A 2.5-Gb/s 15-mW Clock Recovery Circuit

Behzad Razavi, *Member, IEEE*

***Abstract*— This paper describes the design of a 2.5-Gb/s 15-mW clock recovery circuit based on the quadricorrelator architecture. Employing both phase and frequency detection, the circuit combines high-speed operations such as differentiation, full-wave rectification, and mixing in one stage to lower the power dissipation. In addition, a two-stage voltage-controlled oscillator is utilized that incorporates both phase shift elements to provide a wide tuning range and isolation techniques to suppress the feedthrough due to input data transitions. Fabricated in a 20-GHz 1-μm BiCMOS technology, the circuit exhibits an rms jitter of 9.5 ps and a capture range of 300 MHz.**

I. Introduction

HIGH-SPEED low-power clock recovery circuits find wide application in high-performance communication systems [1]. This paper describes the design of a 2.5-Gb/s 15-mW clock recovery circuit (CRC) fabricated in a 20-GHz 1-μm BiCMOS technology. Employing a modified version of the "quadricorrelator" architecture [2], [3], the circuit extracts the clock from a nonreturn-to-zero (NRZ) data sequence using both phase and frequency detection. In addition to low power, an important concern has been to achieve a relatively wide capture range so that the circuit can lock to the input in the presence of temperature and process variations.

The clock recovery circuit has been intended for use at the front-end of fiber-optic receivers operating at 2.5 Gb/s (the SONET OC-48 standard). While the existing implementations at this rate are mostly in III-V technologies [4], it is advantageous to perform this function in mainstream VLSI processes so as to allow higher levels of integration. Shown in Fig. 1 is a block diagram of a typical fiber-optic front end, where NRZ data is received in the form of light and converted to current by a photodetector. The current is then amplified and converted to voltage by a transimpedance amplifier, Z. Since further digital processing requires timing information on the data, a clock recovery circuit extracts the clock and a decision circuit retimes the data. The following data processing can thus be performed synchronously.

While cost, reliability, and performance issues make it desirable to integrate the system of Fig. 1 on a single chip, the total power dissipated in the high-speed blocks often becomes prohibitively large. Thus, it is important that each circuit be designed for minimum power dissipation. This issue becomes even more significant if many receive channels are to be integrated on the same chip [5].

In the next section of the paper, we describe issues related to extracting the clock from NRZ data. In Section III, we present the CRC architecture and in Section IV, circuit implementation of each building block. The experimental results are summarized in Section V.

Manuscript received August 31, 1995; revised November 27, 1995.

The author was with AT&T Bell Laboratories, Holmdel, NJ USA. He is now with Hewlett-Packard Laboratories, Palo Alto, CA 94304 USA.

Publisher Item Identifier S 0018-9200(96)02658-3.

II. Design Issues

Random NRZ data has two properties that directly influence the design of clock recovery circuits. We examine both the frequency-domain and time-domain behavior of the NRZ format to understand these properties and their implications.

For a random binary sequence with bit rate r_b and equal probability of ONE's and ZERO's, the power spectral density is [6]

$$P_x(\omega) = T_b \left[\frac{\sin(\omega T_b/2)}{\omega T_b/2} \right]^2 \tag{1}$$

where $T_b = 1/r_b$. This function exhibits nulls at integer multiples of r_b. Intuitively, we note that the fastest NRZ waveform with bit rate r_b occurs when consecutive bits alternate between ONE and ZERO. The result is a square wave with a frequency equal to $r_b/2$, containing no even-order harmonics [7]. This is illustrated in the MATLAB simulation of Fig. 2(a) for a 2.5-Gb/s NRZ sequence, indicating that no linear, time-invariant operation can extract a 2.5-GHz periodic clock from this data. Thus, a nonlinear function, e.g., edge detection, must be used. One approach to edge detection is to generate a positive impulse for every positive or negative data transition. Fig. 2(b) shows the resulting spectrum if the data undergoes differentiation and full-wave rectification, displaying strong components at 2.5 GHz and its harmonics.

Another property of NRZ data is that it can exhibit long sequences of consecutive ONE's or ZERO's, an important issue if the CRC employs phase-locking (Fig. 3). In the absence of data transitions, the dc component produced by the mixer is zero, and the control voltage of the oscillator, which is stored in the low-pass filter (LPF), gradually diminishes, thereby causing the output frequency to drift. Consequently, the recovered clock suffers from input-dependent jitter. To minimize this effect, the LPF time constant must be sufficiently longer than the maximum length of consecutive ONE's or ZERO's, a remedy that inevitably leads to a small loop bandwidth and hence a narrow capture range.

Since the center frequency of monolithic oscillators varies substantially with temperature and process, the CRC must employ some means of frequency detection and acquisition so as to guarantee locking.

III. Architecture

The clock recovery circuit is based on the "quadricorrelator" architecture, introduced by Richman [2] and modified by

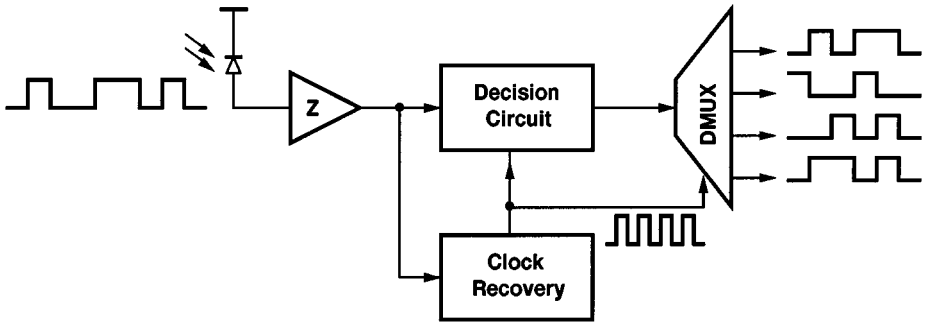

Fig. 1. Front-end of a fiber-optic receiver.

Bellisio [3]. This architecture provides frequency detection of random inputs and is therefore well suited to clock recovery in data communications. Shown in Fig. 4 in simplified form, the quadricorrelator consists of an edge detector and two loops that share the same voltage-controlled oscillator (VCO). Note that the VCO generates quadrature outputs. Loop I multiplies the edge-detected signal by one of the VCO outputs, applies the product to a low-pass filter (LPF), and mixes the resulting low-frequency component with the signal generated by Loop II. The operation of Loop II is similar except that the LPF is followed by a differentiator.

