

Multi-Gb/s Bit-by-Bit Receiver Architectures for 1-D Partial-Response Channels

Masum Hossain and Anthony Chan Carusone, *Senior Member, IEEE*

Abstract—Low-complexity bit-by-bit detection techniques for 1-D partial-response channels are presented. First, a full-rate detection technique is presented which operates at 3.3 Gb/s consuming 40 mA from a 1.8-V supply with a sensitivity of 40-mV differential. The speed of the full-rate architecture is limited by the settling time of a latch circuit which has to be less than 1 UI. To eliminate this limitation, a novel demuxing technique is introduced. Using the proposed technique, a second architecture achieves 5 Gb/s data rate with the same sensitivity and consuming 62 mA (including output buffer) from 1.8-V supply. Both half-rate and full-rate architectures are also studied in 90-nm CMOS targeting chip-to-chip applications. The implemented full-rate architecture operates at 10 Gb/s consuming 32 mW, whereas the simulated half-rate architecture consumes 50 mW and operates at 16.67 Gb/s.

Index Terms—AC coupling, clockless demuxing, decision feedback equalization (DFE), dicode channel, half-rate, peak detection.

I. INTRODUCTION

THERE ARE many new and emerging applications for dicode (1-D) partial-response signaling. Dicode partial-response signaling was applied to magnetic storage channels [1]. More recently, a similar channel response has been observed in multi-Gb/s wireline communication applications such as passive optical networks (PON) and ac-coupled chip-to-chip links that have spectral nulls at dc. The speed of the receivers for these applications is generally limited by the settling time of a latch circuit. This shortcoming is addressed in this paper with two novelties: First, an improved latch circuit provides faster settling time, and second, a parallel architecture permits the positive- and negative-going pulses to be detected separately, thus alleviating the feedback settling-time requirements on the latches.

One interesting area where partial-response signaling has been applied is chip-to-chip links. For example, it was used for a high-speed multidrop bus with magnetically coupled receivers [2]. Capacitive coupling has also been used in chip-to-chip links within a package [3] and over printed circuit board traces of up to 20 cm in length [4]. AC coupling has also been used for bidirectional signaling [5], as a wireless link for modulated data [6] and for power transfer [7].

Manuscript received September 19, 2008; revised January 27, 2009. First published March 27, 2009; current version published February 02, 2010. This work was supported in part by Intel Corporation and in part by Broadcom. This paper was recommended by Associate Editor A. Demosthenous.

The authors are with the Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON M5T 1X2, Canada (e-mail: masum@eecg.utoronto.ca; tcc@eecg.utoronto.ca).

Digital Object Identifier 10.1109/TCSI.2009.2019391

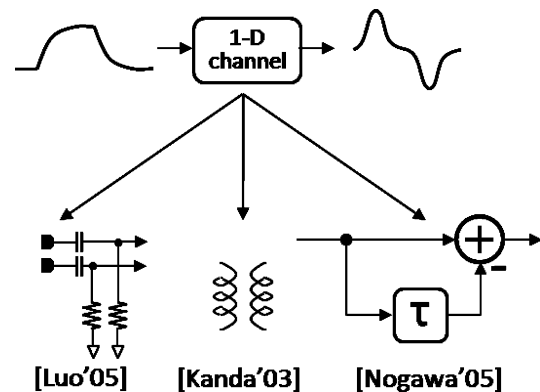


Fig. 1. Recent applications of 1-D partial-response channels.

Another area of interest is burst-mode applications. In a PON system, the receiver at the optical line terminal (OLT) needs to recover data from different optical-network units (ONUs). The packets of data from ONUs arrive in bursts at the OLT end, and their signal strength varies significantly. For high data rates such as 10 Gb/s, the receiver used in the OLT end needs to recover the dc information in less than 1 ns. To avoid the difficulty associated with fast dc extraction, 1-D channel is used to suppress the dc content [8].

Partial-response channel receivers can be broadly classified into two categories: sequence detectors and bit-by-bit detectors. Sequence detectors, such as those using the Viterbi algorithm, make a decision based on a sequence of observations spanning several symbol intervals [9]. Sequence detectors generally outperform bit-by-bit detectors and are now, therefore, dominant in magnetic storage applications. However, they demand sophisticated signal processing and power consumption which are generally intolerable for multi-Gb/s wireline communication applications of (1-D) partial-response signaling. The remainder of this paper will, therefore, focus on bit-by-bit detectors.

All of the multi-Gb/s wireline applications shown in Fig. 1 have behaviorally similar channel responses. The capacitively coupled link in [4], the inductively coupled link in [10], [11], and the burst-mode link in [8] are all dominated by a first-order high-pass characteristic with a cutoff frequency of one to five times the bitrate f_{Bit} . As a result, transitions in the transmitted data appear as narrow electrical pulses at the receiver, while consecutive identical bits result in no signal at the receiver.

Measured and modeled responses of such an ac-coupled channel are shown in Fig. 2 for a 50-fF coupling capacitor and 50- Ω termination resistor. The channel suffers from 40 dB of loss at $0.05 f_{\text{Bit}}$ and more than 15 dB of loss at $2.5 f_{\text{Bit}}$. The measured capacitively coupled channel response closely

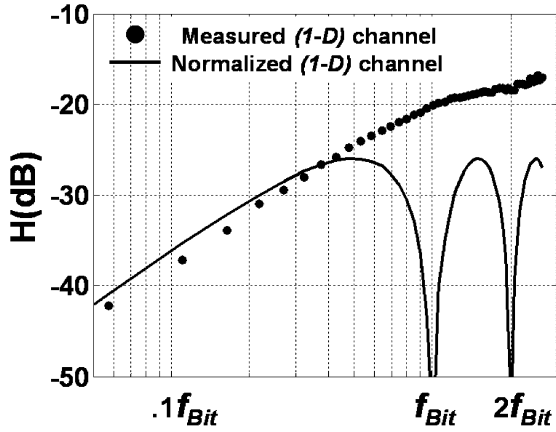


Fig. 2. Frequency responses of (solid line) an ideal 1-D channel and a measured capacitively coupled channel ($R = 50\Omega$, $C = 50$ fF, bit period = 100 ps).

