

A 10-Gb/s CMOS Clock and Data Recovery Circuit with a Half-Rate Linear Phase Detector

Jafar Savoj, *Student Member, IEEE*, and Behzad Razavi, *Member, IEEE*

Abstract—A 10-Gb/s phase-locked clock and data recovery circuit incorporates an 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 in an area of $1.1 \times 0.9 \text{ mm}^2$, the circuit exhibits an RMS jitter of 1 ps, a peak-to-peak jitter of 14.5 ps in the recovered clock, and a bit-error rate of 1.28×10^{-6} , with random data input of length $2^{23} - 1$. The power dissipation is 72 mW from a 2.5-V supply.

Index Terms—Clock recovery, half-rate CDR, optical communication, oscillators, phase detectors, PLLs.

I. INTRODUCTION

WITH THE exponential growth of the number of Internet nodes, the volume of the data transported by its backbone continues to rise rapidly. The load of the global Internet backbone is expected to be as high as 11 Tb/s by the year 2005, indicating that the required bandwidth must increase by a factor of 50 to 100 every seven years.

Among the available transmission media, optical fibers have the highest bandwidth with the lowest loss, serving as an attractive solution for the Internet backbone. However, the electronic interface proves to be the bottleneck in designing high-speed optical systems. In order to push the speed of operation beyond the capabilities of the fabrication processes, a number of transceivers can be fabricated on the same chip. The input and output signals can be carried either over a bundle of fibers, or on a single fiber that uses wave-division multiplexing. In this scenario, both the power dissipation and the complexity of each transceiver become critical. While stand-alone building blocks of optical transceivers have been built in GaAs and silicon bipolar technologies [1], [2], full integration of many transceivers makes it desirable to use CMOS technology.

This paper describes the design of the first 10-Gb/s CMOS clock and data recovery (CDR) circuit. A linear phase detector (PD) is introduced 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 with a pseudorandom sequence of $2^{23} - 1$ while dissipating 72 mW from a 2.5-V supply.

The next section of the paper presents the CDR architecture and design issues. Section III deals with the design of the building blocks. Section IV describes the experimental results.

II. 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 that is 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- μm 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. The concept of the half-rate clock has been used in [2]–[5]. However, [2] and [3] incorporate a bang-bang phase detector, possibly creating a 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. The circuit of [5] benefits from a new linear phase detection scheme, but it may not operate properly with certain data patterns.

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 flip-flop) 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

Manuscript received August 21, 2000; revised December 1, 2000. This work was supported in part by the Semiconductor Research Corporation and in part by Cypress Semiconductor.

The authors are with the Electrical Engineering Department, University of California, Los Angeles, CA 90095-1594 USA (e-mail: razavi@icsl.ucla.edu).

Publisher Item Identifier S 0018-9200(01)03020-7.




Fig. 1. Generic CDR architecture.

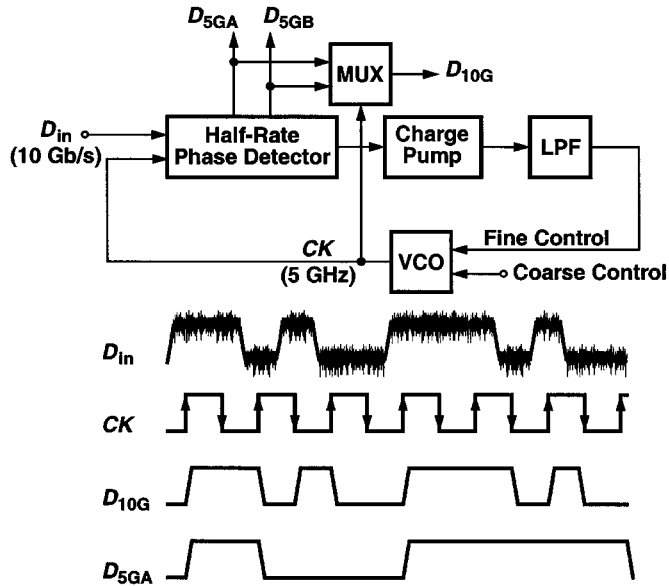


Fig. 2. Half-rate CDR architecture.

presence of a finite skew. For example, the Hogge PD [6] 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 shown in Fig. 2. Here, a half-rate PD produces an error proportional to the phase difference between the 10-Gb/s 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 allowing 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 high speed and low power dissipation while minimizing the ripple on the oscillator control voltage.

