

A $3\times$ Blind ADC-Based CDR for a 20 dB Loss Channel

Mohammad Sadegh Jalali, *Student Member, IEEE*,
Clifford Ting, Joshua Liang, Ali Sheikholeslami, *Senior Member, IEEE*,
Masaya Kibune, and Hirotaka Tamura, *Fellow, IEEE*

Abstract—This paper proposes using a 3-bit ADC to blindly sample the received data from a channel with 20 dB loss at Nyquist at $3\times$ the baud rate. By moving from $2\times$ to $3\times$ sampling, we reduce the required ADC resolution from 5-bit to 3-bit, thereby reducing the overall power consumption by a factor of 2. Measurements from our test chip fabricated in Fujitsu's 65 nm CMOS show a high frequency jitter tolerance of 0.25 UIpp for a 5 Gb/s PRBS31 with a 60'' FR4 channel.

Index Terms—ADC-based CDR, blind-sampling CDR, clock and data recovery, feed-forward CDR.

I. INTRODUCTION

CONVENTIONAL phase-tracking clock and data recovery circuits (CDR) recover a physical clock from the data and use it to sample the data once per unit interval (UI) in baud-rate sampling [1]–[4] or twice per UI in $2\times$ sampling [5], [6]. As shown in Fig. 1(a), binary CDRs sample the input data at the center of the UI with a slicer to resolve the sign of the data. In contrast, analog-to-digital converter (ADC)-based CDRs sample the data with an ADC, as shown in Fig. 1(b), allowing further equalization to be performed in the digital domain. This makes them suitable for applications where channel loss exceeds 25 dB [7]–[9]. Furthermore, digital designs are more robust to PVT variations and power supply noise, and are more easily modified to meet different requirements or ported to newer process nodes, which can reduce design costs. The performance, area and power of digital circuits also benefit more from CMOS process scaling [10] compared to their analog counterparts.

Phase-tracking ADC-based CDRs require feedback from the digital backend to analog clock circuitry such as phase interpolators (PI). Not only is designing a PI with low jitter and high linearity critical for the CDR's performance [11], but co-design of the analog and digital blocks is also required to ensure stability of the feedback loop. In contrast, blind ADC-based CDRs, shown in Fig. 1(c), eliminate this feedback entirely, by oversampling the data with a blind clock (not phase-locked to data), before recovering the clock phase as a digital code [12]–[16]. This feed-forward approach greatly simplifies the design process, eliminating

Manuscript received November 04, 2014; revised March 20, 2015; accepted March 24, 2015. Date of current version May 25, 2015. This paper was recommended by Associate Editor S. Levantino.

M. S. Jalali, C. Ting, J. Liang, and A. Sheikholeslami are with the Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON M5S 3G4, Canada (e-mail: sadegh@eecg.utoronto.ca; cliff.ting@mail.utoronto.ca; ali@ece.utoronto.ca).

M. Kibune and H. Tamura are with Fujitsu Laboratories Limited, Kawasaki 211-8588, Japan (e-mail: kibune.masaya@jp.fujitsu.com; tamura.hirotaka@jp.fujitsu.com).

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

Digital Object Identifier 10.1109/TCSI.2015.2418839

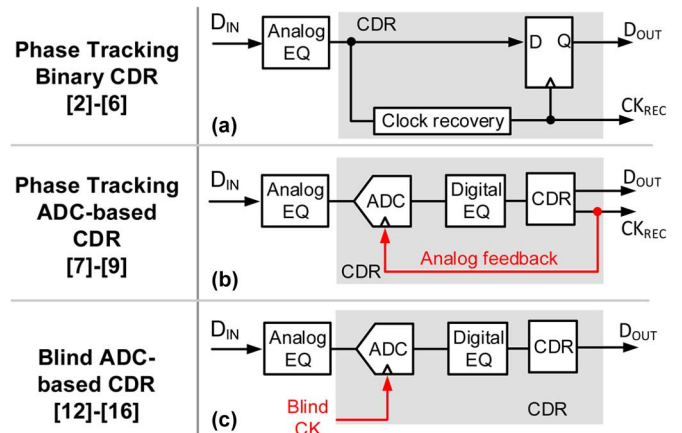


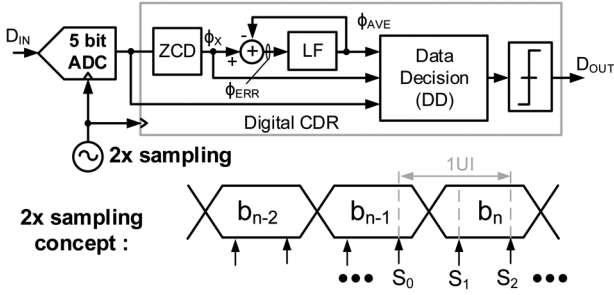

Fig. 1. Basic CDR types: (a) phase-tracking CDR; (b) phase-tracking ADC-based CDR; (c) blind ADC-based CDR.

the PI and allowing the ADC and digital backend to be designed separately, without co-simulation.

Previous blind ADC-based CDRs suffered from high analog power consumption, due to the high resolution of their ADC and the use of oversampling (61% of the total chip power in [12] is consumed in the ADC). To lower the power consumption of blind ADC-based CDRs, we propose $3\times$ oversampling [16] with a 3b flash ADC. Oversampling by $3\times$ instead of $2\times$ [12]–[14] increases the system's phase resolution, allowing ADC voltage resolution to be lowered by 2 bits, reducing the power of the system by a factor of 2. In addition, the proposed architecture reduces the input capacitance of the analog front-end (AFE), lowering the power consumption of any analog blocks driving the ADC (i.e., analog feed-forward equalization, input buffers, etc.) The proposed techniques are also scalable to systems with higher ADC resolutions, allowing them to tolerate channel losses well beyond the 20 dB achieved in this work, with reduced power consumption.

This paper presents a new test-chip that expands our initial work [16] in the following ways: 1) The effect of ADC resolution and oversampling ratio on the CDR performance is analyzed. 2) Linear equalization in both the analog and digital domain is investigated. 3) The performance of the blind CDR in the presence of frequency offset is discussed. 4) ADC offset calibration has been added. While the high offset of the ADC comparators limited the amount of ISI that could be tolerated in [16] to only 6 dB, this work includes offset cancellation circuitry to reduce the ADC offset. We will show that the performance achieved in this work is similar to that of phase-tracking ADC-based CDRs, while the design is greatly simplified.

The remainder of this paper is organized as follows. Section II reviews the basics of blind ADC-based CDRs. Section III explores


 Fig. 2. $2 \times$ sampling: architecture and example.

 Fig. 3. Definition of Φ_X , Φ_{AVE} and Φ_{PICK} .

various options to reduce the power of ADC-based CDRs. Section IV discusses the system architecture and the detailed implementation of each block. In Section V, the measurement results are discussed. Finally, Section VI concludes the paper.

II. BACKGROUND

Fig. 2 shows the basic architecture of a $2 \times$ blind ADC-based CDR [12], where the data is sampled 2 times per UI with a 5-bit ADC. Note that the general architecture does not change if the data is sampled more than two times in each UI or with a higher ADC resolution. As shown in the figure, samples S_1 and S_2 are taken from bit b_n . The blind samples enter the digital CDR in which the zero crossing detector (ZCD) first finds the position of the instantaneous zero crossing phase (Φ_X) with respect to the phase of the blind clock. This is done by performing a linear interpolation between adjacent samples, as shown in Fig. 3. The phase error (Φ_{ERR}) is obtained by subtracting (in modulo-1) Φ_{AVE} from Φ_X , where Φ_{AVE} is the average zero crossing phase. Φ_{ERR} then goes through a third-order loop filter (LF) to update the digital CDR's estimate of the zero crossing location. The phase associated with the eye center, denoted as Φ_{PICK} in Fig. 3, is found by adding 0.5 UI to Φ_{AVE} (modulo-1 addition).