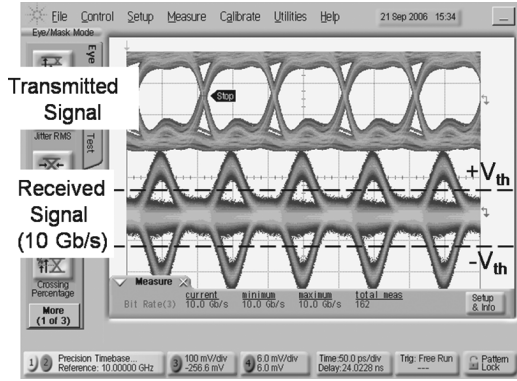


Fig. 3. Measured channel response for a capacitively coupled 1-D channel at 10 Gb/s.

follows an ideal 1-D channel within the band of interest (DC-6 GHz). Corresponding time-domain signals with non-return-to-zero (NRZ) transmitted data are shown in Fig. 3. Note that unlike modern magnetic storage channels, only a small amount of ISI is introduced. However, the sensitivity of the receiver has to be sufficient to capture the small received pulses. In this case, for a 50-fF coupling capacitor and 10-Gb/s data, the signal suffers more than 20 dB of loss, which means that the receiver needs to detect only few tens of millivolts. Since the received pulsewidth is only a fraction of the bit period, receiver circuits will generally require higher gain and bandwidth than an NRZ receiver at the same data rate without ac coupling. Furthermore, the received signal is a three-level signal, so additional decoding is needed to recover the transmitted data.

Different applications of 1-D signaling have different requirements. Chip-to-chip links require low power and area-efficient circuits with moderate sensitivity and dynamic range. On the other hand, a burst-mode application requires higher sensitivity and fast recovery. Without targeting a specific application, we first explore 1-D receiver architectures which can be adapted to meet the requirements of either application. We identify the bottleneck that limits the maximum achievable speed of this type

of receiver. We then propose an improved receiver architecture that obviates this speed limitation.

These receiver architectures are implemented and compared for two particular applications: In the first implementation, we target a 40-mV sensitivity and a >10-dB dynamic range which is required for gigabit PON systems [12]. The two receivers implemented in 0.18- μm CMOS serve as an experimental validation of the theoretical discussion regarding full-rate and half-rate architectures. Their relatively high sensitivity requires large preamp gain, hence increases power consumption. Thus, the implemented full-rate and half-rate prototypes in 0.18 μm CMOS consumes 72 and 110 mW for 3.33 and 5 Gb/s, respectively. Although this power consumption is comparable with other existing burst-mode receivers [13], chip-to-chip links require much lower power consumption.

For chip-to-chip interconnects, we target an 80-mV sensitivity and higher bitrate (10+ Gb/s). To improve power efficiency and achieve higher speed, we implement the proposed receivers in 90-nm CMOS. An implemented full-rate architecture consumes 32 mW and operates up to 10 Gb/s without any equalization. On the other hand, simulated half-rate architecture consumes 50 mW and operates up to 16.67 Gb/s. Achieved power efficiencies of 3.2 and 3.0 mW/Gb/s are comparable with the dc-coupled receivers at these speeds.

We begin our discussion with general bit-by-bit decoding techniques for dicode channels in Section II. The proposed half-rate architecture, introduced at the end of Section II, relaxes the speed bottleneck introduced by feedback in full-rate architectures. In Section III, a full-rate implementation in 0.18 μm CMOS is described. It is based on the architecture introduced in [14] but modified to accommodate threshold adjustability and improve sensitivity. Section IV describes the circuit-level implementation of the half-rate architecture introduced in Section II. Using this technique, the receiver can potentially improve the speed by a factor of two, at the expense of increased power consumption. The 0.18- μm CMOS prototype operates up to 5 Gb/s, 50% faster than the full-rate architecture. Targeting chip-to-chip applications, 90-nm implementation of these architectures are presented in Section V. Finally, these two architectures are compared in the conclusion in Section VI.

II. 1 - D RECEIVER ARCHITECTURE

In this section, dicode (1-D) partial-response bit-by-bit receiver architectures are reviewed.

A. DFE

For uncoded binary data transmitted over a 1-D channel, the data can be recovered using a one-tap decision feedback equalizer, as shown in Fig. 4[15], [16]. In this architecture, the received signal $s(n) = z(n) - z(n-1)$ is compared with a threshold level that is updated based on the immediate previous bit

$$v_{\text{th}}(n) = \beta V(n-1). \quad (1)$$




Fig. 4. Dicode receiver using a DFE for symbol-by-symbol detection.




Fig. 5. 1-D partial-response signaling implemented with a precoder on the transmitter side and a peak detector as a decoder on the receiver side.

A hardware-efficient implementation of this technique is discussed in [3] and [4] where this functionality can be achieved at high speed utilizing a hysteresis latch. Compared with a conventional decision feedback equalization (DFE), this architecture provides several advantages: 1) Since there is no clock required, this architecture can be implemented with less complexity and lower power consumption and 2) since there is no delay flip-flop in the feedback path, it will settle faster than a clocked one-tap DFE. However, there is still a feedback path that must settle, and the highest achievable speed of this architecture is generally also limited by the settling time of that loop which must be less than 1 UI.

B. Full-Rate Precoder and Decoder

A precoding method for 1-D partial-response channel is proposed in [1]. The precoder output y is related to the data to be transmitted z by

$$y(n) = z(n) \oplus y(n-1). \quad (2)$$

One possible architecture is shown in Fig. 5. A similar architecture was presented for a duobinary (1-D) channel in [17] and, more recently, has been discussed in [18]. The decoder converts the three-level received signal $s(n)$ to a two-level binary output using the following decision criteria:

$$|S(t) > |V_{th}| \implies v(n) = 1 \quad (3)$$

$$|S(t) < |V_{th}| \implies v(n) = 0. \quad (4)$$

This is accomplished using a conventional peak-detection receiver shown in Fig. 5. It may be shown that, in the absence of noise causing decision errors, $v(n) = z(n)$.

In some applications, such as burst-mode PON systems, where the transmitter does not provide the precoding functionality, the precoder can be moved to the receiver side as shown in Fig. 6. The thresholds $\pm V_{th}$ are chosen so that the received signal amplitude $|s|$ only exceeds $|V_{th}|$ when there has been a

transition in the transmitted bit stream. [i.e., $z(n) \neq z(n-1)$]. In this case

$$u_1(n) = z(n)\overline{z(n-1)} \quad (5)$$

$$u_2(n) = \overline{z(n)}z(n-1). \quad (6)$$

Hence, the output of the first XOR operation is

$$w(n) = u_1(n) \oplus u_2(n) = z(n) \oplus z(n-1) \quad (7)$$

and the decoder output $v(n)$ is

$$v(n) = v(n-1) \oplus w(n) = v(n-2) \oplus w(n-1) \oplus w(n). \quad (8)$$

Using (7) and (8), the decoded output $v(n)$ can be expressed as a function of the transmitted symbol $z(n)$

$$v(n) = z(n) \oplus v(n-2) \oplus z(n-2). \quad (9)$$

This equation can be iteratively extended back in time to the first transmitted symbol, which we shall assign to time $n = 0$

$$v(n) = z(n) \oplus v(0) \oplus z(0). \quad (10)$$

Thus, if the initial transmitted symbol $z(0)$ and the initial decoder output $v(0)$ are the same, (10) reduces to $v(n) = z(n)$, and the decoder output is indeed equal to the transmitted data. Comparing the transceiver architectures in Figs. 4–6, the highest achievable speed is always limited by the delay of a feedback loop which must be a 1-bit period or less.

