

A 10-Gb/s CMOS Clock and Data Recovery Circuit

Jafar Savoj and Behzad Razavi

Electrical Engineering Department
University of California, Los Angeles

Abstract

A 10-Gb/s phase-locked clock and data recovery circuit incorporates a 5-GHz interpolating voltage-controlled oscillator and a half-rate phase detector. The phase detector provides a linear characteristic while retiming and demultiplexing the data with no systematic phase offset. Fabricated in a 0.18- μm CMOS technology, the circuit exhibits an rms jitter of 6.6 ps in the recovered clock with random data input of length $2^{23}-1$. The power dissipation is 99 mW from a 2.6-V supply.

Introduction

Clock and data recovery (CDR) circuits operating in the 10-Gb/s range have become attractive for the optical fiber backbone of the Internet. While CDR circuits operating at 10-Gb/s and above have been designed in bipolar technologies [1]-[2], cost and integration issues make it desirable to implement these circuits in standard CMOS processes.

This paper describes the design and experimental results of a 10-Gb/s CDR circuit that is realized in 0.18- μm CMOS technology. The speed limitations of the technology are overcome by the CDR architecture. The circuit produces a 5-GHz clock with an rms jitter of 6.6 ps in response to a random data input of length $2^{23}-1$ while consuming 99 mW from a 2.6-V supply.

Architecture

Although exploiting aggressive device scaling, the 0.18- μm CMOS technology used in this work provides marginal performance for 10-Gb/s operation. Even simple digital circuits such as latches fail to operate if clocked at such frequencies in this process. For this reason, the clock frequency is chosen to be half of the input bit rate. The concept of "half-rate clock" has been used in [2]-[4]. However, [2] and [3] incorporate a bang-bang phase detector (PD), possibly creating large ripple on the control line of the oscillator and hence high jitter. The circuit reported in [4] inherently has a smaller output jitter as a result of using a linear phase detector, but it fails to operate at speeds above 6 Gb/s in 0.18- μm CMOS technology.

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. Furthermore, it automatically retimes and demultiplexes the data.

Shown in Fig. 1, the CDR consists of a phase detector, a

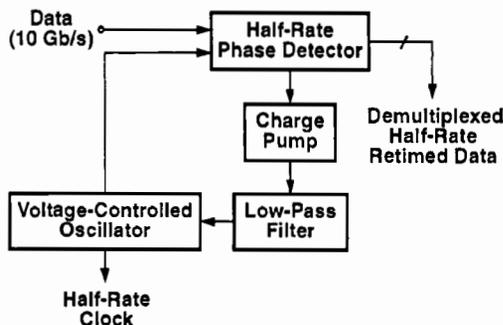


Fig. 1. CDR architecture.

voltage-controlled oscillator (VCO), a charge pump, and a low-pass filter (LPF). The phase detector compares the phase of the input data to that of a half-rate clock, providing an error signal linearly proportional to the phase difference between the clock and data signals. The error signal is fed back to the VCO through the charge pump and the low-pass filter. After phase lock is achieved, the phase of the output clock is within a small offset from the phase of the input data. This guarantees that the clock frequency is exactly equal to one half of the input bit rate. The PD is designed such that in addition to providing information about the phase error, it retimes and demultiplexes the data as well. Consequently, the system has no systematic offset, i.e., inherent misalignment between clock and data edges due to their unidentical paths through the loop does not degrade the quality of detection.

In order to minimize sensitivity to common-mode noise, the CDR circuit incorporates fully differential topologies for all of the building blocks.

Building Blocks

A. VCO

The VCO consists of a three-stage differential ring oscillator with delay interpolation, providing a tuning range wide enough to encompass process and temperature variations [Fig. 2(a)]. Since quadrature clock signals are not required in the CDR circuit, an odd number of stages can be used to achieve reliable operation at 5 GHz.

Figure 2(b) shows the implementation of each delay stage. The fast and slow paths share the load resistors. The delay achieved from this stage varies as the tail currents are partially or fully steered to either the fast or slow paths.

A critical drawback of supply scaling in deep submicron technologies is the increase in the VCO gain for a given tuning range. A higher gain translates the noise on the control line to a higher amount of jitter at the output of the oscillator.

To alleviate this difficulty, the control of the VCO is split between a coarse input and a fine input. 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 in future work. The partitioning of the control allows more than one order of magnitude reduction in the VCO sensitivity.