The data decision (DD) block uses ADC samples, Φ_X , and Φ_{AVE} and determines the sign of the transmitted bit as the sign of the sample closer to Φ_{PICK} and farther from Φ_X [12].

In the next section, we propose techniques to reduce the power consumption of blind ADC-based receivers.

III. PROPOSED SYSTEM

As mentioned earlier, one main advantage of an ADC-based CDR is that it allows equalization to be performed in digital domain. This equalization can be applied to individual blind samples prior to the data decision scheme (Fig. 4(a)), which we refer to as *equalization-first* scheme, or after data interpolation at the eye center (and edge), which we refer to as *interpolation-first* scheme (Fig. 4(b)). The former approach, as proposed in [14], implemented 8 coefficients per DFE tap and would choose one of the coefficients depending on the location of the blind sampling phase (this is because ISI affects each sample differently as each sample is taken at a different UI position). In contrast, the latter approach only needs to implement one coefficient per DFE tap because the eye center is first estimated by interpolating between the blind samples.

While this modification reduces the digital power consumption, the high analog power consumption of the high resolution ADC remains unchanged. In the remainder of this section, we study


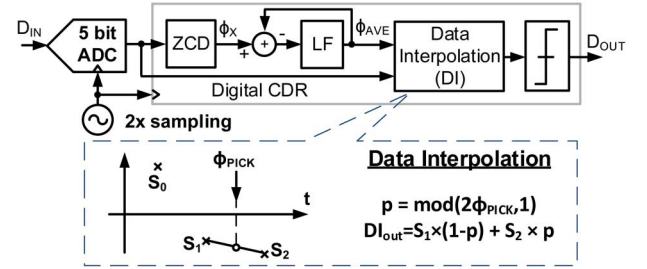


Fig. 4. (a) Equalization-first scheme; (b) interpolation-first scheme.


 Fig. 5. Modified $2 \times$ architecture with data interpolation.

the interpolation-first scheme and the effect of reducing the ADC resolution on the operation of the blocks of the receiver. We then propose a scheme to lower the overall analog power consumption without compromising performance.

A. Interpolation-First Scheme

The use of instantaneous phase in the data decision block in [14] results in having to repeat both the DD and the ZCD blocks when using a loop-unrolled DFE, leading to a power-hungry solution. To reduce the power consumption of the DFE, [17] performs interpolation in the analog domain prior to the ADC. However, performing analog data interpolation (DI) increases the complexity of the design. We perform interpolation in the digital block by replacing the DD block with a DI block, as shown in Fig. 5 (the DFE, following the DI block, is not shown in the figure for the sake of brevity). By estimating the data at the eye center (DI_{OUT}) using linear interpolation between samples on either side of Φ_{PICK} (S_1 and S_2 in the inset of Fig. 5), the equalizer implementation is simplified and the power is reduced.

B. Voltage-Phase Resolution

A large portion of analog power is consumed by the number of comparators used in the ADC. In general, if we sample the received data OSR times per UI with an n -bit flash ADC, the number of comparators used per unit interval will be $M = OSR \times (2^n - 1)$. Here, OSR is the oversampling ratio, corresponding to a *phase resolution* of $1 \text{ UI}/OSR$, and 2^n is the number of levels (L) in the voltage domain, corresponding to a *voltage resolution* of $V_{pp}/(2^n - 1)$ where V_{pp} is the peak-to-peak voltage of the input. For a constant M (corresponding to a constant analog power), we wish to determine the optimum value of OSR (and hence n) that yields the maximum high-frequency jitter tolerance. Unfortunately, it is not easy to determine this analytically. Instead, we resort to simulations in Simulink to find how OSR and n each affect the eye quality at the output of the DI block (i.e., the input to the slicer) and the jitter tolerance.

The eye quality at the DI output is determined in turn by the eye quality at the output of the ADC and the error in Φ_{AVE} , simply because these two are the only inputs to the DI block. Fig. 6(a) shows that the ADC output eye quality is largely affected by n , but not by OSR . This makes sense intuitively because reducing n increases the ADC quantization noise and this results in poor vertical eye opening, but reducing OSR (from 4 to 2 for example) only sub-samples the ADC eye and hence does not introduce error in the eye. Note that in this simulation (and in the other simulations in this section and unless otherwise stated), we assume a peak-to-peak random jitter (equal to 14σ when BER is 10^{-12}) of

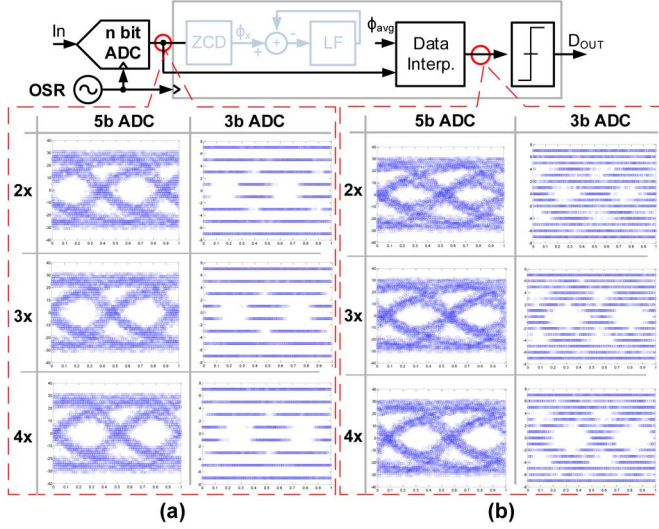


Fig. 6. Eyes at the output of the (a) ADC and (b) DI.

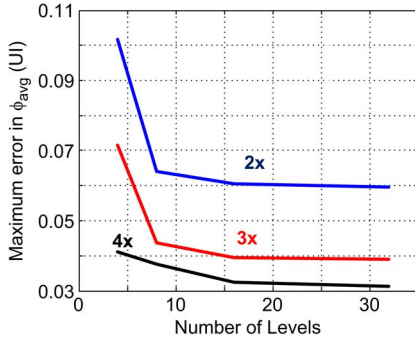


Fig. 7. Error in Φ_{AVE} versus number of ADC levels and OSR .

0.17 UIpp for the TX clock and 0.1 UIpp for the RX clock. Also, the full-scale range of the ADC is chosen such that the peak to peak voltage at the receiver is equal to the full-scale of the ADC. The oversampling ratio influences the eye quality at the DI output through two mechanisms: 1) reducing OSR increases the error in estimating Φ_{AVE} , and 2) reducing OSR increases the error in estimating the data at the center of the eye. Of the two mechanisms, the latter proves to be more critical as we show next.

To find the error in Φ_{AVE} , we apply a known frequency offset between the blind clock and the input data. This frequency offset moves the average zero crossing phase across the UI at a constant rate (i.e., $\Phi_{AVE,ideal}$ will be a ramp). We define the maximum error in Φ_{AVE} as $MAX(|\Phi_{AVE,ideal} - \Phi_{AVE,sim}|)$, where $\Phi_{AVE,sim}$ is the simulated values of Φ_{AVE} over 10^6 UIs. Fig. 7 shows this maximum error as a function of number of ADC levels ($L = 2^n$) and OSR for an 8 dB loss channel. The error in Φ_{AVE} reduces with increasing OSR . For a given OSR , the error reduces with increasing L before leveling off once L exceeds 8. More importantly, note that the maximum error in Φ_{AVE} among all cases is below 0.1 UI. This result may come as a surprise because one expects the error in Φ_X not to be negligible. However, since Φ_{AVE} is obtained by low-pass filtering Φ_X , the error in Φ_X is averaged out and Φ_{AVE} has a low error.

