

Design of Half-Rate Clock and Data Recovery Circuits for Optical Communication Systems

Jafar Savoj^{*}
Electrical Engineering Department
University of California
Los Angeles, CA 90095
jafar@icsl.ucla.edu

Behzad Razavi
Electrical Engineering Department
University of California
Los Angeles, CA 90095
razavi@icsl.ucla.edu

ABSTRACT

This paper describes the design of two half-rate clock and data recovery circuits for optical receivers. Targeting the data rate of 10-Gb/s, the first implementation incorporates a ring oscillator and a linear phase detector whereas the second implementation uses a multiphase LC oscillator and a bang-bang phase/frequency detector. Fabricated in 0.18- μm CMOS technology, the power consumption of each of the circuits is less than 100 mW. The rms jitter of the output clock for the two prototypes is 1 ps and 0.8 ps, respectively, while the latter achieves a capture range of more than 14%.

1. INTRODUCTION

The number of the Internet nodes doubles approximately every 100 days, leading to an average bit rate of a few terabits per second on the backbone. The bandwidth requirements are growing at an extremely fast pace. Applications such as online virtual reality will require data rates that are 10,000 times higher than currently available ones. With fiber optics being the only communication medium capable of handling such high data rates, this trend has suddenly created a widespread demand for high-speed optical and electronic devices, circuits, and systems. The new optical revolution has replaced modular, general-purpose building blocks by end-to-end solutions. Greater levels of integration on a single chip enable higher performance and lower cost. Mainstream VLSI technologies such as CMOS continue to take over the territories previously claimed by GaAs or InP devices.

Clock and data recovery (CDR) circuits operating in the 10-Gb/s range have become attractive for the optical fiber backbone. While CDR circuits operating at 10 Gb/s and above have been designed in bipolar technologies [1],[2], cost

^{*}Now with Transpectrum Technologies, Inc., Los Angeles, CA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

DAC 2001, June 18-22, 2001, Las Vegas, Nevada, USA.
Copyright 2001 ACM 1-58113-297-2/01/0006 ...\$5.00.




Figure 1: Generic CDR architecture.

and integration issues make it desirable to implement these circuits in standard CMOS processes.

This paper describes the design and experimental results of two 10-Gb/s CDR circuits that are realized in 0.18- μm CMOS technology. The speed limitations of the technology are overcome by the CDR architectures. The first circuit benefits from a new linear phase detector (PD) that compares the phase of the incoming data with that of a half-rate clock. The CDR circuit also incorporates a three-stage interpolating ring oscillator to achieve a wide tuning range. Fabricated in a 0.18- μm CMOS technology, the circuit achieves an rms jitter of 1 ps, and a peak-to-peak jitter of 14.5 ps with a pseudorandom sequence (PRBS) of $2^{23} - 1$ while dissipating 72 mW from a 2.5-V supply. The second circuit incorporates a multiphase LC oscillator and a half-rate phase/frequency detector with automatic data retiming. Fabricated in 0.18- μm CMOS technology, the circuit exhibits a capture range of 1.43 GHz, an rms jitter of 0.8 ps, and a peak-to-peak jitter of 9.9 ps with a PRBS of length $2^{23} - 1$. The power dissipation is 91 mW from a 1.8-V supply.

2. A HALF-RATE LINEAR CDR CIRCUIT

2.1 Architecture

The choice of the CDR architecture is primarily determined by the speed and supply voltage limitations of the technology as well as the power dissipation and jitter requirements of the system.

In a generic CDR circuit, shown in Fig. 1, the phase detector compares the phase of the incoming data to the phase of the clock generated by the voltage-controlled oscillator (VCO), producing an error proportional to the phase difference between its two inputs. The error is then applied to a charge pump and a low-pass filter so as to generate the oscillator control voltage. The clock signal also drives a decision circuit, thereby retiming the data and reducing its jitter.

If attempted in a $0.18\text{-}\mu\text{m}$ CMOS technology, the architecture of Fig. 1 poses severe difficulties for 10-Gb/s operation. Although exploiting aggressive device scaling, the CMOS process used in this work provides marginal performance for such speeds. For example, even simple digital latches or three-stage ring oscillators fail to operate reliably at these rates. These issues make it desirable to employ a “half-rate” CDR architecture, where the VCO runs at a frequency equal to half of the input data rate.