It is interesting to note that half-rate architectures do suffer from one drawback: the deviation of the clock duty cycle from 50% translates to bimodal jitter. As depicted in Fig. 3, since both clock edges sample the data waveform, the clock duty cycle distortion pushes both edges away from the midpoint of the bits. Typical duty cycle correction techniques used at lower speeds are difficult to apply here as they suffer from significant dynamic mismatches themselves. Thus, special attention is paid to symmetry in the layout to minimize bimodal jitter.




Fig. 3. Effect of clock duty cycle distortion.

Another important aspect of CDR design is the leakage of data transitions to the oscillator. In Fig. 2, such leakage arises from: 1) capacitive feedthrough from D_{in} to CK in the PD; 2) capacitive feedthrough from D_{5GA} and D_{5GB} to CK through the multiplexer; and 3) coupling of D_{10G} to the oscillator through the substrate. To minimize these effects, the VCO is followed by an isolation buffer and all of the building blocks incorporate fully differential topologies.

III. BUILDING BLOCKS

A. 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 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. 4(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. 4(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 thus improve the jitter performance. Fig. 4(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 $M_1\text{--}M_2$ and $M_3\text{--}M_4$, the current variation is performed through mirror arrangements driven by pMOS differential pairs. Fig. 5 depicts the small-signal gain and phase response of each delay stage. While providing a phase shift of 60° , each stage achieves a gain

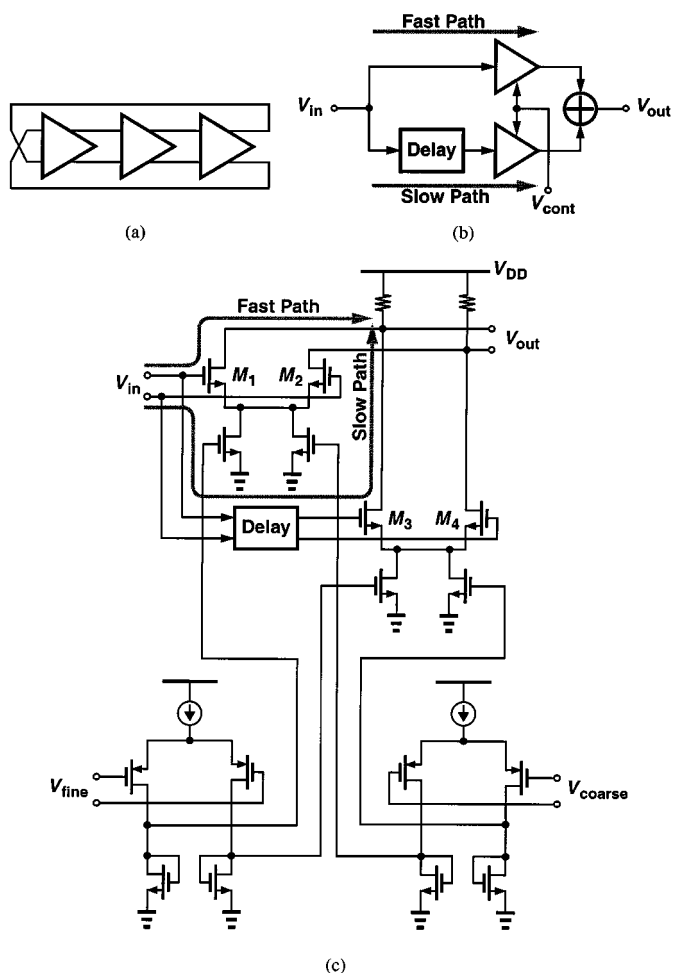


Fig. 4. (a) Three-stage ring oscillator. (b) Implementation of each stage. (c) Transistor-level schematic.

of 5.5 dB at 5 GHz, yielding robust oscillation at the target frequency.