The quadricorrelator of Fig. 4 operates as follows. Suppose the signal at point A has a component at ω_1 while the VCO oscillates at a different frequency, ω_2. Upon mixing the VCO outputs with $\sin\omega_1 t$ and low-pass filtering the results, the circuit yields quadrature beat signals $\sin(\omega_1 - \omega_2)t$ and $\cos(\omega_1 - \omega_2)t$. Next, the latter signal is differentiated and mixed with the former, yielding $(\omega_1 - \omega_2)\cos^2(\omega_1 - \omega_2)t$ at node P. The following LPF extracts the average value of this waveform, $(\omega_1 - \omega_2)/2$, a quantity representing both the polarity and the magnitude of the difference between ω_1 and ω_2. The averaged signal drives the VCO with negative feedback so as to bring ω_2 closer to ω_1. For a large loop gain, the difference between ω_1 and ω_2 eventually drops to a small (but finite) value.

While typical phase/frequency detectors used in charge-pump phase-locked loops (PLL's) [10] produce incorrect outputs if the input signal exhibits missing transitions [11], the quadricorrelator detects the frequency difference properly even if the input signal is random. It can be shown that if the power spectral density at point A in Fig. 4 is symmetric around ω_1, then the averaged value at point Q is proportional to $(\omega_1 - \omega_2)\sigma_{in}^2$, where σ_{in}^2 denotes the input variance [12]. To use this result, we note that the spectrum shown in Fig. 2(b) is approximately symmetric for about $\pm$ 500 MHz around 2.5 GHz but exhibits substantial asymmetry for greater frequency offsets. However, we also note that low-pass filtering after mixing the signal with the VCO output is approximately equivalent to band-pass filtering before down-conversion (Fig. 5). Thus, the quadricorrelator suppresses the asymmetric part of the spectrum, performing correct frequency detection with edge-detected NRZ data.

As mentioned above, the simple quadricorrelator of Fig. 4 reduces the frequency difference between the input and the VCO output to a small but finite value. Since in clock recovery and data retiming, the two frequencies must be precisely equal and the phase of the recovered clock must be aligned with that of the incoming data, the quadricorrelator requires additional means to perform phase detection. As shown in Fig. 6, this can be accomplished by adding a third loop—a simple PLL—to the circuit. Here, as $|\omega_2 - \omega_1|$ drops, Loop III begins to generate an asymmetric signal at node M, assisting the lock process. For $\omega_2 \approx \omega_1$, the dc feedback signal produced by Loops I and II approaches zero and Loop III dominates, locking the VCO output to the input data.

The quadricorrelator has been realized in various analog and digital forms [8], [4], [9]. To achieve a low-power, compact implementation, the architecture is modified as shown in Fig. 7. Here, the edge detector (consisting of a differentiator and a full-wave rectifier) is moved to each arm such that it can be merged with the following mixer and low-pass filter. The resulting circuit performs differentiation, full-wave rectification, mixing, and low-pass filtering in one stage, and is denoted by DRML. Similarly, the second differentiator, mixer, and LPF can be combined (DML in Fig. 7). Note that the high-speed path now includes only three blocks, allowing significant savings in power dissipation. The architecture is fully differential to suppress common-mode noise and provide robust operation from a 3-V supply.

Shown in Fig. 8 is the lock behavior of the CRC simulated at the transistor level. Here, the control voltage of the VCO is plotted as a function of time. At $t = 0$, the input rate is 2.3 Gb/s and the VCO oscillates at 3 GHz, After approximately 400 ns, the circuit locks to the input. At $t = 600$ ns, the input rate is stepped to 2.7 Gb/s, and the CRC takes approximately 400 ns to lock again.

IV. Circuit Implementation

A. DRML

The DRML block of Fig. 7 must differentiate and full-wave rectify the NRZ data, multiply the result by one of the VCO outputs, and low-pass filter the product. The implementation is depicted in Fig. 9.

To differentiate the NRZ data, a bipolar differential pair with capacitive degeneration is used [Fig. 9(a)]. Here, when D_{in} goes high, transistor Q_3 momentarily "hogs" the bias current of Q_4, creating impulses of opposite polarity in their respective collector currents. The operation is similar for

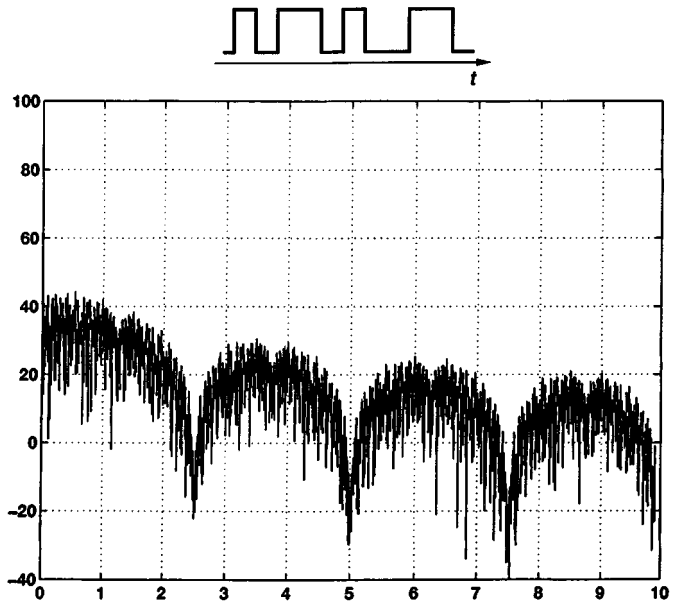

(a)

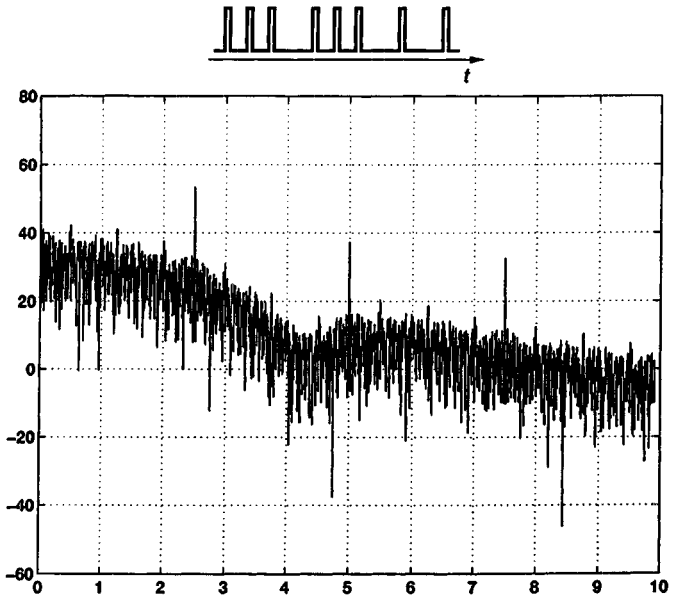

(b)

Fig. 2. Simulated spectrum of NRZ data (a) before and (b) after edge detection. (Horizontal 1 GHz/div., vertical 20 dB/div.)

