

A 10-Gb/s CMOS Clock and Data Recovery Circuit With a Half-Rate Binary Phase/Frequency Detector

Jafar Savoj and Behzad Razavi, *Fellow, IEEE*

Abstract—A 10-Gb/s phase-locked clock and data recovery circuit incorporates a multiphase LC oscillator and a half-rate phase/frequency detector with automatic data retiming. Fabricated in 0.18- μm CMOS technology in an area of $1.75 \times 1.55 \text{ mm}^2$, the circuit exhibits a capture range of 1.43 GHz, an rms jitter of 0.8 ps, a peak-to-peak jitter of 9.9 ps, and a bit error rate of 10^{-9} with a pseudorandom bit sequence of $2^{23} - 1$. The power dissipation excluding the output buffers is 91 mW from a 1.8-V supply.

Index Terms—Bang-bang phase detector, clock and data recovery, frequency detector, ring oscillator, voltage-controlled oscillator (VCO).

I. INTRODUCTION

THE PROBLEM of clock and data recovery (CDR) in high-speed optical communication systems continues to pose interesting challenges at device, circuit, and architecture levels. The design of low-noise oscillators and phase/frequency detectors capable of operating with random data becomes increasingly more difficult as the required speed scales up and the supply voltage scales down, requiring new circuit topologies.

This paper describes the design and experimental verification of a 10-Gb/s phase-locked CDR circuit. Based on a half-rate topology, the architecture incorporates a multiphase LC oscillator and a bang-bang phase/frequency detector with inherent data retiming. The frequency detector guarantees lock for oscillator frequency variations up to 1.43 GHz.

Section II presents the CDR architecture and design issues. Section III deals with the design of the building blocks. Section IV summarizes the experimental results.

II. ARCHITECTURE

With a data rate of 10 Gb/s in a 0.18- μm CMOS technology, many design issues at the circuit and the architecture levels must be considered. While technology-related difficulties such as limited speed and low supply voltage dictate a half-rate CDR architecture [1], the problem of frequency detection imposes additional constraints on the behavior of the overall circuit and its constituent building blocks. More specifically, referenceless frequency acquisition requires that the transition from frequency

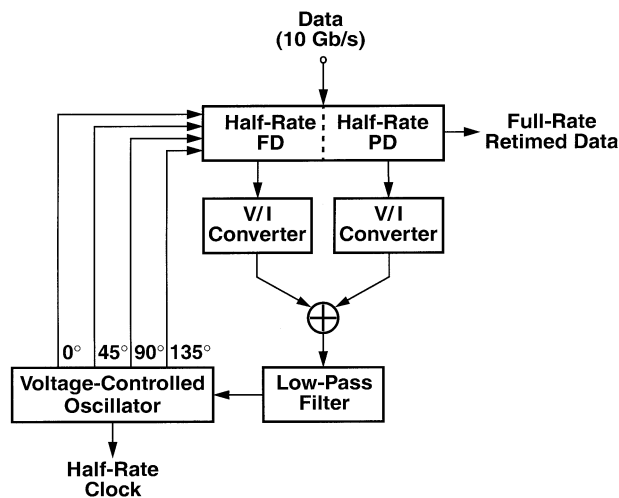


Fig. 1. CDR architecture.

capture to phase capture be “smooth” to avoid glitches that may drive the loop out of lock. This in turn means that, after frequency acquisition is complete: 1) the phase detector (PD) must *automatically* take over and 2) the frequency detector (FD) must produce no output. The design of the PD and even the voltage-controlled oscillator (VCO) is therefore heavily influenced by the FD requirements and vice versa.

Various methods of referenceless frequency detection that have been introduced in prior literature [2], [3] operate with full-rate clocks, failing to generate correct information if used in a half-rate architecture. In this work, a new approach to performing half-rate phase and frequency detection is described. The technique both achieves a high speed and automatically retimes the data.

Shown in Fig. 1, the CDR architecture consists of a PD, an FD, their associated voltage-to-current (V/I) converters, a low-pass filter (LPF), and a VCO. At startup, the FD generates an error signal that drives the VCO frequency toward half of the input data rate, relinquishing the control to the PD when the frequency error is sufficiently small. The PD then locks the VCO phase to the input while producing a retimed output. The inherent retiming capability of the PD proves essential at high speeds as systematic skews in the flipflops would otherwise degrade the tolerance of noise and the quality of detection considerably [4]. Note that the VCO provides half-quadrature phases, which are required for the FD operation. All of the circuits except for the control path of the oscillator are fully differential.

Manuscript received January 3, 2002; revised July 13, 2002.

J. Savoj was with the Electrical Engineering Department, University of California, Los Angeles, CA 90095-1594 USA. He is now with Marvell Semiconductor, Inc., Sunnyvale, CA 94089 USA.