To find the error in Φ_X , Fig. 8 shows the maximum error in estimating Φ_X for a triangular and a rectangular input pattern. Although this figure is drawn for an oversampling ratio of 2, the derivation is performed for a general case. In both these figures, Φ_X and $\hat{\Phi}_X$ represent the actual and the estimated zero crossing phase, respectively. Also we assumed that the ADC samples are




Fig. 8. Error in estimating Φ_X for (a) triangular input (b) rectangular input.

$LSB/2$ away from the actual analog input, where LSB represents the least significant bit of the ADC.

For a triangular input (Fig. 8(a)), using similar triangles,

$$\begin{aligned} \frac{\Phi_X}{UI/OSR} &= \frac{S_0}{S_0 - S_1} \\ \frac{\hat{\Phi}_X}{UI/OSR} &= \frac{S_0 - LSB/2}{S_0 - S_1} \end{aligned} \quad (1)$$

To find the total estimation error, note that if the ADC samples in Fig. 8(a) are above the actual values by $LSB/2$ (instead of below), $\hat{\Phi}_X$ becomes greater (instead of smaller) than Φ_X . Therefore, the total peak to peak error is twice the error shown in Fig. 8(a). The total error in estimating Φ_X can be found (with respect to UI) by:

$$\begin{aligned} Err_X &= 2 \times |\Phi_X - \hat{\Phi}_X| = \frac{2}{OSR} \times \frac{LSB/2}{S_0 - S_1} (UI) \\ &= \frac{LSB}{OSR} \times \frac{OSR}{V_{PP}} (UI) = \frac{V_{PP}}{L} \times \frac{1}{V_{PP}} (UI) = \frac{1}{L} (UI) \end{aligned} \quad (2)$$

where V_{PP} is the peak to peak voltage of the input.

For a rectangular input (Fig. 8(b)), assume that S_0 is taken right before the transition. Therefore, Φ_X is approximately 0. $\hat{\Phi}_X$ can be found by:

$$\frac{\hat{\Phi}_X}{UI/OSR} = \frac{V_{PP}/2 + LSB/2}{V_{PP}} \approx \frac{1}{2} \quad (3)$$

The error in estimating Φ_X for this case is:

$$Err_X = 2 \times |\Phi_X - \hat{\Phi}_X| = \frac{1}{OSR} (UI) \quad (4)$$

Although the error in estimating the zero crossing phase could be large, the loop filter reduces this error to less than 0.1 UI, as shown in Fig. 7. We therefore assume an ideal Φ_{AVE} in the rest of this section.

Fig. 6(b) plots the DI output eye for three values of OSR (2, 3, 4) and two values of n (3, 5). The eye at the DI output is found by opening the CDR loop and manually sweeping Φ_{AVE} across the unit interval. As shown in the figure, the interpolation eye improves with increasing resolution in both the phase and the voltage domain.

Fig. 9 shows the vertical and horizontal eye opening at the output of the DI block versus OSR and L . Both vertical and horizontal eye openings increase significantly when the OSR is increased from 2 to 3, but the 3 \times and 4 \times systems have a similar performance. Also, increasing L beyond 8 does not improve the eye opening for the 3 \times and 4 \times systems. In fact, the horizontal eye opening of the 2 \times , 5b system is slightly less than that of the 3 \times , 3b system and therefore we expect the jitter tolerance of the 2 \times system to be slightly lower than that of the 3 \times system.

We define the optimum oversampling ratio as the OSR for which, given a constant M , the horizontal eye opening at the DI output is maximized. To find the optimum OSR as a function of the number of comparators used per UI (M), we sweep L from 2 to M in steps of 2, and in each case observe the horizontal eye opening. Accordingly, for each M , we find $OSR_{opt} = M/L_{opt}$ where L_{opt} is the L that yields the maximum eye opening. Fig. 10




Fig. 9. Vertical and horizontal eye opening at DI output.

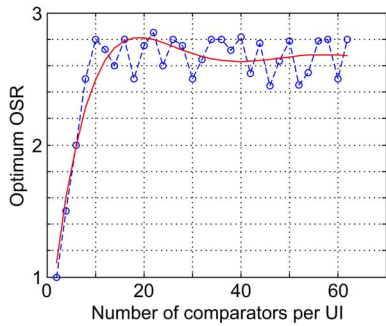


Fig. 10. Optimum oversampling ratio as a function of M .

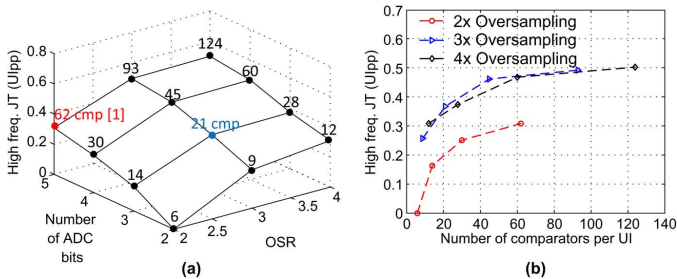


Fig. 11. High frequency jitter tolerance as a function of (a) OSR and n (b) analog power consumption [16].

plots OSR_{opt} (within $\pm 1\%$ error) as a function of M . As M increases from 2 (for which, OSR can only be 1 and the eye is closed) to about 20, the optimum OSR increases from 1 to 2.7 (≈ 3) where it settles.

Putting it all together, Fig. 11(a) shows the jitter tolerance at 500 MHz versus OSR and n for an 8 dB channel with a PRBS7 data [16]. The number of comparators used per UI is also shown for each case. As expected, the jitter tolerance (JT) increases with increasing OSR and n . This is not surprising as increasing n and OSR increases the resolution in both the voltage and the phase space. Fig. 11(b) plots the jitter tolerance, this time as a function of analog power consumption, i.e., the number of comparators per UI. It is clearly observed that the jitter tolerance of a 3× system is higher than that of a 2× system, and almost equal to that of a 4× system. Based on these observations and the results of Fig. 10, and for a JT target above 0.3 UI_{pp}, a 3× system with a 3b ADC is chosen as the optimum design point in this work.

To see the effect of channel attenuation, Fig. 12 plots the vertical and horizontal eye openings for a channel with 12 dB of loss. Again, a jump in performance is observed when going from the 2× to 3× system, while the performance improvement is much less when going from the 3× to the 4× system. Also, we verified that due to the low bandwidth of the CDR, the increased channel attenuation leaves Φ_{AVG} almost unaffected.

Finally, to see how the systems perform in the presence of jitter, Fig. 13 shows the horizontal and vertical eye openings in the presence of random jitter (RJ) with a 12 dB channel. As the OSR increases, the samples used by the DI move closer to the UI center,




Fig. 12. Vertical and horizontal eye opening at DI output for a 12 dB channel.




Fig. 13. Vertical and horizontal eye opening at DI output for a 12 dB channel versus RJ of the receiver clock. A 3b ADC is used in all the systems.




Fig. 14. 3 ×, 3b blind ADC-based receiver.

where the slope of the input is low, reducing the impact of clock jitter on the operation of the DI.