C. Receiver With Half-Rate Decoder

To ease the settling-time requirements of all the previous architectures, we introduce the half-rate decoder shown in Fig. 7. This architecture is a natural progression from the one shown in Fig. 6, where the feedback loop is shifted before the XOR operation. The operation of the half-rate receiver in Fig. 7 is best understood by recognizing that the top path, through u_1 and w_1 , is responsible for receiving positive peaks in s , whereas the bottom path, through u_2 and w_2 , receives only the negative peaks in s . Every positive peak in s (corresponding to a rising edge of z) must be followed by a negative peak (corresponding to a falling edge of z). Hence, the top path can never be active in two consecutive bit periods. Similarly, all negative peaks are followed by a positive peak, so the negative path is never active for two bit periods in a row. Hence, the feedback loops have twice as long to settle, i.e., 2 UI. The front-end of the receiver is unchanged from Fig. 6, so (5) and (6) are still valid.

Now, the decoder output $w_1(n)$ and $w_2(n)$ can be written as

$$w_1(n) = w_1(n-1) \oplus (z(n)\overline{z(n-1)}) \quad (11)$$

$$w_2(n) = w_2(n-1) \oplus (\overline{z(n)}z(n-1)). \quad (12)$$

The full-rate decoded output $v(n)$ is then related to $w_1(n)$ and $w_2(n)$ as follows:

$$v(n) = w_1(n) \oplus w_2(n) \quad (13)$$

$$= (w_1(n-1) \oplus (z(n)\overline{z(n-1)})) \oplus (w_2(n-1) \oplus (\overline{z(n)}z(n-1))). \quad (14)$$




Fig. 6. Modified 1-D partial-response receiver with the precoder moved to the receiver.

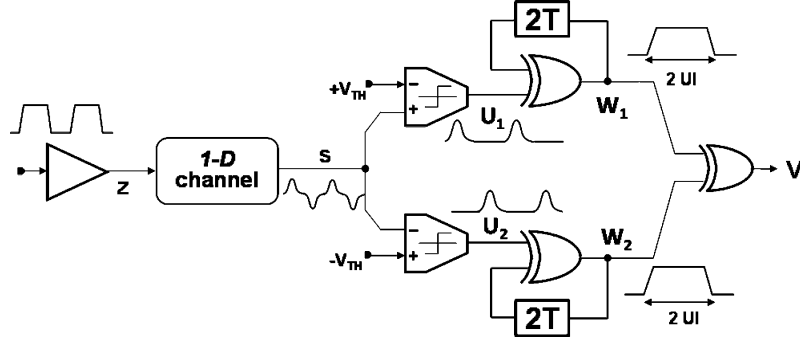


Fig. 7. Proposed half-rate receiver architecture for 1-D partial-response signaling.

Note that

$$\left(z(n) \overline{z(n-1)} \right) \oplus \left(\overline{z(n)} z(n-1) \right) = z(n) \oplus z(n-1). \quad (15)$$

Substituting (15) into (14)

$$v(n) = w_1(n-1) \oplus w_2(n-1) \oplus z(n) \oplus z(n-1) \quad (16)$$

$$= v(n-1) \oplus z(n) \oplus z(n-1) \quad (17)$$

similarly, $v(n-1)$ can be written as

$$v(n-1) = v(n-2) \oplus z(n) \oplus z(n-2) \quad (18)$$

substituting (18) into (17) results in

$$v(n) = z(n) \oplus v(0) \oplus z(0). \quad (19)$$

Thus, the decoder can correctly recover the transmitted symbol $v(n)$ if the initial decoder output and the initial transmitted symbol are the same, $v(0) = z(0)$, the same requirement obtained for the full-rate architecture in Fig. 6. Notice that pulses are generated at $u_1(u_2)$ whenever a rising (falling) edge is observed on the channel data z . Hence, these signals can be used as inputs to a phase detector in a conventional clock-recovery loop. Alternatively, we use them to injection lock an oscillator in the simulations of Section V.

In summary, present low-complexity 1-D decoders generally include feedback loops which must settle in less than 1 UI. However, this requirement can be relaxed by using the half-rate architecture proposed in Fig. 7. The remainder of this paper describes a prototype of the half-rate decoder and compares it with a full-rate decoder in the same technology. Error propagation of such receivers is the same as a conventional DFE. Just as in DFE-based partial-response receivers, so long as a sufficiently low bit error rate (BER) is maintained, there is no observable degradation performance.

III. FULL-RATE BIT-BY-BIT DETECTION

The receiver architecture shown in Fig. 8(a) was introduced in [14]. Notice the linear amplifier in parallel with the hysteresis latch, which improves the receiver's overall speed. With this in place, the receiver's speed is determined by the settling time of the hysteresis latch and the bandwidth of the preamplifier.

A. Implementation

In the hysteresis latch, the received signal is compared with a threshold level provided by a feedback path. The polarity of the threshold is determined by the most recently detected bit. The circuit used for this purpose in [8] is shown in Fig. 8(b). This circuit demonstrates hysteresis if the following condition is satisfied:

$$g_m R_L > 1 \quad (20)$$

where g_m is the small-signal transconductance of the feedback differential pair. In practice, to ensure operation in the presence of noise and process variations, $g_m R_L$ is made nominally greater than two. In addition, to accommodate both large and small inputs, the threshold levels should be adjustable. However, this simple circuit suffers from two main challenges. First, the critical output node is heavily loaded by the capacitance of g_m and the following stages which limits its settling time.

