

Frequency-Division Bidirectional Communication Over Chip-to-Chip Channels

Mike Bichan, *Student Member, IEEE*, Masum Hossain, and Anthony Chan Carusone, *Senior Member, IEEE*

Abstract—Frequency division multiple access is applied to bidirectional communication over chip-to-chip links. Frequency division is implemented by dividing the spectrum into low-frequency (dc) and high-frequency (ac) bands using a simple LC filter. The nonidealities that this filter introduces are compensated for with a transmitter/receiver pair that can recover signals in both bands. The receiver uses a dual-path topology that includes hysteresis to recover data from a signal with no dc content. The transmitter is a 6-tap (FIR) pre-emphasis equalizer with variable tap spacing. In simulation, the transmitter and receiver simultaneously communicate error-free at 8 Gb/s over the ac channel and at 500 Mb/s over the dc channel. Measurements shows that the ac and dc signals can be individually recovered and that the two signals occupy distinct frequency bands.

Index Terms—Complementary metal–oxide–semiconductor (CMOS) integrated circuits, digital communication, equalizers, frequency division multiaccess, hysteresis, microstrip, multichip modules, pulse shaping circuits, transceivers.

I. INTRODUCTION

IN high-speed chip-to-chip communication links, limited channel bandwidth introduces inter-symbol interference (ISI) into the received signal. Equalizers in the transmitter and receiver have been widely used to compensate for loss at high frequency and allow data to be sent at rates higher than the bandwidth of the channel. A common configuration involves a decision feedback equalizer (DFE) in the receiver and a feed-forward equalizer (FFE) in the transmitter [1]–[3]. Another popular technique uses a source-degenerated differential pair to implement tunable receive-side equalization [4]–[7]. The highest data rate achieved over a board-to-board channel by an integrated circuit (IC) implementation is 20 Gb/s [8].

In order to increase the data rate between chips in, for example, CPU-RAM links, usually either the maximum data rate per wire increases or the number of pins dedicated to input/output (I/O) will increase. Alternately, the aggregate data rate can be increased by simultaneously transmitting in both directions over the same wire. Previous work on simultaneous bidirectional (SBD) transceivers used an echo canceler to separate the transmitted signal from the received signal [9].

Manuscript received July 02, 2008; revised December 09, 2008. First published May 08, 2009; current version published May 28, 2009. This work was recommended for publication by Associate Editor W. Beyene upon evaluation of the reviewer comments.

The authors are with the Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON, M5S 3G4 Canada (e-mail: mbichan@eecg.utoronto.ca; masum@eecg.utoronto.ca; tcc@eecg.utoronto.ca).

Digital Object Identifier 10.1109/TADVP.2009.2013454

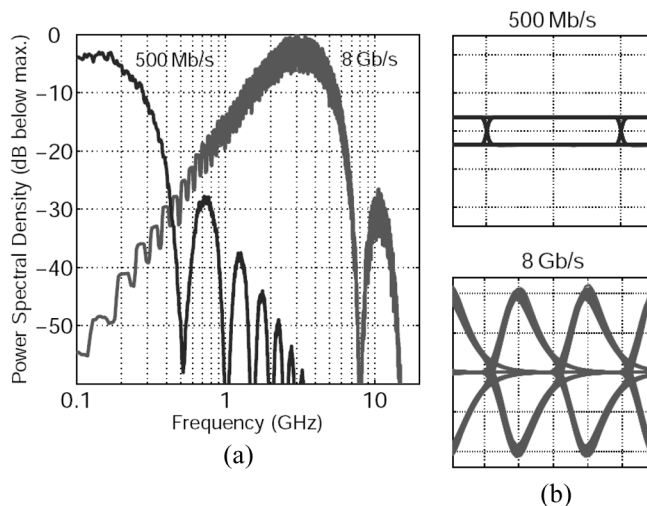


Fig. 1. (a) Power spectra and (b) eye diagrams of the two signals used in the proposed frequency-division bidirectional communication link: 500 Mb/s dc signal and 8 Gb/s ac signal.

Increasing I/O pin density also increases the aggregate data rate. Normally the density is limited by the number of solder balls that can be placed in a given area. However, if the signals are capacitively coupled off-chip no solder ball is required. This ac-coupled signaling can achieve an I/O density an order of magnitude better than area ball bonding [10]–[14].

The ability to communicate in both directions over a pair of microstrip traces without an echo cancellation circuit allows an increase in I/O density without extra power and complexity. In this paper we discuss a method of frequency-division multiplexing for SBD communication which uses passive devices to separate signals moving in the two directions. Previous work in this area has used complex radio-frequency (RF) topologies to enable transmission in multiple frequency bands [15]. The proposed solution uses simple passives to divide the available bandwidth into dc and ac bands. The transmitter and receiver circuits are designed to be programmable so that the same circuits can be used for both the dc and ac bands. The partition of the frequency spectrum is shown in Fig. 1.