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. The coarse control is provided externally in this prototype. The fine control exhibits a gain of 150 MHz/V and the coarse control, 2.5 GHz/V (Fig. 6). The tuning range is 2.7 GHz ($\approx 54\%$).

B. Phase Detector

Phase detectors generally appear in two different forms. Non-linear PDs coarsely quantize the phase error, producing only a positive or negative value at their output. Linear PDs, on the other hand, generate a linearly proportional output that drops to zero when the loop is locked.

Compared to nonlinear PDs, linear PDs result in less charge pump activity, smaller ripple on the oscillator control line, and hence lower jitter. In a linear PD, such as that described in [6],




Fig. 5. Small-signal gain and phase response of each delay stage.

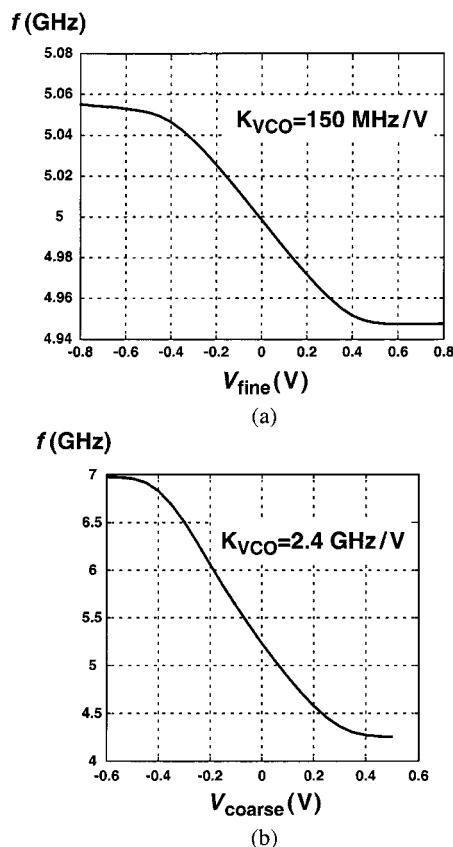


Fig. 6. VCO gain partitioning. (a) Fine control. (b) Coarse control.

the phase error is obtained by taking the difference between the width of two pulses, both of which are generated whenever a data transition occurs. The width of one of the pulses is linearly proportional to the phase difference between the clock and the data, whereas the width of the other is constant. By using a differential error signal, pattern dependency of phase error is can-

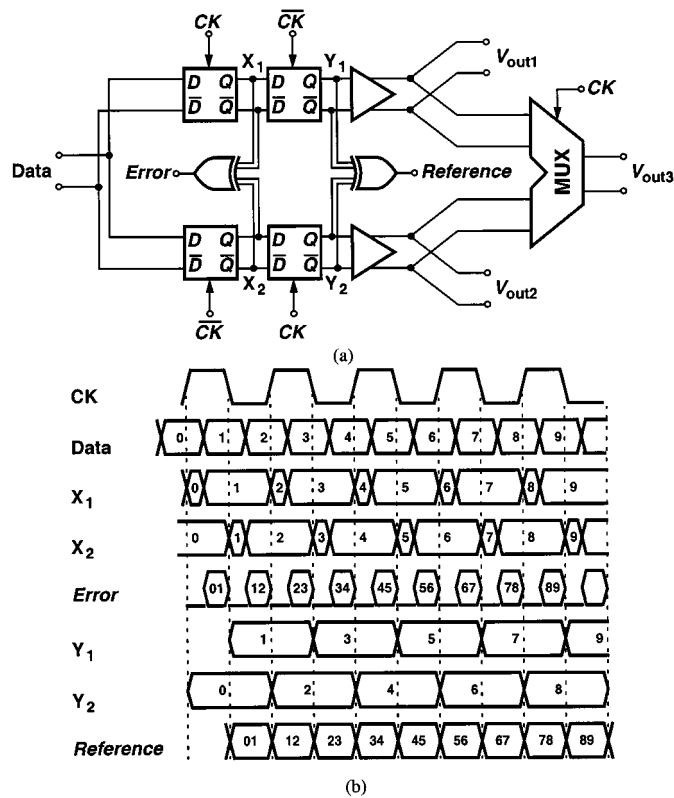


Fig. 7. (a) Phase detector. (b) Operation of the circuit.

celled because both pulses are present only when a data transition occurs.

