Controlled Intersymbol Interference Design Techniques of Conventional Interconnect Systems for Data Rates Beyond 20 Gbps

Wendemagegnehu T. Beyene, *Senior Member, IEEE*, and Amir Amirkhany, *Member, IEEE*

***Abstract*—This paper presents a new design technique of high-speed interconnects with controlled intersymbol interference (ISI) to create efficient signaling over a band-limited channel. Performance of high-speed electrical links is limited by conductor loss, dielectric dispersion, and reflections in the board, package, and connector. These nonidealities result in significant ISI. In current systems, the effect of ISI is either mitigated through complex equalization, signal conditioning, and coding techniques, or through costly impedance control and manufacturing processes. In the proposed approach, instead of eliminating ISI, we shape the response of the channel into a set of channel characteristics with controlled ISI using simple passive structures in the board and the package. The resulting controlled ISI is exploited at the transmitter and receiver to simplify the architecture of the system and to achieve high data rates. The techniques to design interconnects with controlled ISI are reasonably simple to implement in conventional interconnect technologies. Simulation examples are given to demonstrate the validity and advantages of the design technique using duobinary and analog multitone (AMT) signaling methods.**

***Index Terms*—Analog multitone (AMT) signaling, doubinary signaling, intersymbol interference (ISI), partial response signaling.**

I. Introduction

HIGH-SPEED data transmission over passive interconnects is required to support the high-bandwidth demand of current computing and communication systems. The bandwidth of current chip-to-chip and backplane links is limited by the bandwidth of the passive components (package, board, connector), and not by the operating speeds of the active circuitry in the transmitter and receiver. The interconnect systems are band-limited due to the loss and dispersive mechanism inherent to the interconnects and its surroundings. These nonidealities are more pronounced in low-cost packaging, printed circuit board (PCB), and connector technologies. The insertion loss of various length channels in conventional interconnect systems are shown in Fig. 1(a). The received signal energy in long channels such as those in legacy backplanes is very small when the operating frequency exceeds 5 GHz. For medium and short channels, the detectable signal energy is limited to the frequency ranges of up to 10 GHz.

In addition to the attenuation, the discontinuities due to vias, connectors, solder balls, and stubs severely limit the bandwidth of a channel. Some of these band-limiting structures are not part of the signal path and they exist merely because of manufacturing artifacts or mechanical requirements. For example, the via stubs in a conventional backplane and the plating stubs in a wirebond plastic ball grid array (PBGA) package do not serve any electrical purpose [1]. These via and plating stubs significantly reduce the operating bandwidth of the channel by introducing resonances, as shown in Fig. 1(b). As a result, the traces with long stubs show significant reduction in their bandwidth.

The increase in data rate over an interconnect is currently made possible by designing controlled impedance interconnect systems. In addition, the discontinuities created by connectors, solder balls, and stubs have been mitigated so far by careful compensation techniques [2]. However, as the data rate increases, the compensation effect diminishes because of the narrow band nature of the techniques and the increase in the overall channel attenuation. Improved packaging, connector, via technologies, and better materials can further improve transmission characteristics of the channel. However, the cost of the channel will increase significantly.

Another approach to mitigate the effect of ISI is the use of equalization techniques [3], [4]. Unfortunately, for low-cost interconnect technologies with long stubs, the frequency characteristics have nulls that can be problematic for simple equalization techniques. This requires complex on-chip digital signal processing circuitry. An alternative approach is to use multilevel signaling and pack the data into a small bandwidth [5]. Multilevel signaling, however, sacrifices the signal-to-noise ratio (SNR) and adds to the complexity of the timing recovery circuits and to the power consumption of the system.

An alternative to these techniques, which all try to eliminate ISI, is to control ISI and shape it to specific known patterns that can be exploited at the receiver for efficient signal detection [6]. In this alternative approach, the design goal of an interconnect system is to shape the response of the channel using trace and via stubs (contrary to impedance matching), so that the system has the desired characteristics to transmit higher data rates or has simpler transmitter or receiver. Signal transmission technique over controlled ISI channels are well studied in the literature and a number of algorithms exist to choose from [7]–[9].

In Sections II, the state of the art link systems with transmit and receive equalizers are briefly described. Then, two classes of signaling techniques over controlled ISI channel, partial signaling, and analog multitone (AMT) signaling are introduced in Section III. Controlled ISI channel engineering is presented in Section IV. In Section V, a design example of a high data rate signaling and system is presented. Implementation details and

Manuscript received September 11, 2007; revised February 19, 2008. First published May 16, 2008; current version published November 28, 2008. This work was recommended for publication by Associate Editor P. Franzon upon evaluation of the reviewers comments.

The authors are with Rambus Inc., Los Altos, CA 94022 USA (e-mail: wbeyene@rambus.com; amir@rambus.com)

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TADVP.2008.924224

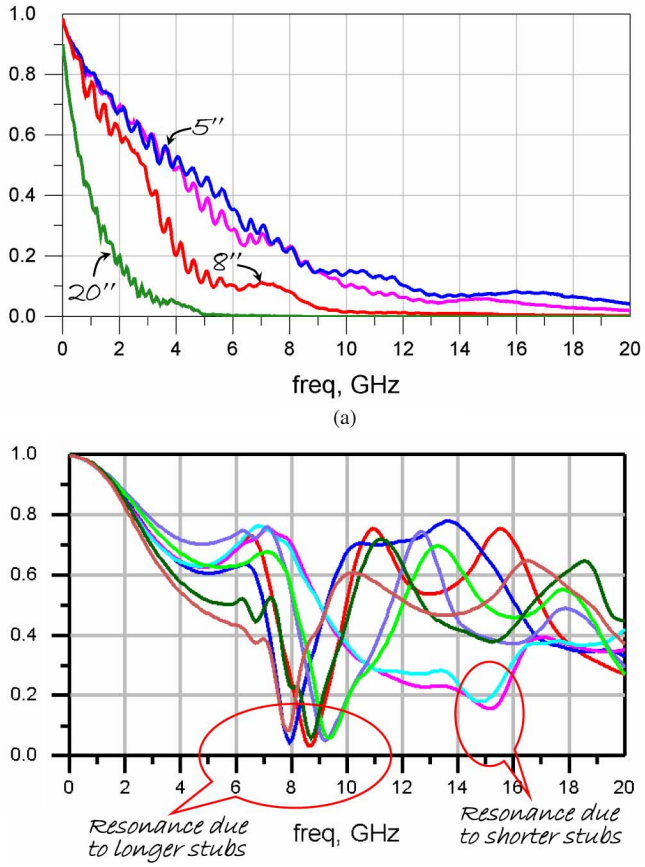

Fig. 1. Insertion loss of lossy and discontinues conventional low-cost interconnects. (a) Insertion loss of typical signal nets with varying length. (b) Insertion loss of wirebond PBGA package nets with 2.6 to 4.5-mm-long plating stubs.

comparisons of various signaling techniques including duobinary and AMT are discussed. Finally, the conclusion is given in Section VII.

II. State of the Art Link Design

In high-speed link design, equalization and signal processing techniques are used to mitigate the effects of ISI. Equalization can compensate channel frequency-dependent loss and dispersion of long traces in boards and packages, and dispersion due to device loadings [3]. Fig. 2 shows a block diagram of a state of the art link. Both linear and decision feedback equalizers often exist in current systems. The linear equalizer is a feed-forward equalizer (FFE) that uses linear filters with adjustable parameters to compensate for channel distortion. This can be implemented as transmitter preemphasis and/or receiver equalization as shown in Fig. 3. Although, receiver equalization has many advantages, the simplest and most cost-effective approach for a multigigabits per second parallel bus is transmit preemphasis. Transmit preemphasis compensates for the low-pass nature of the channel by preemphasizing the high-frequency components of the input signal, as shown in the transfer functions of the transmit equalizer of Fig. 3. However, the transmitter uses part of the signal amplitude budget to generate the preshaping symbols following the main symbol which leads to reduced SNR. The performance of systems with transmit and receive linear equalization can easily be determined using linear analysis in time or frequency domains. The preemphasis filter can be implemented using analog techniques by integrating the filter into each driver module as a parallel transmitter.