The receiver consists of two paths whose outputs are combined using a weighted summer. The summer can be adjusted so that the receiver can recover both regular non return to zero (NRZ) data as well as high-frequency ac pulses.

The transmitter is a 6-tap FIR equalizer with variable tap spacing. Changing the tap spacing allows the transmitter to control the spectrum of the output signal. Using a shorter tap




Fig. 2. Typical chip-to-chip channel.

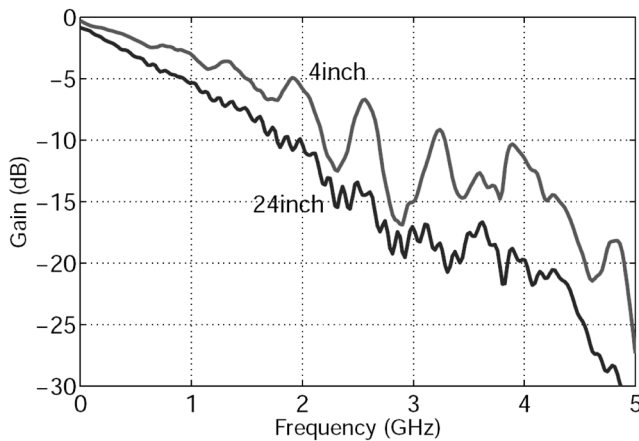


Fig. 3. Frequency response of a chip-to-chip channel.

spacing results in more high frequency content in the output. This feature makes the transmitter ideal for this application where frequency separation of the two signals is important.

Sections II–IV provide background information about chip-to-chip communication, as well as bidirectional and ac coupled links. The simulation results in Section V represent a proof-of-concept of the ac/dc bidirectional transceiver. Results are shown where the ac channel operates at 8 Gb/s and the dc channel operates at 500 Mb/s. Section VI presents some measurements of the transmitter and receiver used together to support the feasibility of the bidirectional technique.

II. CHIP-TO-CHIP COMMUNICATION LINKS

The chip-to-chip channel is a system that consists of either two chips on the same board, or two chips on separate daughtercards attached to a common motherboard. The latter case is pictured in Fig. 2. The frequency response of this channel is shown in Fig. 3 for two different lengths of motherboard.

This frequency response illustrates the main impediment to communication, which is high frequency attenuation. In order to communicate at data rates exceeding the -3 dB frequency of the channel, equalization is required. The difference between the gain of the equalizer at low and high frequencies is known as




Fig. 4. Typical chip-to-chip transceiver with representative eye diagrams.




Fig. 5. Power per bit rate of state-of-the-art transceivers. Transceivers that run at multiple data rates have one (\square) for each rate, connected by lines. The numbers indicate paper references.

the boost. As shown in Fig. 4, a typical chip-to-chip transceiver includes both transmit pre-emphasis and a receive-side equalizer to generate the boost needed to compensate for ISI caused by high frequency attenuation.

Reflections are small, undesired signals that propagate along the channel as a result of impedance discontinuities. The presence of reflections can be seen in the channel frequency response, where they show up as regularly-spaced bumps. They can be lessened to a limited extent by ensuring a good impedance match at the output of the transmitter and the input of the receiver. However, impedance discontinuities in the circuit board itself also cause reflections. Via stubs are a severe source of reflections [16]. Stubs can be counterbored to reduce parasitic inductance and capacitance, but this is an expensive process. Some research focuses on equalizing lossy channels with via stubs that have not been counterbored [2].

When many microstrip traces run alongside one another as in a parallel link, crosstalk occurs among the various traces. The crosstalk path between adjacent traces acts like a differentiator; mainly the high frequency content is transmitted. Because the microstrip trace itself attenuates high frequency content, crosstalk is usually a small problem between transmitters and receivers at opposite ends of the channel—so-called (FEXT). However, when a transmitter and a receiver are colocated the aggressor signal is not attenuated. In this case, (NEXT) closes the received eye and must be ameliorated with crosstalk cancellation [8].

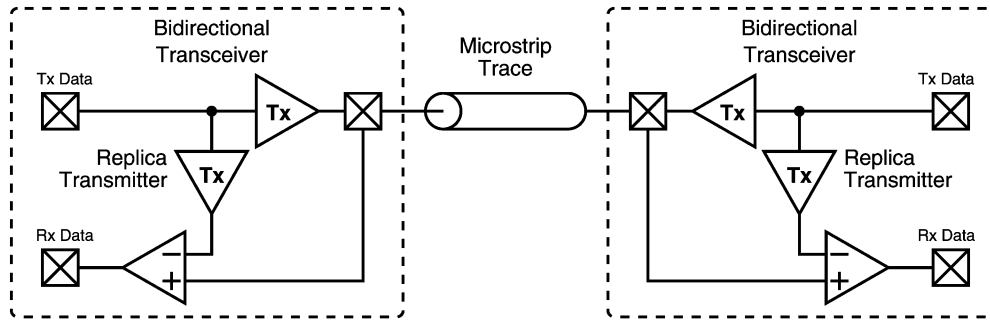


Fig. 6. Typical bidirectional chip-to-chip transceiver.

