



Behzad Razavi

The Design of a Clock and Data Recovery Circuit

High-speed serial links transport data among chips, boards, and servers in data centers. Many such media do not transmit a clock to the destination, requiring that the receiver (RX) perform clock and data recovery (CDR). The CDR can be realized entirely in the analog domain or partially in the digital domain. In this article, we design the former type.

We target the following performance specifications:

- *Data rate, r_b* : 56 Gb/s.
- *Data format*: Non-return-to-zero.
- *Power consumption*: <10 mW.
- *Single-ended input data swing*: $0.3 V_{pp}$.

We simulate the circuit in the slow-slow corner of 28-nm technology at 75 °C with a supply of 0.95 V.

General Considerations

Consider the RX depicted in Figure 1(a). Here a continuous-time linear equalizer (CTLE) partially compensates for the loss of the channel, and a decision-feedback equalizer (DFE) further improves the quality of the data. The

CDR circuit need only recover the clock here; i.e., it acts only as a clock recovery mechanism.

A clock recovery circuit must generate an output satisfying three conditions: 1) it is periodic, 2) it has a frequency equal to the received data rate (or a submultiple thereof), and 3) it bears a proper phase relation with respect to the data so that it can sample the middle of the eye. The first property points to the need for an oscillator, and the last two necessitate phase locking. We thus arrive at the conceptual realization shown in Figure 1(b), where a phase detector (PD) computes the phase error between D_{in} and CK_{out} , delivering the result through a low-pass filter (LPF) to a voltage-controlled oscillator (VCO). After the loop locks, the rising edges of CK_{out} are aligned with the input data transitions, and the falling edges sample the middle of the eye (or the other way around).

If the channel in Figure 1(a) is less hostile to the data, one can omit the DFE and rely on a CDR circuit for recovering the data as well (Figure 2). In such an environment, we prefer, for the sake of simplicity, that the PD in Figure 1(b) produce the retimed data, D_{out} .

We should make two additional remarks here. First, the PD must provide meaningful information while sensing *random* (rather than periodic) data at one of its inputs. Since PDs used in frequency-multiplying phase-locked loops fail here, we must seek other topologies.

Second, we should examine the behavior of the loop in Figure 1(b) when D_{in} does not exhibit transitions for some time (Figure 3). How does the PD respond in such a case? Noting that no phase error information is available during this “long run,” we ask whether the voltage produced by the LPF remains constant or drifts, thereby generating jitter at the VCO output. These aspects of PDs play a central role in the CDR performance.

Early-Late Phase Detection

We select the Alexander PD [1] for this work. This structure is based on two observations. First, if the clock,




FIGURE 2: An RX for low-loss channels.




FIGURE 1: (a) An RX example and (b) a generic CDR architecture. CTLE: continuous-time linear equalizer; DFE: decision-feedback equalizer; LPF: low-pass filter; PD: phase detector; VCO: voltage-controlled oscillator.




FIGURE 3: The effect of a long run on the control voltage.




FIGURE 4: (a) The case of a late clock, (b) the case of an early clock, and (c) an ideal PD characteristic.




FIGURE 5: The Alexander PD.




FIGURE 6: Three samples taken during a long run.

CK, is “late,” i.e., if its rising edge lags data transitions [Figure 4(a)], then the three consecutive data samples taken by the rising and falling edges of CK satisfy this condition: $S_1 = S_2 \neq S_3$. Second, if the clock is “early,” then $S_1 \neq S_2 = S_3$ [Figure 4(b)]. That is, these two distinct conditions reveal whether the phase difference between D_{in} and CK is positive or negative.

At this conceptual level, we recognize that the Alexander PD yields information only in regards to the *polarity* of the phase error, $\Delta\phi$, but not its magnitude. We then predict the characteristic shown in Figure 4(c) even though it is unclear at this point where the “PD output” is located.

Two more remarks are warranted here. First, the decision $S_1 = S_2 \neq S_3$ or $S_1 \neq S_2 = S_3$ must somehow be directed to the VCO so as to adjust

its frequency and phase. Second, S_1 , S_2 , and S_3 occur at different points in time and cannot be compared. That is, we must first store and align these samples.

