Phase-Locking in Wireline Systems: Present and Future[1]

Behzad Razavi

Electrical Engineering Department

University of California, Los Angeles

Abstract

This paper describes the challenges in the design of phase-locked loops and clock and data recovery circuits as speeds approach 80-100 Gb/s. Skew and jitter issues are presented and the effect of reference phase noise, charge pump noise, reference spurs, and loop filter leakage is quantified. The phase noise performance of cascaded loops is analyzed and two new architectures are proposed.

I. INTRODUCTION

Our appetite for higher data rates remains unabated. The ever-growing volume of data in wireless and wireline links demands a proportional increase in the aggregate or serial throughput rates. At the physical layer, other challenges such as distortion of data, power dissipation, and packaging also manifest themselves as higher rates are sought.

This paper deals with phase-locking functions in high-speed wireline transceivers, emphasizing new developments and predicting possible trends for future systems. While modern transceivers employ many other functions as well, phase-locking has not only proved a difficult bottleneck but extended its reach into other building blocks, thus claiming its own place in the top design challenges.

Section II reviews present trends in the field. Section III describes the challenges in transmit (TX) phase-locked loop (PLL) design, including speed and jitter issues and the use of cascaded loops. Section IV deals with clock and data recovery (CDR) circuits, and Section V points to future trends.

II. PRESENT TRENDS

Recent work in phase-locking for wireline systems has entailed a number of interesting trends that are likely to intensify in the future.

(1) Use of RF Design Concepts. The vast research into the problem of RF synthesis has greatly benefited wireline transceivers as well. Examples include the use of on-chip inductors, the design of low-noise oscillators and frequency dividers, methods of phase noise and spur reduction in phase-locked loops, and the use of injection-locked oscillators for clock distribution [1].

(2) Convergence of Functions. The need for global optimization of the performance has discouraged the design of building blocks as independent modules. For example, the transmit PLL and the receive (RX) CDR circuits are now viewed as one entity [2], particularly to avoid mutual injection pulling. Also, the equalization, eye opening, and CDR functions have become intertwined [3, 4].

(3) Convergence of PLLs and DLLs. The potentially lower jitter of delay-locked loops and the synthesis capabilities of PLLs have motivated work on combining the two [5, 6].

(4) Use of Phase Interpolation. The desire to move toward "digital" PLLs and to avoid oscillators within the CDR circuit has translated to extensive phase interpolation, albeit at the cost of routing, mismatch, and jitter issues.

III. TRANSMIT PLL

The TX PLL generates a full-rate clock whose integer sub-multiples drive the multiplexer (MUX) chain. More importantly, it drives a retiming flip flop (FF) in the data path [Fig. 1(a)] so as to remove the output jitter of the MUX. This jitter arises from two sources, namely, the propagation mismatches within the MUX and the duty cycle error at the output of the $\div 2$ stage. Note that duty cycle errors in the full-rate clock are unimportant. It is the need for this retiming that makes the design of the TX PLL and its surrounding circuitry difficult.

A critical issue in the architecture of Fig. 1(a) stems from the delay of the $\div 2$ circuit, ΔT. As shown in Fig. 1(b), this delay displaces the edges of the MUX output, causing sampling closer to the data zero crossings. This effect can be suppressed by inserting an equal delay in the clock path of the FF, but the required full-rate bandwidth complicates the design of such a delay stage. We return to this issue in Section III.D.

Another important issue in the TX of Fig. 1(a) originates from the coupling of the data transitions from the D input of the FF to its clock input. Exemplified by the capacitive path depicted in Fig. 2, this effect heavily corrupts the phase of the oscillator unless a buffer with high reverse isolation is interposed between the PLL and the FF. Unfortunately, such a buffer must incorporate inductors at high frequencies, thereby complicating the routing of the signals.

In high-speed transmitters, the PLL and data path designs are intimately related. As illustrated in Fig. 3, the TX design begins with the output driver (possibly including equalization),
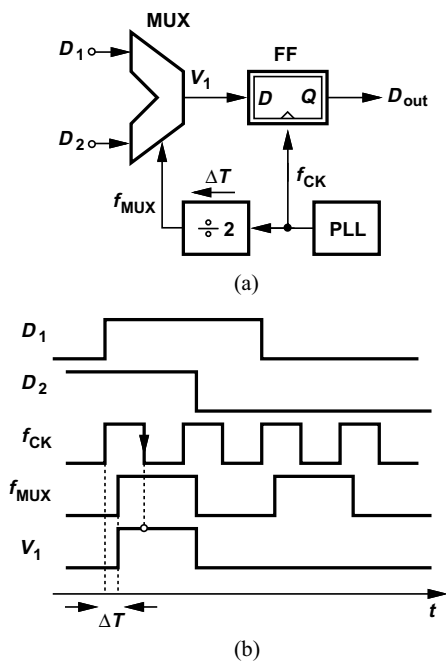
(a)

(b)

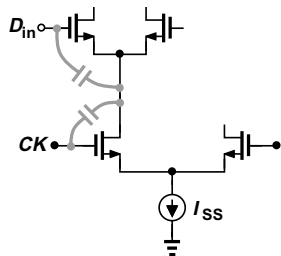Fig. 1. (a) Interface between PLL and data path, (b) effect of divider skew.