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

Digital Object Identifier 10.1109/JSSC.2002.806284

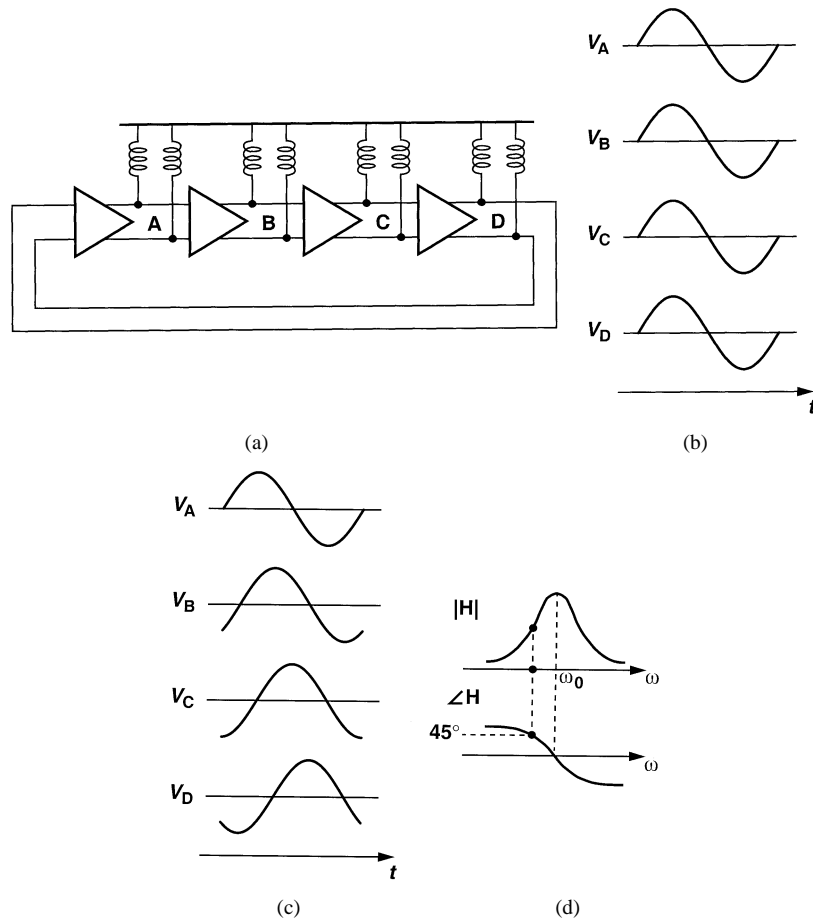


Fig. 2. (a) Four-stage LC -tuned ring oscillator. (b) Waveforms with positive feedback. (c) Waveforms with negative feedback. (d) Amplitude and phase characteristics of each stage.

III. BUILDING BLOCKS

A. VCO

In addition to generating multiple phases, the VCO must satisfy the jitter, tuning range, driving capability, and output swing requirements imposed by the system. An approach to producing half-quadrature phases is to employ a set of coupled oscillators [5]. This technique, however, requires sufficiently strong coupling between adjacent oscillators to avoid parasitic oscillation modes [6], [7]. With no prior knowledge of inductor and oscillator performance in this technology, it was decided to avoid this method. The use of two quadrature oscillators along with phase interpolation suffers from similar issues.

The need for 45° phases naturally leads to a four-stage ring oscillator, but ring topologies using resistive loads suffer from three critical drawbacks: low speed, low Q ($\approx \sqrt{2}$ for four stages [8]), and poor driving capability. The performance of the ring can be substantially improved if the loads are realized as LC tanks [Fig. 2(a)]. Here, the oscillator frequency is a weak function of the number of stages. The open-loop Q of the oscillator is much higher, the load capacitance can be absorbed by the tanks, and the small dc drop across the inductors allows large voltage swings even with a 1.8-V supply.

Unlike resistively-loaded rings, the LC topology of Fig. 2(a) oscillates with positive or negative feedback. If the feedback is positive, then all of the stages operate in unison, thereby pro-

ducing four in-phase signals [Fig. 2(b)].¹ With negative feedback, on the other hand, each stage contributes a phase shift of 45° so that an overall phase of 180° is distributed over the four stages [Fig. 2(c)]. The oscillation frequency is thus offset from $\omega_0 = 1/\sqrt{LC}$ by an amount necessary to yield the requisite 45° phase shift [Fig. 2(d)]. To determine ω_{osc} , we model each tank by a simple parallel network of L , C and R_p , where R_p represents the loss, and set the phase shift to 45°

$$\frac{\pi}{2} - \tan^{-1} \frac{L\omega_{osc}}{R_p(1 - LC\omega_{osc}^2)} = \frac{\pi}{4}. \quad (1)$$

Defining $Q = R_p/L\omega_{osc}$, we have

$$\omega_{osc} = \frac{1}{\sqrt{LC}} \cdot \sqrt{1 - \frac{1}{Q}}. \quad (2)$$

For $Q \approx 4$, $\omega_{osc} \approx 0.87/\sqrt{LC}$. Note that the tank quality factor falls slightly at $\omega_{osc} \neq \omega_0$, limiting the phase noise performance of this topology. This issue can be alleviated by increasing the number of stages so that ω_{osc} approaches ω_0 , but at the cost of a larger number of inductors.