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. Furthermore, 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. 7(a). The circuit consists of four latches and two 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. 7(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. 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. 7(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 pe-

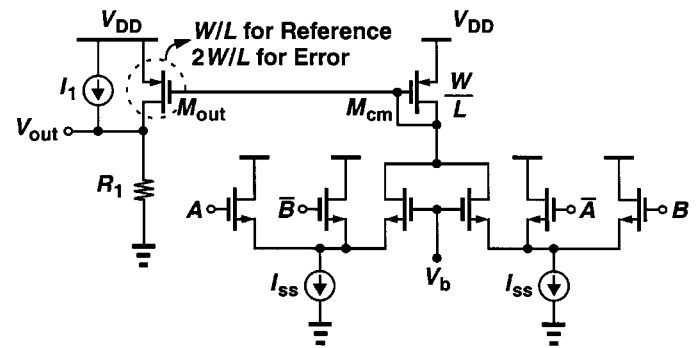


Fig. 8. Symmetric XOR gate.

riodic. 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. 7 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. 8 [7]. The circuit avoids stacking 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.

The gain of the PD is determined by the value of the resistor R_1 and the tail current sources (I_{ss}). The voltage V_b is generated on chip in order to track the variations over temperature and

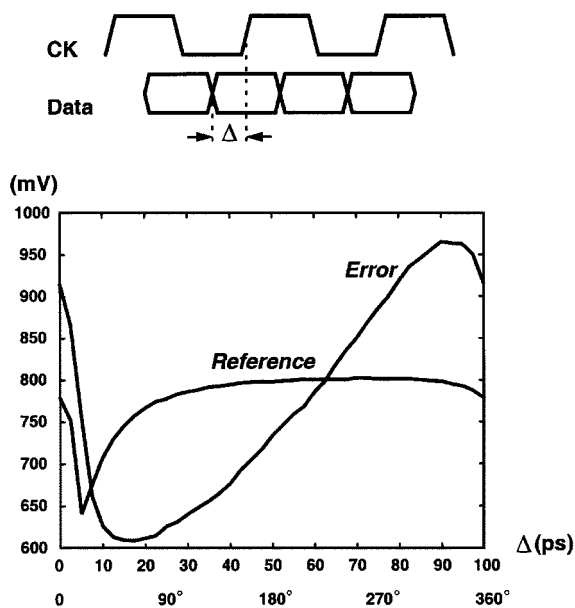


Fig. 9. Determination of PD gain.

process. This voltage equals the output common-mode level of the latches preceding the XOR gate. It is generated using a differential pair that is a replica of the preamplifier section of the latch. Current source I_1 raises the common-mode level of the differential signal formed by the *Error* and *Reference* signals, making V_{out} compatible with the input of the charge pump.

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. Fig. 9 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 systematic 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.

The choice of the logic family used for the XOR gates and the latches is determined by the speed and switching noise considerations. While rail-to-rail CMOS logic achieves relatively high speeds, it requires amplifying the data swings generated by the stage preceding the CDR circuit (typically a limiting amplifier). Furthermore, CMOS logic produces enormous switching noise in the substrate and on the supplies, disturbing the oscillator considerably. For these reasons, the building blocks employ current-steering logic. The phase detector incorporates an input buffer with on-chip resistive matching.

C. Charge Pump and Loop Filter

Fig. 10 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. 10 impact the residual phase error in locked

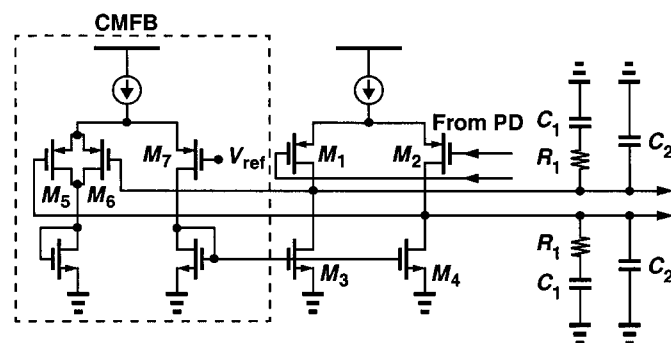


Fig. 10. Charge pump and loop filter.

condition. Thus, their lengths and widths are relatively large to minimize these effects.