IV. RECEIVER DESIGN

Fig. 14 shows a system diagram of the analog front-end that samples the 5 Gb/s data signal at 15 GS/s. The receiver has a single ADC that is composed of eight interleaved sub ADCs. The received signal is fed to a continuous time linear equalizer (CTLE), which is designed to peak at 2.5 GHz. The CTLE output drives the 8 interleaved, 3-bit flash ADCs. The 3-bit ADCs blindly sample the received signal using 8 phases of a 1.875 GHz clock. The samples are demuxed by 32 and are fed to the digital CDR. The blind clock, provided by an off-chip 7.5 GHz source, is divided by a CML shift register into the 8 phases required by the ADCs. One of the 1.875 GHz phases is further divided into a 470 MHz clock that drives the digital CDR. At the beginning of the operation, the ADCs are calibrated to remove their offset.

To study the effect of gain mismatch and timing skew between interleaved ADCs on the performance of the system, Fig. 15 shows the simulated jitter tolerance with a 9 dB channel and a PRBS31 pattern, with and without $\pm 10\%$ gain mismatch as well as $\pm 10\%$ clock skew between the ADCs. The $\pm 10\%$ of gain and phase mismatch reduce the high frequency jitter tolerance by less than 0.1 UI_{pp}. Note that the effect of mismatch between interleaved ADCs is less significant due to the low sampling rate and the low resolution of the ADCs. Furthermore, the low pass nature of the channel reduces the slope of the input, further reducing this effect.

In the rest of this section, we explain the design details of individual building blocks in Fig. 14.

A. Equalization

Linear equalization can be used to flatten the overall frequency response up to the Nyquist frequency. Fig. 16 shows the analog front-end used in this work. The data enters the continuous-time




Fig. 15. Effect of gain mismatch and timing skew between interleaved ADCs on the performance of the system.

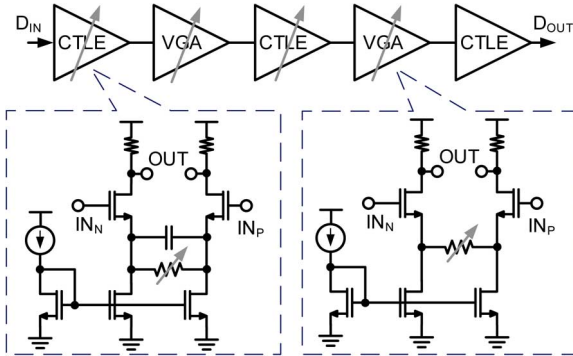


Fig. 16. Analog front-end equalization.

linear equalizer (CTLE) which is a source degenerated differential amplifier, to create peaking. This is followed by a variable gain amplifier (VGA) which compensates for the DC loss of the CTLE. The last stage of CTLE has a fixed capacitive source degeneration to equalize the large load of the ADCs and to provide a flat frequency response up to 2.5 GHz. Since the input load capacitance of the ADCs is large, the linear equalizer is realized using multiple stages.

Although linear equalization can be realized in the digital domain in the form of a feed-forward equalizer (FFE), it suffers from the fact that the peaking required to equalize the channel also amplifies noise. Because a digital FFE combines delayed, scaled versions of the ADC output, quantization noise also propagates through the FFE, limiting its use with lower ADC resolutions. We therefore used a CTLE in this work. Although the CTLE consumes additional power, this solution leads to a better trade-off between power efficiency and channel loss compensation.

We follow the CTLE with a DFE in the digital domain. Note that the DFE performance does not depend on the ADC resolution because the DFE coefficients can have a higher resolution than the ADC. The ADC resolution does however limit the achievable equalization, by requiring that the equalized eye opening (without ADC) be greater than 1 LSB for error-free recovery with the ADC. This sets the minimum eye-opening needed after equalization.

B. Analog to Digital Converter

The 3-bit flash ADC includes seven comparators and RS latches, and a thermometer-to-binary decoder. In order to reduce the power consumption, the clocked comparators directly sample the data signal without preamplifiers. Also, to reduce the loading on the CTLE, small transistors (shortest channel length) are used in the comparator, increasing the offset of the comparator. This necessitates the use of offset cancellation circuitry, which will be explained later in the next section. The latched thermometer code is converted into a binary sample using a Wallace adder [18].

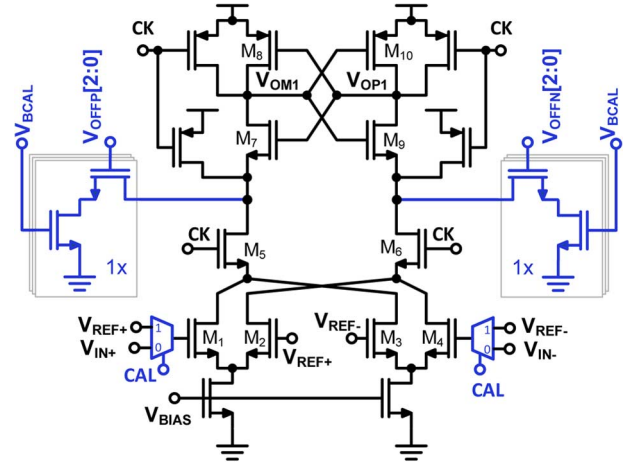


Fig. 17. Modified strongArm comparator with offset cancellation.

This is done by adding the outputs of all the seven RS latches of the ADC.

Fig. 17 shows the modified strongARM comparator which is used for its low power consumption and narrow sampling aperture [19]. To reduce kickback on the data signal and the reference ladder, M5 and M6 are stacked on top of the 4 input transistors (M1–M4). The CTLE output is designed to have a high (and well controlled) common-mode voltage of 0.8 V, which minimizes the impact of common-mode variations on the performance of this comparator. The highlighted offset cancellation transistors, operate by steering some current from the right or the left branch of the comparator to ground, before the current gets to the source of M7 and M9, where the positive feedback is activated. Two 3-bit numbers, V_{OFFFP} and V_{OFFFN} , determine the amount of current that will be subtracted. Note that since current only needs to be subtracted from one branch, one of these two signals are always zero. V_{BCAL} sets the resolution of the offset cancellation circuitry and is set off-chip to be 0.5 V. The detailed algorithm for the ADC offset cancellation will be explained in the following section. Finally, the minimum and the maximum voltages of the ADC ladder are supplied off-chip in this work. In designs where the number of pins are limited, an ADC reference generator [20] can be used to generate these voltages.

C. ADC Calibration

To cancel the offset of the ADC comparators, CAL is set high. This connects both gates of M1 and M2, as well as M3 and M4 in Fig. 17 to the same reference voltage. Both V_{OFFFN} and V_{OFFFP} are set to “000.” The offset in the input transistors causes the output of the comparator to be either “1” or “0.” At this point, V_{OFFFP} is forced to “111” while V_{OFFFN} is kept at “000.” This causes the comparator output to go high. The codes for V_{OFFFP} and V_{OFFFN} are then swept in opposite directions until the comparator output becomes low, at which point, the comparator residual offset would be less than 1LSB of the calibration circuitry (the LSB of the calibration circuitry, nominally set to be 20 mV, can be adjusted by changing the value of V_{BCAL}). CAL is set to zero when calibration is finished, and the comparator operates normally.