Another critical issue in the architecture of Fig. 1 relates to the inherently unequal propagation delays for the two inputs of the phase detector: Most phase detectors that operate properly with random data (e.g. a D flipflop) are asymmetric with respect to the data and clock inputs, thereby introducing a systematic skew between the two in phase-lock condition. Since it is difficult to replicate this skew in the decision circuit, the generic CDR architecture suffers from a limited phase margin - unless the raw speed of the technology is much higher than the data rate.

The problem of the skew demands that phase detection and data regeneration occur in the same circuit such that the clock still samples the data at the midpoint of each bit even in the presence of a finite skew. For example the Hogge PD [3] automatically sets the clock phase to the optimum point in the data eye (but it fails to operate properly with a half-rate clock).

The above considerations lead to the CDR architecture of Fig. 2. Here, a half-rate phase detector produces an error proportional to the phase difference between the 10-Gb/s




Figure 2: Half-rate CDR architecture.

data stream and the 5-GHz output of the VCO. Furthermore the PD automatically retimes and demultiplexes the data, generating two 5-Gb/s sequences D_{5GA} and D_{5GB} . Although the focus of this work is point-to-point communications, a full-rate retimed output, D_{10G} , is also generated to produce flexibility in testing and exercise the ultimate speed of the technology. The VCO has both fine and coarse control lines, the latter provides for inclusion of a frequency-locked loop in future implementations.

In this work, a new approach to performing linear phase detection using a half-rate clock is described. Owing to its simplicity, this technique achieves both a high speed and low power dissipation while minimizing the ripple on the oscillator control voltage.

2.2 Building Blocks

2.2.1 VCO

The design of the VCO directly impacts the jitter performance and the reproducibility of the CDR circuit. While LC topologies achieve a potentially lower jitter, their limited tuning range makes it difficult to obtain a target frequency




Figure 3: (a) Three-stage ring oscillator, (b) implementation of each stage, (c) transistor-level schematic.

without design and fabrication iterations. Since the circuit reported here was our first design in $0.18\text{-}\mu\text{m}$ technology, a ring oscillator was chosen so as to provide a tuning range wide enough to encompass process and temperature variations.

A three-stage differential ring oscillator [Fig. 3(a)] driving a buffer operates no faster than 7 GHz in $0.18\text{-}\mu\text{m}$ CMOS technology. The half-rate CDR architecture overcomes this limitation, requiring a frequency of only 5 GHz.

As shown in Fig. 3(b), each stage consists of a fast and a slow path whose outputs are summed together. By steering the current between the fast and the slow paths, the amount of delay achieved through each stage and hence the VCO frequency can be adjusted. All three stages in the ring are loaded by identical buffers to achieve equal rise and fall times and hence improve the jitter performance. Figure 3(c) shows the transistor implementation of each delay stage. The fast and slow paths are formed as differential circuits sharing their output nodes. The tuning is achieved by reducing the tail current of one and increasing that of the other differentially. Since the low supply voltage makes it difficult to stack differential pairs under M_1 - M_2 and M_3 - M_4 , the current variation is performed through mirror arrangements driven by PMOS differential pairs.

A critical drawback of supply scaling in deep submicron technologies is the inevitable increase in the VCO gain for a given tuning range. To alleviate this difficulty, the control of the VCO is split between a coarse input and a fine input. The partitioning of the control allows more than one order of magnitude reduction in the VCO sensitivity. The idea is that the fine control is established by the phase detector and the coarse control is a provision for adding a frequency detection loop.

2.2.2 Phase Detector

For linear phase comparison between data and a half-rate clock, each transition of the data must produce an “error” pulse whose width is equal to the phase difference. Further-

more, to avoid a dead zone in the characteristics, a “reference” pulse must be generated whose area is subtracted from that of the error pulse, thus creating a net value that falls to zero in lock.

The above observations lead to the PD topology shown in Fig. 4(a). The circuit consists of four latches and two




Figure 4: (a) Phase detector, (b) operation of the circuit.

XOR gates. The data is applied to the inputs of two sets of cascaded latches, each cascade constituting a flipflop that retimes the data. Since the flipflops are driven by a half-rate clock, the two output sequences V_{out1} and V_{out2} are the demultiplexed waveforms of the original input sequence if the clock samples the data in the middle of the bit period.

The operation of the PD can be described using the waveforms depicted in Fig. 4(b). The basic unit employed in the circuit is a latch whose output carries information about the zero crossings of both the data and the clock signal. The output of each latch tracks its input for half a clock period and holds the value for the other half, yielding the waveforms shown in Fig. 4(b) for points X_1 and X_2 . The two waveforms differ because their corresponding latches operate on opposite clock edges. Produced as $X_1 \oplus X_2$, the *Error* signal is equal to ZERO for the portion of time that identical bits of X_1 and X_2 overlap and equal to the XOR of two consecutive bits for the rest. In other words, *Error* is equal to ONE only if a data transition has occurred.