The decision feedback equalizer (DFE) is a nonlinear equalizer that employs previous decisions to eliminate ISI caused by previous detected symbols on the current symbol to be detected. The single-bit response (SBR) of a high-speed interconnect with DFE is shown in Fig. 4. The DFE is not able to cancel the precursor ISI as it is causal. Consequently, the DFE often needs to be paired with a FFE. DFE is the most effective way to cancel postcursor ISI because as opposed to a transmit FFE it does not reduce the transmit peak voltage budget and unlike a receive FFE it does not amplify channel noise. However, a major circuit design challenge in the design of the first (few) postcursor DFE taps is closing the timing for feedback loop in one (or few) unit intervals, as shown in Fig. 5(a). This problem is particularly more pronounced for the first DFE tap, as the received signal has to be detected, multiplied by the appropriate weight coefficient, and subtracted from the incoming signal, all in only one unit interval. A unit interval can be as small as 50 ps in 20-Gbps 2-PAM link. As a result the first DFE tap is often not removed in high-speed links, or look-ahead computation is used to unroll the loop and increase the delay in the feedback loop [10]–[12].

In one-tap loop-unrolled DFE, two decisions are made at each cycle. One comparator decides the received signal as if the previous received signal was a one, and the other comparator decides the received signal as if the previous bit was a zero. Once we know the previous bit, the correct comparator output is chosen. A one-tap loop-unroll DFE is shown in 5(b). The loop-unroll DFE makes two decisions on two conditioned eyes using samplers that are offset by the first postcursor ISI tap magnitude [10]. The lower and upper eye diagrams, are shown in Fig. 6(a) and (b), respectively. The two eyes are split by the amount proportional to the first post-cursor ISI tap. The timing constraint in a loop-unrolled DFE is to accommodate a flip-flop and a multiplexer in one unit interval, which could still be a challenge at 20-Gbps. The number of samplers required for loop-unrolling for more than one tap increases as $2^{\text{Number of taps}}$. As a result, loop-unrolling for more than one tap is generally avoided.

III. Controlled-ISI Link Design

The equalization techniques described in the previous section are targeting arbitrary forms of ISI. However, for the cases where the communication channel has a well-defined pattern, more efficient signal processing techniques are developed that lead to simpler transmitter and receiver architectures. In this section, we describe two of these techniques that are particularly relevant to the link design approach we present later in this paper.

A. Partial Response Signaling

Let us assume the ISI pattern in the communication channel is so severe such that the previous bit gets added to the current bit as the signal travels through the channel. In other words, the received signal at time n (y_n) is given by

$$y_n = x_n + x_{n-1} \tag{1}$$

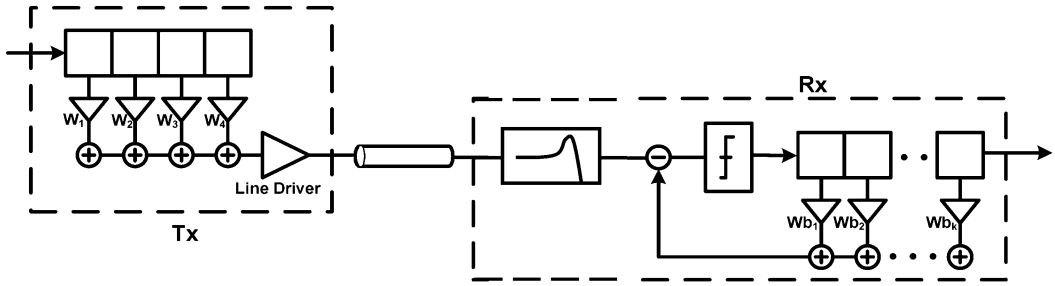

Fig. 2. Block diagram of a state of the art link.

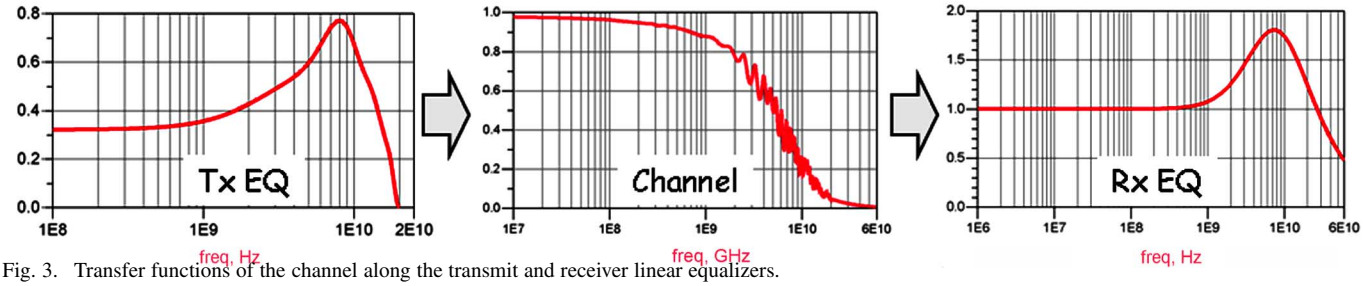

Fig. 3. Transfer functions of the channel along the transmit and receiver linear equalizers.

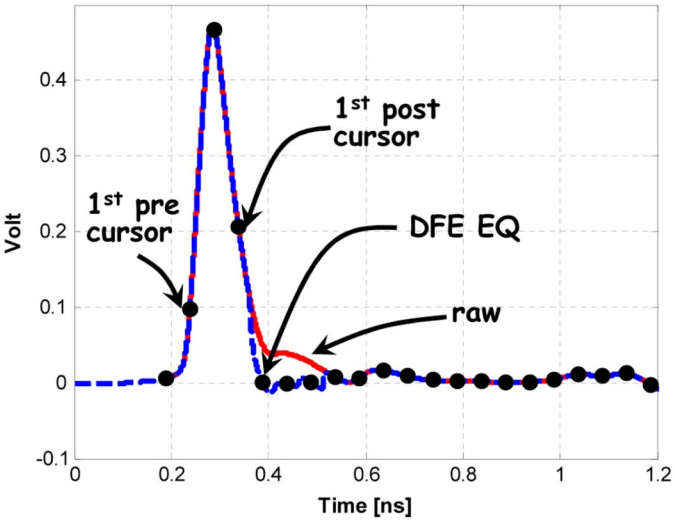

Fig. 4. Single bit response with 3-tap DFE applied at second, third, and fourth postcursors.

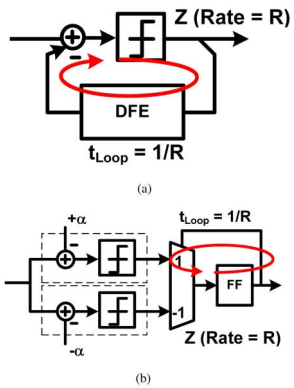

Fig. 5. Delay in the feedback loop of DFE. (a) Standard DFE. (b) Loop-unroll DFE.