The actual implementation of each stage is shown in Fig. 3(a). To achieve a wide tuning range with a low supply voltage, the tank incorporates MOS varactors rather than pn

¹This property may prove useful if the ring is distributed over a large system so as to provide identical clock phases—in a manner similar to that in [9].




Fig. 3. (a) Implementation of each stage. (b) Stacked inductor.




Fig. 4. Modulation of the varactor capacitance by the noise sources in the oscillator core.

junctions. Realized as NMOS transistors placed inside an n-well, the varactors nonetheless provide a wide dynamic range only if their gate–source voltage can assume both negative and positive values. This is accomplished by the level shift resistor R_1 establishing an output common-mode level approximately 0.9 V below V_{DD} . Note that R_1 and the parasitic capacitances of L_1 form a low-pass filter at node P , suppressing the noise at the second harmonic that can be potentially downconverted to ω_{osc} .

An important effect in the VCO that greatly increased the phase noise in the experimental prototype is the modulation of the varactor capacitance by the noise sources in the oscillator core. As illustrated in Fig. 4, the noise currents produced by I_{ss} and R_1 modulate the output common-mode voltage and hence the capacitance of the two varactors. The output common-mode




Fig. 5. Distributed inductor model.

noise $V_{n,CM}$ can simply be placed in series with V_{tune} to predict the resulting phase noise. Specifically,

$$\overline{V_{n,CM}^2} = \left(\overline{I_{n,ss}^2} + \overline{I_{n,R1}^2} \right) R_1^2 \quad (3)$$

if the low-frequency resistance of the inductors is neglected. Using the narrow-band FM approximation [10], we obtain the resulting phase noise spectrum as

$$S_{\Phi}(\omega_0 + \Delta\omega) = 2\overline{V_{n,CM}^2}(\Delta\omega) \left(\frac{K_{VCO}}{2\Delta\omega} \right)^2. \quad (4)$$

We make two observations. First, since the tail current source must consume a voltage headroom of only a few hundred millivolts, it exhibits a large transconductance and a high noise current. Second, the MOS varactor capacitance slope (dC/dV) reaches a maximum in the middle of the tuning range, yielding a high K_{VCO} and raising the phase noise substantially. In other words, this effect further exacerbates the tradeoff between the tuning range and the phase noise.

The above phenomenon has been verified by SpectreRF simulations, which indicate that the flicker noise of the tail current sources and their diode-connected mirror device contribute 90% of the phase noise at 1-MHz offset when the circuit oscillates at 5.0 GHz (near maximum dC/dV). These simulations predict a phase noise of -88.25 dBc/Hz at 1-MHz offset, and the measured value (Section IV) is approximately -86 dBc/Hz.

The inductors are formed using spiral structures. The inductance value required for oscillation at 5 GHz is sufficiently small (≈ 3 nH) to allow the use of a single spiral. However, simulations indicate that the inductor self-resonance frequency is increased if a modified stacked structure is used. Depicted in Fig. 3(b), this topology forms the bottom layer in metal-3 ($M3$) rather than $M4$ or $M5$, substantially reducing the overall capacitance [11]. Note that the end point of the *bottom* spiral is connected to ac ground.

The limited tuning range of LC oscillators demands careful modeling of the tank and layout parasitics so as to predict the oscillation frequency accurately. The stacked inductor is therefore represented by the distributed network shown in Fig. 5, where the spirals and their associated capacitances are decomposed into four π sections. The loss of each section is modeled by R_1 .

Each stage has a tail current of 4 mA, chosen to provide large voltage swings and hence a high slew rate for low-jitter sampling in the phase detector. Utilizing eight inductors, the