Current state of the art chip-to-chip transceivers aim to extend the usefulness of the printed circuit board (PCB) channel by providing ever more sophisticated equalization. For example, some transceivers use modulation techniques such as duobinary signaling [17] or 4-PAM [8] in order to make up for a reduced channel bandwidth with more complex receive-side signal processing. Duobinary signaling involves allowing a certain amount of ISI at the input to the receiver, the idea being that if the ISI is well controlled it can be removed more easily. 4-PAM and other pulse-amplitude modulation schemes can halve the required channel bandwidth by doubling the number of signal levels used. One tradeoff is that the receiver must be able to make decisions using multiple thresholds and smaller vertical eye openings.

Other research improves upon the power dissipation of existing transceivers by scaling power with data rate [6] or by choosing a fixed data rate at which the highest power efficiency can be obtained [19]. The power efficiencies and data rates of recent chip-to-chip transceivers can be seen in Fig. 5.

III. BIDIRECTIONAL COMMUNICATION OVER CHIP-TO-CHIP LINKS

Bidirectional communication can conserve package pins in chip-to-chip links by transmitting in both directions at the same time on one pin. Both the transmitter and the receiver are connected to the same pin, as shown in Fig. 6. A replica transmitter is then used to subtract the transmitted signal from the received signal. The cost of transmitting in both directions is that the signal seen by the receiver is much noisier than for unidirectional communication. The replica transmitter must accurately reproduce the transmitted signal even though it may see a very different impedance [9].

Radio communication links use time and frequency division to divide up access to a channel with finite bandwidth. Time division is also present in chip-to-chip communication in the form of a shared bus, although due to the parasitics added by each additional transceiver on the bus they are not suitable for multi-gigabit communication.

A recent attempt at implementing frequency division for chip-to-chip links has used RF mixers and bandpass filters to divide the channel into several ac bands in addition to baseband [15]. This transceiver achieved 2.4 Gb/s over the baseband channel while transmitting 600 Mb/s over one RF channel.




Fig. 7. Side-view of an ac-coupled link.

Theoretically, multiple orthogonal RF channels could be active at once to allow the transceiver to be used in a bus.

For longer chip-to-chip channels with reduced bandwidth, attenuation above the -3 dB frequency may be too severe to allow multiple bandpass channels. Instead, for multi-gigabit links it may be more practical to divide the available bandwidth into only two bands.

The ratio of the two directions' data rates is also a consideration. Bidirectional techniques with symmetrical data rates lend themselves to high-speed data transfer between chips [9]. However, asymmetrical data rates are also useful in applications such as memory access in a computer as well as adaptation of a transmit-side equalizer using receive-side eye information [3].

IV. AC COUPLED LINKS

AC coupled links use nonconductive means to send signals between chips in close proximity. The earliest such links used capacitive coupling to allow testing of dies before their insertion into costly multichip modules [10], [20]. Shortly thereafter, latch- and hysteresis-based receivers were shown to allow signal transmission up to 200 MHz [11]. Multi-Gbps data rates have since been demonstrated, using both capacitively- and inductively-coupled signaling [12].

AC coupled, or proximity communication, links can increase the aggregate chip-to-chip data rate by increasing the density of I/O pads on-chip. Regular area ball bonds require a pitch on the order of $150 \mu\text{m}$ whereas wireless capacitive connections can achieve a density approximately 60 times greater [14]. Instead of a direct, dc connection, the two chips are placed face-to-face