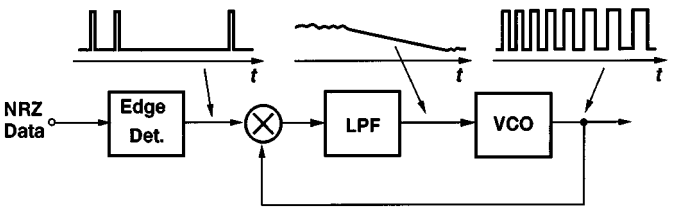

Fig. 3. Simple phase-locked clock recovery circuit.

falling transitions of D_{in}. Determined by the input voltage swing, the value of C_D, and the tail currents, the width of the impulses is approximately equal to 100 ps.

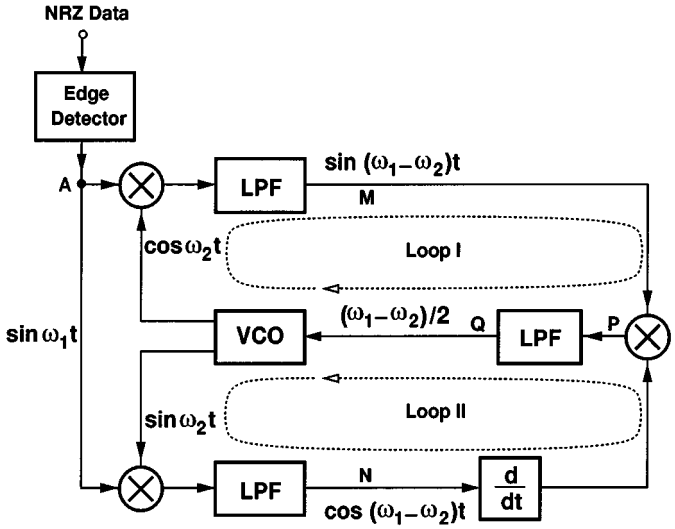

Fig. 4. Simplified quadricorrelator architecture.

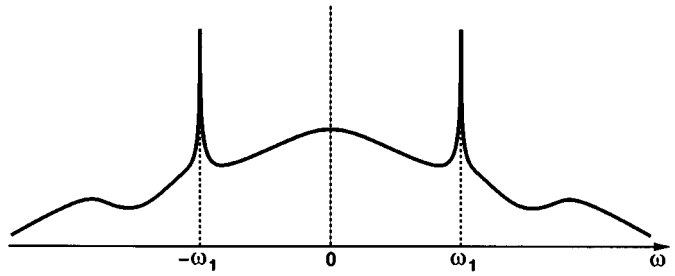

(a)

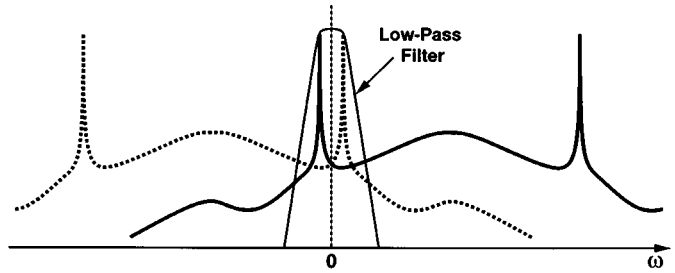

(b)

Fig. 5. Spectrum of edge-detected NRZ data (a) before and (b) after down-conversion and low-pass filtering.

In order to perform full-wave rectification, the current impulses are multiplied by the NRZ data using a Gilbert cell [Fig. 9(b)]. Note that when D_{in} goes high, a positive current impulse flows through Q_3 and Q_5, and a negative current impulse through Q_4 and Q_8. A similar phenomenon occurs when D_{in} goes low. In other words, Q_5 and Q_7 conduct the positive impulses to node X and Q_6 and Q_8 the negative impulses to node Y.

The rectified current impulses are mixed with the VCO output by means of a second Gilbert cell stacked upon the first [Fig. 9(c)]. To avoid saturating Q_5–Q_8, the common-mode level of the VCO output is set close to V_{CC} and that of D_{in} and $\overline{D_{in}}$ one V_{BE} lower. The low-pass filter simply consists of C_1, R_1, and R_2.

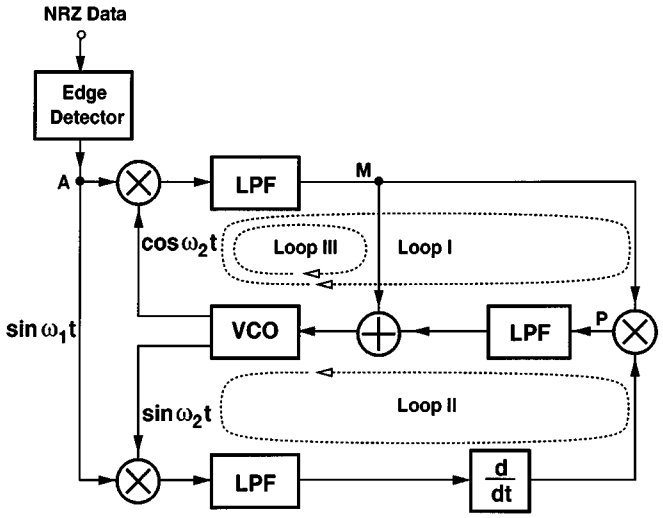

Fig. 6. Quadricorrelator with phase detection.

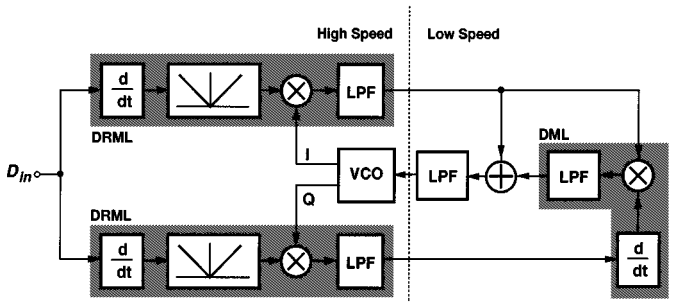

Fig. 7. Modified quadricorrelator.

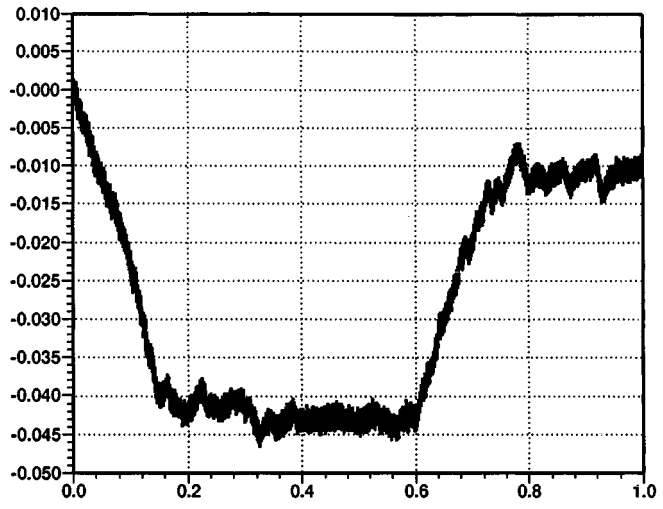

Fig. 8. Simulated capture behavior of the clock recovery circuit. (Horizontal and vertical units are μs and V, respectively.)