The above calibration circuitry requires access to the output of each individual comparator in order to calibrate its offset. However, the only observable output is the thermometer-coded sum of all the comparator outputs. To make the output of individual comparators observable, except for comparator i whose offset is being canceled, V_{OFFFP} and V_{OFFFN} for the other six comparators are respectively forced to “000” and “111.” Therefore, with a high probability, the output of the other comparators are low. The use

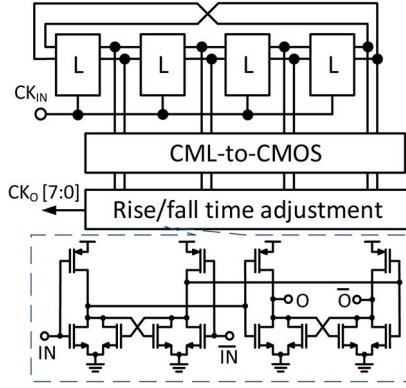


Fig. 18. Clock divider.




Fig. 19. Digital CDR implementation.

of a Wallace tree encoder (as opposed to a bubble error correction type decoder), now makes the output of the comparator under calibration observable. V_{OFFP} and V_{OFFN} for comparator i are then swept until its offset is canceled (the ADC binary output toggles). After this, V_{OFFP} and V_{OFFN} for this comparator are set to “000” and “111,” and the next comparator is calibrated.

D. Clock Divider

Fig. 18 shows the clock division circuitry, in which the high speed CML clock is divided by a factor of four using a CML shift register. The CML outputs of each latch is then converted into CMOS using a CML-to-CMOS converter. Before the CMOS clocks are distributed to the ADCs, a rise/fall time adjuster block, shown in the inset of the figure, corrects for any mismatch between the rise and fall times of the clocks.

E. Digital CDR Design

Fig. 19 shows the detailed implementation of the CDR. Thirty-two demuxed ADC samples corresponding to 10.667 UIs (3 samples per UI) enter the digital CDR. Since the CDR processes an integer number of UIs, the 32 samples first enter the variable UI controller block. The role of this block is to convert three of these 32-sample batches (which arrive in three consecutive clock cycles) into three batches of 30, 33 and 33 samples, corresponding to 10, 11 and 11 UIs. Dummy bits are inserted at the beginning of the first 30-sample group to make it equal-size with the other two groups. A flag ($UI\#$) denotes the number of non-dummy data samples in each group. The dummy bit is discarded in the FIFO.

The 33-sample batch of data then enters the data formatter block, where the ADC output codes are converted from (0 to 7) to (-7 to 7), making the implementation of the ZCD and data decision blocks easier.

To find the instantaneous zero crossing phase (Φ_X), the ZCD divides the UI into three regions, corresponding to the $3 \times$ sampling technique. This is shown in Fig. 20(a), where A_i , B_i , C_i , and D_i are the samples taken by the ADCs and cover one full UI.

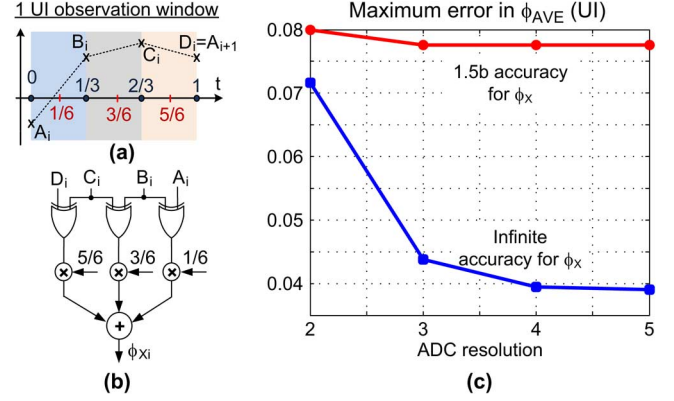


Fig. 20. ZCD (a) operation, (b) implementation, and (c) estimation error.

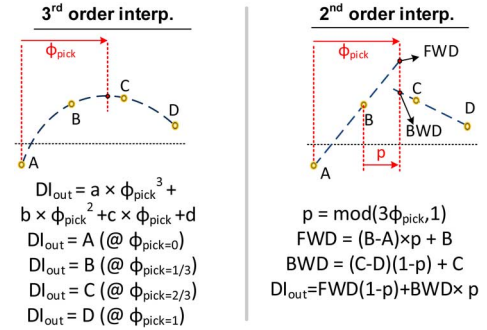


Fig. 21. Implementation of data interpolation.

Note that $D_i = A_{i+1}$. The ZCD XORs the signs of adjacent ADC codes to yield which region Φ_X belongs to. Three levels (i.e., 1/6, 3/6, and 5/6) are used to represent Φ_X , as shown in Fig. 20(b).

While interpolation between adjacent samples can be used to fine tune our estimation of the zero crossing phase, we would like to avoid this to lower the power consumption. Fig. 20(c) shows that although the peak error in estimating Φ_{AVE} doubles when using a 3-level Φ_X , this value is still less than 0.08 UI, leading to a simulated high frequency JT loss of only 0.05 UIpp.

Fig. 21 shows the operation of the data decision block, where two ADC samples before Φ_{PICK} are A and B and two ADC samples after Φ_{PICK} are C and D . The best estimate of the eye value at Φ_{PICK} can be obtained by fitting a third-order polynomial to these four points, as shown in the left inset of Fig. 21. However, this approach is hardware intensive. Instead, we use second-order interpolation, shown in the right inset of Fig. 21. Here, the eye center is estimated by first extrapolating between samples A and B (FWD) and between C and D (BWD) and then performing a weighted sum on the values of these two lines at Φ_{PICK} . Therefore, the value of the eye at Φ_{PICK} can be found by:

$$DI_{out} = (B - A + C - D)p(1 - p) + (C - B)p + B \quad (5)$$

where p is the distance between Φ_{PICK} and B . Implementing the above equation is hardware intensive, but can be simplified by restricting p to discrete values. Our simulations show that limiting the resolution of p to 2b lowers the high frequency jitter tolerance by less than 0.05 UIpp.

We observed that the eye opening improves significantly when performing second-order interpolation instead of first-order interpolation, while the eye opening when performing second and third-order interpolation are similar. Our simulations show a 0.1 UIpp increase in high frequency jitter tolerance when




Fig. 22. Frequency response of the (a) second-order interpolation; (b) 1st order interpolation operating on a 5 Gb/s data.




Fig. 23. CSM operation: (a) $f_{RX} < f_{TX}$, add a bit, (b) $f_{RX} > f_{TX}$, remove a bit, (c) summary of the CSM operation.

using second-order interpolation instead of linear interpolation. Furthermore, Fig. 22(a) shows the frequency response of the proposed second-order interpolation as a function of p . As shown in this figure, the peaking inherent to the interpolation operation provides up to 3 dB of equalization at the Nyquist frequency. This can be intuitively explained by considering Fig. 21 and assuming that a “1” bit occurs in the middle of a string of zero bits (lonely “1”). Depending on the amount of ISI, it is possible to have all four samples below zero. However, since we extrapolate between adjacent samples, the estimated UI center can still be positive, effectively opening the eye. No peaking is observed in the transfer function of linear interpolation, shown in Fig. 22(b).

The DI block nominally outputs 11 digits associated with the 33 samples going into it. However, in the presence of frequency offset, the cycle slip monitor (CSM) block occasionally inserts a digit into the recovered stream if $f_{RX} < f_{TX}$, or removes a digit from it if $f_{RX} > f_{TX}$, as shown in Fig. 23. In the case of $f_{RX} < f_{TX}$, Φ_X decreases over time, decreasing Φ_{AVE} and Φ_{PICK} . This causes Φ_{PICK} to occasionally decrease from $0 + \epsilon$ to $1 - \epsilon$, where ϵ is a small positive number ($\Phi_{PICK} \in [0, 1)$). Fig. 23(a) shows this case, where Φ_{PICK} for batch i is $0 + \epsilon$ and for batch $i + 1$ is $1 - \epsilon$. In this case, none of the 11 digits from the DI block represent the first eye in batch $i + 1$, and hence the CSM block inserts DI_0 into the recovered bit stream [13]. Similarly, in Fig. 23(b), the cycle slip monitor block removes DI_0 from the recovered bit stream when $f_{RX} > f_{TX}$; otherwise the first eye center in batch $i + 1$ would have been selected twice. Fig. 23(c) summarizes the operation of the CSM block for all values of Φ_{PICK} . The DI block outputs 10–12 digits, depending on the position of Φ_{PICK} . This variable length output is absorbed in the elastic buffer following the CDR [21].