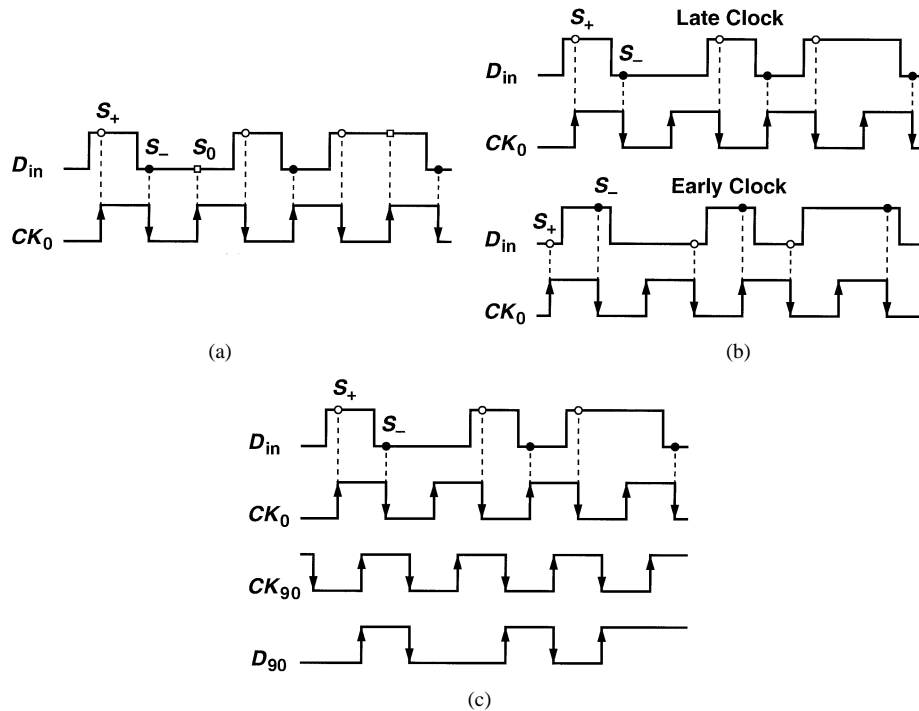


Fig. 6. (a) A DETFF sampling the data by the clock. (b) S_0 samples discarded for late and early clocks. (c) Detecting the absence of data transitions.

VCO occupies a large area ($0.4 \text{ mm} \times 0.5 \text{ mm}$) and introduces a number of issues in the layout and routing. These issues are addressed in the Appendix.

B. Phase/Frequency Detector

The phase detector must sense the phase difference between the input data and the half-rate clock only on data transitions. Furthermore, to allow automatic data retiming, the PD must sample the *data* by the *clock* (rather than the clock by the data), mandating a multipoint sampling scheme. Operating as a bang-bang PD, the topology used in this work is derived from the “data transition tracking loop” (DTTL) described in [12] and [13]. The DTTL utilizes the quadrature phases of the half-rate clock to sample the data in two double-edge-triggered flipflops (DETFFs) (flipflops that sample the data on both edges of the clock). For simplicity, we first present the circuit’s operation for only the rising edges of the clock, considering the cases of early clock and late clock.

To arrive at the PD design, first consider a single DETFF in which the half-rate clock samples the data [Fig. 6(a)]. We note that the samples collectively provide no phase information because they assume positive and negative values with equal probabilities (if 1’s and 0’s occur with equal probabilities). However, let us identify three types of samples: those in the immediate vicinity of positive data edges, S_+ ; those in the immediate vicinity of negative data edges S_- ; and those near no data edges S_0 . We discard the S_0 samples and examine the information carried by S_+ and S_- . As illustrated in Fig. 6(b), S_+ assumes different polarities for early and late clocks, and so does S_- . Nevertheless, S_+ and S_- still fail to represent the phase difference. On the other hand, if the S_- samples are *negated* and used along with the S_+ samples, then the result uniquely determines whether the clock is early or late.

In order to discard the S_0 samples, the PD must detect the absence of data transitions. To this end, we retime D_{in} by a DETFF driven by the quadrature phase of the clock, thereby generating delayed edges of the data [Fig. 6(c)]. A positive transition on D_{90} means an S_+ sample has already been taken and must now be used as the PD output. A negative transition on D_{90} means an S_- sample has already been taken and must be negated and used as the output.

The above observations lead to the PD implementation shown in Fig. 7. Here, two latches operating on opposite clock phases along with a multiplexer form a DETFF [3]. The signal V_I contains the S_+ and S_- samples and the signal V_Q serves as the retimed replica of D_{in} . The third DETFF samples V_I by V_Q , negating S_- on the falling edges of V_Q . As a result, V_{PD} is positive if the clock is late and negative if it is early.

The PD of Fig. 7 operates at high speeds because it employs a half-rate clock. Since in the locked condition the rising and falling edges of the in-phase clock CK_I coincide with data transitions, the quadrature clock samples D_{in} at the optimum point, thus generating at V_Q a full-rate, retimed data stream free of systematic offset.

If the in-phase clock CK_I in Fig. 7 samples the zero crossings of data, then the two latches sensing V_I become metastable, yielding a small differential output at V_{PD} . Thus, the subsequent V/I converter produces a small net output current, allowing the loop filter to retain the voltage developed in previous comparisons.