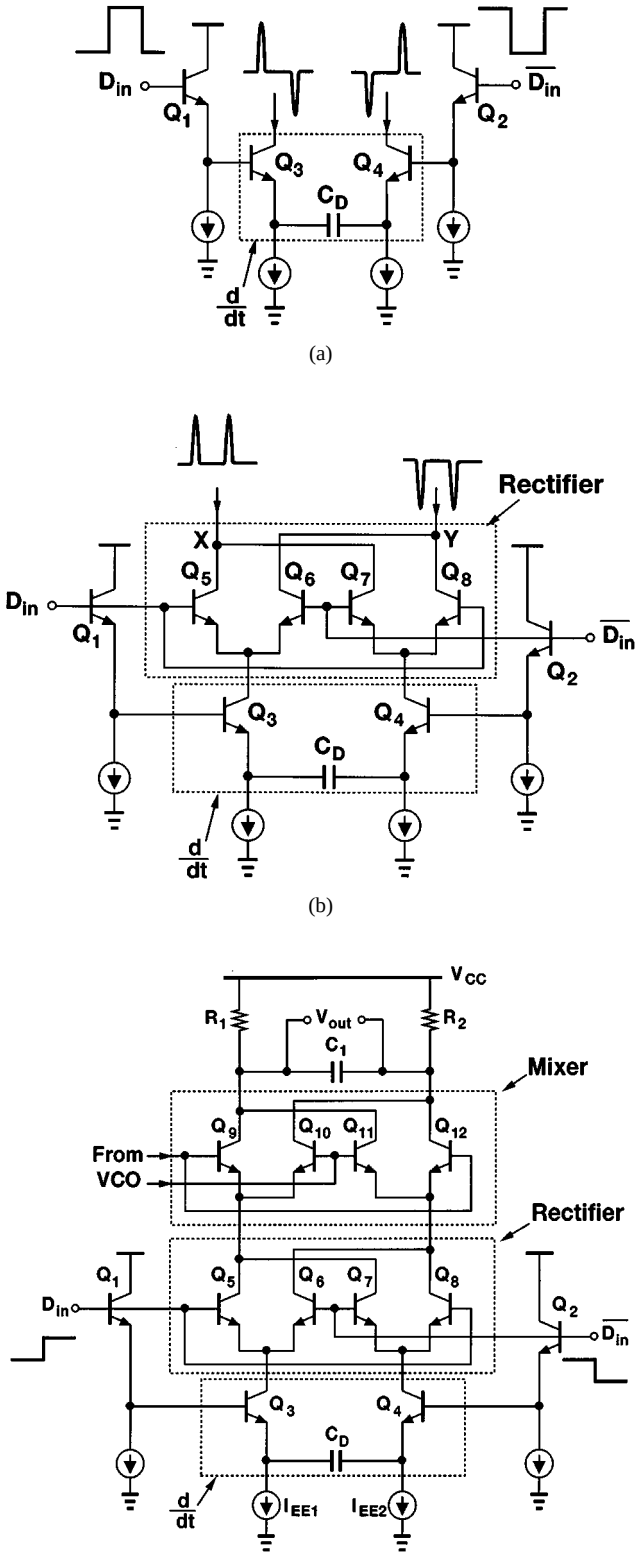

Fig. 9. Differentiator, rectifier, mixer, and LPF.

In the DRML circuit of Fig. 9, Q_1 and Q_2 are biased at 100 μA, $I_{EE1} = I_{EE2} = 200$ μA, $C_D = 60$ fF, $R_1 = R_2 = 1.5$ kΩ, and $C_1 = 200$ fF.

It is instructive to examine the output waveforms of the two DRML circuits in Fig. 7 when the VCO output frequency differs from the input bit rate. With the implementation of Fig. 9(c), the simulated waveforms shown in Fig. 10 are obtained. Here, the input rate is 2.5 Gb/s and the VCO frequency equals 2.7 GHz. Note that the outputs are approximately in quadrature, a behavior similar to that of the digital implementation in [9].

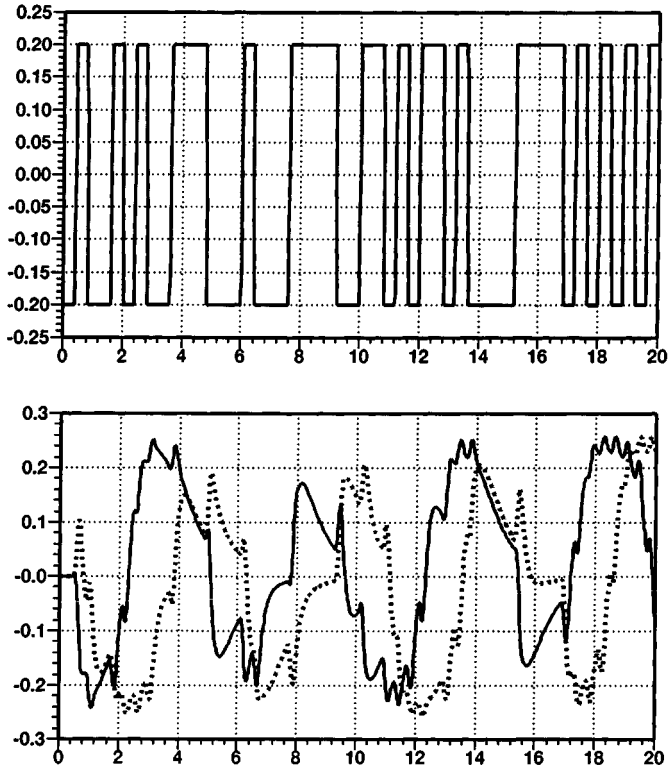

Fig. 10. Simulated input (top window) and outputs (bottom window) of two DRML circuits driven by quadrature VCO signals. (Horizontal and vertical units are μs and V, respectively.)