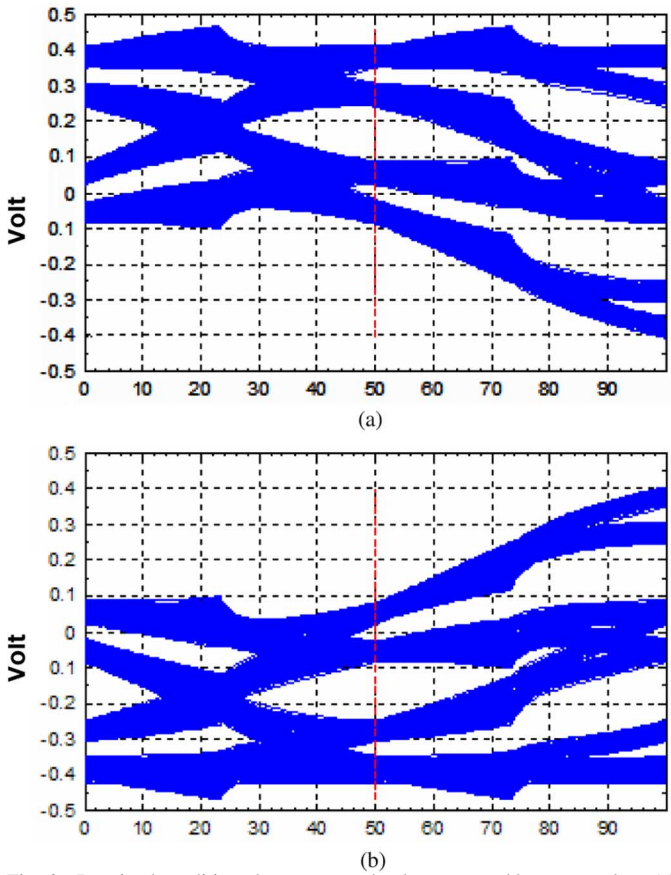

Fig. 6. Received conditioned eyes as seen by the upper and lower samplers. (a) Upper eye. (b) Lower eye.

where x_n is the transmitted symbol at time n. Then for a 2-PAM system, the received signal y_n can either be 0 or 2 if x_n and x_{n-1} are equal, or it is 1 otherwise. Therefore, if we know what x_{n-1} is, we can find out what x_n is using regular DFE or loop-unrolled DFE. Alternatively, we may perform the following simple precoding at the transmitter before transmitting the signal:

$$y_n = y_{n-1} \oplus x_n \tag{2}$$

TABLE I
CHARACTERISTICS OF SOME OF THE PARTIAL-RESPONSE SYSTEMS

Partial-Response System	Characteristic Equation	No. Of Levels
Duobinary	$x_n + x_{n-1}$	3
Dicode	$x_n - x_{n-1}$	3
Modified Duobinary	$x_n - x_{n-2}$	3
Class 2	$x_n + 2x_{n-1} + x_{n-2}$	5

where $\oplus$ represents the XOR operation, and transmit y_n instead of x_n. By doing so, it is easy to verify that levels 2 and 0 at the receiver correspond to $x_n = 0$, and level 1 corresponds to $x_n = 1$ independent of the value of x_{n-1}. Therefore, the encoding on the transmitter eliminated the need for DFE at the receiver without increasing the transmitter voltage headroom requirements since the transmitted sequence still consists of 1 and 0. This signaling technique is called duobinary signaling and was first proposed by Lender [7]. If channel characteristics are not exactly as described above, a linear FFE may precede the channel at the transmitter to equalize the channel to the duobinary ISI pattern. Consequently, in practical systems, duobinary signaling is generally used when channel characteristics are close to a duobinary channel. Duobinary signaling has already been demonstrated at 10 Gbps and beyond over a long FR4 backplane [9]. Other classes of partial-response signaling also exist which are specific to other types of channels. The characteristics of the few common partial response systems, duobinary, dicode, modified duobinary, and class 2, are shown in Table I.

The frequency responses and impulse responses of duobinary, dicode, modified duobinary, and class 2, are also shown in Fig. 7(a) and (b), respectively. The duobinary channel is a low-pass filter with a null frequency at $\Omega = \pi/T$, and the dicode is a high-pass filter with a null frequency at $\Omega = 0$. The modified duobinary is a bandpass filter with null frequencies at both $\Omega = 0$, and $\Omega = \pi/T$. The class 2 system is also a low-pass system with null at a $\Omega = \pi/T$, but with different roll-off shape than that of a duobinary channel.

The eye diagrams of the duobinary and the class 2 systems are shown in Fig. 8(a) and (b), respectively. The duobinary and the class 2 systems have three and five levels, respectively.

B. Multitone Signaling

Partial response methods explained in the previous section mainly exploit the part of channel bandwidth before the first notch in the frequency response of the channel. However, for example, in the case of a duobinary channel with a frequency response of $1 + e^{-j\pi fT}$, channel frequency response bounces back to nonzero values after the first notch frequency. In fact, notches in the frequency response happen at equally spaced distances at $(2k+1)/2T$ and additional channel bandwidth exists for signal transmission between every two notches, which is not utilized by duobinary transmission. In these cases, it is possible to transmit a multitone sequence consisting of a duobinary stream centered at dc and a set of passband duobinary streams centered around nonzero carrier frequencies.

A new multitone architecture suitable for high-speed links, called analog multitone (AMT), was recently proposed in

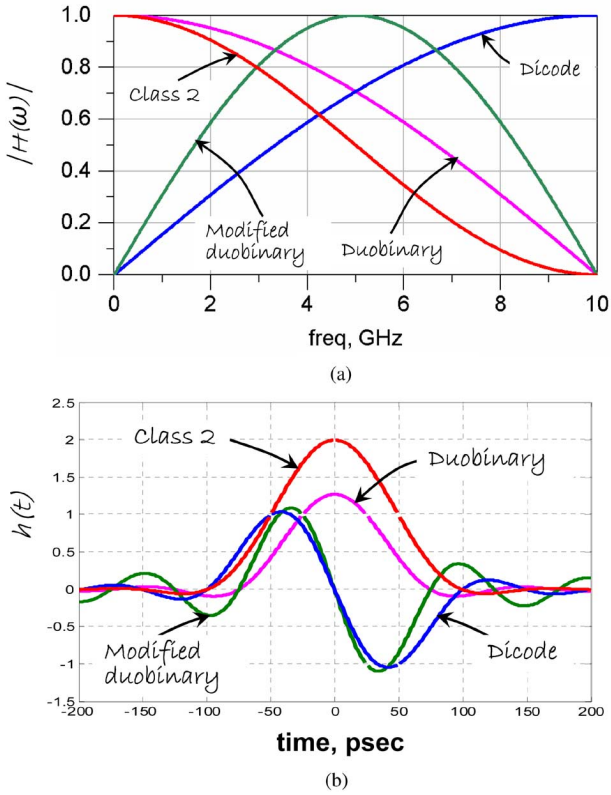

Fig. 7. Frequency-domain characteristics and impulse responses of a few common partial response systems: duobinary (class 1), dicode, modified-duobinary, and class 2. (a) Frequency-domain response. (b) Time-domain response.