Fig. 2. Coupling of data to clock through FF capacitances.

aiming to deliver the necessary output swing to the load (a laser or a transmission line) and an on-chip back-termination resistor. Since bandwidth requirements limit the number of stages used in the driver [7], its input capacitance, C_{dr}, tends to be large, thus dictating a flip flop with high drive capability and hence high currents. To accommodate such currents with a limited voltage headroom, the FF itself employs wide transistors and exhibits large input capacitances, C_D and C_{CK}. This in turn necessitates high currents in the MUX and the VCO
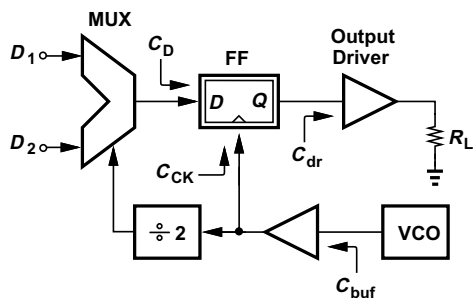

Fig. 3. Propagation of driver capacitance to VCO.

buffer, creating large capacitive loads for the VCO and the $\div 2$ stage. In other words, the output driver's input capacitance "propagates" to the VCO.

Due to the propagation of the output driver's input capacitance, the VCO must typically drive a large capacitance itself even if the path through its buffer and the FF may employ some tapering. At frequencies approaching 100 GHz, the design of such a VCO becomes exceedingly difficult as it requires very small inductors, e.g., around 50 pH.

Figure 4(a) shows an oscillator topology that exhibits poten-

(a)

(b)

Fig. 4. (a) Oscillator based on inductive feedback, (b) basic amplifier circuit.
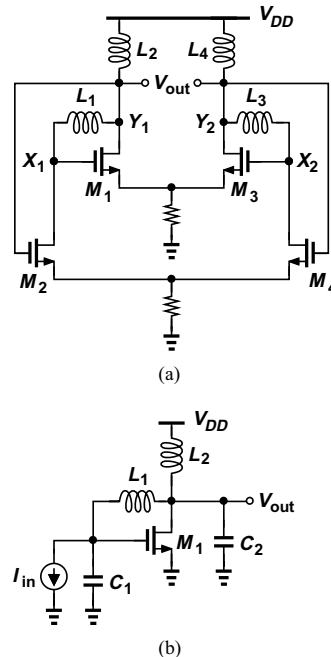
tial for high-speed operation [8]. The circuit is derived from the single-ended transimpedance amplifier illustrated in Fig. 4(b), whose transfer function exhibits two imaginary poles at

$$\omega_{p2,p2'} = \pm j\sqrt{\frac{3+\sqrt{5}}{2LC}}. \tag{1}$$

The magnitude of this pole is about 0.62% higher than that of a second-order LC tank, allowing operation at high speeds. For example, oscillation frequencies in the range of 80 GHz to 128 GHz have been achieved in 90-nm CMOS technology [8]. To tune the frequency, varactors must be tied from nodes X_1, X_2, Y_1, and Y_2 to the control voltage. By virtue of inductive feedback, the circuit can drive heavy capacitive loads and operate from a low supply voltage.

Figure 5 compares the simulated phase noise of this oscillator with that of a standard cross-coupled topology at 80 GHz, assuming a given inductor design, a given power consumption, and a given buffer. Interestingly, the inductive feedback suppresses the flicker noise contribution of M_1 and M_3 to the phase noise because a low-frequency voltage perturbation in
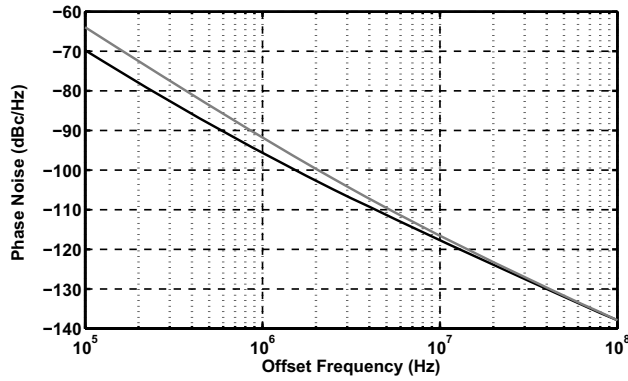
Fig. 5. Simulated phase noise of the inductive-feedback oscillator (black line) and cross-coupled oscillator (gray line).

Fig. 6. Miller divider based on inductive-feedback amplifier.

series with the gate of, say, M_1, cannot change the phase difference between V_{X1} and V_{Y1}. Note that the oscillator of Fig. 4(a) is loaded with much greater capacitance so that it operates at the same frequency as the cross-coupled topology.

Even with the topology of Fig. 4(a), the design of VCOs at frequencies approaching 100 GHz faces other critical challenges. First, the quality factor, Q, of inductors does not scale linearly with frequency, beginning to saturate above 50 GHz. For example, [10] reports a Q of 12 for 180-pH inductors at 60 GHz, and [11] a Q of 17 for 400-pH inductors at 50 GHz. Second, the Q of varactors is likely to fall *below* the Q of inductors at these frequencies. Both effects exacerbate the trade-offs among phase noise, power dissipation, tuning range, and output swings.

