A 20 Gb/s CMOS Optical Receiver With Limited-Bandwidth Front End and Local Feedback IIR-DFE

Alireza Sharif-Bakhtiar, *Student Member, IEEE*, and Anthony Chan Carusone, *Senior Member, IEEE*

Abstract—Implementation of highly integrated optical receivers in CMOS promises low cost, but combining high gain, low noise, high bandwidth, and low power in a CMOS transimpedance amplifier is a challenge. Fortunately, the sensitivity of an optical receiver is improved by limiting its front-end bandwidth far below the symbol rate and using equalization to eliminate the resulting intersymbol interference (ISI). Analysis reveals that when using a decision-feedback equalizer (DFE) to cancel all postcursor ISI, receiver sensitivity is optimized by taking a front-end bandwidth as low as $0.12 f_{bit}$, depending upon the frequency response and noise spectrum assumed for the front end. This paper presents a 20 Gb/s optical receiver with a front-end bandwidth of 3 GHz. The front end is designed to have an approximately first-order response, ensuring only postcursor ISI, which may be efficiently canceled with a first-order infinite-impulse response DFE (IIR-DFE). An IIR-DFE circuit is also proposed that obviates the need for an explicit full-rate multiplexor. Fabricated in 65 nm CMOS, the receiver achieves 0.705 pJ/b efficiency with the IIR-DFE consuming 150 fJ/b. Using a photodiode with 12 GHz analog bandwidth and responsivity of 0.5 A/W, the receiver has a sensitivity of −5.8 dBm optically modulated amplitude.

Index Terms—Decision-feedback equalizer (DFE), infinite-impulse response, low power, optical interconnects, optical receiver, vertical cavity surface-emitting laser (VCSEL).

I. INTRODUCTION

THE improving power efficiency and cost of optical communication links have made them a suitable candidate to replace copper for 10+Gb/s links as short as 10 m or even less in high-performance computing and networking applications. In particular, links-based upon direct modulation of vertical cavity surface-emitting lasers (VCSELs) provide the lowest cost optoelectronic components and packaging. Typically, the directly modulated VCSEL is coupled to a multimode fiber, which transfers the light to a discrete photodiode. Many state-of-the-art VCSEL-based links rely upon SiGe circuits [1], [2] necessitating separate dies for the front-end circuits and CMOS digital processing, increasing packaging complexity
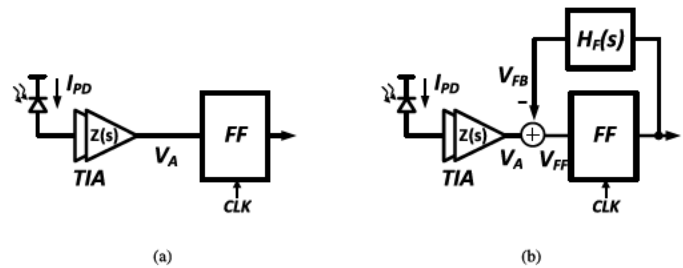
Fig. 1. (a) Conventional optical receiver. (b) Proposed low-bandwidth front end followed by DFE.

and cost. This work seeks compact receiver circuits that may be integrated onto the large nanoscale CMOS chips inside high-performance computing and networking equipment.

The front end of an optical receiver circuit comprises a transimpedance, $Z(s)$, that translates the small receiver input current, I_{PD}, into a voltage signal suitable for retiming. Fig. 1(a) shows a conventional arrangement where good signal integrity and low jitter at the retimer input, V_A, are maintained without equalization by designing the receiver's input response, $Z(s)$, to have a bandwidth of at least $(2/3) f_{bit}$, where f_{bit} is the input data rate [3]. This may be a sensible approach when separate dies are used for the front-end amplifier (possibly including a limiting amplifier) and subsequent retimer, but is less so in an integrated receiver as it requires a combination of low noise, high gain, and high bandwidth that is challenging in CMOS. Examples of this approach include [4]–[6], which were designed for 10, 25, and 25 Gb/s data rates and consumed 8 m, 49 mW, and 91 mW, respectively.

It has been long understood that the lowest possible noise and the best possible sensitivity are achieved by optical receivers with input bandwidths far less than f_{bit} [7]. A practical challenge recognized even in that early work was the need for an equalizer following the bandwidth-limiting front end. For example, a passive linear equalizer was employed in [8]. More recently, the receiver in [9] uses a 300 MHz input bandwidth for an 18.6 Gb/s input data rate. A feedforward equalizer (FFE) with a dynamic offset modulation, which is effectively a two-tap FFE, is used to recover the data. In both cases, linear equalizers amplify the noise relative to the signal power, which in turn reduces the receiver's sensitivity. In addition, the equalizer in [9] requires baud-rate samples of the received waveform, which can be a challenge at such high data rates. It achieves a power efficiency of 0.4 pJ/b (excluding clocking circuitry) with a sensitivity

of −4.7 dBm optically modulated amplitude (OMA). Moreover, the circuit's output is in a return-to-zero format, which would require additional high-speed RS or similar latches to restore the recovered digital outputs. To reduce the input-referred noise of the sampler and increase the sampled signal amplitude, a transimpedance amplifier (TIA) is introduced in [10] and [11] with a bandwidth less than 1 GHz between the photodiode and the sampler for 25 Gb/s data. Excellent energy efficiency is achieved, down to 0.17 pJ/b in [10], illustrating the benefits of using advanced 28 nm CMOS technology and a silicon photonic photodiode at 1310 nm with only 8 fF capacitance (C_{PD}) and 0.8 A/W responsivity. The receiver sensitivity is −4 dBm OMA.

Continuous-time equalization and FFE restore the low-pass filtered data by boosting the high-frequency spectral content relative to the low-frequency content. This causes the amplification of the high-frequency noise relative to the signal amplitude and, hence, the degradation of the signal-to-noise ratio (SNR). However, the decision-feedback equalizer (DFE) uses the "noiseless" reconstructed signal at the output of the slicer to remove the postcursor intersymbol interference (ISI). Thus, the DFE removes ISI without boosting high-frequency noise. The work presented in [12] uses a 1.1 GHz bandwidth front-end TIA for 4 Gb/s input data and uses a two-tap DFE to recover the data. The receiver achieves a remarkable −25 dBm OMA optical sensitivity with 1.1 pJ/b power efficiency. The work in [13] replaces the TIA with a resistor. The resistor and the parasitic capacitance at the input of the receiver (including the photodiode capacitance) limit the signal bandwidth to approximately $0.1 f_{bit}$. An IIR-DFE recovers the signal from the low-bandwidth node. The receiver achieves a −5 dBm OMA sensitivity with 0.93 pJ/b power efficiency at 9 Gb/s. The photodiodes in both of these works have the parasitic capacitance of 140 fF and the responsivity of 0.55 A/W. The speed of the receiver is limited by the maximum operating speed of the DFE.

This paper analyzes the receiver's eye opening in the presence of ISI and noise as a function of receiver front-end bandwidth with and without a DFE present. It is shown that the best sensitivity is achievable with front-end bandwidths as low as $0.12 f_{bit}$ as long as the resulting ISI can be eliminated by a DFE [Fig. 1(b)]. In practice, reducing the required front-end bandwidth permits it to be designed in CMOS with higher gain and lower power. Increasing the TIA gain also reduces the input-referred noise of the following stages. In Section II, it is shown that this approach results in power savings and improvements in the noise performance of the receiver.

The primary challenge in such an architecture is to design the high-speed DFE to equalize for the front-end bandwidth limitation while consuming low power. A first-order low-pass response with a bandwidth of $0.12 f_{bit}$ requires at least three discrete DFE taps to remove most of the ISI. However, one IIR tap can remove multiple postcursor ISI terms, potentially offering lower power consumption. The IIR-DFE approximates a long tail of the pulse response using an analog passive RC circuit in the feedback, and subtracts the approximated tail from the input signal. If the input front end exhibits a first-order response, an IIR-DFE is particularly attractive,
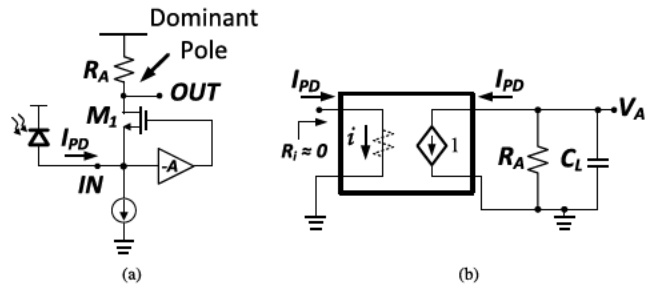

Fig. 2. (a) RGC input stage. (b) Simplified small-signal model of the RGC.

since a simple first-order RC feedback network can provide accurate ISI cancelation.