The design of the loop filter is based on a linear time-invariant model of the loop and is performed in continuous time domain. The loop is in general a nonlinear time-variant system and can only be assumed linear if the phase error is small. The time-invariant analysis is valid if the averaging behavior of the loop rather than its single-cycle performance is of interest, i.e., the loop can be analyzed by continuous-time approximation if the loop bandwidth is small. Under this condition, the state of the CDR changes by only a small amount on each cycle of the input signal.

A low-pass jitter transfer function with a given bandwidth and a maximum gain in the passband is specified for a SONET system. The closed-loop transfer function of the CDR has a zero at a frequency lower than the first closed-loop pole. This results in jitter peaking that can never be eliminated. But the peaking can be reduced to negligible levels by overdamping the loop.

As derived in [8], the closed-loop unity-gain bandwidth is approximated as

$$\omega_{-3\text{dB}} = R_1 \cdot K_{VCO} \cdot K_{PD} \cdot g_m \quad (1)$$

where K_{VCO} and K_{PD} are the gains of the VCO and PD, respectively, and g_m denotes the conversion gain of the charge pump. Equation (1) can be used to determine the value of R_1 . The amount of the jitter peaking in the closed-loop transfer function can be approximated as

$$JP = 1 + \frac{1}{\omega_{-3\text{dB}} \cdot R_1 \cdot C_1} \quad (2)$$

Equation (2) yields the required value of C_1 . In order to obtain greater suppression of high-frequency jitter, a second capacitor is added in parallel with the series combination of R_1 and C_1 . These components are added externally to achieve flexibility in defining the closed-loop characteristics of the circuit.

Another advantage of linear PDs over their bang–bang counterparts is that their jitter transfer characteristics is independent of the jitter amplitude. It should also be mentioned that if the CDR is followed by a demultiplexer, the tight specifications for jitter peaking need not to be satisfied because such specifications are defined for cascaded regenerators handling full-rate data.

Fig. 11 depicts the simulated behavior of the CDR circuit at the transistor level. The voltage across the filter is initialized to




Fig. 11. Lock acquisition.




Fig. 12. Chip photograph.

a value relatively close to its value in phase lock. The loop goes through a transition of 350 ns before it locks. The ripple on the control line in phase lock is approximately 1 mV.

IV. EXPERIMENTAL RESULTS

The CDR circuit has been fabricated in a 0.18- μm CMOS process. Fig. 12 shows a photograph of the chip, which occupies an area of $1.1 \times 0.9 \text{ mm}^2$. Electrostatic discharge (ESD) protection diodes are included for all pads except the high-speed lines. Nonetheless, since all of these lines have a 50- Ω termination to V_{DD} , they exhibit some tolerance to ESD. 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. While such a high supply voltage creates hot-carrier effects in rail-to-rail CMOS circuits, it is less detrimental in this design because no transistor in the circuit experiences a gate-source or drain-source voltage of more than 1 V. This issue is nonetheless resolved in a second design [9] by proper choice of resistor dimensions. The circuit is brought close to lock with the aid of the VCO coarse control before phase locking takes over.

Fig. 13(a) shows the spectrum of the clock in response to a 10-Gb/s data sequence of length $2^{23} - 1$. The effect of the noise shaping of the loop can be observed in this spectrum. The phase




Fig. 13. (a) Spectrum of the recovered clock. (b) Recovered clock in the time domain.