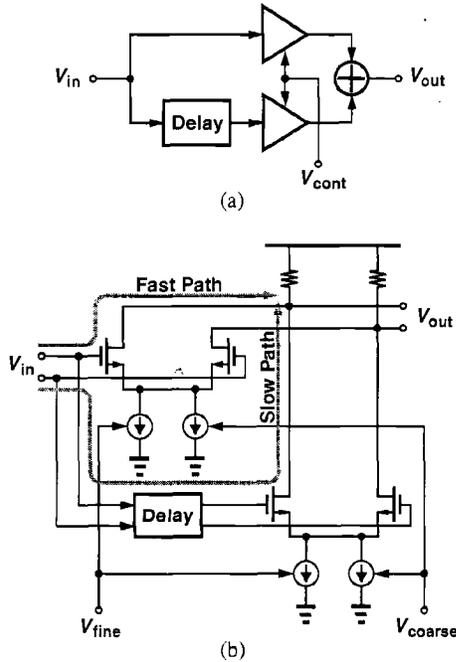


Fig. 2. (a) Delay interpolation in one stage of the VCO, (b) implementation of interpolating delay stage.

B. Phase Detector

Compared to bang-bang phase detectors, linear PDs result in less charge pump activity, smaller ripple on the oscillator control line, and hence lower jitter.

Shown in Fig. 3, the phase detector consists of four latches and two XOR gates. The data is applied to the inputs of two sets of cascaded latches. Each cascade constitutes a flipflop that performs the task of retiming 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.

The operation of the phase detector can be described using the waveforms depicted in Fig. 4. 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 remains constant for the other half, yielding the waveforms shown in Fig. 4 for points A and B. The two waveforms differ

because their corresponding latches operate on opposite clock

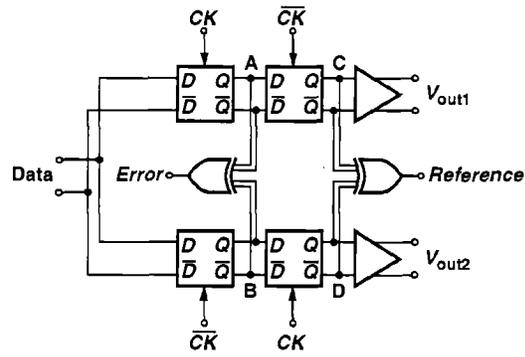


Fig. 3. Phase Detector.

edges. Produced as $A \oplus B$, the *Error* signal consists of ZEROS for the portion of time that identical bits of A and B overlap and the XOR of two consecutive bits for the rest of the time.

If the clock samples the data in the middle of the bit period, the XOR of two consecutive bits in the *Error* signal has a width equal to one-fourth of the clock period. Therefore, during a long span, the average value of the error waveform equals $0.75 \times 0 + 0.25 \times 1 = 0.25$. If the logical ONE is represented by a voltage V_1 , this value is equal to $0.25V_1$. This derivation assumes that the XOR of two consecutive bits has an equal probability of being a ONE or a ZERO. Note that the average value of the error is linearly proportional to the phase difference between the data and the half-rate clock.

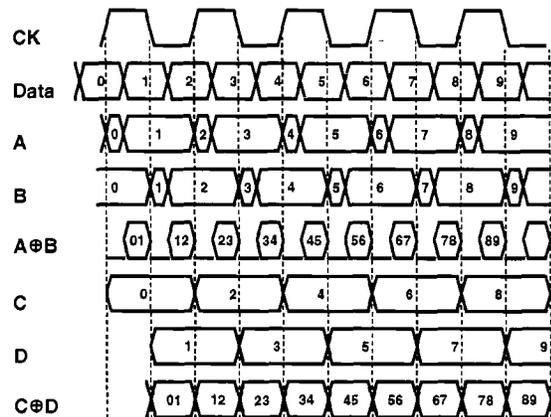


Fig. 4. Operation of the phase detector.

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 the error signal pattern dependent. For this reason, a reference signal must also be generated whose average conveys this dependence.

The two waveforms C and D in Fig. 4 carry the samples of the data at the falling and rising edges of the clock, respec-

tively. That is, the input bits are split into two groups of even and odd bits on these two lines. If the two waveforms are applied to an XOR gate, the resulting signal, denoted by *Reference*, is equal to the XOR of consecutive bits for an interval equal to one half of the clock period. The average value of *Reference* is therefore given by $0.5 \times 0 + 0.5 \times 1 = 0.5$ and, if the logical ONE is represented by V_2 , equal to $0.5V_2$. Thus, in lock condition, the difference between the averages of *Error* and *Reference* goes to zero if V_1 is chosen to be twice V_2 . The phase error with respect to the optimum point is then linearly proportional to the difference between *Error* and *Reference*.

