

11.7 A 56Gb/s 50mW NRZ Receiver in 28nm CMOS

Atharav Atharav, Behzad Razavi

University of California, Los Angeles, CA

The power consumption of wireline transceivers has become increasingly critical as higher data rates and a larger numbers of lanes per chip are sought [1-6]. While attractive for lossy channels, PAM-4 signaling has mostly dictated ADC-based receivers (RXs) and relatively high power consumption [1,2]. Non-return-to-zero (NRZ) receivers, on the other hand, can be realized in the analog domain, potentially consuming less power, but they must deal with a greater loss. This paper introduces an NRZ RX that achieves more than a twofold reduction in power while exhibiting BER 10^{-12} for a channel loss of 25dB at 28GHz. The proposed design can compete with PAM-4 counterparts and/or serve in 112Gb/s systems that must also support 56Gb/s reception. Figure 11.7.1 shows the RX architecture. The data path consists of a CTLE core, a DFE core, a discrete-time linear equalizer (DTLE) [4], and a DMUX. The receiver performance is greatly improved by a number of feedforward and feedback paths. Also proposed is a half-rate “band-pass” CDR that avoids loading the main data path and the use of quadrature VCOs.

The CTLE must provide a wide bandwidth, requiring inductive peaking, and a high boost, typically demanding several stages and high power. We propose the topology shown in Fig. 11.7.2, where the main path comprises two simple, degenerated differential pairs, G_{m1} and G_{m2} , but the boost is considerably raised by the three feedforward paths. The high-pass action of these paths manifests itself at frequencies beyond the main path’s bandwidth, thus helping “invert” the channel and approach a somewhat flat frequency response. As shown in Fig. 11.7.2 for one feedforward path, the resonance frequency of G_{m1} (ω_{p1}) can be placed well above the pole of the main path (ω_{p1}), extending the bandwidth of the overall equalizer. As a result, the vertical and horizontal eye openings at the DFE summing node increase by 110mV (55%) and 9ps, respectively. The CTLE provides a boost factor of 12dB at 28GHz while drawing 9mW.

With a CTLE boost of 12dB, the 56Gb/s data arriving at the DFE input still suffers from complete eye closure. The DFE thus bears a heavy burden of equalization. Figure 11.7.3 depicts the dual-loop DFE, highlighting the proposed high-pass feedforward and feedback paths, which are modeled by $G(s) = \alpha s$ and $H(s) = \beta s$, respectively. Since $D_{sum}(n) = (1 + \alpha s) D_{in}(n) - (k_1 + \beta s) D_{out}(n-1)$, the new paths add $\alpha dD_{in}/dt$ to the input and $\beta dD_{out}/dt$ to the feedback signal. Both derivatives sharpen the transitions in D_{sum} , obviating the need for inductive peaking at the DFE summing node.

The two high-pass paths in Fig. 11.7.3 can be naturally created by utilizing the inductors within the CTLE. Shown in the same figure is the actual DFE implementation, including the first tap. The sections in black form a half-rate/quarter-rate topology [4], where the CTLE output is DMUXed to generate D_{odd} and D_{even} and the results are applied to the top and bottom paths. The four latches, L_1 - L_4 , sense and demultiplex the DFE summing nodes and are driven by the quarter-rate clock. These latch outputs are then multiplexed and returned to the other path, completing the DFE loops. The first tap coefficient is adjusted by the value of k . We remark that the 56Gb/s NRZ data arriving from the CTLE is demultiplexed by DMUX₀ and then by DMUX₁ and DMUX₂. The top and bottom DFE loops thus carry a rate of 28Gb/s at their summing nodes.

The high-pass action characterized by $G(s)$ and $H(s)$ is implemented by the blocks in red and the CTLE inductor, L_{D2} . We note that node P carries the derivative of the CTLE output and applies it to the DFE summing nodes through the α stages. Also, the DFE 14Gb/s outputs are multiplexed and return to this node with a weighting factor of $-\beta/\alpha$. As a result, a signal proportional to $\alpha L_{D2} sD_{in}(n) - \beta L_{D2} sD_{out}(n-1)$ is injected into the summing nodes. The second tap of the DFE employs similar concepts. The proposed DFE techniques increase the vertical and horizontal eye openings at the summing nodes by 60mV (30%) and 4ps, respectively. Overall, the CTLE and DFE feedforward and feedback techniques increase the eye openings by 140mV (70%) and 10ps. Employing charge-steering techniques [4], the DFE and DMUXes consume a total of 5.8mW.