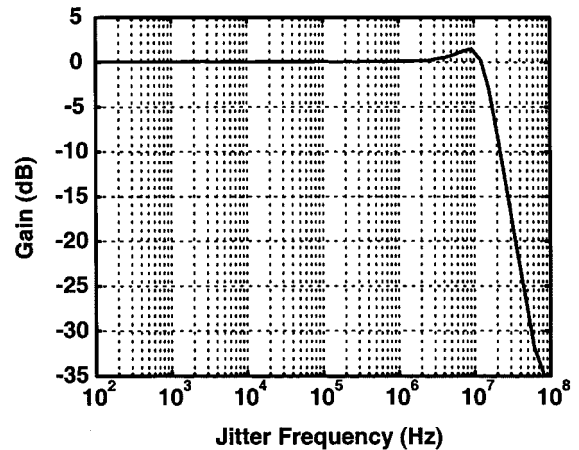


Fig. 14. Measured jitter transfer characteristic.

noise at 1-MHz offset is approximately equal to -106 dBc/Hz . Fig. 13(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 reduced to 4.4 and 0.6 ps, respectively, for a random sequence of length $2^7 - 1$.

The measured jitter transfer characteristics of the CDR is shown in Fig. 14. The jitter peaking is 1.48 dB and the 3-dB bandwidth is 15 MHz. The loop bandwidth can be reduced to

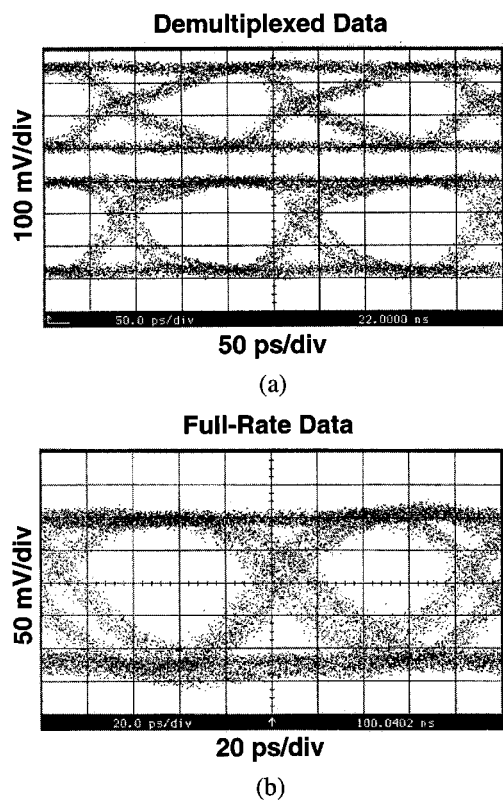


Fig. 15. (a) Recovered demultiplexed data. (b) Recovered full-rate data.

the SONET specifications, but the jitter analyzer must then generate large jitter and drives the loop out of lock. The loop bandwidth can be reduced to the SONET specifications if a means of frequency detection is added to the loop [9]. The circuit is then much less susceptible to loss of lock due to the jitter generated by the analyzer.

Fig. 15 depicts the retimed data. The demultiplexed data outputs are shown in Fig. 15(a). The difference between the waveforms results from systematic differences between the bond wires and traces on the test board. Fig. 15(b) depicts the 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} . This BER can be reduced if the bandwidth of the output buffer driving the 10-Gb/s data is increased. Furthermore, if the value of the linear resistors is adjusted to their nominal value, the increased operating speed of the back-end multiplexer results in an improved BER [9].

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, 33.2, and 18.1 mW, respectively.

V. CONCLUSION

CMOS technology holds great promise for optical communication circuits. The raw speed resulting from aggressive scaling along with high levels of integration provide a high performance at low cost. A 10-Gb/s clock and data recovery circuit designed in 0.18- μm CMOS technology performs phase locking, data regeneration, and demultiplexing with 1 ps of RMS jitter.