A. Frequency Dividers

High-speed frequency dividers pose another serious challenge to the design of PLLs. In addition to speed, other important parameters of dividers include the minimum required input voltage swing, the input capacitance, the output drive capability, and the complexity, especially the number of inductors required in each topology. Static (flip-flop-based) topologies suffer from a limited speed and injection-locked dividers (ILDs) exhibit a limited lock range, placing the overall PLL design at risk. This is because, due to modeling inaccuracies and process variations, the lock range of ILDs may not enclose the desired frequency, thus causing lock failure or false lock [9]. Two other topologies, namely, the Miller divider and the "heterodyne PLL" [9] can achieve a wider lock range at the cost of greater complexity.

The circuit technique illustrated in Fig. 4(b) can also improve the speed of frequency dividers. Shown in Fig. 6 is a Miller topology employing the inductive feedback configuration [8]. The cross-coupled pair increases the loop gain and hence the lock range. Also, M_1 and M_2 form a differential "sampling mixer," which presents less loading to the amplifier than conventional double-balanced passive mixers. Specifically, the capacitance at node P switches periodically between X and Y in a conventional mixer, thereby introducing a re-
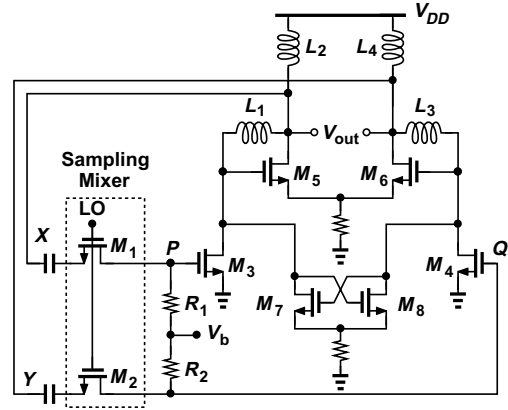
sistance between these two nodes and lowering the gain of the amplifier. Here, on the other hand, the voltage is simply stored on the capacitance for a half cycle (if R_1 and R_2 are sufficiently large).

The circuit of Fig. 6 achieves high speeds even in 90-nm CMOS technology. For example, one choice of inductor values provides a lock range of 88 to 104 GHz.

Heterodyne phase-locking is another candidate for high-speed dividers. Depicted in Fig. 7 in its simplest form, a
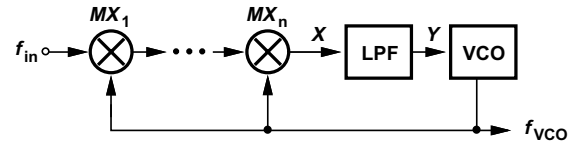

Fig. 7. Basic heterodyne PLL.

heterodyne PLL mixes the input with the VCO output n times, generating a frequency component at X given by $f_{in} - nf_{VCO}$. If the circuit locks, this component must be equal to zero, and f_{VCO} equal to f_{in}/n. Other divide ratios can be realized by inserting dividers in the feedback loop and/or at the input ports of the mixer [9]. A prototype realized in 0.13-μm CMOS technology operates from 64 to 70 GHz while consuming 6.5 mW.

The use of consecutive mixers in a heterodyne PLL raises the possibility of false lock due to unwanted mixing products. However, it can be shown that for divide ratios up to 4, the limited VCO tuning range prohibits false lock.

B. Jitter Issues

The TX PLL produces the dominant jitter in the transmitted data if the flip flop and output driver in the data path exhibit sufficient bandwidth. Jitter becomes much more pronounced as we approach 80-100 GHz because (1) the Q of inductors begins to saturate and the Q of varactors is likely to be even lower; (2) the very large frequency multiplication factor realized by the PLL dramatically amplifies the reference phase noise, S_{REF}. The PLL loop bandwidth must therefore be chosen carefully.

Reference Phase Noise The choice of the loop bandwidth is governed by the available frequencies, phase noise, and cost of crystal oscillators. Since low-noise, low-cost crystal oscillators typically operate at frequencies no higher than 100 MHz, we assume a PLL multiplication factor, N, of roughly 1000 and hence a 60-dB amplification of the reference phase noise within the loop bandwidth. As shown in Fig. 8, a natural
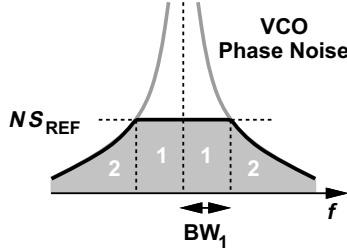

Fig. 8. Effect of reference phase noise.

choice for the loop bandwidth, BW_1, is given by the intersection of the amplified reference phase noise and that of the free-running VCO. For example, a 100-MHz crystal oscillator displays a constant phase noise of about -150 dBc/Hz beyond 100-kHz offset, suggesting that we seek the offset frequency at which the VCO phase noise falls to -90 dBc/Hz. Assuming a $1/f^2$ roll-off beyond the loop bandwidth, noting that regions 1 and 2 have equal areas, and integrating the phase noise we obtain the rms jitter as

$$\text{Jitter} = \frac{\sqrt{4NS_{REF} \cdot BW_1}}{2\pi} T_{CK}, \qquad (2)$$

where the factor of 4 accounts for the areas in regions 1 and 2 on both sides of the carrier, and T_{CK} denotes the carrier period. If $NS_{REF} = -90$ dBc/Hz and $BW_1 = 2$ MHz, then the rms jitter is equal to 1.42% of the clock period, a marginally acceptable value.