A one-tap loop-unrolled DFE equalizes the interpolated eye center. Fig. 24(a) shows the operation of the DFE. First assume that due to ISI, all four samples are negative when a lonely “1” is received. Although second-order interpolation somewhat opens the eye, this is not enough, as the interpolated eye is still below zero. By adding the ISI of the previous bit, α , to the interpolated eye center, the equalized eye center is positive and the bit is recovered correctly. A similar situation exists when a lonely “0” is received. Fig. 24(b) shows the implementation of the loop-unrolled

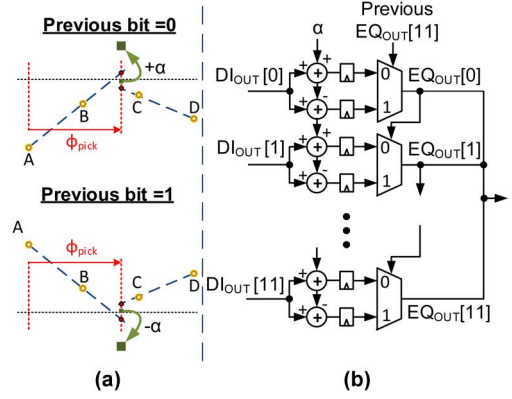


Fig. 24. (a) DFE operation. (b) DFE implementation.




Fig. 25. Simulated (a) vertical eye opening versus channel loss; (b) bathtub curve with an 11 dB channel.

DFE, where α is both added and subtracted from the DI output. The value of the previous bit chooses the correct signal as the DFE output. By subtracting the channel ISI from the interpolated eye center, the DFE is able to increase the eye opening to 1LSB for an 11 dB channel. For a total loss of 22 dB (and with the CTLE providing about 10 dB of boost), the remaining first post-cursor ISI is roughly 3 times smaller than the main cursor. The range of the first post cursor tap was designed accordingly.

To verify the performance of the DFE, Fig. 25(a) shows the vertical eye opening of data before and after equalization versus channel attenuation (at the Nyquist frequency). Fig. 25(b) shows the bathtub curve of the digital CDR with DFE on and off for the 11 dB channel. The minimum verifiable BER is 10^{-6} due to limited simulation speed. Both simulations are done with a PRBS31 pattern.

V. MEASUREMENT RESULTS

The chip, shown in Fig. 26, is fabricated in Fujitsu's 65 nm CMOS process. The area of each block is shown. To measure the chip, a PRBS generator (Centellax TG1B1-A) is clocked with a 5 GHz source (Centellax TG1C1-A), and is connected to the chip through an FR4 channel, a Tyco backplane, or a combination of both. For jitter tolerance measurements, sinusoidal jitter (SJ) is inserted on the clock of this PRBS generator. A 7.5 GHz clock provides the blind clock to the chip. An FPGA programs the chip and a logic analyzer monitors the output of the chip.

A. Calibration

As previously mentioned, the size of the ADC comparator input transistors is minimized to reduce loading on the CTLE output (no pre-amplifier is used to save power). This leads to a large measured offset, on the order of 100 mV, between the input pairs, which is significantly reduced after calibration. The full-scale range is 500 mV for the ADCs, which is the same as the peak to peak voltage at the ADC input.

To see the effect of calibration on the system performance, Fig. 27 shows the measured superimposed eyes of all the ADCs with a PRBS31 pattern before and after ADC calibration. Eight




Fig. 26. Chip photo.

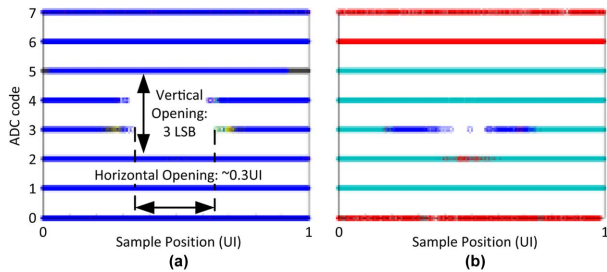


Fig. 27. Measured ADC eye with PRBS31 pattern through a 32'' FR4 channel: (a) after calibration; (b) before calibration.




Fig. 28. Measured INL and DNL of the ADCs before and after calibration.

colors are used to show the output of the eight ADCs. In this measurement, the 5 Gb/s data goes through a 32'' FR4 channel. The CDR achieves error free ($BER < 10^{-12}$) operation after calibration, but makes occasional mistakes before that. Fig. 28 shows the measured integral non-linearity (INL) and differential non-linearity (DNL) of one of the eight ADCs before and after calibration. Without offset cancellation, the INL is as large as 1LSB. Calibration is performed at start-up, prior to all of the remaining measurements.

B. Digital Receiver Without CTLE

To verify the basic operation of the CDR, the data is directly connected to the chip through a 48'' SMA cable. Both DFE and CTLE are disabled. For a PRBS7 pattern, the jitter tolerance is 0.63 UIpp (equipment limit) at 100 MHz, while for a PRBS31 pattern, the jitter tolerance decreases to 0.52 UIpp at 100 MHz. The low frequency jitter tolerance is 16 UIpp at 100 kHz and 32 UIpp at 50 kHz for both patterns. The Centellax TG1C1-A clock source can generate a maximum high frequency SJ of 0.63 UIpp, and a maximum low frequency SJ of 16 UIpp at 100 kHz and 32 UIpp at 50 kHz and below.

To verify the operation of the DFE, the PRBS31 data is connected to the chip through a 16'' FR4 channel. Fig. 29(a) shows the jitter tolerance with and without the DFE. CTLE is disabled in both measurements. The high frequency jitter tolerance increases by about 0.15 UIpp to 0.5 UIpp when the DFE is turned on. Due to equipment limits, the low frequency jitter tolerance curve is flat




Fig. 29. Measured jitter tolerance for a PRBS31 pattern at 5 Gb/s over 16'' FR4 channel: (a) with and without DFE and with no frequency offset; (b) with 500 ppm and 1000 ppm of frequency offset and with DFE on.




Fig. 30. (a) Measured jitter tolerance at 100 MHz in the presence of ± 1500 ppm of frequency offset (5 Gb/s PRBS31 over 16'' FR4 channel). (b) Measured jitter tolerance comparison between a 6 Gb/s and a 5 Gb/s receiver.

below 50 kHz. We denote this by a dotted line in all jitter tolerance figures in this section.

The high frequency jitter tolerance is decreased in the presence of frequency offset. This is shown in Fig. 29(b), assuming a frequency offset of 500 ppm and 1000 ppm, for the 16'' FR4 channel. The loss of this channel and the probe card is 6 dB at the Nyquist frequency. Fig. 30(a) shows the measured JT at 100 MHz in the presence of frequency offset. The decrease in JT is caused by both the mismatch between the ADCs as well as by the offset of the worst ADC, as in the presence of frequency offset and over time, all the ADCs will be used to estimate the eye center. The maximum tolerable frequency offset is limited by the bandwidth of the loop, which determines the accuracy with which Φ_{AVE} tracks frequency offset. The loop filter was designed to tolerate a frequency offset of at least 1500 ppm.