Alexander PD

The principles described in the “Early-Late Phase Detection” section indicate that an early-late PD requires the following functions: 1) a means of sampling the data to obtain S_1 , S_2 , and S_3 [i.e., flipflops (FFs)]; 2) a means of aligning S_1 , S_2 , and S_3 (i.e., FFs); and 3) a means of comparing S_1 with S_2 and S_2 with S_3 [i.e., exclusive-OR (XOR) gates].

We thus construct the Alexander PD topology illustrated in Figure 5 [1]. We note that FF₁ samples D_{in} on the rising of CK, forming S_1 in Figure 4(a). This sample is delayed by FF₃ for alignment with S_2 . Moreover, FF₂ samples D_{in} on the falling edge of CK, forming S_2 , which is also delayed by FF₄. The third sample in Figure 4(a) is taken by the second rising edge of CK, as performed by FF₁ in Figure 5. Since the outputs of FF₁, FF₃, and FF₄ change at the same time, we consider S_1 , S_2 , and S_3 properly aligned

for comparison. The two XOR gates produce $A = S_1 \oplus S_2$ and $B = S_2 \oplus S_3$, yielding $AB = 01$ if $S_1 = S_2 \neq S_3$ or $AB = 10$ if $S_1 \neq S_2 = S_3$. We have therefore created unique signatures for AB corresponding to positive and negative phase errors. For this PD’s detailed waveforms, the reader is referred to [2].

Recall from the general considerations described previously that a PD’s behavior in the absence of data transitions is of great interest. In the Alexander PD of Figure 5, a long run yields $S_1 = S_2 = S_3$ (Figure 6) and hence $A = B = 0$. We then interpose a stage between the PD and the LPF in Figure 1(b) that does not “respond” to $A = B = 0$, allowing V_{cont} to remain constant. This is a key advantage of the Alexander PD.

Another useful property of this PD is that the outputs of FF₁ or FF₃ in Figure 5 can serve as recovered data (after the CDR loop locks). This satisfies the example depicted in Figure 2.

CDR Architecture

The foregoing thoughts lead to the CDR architecture shown in Figure 7, where a transconductance stage, G_{m1} , senses the voltage difference between A and B and injects a proportional current into the filter consisting of R_1 and C_1 . The loop adjusts the VCO’s output phase until the desired control voltage and hence output frequency are obtained.

In this article, we pursue a “full-rate” CDR architecture, one employing a VCO whose frequency is equal to the data rate, r_b . A “half-rate” design, on the other hand, would operate the VCO at half of r_b , simplifying clock distribution but demanding greater CDR complexity.

The CDR circuit of Figure 7 merits several remarks. First, in the absence of data transitions, A and B are equal, commanding G_{m1} to deliver a zero current to the filter. That is, V_{cont} , in principle, remains constant. Second, G_{m1} must respond to only the *average* difference between the voltages at A and B. These nodes therefore need not provide a wide bandwidth (BW) for the

XOR output current pulses, a great design simplification.

Third, the XOR gates in Figure 7 must present equal propagation delays in their two input paths to avoid phase offset in the locked condition. We thus prefer a symmetric XOR topology.

Choice of CDR BW

The CDR BW is governed by four factors. First, a wide BW is desirable so as to suppress the VCO phase noise to a greater extent. Second, a wider BW grants the CDR greater “agility,” allowing the recovered clock to track more closely the jitter present in the input data. As illustrated in Figure 8(a), this means that the clock continues to sample the middle of the data eye even though the data period, T_b , fluctuates.

Of course, fast jitter components, i.e., rapid changes in the bit period, do not travel to the CDR output, displacing the middle of the eye from its ideal point [Figure 8(b)].