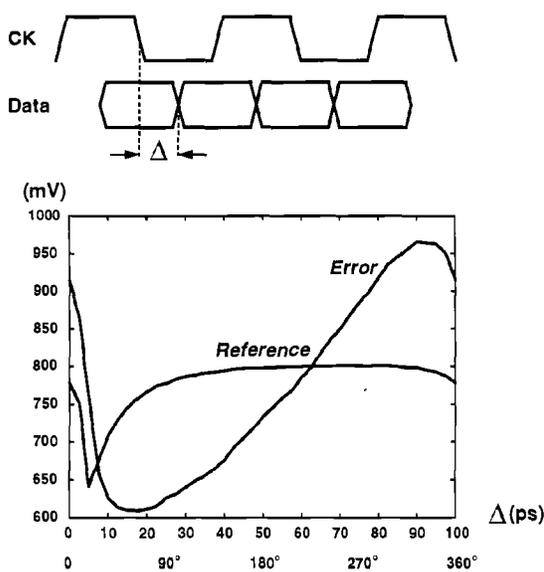


Fig. 5. Determination of PD gain.

In order to obtain the input/output characteristic of the phase detector, the average value of *Error* and *Reference* is determined for a given phase difference between the data edge and the clock edge. Figure 5 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. *Error* and *Reference* cross at a phase difference approximately 55 ps from the metastable point, indicating that the systematic offset between the data and the clock is very small.

The latches used in the phase detector of Fig. 3 operate with 10-Gb/s data and a 5-GHz clock. The limited data and clock swings at these frequencies mandate a current-steering topology for each building block of the PD. Simulations suggest that the operation speed of these circuits is maximized if the device dimensions and bias currents are chosen to maximize the small-signal f_T .

A critical issue in generating the *Error* and *Reference* signals in Fig. 3 is the symmetry of the XOR gates. A Gilbert cell realization suffers from a very large delay difference between

its two inputs, and cross-coupling two Gilbert cells still requires power-hungry level shift circuits.

The XOR implementation used in this work is based on the topology described in [5]. Shown in Fig. 6, the circuit avoids stacking 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 XORs in the phase detector are sensed with respect to each other, thus acting as a differential drive for the charge pump.

Recall that the output of the XOR gate generating the *Error* must be twice that of the XOR gate producing the *Reference*. This is accomplished by scaling the PMOS current mirror dimensions in Fig. 6. The gain of the phase detector is determined primarily by I_{SS} and R_1 , and I_{S1} provides more flexibility in establishing the proper common-mode (CM) level at the output.

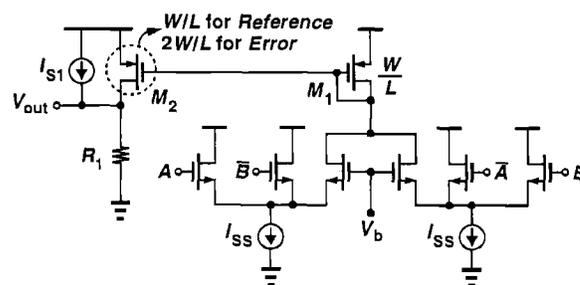


Fig. 6. Symmetric XOR.

C. Charge Pump

Figure 7 shows the implementation of the differential charge pump. The common-mode feedback (CMFB) circuit senses the output CM level by M_5 and M_6 , providing correction through M_3 and M_4 .

Both the matching and channel-length modulation of M_1 - M_4 in Fig. 7 impact the residual phase error in locked condition. Thus, their lengths and widths are relatively large to minimize these effects.

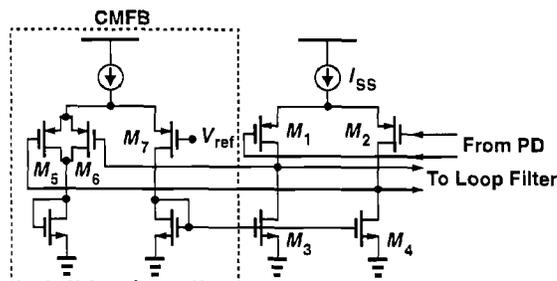


Fig. 7. Charge pump.

Experimental Results

The CDR circuit has been fabricated in a 0.18- μm CMOS technology. The prototype was tested the night before the paper submission deadline. The preliminary results are summarized here.

Figure 8 shows the spectrum of the recovered 5-GHz clock with a 10-Gb/s data sequence of length $2^{23}-1$. The phase noise at 1-MHz offset is approximately equal to -106 dBc/Hz. The noise-shaping of the loop can be improved by adjusting the value of the loop filter components.

Figure 9 depicts the recovered clock and one of the demultiplexed outputs in the time domain. The clock jitter is equal to 6.6 ps,rms and 46.7 ps,pp with a data sequence of $2^{23}-1$, dropping to 4.6 ps,rms and 33.3 ps,pp for a sequence of 2^7-1 . This indicates that the external loop filter capacitor must be increased to reduce the pattern-dependent jitter.