[13]–[15]. A simplified three-channel AMT system is shown in Fig. 9(a). The input bit stream is parallelized to three streams with each substream running at one-third of the total bit-rate. Each substream is consequently modulated to its respective carrier frequency and the combined signal is sent over the line. Fig. 9(b) symbolically shows the frequency responses of the individual subchannels at the receiver input. In an AMT system, all carrier frequencies are integer multiples of the substream symbol rate. When no ISI exists, the substreams are separated from each other at the receiver using mixers and integrators. In practical systems with ISI, feed-forward equalizers are placed at the transmitter per substream to maintain the orthogonality between the substreams at the receiver input. The mixers at the transmitter are also combined with the transmit equalizers and performed in the digital domain. Similar to a conventional nonreturn-to-zero (NRZ) system, every substream in an AMT system can have a DFE at the receiver, and DFEs can even exist between the substreams to cancel postcursor interchannel interference (ICI). However, the DFE in AMT runs at the substream rate, which is a fraction of the total system bit rate. Therefore, timing constraints are relieved. Over a duobinary channel $1 + e^{-j\pi fT}$, if the substream rate is set equal to T, the channel delays the combined signal at the transmitter output (and consequently the individual constituent substreams) by a full substream cycle and adds it to itself. However, since all carrier frequencies are integer multiples of $1/T$, this operation does not affect the orthogonality between the substreams. As a result, after mixing and integration at the receiver, each

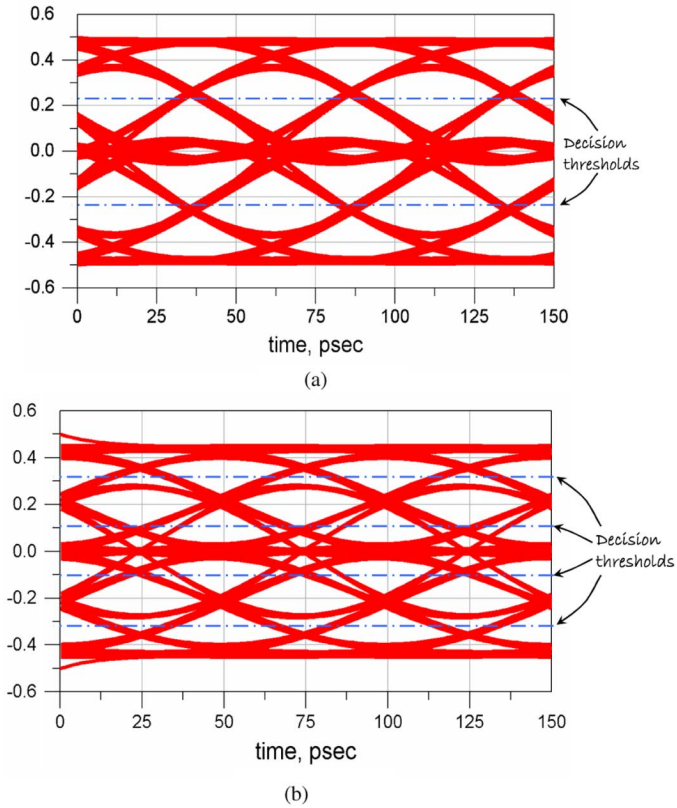

Fig. 8. Eye diagrams of partial response signals at 20 Gbps data rate for (a) duobinary (class 1: three levels) and (b) class 2 (five levels). (a) Doubinary. (b) Class 2.

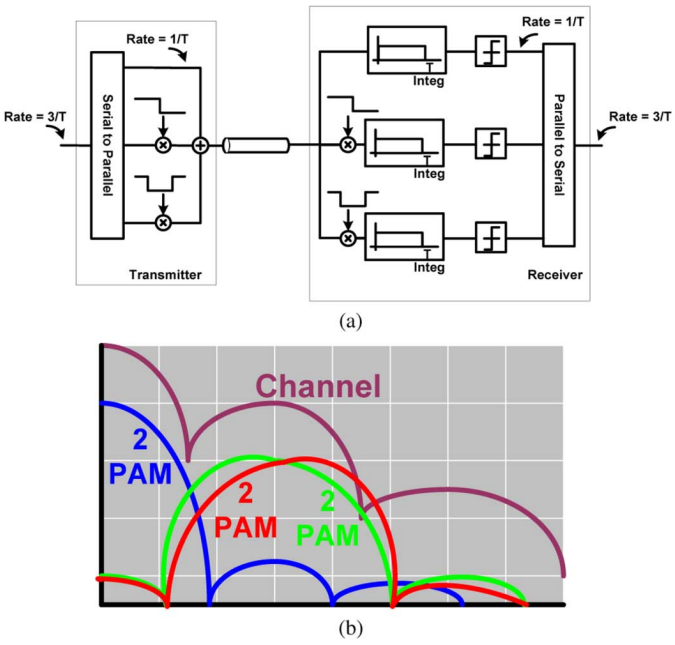

Fig. 9. Three-channel AMT system and the subchannel frequency responses. (a) Three-channel AMT system. (b) Subchannel frequency responses.

substream is separated as a duobinary sequence. Consequently, duobinary encoding can be performed per substream at the transmitter to simplify detection per substream.

Even though we based the argument in this section on a duobinary channel, all the arguments can be extended to all the variations of partial response signaling methods introduced in the previous section.

IV. Controlled ISI Channel Engineering

High-speed link channels are generally such that eliminating ISI with the aid of equalization techniques leads to fairly complex systems and substantial power-consumption. On the other hand, link channels are generally low-pass with sharp rolloffs. Therefore, they are close, but not exactly the same as the low-pass partial response channels in Section III. Even if the channel has a notch similar to duobinary system, the notch frequencies may not correspond to the target signaling rate. Therefore, to create an equivalent partial response channel certain amount of equalization is necessary at the transmitter which leads to increased power consumption and reduces SNR. In this paper, we propose an alternative method to create partial response channels out of regular link channels by adding passive waveguide structures to the PCB trace and on the package.

Fig. 10(a) shows a simple interconnect system with a single stub to tune its transfer characteristic. The response of the system can be shaped by changing the stub length and impedance. The length of the stub determines the frequency nulls, whereas the impedance of the stub affects the roll-off, as shown in Fig. 10(b) and (c), respectively. The length of the stub and the first null frequency are related by

$$f_0 = \frac{1}{4l}\frac{c_0}{\sqrt{\epsilon_r}} \tag{3}$$

where c_0 is the speed of light, ϵ_r is the relative permittivity or dielectric constant of the material, and l is the length of the stub. For a low-loss material, the channel magnitude responses are relatively insensitive to the position of the stub.

By having multiple stubs, a desired spectral shape can be closely approximated. The channel characteristic is modified at multiple locations by two adjustable parameters, the length and impedance of the stubs. Thus, at each stub location, there can exist a change in the characteristic impedance of the channel. Fig. 11(a) shows an interconnect system with two stubs where one stub, for example, can represent a discontinuity which is part of the channel and cannot be removed. The other stub can be intentionally added to change the response of the channel. The lengths of the traces before and after stubs are $l_1 = 50$ mm and $l_3 = 25$ mm, respectively. The lengths of both stubs are $l_4 = l_5 = 14.5$ mm. By varying the distance between the stubs, the transfer function of the interconnect system is significantly modified, as shown in Fig. 11(b). Curve 1, shown in Fig. 11(b), is the channel response before adding the second stub. When the distance between the stubs is 10 mm, a flat channel response is obtained as shown by curve 2. On the other hand, when the distance is set to 5.0 and 2.5 mm, the channel responses show peaking at higher frequencies as shown by curves 3 and 4, respectively. Therefore, the overall channel behavior of an interconnect system can considerably be modified by the use of stubs and the interaction between the reflections from the stubs. This effect caused by the partial reflections from the stubs can be