An interesting point that arises here relates to the frequency of jitter. As observed in the above calculations, most of the jitter originates from offsets up to a few tens of megahertz. For date rates of 80 to 100 Gb/s, this low-frequency jitter is readily removed by the RX CDR circuit and, therefore, proves unimportant. However, the transmit data mask may still be violated.

Charge Pump Noise The large multiplication factor also raises concern regarding the charge pump (CP) noise. We introduce an analysis here to determine the CP-induced phase noise at the output within the loop bandwidth.

First, suppose the Up and Down currents in the charge pump exhibit a mean value of I_P and a static mismatch of ΔI. It can be shown that such a mismatch gives rise to an input static phase error of

$$\Delta T = \frac{\Delta I}{I_P} T_{RST}, \qquad (3)$$

where T_{RST} denotes the width of the Up and Down pulses in the locked condition (approximately equal to five gate delays).

If expressed in picoseconds—rather than in radians—this error appears at the PLL output without multiplication. To convert the output error to a phase quantity, we normalize ΔT to T_{CK} and multiply the result by 2π:

$$\Delta\phi_{out} = 2\pi \frac{\Delta I}{I_P} \cdot \frac{T_{RST}}{T_{CK}}. \qquad (4)$$

Next, let us consider the noise of each current source, $\overline{I_n^2}$, as a mismatch between the two. Since the noise powers of the two current sources add, we have, within the loop bandwidth,

$$\overline{\Delta\phi_{out}^2} = 4\pi^2 \frac{2\overline{I_n^2}}{I_P^2} \frac{T_{RST}^2}{T_{CK}^2}. \qquad (5)$$

In addition, the harmonics of the Up and Down pulses downconvert high-frequency noise to baseband. Since the harmonics have a sinc envelope that crosses zero at $1/T_{RST}$, we assume that the number of harmonics is equal to $(1/T_{RST})/f_{REF}$ and they have roughly equal amplitudes. Since each harmonic downconverts two noise sidebands and since all of the sidebands are uncorrelated, we must multiply (5) by $2(1/T_{RST})/f_{REF}$:

$$\overline{\Delta\phi_{out,tot}^2} = 4\pi^2 \frac{2\overline{I_n^2}}{I_P^2} \frac{2NT_{RST}}{T_{CK}}, \qquad (6)$$

where N denotes the PLL multiplication factor. For thermal noise of a MOSFET, $\overline{I_n^2} = 4kT\gamma g_m = 4kT\gamma(2I_P)/(V_{GS} - V_{TH})$, yielding

$$\overline{\Delta\phi_{out,tot}^2} = 4\pi^2 \frac{16kT\gamma}{(V_{GS} - V_{TH})I_P} \frac{2NT_{RST}}{T_{CK}}. \qquad (7)$$

For example, with $\gamma = 1$, $I_P = 1$ mA, $V_{GS} - V_{TH} = 100$ mV, $N = 1000$, $T_{RST} = 30$ ps, and $T_{CK} = 10$ ps, we obtain $\overline{\Delta\phi_{out}^2} = -98$ dBc/Hz.

Effect of Reference Spurs The trade-off between the loop bandwidth and the level of output reference spurs creates another constraint in the design. Assuming the first harmonic of the ripple on the control voltage is expressed as $V_m \cos\omega_{REF}t$, we write the output as

$$V_{out} = V_0 \cos\left(\omega_{CK}t + \frac{V_m K_{VCO}}{\omega_{REF}} \sin\omega_{REF}t\right). \qquad (8)$$

The zero crossings therefore deviate from their ideal points by a maximum of $\pm V_m K_{VCO}/\omega_{REF}$ radians, exhibiting a peak-to-peak jitter equal to

$$J_{pp} = \frac{1}{2\pi} \frac{2V_m K_{VCO}}{\omega_{REF}} T_{CK}. \qquad (9)$$

Since the relative magnitude of the sidebands in the output spectrum is given by $V_m K_{VCO}/(2\omega_{REF})$, we conclude that the relative jitter, J_{pp}/T_{CK}, and the relative sideband magnitude are nearly equal. For J_{pp}/T_{CK} to remain below 1%, the

sidebands must fall to -40 dBc, a relatively relaxed requirement.

Effect of Capacitor Leakage The gate leakage current has reached significant values in 45-nm technology. Plotted in Fig. 9 is the simulated leakage for a 10 μm/0.5 μm device
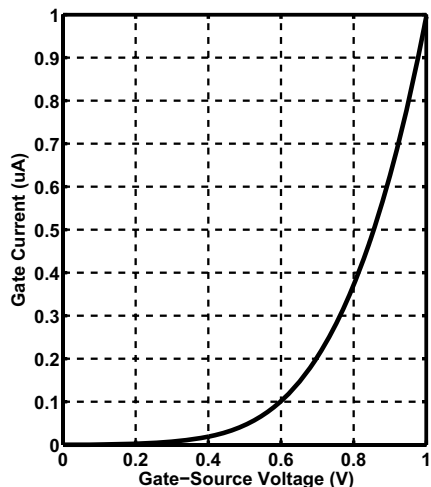

Fig. 9. Gate leakage in 45-nm technology.

with a gate dielectric thickness of 20 Å. (The source and drains are grounded.) Note that the strong dependence of the leakage on the gate-source voltage makes cancellation difficult.