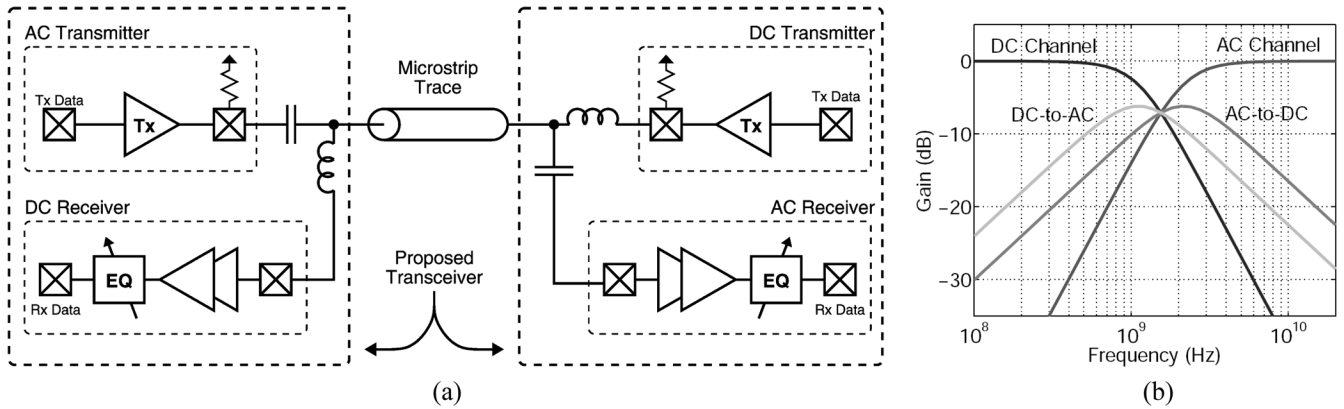


Fig. 8. (a) Proposed ac/dc bidirectional link, (b) frequency response of different paths in the ac/dc bidirectional link with $L = 15 \text{ nH}$ and $C = 4 \text{ pF}$.

and the signal is sent over a capacitor made from two metallic plates. One plate on the transmitter side and one on the receiver side. ac coupled links are especially attractive for multi-chip modules where one of the chips maintains a conventional, dc wirebonded connection. An illustration of this technique is shown in Fig. 7.

Aside from mechanical alignment of the chips, ac coupled links face another challenge. Because the channel does not transmit any energy at dc, conventional broadband transceivers cannot be used. When an NRZ signal is sent across the channel, what appears at the receiver is a series of ac pulses corresponding to the data transitions at the transmitter. The receiver must be able to recover the original NRZ data from these ac pulses. Since the dc content is not received in any case, it can be eliminated from the transmitted signal as well.

AC coupled receivers typically include a memory element that stores the previously received bit until a transition is detected. When the received signal is corrupted by noise and a bit error occurs, the receiver continues to make errors until the next transition occurs. The impact of these error runs can be minimized by coding the data stream to contain frequent transitions [21].

V. AC/DC BIDIRECTIONAL LINK

To make use of the high frequency channel bandwidth we propose an ac channel using the same physical wire as the broadband channel. The ac and dc channels are frequency-separated by a simple filter consisting of passive elements. The two channels can then be used as a bidirectional link by sending the frequency-separated signals along the wire in opposite directions. A schematic of the bidirectional link is shown in Fig. 8(a).

The data rates of the ac and dc channels must be chosen so as to maximize the aggregate data rate while still allowing the two receivers to recover independent data simultaneously. As seen in Fig. 1, both signals contain negligible energy at frequencies equal to their respective data rates. So intuitively it is attractive to set the data rate of the ac signal to twice that of the dc signal to place the frequency of maximum ac power at a null in the dc

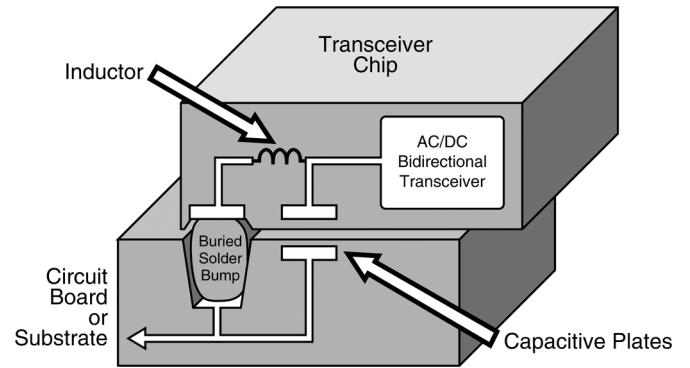


Fig. 9. AC/DC bidirectional link using buried-bump technology.

power. In practice, the data rate of the ac signal must be several octaves higher than that of the dc signal in order to allow filtering of the two frequency bands.

Frequency selection is performed by simple inductor and capacitor elements. Hence, the filters do not strongly separate the two frequency bands. The frequency responses of the various paths through the system are shown in Fig. 8(b). There remains significant NEXT between the ac (dc) transmitter and dc (ac) receiver. The -3 dB frequency of the filters should be chosen to limit this interference and permit signal detection at both ends of the link.

In addition, to produce a maximally-flat frequency response for both the ac and dc channels, the impedance seen by the transmitter should be equal to $50 \text{ }\Omega$ across all frequencies. Unfortunately, the -3 dB frequencies of the ac and dc channels cannot be more than one octave apart without affecting the characteristic impedance of the channel. The -3 dB frequency of the ac channel should be chosen first to allow recovery of the high-frequency ac data. Then, the -3 dB frequency of the dc channel should be chosen to satisfy the requirement of $50 \text{ }\Omega$ impedance across all frequencies. The remaining low-frequency bandwidth is then used for dc data. The relationship between the data rates and -3 dB frequencies is discussed further with the simulation results in this section.