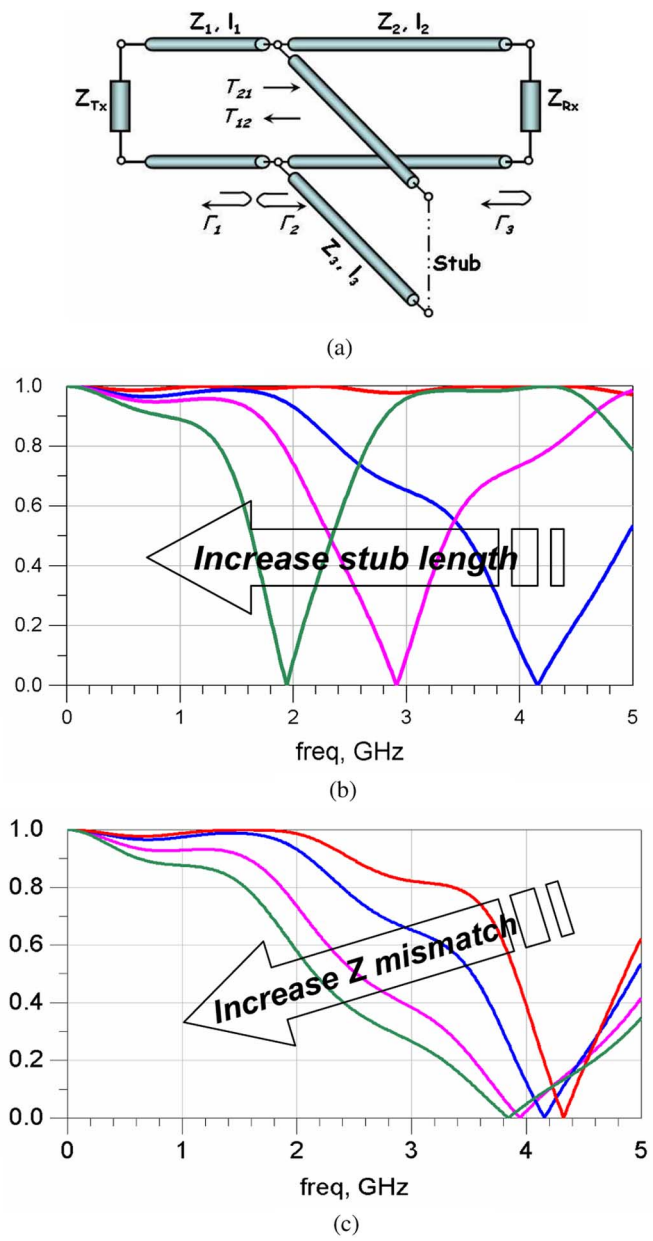

Fig. 10. Shaping the channel transfer function with one stub. (a) A single stub addition. (b) $H(f)$ as a function of stub length. (b) $H(f)$ as a function of stub impedance.

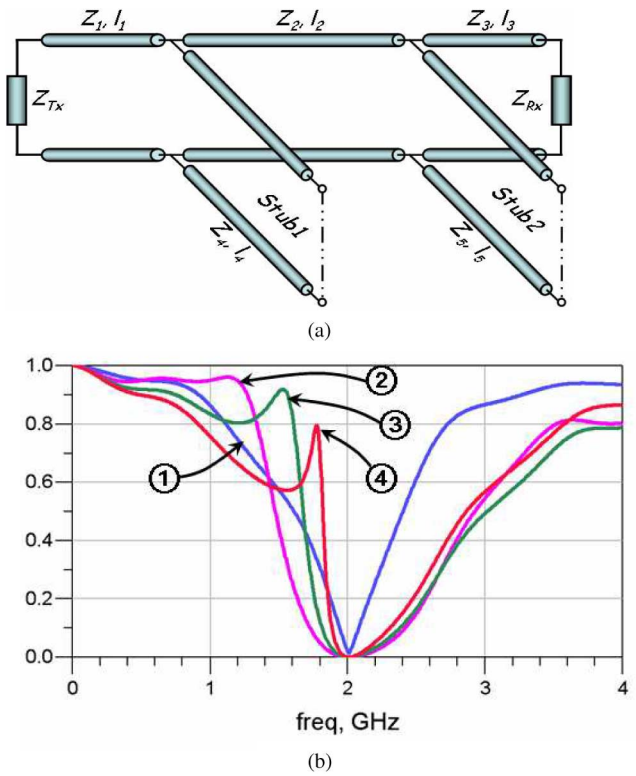

Fig. 11. Shaping the channel transfer function with multiple stubs. (a) The use of multiple stubs to shape a channel characteristic. (b) $H(f)$ as a function of the distance between the stubs. Curve 1 is for only one stub, 2, 3, and 4 for two stubs when the distance l_2 between the stubs are 10, 5, and 2.5 mm, respectively.

approximated using the theory of small reflections [16], [17]. However, in practice since only smell number of design parameters exist to adjust, it easier to find the optimal through simulation to tune few key parameters.

V. Case Study: A 6-in FR4 Chip-to-Chip Interconnect

A diagram of the chip-to-chip interconnect system that we are going to study in this paper is shown in Fig. 12(a). The interconnect consists of a 6-in-long FR4 PCB trace, two low-cost plastic packages each with substrate traces up to 20-mm-long, PCB and package vias, and device parasitics (Ci's and Ri's), as well as stubs. All these components contribute to signal attenuation and dispersion as it propagates through the channel. In order to analyze the channel at multigigahertz frequencies, accurate models of the passive and active components in the channel are generated. Fig. 12(b) shows a circuit representation of a point-to-point differential interconnect system with transmission lines, terminations, and a simple model for the main driver We will study the performance of different transmission methods over this communication channel to demonstrate the effectiveness of the proposed techniques.

A. Conventional Transmit and Receive Equalization

The first approach to mitigate the effect of ISI is the use of equalization techniques. The transfer function of the system before equalization and with power-constrained transmit equalization for a target data rate of 20 Gb/s is shown in Fig. 13(a). The plating stub on the low-cost plastic package adds a stub to the transfer function of the channel at around 14 GHz. For this reason, the transmit equalizer has to significantly attenuate the low-frequency content of the transmit data to make the overall equalized transfer function flat. This leads to significant loss of SNR at the receiver. Fig. 13(b) shows the original and the equalized SBR of the system. The equalized eye diagrams using a standard DFE and a one-tap loop-unroll DFE are shown in Fig. 14(a) and (b), respectively. Although the received eye diagram using a standard DFE shows some eye opening, the voltage and timing margins of the system are clearly improved using a loop-unroll DFE.

B. Channel Engineering With Doubinary Signaling

The second approach is to modify the characteristics of the channel to match the characteristics of a duobinary system by

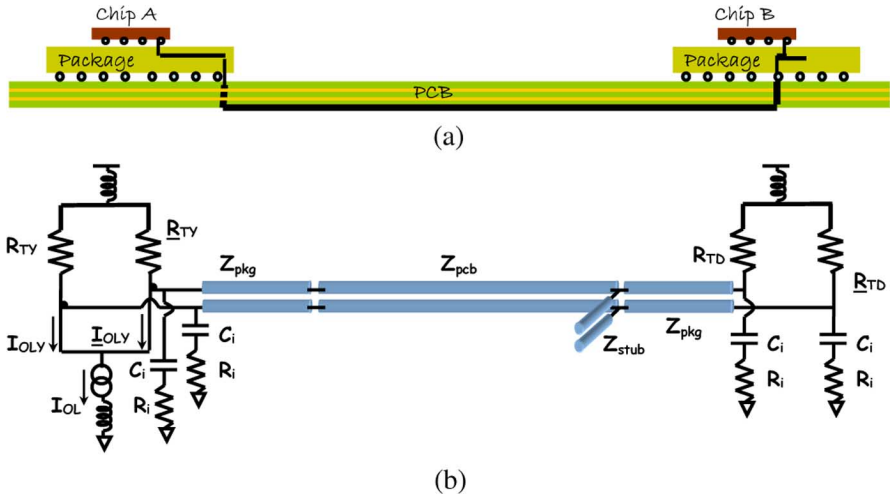

Fig. 12. Chip-to-chip interconnect system and the channel model. (a) Interconnect system. (b) Channel model.