Figure 10 shows both demultiplexed data outputs. The difference between the waveforms results from systematic differences between the bond wires and traces on the test board.

In the preliminary test, the prototype consumes 99 mW from a 2.6-V supply. Due to an unexpectedly low VCO center frequency at the nominal supply voltage, the supply had to be increased to achieve reliable operation at 10 Gb/s.

Acknowledgments

The authors wish to thank NewPort Communications for layout, fabrication, and test support. They are particularly indebted to Armond Hairapetian, German Gutierrez, Afshin Momtaz, Mohammad Tabatabaei, and Pascal Tran.

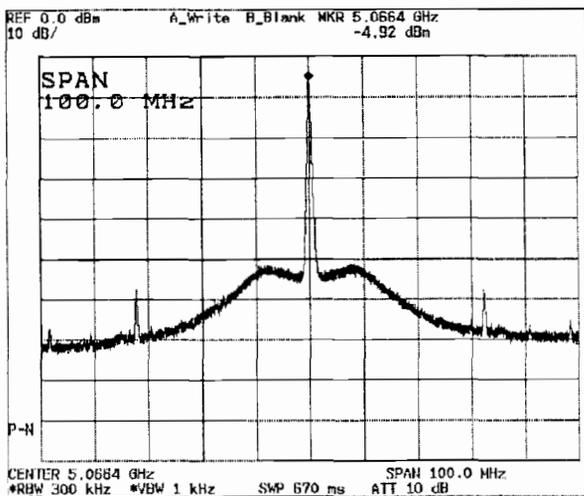


Fig. 8. Spectrum of the recovered clock.

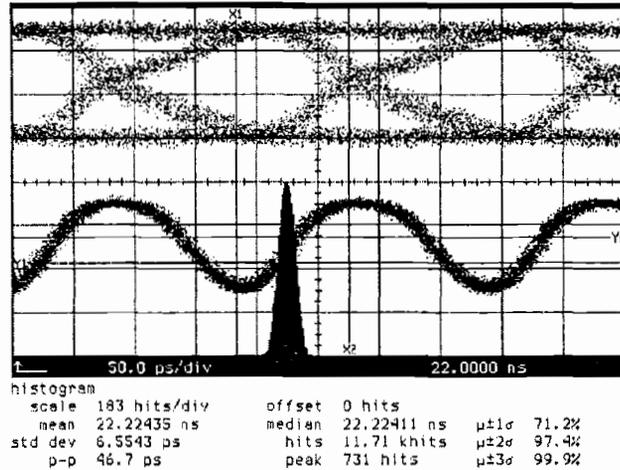


Fig. 9. Recovered clock and one of the demultiplexed outputs in the time domain. (Vert. scale: 100 mV/div.)

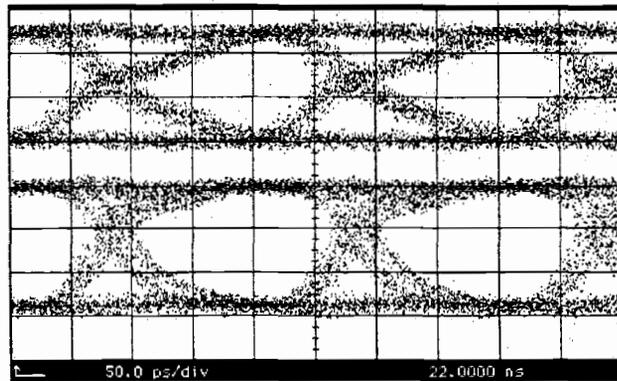


Fig. 10. Demultiplexed outputs. (Vert. scale: 100 mV/div.)

References

- [1] Greshishchev, Y. M., P. Schvan, "SiGe Clock and Data Recovery IC with Linear Type PLL for 10 Gb/s SONET Application," *Proc. IEEE BCTM*, Oct., 1999.
- [2] Wurzer, M., et al., "40-Gb/s Integrated Clock and Data Recovery Circuit in a Silicon Bipolar Technology," *Proc. IEEE BCTM*, pp.136-139, Oct., 1998.
- [3] Rau, M., et al., "Clock/Data Recovery PLL Using Half-Frequency Clock," *IEEE Journal of Solid-State Circuits*, Vol. 32, pp.1156-1159, July 1997.
- [4] Nakamura, K., et al., "A 6 Gb/s CMOS Phase Detecting DEMUX Module Using Half-Frequency Clock," *Digest of Symposium on VLSI Circuits*, pp.196-197, June 1998.
- [5] Razavi, B., 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.