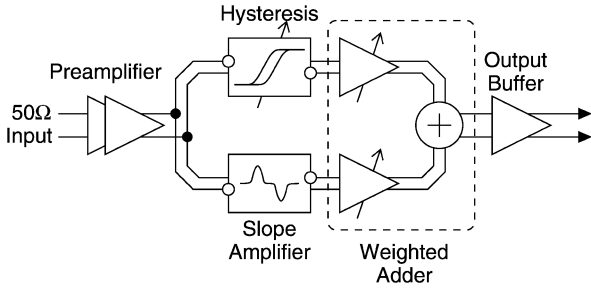


Fig. 10. Hysteresis-based dual-path receiver.

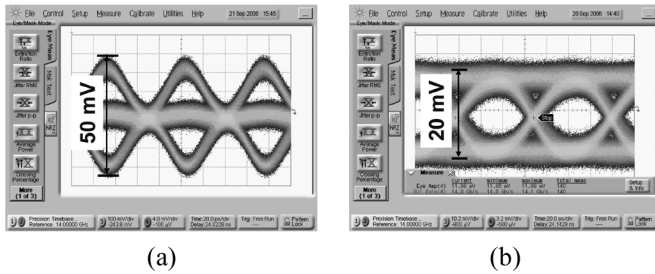


Fig. 11. Measurements of the receiver presented in [23] recovering 14 Gb/s data. (a) Received signal. (b) Equalized signal.

In [22], a technique was presented that allows both dc and ac connections to the same chip. This technique, shown in Fig. 9, uses capacitive plates to send the ac signal from one chip to another. The dc contact is made by a solder bump. Alignment of the capacitive plates is made possible because the solder bump is buried in the substrate, allowing the capacitive plates to be sufficiently close together. The buried-bump technique fits our application because the ac coupled interconnect provides the required capacitance for free while the dc contact through the buried solder bump leaves us free to set the series inductance on-chip.

A. Receiver

The receiver passes the incoming signal through a preamplifier before equalizing it using the weighted sum of two paths. The first path is a hysteresis circuit that can recover a full-swing NRZ signal from low-swing ac pulses. The second path is a linear amplifier that amplifies the data transitions to compensate for the reduced bandwidth of the hysteresis circuit. A schematic of the receiver is shown in Fig. 10.

This receiver was first presented in [23] as an ac-coupled receiver. Because of the adjustable weighted summer, when no equalization is required the hysteresis path can be turned off which increases the bandwidth and reduces jitter. One limitation of this hysteresis-based receiver is that jitter is unaffected by the equalization. That is, it can increase the vertical, but not the horizontal, eye opening. Measurement results of the receiver recovering data from 14 Gb/s ac pulses are shown in Fig. 11.




Fig. 12. FIR transmitter with variable tap delay.




Fig. 13. Measured 5 Gb/s output of the transmitter presented in [24] for the 24-in channel pictured in Fig. 2. (a) Unequalized channel output. (b) Channel output with transmit pre-emphasis enabled.

B. Transmitter

The transmitter is a 6-tap FIR equalizer with 4-bit adjustable tap weights and variable delay between taps. The tap delay can be set to fractional values as well as the usual baud rate spacing. The transmitter can be used to pre-emphasize the signal before it reaches the channel, ensuring that the signal seen at the receiver is more easily recoverable. The flexibility of the transmitter means that it can be used to equalize a variety of channels. A schematic of the transmitter is shown in Fig. 12 [24]. Measurements of the transmitter equalizing a 24-in link at 5 Gb/s are shown in Fig. 13.

This transmitter is suited to the frequency-division bidirectional application because it can be used as the transmitter for both the ac and dc channels. For the ac channel, a short tap delay can be used to send an ac pulse with accentuated high frequency content. For the dc channel, the tap delay can be set longer and the rise time of the signal controlled so that the high frequency content is reduced. This control over the pulse shape allows the simple inductance–capacitance LC filter to separate the two signals so that they can be recovered at the receiver.

C. Link Performance

Unidirectional simulations of the ac and dc channels are shown in Fig. 14. These are the outputs of the ac and dc channels when the transmitter is set up to send the signals shown in Fig. 1(b). The limited bandwidth of the dc channel causes ISI at the output of the channel. The transmitted ac pulse has been decreased in amplitude by the attenuation of the channel.




Fig. 14. Simulated eye diagrams from (a) the 500 Mb/s input to the dc receiver when the ac transmitter is off and (b) the 8 Gb/s input to the ac receiver when the dc transmitter is off.