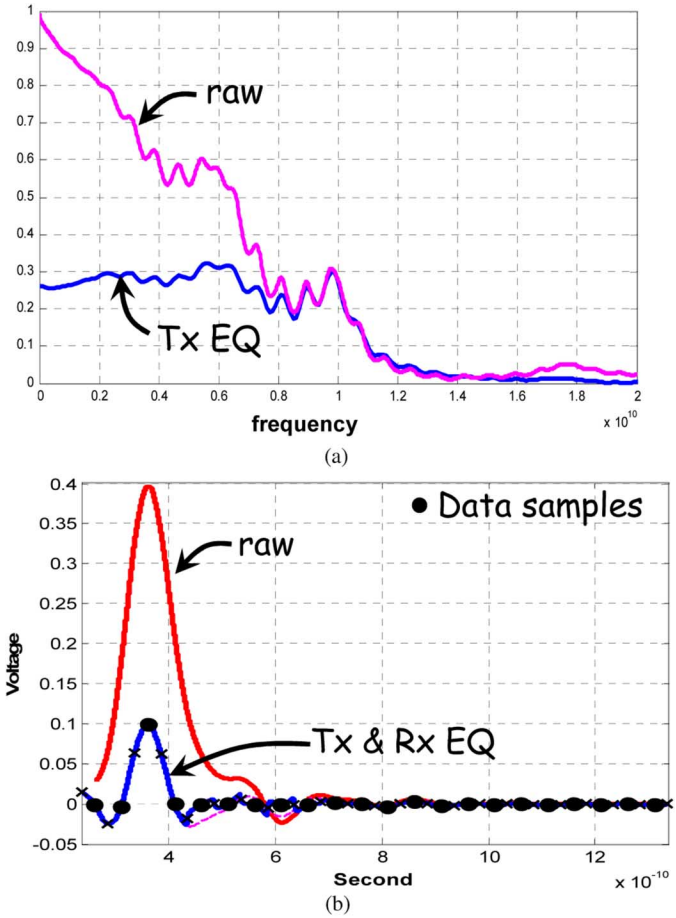

Fig. 13. Transmit and receive equalization. (a) Frequency responses. (b) Single-bit responses.

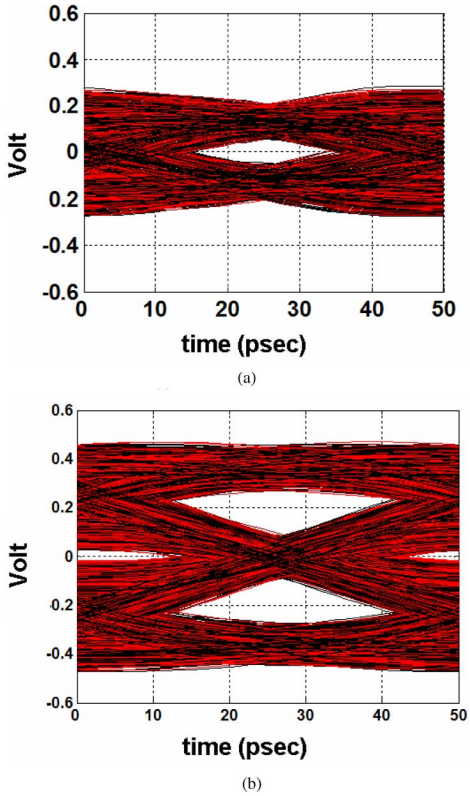

Fig. 14. Eye diagrams with transmit and receive equalization. (a) Using standard DFE. (b) Using one-tap loop-unroll DFE.

designing trace stubs on the board and package, as shown in Fig. 15. These stubs are designed to place nulls at the maximum operating frequency of the system. The length and impedance of the traces are tuned to control the shape of the frequency characteristics of the response. The trace length of the stub is 3.75 mm.

The transfer functions of the ideal duobinary, original and modified interconnect systems are shown in Fig. 16(a). The modified transfer function closely matches that of the duobinary system and it has a null at the Nyquist frequency of 10 GHz. Even though the transfer function of the original system shows that more energy can be transferred at 10 GHz, the ISI is so severe that a data rate of 20 Gbps cannot be transmitted reliably. The modified system, however, introduces a controlled amount of ISI and hence the system shows much improved voltage and timing margins as shown in Fig. 16(b) without any need for

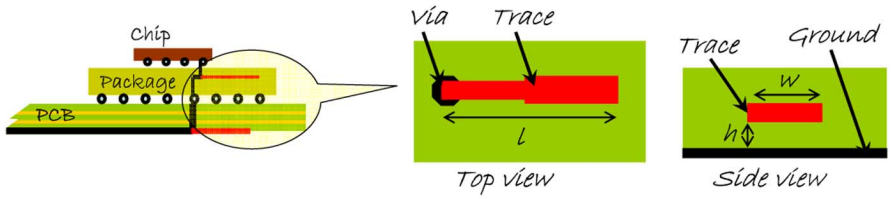

Fig. 15. An example of package and PCB traces that can be designed in to modify the characteristic of the overall interconnect system.

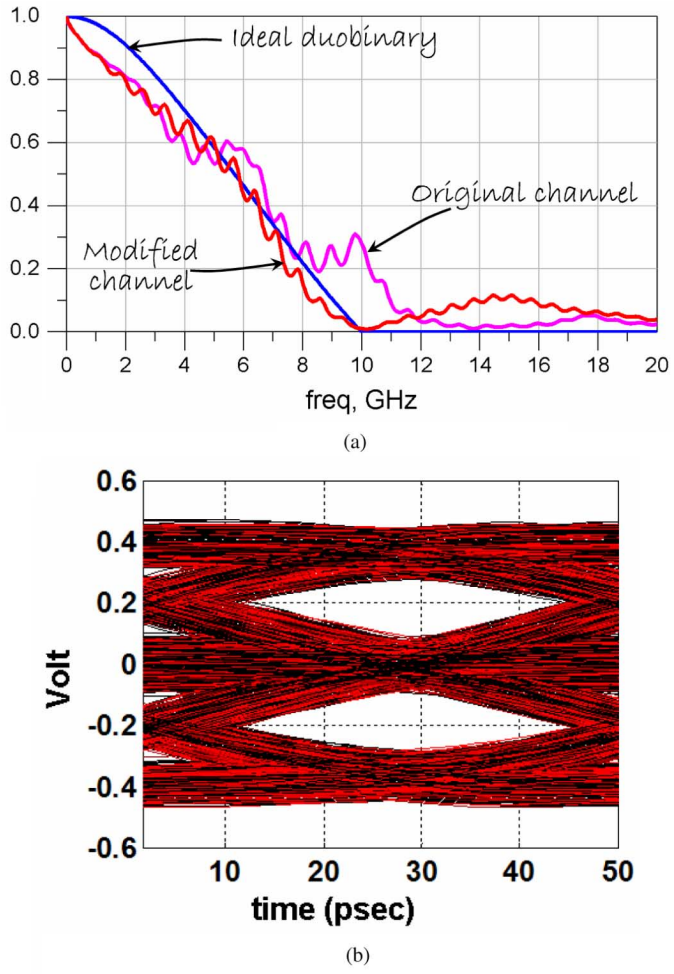

Fig. 16. (a) Transfer functions of the ideal duobinary, original and modified channels, and the eye diagrams of (b) the modified channels. (a) Channel transfer functions. (b) Modified eye diagram.