To reduce the time constant at the critical nodes, cascode devices were used. Due to this transistor stacking, VDD was increased to 2.5 V in a 0.13- μm CMOS process. In this paper, the time constant is improved within the process nominal VDD. The second challenge with using the hysteresis circuit in Fig. 8(b) is that adjusting the threshold level will also effect other aspects of the design, such as its settling time. Hence, to provide programmability, several copies of the circuit were operated in parallel in [13]. The proposed circuit is shown in Fig. 8(c). An additional differential pair g_{m1} is introduced in the latch which

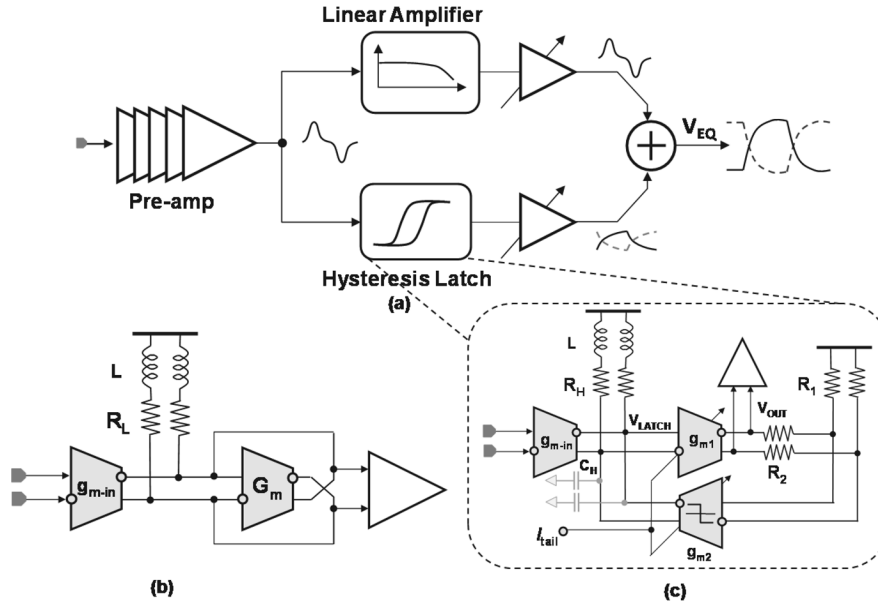


Fig. 8. (a) Receiver architecture from [14]. (b) Hysteresis latch from [8]. (c) Proposed hysteresis latch.

provides several advantages: First, note that the condition for hysteresis is now

$$g_{m1}R_1g_{m2}R_H > 1. \quad (21)$$

Compared with the condition for the previous circuit in (22), there is additional flexibility to choose the gain of each stage, R_1 and R_H , to minimize the settling time. Second, g_{m1} also works as a buffer between the critical node and the following stages. Finally, this architecture also allows adjustment of the threshold levels, as shown in Fig. 9(a).

In addition, note the use of a “split load” at the output [19] so that the feedback is taken from the fast-settling node with low impedance, while the output is taken from the node with larger swing. Hence, the feedback-loop settling time is dominated by the time constant $R_H C_H$, whereas the “output settling time” is dependent on the time constant $(R_1 + R_2)C_L$. This allows design flexibility and relaxes the tradeoff between speed, sensitivity, and noise immunity. Simulations of the hysteresis latch in Fig. 9(b) indicate that the split load improves the settling time by >20% in this circuit. The targeted sensitivity of the receiver is a 40-mV differential input. The hysteresis comparator thresholds can be adjusted from 150 to 400 mV differential input. In PON applications, the receiver threshold can be adjusted based on a training preamble which precedes each data burst.

A five-stage preamplifier providing 24-dB differential gain is used in front of the hysteresis latch. Budgeting 30 mW of power for the preamplifier, without inductive peaking, the achieved bandwidth is only 2.5 GHz, resulting in excessive data-dependent jitter. Hence, inductive peaking was used to extend the bandwidth to 3.5 GHz.

B. Experimental Results

A die photograph of the receiver front end is shown in Fig. 10. Measurements were made with a channel comprising an approximately 3-ft long SMA cable and a 50-fF ac-coupling ca-

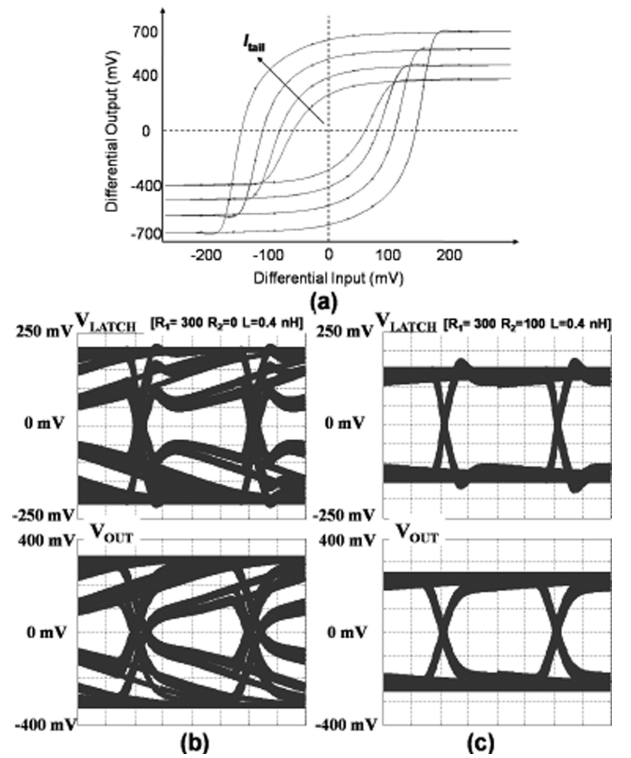

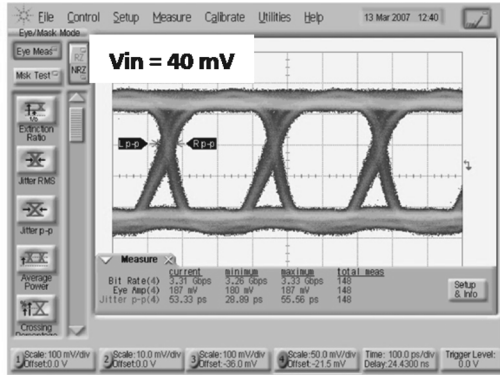
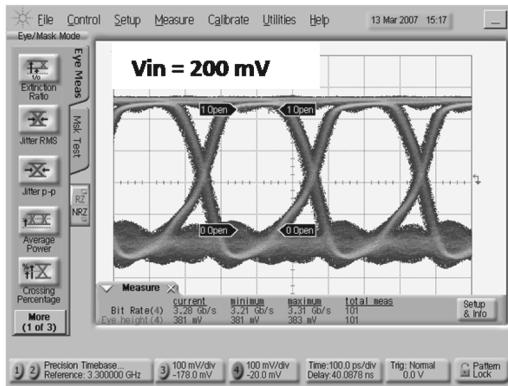


Fig. 9. Hysteresis latch simulations: (a) Threshold adjustments by changing I_{tail} . (b)–(c) Improvement of threshold and output settling time with resistor splitting: (b) without resistor splitting and (c) with resistor splitting.

pacitor on-chip which, together with the 50- Ω on-chip termination, forms the high-pass filter characterized in Figs. 2 and 3. The measured results are obtained with single-ended excitation only. The receiver’s dynamic range was tested by varying the input amplitude from 40 to 200 mV. For a 40-mV input, the threshold level was adjusted to ± 70 mV, and for a 180 mV input, the threshold level was adjusted to ± 180 mV. The receiver demonstrated error-free data recovery at 3.3 Gb/s for a

Fig. 10. Full-rate receiver die photograph in 0.18 μm CMOS.