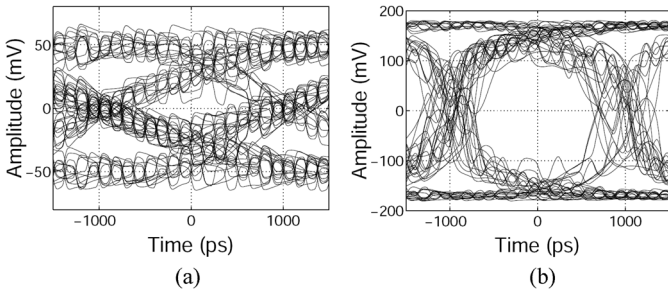


Fig. 15. Simulated 500 Mb/s eye diagrams from the (a) input to the dc receiver and (b) equalized output. Both signaling directions are operating.

Bidirectional simulation results for the ac and dc channels are shown in Figs. 15 and 16. The maximum achievable data rates are 500 Mb/s in the dc band and 8 Gb/s in the ac band, limited by crosstalk between the two links. In both cases, the receiver is set up so that most of the signal comes through the hysteresis path, with only a small amount coming through the linear path. Comparing Fig. 14(a) to Fig. 15(a) and Fig. 14(b) to Fig. 16(a) shows the interference introduced by bidirectional operation.

The link was been simulated using a 10 cm microstrip line and filter values $L = 80$ nH and $C = 10$ pF. These values set the cutoff frequency separating the two bands to 180 MHz. This frequency cannot be increased to permit a higher data rate in the dc link because the ac signal still has significant energy down to a few hundred megahertz. When this energy is filtered out, the received ac signal becomes unrecoverable. In contrast, the ac data rate could be reduced arbitrarily as long as the pulses sent contained sufficient high frequency content.

To allow a lower ratio of forward- to back-channel data rates, the cutoff frequency would have to be moved higher towards the ac data rate. Recovery of ac pulses with significant low-frequency energy removed requires a topology more sophisticated than that of the dual-path receiver shown here.

VI. MEASUREMENT RESULTS

To show the feasibility of the proposed bidirectional communication channel, we present measurement results of the receiver recovering data sent by the transmitter over a short SMA cable




Fig. 16. Simulated 8 Gb/s eye diagrams from the (a) input to the ac receiver and (b) equalized output. Both signaling directions are operating.

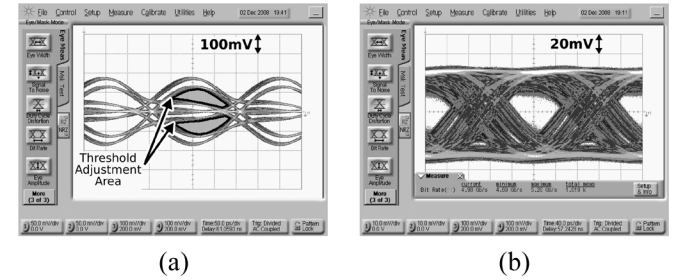


Fig. 17. Measured 5 Gb/s eye diagrams for the (a) input to the ac receiver and the (b) equalized output. Separate transmitter and receiver chips were flip-chip bonded to a printed circuit board. Hysteresis threshold must be set within the threshold adjustment area shown in (a).

channel. In these results the passive filter is imitated by changing the pulse shape used by the programmable transmitter.

For bidirectional operation to be feasible, we must show:

- the receiver can recover high-speed ac data;
- the receiver can recover low-speed dc data;
- the power spectra of the dc and ac signals are sufficiently separated to allow filtering by simple passive elements.

The eye diagrams in Fig. 17 show a 5 Gb/s ac pulse input to the receiver and the open eye at the output of the receiver. To recover the data with no errors, the adjustable hysteresis threshold of the receiver must be set precisely in the threshold adjustment area shown in Fig. 17(a). This hysteresis threshold allows the receiver to distinguish between genuine pulses that should be amplified and noise or reflections that should be attenuated.

For a linear equalizer, an open eye at the output of the receiver is a good indication that no errors are being made. However, because our receiver contains a nonlinear hysteresis block, it is possible to obtain an open eye even when bit errors are present. To verify correct operation, a complete pseudo-random bit stream (PRBS) sequence of length $2^7 - 1$ is shown in Fig. 18 for both the input to the ac receiver and the recovered data.

Recovery of low-speed data can be easily accomplished by the dc receiver. In addition, the spectrum of the transmitted dc signal can be shaped by adjusting the pulse shape produced by the transmitter. Fig. 19(a) and (b) shows the eye diagram of a normal dc pulse along with its power spectrum. By shaping the transmitted pulse to produce longer rise and fall times as shown in Fig. 19(c), the high-frequency content of the dc pulse can be suppressed by 12 dB as shown in Fig. 19(d).