IIR-DFEs at or above 10 Gb/s have been reported in [15]–[17]. In each case, the feedback loop is comprised of a flip-flop, multiplexer, the IIR filter, and a summing node. The delay in the feedback loop has to be less than 1-UI for the first postcursor to be removed. This criterion poses an upper limit on the maximum speed of the IIR-DFE. The delay of the multiplexer and IIR filter is particular challenges. Hence, several works have combined one or more IIR-DFE taps with an FIR tap [15], [17]. In this paper, we prefer to avoid the additional power consumption of the added FIR tap, and focus instead on maximizing the operating speed of the IIR-DFE tap. To do so, a novel local feedback RC circuit is used to provide the IIR-DFE. As explained in Section III-B, the local feedback reduces the delay of the feedback loop to the settling time of the decision circuit. The approach is incorporated into a half-rate receiver architecture.

This paper is organized as follows. Section II presents modeling and analysis of the receiver performance as a function of the front-end bandwidth. Section III describes the circuit implementation of different blocks in the receiver, including the front end, IIR-DFE, and a technique to reduce the effect of the supply noise. Section IV presents the measurement results for a 20 Gb/s receiver based on this approach. Finally, Section V provides the conclusions.

II. BANDWIDTH-LIMITED FRONT-END DESIGN

This section presents an analysis of the signal, ISI, and noise at the output of a receiver front end as a function of its bandwidth. Note that a relatively short optical fiber link is presumed where modal and other dispersion are negligible, so that the input to the receiver is essentially free from ISI; all of the ISI considered here is introduced by the receiver front end. However, it is conceivable that a similar analysis and conclusions may be applied to the design of receivers for bandwidth-limited channels.

Optical receiver front ends convert photodetector current, I_{PD}, into a detectable voltage by passing I_{PD} through an impedance. This may be done simply with a passive resistor connected directly to the photodetector, as in [13]. In that case, front-end bandwidth is determined by the input time constant: the product of the passive resistor and the parallel capacitances of the photodetector, pads, ESD, and proceeding amplifier input. Alternatively, a regulated cascode (RGC), pictured in Fig. 2(a), is commonly used [6], [18]. An RGC
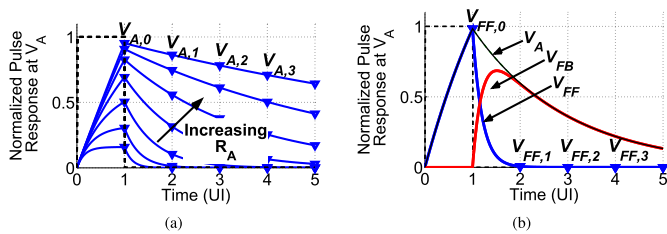
Fig. 3. Pulse responses of a first-order front end for different R_A values, normalized to $I_0/(C_L f_{bit})$ (3). (a) Without DFE. (b) With DFE.

employs a feedback amplifier $(-A)$ to boost the apparent transconductance of an input common-gate transistor, M_1. As a result, the RGC small-signal input resistance is so small that the circuit's dominant time constant is at the output node, which is isolated from the relatively large photodetector, pad, and ESD capacitances. This permits the use of a larger resistor, R_A, and hence larger gain compared with a simple passive resistor load designed for the same bandwidth. In either case, the front-end response $Z(s)$ is the first-order low pass. Section II-A analyzes the first-order receiver front end, with a focus on RGC inputs, neglecting noise to highlight the benefits of limiting front-end bandwidth far below the input bit rate. In Section II-B, noise is included in the analysis. Another method to realize transimpedance gain with higher bandwidth is to connect the resistor in feedback around an amplifier. Such TIAs often have the second-order responses, which are analyzed in Section II-C.

A. First-Order Noiseless Model

A simplified small-signal model for an optical front end utilizing an RGC input is shown in Fig. 2(b). The transimpedance from I_{PD} to V_A has a first-order response, $Z_1(s)$ with a bandwidth $f_A = (1/2\pi R_A C_L)$

$$Z_1(s) = \frac{R_A}{1 + s/(2\pi f_A)}. \tag{1}$$

Fig. 3(a) shows the front-end's response to a input current pulse for different values of R_A assuming C_L is fixed

$$V_A(t) = \begin{cases} R_A I_0 (1 - e^{-2\pi f_A t}) & 0 \leq t \leq 1/f_{bit} \\ R_A I_0 (1 - e^{-2\pi f_A/f_{bit}}) e^{-2\pi f_A t} & 1/f_{bit} \leq t. \end{cases} \tag{2}$$

The plots are normalized with respect to the pulse amplitude I_0 and duration $(1/f_{bit})$, and baud-spaced samples are indicated forming a discrete-time sequence

$$V_{A,i} = R_A I_0 (1 - e^{-2\pi f_A/f_{bit}}) e^{-2\pi i f_A/f_{bit}}, \quad i \geq 0. \tag{3}$$

With $V_{A,0}$ be the main cursor and $V_{A,i}(i > 0)$ the postcursor ISI. It can be seen that a small value of R_A maximizes bandwidth and hence minimizes postcursor ISI ($V_{A,i>0}$), however, results in a small main sample ($V_{A,0}$). On the other hand, a large value of R_A makes the main sample larger, but also increases postcursor ISI. If not canceled, a worst case data pattern causes all postcursor ISI to add constructively reducing

vertical eye opening at V_A to

$$S_A = V_{A,0} - \sum_{i \neq 0} |V_{A,i}|. \tag{4}$$

Substituting (3) into (4), we arrive at the eye opening for the first-order front end without DFE

$$S_{A,\text{1st-order}} = R_A I_0 (1 - 2e^{-2\pi f_A/f_{bit}}). \tag{5}$$

Note that with $(f_A/f_{bit}) < 0.11$, (5) results in $S_{A,\text{1st-order}} < 0$, indicating that ISI overwhelms the main sample and the eye diagram at V_A is closed. Moreover, differentiating (5) with respect to $(f_A/f_{bit}) = (1/2\pi R_A C_L f_{bit})$ and equating the result with 0 results in $(f_A/f_{bit}) = 0.267$, which is the front-end bandwidth providing the best possible eye opening without equalization (neglecting noise). However, improved eye opening can be achieved by removing postcursor ISI with an equalizer. Whereas feedforward linear equalizers amplify noise to do so [8], [9], a DFE can, in principle, cancel postcursor ISI noiselessly and is, therefore, presumed here. Fig. 3(b) shows the normalized pulse responses of the front end prior to the DFE, V_A in Fig. 1(b), the feedback signal for the cancelation of the postcursor ISI, V_{FB}, and the resulting sampled pulse response $V_{FF,i}$. An ideal infinite-length DFE is assumed wherein all postcursor ISI ($V_{A,i}$, $i > 0$) is precisely canceled. Hence, the worst case eye opening at V_{FF} after DFE cancelation includes only the precursor terms

$$S_{FF} = V_{A,0} - \sum_{i < 0} |V_{A,i}|. \tag{6}$$

Since a first-order front end introduces no precursor ISI it results in an eye opening at V_{FF} of

$$S_{FF,\text{1st-order}} = V_{A,0} = R_A I_0 (1 - e^{-2\pi f_A/f_{bit}}). \tag{7}$$

The maximum value of (7), $S_{FF,\text{1st-order}} = (I_0/C_L f_{bit})$, is achieved as $R_A \to \infty$ and $f_A \to 0$, corresponding to an integrating front end. However, diminishing improvement in $S_{FF,\text{1st-order}}$ is observed for the values of f_A below about $0.2 f_{bit}$ owing to the decaying exponential in (7) [13]. This will lead us to prefer a value $f_A > 0$ once noise is properly considered in Section II-B.

Fig. 4 shows the eye opening with DFE $S_{FF} = V_{A,0}$ from (7), the sum of all ISI terms $|V_{A,i}|$ ($i \neq 0$), and the eye opening without the DFE, $S_{A,\text{1st-order}}$ from (5), as amplifier gain R_A, and, hence, bandwidth f_A are swept. It shows that when postcursor ISI is fully removed by a DFE and noise is neglected, very low front-end bandwidths can offer up to $2.7\times$ more eye opening, or 4.3 dB better optical sensitivity, than is possible without equalization.

B. First-Order Noisy Model

Next noise is introduced into the analysis, revealing the optimal bandwidth for a first-order receiver front end followed by a DFE. The input-referred power spectral density $I_n^2(f)$ is filtered by the first-order $Z_1(s)$ resulting in rms noise at the output

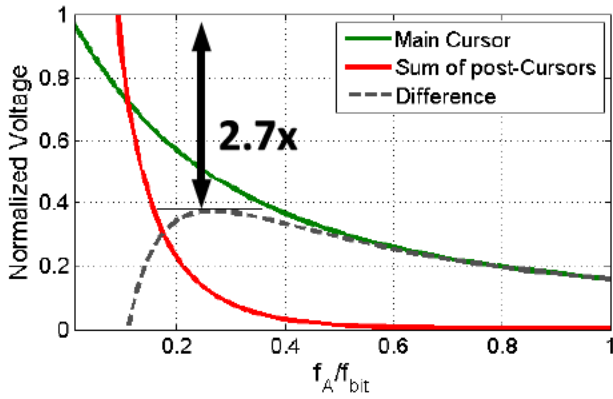$$\overline{V_{n,A}} = \sqrt{\int_0^\infty |Z(f)|^2 I_n^2(f) df}. \tag{8}$$

Fig. 4. Main sample, ISI, and worst case eye opening as a function of f_A for a first-order noiseless front end having fixed load capacitance, C_L.