(a)




(b)

Fig. 11. Measured output eye of the full-rate receiver at 3.3 Gb/s for different input amplitude (a) 40 mV (b) 200 mV.

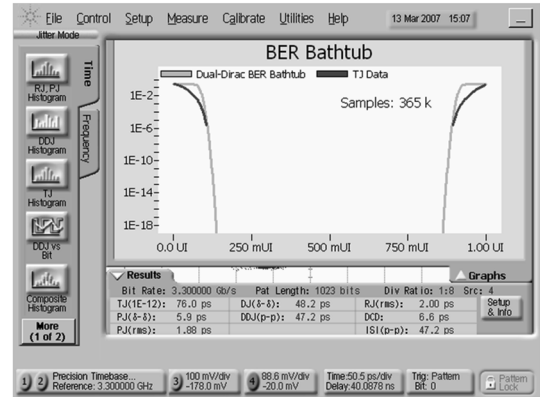
pseudorandom binary sequence (PRBS) $2^{10} - 1$ pattern at both signal amplitudes [Fig. 11(a) and (b)].

IV. HALF-RATE DETECTION

The speed of the architecture in Section III is limited by the finite bandwidth of the preamplifier and the threshold settling time. To further increase the speed of the (1-D) partial-response receivers, the parallel half-rate architecture described in Section II is used. The block diagram of a CMOS implementation of this architecture is shown in Fig. 13. The front end is composed of two major circuit blocks: a slicer and a toggle flip-flop (T-FF). The T-FF provides the feedback and XOR operation shown in each path of Fig. 7. The circuit outputs $Demux1$ and $Demux2$ correspond to w_1 and w_2 in Fig. 7. These may be XORed to recover the full-rate data or further demultiplexed for



(a)



(b)

Fig. 12. Results for a PRBS $2^{10} - 1$ pattern: (a) A segment of the transmitted and recovered sequences and (b) BER bathtub plot.

digital decoding at a much slower rate. One possible implementation is shown in Fig. 13, where recovered half-rate clock is used to further demultiplex $Demux1$ and $Demux2$. These demuxed bitstreams are then XORed at a half-rate to decode the even and odd bitstreams. Half-rate clock can be recovered from the transition information provided by u_1 and u_2 .

A. Implementation

The first stage of the slicer is a differential difference amplifier that compares the input with $\pm V_{th}$. The detected pulses are then passed through five inductively peaked amplifier stages providing 26-dB gain. Fortunately, due to the half-rate architecture, lower bandwidth can be tolerated here than in the full-rate preamplifier. Hence, the total current consumption is only 9 mA from an 1.8-V supply for each amplifier chain providing a 2.2-GHz bandwidth.

High-speed T-FFs have been widely used as dividers in both wireline and wireless applications. Conventional current-mode logic T-FFs employ two back-to-back D latches as shown in Fig. 13(d). A typical implementation of the D latch is shown in Fig. 14(a). This type of T-FF exhibits self-oscillation which allows it to operate as a high-frequency divider. A typical sensitivity curve is shown in Fig. 14(a). Unfortunately, noise around the self-oscillation frequency can cause the output to toggle erroneously during periods when there is no transition in the received data, resulting in bit errors in the decoded sequence. Thus, self-oscillation in the T-FF must be avoided to use it as a

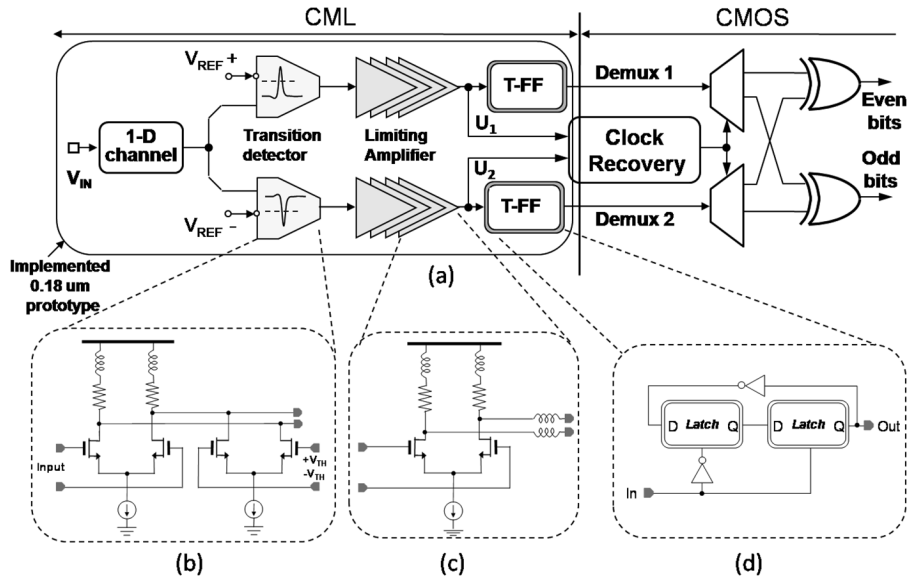


Fig. 13. (a) Proposed half-rate receiver architecture. (b) Transition detector circuit. (c) Building block of the five-stage preamp (d) T-FF.