REFERENCES

- [1] Y. M. Greshishchev and P. Schvan, "SiGe clock and data recovery IC with linear type PLL for 10-Gb/s SONET application," in *Proc. 1999 Bipolar/BiCMOS Circuits and Technology Meeting*, Sept. 1999, pp. 169–172.
- [2] M. Wurzer *et al.*, "40-Gb/s integrated clock and data recovery circuit in a silicon bipolar technology," in *Proc. 1998 Bipolar/BiCMOS Circuits and Technology Meeting*, Sept. 1998, pp. 136–139.
- [3] M. Rau *et al.*, "Clock/data recovery PLL using half-frequency clock," *IEEE J. Solid-State Circuits*, vol. 32, pp. 1156–1159, July 1997.
- [4] K. Nakamura *et al.*, "A 6 Gb/s CMOS phase detecting DEMUX module using half-frequency clock," in *Dig. Symp. VLSI Circuits*, June 1998, pp. 196–197.
- [5] E. Mullner, "A 20-Gb/s parallel phase detector and demultiplexer circuit in a production silicon bipolar technology with $f_T = 25$ GHz," in *Proc. 1996 Bipolar/BiCMOS Circuits and Technology Meeting*, Sept. 1996, pp. 43–45.
- [6] C. Hogge, "A self-correcting clock recovery circuit," *J. Lightwave Technol.*, vol. LT-3, pp. 1312–1314, Dec. 1985.
- [7] B. Razavi, Y. Ota, and R. G. Swarz, "Design techniques for low-voltage high-speed digital bipolar circuits," *IEEE J. Solid-State Circuits*, vol. 29, pp. 332–339, Mar. 1994.
- [8] L. M. De Vito, "A versatile clock recovery architecture and monolithic implementation," in *Monolithic Phase-Locked Loops and Clock Recovery Circuits, Theory and Design*, B. Razavi, Ed. New York: IEEE Press, 1996.
- [9] J. Savoj and B. Razavi, "A 10-Gb/s CMOS clock and data recovery circuit with frequency detection," in *Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2001, pp. 78–79.



Jafar Savoj (S'98) was born in Tehran, Iran, in 1974. He received the B.S.E.E. degree from Sharif University of Technology, Tehran, in 1996 and the M.S.E.E. degree from the University of California, Los Angeles (UCLA), in 1998. He is currently working toward the Ph.D. degree at UCLA.

He spent the summer of 1998 with Integrated Sensor Solutions, San Jose, CA, working on the design of high-precision interfaces for sensor applications. During the summer of 1999, he was with Newport Communications, Irvine, CA, developing CMOS clock and data recovery circuits for the SONET OC-192 standard.

Mr. Savoj received the IEEE Solid-State Circuits Society Predoctoral Fellowship for 2000–2001, and the Beatrice Winner Award for Editorial Excellence at the 2001 ISSCC. He is also a recipient of the Design Contest Award of the 2001 Design Automation Conference.



Behzad Razavi (S'87–M'90) received the B.Sc. degree in electrical engineering from Sharif University of Technology, Tehran, Iran, in 1985 and the M.Sc. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1988 and 1992, respectively.

He was with AT&T Bell Laboratories, Holmdel, NJ, and subsequently Hewlett-Packard Laboratories, Palo Alto, CA. Since September 1996, he has been an Associate Professor of electrical engineering at the University of California, Los Angeles. His current re-

search includes wireless transceivers, frequency synthesizers, phase-locking and clock recovery for high-speed data communications, and data converters. He was an Adjunct Professor at Princeton University, Princeton, NJ, from 1992 to 1994, and at Stanford University in 1995. He is a member of the Technical Program Committees of the Symposium on VLSI Circuits and the International Solid-State Circuits Conference (ISSCC), in which he is the chair of the Analog Subcommittee. He is an IEEE Distinguished Lecturer and the author of *Principles of Data Conversion System Design* (New York: IEEE Press, 1995), *RF Microelectronics* (Upper Saddle River, NJ: Prentice-Hall, 1998), and *Design of Analog CMOS Integrated Circuits* (New York: McGraw-Hill, 2000), and the editor of *Monolithic Phase-Locked Loops and Clock Recovery Circuits* (New York: IEEE Press, 1996).

Dr. Razavi received the Beatrice Winner Award for Editorial Excellence at the 1994 ISSCC, the Best Paper Award at the 1994 European Solid-State Circuits Conference, the Best Panel Award at the 1995 and 1997 ISSCC, the TRW Innovative Teaching Award in 1997, and the Best Paper Award at the IEEE Custom Integrated Circuits Conference in 1998. He has also served as Guest Editor and Associate Editor of the IEEE JOURNAL OF SOLID-STATE CIRCUITS and IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS and *International Journal of High Speed Electronics*.