can be difficult to constrain all of the signal energy to propagate through just the connectors and board traces. Channels can and usually do behave as antennas and radiate some of the signal power. Similarly, the antenna effect also works in the receiver mode. That is, the antenna picks up and propagates stray radiated signals, along with the intended signals, to a receiver. Because most backplanes have a multitude of simultaneously operating channels, the issue of channel crosstalk is critical. Crosstalk-measurement schemes verify how poorly channels radiate, as well as how susceptible channels are to radiated signals. The issue of crosstalk is not new and is a problem even for older, lower speed backplanes. The biggest contributors to crosstalk in backplane applications are often the connectors. A 10-Gbps "aggressor" signal's power can approach a "victim" signal's power at the 10-Gbps signal's 5-GHz Nyquist frequency. You must test the receiver for interference tolerance to prove that it is robust in the backplane environment.

DIFFERENTIAL SIGNALING

Differential-signaling schemes can reduce the level of signal radiation. The fields from the complementary signals can cancel each other out and reduce the overall radiation. From the receiver perspective, when signals do invade a channel, the same signal is present on both lines of a differential channel. A receiver with high common-mode rejection can better deal with crosstalk signals. The challenge for the very-high-speed backplane is that, because of cost constraints, a larger burden falls on the receiver to correct crosstalk problems. New receiver designs are necessary to yield higher degrees of tolerance to invasive signals.

The system's required ultrahigh data integrity necessitates intelligent failure prediction and good debugging information. In response to these requirements, some ASIC vendors provide built-in channel-signal-integrity analysis and performance tuning that provide both insight into the quality of the signal that the receiver sees and the ability to tune the transmitter and receiver to quickly detect and repair marginal interconnects and voltage or temperature anomalies.

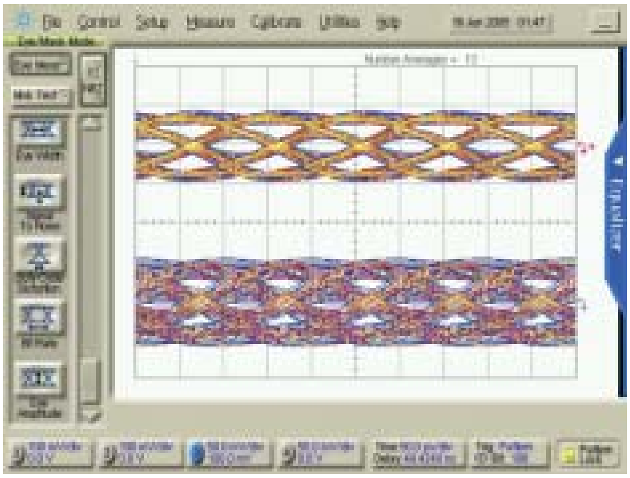

Figure 5 At the output of a 20-in. backplane, the upper eye diagram has passed through a virtual four-tap linear equalizer implemented within the digital-communications analyzer (oscilloscope).

Separate vendors may independently manufacture and test transmitters, connectors, pc boards, and receivers. This fact complicates the construction of a working design. Different backplanes may present a wide range of responses, and a one-size-fits-all design is likely to be insufficient for many channels. Trace geometries and layouts, material variance, and stub designs can lead to many levels of channel performance. Yet, system integrators expect the system to work when all of its components are assembled. The task of specifying the elements to guarantee a system-level specification can become difficult. Also, compensation schemes at the transmitter and receiver may need to be adaptive to allow usage with a broad range of backplane channels and connectors. Nevertheless, with several standards bodies tackling the problem, you can expect to soon see backplanes operating at 10 Gbps.

AUTHORS' BIOGRAPHIES

Shannon Sawyer is an R&D engineer at Intel's Fort Collins, CO, Design Center, where he works on signal and power integrity for the Itanium family of processors. Sawyer is also an active member of the IEEE 802.3ap 10G Backplane Task Force, Local IEEE Solid State Circuits Society, and Microwave Theory and Techniques Society. Before joining Intel, he worked for 10 years at HP/Agilent/Avago in high-speed-pc-board design and interconnect modeling and design, medical-instrumentation, and image-sensor ICs. He holds a bachelor's degree with honors in electrical engineering from California State Polytechnic University—San Luis Obispo.

Greg Le Cheminant has worked for Agilent/HP since 1985, first as a manufacturing engineer for microwave instrumentation and then as a marketing engineer. He is involved in the development of measurement applications and tools for high-speed digital communications. He also participates in several standards committees, including IEEE 802.3. He holds a bachelor's degree in electrical-engineering technology and a master's degree in electrical engineering from Brigham Young University (Provo, UT).