The half-rate CDR circuit in Fig. 11.7.1 presents three challenges. First, if tied to the main CTLE output, the CDR input capacitance - that of three flipflops - severely limits the CTLE bandwidth. This is resolved by feeding the CDR from node Q of the CTLE. The proposed architecture is shown in Fig. 11.7.4, where L_{D1} and C_{in} resonate at about 26GHz. The band-pass nature of the data at Q leads to interesting CDR characteristics. It can be shown that the CDR locks for channel losses greater than approximately 6dB at Nyquist. For lower losses, the CDR can sense the CTLE input instead. Also, in the presence of long runs, the phase detector (PD) flipflops become metastable; fortunately, the two

XORs in Fig. 11.7.4 produce a zero difference for G_{m1} in such a case, as expected of the Alexander PD.

The second challenge is that the half-rate Alexander PD in Fig. 11.7.4 requires quadrature clock phases, CK_1 and CK_0 . In this work, CK_1 is simply delayed by an inverter to generate CK_0 . The key point here is that, if CK_1 samples the data transitions, then the inverter delay variation only moves CK_0 in the middle of the eye. According to simulations, the PD gain changes negligibly if the delay varies from 6.5 to 10.5ps. The data samples taken by CK_0 can afford to be erroneous occasionally as the CDR is only responsible for recovering the clock and not the data.

The third challenge relates to the small data swings delivered by the CTLE (around 20mV_{pp} for a 1010 sequence) to the CDR, which would translate to a low phase detector gain and a narrow loop bandwidth or simply failure to lock. To resolve this issue, the PD flipflops must achieve both a high sensitivity and a wide bandwidth, would require inductive peaking in the PD’s six latches, L_1 - L_6 , and impose exceedingly difficult signal distribution in the receiver. We instead insert simple differential pairs A_1 - A_6 in Fig. 11.7.4 so as to obtain acceptably large voltage swings at the XOR inputs without inductive peaking. Our method raises the height of this eye opening by about 50%. To control the phase difference between the main path and the CDR path, the output buffer of the CDR employs delay control by means of programmable capacitors. The CDR draws 19.3mW and the ± 2 circuit in Fig. 11.7.1 draws 5.1mW.

The complete receiver is fabricated in 28nm CMOS technology. Figure 11.7.7 shows the die photograph with a core area of 250 μ m \times 275 μ m. All measurements are conducted with a PRBS of 2⁻¹ and a channel loss of 25dB at 28GHz unless otherwise stated. The RX has been tested on a probe station in two different conditions: 1) with the CDR disabled and an external clock applied to the DFE so that the phase difference between D_{in} and CK can be adjusted and the bathtub curve can be measured; and 2) with the CDR enabled. The bit-error-rate tester (BERT), Keysight’s M8040A, provides the capability to emulate a 2-tap transmit feedforward equalizer (FFE) in the data applied to the channel. Figure 11.7.5 plots the measured bathtub curves of the receiver for two cases: 1) for channel A, which has a loss of 25dB at 28GHz, (with no FFE); and 2) for channel B, which has a loss of 30dB at 28GHz, while the BERT is set to an FFE function of the form $-0.2 + 0.8z^{-1}$ to emulate a typical transmitter. The horizontal eye openings are 0.4UI and 0.3UI, respectively.

The remaining measurements include the CDR circuit as well. The measured BER is less than 10^{-12} . Figure 11.7.5 plots the phase noise of the recovered half-rate for offsets up to 100MHz (due to equipment limitation), at which it is -124 dBc/Hz. For greater offsets, the phase noise is directly read from the spectrum, yielding -126 dBc/Hz at 200MHz and -127 dBc/Hz beyond. The jitter integrated from 100Hz to 14GHz is equal to 500fs. The jitter transfer and tolerance have also been measured, revealing a -3 dB loop bandwidth of 55MHz and tolerance of 1UI_{pp} at approximately 9MHz.

Figure 11.7.6 summarizes this receiver performance and compares it with that of the prior art. Note that some receivers do not include clock generation.

Acknowledgements:

Work supported by Texas Instruments, Realtek Semiconductor, and TSMC University Shuttle Program.

References:

- [1] J. Im et al., “A 112Gb/s PAM-4 Long-Reach Wireline Transceiver Using a 36-Way Time-Interleaved SAR-ADC and Inverter-Based RX Analog Front-End in 7nm FinFET,” *ISSCC Dig. Tech. papers*, pp. 116-118, Feb. 2020.
- [2] T. Ali et al., “A 460mW 112Gbps DSP-Based Transceiver with 38dB Loss Compensation for Next Generation Data Centers in 7nm FinFET technology,” *ISSCC Dig. Tech. papers*, pp. 118-120, Feb. 2020.
- [3] P. J. Peng, J. F. Li, L. Y. Chen and J. Lee, “A 56Gb/s PAM-4/NRZ transceiver in 40nm CMOS,” *ISSCC Dig. Tech. Papers*, pp. 110-111, Feb. 2017.
- [4] A. Manian and B. Razavi, “A 40-Gb/s 14-mW CMOS Wireline Receiver,” *IEEE J. Solid-State Circuits*, vol. 52, pp. 2407-2421, Sept. 2017.
- [5] E. Depaoli et al., “A 64 Gb/s Low-Power Transceiver for Short-Reach PAM-4 Electrical Links in 28-nm FDSOI CMOS,” *IEEE J. Solid-State Circuits*, vol. 54, pp. 6-17, Jan. 2019.
- [6] A. Cevrero et al., “A 100Gb/s 1.1pJ/b PAM-4 RX with Dual-Mode 1-Tap PAM-4 / 3-Tap NRZ Speculative DFE in 14nm CMOS FinFET,” *ISSCC Dig. Tech. papers*, pp. 112-114, Feb. 2019.