Fig. 5. Bandwidth dependence of signal integrity at the output of a first-order receiver front end (Γ). (a) Assuming white input-referred noise. (b) Assuming white noise below $0.3 f_{bit}$, beyond which the input-referred noise spectrum increases in proportion to f^2.

First, consider the case where white noise is dominant, so that the input-referred noise spectral density is constant. Substituting $I_n^2(f) = I_{n0}^2$ and (1) into (8), the rms output noise is

$$\overline{V_{n,A}} = R_A I_{n0} \sqrt{\frac{\pi}{2} f_A}. \tag{9}$$

Signal integrity in the presence of noise and ISI is quantified by the ratio of vertical eye opening at the input of the retimer (S_A without DFE or S_{FF} with DFE) to the rms noise at the retimer input. Without DFE, using (5) and (9), the ratio is

$$\frac{S_A}{\overline{V_{n,A}}} = \frac{I_0}{I_{n0}} \sqrt{\frac{2}{\pi f_{bit}}} \frac{(1 - 2e^{-2\pi f_A/f_{bit}})}{\sqrt{f_A/f_{bit}}}$$

$$\equiv \frac{I_0}{I_{n0}} \sqrt{\frac{2}{\pi f_{bit}}} \Gamma_A(f_A/f_{bit}). \tag{10}$$

With an ideal DFE, there is no residual ISI, so the result is simply the sampled SNR. Using (7) and (9)

$$\frac{S_{FF}}{\overline{V_{n,A}}} = \frac{I_0}{I_{n0}} \sqrt{\frac{2}{\pi f_{bit}}} \frac{(1 - e^{-2\pi f_A/f_{bit}})}{\sqrt{f_A/f_{bit}}}$$

$$\equiv \frac{I_0}{I_{n0}} \sqrt{\frac{2}{\pi f_{bit}}} \Gamma_{FF}(f_A/f_{bit}). \tag{11}$$

In each case, the ratio is a product of $(I_0\sqrt{2}/I_{n0}\sqrt{\pi f_{bit}})$ and the following functions of normalized amplifier bandwidth:

$$\Gamma_A(f_A/f_{bit}) = \frac{(1 - 2e^{-2\pi f_A/f_{bit}})}{\sqrt{f_A/f_{bit}}} \tag{12}$$

$$\Gamma_{FF}(f_A/f_{bit}) = \frac{(1 - e^{-2\pi f_A/f_{bit}})}{\sqrt{f_A/f_{bit}}}. \tag{13}$$

Assuming I_{n0} does not depend on f_A, the functions Γ are here used to determine the optimal front-end bandwidth. Fig. 5(a) shows that Γ_A peaks at $0.4 f_{bit}$ whereas Γ_{FF} peaks at $0.2 f_{bit}$ providing 1.8 dB improvement in optical sensitivity. Moreover, at their optimal points, the front end with DFE has a gain more than $2\times$ greater than the front end without DFE, also reducing the impact of subsequent receiver circuits.
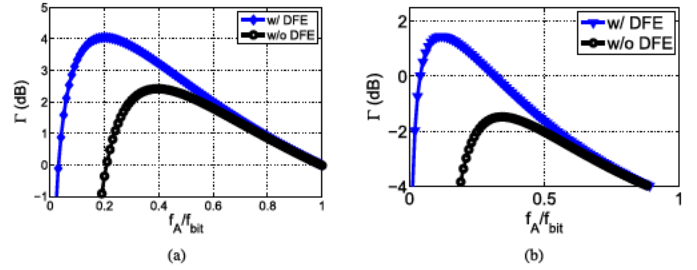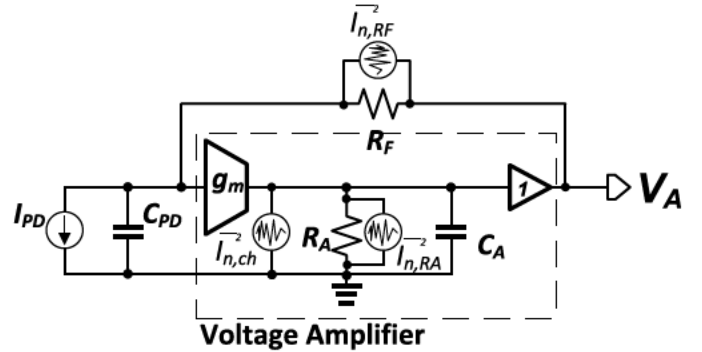

Fig. 6. Simplified model of a second-order feedback TIA.

Including the noise of the feedback amplifier and R_A, the input-referred noise spectrum is no longer white, but rather increases beyond some corner frequency f_c at +20 dB/decade

$$I_n^2(f) = I_{n0}^2 \left(1 + \frac{f^2}{f_c^2}\right). \tag{14}$$

Substituting (1) and (14) directly into (8) would result in infinite output noise. It is, therefore, assumed that the output noise is, in fact, bandlimited by a second pole at f_{bit}. At this frequency, the second pole has negligible impact on the eye opening, still given by (5) and (7).

The f^2 noise term in (14) increases the impact of high-frequency noise, so the optimal front-end bandwidth f_A shifts to even lower frequencies than in the case of white input-referred noise. Also, the performance gap between receivers with and without DFE widens. For example, assuming $f_c = 0.3 f_{bit}$ consistent with our design, Fig. 5(b) shows that the optimum bandwidth with DFE is $f_A = 0.12 f_{bit}$, offering 2.9 dB improvement in sensitivity compared with no DFE.

C. Second-Order Noisy Model

This section extends the noise analysis to front ends having a feedback TIA connected directly to the photodiode [5], [19]–[21]. Consider the second-order feedback TIA in Fig. 6 in which the voltage amplifier has an open-loop dc voltage gain of $A = g_m R_A$. The voltage amplifier has a single pole at $f_A = 1/(2\pi R_A C_A)$ resulting in gain-bandwidth product of $f_0 = A f_A$, which is roughly proportional to the technology f_T. Defining $f_{in} = 1/(2\pi R_f C_{PD})$, the transfer
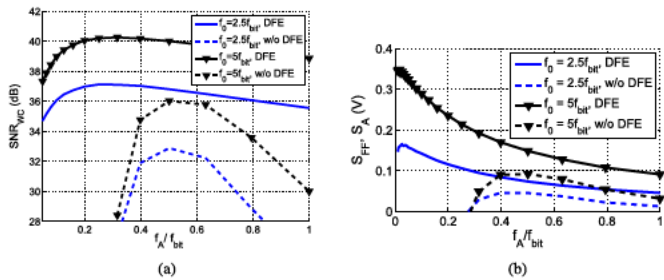
Fig. 7.   Comparison of the systems with DFE ($f_{in} = f_{bit}/60$) and without DFE with $C_{PD} = 200$ fF, $f_{bit} = 20$ GHz, and $f_0 = 2.5f_{bit}$, $5f_{bit}$, and f_A is swept by changing R_A at a fixed C_A. (a) SNR_{WC}. (b) Worst case vertical eye opening.
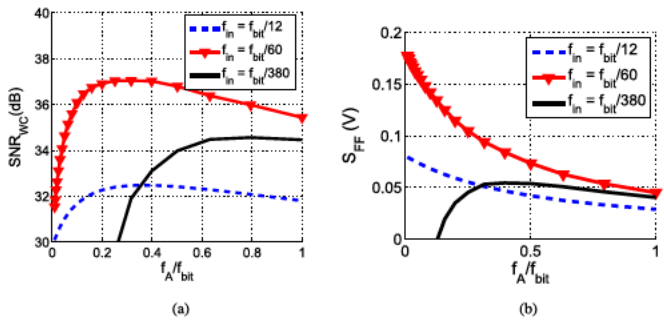

Fig. 8.   (a) SNR_{WC}. (b) Worst case vertical eye opening for different f_{in} values for the receiver with DFE. $C_{PD} = 200$ fF, $f_{bit} = 20$ GHz, and f_A is swept by changing R_A at a fixed C_A.

function of the feedback TIA is given by

$$Z_2(s) = \frac{R_f A}{1+A} \cdot \frac{1}{s^2/\omega_n^2 + \omega_n/Qs + 1} \quad (15)$$

$$\omega_n = 2\pi\sqrt{f_A f_{in}(\frac{f_0}{f_A}+1)}, \quad Q = \frac{\sqrt{f_A f_{in}(\frac{f_0}{f_A}+1)}}{f_A + f_{in}}. \quad (16)$$

The noise sources in the amplifier are shown in Fig. 6 with $I_{n,Rf} = 4kT/R_f$, $I_{n,RA} = 4kT/R_A$, and the thermal noise of the channel given by $I_{n,ch} = 4kT\gamma/g_m$ (with γ assumed to be 2 in the following simulations). Similar to Section II-B, we use $S_A/V_{n,rms}$ in the absence of a DFE and $S_{FF}/V_{n,rms}$ with DFE, both hereby referred to as the worst case SNRs (SNR_{WC}), to quantify the noise performance of the front end. Terms S_A and S_{FF} are as defined in (4) and (6) and $V_{n,rms}$ is the total integrated rms noise at the output of the closed-loop TIA (V_A). In particular, a second-order TIA is considered with $f_{bit} = 20$ GHz, $C_{PD} = 200$ fF, and $f_0 = 2.5f_{bit}$. The simulations in Figs. 7 and 8 assume the $I_0 = 100$ μA input current signal and f_A is swept by sweeping R_A with fixed C_A. Changing the values of I_0 and f_{bit} only scales the axes of Figs. 7 and 8 without changing the optical choice of f_{in} and f_A.