The binary PD used in this work is somewhat similar to a simple D flipflop PD [14] in that it does not assume a tristate output in the absence of data transitions. In other words, even without any data edges, the PD continues to generate a high or low level, thus charging or discharging the loop filter and forcing the VCO frequency to drift. Nevertheless, with a loop

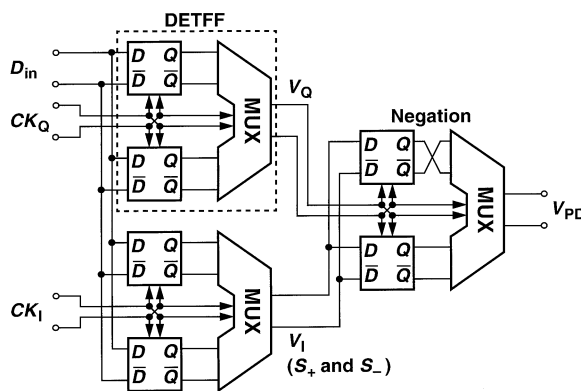




Fig. 7. Phase detector.



(a)

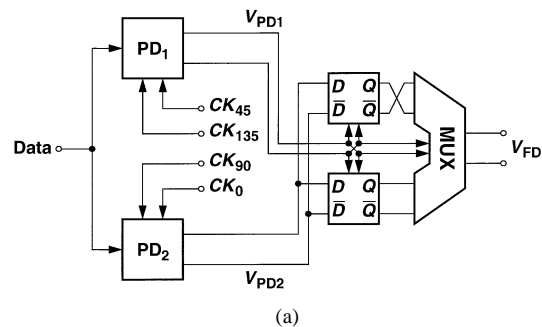


(b)

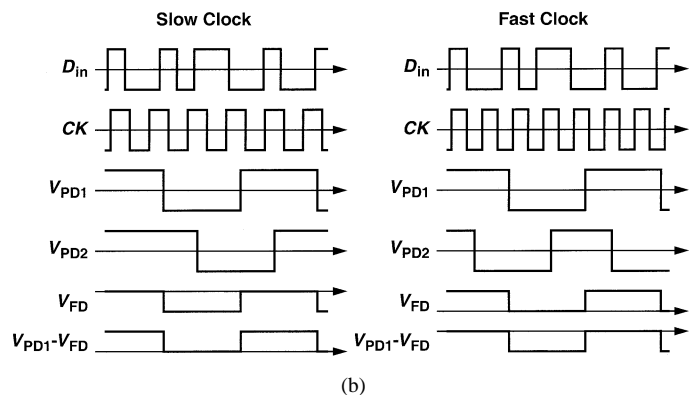
Fig. 8. (a) Lock acquisition. (b) Clock phase in presence of 72 consecutive bits.

bandwidth of about 5 MHz, the CDR circuit can tolerate relatively long runs. Fig. 8 plots the simulated control voltage and oscillator output phase versus time as the loop acquires lock and experiences run lengths of 72 bits (at $t = 350$ ns and 450 ns). The maximum phase deviation after each run is approximately 10% of the clock period.

As with other half-rate phase detectors, the PD of Fig. 7 produces a beat component when sensing a data rate unequal to twice the clock frequency. However, the dc content of the beat is still zero, failing to represent the frequency error. In this work, we extend the concept of full-rate frequency detection by means of quadrature sampling (i.e., as in “quadricorrelator” [15] and its digital version [3]) to half-rate frequency detection through the use of half-quadrature phases. Depicted in Fig. 9(a), the FD incorporates two instances of the half-rate PD, with the quadrature clock phases applied to PD₂ leading those applied to PD₁



(a)



(b)

Fig. 9. (a) Phase and frequency detector. (b) Timing diagram in the PFD for slow and fast clock signals.

by 45°. From the waveforms shown in Fig. 9(b), we recognize the following.

- If the clock is slow, then V_{PD1} leads V_{PD2} . Therefore, if V_{PD2} is sampled by the rising (falling) edges of V_{PD1} , the result is negative (positive).
- If the clock is fast, then V_{PD1} lags V_{PD2} . Therefore, if V_{PD2} is sampled by the rising (falling) edges of V_{PD1} , the result is positive (negative).

To avoid the ambiguity associated with falling and rising edges in the two cases, the FD of Fig. 9(a) employs a modified DETFF similar to that used in the PD of Fig. 7. The idea is that the falling edges of V_{PD1} sample and negate V_{PD2} before routing it to V_{FD} . The signals V_{PD1} and V_{FD} can therefore serve as the phase error and frequency error signals, respectively.

As shown in Fig. 9(b), if the PFD is designed such that V_{FD} has a ternary output, then the difference between V_{PD1} and V_{FD} contains positive and negative pulses for slow and fast clock signals, respectively. The polarity of these pulses determines the sign of the frequency error. The multiplexer in the back-end flipflop of the PFD can be modified to produce ternary pulses [3] (Fig. 10). In the two normal states of operation, current is steered to one of the two load resistors. In the third state, current does not flow through any of the resistors and the output settles at the common-mode level. This characteristic also tristates V_{FD} when phase lock is achieved. Each of the three inputs to the multiplexer (V_{in1} , V_{in2} , and V_{CK}) has a differential peak-to-peak swing of 1.2 V.