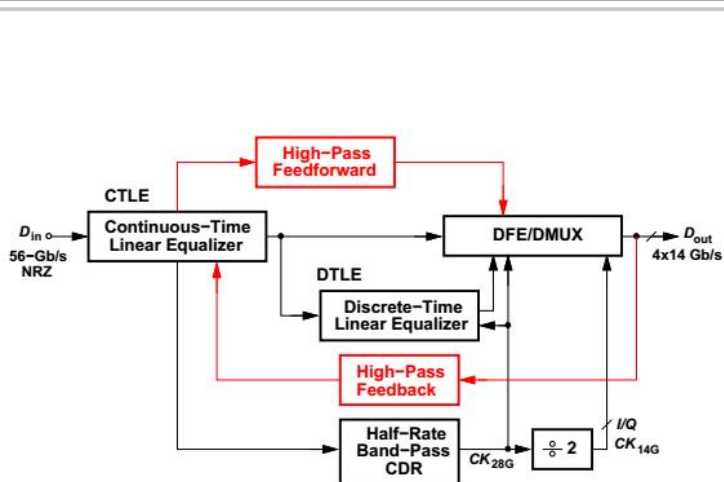


Figure 11.7.1: Conceptual receiver architecture.

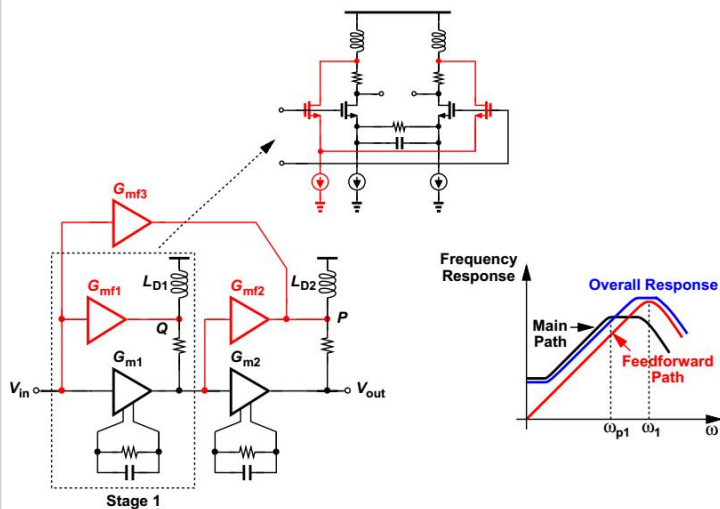


Figure 11.7.2: Proposed CTLE architecture.

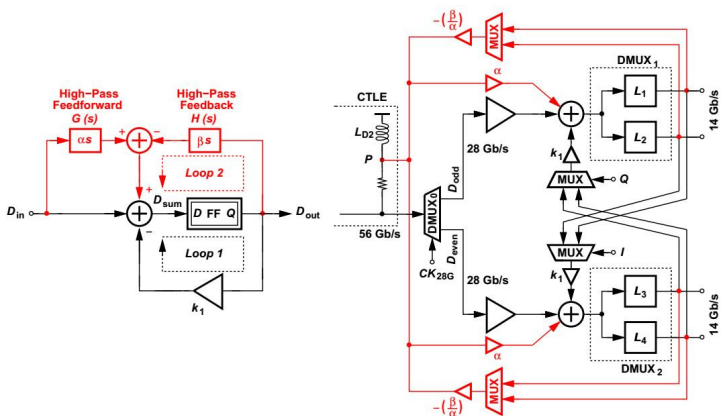


Figure 11.7.3: Proposed dual-loop DFE architecture.