The gate leakage readily manifests itself if the loop filter incorporates MOS capacitors. As illustrated in Fig. 10 [12], the leakage current I_G discharges the loop filter while the
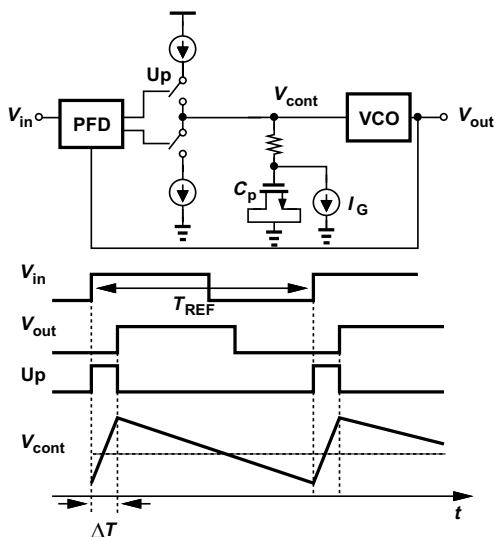

Fig. 10. Effect of gate leakage on PLL performance.

charge pump is off. In the steady stage, the PLL develops a phase offset, ΔT, during which the CP replenishes the charge drained by I_G. Thus, the peak-to-peak ripple on the control voltage is given by $(I_G/C_P)T_{REF}$, where it is assumed $\Delta T \ll T_{REF}$.

Interestingly, the "self-droop" rate I_G/C_P is independent of the MOS lateral dimensions and hence a constant of the technology, reaching 1.2 mV/ns for 45-nm devices at $V_{GS} = 0.6$ V. For example, if $T_{REF} = 10$ ns, then a ripple of 12 mV$_{pp}$ appears on the control voltage, yielding large sidebands at the VCO output. If the sidebands are to remain 40 dB below the carrier (as shown above), then $K_{VCO} < 520$ MHz, a very difficult condition to meet for a VCO running at 80-100 GHz.

C. Cascaded PLLs

The large multiplication factor required to translate f_{REF} to f_{CK} makes cascaded PLLs [13] a plausible alternative. Specifically, we seek the conditions under which such a cascade exhibits less jitter than a single PLL.

Consider the cascade shown in Fig. 11, where we assume the following are constant: f_{REF}, N_1N_2, the (free-running)
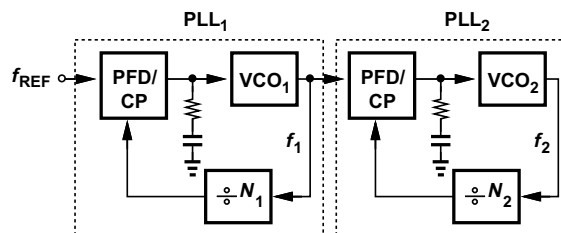

Fig. 11. Cascaded PLLs.

phase noise of VCO$_2$ (S_2), and the phase noise of the reference (S_{REF}). We seek the optimum choice of N_1, N_2, and the loop bandwidths of the two PLLs, BW$_1$ and BW$_2$. As we will see, the utility of cascaded PLLs directly depends on the phase noise of VCO$_1$ (S_1) relative to that of VCO$_2$. We denote the Q's of the two oscillators by Q_1 and Q_2.

We analyze three scenarios for the two VCOs. (1) The phase noise (at a given frequency offset) directly scales with frequency, $S_2 = N_2S_1$. This occurs if the Q remains relatively *constant* from f_1 to f_2. (2) The phase noise is constant, $S_2 = S_1$, requiring that $Q_2 = N_2Q_1$. (3) Due to tuning range limitations at f_2, VCO$_2$ is much more difficult to design than VCO$_1$, and $S_2 > N_2S_1$.

Figure 12(a) plots on a log scale the single-sideband (SSB) profiles of S_1 and S_2 for the first scenario. If BW$_1$ is chosen at the intersection of S_1 and N_1S_{REF} and BW$_1$ =BW$_2$, then the amplified phase noise of PLL$_1$ adds to the shaped phase noise of VCO$_2$, resulting in the "humps" shown in S_{out} and hence higher jitter than that of a single PLL (the gray curve). It can be shown that BW$_2$ >BW$_1$ or BW$_2$ <BW$_1$ yield even higher jitter. In other words, this scenario provides no jitter advantage over a single loop.

Figure 12(b) illustrates the second scenario, where the amplified phase noise of PLL$_1$ dominates, thereby producing a larger jitter than does the first scenario. This holds regardless of the choice of BW$_2$.