As mentioned in the introduction section, one major advantage of blind ADC-based CDRs is its portability to different data rates and technologies. By increasing the data rate to 6 Gb/s and the frequency of the blind clock to 9 GHz, we can transform the receiver into one that operates at 6 Gb/s. Fig. 30(b) compares the measured jitter tolerance results of a 5 Gb/s and 6 Gb/s receiver for a 32'' FR4 channel with a PRBS31 pattern. The loss of this channel and the probe card is 9 dB and 10.5 dB at 2.5 GHz and 3 GHz, respectively. As seen here, the high frequency jitter tolerance drops to roughly 0.3 UIpp at 6 Gb/s, which is due to an increase in the channel ISI and a decrease in the ADC ENOB (this is caused by there being less time for the drains of the input transistors to settle before the comparator has to make its decision). Note that the power consumption scales linearly with data rate. Also, for this channel, the operation is not error free when the DFE is disabled. In the rest of this section, the DFE is kept on and manually adjusted for best performance.

C. Digital Receiver With CTLE

To verify the performance of the CTLE together with the blind receiver, we tested our design with various length of FR4 channels and Tyco channels. Fig. 31 shows the measured frequency response of two of the channels used for measurements. The 60'' FR4 channel has a loss of -14 dB at 2.5 GHz while the 34'' Tyco

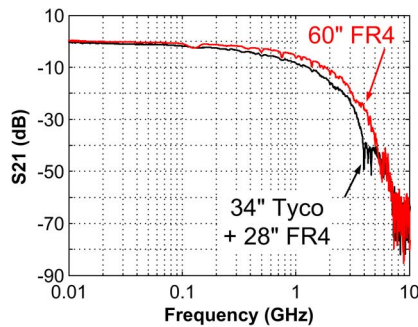


Fig. 31. Measured frequency response of the channel.




Fig. 32. Measured jitter tolerance for a (a) 51'', 55'' and 60'' FR4 channel (b) 34'' Tyco, with 16'' and 28'' FR4 channels.

together with the 28'' FR4 has a measured loss of -19 dB at the Nyquist frequency. The probes have a measured loss of 1 dB at 2.5 GHz. Since the 60'' channel was obtained by cascading three 16'' and one 12'' FR4 channel in series, the reflections are high. S11 is as high as -5 dB at 125 MHz for this channel. For the other channel, S11 stays below -10 dB until 2.5 GHz.

Fig. 32(a) shows the resulting jitter tolerance for the 51'', 55'' and 60'' FR4 channels. The measured eye is closed for the 55'' and 60'' FR4 channels, while barely open for the 51'' channel. Different CTLE and DFE coefficients were used in each measurement. The chip is then tested with a 34'' Tyco backplane, a 34'' Tyco backplane and a 16'' FR4 channel, a 34'' Tyco backplane and a 28'' FR4 channel. Fig. 32(b) shows the jitter tolerance of all these cases. The ADC eye (not shown) is closed in the last case. The high frequency JT of the cascade of the 34'' Tyco backplane and the 16'' FR4 channel is 0.2 UIpp.

The ADC and DEMUX consume 40.8 mW (40 μ W of which is consumed in the ladder), the clock divider consumes 14.4 mW, the CTLE and the input buffers consumes 24.2 mW and the digital CDR consumes 27 mW. Table I summarizes the results and compares this work against previous work. Also unless specified, PRBS7 pattern is used to verify operation.

Fig. 33 shows the power efficiency (in mW/Gb/s) versus channel loss, where the dashed lines show the two main trends. Compared to the other blind ADC-based CDRs, this work achieves a much lower power consumption while tolerating a higher channel attenuation, while compared to the previous phase-tracking architectures, this work achieves a similar power consumption, without having the complexity of a phase tracking loop.

VI. CONCLUSION

By moving from $2\times$ to $3\times$ sampling, we manage to reduce the ADC power by a factor of 3 through lowering its required resolution from 5 to 3 bits. In addition, by redesigning the CDR, we reduce the digital power consumption in two ways: 1. In this work, the DD block only uses Φ_{AVE} , while in [12]–[14] both Φ_{AVE} and Φ_X are used to make a decision. By dropping Φ_X from the decision making process, we can afford to lower the accuracy in es-

TABLE I
SUMMARY AND COMPARISON WITH PREVIOUS WORK

CDR	Data Rate (Gb/s)	Tech. (nm)	AFE power (mW)	Dig. power (mW)	Max. ISI (dB)	ADC res. (n)	Total power (mW)
[2] [‡]	16	40	NA	NA	34	NA	235.2*
[3] [†]	16	22	NA	NA	34	NA	59.2
[4] [◊]	11.8	32	NA	NA	25	NA	78
[5] [‡]	11.1	65	NA	NA	16	NA	78
[6] [‡]	10	0.13	NA	NA	18	NA	133
[7]	7.5	65	NA	NA	24	4.5	330*
[8] [†]	10	65	NA	NA	26	6	500
[9]	10.3	40	195	NA	34	6	NA
[12]	5	65	110	68.4	10	5	178.4
[13]	5	65	NA	NA	15	5	280
[14]	5	65	NA	57.6	13.3	5	211.2
[15]	10	65	109	111.6	NA	5	306
[16] [†]	5	65	38.4	42	6	3	94.8
[17]	32	65	234.4	74	22	NA	308.4
This work w/o CTLE[†]	5	65	55.2	27	10	3	82.2
This work w CTLE[†]	5	65	79.4	27	20	3	106.4

* Both TX and RX included in the power number

[†] PRBS31

[◊] PRBS23

[‡] Non-ADC based

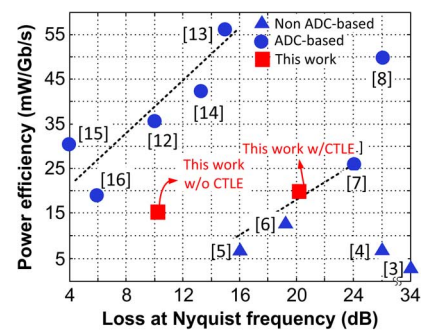


Fig. 33. Power efficiency versus channel loss.

timating Φ_X to three levels (corresponding to 3 samples per UI) and simplify the PD design. Any high frequency error in Φ_X is heavily attenuated and filtered by the ensuing LPF, maintaining a high accuracy for Φ_{AVE} . 2. Since we have access to the interpolated data at the eye center, we can directly equalize it. By moving linear equalization from the digital domain to the analog domain, we demonstrate a BER better than 10^{-12} for a PRBS31 pattern going through a 20 dB loss channel.