It may seem that the *Error* signal uniquely represents the phase difference, but that would be true only if the data were periodic. The random nature of the data and the periodic behavior of the clock in fact make the average value of *Error* pattern dependent. For this reason, a reference signal must also be generated whose average conveys this dependence.

The two waveforms Y_1 and Y_2 contain the samples of the data at the rising and falling edges of the clock. Thus, $Y_1 \oplus Y_2$ contains pulses as wide as half the clock period for every data transition, serving as the reference signal.

While the two XOR operations provide both the *Error* and the *Reference* pulses for every data transition, the pulses in *Error* are only half as wide as those in *Reference*. This means that the amplitude of *Error* must be scaled up by a factor of two with respect to *Reference* so that the difference between their averages drops to zero when clock transitions are in the middle of the data eye. The phase error with respect to this point is then linearly proportional to the difference between the two averages.

In order to generate a full-rate output, the demultiplexed sequences are combined by a multiplexer that operates on the half-rate clock as well. This output can also be used for testing purposes in order to obtain the overall bit error rate (BER) of the receiver.

It is important to note that the XOR gates in Fig. 4 must be symmetric with respect to their two differential inputs. Otherwise, differences in propagation delays result in systematic phase offsets. Each of the XOR gates is implemented as shown in Fig. 5 [4]. The circuit avoids stacking




Figure 5: Symmetric XOR gate.

stages while providing perfect symmetry between the two inputs. The output is single-ended but the single-ended *Error* and *Reference* signals produced by the two XOR gates in the phase detector are sensed with respect to each other, thus acting as a differential drive for the charge pump. The operation of the XOR circuit is as follows. If the two logical inputs are not equal, then one of the input transistors on the left and one of the input transistors on the right turn on, thus turning M_{cm} off. If the two inputs are identical, one of the tail currents flows through M_{cm} . Since the average current produced by the *Error* XOR gate is half of that generated by the *Reference* XOR gate, transistor M_{out} is scaled differently, making the average output voltages equal for zero phase difference. Channel length modulation of transistor M_{out} reduces the precision of current scaling between the two XOR gates. This effect can be avoided by increasing the length of the device.

It is instructive to plot the input/output characteristic of the PD to ensure linearity and absence of dead zone. This is accomplished by obtaining the average values of *Error* and *Reference* while the circuit operates at maximum speed. Figure 6 shows the simulated behavior as the phase difference varies from zero to one bit period. The *Reference* average exhibits a notch where the clock samples the metastable points of the data waveform. The *Error* and *Reference* signals cross at a phase difference approximately 55 ps from the metastable point, indicating that the system-




Figure 6: Determination of PD gain.

atic offset between the data and the clock is very small. The linear characteristic of the phase detector results in minimal charge pump activity and small ripple on the control line in the locked condition.

2.3 Experimental Results

The CDR circuit has been fabricated in a 0.18- μm CMOS process. The chip occupies an area of $1.1 \times 0.9 \text{ mm}^2$. The circuit is tested in a chip-on-board assembly. In this prototype, the width of the poly resistors was not sufficient to guarantee the nominal sheet resistance. As a result, the fabricated resistor values deviated from their nominal value by 30%, and the VCO center frequency was proportionally lower than the simulated value at the nominal supply voltage (1.8 V). The supply was increased to 2.5 V, to achieve reliable operation at 10 Gb/s.

Figure 7(a) shows the spectrum of the clock in response to a 10-Gb/s data sequence of length $2^{23} - 1$. The effect of




Figure 7: (a) Spectrum of the recovered clock, (b) recovered clock in the time domain, (c) measured jitter transfer characteristic, (d) recovered demultiplexed data.

the noise shaping of the loop can be observed in this spectrum. The phase noise at 1-MHz offset is approximately

equal to -106 dBc/Hz . Figure 7(b) depicts the recovered clock in the time domain. The time-domain measurements using an oscilloscope overestimate the jitter, requiring specialized equipment, e.g., the Anritsu MP1777 jitter analyzer. The jitter performance of the CDR circuit is characterized by this analyzer. A random sequence of length $2^{23} - 1$ produces 14.5 ps of peak-to-peak and 1 ps of rms jitter on the clock signal. These values are respectively reduced to 4.4 ps and 0.6 ps for a random sequence of length $2^7 - 1$.