Without DFE: Fig. 7 shows that $f_A = f_{bit}/2$ and $\omega_n = 2\pi \cdot 0.4 f_{bit}$ (resulting in $f_{in} = 0.05 f_{bit}$ and $Q = 0.7$) give the optimum SNR_{WC} of 33 dB and S_A of 50 mV is reached when $f_0 = 2.5f_{bit}$. This is in contrast with the case of constant white input-referred noise where the optimum occurs at $\omega_n = 2\pi \cdot 0.7 f_{bit}$ [3] where $Q = 0.7$ is assumed to minimize ISI for a given bandwidth. The reason for the discrepancy

is that the output noise spectrum of the TIA in Fig. 6 has the form of (14), which pushes the optimum bandwidth lower.

With DFE: To study the optimal design choices of f_A and f_{in} when TIA is followed by an ideal (infinite-length) multitap DFE, Fig. 8 plots S_{FF} and SNR_{WC} as a function of f_A for different f_{in} values. For $f_A \gg f_{bit}$, the optimum SNR_{WC} is achieved with $f_{in} \approx 0$ ($R_f \rightarrow \infty$). Intuitively, in this case, $V_{n,rms}$ is the dominated by $I_{n,ch}$ times the output impedance of the TIA and is not filtered by the low-frequency input pole. So to maximize SNR_{WC}, the signal amplitude must be maximized by letting $R_f \rightarrow \infty$. However, $f_A \gg f_{bit}$ at a given f_0 is not desirable, since it reduces the voltage swing at the output of the voltage amplifier by reducing $A = (f_0/f_A)$. In Fig. 8, this case is shown by setting f_{in} to the small value of $f_{bit}/380$. It can be seen that, reducing f_A below f_{bit} results in a drop in SNR_{WC} due to the increase in the precursor ISI.

The precursor can be reduced by increasing f_{in} to $f_{bit}/60$. As f_A is decreased (A increases) at a given f_{in}, the input resistance of the TIA decreases, which in turn increases the closed-loop pole frequency at the input of the TIA. This allows for a larger S_{FF} without much increase in noise, which in turn improves SNR_{WC}. Decreasing f_A continues to improve SNR_{WC} until $f_A \approx 2A f_{in}$ where $Q \approx 0.7$. Decreasing f_A further degrades SNR_{WC} by causing excess high-frequency peaking in the output noise spectrum. It turns out that for the best SNR_{WC}, R_f must be chosen so the closed-loop Q is about 0.7 with $f_A = 0.3f_{bit}$. Using (16), this optimum translates to $f_{in} \approx f_A/(2A) = f_{bit}^2/(22f_0)$, which can be used to find the optimum R_f for a given C_{PD}.

Increasing f_{in} beyond $f_{bit}/60$ (by choosing a smaller R_f) reduces SNR_{WC} by reducing the transimpedance gain of the TIA and increasing the high-frequency noise at the output. For example, this can be seen when $f_{in} = f_{bit}/12$ in Fig. 8. Note that for this plot, $Q = 0.7$ happens at $f_A = 0.7f_{bit}$ being higher than the case with $f_{in} = f_{bit}/60$.

Note that if photodiode parasitic capacitance is doubled, the optimum R_f for $f_{in} \approx f_{bit}^2/(22f_0)$ must be halved. This does not affect the output voltage noise of the TIA (assuming the noise is dominated by the transistor's thermal noise) but halves the signal swing degrading SNR_{WC} by 6 dB. Therefore, C_{PD} has to be kept as small as possible. Due to the relatively large $C_{PD} = 200$ fF used in this paper, a first-order front end provides a superior performance and was, therefore, adopted.

Going to a faster technology would provide a higher gain bandwidth, f_0. This allows g_m and R_f to be increased, roughly in proportion, while maintaining the same TIA ω_n and Q. As a result, both the noise contribution of $I_{n,ch}$ and the closed-loop gain increase, with a net benefit in terms of both SNR_{WC} and S_{FF}. The benefits are shown in Fig. 7 for the case $f_0 = 5f_{bit}$.

In summary, a front end without a DFE reaches $SNR_{WC} = 33$ dB with $S_A = 0.1$ V when $f_A = 0.5f_{bit}$ and $Q = 0.7$. Adding the DFE reduces the optimum f_A to $0.3f_{bit}$ and $f_{in} \approx f_{bit}^2/(22f_0)$ with $SNR_{WC} = 37$ dB and $S_{FF} = 0.2$ V. Comparing these numbers, a 4 dB improvement in SNR_{WC} and a factor of two in vertical eye opening can be gained by utilizing an ideal DFE.
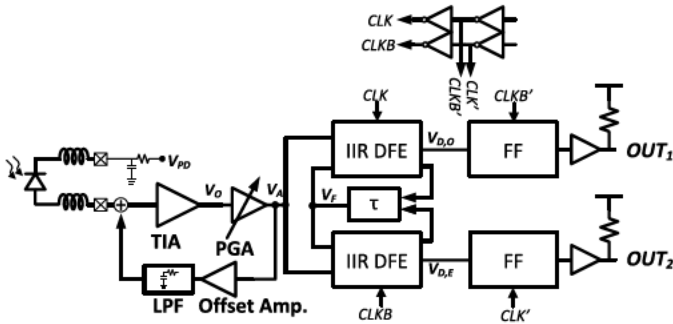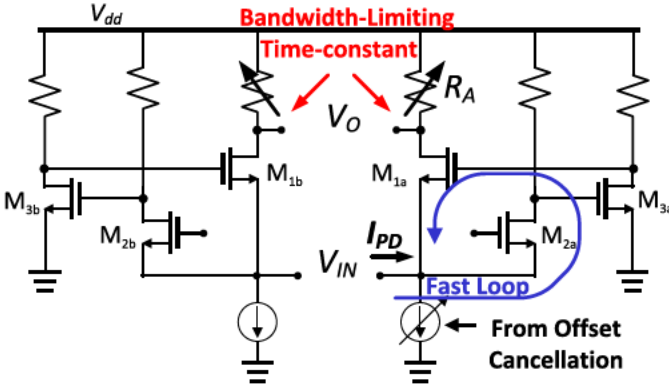
Fig. 9. System block diagram.

Fig. 10. TIA schematic.

III. CIRCUIT PROTOTYPE

A prototype receiver is developed in 65 nm CMOS demonstrating the combination of a limited-bandwidth first-order front end followed by an IIR-DFE at 20 Gb/s. The block diagram is shown in Fig. 9. The front end is comprised of a pseudodifferential RGC input stage, followed by a programmable-gain amplifier (PGA), an offset compensation loop, and half-rate IIR-DFE. The RGC provides relatively low noise and a first-order response whose bandwidth is not dependent upon the large and variable photodiode capacitance ($C_{PD} \approx 200$ fF in this paper). The resulting single time constant pulse response is well suited for cancelation by a simple first-order IIR-DFE. The PGA permits the control of the signal swing at the input of the DFE. Offset cancelation balances the pseudodifferential input and compensates for mismatch through the front end. As a tradeoff between low-pass filter area and the number of tolerable consecutive identical digits (CIDs), the lower cutoff frequency of the the front end is made 4 MHz sufficient to accommodate a PRBS7 sequence. Patterns with longer CID sequences would either need a lower cutoff frequency to avoid baseline wander or some coding to increase the transition density. For instance, PRBS31 requires a cutoff frequency smaller than 1 MHz, which makes the filter roughly four times bigger. Finally, the integrated IIR-DFE retimes and demultiplexes the input data, providing two half-rate outputs.