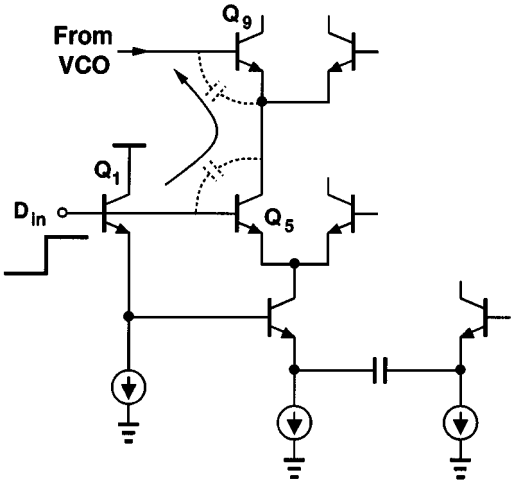

Fig. 11. Feedthrough from D_{in} to VCO.

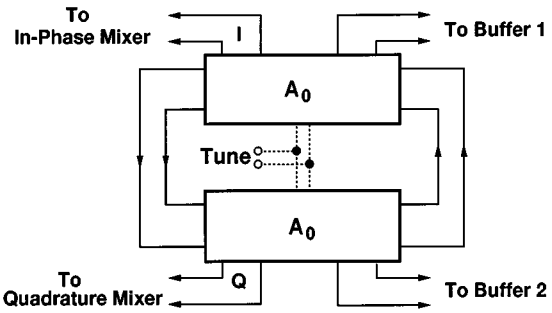

Fig. 12. VCO block diagram.

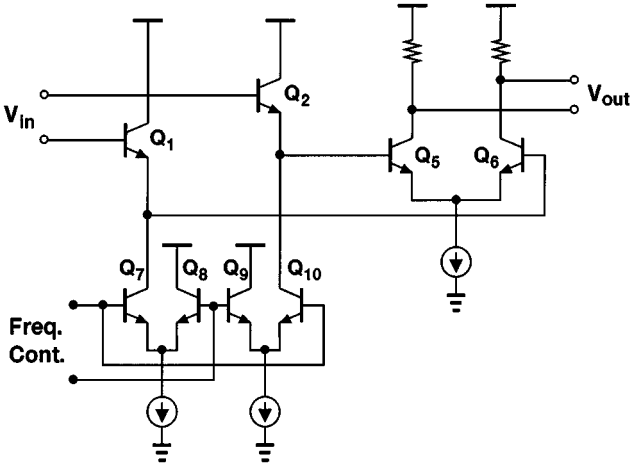

Fig. 13. Simple gain stage.

The interface between the VCO and the DRML stage entails an important issue: the feedthrough of the input signal to the VCO circuit. As illustrated in Fig. 11, data transitions at the input propagate through the base-collector capacitance of Q_5 and base-emitter capacitance of Q_9, thus disturbing the oscillation in the VCO and introducing input-dependent jitter. While differential operation lessens this effect to some extent, capacitance nonlinearities still cause significant feedthrough. Therefore, the oscillation path of the VCO must be sufficiently isolated from this disturbance.

B. VCO

The VCO is a critical part of the clock recovery circuit. Trade-offs among speed, jitter, and power dissipation make the design of this block especially difficult.

The VCO is configured as a ring oscillator with an even number of stages so as to provide quadrature outputs. Since a four-stage oscillator does not achieve the required speed in our technology, we utilize only two stages (Fig. 12). Note that each stage must drive three other circuits: the other stage in the ring, a mixer, and an output buffer. (Both stages are loaded with buffers to maintain symmetry.) In the VCO of Fig. 12, these outputs are taken from different ports to distribute the loading as well as ensure isolation of the VCO from the feedthrough noise.

To arrive at a suitable gain stage for the VCO, we begin with the simple variable-delay stage shown in Fig. 13. A VCO using this circuit exhibits a limited tuning range because the gain (at 2.5 GHz) drops considerably as the bias current of the emitter followers is reduced so as to decrease the frequency of oscillation. To resolve this issue, the circuit is modified as shown in Fig. 14(a), where cross-coupled devices Q_3-Q_4 are interposed between the emitter followers and the differential pair. The role of Q_3 and Q_4 can be understood with the aid of Fig. 14(b), where C_{in} represents the input capacitance of Q_5-Q_6 and C_j denotes the junction capacitance due to Q_7 and Q_{10}. At 2.5 GHz, these capacitances provide a relatively

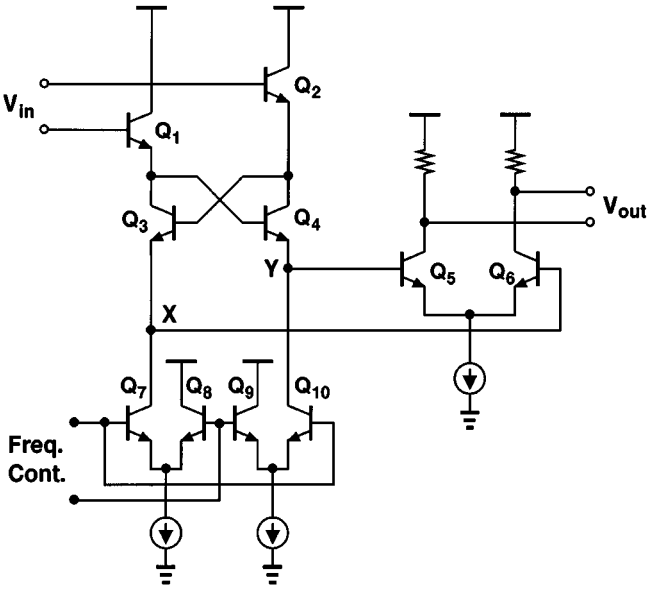

(a)

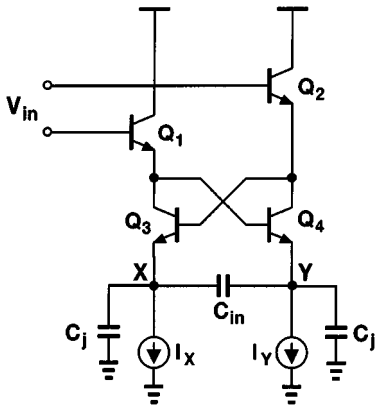

(b)

Fig. 14. (a) Gain stage with cross-coupled pair and (b) input branch with equivalent capacitances.

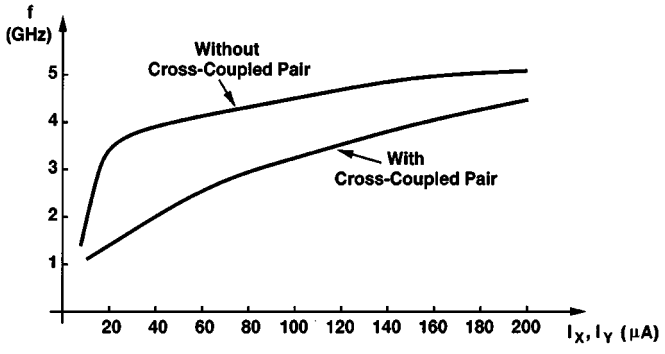

Fig. 15. Simulated VCO frequency with and without cross-coupled pair.

low equivalent impedance from X and Y to ground, allowing the positive feedback around Q_3-Q_4 to contribute significant gain (even at low bias currents) and phase shift. Thus, the VCO output frequency can be varied over a wide range with little degradation in the amplitude. Plotted in Fig. 15 is the simulated VCO frequency as a function of I_X and I_Y, indicating a smoother characteristic with the cross-coupled pair present.