C. V/I Converter

Fig. 11 shows the implementation of the V/I converter. Since the circuit drives the single-ended control of the varactors, it is designed to provide a single-ended output. In phase lock, the





Fig. 10. Modified multiplexer.

Fig. 11. V/I converter.

differential frequency error signal falls to zero. Therefore, I_2 is equally split between M_3 and M_4 , affecting V_{out} negligibly. In order to reduce the ripple on the oscillator control voltage, a relatively small current ($\approx 125 \mu\text{A}$) is used.

Simulation results indicate that the capture range of the circuit is smaller than the VCO tuning range because the output of the V/I converter cannot swing from rail to rail. All of the transistors in the V/I converter have small overdrive voltages to alleviate this issue.

D. Output Buffers

The output buffer delivering high-speed data to the external load requires a wide band with well-behaved phase response so as to introduce negligible intersymbol interference (ISI). In order to drive doubly terminated loads with reasonable voltage swings, the buffer must carry a large current, thus demanding large transistors and hence exhibiting substantial capacitance at its input. Inductive peaking has proved an efficient approach to increasing the bandwidth in CMOS amplifiers [16], [17]. Shown in Fig. 12, the buffer utilizes inductive peaking so as to drive the output differential pair. Simulations indicate that the use of inductive peaking in every other stage is sufficient to provide ISI-free operation, reducing the required area. Each differential pair is tapered up by a factor of two with respect to its preceding stage. Since the quality factor of the inductors is not critical here, the line width of the spiral structures is reduced to the minimum necessary for electromigration. The width is therefore set to $4 \mu\text{m}$ to achieve a high self-resonance frequency. Table I displays the component values used in the design.

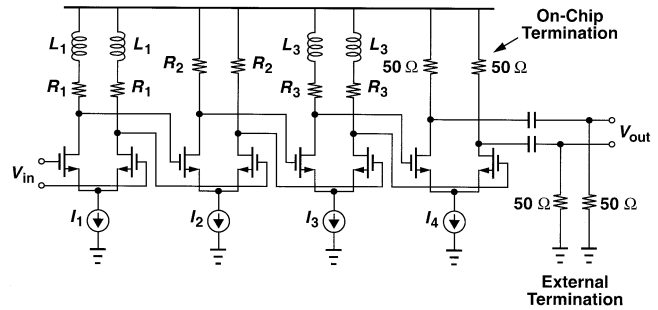


Fig. 12. Output buffer.

TABLE I
OUTPUT BUFFER COMPONENT VALUES

I_1	1 mA	R_1	600 Ω	L_1	3.5 nH
I_2	2 mA	R_2	300 Ω	L_3	1.5 nH
I_2	4 mA	R_3	150 Ω		
I_4	10 mA				




Fig. 13. Chip photograph.

IV. EXPERIMENTAL RESULTS

The CDR circuit has been fabricated in a $0.18\text{-}\mu\text{m}$ CMOS process. Fig. 13 shows a photograph of the chip, which occupies an area of $1.75 \times 1.55 \text{ mm}^2$. ESD protection diodes are included for all but high-speed pads. Nonetheless, since all of the high-speed lines have a $50\text{-}\Omega$ termination to V_{DD} , they exhibit some tolerance to ESD. The circuit is tested in a chip-on-board assembly while running from a 1.8-V supply.

Fig. 14(a) depicts the measured VCO tuning characteristic. The VCO achieves a tuning range of 1.2 GHz ($\approx 24\%$). The maximum gain of the VCO is quite high, about 1.22 GHz/V , be-




Fig. 14. (a) VCO tuning range. (b) Phase noise over tuning range.

cause the circuit was conservatively designed to achieve an oscillation frequency of 5 GHz despite poor models of active and passive devices. Experimental results on six chips indicate that the tuning characteristic of the VCO varies by about 1%.² This variation may reach a few percent for oscillators fabricated in different lots and operating at temperature extremes, suggesting that the tuning range can be reduced to achieve a smaller VCO gain and a better phase noise performance. The VCO achieves the highest signal purity at the lower end of its tuning range (-102.35 dBc/Hz at 1-MHz offset). The open-loop VCO phase noise at 5 GHz is -86 dBc/Hz. As explained in Section III, this degradation is attributed to the modulation of the varactor capacitance by the common-mode noise in the VCO core. Despite the relatively high phase noise, the closed-loop CDR suppresses the in-band components, providing a low jitter.