A. Receiver Front End

The input stage is shown in Fig. 10 [22]. When loaded with the photodiode ($C_{PD} = 200$ fF), the inner "fast" loop has a bandwidth of ≈ 20 GHz. The bandwidth of the overall
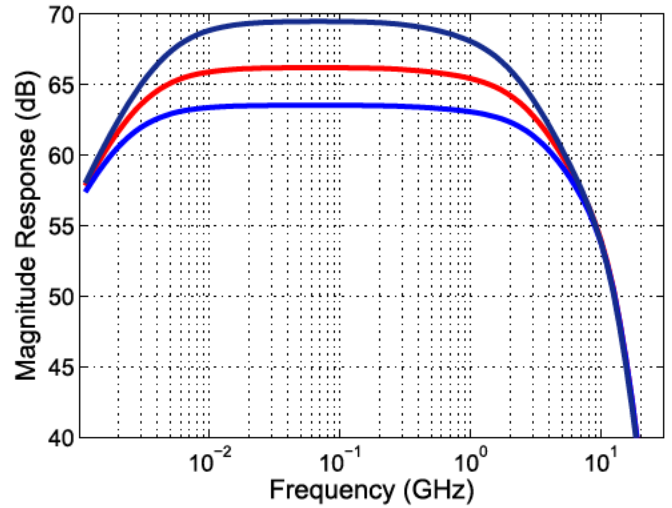
(a)
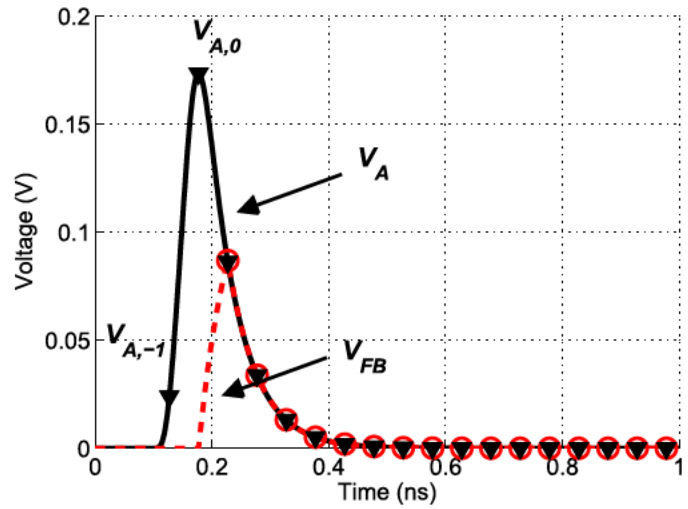
(b)

Fig. 11. (a) Electrical response of the receiver up to the output of the PGA for three TIA gain/bandwidth settings and PGA gain set to maximum. (b) Receiver front-end pulse response in the presence of high-frequency poles and the ideal IIR-DFE's corresponding feedback signal (V_{FB}).

response V_O/V_{IN} is, therefore, limited by the time constant at the output node, V_O. The bias current through M_1 is nominally 300 μA; therefore, R_A can be as large as 1.3 kΩ without running into voltage headroom problems under a 1 V supply. The value of R_A is coarsely programmable over the range 650–1300 Ω corresponding to simulated bandwidths of 2–4 GHz permitting the exploration of front-end gain-bandwidth tradeoffs. It is worth comparing the performance of the RGC front end to what could have been obtained using a simple passive resistor as a first-order transimpedance [13]. The noise and power consumption associated with M_{1-3} is of course avoided using a passive front end. However, for $C_{PD} = 200$ fF and a bandwidth of $0.15\times$, the targeted data rate of $f_{bit} = 20$ Gb/s and a maximum resistance of only 260 Ω could have been used, perhaps even lower after accounting for circuit bond-pad capacitances and ESD-protection. An additional 14 dB of broadband gain would, therefore, have been required in the front end, introducing significant
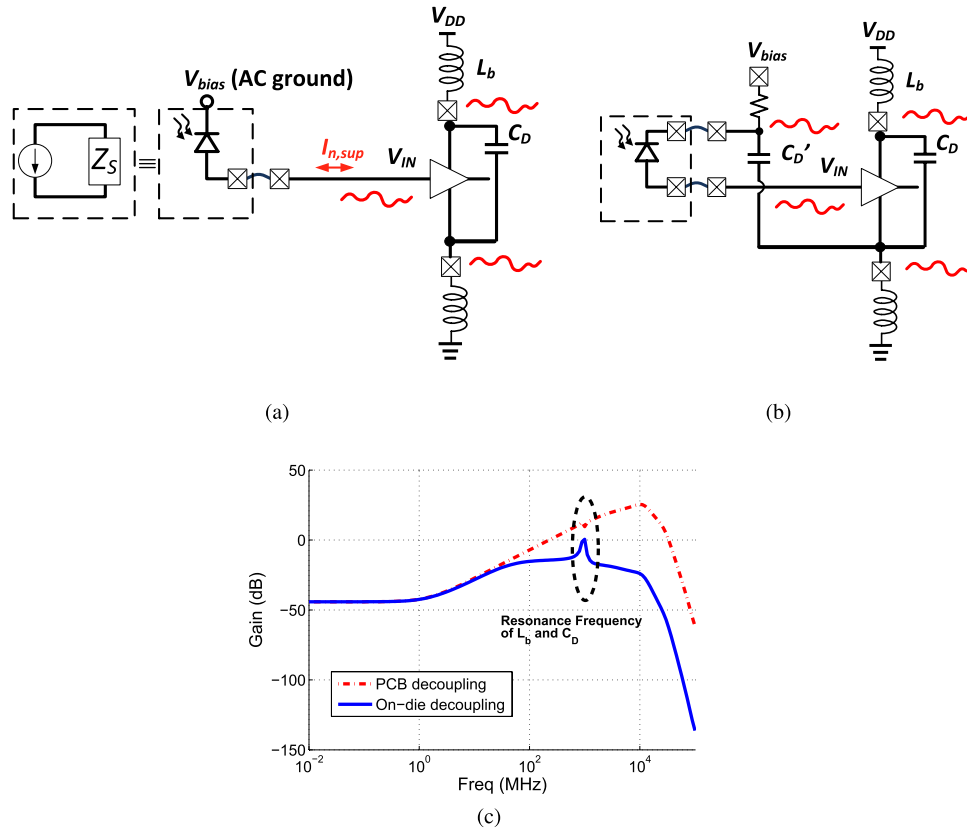
Fig. 12. (a) Noise introduced by supply/ground noise modulates the input current. (b) Decoupling the photodiode on-chip prevents the supply noise from inducing input current. (c) Simulated frequency response from the chip ground net to the differential output in both cases.

additional noise and power consumption. Furthermore, it is unlikely a full-rate retimer with DFE feedback all the way to the input could have been accommodated in 65 nm CMOS at 20 Gb/s, as was done at 9 Gb/s in [13].

The PGA that follows the RGC comprises a single-stage resistively loaded nMOS differential pair, with digitally programmable degeneration resistance. The simulated gain is programmable over the range approximately 1–8 dB while maintaining over 18 GHz bandwidth, so that the front-end response remains dominantly first order. Fig. 11(a) shows the frequency response of the receiver upto the output of the PGA. The receiver front end consumes 4.5 mW, including RGC, PGA, and offset cancellation circuitry.

It is worth noting that there are other high-frequency poles due to the TIA feedback loop, PGA output pole, and PD's intrinsic bandwidth. Fig. 11(b) shows the pulse response in the presence of these high-frequency poles. The postcursor ISI is still determined primarily by the low-frequency pole at the output of the TIA and, therefore, well approximated by the first-order IIR feedback. However, the high-frequency poles cause a precursor ISI ($V_{A,-1}$), which cannot be eliminated by a DFE. Thus, the additional poles need to be kept at a frequency above $f_{\rm bit}$ to minimize the precursor ISI penalty.

The signal amplitude at the input of the receiver can be as low as 5 mV, and is single ended, making it particularly sensitive to supply noise. If the photodiode is connected, as shown in Fig. 12(a), its bias voltage is not decoupled to the same ac ground as the input stage. Hence, ac voltages appear across the finite source impedance, Z_S, inducing an input-referred noise current $I_{n,\rm sup}$. Biasing capacitance for twin p-i-n photodiodes was suggested to reduce the impact of packaging inductance in [14]. The presented receiver uses a similar technique to increase ground noise rejection. Fig. 12(b) shows the biasing method used in this prototype, whereby the photodiode bias voltage is also decoupled on-die to the same ac ground as the input stage. With this scheme, both photodiode terminals are modulated by the same supply noise as the input stage, and no noise current arises. Fig. 12(c) compares the frequency response from the chip ground net to the output of the TIA in both cases. This is crucial, since, for instance, at 20 Gb/s, the clock buffers and the half-rate comparators cause a 10 GHz tone on the ground. Without the on-die decoupling, even 1 mV amplitude of this 10 GHz noise alone causes 10 mV noise at the output of the TIA. On-die decoupling reduces this noise at the output of the TIA to 0.1 mV.

To roughly compare the given front end and one with wider bandwidth and no DFE (for a given latch sensitivity), the TIA bandwidth is increased to $f_A \approx 0.35 f_{\rm bit}$ by reducing R_A. This requires the PGA gain to increase from 2.5 to 5 (Fig. 4) while maintaining $\approx$20 GHz of bandwidth, which can be done by adding differential-pair gain stages. Based on simulations in 65 nm CMOS technology, this results in an additional 5 mW in power and 2.2 dB sensitivity reduction.