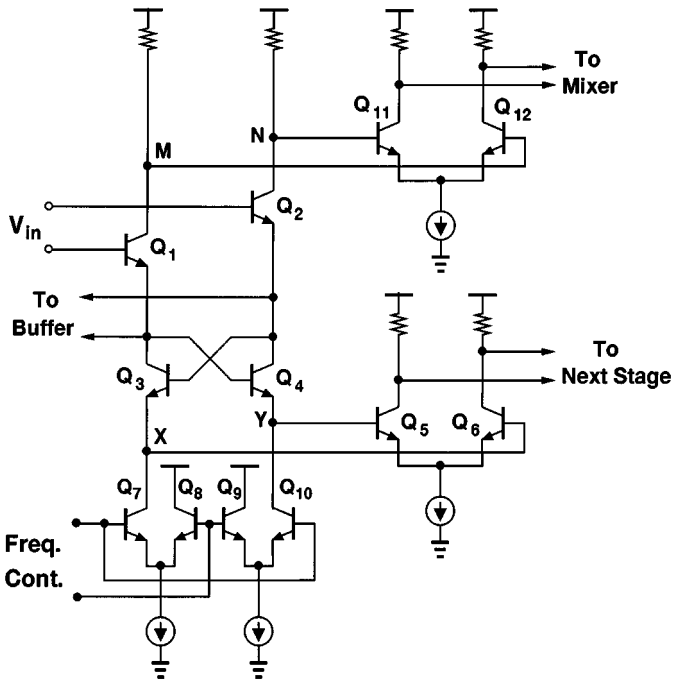

Fig. 16. Gain stage with reverse isolation.

In order to suppress the effect of data feedthrough noise, the gain stage is further modified as shown in Fig. 16, where the signal driving each mixer is taken from the *collectors* of Q_1 and Q_2 and buffered by the pair Q_{11}-Q_{12}. At 2.5 GHz, the voltage gain from V_{in} to nodes M and N is approximately equal to the total capacitance seen at X and Y divided by that at M and N—roughly unity. Thus, the voltage swings at the collectors of Q_1 and Q_2 are nearly equal to those in the main oscillation path while these transistors provide substantial reverse isolation. Note that the data feedthrough arriving from the mixer must go through the base-collector capacitance of Q_{11}-Q_{12} and Q_1-Q_2 before it disturbs the oscillation path. Simulations indicate that without these precautions, input feedthrough can introduce as much as 8 ps of jitter in the recovered clock.

The port at the emitter of Q_1 and Q_2 provides a relatively low output impedance and proper common-mode level for the output buffer.

In the gain stage of Fig. 16, the tail currents of Q_7-Q_8 and Q_9-Q_{10} are 200 μA and those of Q_5-Q_6 and Q_{11}-Q_{12} are 400 μA. Two additional fixed currents of 75 μA (not shown in Fig. 16) are drawn from nodes X and Y to avoid "starving" Q_1-Q_4 during loop transients. All the resistors are equal to 1 kΩ.

While the current sources in the present VCO are replicas of a constant (external) current source, in practice it may be advantageous to make the currents depend on the temperature in such a way that the frequency of oscillation becomes relatively insensitive to temperature variations. Typical bandgap techniques can be used to create such dependency.

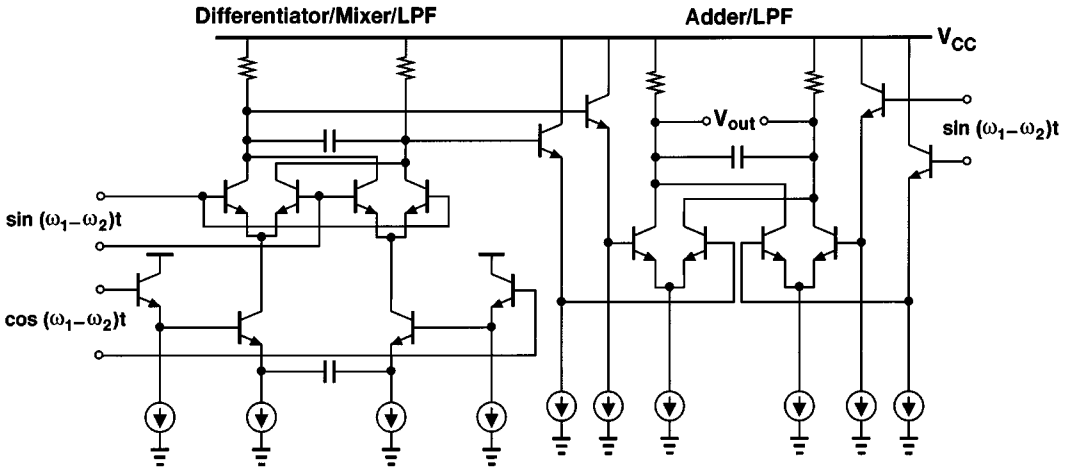

Fig. 17. Implementation of DML and adder of Fig. 7.

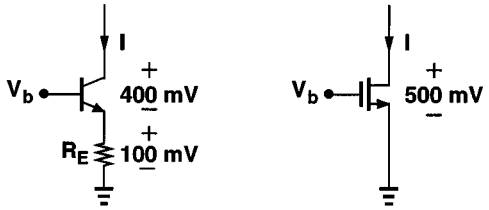

Fig. 18. Bipolar and MOS current sources.

C. DML and Adder

The DML and adder circuits of Fig. 7 are realized as depicted in Fig. 17. The DML stage is similar to the DRML of Fig. 9(c), but with the top mixer removed, while the adder sums the output currents of two differential pairs. Since these circuits operate at frequencies of only a few hundred megahertz, they dissipate a small percentage of the overall power.

D. MOS and Bipolar Current Sources

All the current sources in the clock recovery circuit are implemented with NMOS devices. The choice between MOS and (degenerated) bipolar current sources depends on the trade-offs among three parameters: the associated capacitance, the required voltage headroom, and the noise current. For a typical current of 200 μA, we use an NMOSFET with $W/L = 16\ \mu\text{m}/1\ \mu\text{m}$ to limit the required headroom to 0.5 V. The capacitance introduced by this device in the signal path is approximately equal to that of a minimum-size bipolar transistor in our technology.