The measured jitter transfer characteristics of the CDR is shown in Fig. 7(c). The jitter peaking is 1.48 dB and the 3-dB bandwidth is 15 MHz. The loop bandwidth can be reduced to the SONET specifications, but the jitter analyzer must then generate large jitter and drives the loop out of lock.

Figure 7(d) depicts the retimed demultiplexed data. The difference between the waveforms results from systematic differences between the bond wires and traces on the test board. The circuit also generates a full-rate output. Using this output, the BER of the system can be measured. With a random sequence of $2^7 - 1$, the BER is smaller than 10^{-12} . However, a random sequence of $2^{23} - 1$ results in a BER of 1.28×10^{-6} .

The CDR circuit exhibits a capture range of 6 MHz and a tracking range of 177 MHz. The total power consumed by the circuit excluding the output buffers is 72 mW from a 2.5-V supply. The VCO, the PD, and the clock and data buffers consume 20.7 mW, 33.2 mW and 18.1 mW, respectively.

3. A HALF-RATE BINARY CDR CIRCUIT

3.1 Architecture

In contrast to previous circuit, this design incorporates an LC oscillator to reduce the jitter as well as a phase/frequency detector to achieve a wide capture range. Shown in Fig. 8, the CDR consists of a phase/frequency detector (PFD), a VCO, a charge pump, and a low-pass filter. The PFD com-




Figure 8: CDR architecture.

pares the phase and frequency of the input data to that of a half-rate clock, providing two binary error signals for phase and frequency. These error signals are fed back to the VCO through the charge pump and the low-pass filter. The PFD is designed such that, in addition to providing information about the phase error, it retimes the data as well. Consequently, the CDR exhibits no systematic offset, i.e., inherent skews between clock and data edges due to their unidentical paths through the loop do not degrade the quality of detection. The VCO provides four differential half-quadrature

phases over the full tuning range. All building blocks are fully differential.

3.2 Building Blocks

3.2.1 VCO

Since the half-rate PFD requires clock phases that are integer multiples of 45° , the 5-GHz VCO is designed as a ring structure consisting of four LC-tuned stages [Fig. 9(a)]. If the dc feedback around the ring is positive, all stages




Figure 9: (a) Four-stage LC-tuned ring oscillator, (b) implementation of each stage.

operate in-phase at the resonance frequency defined by the LC tanks. On the other hand, if the dc feedback is negative, the frequency shifts by a small amount so as to allow each stage to contribute 45° of phase.

The proposed oscillator topology has two advantages over resistive-load ring oscillators. First, owing to the phase slope (Q) provided by the resonant loads, it exhibits less phase noise. Second, its frequency of oscillation is only a weak function of the number of stages, generating multiple phases with no speed penalty. By comparison, a four-stage resistive-load ring operates at a lower frequency.

Figure 9(b) shows the implementation of each stage. The loads are formed using on-chip spiral inductors and MOS varactors. Resistor R_1 provides a shift in the output common-mode level. This allows both positive and negative voltages across the varactors and thus maximizes the tuning range. Modeling each tank by a parallel network, we note that the required 45° phase shift slightly “detunes” the circuit. It can be proved that the frequency of oscillation is then given by $\omega_{osc} = (\sqrt{LC})^{-1}(\sqrt{1-1/Q})$, where Q denotes the quality factor of each stage at the frequency of oscillation.

3.2.2 Phase Detector

The phase detector (PD) is derived from the data transition tracking loop (DTTL) described in [5] and [6]. In this PD, in-phase and quadrature phases of a half-rate clock signal sample the data in two double-edge-triggered flipflops. As shown in Fig. 10, four distinct possibilities can be identified for the cases when the clock is early or late, whether the differential data pulse is positive or negative. For a pos-




Figure 10: In-phase and quadrature samples of data pulses with early and late clock signals.

itive pulse, if the clock is early, the quadrature sample is negative and if the clock is late, the quadrature sample is positive. When the data pulse is negative, the polarity of the quadrature samples is reversed. Thus, for a rising data transition, the quadrature sample is routed to the output, and for a falling transition, its complement is.

Figure 11 shows the implementation of the PD. Two latches




Figure 11: Phase detector.

operating on opposite clock phases and a multiplexer form a DETFF that samples the data using both the positive and negative transitions of a half-rate clock. The two signals V_1 and V_2 are therefore the in-phase and quadrature samples of data, respectively, and one is used to route the other or its complement.