Fig. 15(a) shows the spectrum of the clock in response to a 9.953 28-Gb/s data sequence of length $2^{23} - 1$. The phase noise at 1-MHz offset is approximately equal to -107 dBc/Hz. Fig. 15(b) depicts the recovered clock in the time domain. The jitter performance of the CDR circuit is characterized by the Anritsu MP1777 jitter analyzer. A random sequence of length $2^{23} - 1$ produces 9.9 ps of peak-to-peak jitter and 0.8 ps of rms jitter on the clock signal. These values are reduced to 2.4 and 0.4 ps, respectively, for a random sequence of length $2^7 - 1$. SONET OC-192 specifies 0.1 unit interval (10 ps) as the max-

²Note that the inductance is defined by lithography and the varactor capacitance by the gate oxide thickness, both of which are controlled very tightly.




Fig. 15. (a) Spectrum of the recovered clock. (b) Recovered clock in the time domain.




Fig. 16. Measured jitter transfer characteristic.

imum peak-to-peak jitter on the clock signal if a bandpass filter is placed between the CDR circuit and the measuring equipment [18]. The measured jitter number is less than the number specified by the standard in the absence of the bandpass filter.

The measured jitter transfer characteristic of the CDR using a $2^{23} - 1$ pseudorandom bit sequence (PRBS) is shown in Fig. 16. The jitter peaking is 0.04 dB and the -3 dB bandwidth is 5.2 MHz.

Fig. 17 depicts the full-rate retimed data. With a PRBS of $2^{23} - 1$ and no input jitter, the BER is equal to 10^{-9} . The circuit therefore does not pass the SONET jitter tolerance test. The limited bandwidth of the output buffer appears to cause this failure.

Despite the small loop bandwidth, the frequency detector provides a capture range of 1.43 GHz, obviating the need for ex-




Fig. 17. Recovered clock and data.

ternal references. Note that this number is smaller than twice the VCO tuning range because the V/I converter driving the control line of the VCO cannot swing from rail to rail. The total power consumed by the circuit excluding the output buffers is 91 mW from a 1.8-V supply. The VCO, the PFD, and the clock and data buffers consume 30.6, 42.2, and 18.2 mW, respectively.

V. CONCLUSION

A 10-Gb/s CDR circuit designed in 0.18- μm CMOS technology performs frequency acquisition, phase locking, and data regeneration. Achieving an rms jitter of 0.8 ps, this circuit is the first CMOS CDR circuit to meet the jitter generation requirements defined by SONET. The power consumption of this circuit is much smaller than the power consumption of similar circuits fabricated in bipolar or GaAs processes. The low power dissipation, high integration, and low cost of the CMOS process hold great promise for implementation of optical communication circuits in this technology.

APPENDIX LAYOUT ISSUES

The VCO occupies a large chip area (0.4 mm \times 0.5 mm) as it incorporates eight spiral inductors. Therefore, the metal lines carrying the multiphase clock signals from the core of the VCO to the surrounding circuits are very long. These interconnects are laid out using wide traces of the top metal layer in order to reduce the resistance of the wire. This results in a large routing capacitance because the fringe capacitance of the top metal layer in a 0.18- μm CMOS technology is several times the bottom-plate capacitance. Shown in Fig. 18 are two possible arrangements for the VCO and its buffer stages. In Fig. 18(a), each buffer is placed next to each stage in the VCO, thereby minimizing the capacitance seen by the VCO core but presenting a large load capacitance to each buffer. In Fig. 18(b), on the other hand, the buffers are close to the phase detector, thus requiring that the VCO core drive significant interconnect capacitance. Simulations using distributed models for the interconnects indicate that the topology of Fig. 18(a) considerably attenuates the clock signals that reach the PFD (primarily because similar issues prohibit the buffers themselves from using inductors). Thus, the latter configuration is employed at the cost of limiting the tuning range and creating uncertainty in the oscillation frequency.

The differential interconnects suffer from a large fringe capacitance. To remedy this issue, the routing can be designed




Fig. 18. Placing the clock buffers (a) next to and (b) far from the VCO core.

such that adjacent lines carry signals that are close in phase, e.g., 45° rather than 180° . But, as is evident from Fig. 18(a), such an arrangement yields unequal line lengths and hence mismatch between the clock phases. A spacing of 5 μm is therefore used between differential lines to minimize the fringe component.

The metal trace connecting the control line of the VCO to the pad is shielded by two metal layers that are connected to the VCO supply. This minimizes noise coupling from the environment.