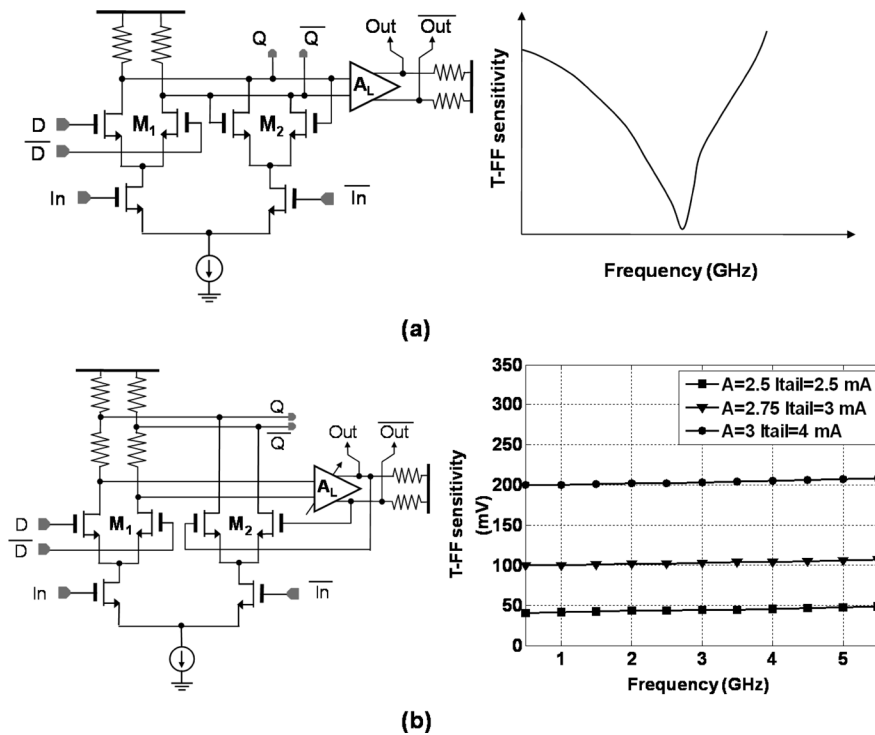


Fig. 14. (a) Conventional D latch and corresponding T-FF sensitivity. (b) Proposed D Latch and simulated T-FF sensitivity.

decoder in this application. In addition, the buffer A_L is needed to drive a capacitive load without loading the latch nodes.

To alleviate both of these problems, we bring the buffer within the feedback loop as shown in Fig. 14(b). The gain of A_L is easily made adjustable to allow variable latching strength and, hence, T-FF sensitivity. The modified architecture provides frequency-independent sensitivity characteristics, as shown in Fig. 14(b). Furthermore, stage A_L effectively buffers the critical latch node and eliminates the requirement of an additional buffer.

Thus, the proposed latch circuit does not consume additional power compared with a conventional latch implementation.

B. Experimental Results

A prototype receiver in 0.18- μm CMOS is shown in Fig. 15. In this implementation, we used the same 50-fF coupling capacitor and 50- Ω resistance as a high-pass filter to provide the 1-D partial response. The receiver provided error-free operation for a PRBS $2^{10} - 1$ pattern up to 5 Gbits/s. Eye diagrams of the demultiplexed outputs at 3.33 and 5 Gb/s are shown in Figs. 16




Fig. 15. Half-rate receiver die photograph in 0.18- μm CMOS.

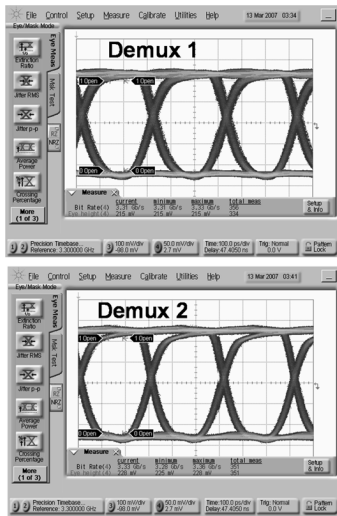


Fig. 16. Measured demuxed eye at 3.3 Gb/s.

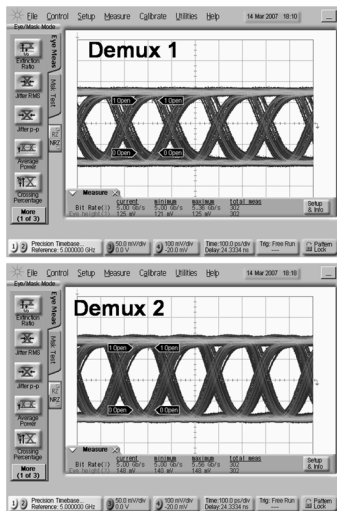


Fig. 17. Measured demuxed eye at 5 Gb/s.

and 17, respectively. Portions of actual PRBS $2^{10} - 1$ transmitted and recovered sequences are shown in Fig. 18. Note that none of the eye diagrams show any ringing which indicates that the proposed T-FF implementation strongly suppresses self-oscillation.

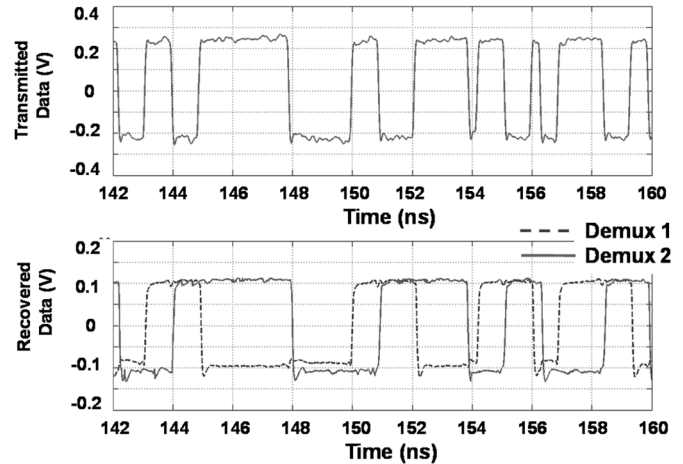


Fig. 18. Transmitted and demuxed data streams: demuxed data streams are overlaid to demonstrate the decoding functionality.

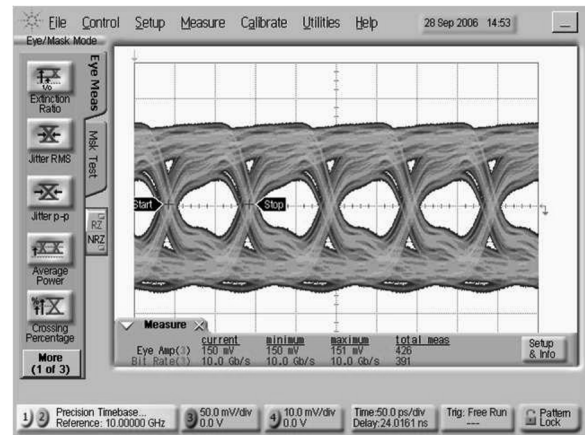


Fig. 19. Recovered 10 Gb/s NRZ eye from full-rate receiver implemented in 90-nm CMOS.