Shown in Fig. 12(c) is the third scenario. Here, the amplified phase noise of PLL$_1$ reaches $N_1N_2S_{REF}$ but extends only to BW$_1$. Thus, if BW$_2$ is maximized (subject to conditions such
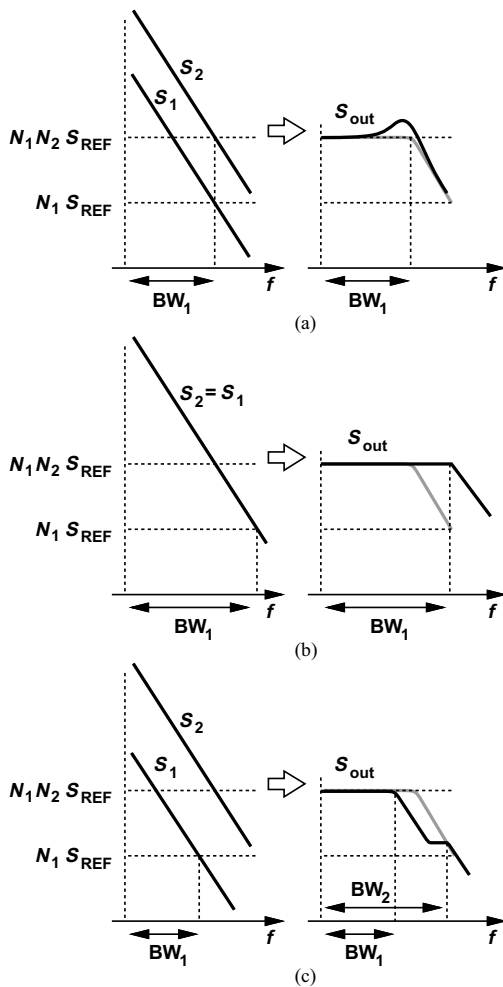
Fig. 12. Output phase noise of cascaded PLLs for (a) $S_2 = N_2 S_1$, (b) $S_2 = S_1$, and (c) $S_2 > N_2 S_1$.

as $BW_2 < 0.1 N_1 f_{REF}$), then the cascade exhibits less jitter than a single PLL.

D. A New Approach

We propose a TX/PLL interface that can relax many of the issues described above. Illustrated in Fig. 13(a), the idea is to employ a half-rate PLL and a frequency doubler to generate the full-rate clock. The twofold reduction in the required PLL speed greatly eases the design of the VCO and frequency dividers. More importantly, the architecture eliminates the troublesome divider delay depicted in Fig. 1. Instead, the delay through the doubler must match that through the MUX, a simpler task because they have the same polarity.

The proposed approach nonetheless entails two important issues. First, the doubler must avoid attenuation so that it only doubles the phase noise and generates large output swings that can directly clock the FF. This can be accomplished by means of the inductively-loaded symmetric XOR gate shown in Fig. 13(b) [14]. Second, duty-cycle distortion in the PLL
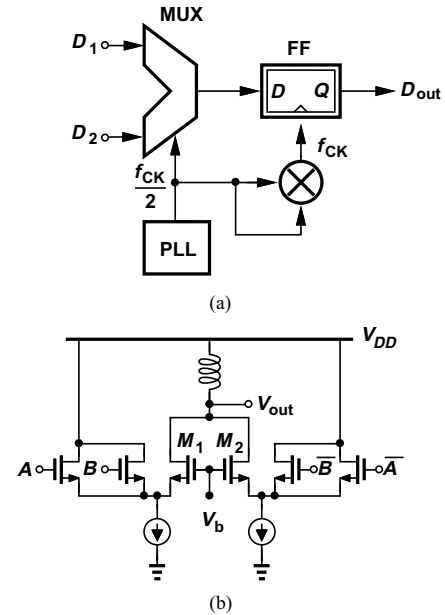

Fig. 13. (a) Proposed TX PLL interface, (b) doubler implementation as a low-voltage symmetric XOR .

output translates to displacement of every other falling (or rising) edge at the doubler output. Fortunately, the inductive (resonant) load in Fig. 13(b) reduces this effect by about 12 dB. Note also that the common-gate transistors M_1 and M_2 provide a high reverse isolation, suppressing the coupling of data transitions to the oscillator.

IV. RX CDR CIRCUIT

The CDR circuit presents its own challenges in receiver design, especially if it is integrated along with the transmitter. In addition to recovering the clock with a small jitter and retiming the data optimally, the circuit must remain immune to both the noise generated by the data edges at the TX output and the injection-pulling effected by the TX PLL. Arising from coupling through the substrate and the package, these two types of corruption directly determine the choice of the transceiver architecture.

A. Phase-Interpolating CDR Circuits

A class of CDR circuits dealing with these issues is based on shared PLLs with phase interpolation [2]. Illustrated in Fig. 14 [2], the idea is to perform phase comparison with the input data through the use of interpolated phases obtained from the TX PLL. The phase detector (PD) selects the interpolated edges so as to optimally sample the data (in the middle and at the zero crossings of the eye). In the presence of an offset between f_{VCO} and the data rate, the interpolated phases rotate at a rate equal to the offset. Thus, if the CDR loop bandwidth remains sufficiently greater than the frequency offset, the data is sampled properly.

The above architecture can employ discrete or continuous phase interpolation. With the former, the phase rotation in
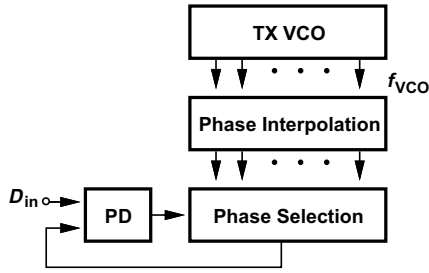
Fig. 14. CDR loop using interpolated phases of TX VCO.

the presence of a frequency offset occurs in discrete steps, failing to retime the data optimally. For this reason, a large number of finely-spaced edges must be produced, leading to a complex layout and mismatches and unwanted couplings in the routing. Continuous interpolation is therefore better suited to high-speed design as it requires only quadrature phases.