REFERENCES

- [1] J. Savoj and B. Razavi, "A 10-Gb/s CMOS clock and data recovery circuit with a half-rate linear phase detector," *IEEE J. Solid-State Circuits*, vol. 36, pp. 761–768, May 2001.
- [2] 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.
- [3] A. Pottbacker, U. Langmann, and H. U. Schreiber, "A Si bipolar phase and frequency detector IC for clock extraction up to 8 Gb/s," *IEEE J. Solid-State Circuits*, vol. 27, pp. 1747–1751, Dec. 1992.
- [4] C. Hogge, "A self-correcting clock recovery circuit," *J. Lightwave Technol.*, vol. LT-3, pp. 1312–1314, Dec. 1985.
- [5] J. Kim and B. Kim, "A low phase-noise CMOS LC oscillator with a ring structure," in *ISSCC Dig. Tech. Papers*, Feb. 2000, pp. 430–431.
- [6] C. Lam and B. Razavi, "A 2.6 GHz/5.2 GHz CMOS voltage-controlled oscillator," in *ISSCC Dig. Tech. Papers*, Feb. 1999, pp. 402–403.
- [7] T.-P. Liu, "A 6.5 GHz monolithic CMOS voltage-controlled oscillator," in *ISSCC Dig. Tech. Papers*, Feb. 1999, pp. 404–405.
- [8] B. Razavi, "A study of phase noise in CMOS oscillators," *IEEE J. Solid-State Circuits*, vol. 31, pp. 331–343, Mar. 1996.
- [9] J. Wood, T. C. Edwards, and S. Lipa, "Rotary traveling wave oscillator arrays: A new clock technology," *IEEE J. Solid-State Circuits*, vol. 36, pp. 1654–1665, Nov. 2001.
- [10] L. W. Couch, *Digital and Analog Communication Systems*, 4th ed. New York: Macmillan, 1993.
- [11] A. Zolfaghari, A. Chan, and B. Razavi, "Stacked inductors and transformers in CMOS technology," *IEEE J. Solid-State Circuits*, vol. 36, pp. 620–628, Apr. 2001.
- [12] T. O. Anderson, W. J. Hurd, and W. C. Lindsey, "Transition Tracking Bit Synchronization System," U.S. Patent 626 298, Dec. 1971.
- [13] A. W. Buchwald, "Design of Integrated Fiber-Optic Receivers Using Heterojunction Bipolar Transistors," Ph.D. dissertation, Univ. of California, Los Angeles, 1993.
- [14] M. Soyuer, "A monolithic 2.3-Gb/s 100-mW clock and data recovery circuit in silicon bipolar technology," *IEEE J. Solid-State Circuits*, vol. 28, pp. 1310–1313, Dec. 1993.

- [15] D. Richman, "Color carrier reference phase synchronization accuracy in NTSC color television," *Proc. IRE*, vol. 42, pp. 106–133, Jan. 1954.
- [16] J. Savoj and B. Razavi, "A CMOS interface circuit for detection of 1.2 Gb/s RZ data," in *ISSCC Dig. Tech. Papers*, Feb. 1999, pp. 278–279.
- [17] S. Mohan, M. Hershenson, S. Boyd, and T. Lee, "Bandwidth extension in CMOS with optimized on-chip inductors," *IEEE J. Solid-State Circuits*, vol. 35, pp. 346–355, Mar. 2000.
- [18] "SONET OC-192 transport system generic criteria," *GR-1377-CORE*, no. 5, Dec. 1998.



Behzad Razavi (S'87-M'90-SM'00-F'03) received the B.Sc. degree in electrical engineering from Sharif University of Technology, 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 and Hewlett-Packard Laboratories until 1996. Since September 1996, he has been an Associate Professor and subsequently Professor of electrical engineering at the University of California, Los Angeles. His current research 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 served on the Technical Program Committee of the International Solid-State Circuits Conference (ISSCC) from 1993 to 2002 and is presently a member of the Technical Program Committee of Symposium on VLSI Circuits. He is the author of *Principles of Data Conversion System Design* (New York: IEEE Press, 1995), *RF Microelectronics* (Upper Saddle River, NJ: Prentice-Hall, 1998) (also translated to Japanese by Tadahiro Kuroda), *Design of Analog CMOS Integrated Circuits* (New York: McGraw-Hill, 2001), and *Design of Integrated Circuits for Optical Communications* (New York: McGraw-Hill, 2002), and the editor of *Monolithic Phase-Locked Loops and Clock Recovery Circuits* (New York: IEEE Press, 1996).

Prof. 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 was the co-recipient of both the Jack Kilby Outstanding Student Paper Award and the Beatrice Winner Award for Editorial Excellence at the 2001 ISSCC. He is an IEEE Distinguished Lecturer and has also served as Guest Editor and Associate Editor of the IEEE JOURNAL OF SOLID-STATE CIRCUITS, IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS, and *International Journal of High Speed Electronics*.

Jafar Savoj received the B.Sc. degree in electrical engineering from Sharif University of Technology, Iran, in 1996 and the M.Sc. and Ph.D. degrees in electrical engineering from the University of California, Los Angeles, in 1998 and 2001, respectively.

He was with Transpectrum, Los Angeles, CA, and is presently with Marvell Semiconductor, Inc., Sunnyvale, CA. He serves as the Panel Chair of the Custom Integrated Circuits Conference (CICC) and is the author of *High-Speed CMOS Circuits for Optical Receivers* (Norwell, MA: Kluwer, 2001).

Dr. 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.