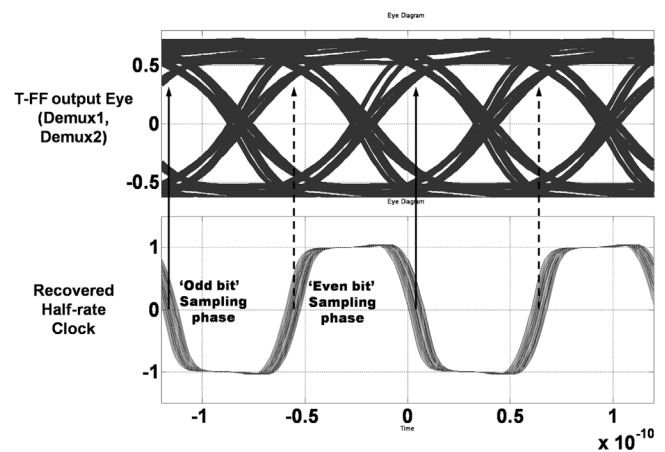


Fig. 20. Transistor level simulation results of the 90-nm half-rate receiver at 16.67 Gb/s. Even samples are generated by (dashed arrow) the rising edge of the recovered clock and odd samples are generated by (solid arrow) the falling edge.

V. 90-nm IMPLEMENTATION

For chip-to-chip applications, the sensitivity and dynamic range requirements are relaxed. We target an 80-mV sensitivity in a 90-nm CMOS process resulting in greatly improved

TABLE I
COMPARISON OF STATE-OF-ART 1-DRECEIVERS

	[13]	[4]	This work	This work	This work	This work
	Full-rate	Full-rate	Full-rate	Half-rate	Full-rate	Half-rate (Simulated)
Technology	0.13 μ m	0.18 μ m	0.18 μ m	0.18 μ m	90nm	90nm
Channel	On-chip L-R	Proximity coupled	On-chip C-R	On-chip C-R	On-chip C-R	On-chip C-R
Pre-amp Gain	26 dB	—	23 dB	22 dB	17 dB	17 dB
Bit-rate	10 Gb/s	3 Gb/s	3.33 Gb/s	5Gb/s	10 Gb/s	16.67 Gb/s
VDD	2.5 V	1.8 V	1.8 V	1.8 V	1.2 V	1.2 V
Receiver Power Consumption	500 mW Includes Logic+ Buffer power	10 mW	72 mW Includes buffer power	110 mW Includes buffer power	32 mW	50 mW (With Clock Recovery 110 mW)
Area	2.72mm ²	—	0.97mm ²	1.08mm ²	0.045mm ²	—
Sensitivity	40 mV	120 mV	40 mV	40 mV	80 mV	80 mV
FoM(pJ/Bit)	50	3.33	21.8	22	3.2	3.0

power-efficiency. No inductors were used in these designs as chip-to-chip applications demand compact circuitry.

A. Full-Rate Bit-by-Bit Detection

The full-rate architecture of Section III achieves an 80-mV sensitivity with a five-stage preamplifier that consumes only 10 mW. The hysteresis latch consumes 7 mW and operates up to 10 Gb/s. For experimental study, we used the implemented full-rate receiver in [14]. Similar to Section III, only nonlinear path is used for NRZ recovery. Due to the relaxed dynamic range, we can use a fixed threshold in the latch. Recovered NRZ data measured with the same channel as in Sections III and IV are shown in Fig. 19 at 10 Gb/s. The total power consumption is 32 mW from a 1.2-V supply.

B. Half-Rate Receiver

Simulations of the half-rate architecture in 90-nm CMOS are used to demonstrate the following: 1) A speed improvement over the full-rate architecture, commensurate with that observed in 0.18- μ m CMOS, is possible and 2) clock recovery and 1:2 demultiplexing is readily feasible within this architecture. The same circuits described in Fig. 13 are ported to a 90-nm process. Following the T-FF, all remaining circuitry is implemented using full-swing CMOS logic. The spectrum of the signals u_1 and u_2 contain tones at the baud rate which can be utilized for clock recovery using a phase-locked loop. Phase locking can also be done using an injection-locked oscillator or gated voltage-controlled oscillator (VCO) which provides the fast locking required for burst-mode applications. Injection locking a half-rate clock relaxes the VCO design compared with a full-rate VCO. In the proposed architecture, the signals u_1 and u_2 are used to injection lock a half-rate ring oscillator which operates at 8.33 GHz. The recovered half-rate clock is then used to demultiplex and retime the data. Proper recovery of the even and odd data is demonstrated using this technique in simulations at 16.67 Gb/s in Fig. 20. This represents a 67% increase in data rate over the full-rate measurements, which is very comparable with the measurement results from the 0.18- μ m prototypes, where the half-rate architecture offered a

50% increase in data rate from 3.3 to 5 Gb/s. The total simulated power consumption, including clock recovery, demultiplexing, and required logic, is 110 mW from a 1.2-V supply.

VI. CONCLUSION

In recent years, there has been significant effort to improve sensitivity and speed of ac coupled receivers. This trend is driven by desires to use small ac-coupling capacitors, achieve higher data rates, and/or accommodate lossy channels. Thus, high preamp gain and bandwidth are required at the cost of additional power and area. A more detailed comparison of the proposed receivers and state-of-the-art receivers with high sensitivity is given in Table I. Sensitivity is measured by the minimum signal amplitude required for error-free detection at the receiver. For comparison of different implemented receivers, we used a figure of merit (FoM), which is defined as

$$\text{FoM(pJ/Bit)} = \frac{\text{Power consumption}}{\text{Bitrate}}. \quad (22)$$

Compared with the 10-Gb/s receiver, the presented full-rate and half-rate receivers achieve similar sensitivity with significant power reduction. On the other hand, power and area efficiency can be further improved in the 90-nm implementation. Compared with full-rate architectures, the proposed half-rate architecture can potentially achieve twice the speed at the cost of additional hardware complexity and power. In this paper, a 50% improvement in speed is achieved at the cost of a 30% increase in power. Clearly, the half-rate architecture is particularly useful when the targeted speed is not achievable using full-rate architectures.

Another potential application of the half-rate architecture is for clockless demultiplexing which was proposed for burst-mode applications in [20] to relax the lock-time requirements of the subsequent clock and data-recovery circuitry. In [20], a finite state machine performs sophisticated processing resulting in high power consumption. On the other hand, the half-rate receiver proposed in this paper demultiplexes the bitstream based on edge detection which is actually performed by the passive 1-D channel, thus providing reduced power consumption.

ACKNOWLEDGMENT

The authors would like to thank CMC for providing the fabrication facilities.