REFERENCES

- [1] M. Jalali *et al.*, "An 8 mW frequency detector for 10 Gb/s half-rate CDR using clock phase selection," in *Proc. CICC*, Sep. 2013, pp. 1–4.
- [2] A. Joy *et al.*, "Analog-DFE-based 16 Gb/s SerDes in 40 nm CMOS that operates across 34 dB loss channels at Nyquist with a baud rate CDR and 1.2 Vpp voltage-mode driver," in *ISSCC Dig. Tech. Papers*, Feb. 2010, pp. 350–351.
- [3] P. Francese *et al.*, "A 16 Gb/s 3.7 mW/Gb/s 8-tap DFE receiver and baud rate CDR with 30 kppm tracking bandwidth," in *ASSCC Dig. Tech. Papers*, Nov. 2013, pp. 33–36.
- [4] F. Spagna *et al.*, "A 78 mW 11.8 Gb/s serial link transceiver with adaptive RX equalization and baud-rate CDR in 32 nm CMOS," in *ISSCC Dig. Tech. Papers*, Feb. 2010, pp. 366–367.
- [5] J. Bulzacchelli *et al.*, "A 78 mW 11.1 Gb/s 5-tap DFE receiver with digitally calibrated current-integrating summers in 65 nm CMOS," in *ISSCC Dig. Tech. Papers*, Feb. 2009, pp. 368–369.
- [6] S. Gondi and B. Razavi, "Equalization and clock and data recovery techniques for 10-Gb/s CMOS serial-link receivers," *IEEE J. Solid-State Circuits*, vol. 42, no. 9, pp. 1999–2011, Sep. 2007.
- [7] M. Harwood *et al.*, "A 12.5 Gb/s SerDes in 65 nm CMOS using a baud-rate ADC with digital receiver equalization and clock recovery," in *ISSCC Dig. Tech. Papers*, Feb. 2007, pp. 436–437.

- [8] J. Cao *et al.*, "A 500 mW digitally calibrated AFE in 65 nm CMOS for 10 Gb/s serial links over backplane and multimode fiber," in *ISSCC Dig. Tech. Papers*, Feb. 2009, pp. 370–371.
- [9] B. Zhang *et al.*, "A 195 mW/55 mW dual-path receiver AFE for multi-standard 8.5-to-11.5 Gb/s serial links in 40 nm CMOS," in *ISSCC Dig. Tech. Papers*, Feb. 2013, pp. 34–35.
- [10] Y. Chiu *et al.*, "Scaling of analog-to-digital converters into ultra-deep-submicron CMOS," in *Proc. CICC*, Sept. 2005, pp. 375–382.
- [11] B. Casper and F. O'Mahony, "Clocking analysis, implementation and measurement techniques for high-speed data links—A tutorial," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 56, no. 1, pp. 17–39, Jan. 2009.
- [12] O. Tyshchenko *et al.*, "A 5 Gb/s ADC-based feed-forward CDR in 65 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 45, no. 6, pp. 1091–1098, Jun. 2010.
- [13] H. Yamaguchi *et al.*, "A 5-Gb/s transceiver with an ADC-based feed-forward CDR and CMA adaptive equalizer in 65-nm CMOS," in *ISSCC Dig. Tech. Papers*, Feb. 2010, pp. 168–169.
- [14] S. Sarvari *et al.*, "A 5 Gb/s speculative DFE for $2 \times$ blind ADC-based receivers in 65-nm CMOS," in *IEEE Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2010, pp. 69–70.
- [15] C. Ting *et al.*, "A blind baud-rate ADC-based CDR," in *ISSCC Dig. Tech. Papers*, Feb. 2013, pp. 122–123.
- [16] M. Jalali *et al.*, "A $3 \times$ blind ADC-based receiver," in *ASSCC Dig. Tech. Papers*, Nov. 2013, pp. 349–352.
- [17] Y. Doi *et al.*, "32 Gb/s data-interpolator receiver with 2-tap DFE in 28 nm CMOS," in *ISSCC Dig. Tech. Papers*, Feb. 2013, pp. 36–37.
- [18] C. S. Wallace, "A suggestion for a fast multiplier," *IEEE Trans. Electron. Comput.*, vol. EC-13, no. 1, pp. 14–17, Feb. 1964.
- [19] M. El-Chammas and B. Murmann, "A 12-GS/s 81-mW 5-bit time-interleaved flash ADC with background timing skew calibration," *IEEE J. Solid-State Circuits*, vol. 46, no. 4, pp. 838–847, Apr. 2011.
- [20] M. Chahardori *et al.*, "A 4-bit, 1.6 GS/s low power flash ADC, based on offset calibration and segmentation," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 60, no. 9, pp. 2285–2297, Sep. 2013.
- [21] A. Sheikholeslami and H. Tamura, "An design metrics for blind ADC-based wireline receivers," in *Proc. CICC*, Sep. 2013, pp. 1–8.



Mohammad Sadegh Jalali received the B.S. degree (with honors) in electrical engineering from the University of Tehran, Iran, the M.S. degree from the University of British Columbia, Canada, and the Ph.D. degree from the University of Toronto, Canada, in 2008, 2010, and 2014, respectively. In 2014 he joined Semtech-Snowbush IP, and has been engaged in the development of multistandard SerDes IP.



Clifford Ting received the B.A.Sc. and M.A.Sc. degrees in electrical engineering from the University of Toronto, Canada, in 2007 and 2013, respectively. His research interests are in the design of integrated circuits for high-speed chip-to-chip communications, including clock-and-data recovery blocks and equalizers. In 2013, he joined Intel Corporation and has been engaged in the design of high-speed IO.



Joshua Liang received the B.A.Sc. degree in engineering science and M.A.Sc. degree in electrical engineering from the University of Toronto, Canada, in 2007 and 2009, respectively. From 2009 to 2011 he was an Analog Designer with Zarlink Semiconductor (now Microsemi), where he worked on circuits for low-jitter clock synthesis. Since 2012 he has been working toward the Ph.D. degree in electrical engineering at the University of Toronto, Canada, in the area of circuit design for high-speed wireline and optical communications.



Ali Sheikholeslami (S'98–M'99–SM'02) received the B.Sc. degree from Shiraz University, Iran, in 1990 and the M.A.Sc. and Ph.D. degrees from the University of Toronto, Canada, in 1994 and 1999, respectively, all in electrical engineering. In 1999, he joined the Department of Electrical and Computer Engineering at the University of Toronto, where he is currently a Professor. He was on research sabbatical with Fujitsu Labs in 2005–2006, and with Analog Devices in 2012–2013. His research interests are in analog and digital integrated circuits, high-speed

signaling, and VLSI memory design. He has coauthored over 50 journal and conference articles and 8 patents.

He served on the Memory, Technology Directions, and Wireline Subcommittees of the ISSCC in 2001–2004, 2002–2005, and 2007–2013, respectively. He is currently an Associate Editor for the *Solid-State Circuits Magazine* and the Educational Events Chair for ISSCC. He was an Associate Editor for the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS—PART I: REGULAR PAPERS for 2010–2012, and the program chair for the 2004 IEEE ISMVL. He is a registered professional engineer in Ontario, Canada.

Dr. Sheikholeslami has received numerous teaching awards including the 2005–2006 Early Career Teaching Award and the 2010 Faculty Teaching Award, both from the Faculty of Applied Science and Engineering at the University of Toronto.



Masaya Kibune was born in Kanagawa, Japan, in 1973. He received the B.S. and M.S. degrees in applied physics from Tokyo University, Tokyo, Japan, in 1996 and 1998, respectively.

In 1998, he joined Fujitsu Laboratories, Ltd., Kanagawa, Japan. He has been engaged in research and design of high-speed IO with CMOS.



Hirotaka Tamura (M'02–SM'10–F'13) received his B.S., M.S., and Ph.D. degrees in electronic engineering from Tokyo University, Tokyo, Japan, in 1977, 1979, and 1982, respectively. He joined Fujitsu Laboratories, Japan, in 1982. After being involved in the development of different exploratory devices such as Josephson junction devices and high-temperature superconductor devices, he moved into the field of CMOS high-speed signaling in 1996. His first contribution to this area was in the designing of a receiver front-end for DRAM-to-processor

communications. Then, he got involved in the development of a multi-channel high-speed I/O for server interconnects. Since then he has been working in the area of architecture- and transistor-level design for CMOS high-speed signaling circuits.