The relationship between voltage headroom and noise current can be derived with the aid of Fig. 18. If the base resistance is neglected, the total noise current of the bipolar current source per unit bandwidth is

$$\overline{I_{n,bi}^2} = 2kT\frac{1+2g_mR_E}{(1+g_mR_E)^2}g_m \tag{2}$$

$$\approx 0.74kT\frac{I}{V_T} \tag{3}$$

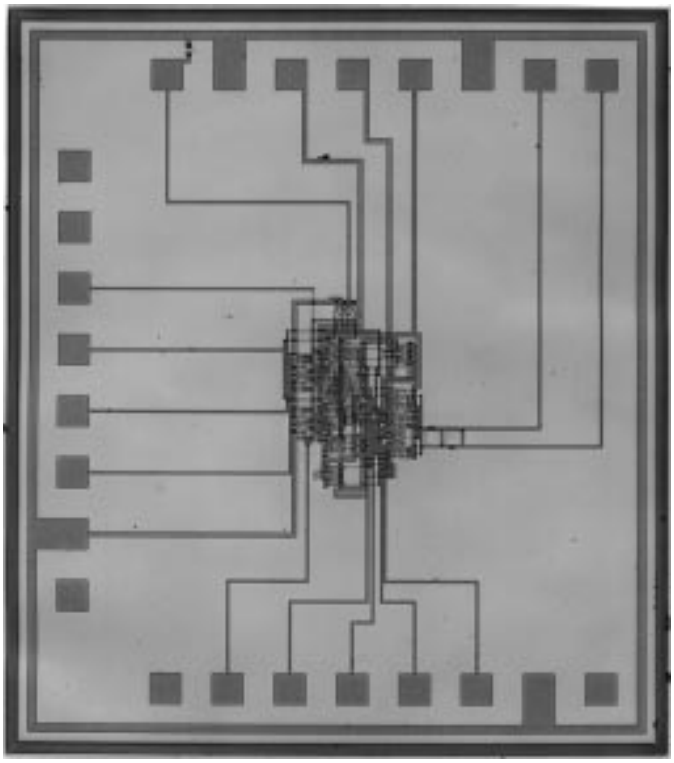

Fig. 19. CRC die photograph.

whereas for a square-law MOSFET

$$\overline{I_{n,MOS}^2} = 4kT\frac{2}{3}g_m \tag{4}$$

$$= 4kT\frac{4}{3}\frac{I}{V_{GS}-V_{TH}}. \tag{5}$$

If the MOS device is to remain in saturation for $V_{DSmin} = 500$ mV, then $V_{GS} - V_{TH} = 500$ mV$\approx 19.2V_T$. Therefore

$$\overline{I_{n,MOS}^2} \approx 0.28kT\frac{I}{V_T}. \tag{6}$$

Thus, for the current and voltage levels used here, MOS current sources exhibit approximately 2.7 times less noise than their bipolar counterparts while introducing a comparable

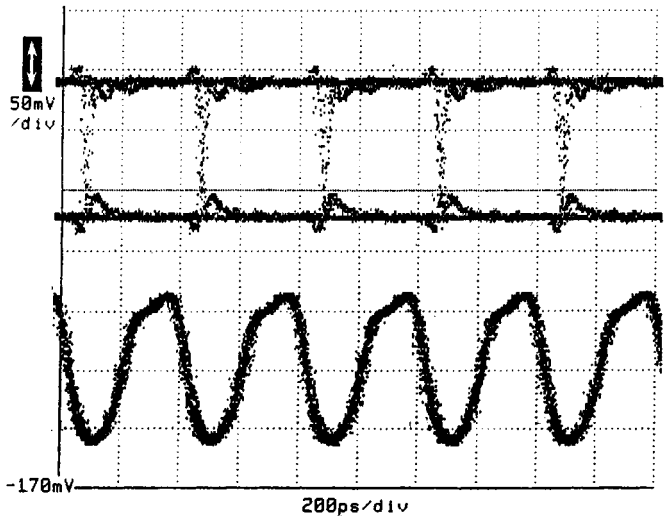

(a)

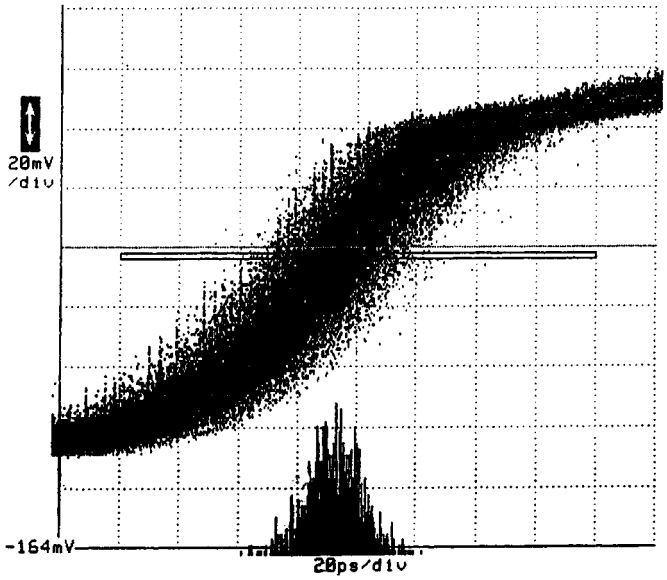

(b)

Fig. 20. Measured waveforms in time domain: (a) input eye diagram and recovered clock and (b) jitter histogram of the clock.

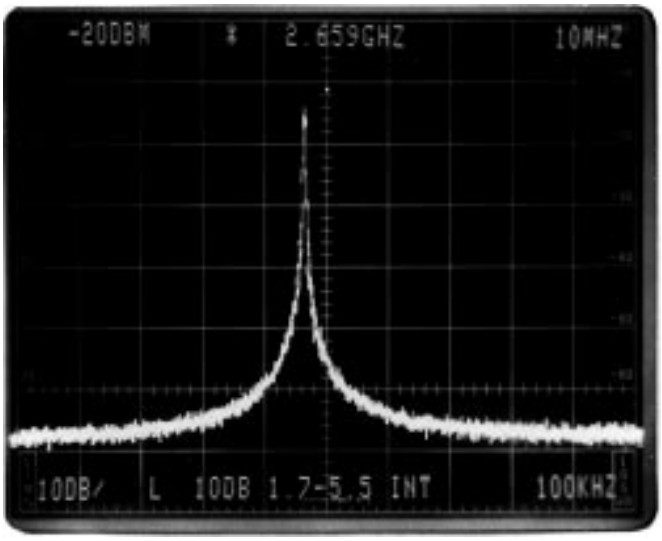

(a)