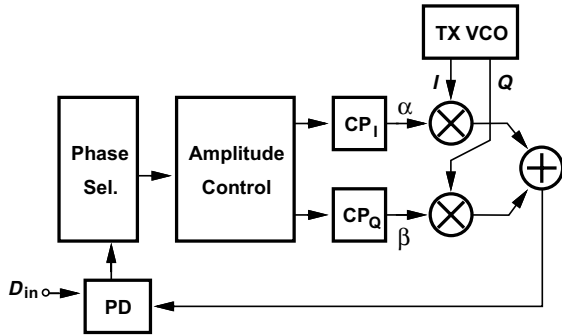Figure 15 shows an example of CDR with continuous phase

Fig. 15. CDR loop using continuous phase interpolation.

interpolation [2]. Here, quadrature phases provided by the TX PLL are summed with weights α and β and the result is applied to a bang-bang phase detector. From the phase error, the PD commands the two charge pumps CP_I and CP_Q to adjust α and β, respectively. The loop locks when every other edge of the interpolated clock samples the data at its zero crossings.

The CDR loop of Fig. 15 employs an amplitude control circuit to avoid a degenerate state. Since the phase of the interpolated clock only depends on α/β, these weights may diminish toward zero while, in principle, maintaining the proper phase. As a result, the interpolated clock amplitude continues to decline, eventually causing lock failure. The amplitude control circuit monitors α and β individually, adjusting their value if they fall or rise excessively.

PLL sharing with phase interpolation suffers from two critical drawbacks at very high speeds. First, the required quadrature or multiphase VCO topology inevitably degrades the phase noise of the TX PLL—where jitter is most important. Second, with the large layout dimensions of the TX and the RX —especially if many peaking inductors are used—it becomes exceedingly difficult to route the TX PLL outputs to the RX CDR circuit. The issue of sharing a PLL between TX and RX paths has already manifested itself in ultra-wideband RF transceivers operating at 10 GHz and in new RF designs targeting 60 GHz. In these cases, it is advantageous to employ two

independent synthesizers so as to avoid long interconnects.

B. VCO-Based CDR Circuits

With the foregoing issues plaguing phase-interpolating CDR circuits, the conventional, VCO-based architectures appear a more feasible approach at data rates of 80-100 Gb/s. Except for regenerator units required in long-haul optical communication, most applications make it desirable to employ a half-rate or quarter-rate CDR topology. This is because (a) mutual injection pulling [15] prohibits equal nominal frequencies for the TX VCO and the RX VCO; and (b) data demultiplexing in the receive path is greatly simplified if it is realized within the CDR circuit.

The choice between half-rate and quarter-rate architectures is determined by a number of factors: (1) if the TX PLL operates at half rate [e.g., as in the architecture of Fig. 13(a)], then the CDR circuit must run at quarter rate or lower; (2) the VCO design and layout becomes more complex as the number of required phases increases [16].

C. A New Approach

In order to avoid injection pulling between the TX PLL and the RX CDR circuit, it is desirable to choose a *non-integer* relationship between their VCO frequencies. We propose a CDR architecture that can coexist with a full-rate or half-rate TX PLL with minimal pulling. Shown in Fig. 16(a), the loop incorporates a VCO running at $f_{CK}/3$ (f_{CK} denotes the

(a)

(b)
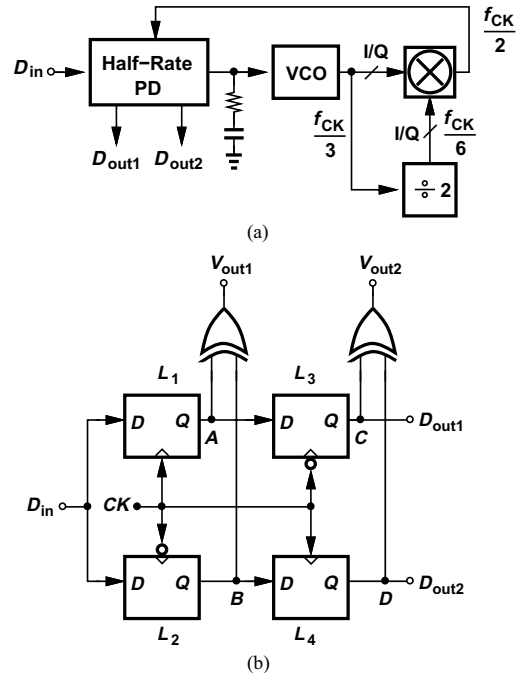
Fig. 16. (a) Proposed CDR architecture, (b) realization of the PD.

full-rate clock) and a $\div 2$ circuit generating $f_{CK}/6$. These two frequencies are applied to a single-sideband mixer so as to yield a half-rate clock. A half-rate PD performs phase comparison with the input data while producing demultiplexed outputs D_{out1} and D_{out2}.

The design of the VCO for operation at $f_{CK}/3$ is relatively relaxed even though it must provide quadrature outputs for SSB mixing. However, since the mixer does not generate quadrature phases of the half-rate clock, the PD topology must be chosen accordingly. A half-rate PD that does not require quadrature clock phases is reported in [17] and shown in Fig. 16(b). Here, latches L_1-L_4 serve as the phase detector, and V_{out1} and V_{out2} are applied to a charge pump or V/I converter.