B. IIR-DFE

All previously reported IIR-DFEs require the full-rate data pattern to be reproduced at receiver internal nodes, and
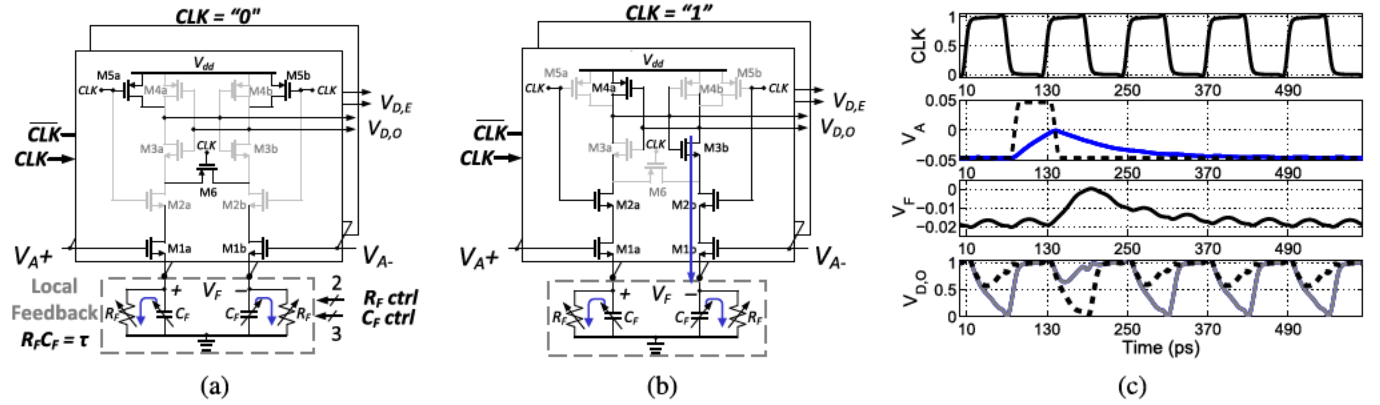
Fig. 13. DFE schematic (a) when clock is "0" and (b) when clock is "1." (c) Simulated voltage waveforms in the DFE. R_f can be programmed in the range of 20–160 Ω and C_F can be programmed in the range of 1–2 pF.

then passed through an analog filter. This requires either a full-rate retimer [13], or in most cases, a full-rate multiplexer [15]–[17] consuming additional power and adding delay to the DFE feedback path, and has limited their operating speed to 10 Gb/s. Here, no full-rate data signal is reproduced. Instead, the passive IIR filter is multiplexed between half-rate signal paths. The IIR-DFE schematic is shown in Fig. 13(a). A single differential IIR filter, R_F and C_F, degenerates two half-rate latches. Transistors M_1 are the input transistors, serving as the DFE summer. They act upon their gate–source voltage: the difference between the front-end output V_A, and IIR feedback voltage V_F. Transistors (M_2) are clocked to alternately connect each of the half-rate latches (M_{3-4}) to the input transistors, effectively multiplexing the IIR filter between latches. When the clock is low, M_2 disconnects the latch from the input and feedback, and precharges the output nodes to V_{DD}.

When the clock goes high, M_2 injects a differential current ($I_{D,\text{diff}}$) proportional to ($V_A - A_F V_F$) tripping the latch. As derived in the Appendix, A_F is given by

$$A_F \approx 1 + \frac{2(W/L)_1(V_{GS1} - V_t)^2}{(W/L)_2(V_{DD} - V_t)^2}. \tag{17}$$

The polarity of $I_{D,\text{diff}}$ determines, which DFE output will be pulled low.

As the latch resolves, one of the output nodes is pulled low, and charge stored on its output capacitance in the precharge phase passes through M_1 and M_2 to the corresponding local feedback capacitor C_F. When the clock phase is complete, the resulting differential voltage pulse on V_F is immediately available to cancel postcursor ISI by degenerating the other half-rate latch. Meanwhile, C_F is continuously discharging via R_F, producing the exponentially decaying waveform on V_F that is characteristic of a first-order IIR-DFE. Local feedback obviates the need to multiplex the high-speed CMOS outputs $V_{OUT,E}$ and $V_{OUT,O}$ back to a full data rate pattern within a feedback loop, as in previous IIR-DFEs [15]–[17]. Simulated waveforms illustrate the DFE operation for an isolated input pulse in Fig. 13(c). The lone "1" bit results in a pulse on V_F and subsequent decay. Note that the output decision changes polarity as it should because of the DFE local feedback, even though the differential input to the latch V_A never crosses zero.
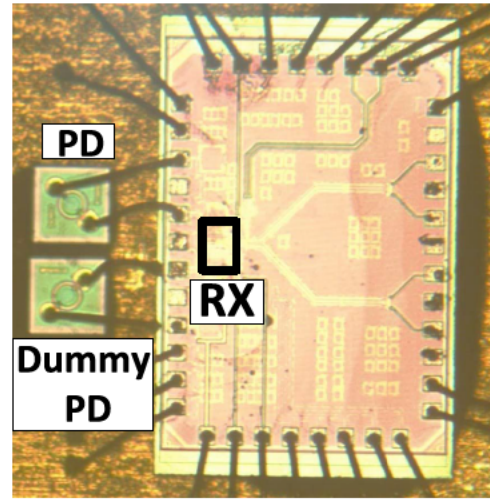

Fig. 14. CMOS chip wire bonded to the main photodiode. The dummy photodiode wire bonds disconnected.

The voltages on V_F are less than 50 mV, having little impact on the regeneration speed of the latch.

The IIR-DFE tap gain depends upon the parameters in (17), and upon the amplitude of the voltage pulses arising at V_F in response to each decision. The latter is inversely proportional to C_F, which is made programmable to provide control over the tap gain. The DFE time constant is $\tau = R_F C_F$, and is controlled by a digitally programmable R_F. Variations in the supply voltage also affect the gain of the A_F in (17) by changing V_{DD} and V_{GS1}. However, variations in V_{DD} and V_{GS1} counteract each other reducing the effect on A_F. Latch offset compensation is also incorporated into the latch via a differential pair in parallel with M_3 (not shown).

IV. EXPERIMENTAL RESULTS

A prototype was fabricated in 65 nm CMOS technology. The prototype die and photodiode were copackaged in a QFN package and directly wire bonded together. A second dummy photodiode was included in the package and bonded to the other side of the receiver's pseudodifferential input for some tests to examine the impact of a balanced source impedance
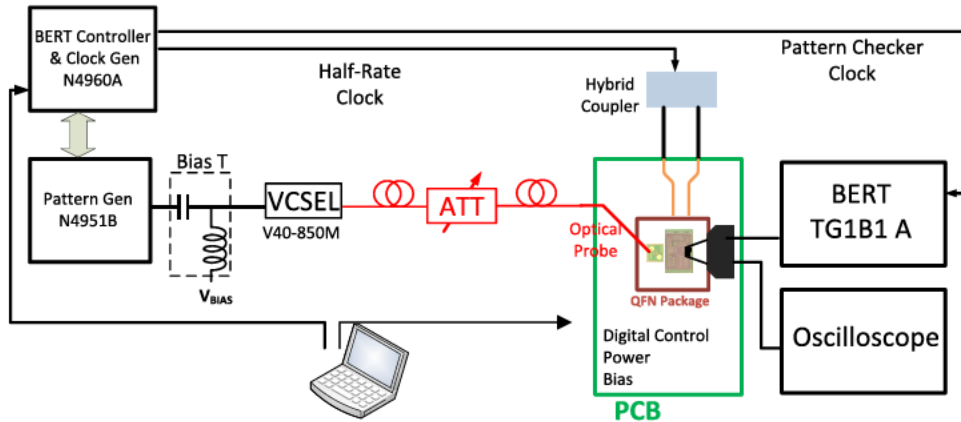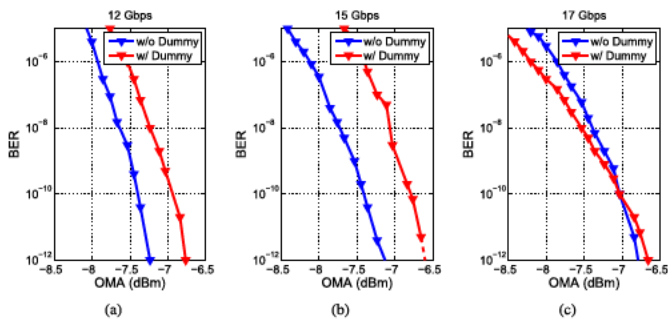
Fig. 15.   Test setup.

Fig. 16.   Waterfall curves with and without the dummy photodiode. (a) 12 Gb/s. (b) 15 Gb/s. (c) 17 Gb/s.

Fig. 17.   Bathtub curves for different DFE time-constant settings at 17 Gb/s.

on the receiver's performance. The die and package are shown in Fig. 14.

The test setup is shown in Fig. 15. The photodiodes used are Cosemi BPD2010 having a typical bandwidth of 12 GHz, $C_{PD} = 200$ fF, and responsivity of 0.5 A/W. The input optical data pattern was generated by directly modulating an 850 nm VCSEL with a PRBS length-(2^7-1) pattern generator, and the input optical power controlled by a variable optical attenuator. The receiver bit-error rate (BER) is measured on each of its half-rate outputs.