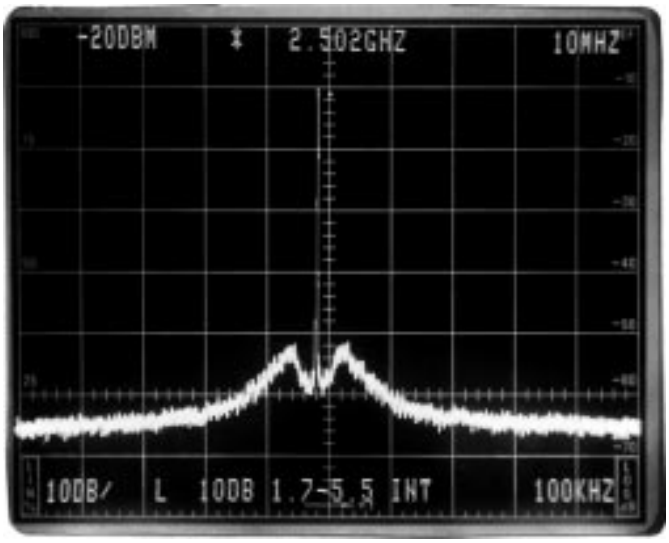

(b)

Fig. 21. Output spectrum (a) with free-running VCO and (b) locked to 2.5-Gb/s data. (Horizontal 10 MHz/div., vertical 10 dB/div.)

parasitic capacitance. This is particularly beneficial in the VCO, where noise in the current sources contributes to jitter. While MOS devices exhibit much higher $1/f$ noise than do bipolar transistors, low-frequency jitter components are suppressed within the bandwidth of Loop III in Fig. 6. Note that since the bipolar current source of Fig. 18 is very close to deep saturation, it may slow down the circuit during transients. The MOS device, on the other hand, is quite benign in this respect.

V. Experimental Results

The clock recovery circuit has been fabricated in a 20-GHz 1-μm BiCMOS technology [13]. Even though the circuit can be implemented in a pure bipolar process with minor modifications, the use of BiCMOS allows for higher levels of integration if additional low-speed processing is to be performed in CMOS. Fig. 19 shows a photograph of the die, whose active area is approximately 500 $\times$ 500 μm^2. Tested on wafer using high-speed Cascade probes, the entire CRC (excluding the input and output buffers) dissipates 15 mW with a 3-V supply. Of this power, approximately 4 mW is

TABLE I
Performance of the Clock Recovery Circuit

Center Frequency	**2.5 Gb/sec**
Tracking Range	**800 MHz**
Capture Range	**300 MHz**
Jitter	**9.5 psec, rms**
Power Dissipation	**15 mW**
Supply Voltage	**3 V**
Technology	**20-GHz, 1-μm BiCMOS**

consumed by the two DRML circuits and 9 mW by the VCO. For on-chip circuitry following the CRC, the output buffers need not drive a 50-Ω load and can be implemented with a pair of emitter followers dissipating an additional 1.5 mW.

Shown in Fig. 20 are the circuit's measured waveforms in the time domain in response to a pseudorandom binary sequence at 2.5 Gb/s with length $2^{20} - 1$. In Fig. 20(a), the top trace is the eye diagram of the input and the bottom trace the recovered clock. In Fig. 20(b), the jitter histogram of the clock is shown, indicating rms and peak-to-peak values of 9.5 ps and 60 ps, respectively. The CRC achieves a capture range of 300 MHz and a tracking range of 800 MHz.

The output has also been examined in the frequency domain. Displayed in Fig. 21(a) is the output spectrum of the free-running VCO, exhibiting a phase noise of -80 dBc/Hz at 5-MHz offset. Fig. 21(b) shows the output spectrum when the CRC is locked to a 2.5-Gb/s pseudorandom sequence. The phase noise is suppressed within the bandwidth of the loop, taking on its natural shape for offsets greater than 5 MHz.

Table I summarizes the performance of the clock recovery circuit.

VI. Conclusion

All-analog implementations of the quadricorrelator architecture can perform high-speed clock recovery with low power dissipation. Combining operations such as differentiation, rectification, and mixing in a compact stage, a modified version of this architecture recovers the clock from 2.5-Gb/s NRZ data while dissipating 15 mW. The circuit also employs a two-stage ring oscillator with additional phase shift and isolation techniques to ensure reliable operation across a wide frequency range as well as to minimize the data transition feedthrough noise.

References

[1] 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.
[2] D. Richman, "Color-carrier reference phase synchronization accuracy in NTSC color television," *Proc. IRE*, vol. 42, pp. 106–133, Jan. 1954.
[3] J. A. Bellisio, "A new phase-locked loop timing recovery method for digital regenerators," in *IEEE Int. Conf. Rec.*, vol. 1, June 1976, pp. 10–17.
[4] H. Ransijn and P. O'Connor, "A PLL-based 2.5-Gb/s clock and data regenerator IC," *IEEE J. Solid-State Circuits*, vol. 26, pp. 1337–1344, Oct. 1991.
[5] F. B. McCormick *et al.*,"Five-stage free-space optical switching network with field-effect transistor self-electrooptic-effect-device smart-pixel arrays," *Applied Optics*, vol. 33, pp. 1601–1618, Mar. 10, 1994.
[6] S. K. Shanmugam, *Digital and Analog Communication Systems.* New York: Wiley, 1979.
[7] B. Razavi, "Design of monolithic phase-locked loops and clock recovery circuits—a tutorial," in *Monolithic Phase-Locked Loops and Clock Recovery Circuits*, B. Razavi, Ed. Piscataway, NJ: IEEE Press, 1996.
[8] R. R. Cordell *et al.*, "A 50 MHz phase- and frequency-locked loop," *IEEE J. Solid-State Circuits*, vol. SC-14, pp. 1003–1009, Dec. 1979.
[9] A. Pottbacker and U. Langmann, "An 8 GHz silicon bipolar clock recovery and data regenerator IC," *IEEE J. Solid-State Circuits*, vol. 29, pp. 1572–1576, Dec. 1994.
[10] I. A. Young, J. K. Greason, and K. L. Wong, "A PLL Clock Generator with 5 to 110 MHz of Lock Range for Microprocessors," *IEEE J. Solid-State Circuits*, vol. 27, pp. 1599–1607, Nov. 1992.
[11] D. Wolaver, *Phase-Locked Loop Circuit Design.* Englewood Cliffs, NJ: Prentice-Hall, 1991.
[12] D. G. Messerschmitt, "Frequency detectors for PLL acquisition in timing and carrier recovery," *IEEE Trans. Commun.*, vol. COM-27, pp. 1288–1295, Sept. 1979.
[13] J. Sung *et al.*, "BEST2—A high performance super self-aligned 3V/5V BiCMOS technology with extremely low parasitics for low-power mixed-signal applications," in *Proc. IEEE CICC*, May 1994, pp. 15–18.

Behzad Razavi (S'87–M'90), for a photograph and biography, see p. 343 of the March issue of this JOURNAL.