The foregoing intuition can be formulated as follows. Consider the data recovery arrangement shown in Figure 9(a), where the rising edge of CK incurs an undesirable phase error, ϕ_e , with respect to the middle of the data eye. Constructing the phase-domain model illustrated in Figure 9(b), we note that $\phi_e = \phi_{in} - \phi_{CK}$ is ideally zero. Of particular interest to us is the case where ϕ_{in} fluctuates, slowly or rapidly, with time. For a type-I architecture, we have

$$\phi_{CK} = \frac{\omega_n^2}{s^2 + 2\zeta\omega_n s + \omega_n^2} \phi_{in} \quad (1)$$

where $\zeta = \sqrt{\omega_{LPF}/(K_{PD}K_{VCO})}$ and $\omega_n = \sqrt{K_{PD}K_{VCO}\omega_{LPF}}$. It follows that

$$\phi_e = \phi_{in} - \phi_{CK} \quad (2)$$

$$= \frac{s^2 + 2\zeta\omega_n s}{s^2 + 2\zeta\omega_n s + \omega_n^2} \phi_{in}. \quad (3)$$

That is, the jitter experiences a high-pass response as it reaches the rising clock edge in Figure 9(b) and (c). For this reason, we wish to maximize the CDR loop BW.




FIGURE 7: The CDR architecture.




FIGURE 8: The effects of (a) slow and (b) fast jitter components on the clock edge placement.




FIGURE 9: (a) An illustration of data retiming by an FF, (b) its phase-domain model, and (c) its magnitude response.




FIGURE 10: The CML latch used in the PD.




FIGURE 11: A symmetric XOR topology.




FIGURE 12: The PD characteristic.

The data tracking ability of the recovered clock also translates to higher “jitter tolerance,” a point beyond the scope of this article.

In practice, as we increase the CDR BW, the transfer function, ϕ_{CK}/ϕ_{in} , begins to exhibit substantial peaking, imparting *greater* jitter to CK than that of D_{in} . For this reason, the BW is typically limited to a few tens of megahertz.

PD Design

Our full-rate CDR architecture requires that the PD’s FFs run at 56 GHz. Since rail-to-rail logic in 28-nm technology fails at these rates, we turn to the current-mode logic (CML) latch from a previous article [3], shown in Figure 10.

For the XOR gates, we incorporate the symmetric topology shown in Figure 11(a) [4], where V_b is chosen equal to the common-mode level of the inputs. The output current, I_{out} , is equal to I_1 or I_2 when $A = B$ and equal to zero when $A \neq B$. This structure generates a single-ended output, but the two XOR gates in Figure 7 provide a differential signal to the G_m stage.

The XOR design is governed by two considerations. First, we wish M_1 – M_6 to experience complete switching so as to maximize the difference between the output current for $A = B$ and that for $A \neq B$ (ideally, I_{ss}). This necessitates wide transistors and hence large capacitances presented to the FFs in Figure 7. We can then increase the widths of M_1 – M_6 only to the point where the FFs still operate properly at 56 Gb/s.

Second, to maximize the difference between the high and low levels of V_{out} , we also wish to maximize the value of R_D in Figure 11, while maintaining M_3 and M_6 in saturation so that they can be switched by the other transistors. We thus target an output CM level around 500 mV and replace R_D with a 25- $\mu\text{m}/120\text{-nm}$ diode-connected PMOS device.

The overall PD design must ensure that the input data in Figure 7 properly travels through the FFs. In addition, we must construct the PD’s input-output characteristic and obtain its gain, K_{PD} . Returning to the behavior depicted in Figure 4(c), the reader may wonder whether $K_{PD} \rightarrow \infty$ at $\Delta\phi = 0$. In reality, it does not. Effects such as FF metastability and clock and data jitter lead to a less abrupt transition [5].

Figure 12 plots the simulated PD characteristic, where the average difference between A and B in Figure 5 is measured as the clock phase varies. Note that the average should be computed over thousands




FIGURE 13: A VCO circuit diagram.

of cycles to obtain consistent results with random input data.

The characteristic reveals a number of interesting phenomena. First, the linear regime spans about 50° ($\equiv 2.5$ ps). If jitter in D_{in} and CK creates a phase difference comparable to this value, then the linear range becomes wider.

Second, the characteristic is asymmetric around the zero average output. That is, equal and opposite phase differences between the clock and the data yield different magnitudes for the output average. This behavior is attributed to the unmatched clock and data paths in the latch of Figure 10.

Third, the plot exhibits a narrow “dead zone” in the vicinity 60° , primarily because the FFs in Figure 7 become metastable for a small $\Delta\phi$ and, at 56 GHz, deliver low-voltage swings to the XOR gates. Due to the small loop gain, the recovered clock can accumulate as much as 10° ($\equiv 0.5$ ps) of random displacement with respect to the middle of the data eye.