In the SSB mixer of Fig. 16(a), mismatches produce a fraction of the unwanted sideband at $f_{CK}/3 - f_{CK}/6 = f_{CK}/6$, which is 25 to 30 dB below the wanted component. Fortunately, an inductively-loaded mixer such as that in Fig. 13(b) suppresses this component by about 20 dB. Note that third-order nonlinearity at the input ports of the mixer is benign because it results in a component given by $3(f_{CK}/3) - 3(f_{CK}/6) = f_{CK}/2$.

V. FUTURE TRENDS

As wireline transceivers target speeds greater than 40 Gb/s, a number of trends are likely to emerge.

(1) Millimeter-wave device modeling and circuit techniques will find increasingly broader usage in wireline designs.

(2) While conventional transmission standards typically do not distinguish among different jitter frequencies in the TX mask, future standards may add a spectral mask to reveal components that are more difficult to remove in the receiver.

(3) Future "circuit-friendly" standards may allow a greater bandwidth for CDR circuits so that the recovered clock simply tracks the data edges rather than ignores the input jitter. This larger bandwidth will allow suppressing the phase noise of the CDR VCO to a greater extent.

(4) Layout and packaging of high-speed transceivers will also draw upon the techniques developed in millimeter-wave systems.

VI. CONCLUSION

High-speed transceivers continue to present interesting challenges, requiring more than ever TX data path and PLL co-design, TX PLL and RX CDR co-design, and the use of RF and millimeter-wave techniques. It appears that the reference phase noise will play a major role in the jitter of PLLs, and cascaded loops will offer marginal improvement. A TX/PLL interface and a CDR architecture are proposed as a means of relaxing some of these issues.

REFERENCES

[1] M. Sasaki, "A 9.5GHz 6ps-Skew Space-Filling-Curve Clock Distribution with 1.8V Fill-Swing Standing-Wave oscillators," *ISSCC Dig. Tech. Papers,* pp. 518-519, Feb. 2008.

[2] F. Yang et al, "A CMOS low-power multiple 2.5-3.125-Gb/s serial link macrocell for high IO bandwidth network ICs," *IEEE J. Solid-State Circuits*, vol. 37, pp. 1813-1821, Dec. 2002.

[3] S. Gondi and B. Razavi, "Equalization and Clock and Data Recovery Techniques for 10-Gb/s CMOS Serial Links," *IEEE J. Solid-State Circuits*, vol. 42, pp. 1999-2011, Sept. 2007.

[4] H. Noguchi et al, "A 40Gb/s CDR Cicruit with Adaptive Decision-Point Control Using Eye-Opening Monitor Feedback," *ISSCC Dig. Tech. Papers,* pp. 228-229, Feb. 2008.

[5] R. Farjad-Rad et al, "A low-power multiplying DLL for low-jitter multigigahertz clock generation in highly integrated digital chips," *IEEE J. Solid-State Circuits*, vol. 37, pp. 1804-1812, Dec. 2002.

[6] S. Gierkink, "An 800MHz-122dBc/Hz-at-200kHz Clock Multiplier Based on a Combination of PLL and Recirculating DLL," *ISSCC Dig. Tech. Papers,* pp. 454-455, Feb. 2008.

[7] S. Galal and B. Razavi, "10-Gb/s Limiting Amplifier and Laser/Modulator Driver in 0.18um CMOS Technology," *IEEE J. Solid-State Circuits*, vol. 38, pp. 2138-2146, Dec. 2003.

[8] B. Razavi, "A Millimeter-Wave Circuit Technique," to appear in *IEEE J. Solid-State Circuits*, Sept. 2008.

[9] B. Razavi, "Heterodyne Phase Locking: A Technique for High-Speed Frequency Division," *IEEE J. Solid-State Circuits*, vol. 42, pp. 2877-2892, Dec. 2007.

[10] K. Scheir et al, "Design and Analysis of Inductors for 60 GHz Applications in a Digital CMOS Technology," *Proc. 69th ARFTG Microwave Measurement Conference,* June 2007.

[11] T. Dickson et al, "30-100 GHz Inductors and Transformers for Millimeter-Wave (BI)CMOS Integrated Circuits," *IEEE Trans. Microwave Theory and Techniques*, vol. 53, pp. 123-133, Jan. 2005.

[12] B. Razavi, "Design Considerations for Future RF Circuits," *Proc. International Conference on Circuits and Systems*, pp. 741-744, May 2007, New Orleans.

[13] M. Kossel et al, "A low-jitter wideband multiphase PLL in 90nm SOI CMOS technology," *ISSCC Dig. Tech. Papers*, pp. 414-415, and slide supplement, Feb. 2005.

[14] B. Razavi, K. F. Lee, and R. H. Yan, "Design of High-Speed Low-Power Frequency Dividers and Phase-Locked Loops in Deep Submicron CMOS," *IEEE J. Solid-State Circuits*, vol. 30, pp. 101-109, Feb. 1995.

[15] B. Razavi, "Mutual Injection Pulling Between Oscillators," *Proc. IEEE Custom Integrated Circuits Conference*, pp. 675-678, Sept. 2006.

[16] J. Lee and B. Razavi, "A 40-Gb/s clock and data recovery circuit in 0.18um CMOS technology," *IEEE J. Solid-State Circuits*, vol. 38, pp. 2181-2190, Dec. 2003.

[17] 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.