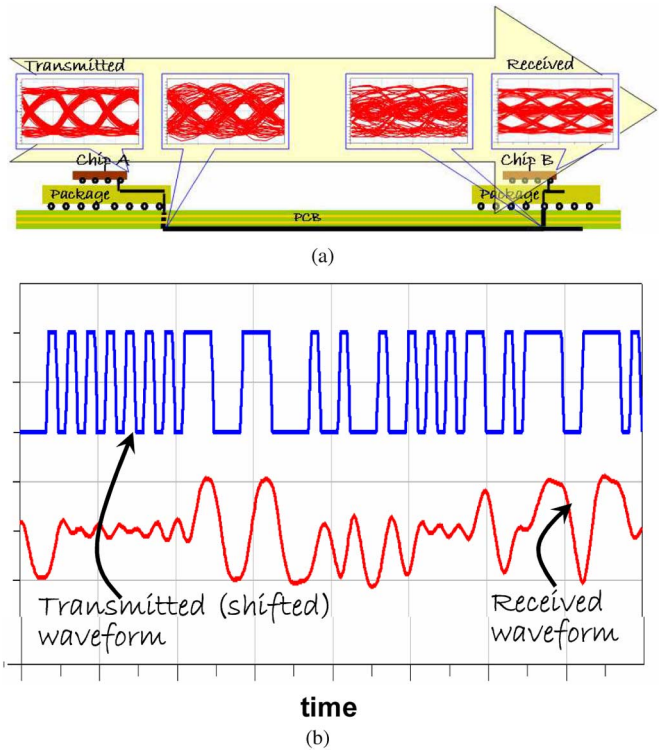

Fig. 17. The modified chip-to-chip interconnect system converts binary into ternary signal. The eye diagrams along the signal path and the transmitted waveform (shifted by the delay of the channel) and the waveform at the receiver. (a) Modified chip-to-chip interconnect system. (b) Transmitted and received waveforms.

transmit or receive equalization. The partial response signaling, such as duobinary signaling also has a better crosstalk, and reflection immunity as well as low electromagnetic interference (EMI) due to the reduced spectral content at high frequencies. The required maximum frequency can be reduced because a partial-response signaling allows a controlled amount of ISI.

Fig. 17(a) shows the eye diagrams along the signal path from the input to the output of the chip-to-chip system. The transmission medium converts the standard binary NRZ bits to duobinary or class 2 multilevel correlated bits. The transformed binary data pattern and the received waveforms are shown in Fig. 17(b). The input waveform is shifted by the delay of the channel to align the waveforms and show the effect of the channel on the input. The output bit is the sum of the current and the previous bits sent by the transmitter.

C. Channel Engineering With AMT Signaling

Fig. 18 shows the frequency response of the original channel in decibels together with the loss tangent. Careful inspection of the frequency response indicates that the channel response starts to bounce back to reach the loss tangent after the first notch frequency (caused by the plating stub), similar to Fig. 10(b). However, the response is dampened due to the existence of other discontinuities in the signal path. If it was not for these other discontinuities, the response of the channel would bounce back to the loss tangent after the first notch, to have less than 20 dB attenuation below 20 GHz. Therefore, part of the channel useful transmission capacity is wasted beyond the first notch. The third approach investigated in this section is to extend the length of the plating stub on the package to move the dominant notch frequency to lower frequencies, as shown in Fig. 18. Even though this modification reduces the

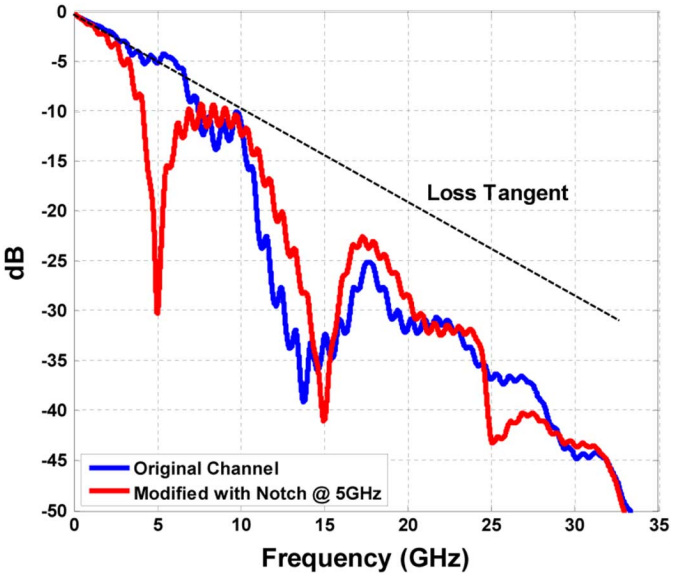

Fig. 18. Original and modified frequency responses.

bandwidth of the channel before the first notch frequency, it allows the channel response to bounce back to the loss tangent, before being dampened by other discontinuities. Since now the frequency response of the channel is mostly determined by a single stub over useful frequencies, and stubs have periodic frequency responses, a second notch is created at 15 GHz. As a result, a whole pass-band channel is opened up between 5 and 15 GHz for pass-band signal transmission using an AMT transmitter, similar to Fig. 9(b). The AMT system would require three-channels (a 10 Gb/s duobinary baseband channel, and two quadrate 10 Gb/s duobinary passband channel) to achieve an aggregate data-rate of 30 Gb/s. Fig. 19(a)–(c) shows the eye diagrams of the three substreams of the optimized AMT system over this channel.

Since the modified channel does not completely bounce back to 0 dB, as an ideal duobinary channel would do, the AMT system requires certain amount of transmit equalization or alternatively would need a linear equalizer with 10 dB gain at 10 GHz at the receiver front-end. The receiver of the AMT system also requires a one tap DFE per substream with DFE loop time of 200 ps to cancel the second postcursor ISI for each substream. With this additional complexity, 50% higher data rate compared to the previous method is achievable. Fig. 20 shows the equalized response of the three substreams at the receiver input before down-conversion.

VI. Conclusion

In this paper, a controlled ISI design technique of interconnects of high-speed chip-to-chip communications is presented. The nonidealities in conventional packaging and board technologies are utilized to shape the response of the interconnect system. This enables conventional interconnects to transmit high data rates with minimum bandwidth and reduced timing uncertainties using partial response and multitone signaling. The partial response signaling, such as duobinary signaling, also shows better crosstalk, EMI, and reflection immunity.

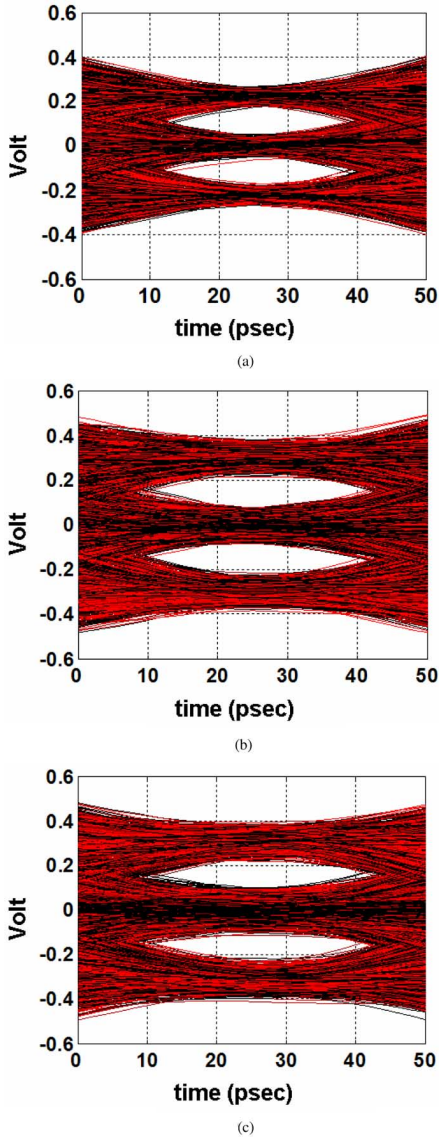

Fig. 19. (a) Eye diagrams of the first, second, and third substreams. (a) First channel. (b) Second channel. (c) Third channel.