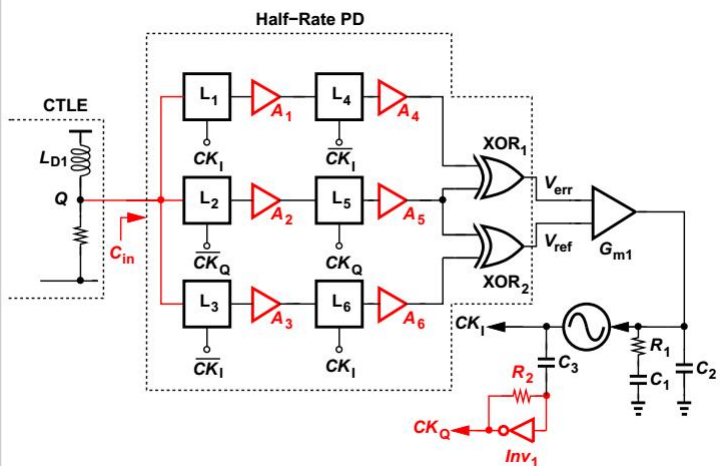


Figure 11.7.4: Proposed CDR architecture.




Figure 11.7.5: Measured bathtub curve and phase noise of the recovered clock at frequency offsets of 100Hz and 100MHz.

Reference	Shibasaki [ISSCC 2016]	Peng [ISSCC 2017]	Han [JSSC Dec. 2017]	Depaoli [JSSC Jan. 2019]	Cevrero [ISSCC 2019]	Ashkan [JSSC Mar. 2019]	This Work
Modulation	NRZ	PAM4	NRZ	PAM4	NRZ	PAM4	NRZ
Data Rate (Gb/s)	56	56	60	64	56	56	56
Architecture	CTLE 1-tap DFE	CTLE 3-tap DFE	2-tap RX FFE CTLE 3-tap DFE	CTLE	CTLE 3-tap DFE	CTLE 1-tap FIR DFE 1-tap IIR DFE	CTLE with High-Pass FF Path DTLE Dual-loop DFE 2 conv. & 2 HP taps
Channel Loss	18.4 dB* @ 28 GHz	24 dB** @ 14 GHz	21 dB** @ 30 GHz	16.8 dB*** @ 16 GHz	37.8 dB* @ 28 GHz	20.8 dB* @ 28 GHz	30 dB* @ 28 GHz 16.5 dB* @ 14 GHz 25 dB @ 28 GHz 13.5 dB @ 14 GHz
Horizontal Eye (UI)	0.28 @ 10 ⁻⁹	0.25 @ 10 ⁻¹²	0.3 @ 10 ⁻¹²	0.19 @ 10 ⁻⁶	0.44 @ 10 ⁻¹²	0.19 @ 10 ⁻¹²	0.4 @ 10 ⁻¹²
Clock Jitter (fs, rms)	-	688 (100 Hz -1 GHz)	-	-	-	-	500 (100 Hz -14 GHz)
PRBS	15	7	7	Q13	15	15	7
Power (mW)	Incl. [§] Excl. ^{§§}	141.7 382	136	- 180	- 112	259	49.56 43.6
Power Eff. (pJ/bit)	Incl. [§] Excl. ^{§§}	2.53 6.82	2.26	- 2.81	- 2	4.63	0.88 0.77
Area (mm ²)	-	1.4 [¶] 1.26	2.03	0.32	0.053	0.51	0.102
Technology	28-nm CMOS	40-nm CMOS	65-nm CMOS	28-nm FDSOI	14-nm FINFET	65-nm CMOS	28-nm CMOS

* Includes 2-tap TX FFE
 ** Includes 3-tap TX FFE
 *** Includes 4-tap TX FFE
 § Includes TX area
 § Includes Clock Gen.
 §§ Excludes Clock Gen.

Figure 11.7.6: Performance summary and comparison to prior art.

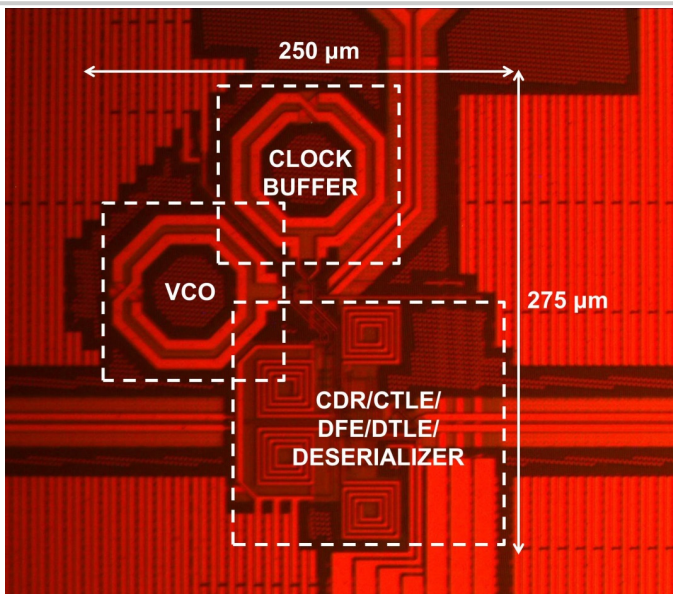


Figure 11.7.7: Receiver die photograph.