Fig. 18. (a) Received and (b) recovered 5 Gb/s ac-coupled PRBS sequence of length $2^7 - 1$. The eye diagrams for these signals are shown in Fig. 17.




Fig. 19. Eye diagrams and power spectra of two different 500 Mb/s dc pulse shapes. High-frequency content is reduced for the slew-rate limited signal in (c).




Fig. 20. Power spectra of (a) a 500 Mb/s dc signal and (b) a 5 Gb/s ac signal.

By comparing the power spectrum of this 500 Mb/s slew-rate limited dc signal with the power spectrum of a 5 Gb/s ac signal we can see that the majority of the power in each signal occupies a frequency band that is less used by the other signal. The relevant spectra are shown in Fig. 20. The spectrum of the 5 Gb/s ac signal consists of a number of peaks rising out of the noise floor because the signal is a PRBS sequence that repeats every $2^7 - 1$ bits. The peaks occur at frequencies that are integer multiples of $5 \text{ GHz}/(2^7 - 1) = 39.4 \text{ MHz}$. If we plot only the peaks of the two power spectra against a logarithmic frequency scale, we obtain the graph shown in Fig. 21 where we can see that the ac and dc signals occupy distinct frequency bands.




Fig. 21. DC and ac spectra from Fig. 20, plotted on a logarithmic frequency scale.

VII. CONCLUSION

A technique was presented for frequency-division bidirectional communication over chip-to-chip channels. The up- and down-stream signals were separated using simple *LC* elements. The transmitter and receiver circuits employed have the flexibility to be used in both the ac and dc links. In simulation, the bidirectional communication system achieves simultaneous communication of 8 Gb/s over the ac channel and 500 Mb/s over the dc channel. In measurements with a single transmitter-receiver pair, the three objectives of dc communication, ac communication, and frequency separation of the two signals were demonstrated. Future work will focus on a prototype of the complete bidirectional system to explore the limits of this technique.

ACKNOWLEDGMENT

The authors would like to thank CMC Microsystems and Intel Corporation for their help and support.

REFERENCES

- [1] J. F. Bulzacchelli *et al.*, "A 10-Gb/s 5-tap DFE/4-tap FFE transceiver in 90-nm CMOS technology," *IEEE J. Solid-State Circuits*, vol. 41, no. 12, pp. 2885–2900, Dec. 2006.
- [2] K. Krishna *et al.*, "A multigigabit backplane transceiver core in 0.13- μm CMOS with a power-efficient equalization architecture," *IEEE J. Solid-State Circuits*, vol. 40, no. 12, pp. 2658–2666, Dec. 2005.
- [3] V. Stojanovic *et al.*, "Autonomous dual-mode (PAM2/4) serial link transceiver with adaptive equalization and data recovery," *IEEE J. Solid-State Circuits*, vol. 40, no. 4, pp. 1012–1026, Apr. 2005.
- [4] S. Gondi and B. Razavi, "Equalization and clock and data recovery techniques for 10-Gb/s CMOS serial-link receivers," *IEEE J. Solid-State Circuits*, vol. 42, no. 9, pp. 1999–2011, Sep. 2007.
- [5] J. Jaussi *et al.*, "A 20 Gb/s embedded clock transceiver in 90 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2006, pp. 1334–1343.