REFERENCES

- [1] M. Kobayashi and D. T. Tang, "Application of partial-response channel coding to magnetic recording systems," *IBM J. Res. Develop.*, vol. 14, no. 4, pp. 368–375, Jul. 1970.
- [2] J. Benham, R. Amirtharajah, J. L. Critchlow, T. Simon, and T. F. Knight, Jr., "An alignment insensitive separable electromagnetic coupler for high-speed digital multidrop bus applications," *IEEE Trans. Microw. Theory Tech.*, vol. 51, no. 12, pp. 2597–2603, Dec. 2003.
- [3] R. Drost, R. Hopkins, R. Ho, and I. Sutherland, "Proximity communication," *IEEE J. Solid-State Circuits*, vol. 39, no. 9, pp. 1529–1535, Sep. 2004.
- [4] L. Luo, J. M. Wilson, S. E. Mick, J. Xu, L. Zhang, and P. D. Franzon, "A 3 Gb/s AC coupled chip-to-chip communication using a low swing pulse receiver," in *Proc. IEEE ISSCC Dig. Tech. Papers*, Feb. 2005, pp. 522–523.
- [5] A. Fazzi, R. Canegallo, L. Ciccarelli, L. Magagni, F. Natali, E. Jung, P. L. Rolandi, and R. Guerrieri, "3D capacitive interconnections with mono- and bidirectional capabilities," in *Proc. IEEE ISSCC Dig. Tech. Papers*, Feb. 2007, pp. 356–357.
- [6] Q. Gu, Z. Xu, J. Ko, and M.-C. F. Chang, "Two 10 Gb/s/pin low-power interconnect methods for 3-D ICs," in *Proc. IEEE ISSCC Dig. Tech. Papers*, Feb. 2007, pp. 448–449.
- [7] C. Calurciello and A. G. Andreou, "Capacitive inter-chip data and power transfer for 3-D VLSI," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 53, no. 12, pp. 1348–1352, Dec. 2006.
- [8] M. Nogawa, K. Nishimura, S. Kimura, T. Yoshida, T. Kawamura, M. Togashi, K. Kumozaki, and Y. Ohtomo, "A 10 Gb/s burst-mode CDR IC in 0.13 μm CMOS," in *Proc. ISSCC Dig. Tech. Papers*, Feb. 2005, pp. 228–229.
- [9] G. D. Forney, "Maximum-likelihood sequence estimation of digital sequences in the presence of intersymbol interference," *IEEE Trans. Inf. Theory*, vol. IT-18, no. 3, pp. 363–378, May 1972.
- [10] K. Kanda, D. Antono, K. Ishida, H. Kawaguchi, T. Kuroda, and T. Sakura, "1.27 Gb/s/pin 3 mW/pin wireless superconnect (WSC) interface scheme," in *Proc. IEEE ISSCC Dig. Tech. Papers*, Feb. 2003, pp. 186–187.
- [11] N. Miura, Y. Kohama, Y. Sugimori, H. Ishikuro, T. Sakurai, and T. Kuroda, "An 11 Gb/s inductive-coupling link with burst transmission," in *Proc. IEEE ISSCC Dig. Tech. Papers*, Feb. 2008, pp. 298–299.
- [12] T. Nakanishi, K.-I. Suzuki, Y. Fukada, N. Yoshimoto, M. Nakamura, K. Kato, K. Nishimura, Y. Ohtomo, and M. Tsubokawa, "High sensitivity APD burst-mode receiver for 10 Gb/s TDM-PON system," *IEICE Electron. Express*, vol. 4, no. 19, pp. 588–592, Oct. 2007.
- [13] M. Nogawa, Y. Ohtomo, S. Kimura, K. Nishimura, T. Kawamura, and M. Togashi, "A 10 Gb/s burst-mode adaptive gain select limiting amplifier in 0.13 μm CMOS," in *Proc. ISSCC Dig. Tech. Papers*, Feb. 2006, pp. 940–949.
- [14] M. Hossain and A. C. Carusone, "A 14 Gb/s 32 mW AC coupled receiver in 90-nm CMOS," in *Proc. VLSI Circuits Symp.*, Kyoto, Japan, Jun. 2007, pp. 32–33.
- [15] R. W. Wood and R. W. Donaldson, "Decision feedback equalization of the DC null in high-density digital magnetic recording," *IEEE Trans. Magn.*, vol. MAG-14, no. 4, pp. 218–222, Jul. 1978.
- [16] J. W. M. Bergmans, "Decision feedback equalization for digital magnetic recording systems," *IEEE Trans. Magn.*, vol. 24, no. 1, pp. 683–688, Jan. 1988.
- [17] J. H. Sinsky, A. Adamiecki, and M. Duelk, "10 Gb/s electrical backplane transmission using duobinary signaling," in *Proc. IMS Dig. Tech. Papers*, Jun. 2004, pp. 109–112.
- [18] M. R. Ahmadi, J. Moon, and R. Harjani, "Constrained partial response receivers for high-speed links," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 55, no. 10, pp. 1006–1010, Oct. 2008.
- [19] Y. M. Greshishchev and P. Schvan, "A 60 dB gain 55 dB dynamic range 10 Gb/s broadband SiGe HBT limiting amplifier," in *Proc. IEEE ISSCC Dig. Tech. Papers*, Feb. 1999, pp. 382–383.
- [20] B. Analui and A. Hajimiri, "Instantaneous clockless data recovery and demultiplexing," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 52, no. 8, pp. 437–441, Aug. 2005.



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

From September 2007 to January 2008, he was with Intel Circuit Research Lab as a Graduate Intern. Currently, he is with the Analog and Mixed Signal Division, Gennum Corporation, Toronto. His research interest includes mixed-signal circuits for high-speed chip-to-chip communications, low power voltage-controlled oscillator, phase interpolator, and clock-recovery techniques.

Mr. Hossain won the Best Student Paper award in the 2008 IEEE Custom Integrated Circuits Conference.



Anthony Chan Carusone (SM'96–M'02–S'08) received the B.A.Sc. and Ph.D. degrees from the University of Toronto, Toronto, ON, Canada, in 1997 and 2002, respectively.

In 2008, he was a Visiting Researcher with the University of Pavia, Pavia, Italy, and later with the Circuits Research Laboratory, Intel Corporation, Hillsboro, OR. Since 2001, he has been with the Department of Electrical and Computer Engineering, University of Toronto, where he is currently an Associate Professor.

Prof. Carusone is a member and past Chair of the Analog Signal Processing Technical Committee for the IEEE Circuits and Systems Society, a member and past Chair of the Wireline Communications subcommittee of the Custom Integrated Circuits Conference, and an appointed member of the Administrative Committee of the IEEE Solid-State Circuits Society. He has served as a Guest Editor for both the IEEE JOURNAL OF SOLID-STATE CIRCUITS and the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS I—REGULAR PAPERS. He is currently Editor-in-Chief of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS II—EXPRESS BRIEFS. While with the University of Toronto, he was the recipient of the Governor General's Silver Medal. As a coauthor, he was the recipient of the Best Paper award at the 2005 Compound Semiconductor Integrated Circuits Symposium and the Best Student Papers award at both the 2007 and 2008 Custom Integrated Circuits Conferences.