The receiver BER is plotted versus input optical power with and without the dummy photodiode in Fig. 16. An asymmetrical source impedance without dummy photodiode permits supply noise to appear differentially. Hence, the use of a dummy photodiode to make the pseudodifferential input stage as symmetric as possible is generally seen as beneficial. However, as explained in Section III-A, decoupling the bias of the photodiode to the chip ground provides very good supply rejection even without the dummy photodiode. Moreover, the additional dummy photodiode's capacitance sinks more of the circuit's thermal noise current and, therefore, increases the input stage's thermal noise. As a result, Fig. 16 shows that removing the dummy photodiode improves sensitivity by 0.5, 0.5, and 0.15 dB at 12, 15, and 17 Gb/s, respectively. Sensitivity improvement at 17 Gb/s is slightly less than the former two cases, which could be caused by some high-frequency supply glitches whose timing impact the sensitivity. The difference in sensitivity improvement is
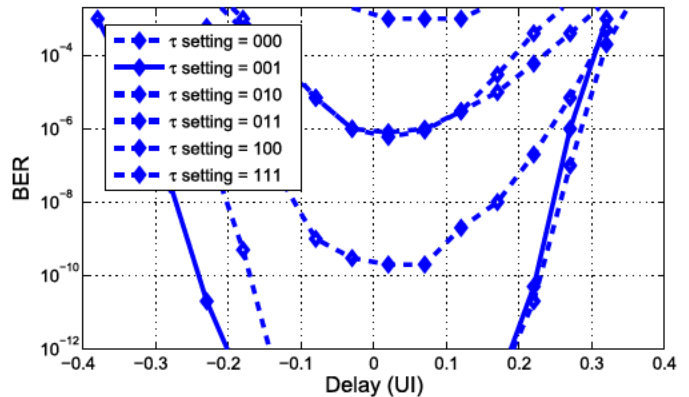
so small (0.4 dB) that it is hard to quantify. Sensitivities between -7.3 and -6.8 dBm OMA at a BER of 10^{-12} are measured. Sensitivity numbers are measured at the PD; however, the loss due to the optical probe was measured to be negligible. To verify the functionality of the DFE, the receiver is tested with a 17 Gb/s optical input at -3.5 dBm OMA (equivalent to 230 μA_{pp} input current). The front-end bandwidth is set to approximately 2 GHz $= 0.12 f_{bit}$. Note that at this bandwidth, the eye diagram at the output of the front end (V_A) is completely closed. Bathtub curves for different DFE time constants are shown in Fig. 17. At the optimal time constant, the eye opening is 0.4 UI at a BER of 10^{-12}. Also when τ is significantly different than the optimal value, even though input power is much greater than the receiver noise, the receiver becomes ISI-limited and does not reach BER $= 10^{-12}$. Increasing input power does not improve the SNR in the ISI-limited receiver.

Operation up to a maximum data rate of 20 Gb/s was measured. The input optical eye diagram before the variable optical

TABLE I
BREAKDOWN OF POWER CONSUMPTION WHEN OPERATED AT 20 Gb/s

TIA + PGA + Offset Cancellation	4.5 mW
DFE + Flip-Flops	3.0 mW
Clock Buffers	6.6 mW
Total	**14.1 mW**

TABLE II

COMPARISON OF CMOS OPTICAL RECEIVERS WITH INTEGRATED RETIMER CIRCUIT FOR DISCRETE PHOTODIODES, EITHER AT SIMILAR DATA RATES OR UTILIZING DFE

	[23]	[9]	[12]	[13]	This Work
CMOS Technology	32nm SOI	65nm	90nm	90nm	65nm
Equalization	Tx FFE	Rx FFE	2-tap DFE	IIR-DFE	IIR-DFE
Max. Data Rate (Gb/s)	28	18.6	4	9	20
Photodiode Capacitance (fF)	85	150	140	140	200
Photodiode Responsivity (A/W)	0.55	1.0	0.55	0.55	0.5
Sensitivity (dB OMA)	-3	-4.7*	-22	-5	-5.8
Receiver Power Efficiency (pJ/bit)	1.95	0.4**	1.15	0.93	0.71
Area (mm^2)	0.012	0.0028	0.0045	0.004	0.027

* Coupling loss de-embedded
** Does not include clock generation or SR-latch

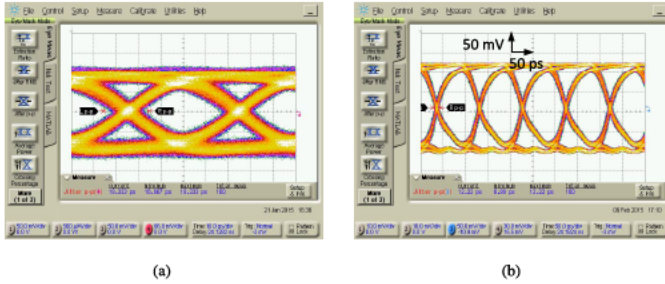
(a)   (b)

Fig. 18.   (a) Optical eye diagram obtained from the 850 nm VCSEL before the optical attenuator at 20 Gb/s. Having peak to peak jitter of 19.3 ps and extinction ratio of 3.6 dB. (b) Half-rate output at 10 Gb/s.
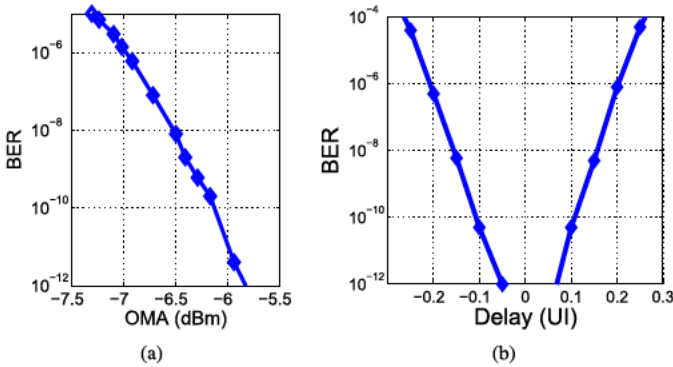

(a)   (b)

Fig. 19.   Measurements at 20 Gb/s with a TIA bandwidth of 3 GHz. (a) Bathtub curve at −5.1 dBm OMA shows 0.12-UI timing margin at BER = 10^{-12}. (b) Measured BER versus OMA illustrates a sensitivity of −5.8 dBm at BER = 10^{-12}.

attenuator and the receiver's half-rate output eye diagrams are shown in Fig. 18. The measurements at 20 Gb/s are taken with the dummy photodiode removed, the front-end bandwidth set to approximately 3 GHz = $0.15 f_{\text{bit}}$, and the DFE time constant adjusted accordingly. The sensitivity for a BER of 10^{-12} was −5.8 dBm (OMA). The bathtub curve for −5.1 dBm input OMA shows 0.12 UI timing margin at a BER of 10^{-12} (Fig. 19). The receiver consumes 14.1 mW

total power, with a detailed breakdown in Table I. At 3 mW, the DFE consumes only 0.15 pJ/b at 20 Gb/s, which compares favorably with previous IIR-DFEs [15]–[17].

V. CONCLUSION

In summary, this paper studied the power and noise performance benefits of limiting the front-end bandwidth of an optical receiver having the first- and second-order responses, particularly in the presence of a DFE. A novel low-power IIR-DFE circuit was introduced to remove postcursor ISI from the bandwidth-limited front end without introducing an explicit multiplexer in the feedback path. The results of a prototype in 65 nm CMOS are summarized in Table II along with other state-of-the-art CMOS optical receivers that operate with discrete photodiodes having integrated retimers. Examples operating at comparable data rates [9], [23] and those employing DFE [12], [13] are included. Note that the photodiode used for testing here has larger capacitance and lower responsivity than any of the comparison works, which negatively impacts the maximum data rate and sensitivity. Nevertheless, this paper exhibits the best sensitivity among the compared works, excepting [12], which also uses a low-bandwidth front end followed by a discrete-time DFE to cancel postcursor ISI. (Operating at only 4 Gb/s, its front end can benefit from both higher gain and lower noise equivalent bandwidth than is possible here.) By comparison, works that focus upon the design of a CMOS front end combining low noise and high bandwidth result in either significantly higher power consumption (e.g., 3.6 pJ/b excluding clocking and retimers in [6]) or lower sensitivity (e.g., −3 dBm OMA in [23]). Moreover, this paper demonstrates an IIR-DFE at higher data rates than in previous optical receivers with a power efficiency, including clock buffers and latches, of 0.705 pJ/b.