- [6] G. Balamurugan *et al.*, "A scalable 5–15 Gbps 14–75 mW low power I/O transceiver in 65 nm CMOS," in *Symp. VLSI Circuits Dig. Tech. Papers*, 2007, pp. 270–271.
- [7] B. Casper *et al.*, "A 20 Gb/s forwarded clock transceiver in 90 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2006, pp. 263–272.
- [8] Y. Hur *et al.*, "Equalization and near-end crosstalk (NEXT) noise cancellation for 20-Gb/s backplane serial I/O interconnections," *IEEE Trans. Microwave Theory Tech.*, vol. 53, no. 1, pp. 246–255, Jan. 2005.
- [9] R. J. Drost and B. A. Wooley, "An 8-Gb/s/pin simultaneously bidirectional transceiver in 0.35- μm CMOS," *IEEE J. Solid-State Circuits*, vol. 39, no. 11, pp. 1894–1908, Nov. 2004.
- [10] D. Saltzman and T. Knight, "Capacitive coupling solves the known good die problem," in *IEEE Multi-Chip Module Conf. (MCMC)*, Mar. 1994, pp. 95–100.
- [11] S. A. Kuhn, M. B. Kleiner, R. Thewes, and W. Weber, "Vertical signal transmission in three-dimensional integrated circuits by capacitive coupling," in *Proc. IEEE Int. Symp. Circuits Syst.*, 1995, vol. 1, pp. 37–40.
- [12] S. Mick, J. Wilson, and P. Franzon, "4 Gbps high-density ac coupled interconnection," in *Proc. IEEE Custom Integrated Circuits Conf.*, 2002, pp. 133–140.
- [13] D. Hopkins *et al.*, "Circuit techniques to enable 430 Gb/s/mm² proximity communication," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2007, pp. 368–369.
- [14] R. J. Drost, R. D. Hopkins, R. Ho, and I. E. Sutherland, "Proximity communication," *IEEE J. Solid-State Circuits*, vol. 39, no. 9, pp. 1529–1535, Sep. 2004.
- [15] J. Ko *et al.*, "An RF/baseband FDMA-interconnect transceiver for reconfigurable multiple access chip-to-chip communication," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2005, vol. 1, pp. 338–339.
- [16] S. Deng *et al.*, "Effects of open stubs associated with plated through-hole vias in backpanel designs," in *IEEE Int. Symp. Electromagn. Compatibil.*, Aug. 2004, vol. 3, pp. 1017–1022.
- [17] K. Yamaguchi *et al.*, "12 Gb/s duobinary signaling with x2 oversampled edge equalization," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2005, pp. 70–71.
- [18] E. Prete, D. Scheideler, and A. Sanders, "A 100 mW 9.6 Gb/s transceiver in 90 nm CMOS for next-generation memory interfaces," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2006, pp. 253–262.
- [19] J. Poulton *et al.*, "A 14-mW 6.25-Gb/s transceiver in 90-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 42, no. 12, pp. 2745–2757, Dec. 2007.
- [20] T. F. Knight and D. B. Salzman, "Method and apparatus for non-conductively interconnecting integrated circuits," U.S. Patent 6 728 113, Apr. 2004.
- [21] A. X. Widmer and P. A. Franaszek, "A dc-balanced, partitioned-block, 8 B/10 B transmission code," *IBM J. Res. Develop.*, vol. 27, no. 5, pp. 440–451, Sep. 1983.
- [22] J. Wilson *et al.*, "Fully integrated ac coupled interconnect using buried bumps," *IEEE Trans. Adv. Packag.*, vol. 30, no. 2, pp. 191–199, May 2007.
- [23] M. Hossain and A. Chan Carusone, "A 14-Gb/s 32 mW ac coupled receiver in 90-nm CMOS," in *Symp. VLSI Circuits Dig. Tech. Papers*, 2007, pp. 32–33.
- [24] M. Bichan and A. Chan Carusone, "A 6.5 Gb/s backplane transmitter with 6-tap FIR equalizer and variable tap spacing," in *IEEE Custom Integrated Circuits Conf. (CICC) Dig. Tech. Papers*, 2008, pp. 611–614.



Mike Bichan (S'04) received the B.A.Sc. degree from the Division of Engineering Science and the M.A.Sc. degree from the Department of Electrical and Computer Engineering, in 2003 and 2006, respectively, from the University of Toronto, Toronto, ON, Canada, where he is currently working toward the Ph.D. degree.

His research is focused on power-scalable transceivers for chip-to-chip communication.



Masum Hossain received the B.Sc. degree in electrical engineering from Bangladesh University of Engineering and Technology, Bangladesh, and the M.Sc. degree from Queen's University, Kingston, ON, Canada, in 2002 and 2005, respectively. During his M.Sc. studies, he worked on K-band wireless receiver in CMOS. He is currently working towards the Ph.D. degree in electrical engineering at University of Toronto, Toronto, ON, Canada.

From September 2007 to January 2008, he was with Intel Circuit Research Laboratory as a graduate intern. His research interests include mixed signal circuits for high-speed chip-to-chip communications, low power VCO, phase interpolator, and clock recovery techniques.



Anthony Chan Carusone (S'96–M'02–SM'08) received the B.A.Sc. and Ph.D. degrees at the University of Toronto, Toronto, ON, Canada, in 1997 and 2002, respectively, during which time he received the Governor-General's Silver Medal.

Since 2001, he has been with the Department of Electrical and Computer Engineering at the University of Toronto where he is currently an Associate Professor. He was a visiting researcher at the University of Pavia, Pavia, Italy, in 2008. He has also served as an occasional consultant to industry in the area of integrated circuits and high-speed communication. In 2008, he worked at Intel Corporation on high-speed chip-to-chip interfaces.

Prof. Chan Carusone was a coauthor of the best paper at the 2005 Compound Semiconductor Integrated Circuits Symposium and the best student paper at the 2007 Custom Integrated Circuits Conference. From 2002 to 2007 he held the Canada Research Chair in Integrated Systems. He is a past Chair of the Analog Signal Processing Technical Committee for the IEEE Circuits and Systems Society, a member of the technical program committee for the Custom Integrated Circuits Conference, and is currently the Editor-in-Chief of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS II: EXPRESS BRIEFS.