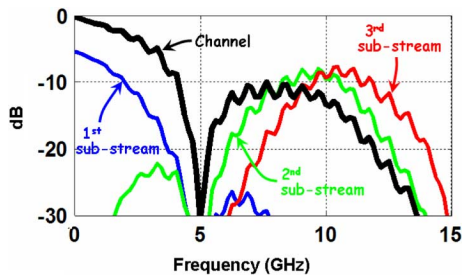

Fig. 20. Equalized response of the three AMT substreams at the receiver input before down-conversion.

Using the proposed design approach, low-cost conventional package and board technologies can be used to run next-generation links beyond 20 Gbps. Design examples are given to demonstrate the validity and advantages of the technique using duobinary and analog multitone signaling methods.

REFERENCES

[1] W. T. Beyene, C. Yuan, R. Dhat, and D. Secker, "Effects of plating stubs on the electrical performance of wirebond PBGA packages," in *Proc. Int. Conf. Adv. Packag. Syst. (ICAPS)*, Reno, NA, Mar. 10–13, 2002, pp. 17–22.

[2] D. Kam, J. Yu, H. Choi, K. Bae, C. Lee, and J. Kim, "Design of low cost package for 40-gb/s serial link using wire-bonded plastic ball grid array," *IEEE Design Test Computers*, pp. 212–219, May–Jun. 2006.

[3] W. J. Dally and J. W. Poulton, *Digital Systems Engineering*. Cambridge, U.K.: Cambridge Univ. Press, 1998.

[4] E. Lee and D. Messerschmitt, *Digital Communication*, 2nd ed. Boston, MA: Kluwer, 1994.

[5] J. Zerbe *et al.*, "Design, equalization and clock recovery for a 2.5–10 Gb/s 2-PAM/4-PAM backplane transceiver cell," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, San Francisco, CA, Feb. 2003, pp. 80–81.

[6] W. T. Beyene, "Controlled inter-symbol interference design techniques of conventional interconnect systems for data rates beyond 20 gbps," in *IEEE 15th Topical Meeting Electr. Performance Electron. Packag.*, Scottsdale, AZ, Oct. 23–25, 2006, pp. 159–162.

[7] A. Lender, "The duobinary technique for high-speed data transmission," *IEEE Trans. Commun. Electron.*, vol. CE-82, pp. 214–218, May 1963.

[8] P. Kabal and S. Pasupathy, "Partial-response signaling," *IEEE Trans. Commun.*, vol. C-23, no. 9, pp. 921–934, Sep. 1975.

[9] J. H. Sinky, M. Duelk, and A. Adamiecki, "High-speed electrical backplane transmission using duobinary signaling," *IEEE Trans. Microw. Theory Tech.*, vol. 53, no. 1, pp. 152–160, Jan. 2005.

[10] Kasturia and J. H. Winters, "Techniques for high-speed implementation of nonlinear cancellation," *IEEE J. Sel. Areas Commun.*, vol. 9, no. 5, pp. 711–717, Jun. 1991.

[11] V. Stojanović, A. Ho, B. Garlepp, F. Chen, J. Wei, E. Alon, C. Werner, J. Zerbe, and M. A. Horowitz, "Adaptive equalization and data recovery in dual-mode (PAM2/4) serial link transceiver," in *IEEE Symp. VLSI Circuits*, Jun. 2004, vol. 9, no. 5, pp. 348–351.

[12] A. Garg, A. Chan Carusone, and S. Voinigescu, "A 1-tap 40-gbps look-ahead decision feedback equalizer in 0.18 μm sige bicmos technology," *IEEE J. Solid-State Circuits*, vol. 41, no. 10, pp. 2224–2232, Oct. 2006.

[13] A. Amirkhany, A. Abbasfar, V. Stojanović, and M. A. Horowitz, "Practical limits of multi-tone signaling over high-speed backplane electrical links," in *2007 Symp. VLSI Technol. Circuits*, Japan, Jun. 14–16, 2007, pp. 38–39.

[14] A. Amirkhany, A. Abbasfar, V. Stojanović, and M. A. Horowitz, "Analog multi-tone signaling for high-speed backplane electrical links," in *GlobeCom*, Nov. 2006, pp. 1–6.

[15] A. Amirkhany, A. Abbasfar, J. Savoj, M. Jeeradit, B. Garlepp, V. Stojanović, and M. A. Horowitz, "A 24 Gb/s software programmable multi-channel transmitter," in *VLSI Symp. VLSI Circuits*, Kyoto, Japan, Jun. 2007, pp. 38–39.

[16] R. E. Collin, *Foundations For Microwave Engineering*, 2nd ed. New York: McGraw-Hill, 1992.

[17] D. M. Pozar, *Microwave Engineering*, 3rd ed. New York: Wiley, 2005.

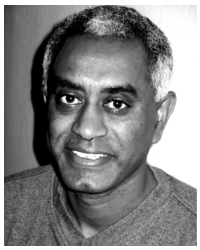

Wendemagegnehu (Wendem) T. Beyene (S'87–M'88–SM'05) was born in Addis Ababa, Ethiopia. He received the B.S. and M.S. degrees in electrical engineering from Columbia University, in 1988 and 1991 respectively, and the Ph.D. degree in electrical and computer engineering from University of Illinois, Urbana-Champaign, in 1997.

From 1988 to 1994, he was with the IBM, Microelectronics Division, Fishkill, NY, where he worked on design and electrical characterization of advanced multilayer packages. From 1997 to 2000, he was with Hewlett-Packard Company and Agilent Technologies EEsof EDA, Westlake Village, CA working on analog and RF circuit simulation tools. In 2000, he joined Rambus Inc., Los Altos, CA and is currently a Senior Principal Engineer responsible for signal integrity of multigigabit memory and logic-to-logic interfaces. His professional interests are in the general area of simulation and optimization of deterministic and stochastic systems. Specific current interests include efficient circuit simulation of large distributed networks, noise analysis of autonomous circuits, and analysis of parameter variations in high-speed links.

Dr. Beyene is a member of Eta Kappa Nu, Tau Beta Pi, and SIAM.

Amir Amirkhany (S'04–M'08) received the B.Sc. degree from Sharif University of Technology, Iran, in 1999, the M.Sc. degree from the University of California, Los Angeles, in 2002, and the Ph.D. degree from Stanford University, Stanford, CA, in 2008, all in electrical engineering.

Since July 2007, he has been with Rambus Inc., Los Altos, CA. From 2003 to 2007, he was a Research Assistant with the VLSI group, Stanford University, where he was involved with the design of chip-to-chip electrical links, in close collaboration with Rambus Inc. From 2001 to 2002, he was with Sequoia Communications, Los Angeles, CA, working on the ASIC design of WCDMA systems. His main research interests include the design and implementation of communication systems, VLSI circuit design, and application of communication and signal processing techniques to the design of low power circuits.

Dr. Amirkhany was a recipient of a Best Student Paper Award at the IEEE Global Communications Conference in 2006 for his work on the design and analysis of an analog multitone system for chip-to-chip interconnects.