The proposed phase detector operates at high speeds because it uses a half-rate clock. Since in the locked condition, the rising and falling edges of the quadrature clock coincide with data transitions, the in-phase clock transitions sample the data at its optimum point with no systematic offset, generating a full-rate output stream. Also, since the phase-error signal is revalidated only at data transitions, it incurs little ripple. Note that the output is independent of the data transition density, resulting in reduction of pattern-dependent jitter.

3.2.3 Frequency Detector

With the very small CDR loop bandwidths specified by optical standards, circuits employing only phase detection suffer from an extremely narrow capture range, e.g., about 1% of the center frequency. For this reason, a means of frequency detection is necessary to guarantee lock to random data.

As with other phase detectors, the half-rate PD of Fig. 11 generates a beat frequency equal to the difference between




Figure 12: (a) Phase and frequency detector, (b) timing diagram in the PFD.

the data rate and twice the VCO frequency. However, it does not provide knowledge of the polarity of this difference. Figure 12(a) depicts the half-rate phase and frequency detector introduced in this work. A second PD is added and driven by phases that are 45° away from those in the first PD. From the waveforms shown in Fig. 12(b), we make the following observations. (1) If the clock is slow, V_{PD1} leads V_{PD2} . Therefore, if V_{PD2} is sampled by the rising and falling edges of V_{PD1} , the results are negative and positive, respectively. (2) If the clock is fast, V_{PD1} lags V_{PD2} . Therefore, if V_{PD2} is sampled by the rising and falling edges of V_{PD1} , the results are the reverse of the previous case. Thus, for a rising transition of V_{PD1} , V_{PD2} is routed to the output, and for a falling transition, its complement is.

3.3 Experimental Results

The CDR circuit has been fabricated in a $0.18\text{-}\mu\text{m}$ CMOS technology. The chip occupies an area of $1.75 \times 1.55 \text{ mm}^2$. The circuit is tested in a chip-on-board assembly while operating with a 1.8-V supply.

Figure 13(a) shows the spectrum of the clock in response to a 9.95328-Gb/s data sequence of length $2^{23} - 1$. The phase noise at 1-MHz offset is approximately equal to -107 dBc/Hz . Figure 13(b) depicts the recovered clock in the time domain. A pseudo-random sequence of length $2^{23} - 1$ produces 9.9 ps of peak-to-peak and 0.8 ps of rms jitter on the clock signal. These values are reduced to 2.4 ps and 0.4 ps for a PRBS of length $2^7 - 1$. The jitter characteristics are measured by the Anritsu MP1777 jitter analyzer.

The measured jitter transfer characteristic of the CDR is shown in Fig. 13(c). The jitter peaking is 0.04 dB and the




Figure 13: (a) Spectrum of the recovered clock, (b) recovered clock in the time domain, (c) measured jitter transfer characteristic, (d) recovered data and clock.

3-dB bandwidth is 5.2 MHz. Despite the small loop bandwidth, the frequency detector provides a capture range of 1.43 GHz, obviating the need for external references. Figure 13(d) depicts the retimed data at 10 Gb/s.

The total power consumed by the circuit excluding the output buffers is 91 mW from a 1.8-V supply.

4. REFERENCES

1. Y. M. Greshishchev and P. Schvan, "SiGe Clock and Data Recovery IC with Linear Type PLL for 10 Gb/s SONET Application," *Proceedings of the 1999 Bipolar/BiCMOS Circuits and Technology Meeting*, pp. 169-172, Sept. 1999.
2. M. Wurzer, et al., "40-Gb/s Integrated Clock and Data Recovery Circuit in a Silicon Bipolar Technology," *Proceedings of the 1998 Bipolar/BiCMOS Circuits and Technology Meeting*, pp. 136-139, Sept. 1998.
3. C. Hogge, "A Self-Correcting Clock Recovery Circuit," *IEEE Journal of Lightwave Technology*, Vol. LT-3, pp. 1312-1314, Dec. 1985.
4. B. Razavi, Y. Ota, R. G. Swarz, "Design Techniques for Low-Voltage High-Speed Digital Bipolar Circuits," *IEEE Journal of Solid-State Circuits*, Vol. 29, pp. 332-339, March 1994.
5. T. O. Anderson, W. J. Hurd, and W. C. Lindsey, "Transition Tracking Bit Synchronization System," U.S. Patent No. 3,626,298, Dec. 1971.
6. A. W. Buchwald, *Design of Integrated Fiber-Optic Receivers Using Heterojunction Bipolar Transistors*, Ph.D. Thesis, University of California, Los Angeles, Jan. 1993.