APPENDIX

The output of the DFE is largely determined by the polarity of the differential drain current $I_{D,\text{diff}} = I_{D2a} - I_{D2b} =$

$I_{D1a} - I_{D1b}$ immediately after the clock signal (CLK) transitions to high. In this instant, transistors $M_{1a,b}$ are in triode and $M_{2a,b}$ are in saturation. The drain current of M_1 is, therefore, given by $I_{D1} = \mu_n C_{ox}(W/L)_1(v_{GS1} - V_t)v_{DS1}$. The gate–source voltage may be expanded into $v_{GS1} = V_{GS1} + V_A/2 - V_F/2$ where V_{GS1} is its common-mode value when $V_A = V_F = 0$. Similarly, $v_{DS1} = V_{DS1} - V_F/2$ (assuming $r_{ds1} \gg 1/g_{m2}$). Substituting these provides the following expression for the differential current:

$$I_{D,\text{diff}} \approx \mu_n C_{ox}\left(\frac{W}{L}\right)_1 (V_{DS1}V_A - (V_{GS1} - V_t + V_{DS1})V_F)). \tag{18}$$

The term V_{GS1} depends upon the common-mode voltage applied to the latch input at V_A, whereas V_{DS1} is determined by the saturation current of M_2 passing through triode resistance M_1

$$V_{DS1} = \frac{(W/L)_2(V_{GS2} - V_t)^2}{2(W/L)_1(V_{GS1} - V_t)}. \tag{19}$$

Substituting (19) into (18) results in

$$I_{D,\text{diff}} \approx \mu_n C_{ox}\left(\frac{W}{L}\right)_1 (V_A - A_F V_F)V_{DS1} \tag{20}$$

where $A_F = \frac{V_A}{V_F} = 1 + \frac{2(W/L)_1(V_{GS1}-V_t)^2}{(W/L)_2(V_{DD}-V_t)^2)}$.

REFERENCES

[1] D. M. Kuchta *et al.*, "64Gb/s transmission over 57m MMF using an NRZ modulated 850nm VCSEL," in *Proc. Opt. Fiber Commun. Conf.*, Mar. 2014, pp. 1–3.

[2] Y. Tsunoda, M. Sugawara, H. Oku, S. Ide, and K. Tanaka, "A 40Gb/s VCSEL over-driving IC with group-delay-tunable pre-emphasis for optical interconnection," in *IEEE Int. Solid-State Circuits Conf. Dig.*, Feb. 2014, pp. 154–155.

[3] E. Säckinger, *Broadband Circuits for Optical Fiber Communication*. Hoboken, NJ, USA: Wiley, 2005.

[4] D. Guckenberger, J. D. Schaub, D. Kucharski, and K. T. Kornegay, "1V, 10mW, 10Gb/s CMOS optical receiver front-end," in *IEEE Radio Freq. Integr. Circuits (RFIC) Symp. Dig. Papers*, Jun. 2005, pp. 309–312.

[5] P.-C. Chiang, J.-Y. Jiang, H.-W. Hung, C.-Y. Wu, G.-S. Chen, and J. Lee, "4 × 25 Gb/s transceiver with optical front-end for 100 GbE system in 65 nm CMOS technology," *IEEE J. Solid-State Circuits*, vol. 50, no. 2, pp. 573–585, Feb. 2015.

[6] T. Takemoto, H. Yamashita, T. Yazaki, N. Chujo, Y. Lee, and Y. Matsuoka, "A 4 × 25-to-28Gb/s 4.9mW/Gb/s −9.7dBm high-sensitivity optical receiver based on 65nm CMOS for board-to-board interconnects," in *IEEE Int. Solid-State Circuits Conf. Dig.*, Feb. 2013, pp. 118–119.

[7] S. D. Personick, "Receiver design for digital fiber optic communication systems, I," *Bell Syst. Tech. J.*, vol. 52, no. 6, pp. 843–874, Jul. 1973.

[8] J. E. Goell, "An optical repeater with high-impedance input amplifier," *Bell Syst. Tech. J.*, vol. 53, no. 4, pp. 629–643, Apr. 1974.

[9] M. H. Nazari and A. Emami-Neyestanak, "A 24-Gb/s double-sampling receiver for ultra-low-power optical communication," *IEEE J. Solid-State Circuits*, vol. 48, no. 2, pp. 344–357, Feb. 2013.

[10] S. Saeedi and A. Emami, "A 25Gb/s 170μW/Gb/s optical receiver in 28nm CMOS for chip-to-chip optical communication," in *Proc. IEEE Radio Freq.Integr. Circuits Symp.*, Jun. 2014, pp. 283–286.

[11] S. Saeedi, S. Menezo, G. Pares, and A. Emami, "A 25Gb/s 3D-integrated CMOS/silicon-photonic receiver for low-power high-sensitivity optical communication," *J. Lightw. Technol.*, vol. 34, no. 12, pp. 2924–2933, Jun. 15, 2016.

[12] A. V. Rylyakov, C. L. Schow, and J. A. Kash, "A new ultra-high sensitivity, low-power optical receiver based on a decision-feedback equalizer," in *Proc. Opt. Fiber Commun. Conf. Expo. Nat. Fiber Opt. Eng. Conf. (OFC/NFOEC)*, Mar. 2011, pp. 1–3.

[13] J. Proesel, A. Rylyakov, and C. Schow, "Optical receivers using DFE-IIR equalization," in *IEEE Int. Solid-State Circuits Conf. Dig.*, Feb. 2013, pp. 130–131.

[14] N. Takachio, K. Iwashita, S. Hata, K. Katsura, K. Onodera, and H. Kikuchi, "A 10 Gb/s optical heterodyne detection experiment using a 23 GHz bandwidth balanced receiver," in *IEEE MTT-S Int. Microw. Symp. Dig.*, vol. 1, Dallas, TX, USA, May 1990, pp. 149–151.

[15] B. Kim, Y. Liu, T. O. Dickson, J. F. Bulzacchelli, and D. J. Friedman, "A 10-Gb/s compact low-power serial I/O with DFE-IIR equalization in 65-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 44, no. 12, pp. 3526–3538, Dec. 2009.

[16] O. Elhadidy and S. Palermo, "A 10 Gb/s 2-IIR-tap DFE receiver with 35 dB loss compensation in 65-nm CMOS," in *Proc. Symp. VLSI Circuits (VLSIC)*, Jun. 2013, pp. C272–C273.

[17] S. Shahramian and A. Chan Carusone, "A 0.41 pJ/bit 10 Gb/s hybrid 2 IIR and 1 discrete-time DFE tap in 28 nm-LP CMOS," *IEEE J. Solid-State Circuits*, vol. 50, no. 7, pp. 1722–1735, Jul. 2015.

[18] S. M. Park and H.-J. Yoo, "1.25-Gb/s regulated cascode CMOS transimpedance amplifier for gigabit Ethernet applications," *IEEE J. Solid-State Circuits*, vol. 39, no. 1, pp. 112–121, Jan. 2004.

[19] J. Proesel, C. Schow, and A. Rylyakov, "25Gb/s 3.6pJ/b and 15Gb/s 1.37pJ/b VCSEL-based optical links in 90nm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2012, pp. 418–420.

[20] J. E. Proesel, B. G. Lee, C. W. Baks, and C. L. Schow, "35-Gb/s VCSEL-based optical link using 32-nm SOI CMOS circuits," in *Proc. Opt. Fiber Commun. Conf. Expo. Nat. Fiber Opt. Eng. Conf. (OFC/NFOEC)*, Mar. 2013, pp. 1–3.

[21] A. Rylyakov *et al.*, "A 25 Gb/s burst-mode receiver for low latency photonic switch networks," *IEEE J. Solid-State Circuits*, vol. 50, no. 12, pp. 3120–3132, Dec. 2015.

[22] C. Kromer *et al.*, "A 100mW 4 × 10Gb/s transceiver in 80nm CMOS for high-density optical interconnects," in *IEEE Int. Solid-State Circuits Conf. Dig.*, vol. 1. Feb. 2005, pp. 334–602.

[23] B. G. Lee *et al.*, "Latch-to-latch CMOS-driven optical link at 28 Gb/s," in *Proc. Conf. Lasers Electro-Opt. (CLEO)*, Jun. 2014, pp. 1–2.

Alireza Sharif-Bakhtiar (S'06) received the B.S. degree in electrical engineering from the Sharif University of Technology, Tehran, Iran, and the M.S. degree from the University of British Columbia, Vancouver, BC, Canada, in 2008 and 2011, respectively. He is currently pursuing the Ph.D. degree in electrical engineering at the University of Toronto, Toronto, ON, Canada.

His current research interests include high-speed optical interconnects.

Anthony Chan Carusone (S'01–M'07–SM'16) received the Ph.D. degree from the University of Toronto, ON, Canada, in 2002.

He is currently a Professor with the Department of Electrical and Computer Engineering, University of Toronto. He is also an occasional consultant to industry in the areas of integrated circuit design, clocking, and digital communication.

Prof. Chan Carusone co-authored the Best Student Papers at the 2007, 2008 and 2011 Custom Integrated Circuits Conferences, the Best Invited Paper at the 2010 Custom Integrated Circuits Conference, the Best Paper at the 2005 Compound Semiconductor Integrated Circuits Symposium, and the Best Young Scientist Paper at the 2014 European Solid-State Circuits Conference. He also co-authored, along with David Johns and Ken Martin, the second edition of the classic textbook *Analog Integrated Circuit Design*. He was Editor-in-Chief of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS II: EXPRESS BRIEFS in 2009, and has served on the Technical Program Committee for the Custom Integrated Circuits Conference, and the VLSI Circuits Symposium. He currently serves on the Editorial Board of the IEEE JOURNAL OF SOLID-STATE CIRCUITS, as a member of the Technical Program Committee of the International Solid-State Circuits Conference, and as a Distinguished Lecturer for the IEEE Solid-State Circuits Society.