VCO Design

The 56-GHz VCO is configured as an oscillator core followed by a buffer (Figure 13), the latter isolating the former from the kickback noise of the PD FFs. The CM level is set to about $V_{DD}/2$. Producing rail-to-rail voltage swings, the circuit draws 3 mA from the supply.

Simulation Results

We select $R_l = 1$ k Ω , $C_l = 4$ pF, and $G_{m1} = 0.6$ mS in Figure 7. Figure 14 plots the VCO control voltage as the CDR settles. The circuit locks in about 100 ns. This voltage exhibits a ripple of about 70 mV_{pp}, but a great deal of this perturbation arises from the coupling of the voltage swings at X and Y in Figure 13 through the varactors. This effect is benign as it modulates the VCO at twice its frequency.

Figure 15 displays the recovered clock eye diagram, revealing a peak-to-peak jitter of about 400 fs. Similarly, Figure 16 presents the data eye diagram at the output of FF_3 in Figure 7. The jitter is about 520 fs_{pp}.

The “jitter generation” of the CDR is also of interest and is defined as the clock jitter when the circuit senses a periodic input at 28 GHz. In this test, we run a transient noise

simulation while allowing electronic noise components up to 200 GHz. Figure 17 shows the result, suggesting a peak-to-peak jitter of 150 fs for an observation period of 100 ns.




FIGURE 14: CDR control voltage settling behavior.




FIGURE 15: A recovered clock eye diagram.




FIGURE 16: A data eye diagram.

(continued on p. 116)

Top Accessed Articles in January 2026



- *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*
 - An Updated Assessment of Reinforcement Learning for Macro Placement
 - MNSIM 2.0: A Behavior-Level Modeling Tool for Processing-In-Memory Architectures
 - RTL Coder: Fully Open-Source and Efficient LLM-Assisted RTL Code Generation Technique



- *IEEE Design & Test*
 - A Survey of Silicon Photonics for Energy-Efficient Manycore Computing
 - STAR: A Mixed Analog Stochastic In-DRAM Convolutional Neural Network Accelerator
 - Ethical Design of Computers: From Semiconductors to IoT and Artificial Intelligence



- *IEEE Embedded Systems Letters*
 - CHISEL: Compression-Aware High-Accuracy Embedded Indoor Localization With Deep Learning
 - SANGRIA: Stacked Autoencoder Neural Networks With Gradient Boosting for Indoor Localization
 - On-Device Personalization for Human Activity Recognition on STM32



Contact CEDA Vice President of Publications Ian O'Connor at ian.oconnor@ec-lyon.fr or the journal's respective editor-in-chief for questions regarding CEDA-sponsored publications.

Share What's New

We thank you for forwarding this content to others who might be interested including other members of your group or organization.



THE ANALOG MIND (continued from p. 15)




FIGURE 17: A recovered clock including CDR electronic noise sources.

References

- [1] J. D. H. Alexander, "Clock recovery from random binary data," *Electron. Lett.*, vol. 11, no. 22, pp. 541–542, 1975, doi: [10.1049/el:19750415](https://doi.org/10.1049/el:19750415).
- [2] B. Razavi, *Design of CMOS Phase-Locked Loops*. Cambridge, U.K.: Cambridge Univ. Press, 2020.
- [3] B. Razavi, "The design of an equalizer - Part two [The Analog Mind]," *IEEE Solid-State Circuits Mag.*, vol. 14, no. 1, pp. 7–12, Winter 2022, doi: [10.1109/MSSC.2021.3126997](https://doi.org/10.1109/MSSC.2021.3126997).
- [4] B. Razavi, Y. Ota, and R. G. Swartz, "Design techniques for low-voltage high-speed digital bipolar circuits," *IEEE J. Solid-State Circuits*, vol. 29, no. 3, pp. 332–339, Mar 1994, doi: [10.1109/4.278358](https://doi.org/10.1109/4.278358).
- [5] J. Lee and B. Razavi, "Analysis and modeling of bang-bang clock and data recovery circuits," *IEEE J. Solid-State Circuits*, vol. 39, no. 9, pp. 1571–1580, Sep. 2004, doi: [10.1109/JSSC.2004.831600](https://doi.org/10.1109/JSSC.2004.831600).

SSC