

# Design of High-Speed Circuits for Optical Communication Systems

Behzad Razavi  
Electrical Engineering Department  
University of California, Los Angeles

## Abstract

This paper presents the design of circuits and architectures for optical communication transceivers. First, a system overview illustrating the challenges in high-speed implementations is given. Next, the design of transimpedance amplifiers and limiters is discussed and the problem of clock and data recovery is addressed. Finally, jitter issues and methods of estimating the jitter are introduced.

## I. INTRODUCTION

Recent statistics indicate that the number of internet nodes doubles approximately every 100 days, leading to backbone data rates exceeding 10 Tb/s by the year 2005. Revitalizing optical communication, this trend has suddenly created a widespread demand for high-speed optical and electronic devices, circuits, and systems.

The new optical revolution is reminiscent of the monumental change that the RF design paradigm began to experience in the early 1990s: modular, general-purpose building blocks are gradually replaced by end-to-end solutions that benefit from device/circuit/architecture codesign; greater levels of integration on a single chip enable higher performance and lower cost; and mainstream VLSI technologies such as BiCMOS and CMOS continue to take over the territories thus far claimed by GaAs and InP devices.

Optical systems are nonetheless different from RF systems in many respects. In addition to speeds in the range of tens of gigabits per second, other issues such as the random nature of data, the very broad spectrum of signals, and the existence of very low frequency components in the data demand vastly different circuit techniques. Furthermore, the large amount of high-speed digital circuitry required in optical transceivers makes full integration difficult.

This paper deals with the design of high-speed circuits for optical communications. The next section provides an overview of a typical system and its critical challenges. Section III presents the design of transimpedance amplifiers (TIAs) and limiters and Section IV describes clock and data recovery architectures. Section V reviews jitter issues and proposes a method of estimating the jitter due to oscillator phase noise.

## II. SYSTEM OVERVIEW

Figure 1 shows a typical optical system. In the transmitter (TX), a number of channels are multiplexed into a high-speed

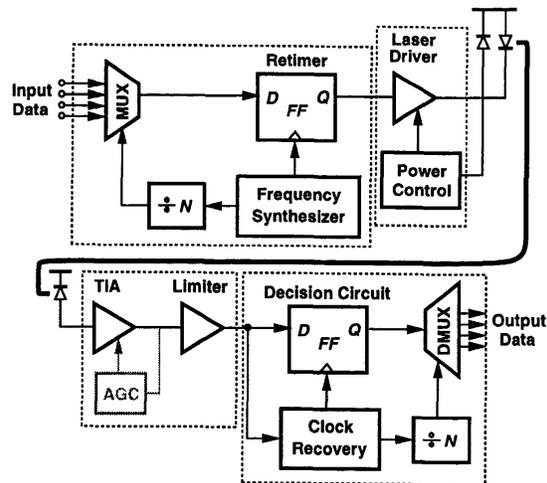


Fig. 1. Optical transceiver system.

data stream, the result is retimed and applied to a laser driver, and the optical output thus produced is delivered to the fiber. A frequency synthesizer generates clocks for both the multiplexer (MUX) and the retiming flipflop (FF). Also, since the laser output power varies with temperature and aging, a monitor photodiode (PD) and a power control circuit continuously adjust the output level of the driver.

In the receiver (RX), a photodiode converts the optical signal to a current and a transimpedance amplifier and a limiter raise the signal swing to logical levels. [The TIA may incorporate automatic gain control (AGC) to accommodate a wide range of input currents.] Subsequently, a clock recovery circuit extracts the clock from the data with proper edge alignment and retimes the data by a “decision circuit.” The result is then demultiplexed, thereby producing the original channels.

While the system topology of Fig. 1 has not changed much over the past several decades, the design of its building blocks and the levels of integration have. Motivated by the evolution and affordability of IC technologies as well as the demand for higher performance, this change has created new challenges, necessitating new circuit and architecture techniques. We review some of the challenges here.

The transmitter of Fig. 1 entails several issues that manifest themselves at high speeds and/or in scaled IC technologies. Since the jitter of the transmitted data is determined by pri-

marily that of the synthesizer, a robust, low-noise phase-locked loop (PLL) with high supply and substrate rejection becomes essential. Furthermore, the design of skew-free, synchronous multiplexers proves difficult at high data rates.

Another critical challenge arises from the laser driver, a circuit that must deliver tens of milliamperes of current with very short rise and fall times. Since laser diodes may experience large voltage swings between on and off states, the driver design becomes more difficult as scaled technologies impose lower supply voltages. The package parasitics also severely limit the speed with which such high currents can be switched to the laser [1].

The optical components in Fig. 1, namely, the laser diode, the fiber, and the photodiode, introduce their own nonidealities in the systems, requiring close interaction between electronic and optical design. Effects such as chirp, dispersion, attenuation, and efficiency play a major role in the overall link budget.

The receiver of Fig. 1 also presents many problems. The noise, gain, and bandwidth of the TIA and the limiter directly impact the sensitivity and speed of the overall system, raising additional issues as the supply voltage scales down. Moreover, the clock and data recovery functions must provide a high speed, tolerate long runs (sequences of identical bits), and satisfy stringent jitter and bandwidth requirements.

Full integration of the transceiver shown in Fig. 1 on a single chip raises a number of concerns. The high-speed digital signals in the MUX and DMUX may corrupt the receiver input or the oscillators used in the synthesizer and the CDR circuit. The high slew rates produced by the laser driver may lead to similar corruptions and also desensitize the TIA. Finally, since the VCOs in the transmit synthesizer and the receive CDR circuit operate at slightly different frequencies (with the difference given by the mismatch between the crystal frequencies in two communicating transceivers), they may pull each other, generating substantial jitter.

The above issues have resulted in multichip solutions that integrate the noisy and sensitive functions on different substrates. The dashed boxes in Fig. 1 indicate a typical partitioning, suggesting the following single-chip blocks: the synthesizer/MUX circuit (also called the “serializer”), the laser driver along with its power control circuitry, the TIA/limiter combination, and the CDR/MUX circuit (also called the “deserializer”). Recent work has integrated the serializer and deserializer (producing a “serdes”) but the TX and RX amplifiers remain in isolation.

### III. TIAs AND LIMITERS

#### A. TIAs

Transimpedance amplifiers play a critical role in optical receivers. Trade-offs between noise, speed, gain, and supply voltage present many challenges in TIA design. As TIAs experience a tighter performance envelope with technology scaling at the device level and speed scaling at the system level, it becomes necessary to design the cascade of the TIA,

the limiter, and the decision circuit concurrently. The TIA bandwidth is typically chosen to be equal to 0.7 times the bit rate - a reasonable compromise between the total integrated noise and the intersymbol interference (ISI) resulting from limited bandwidth.

Shown in Fig. 2, the common-gate (or common-base) topology is a candidate for TIAs as it provides a relatively low input

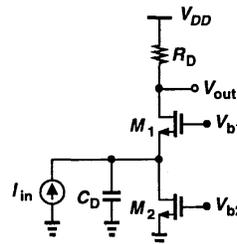


Fig. 2. Common-gate TIA.

impedance, a broad band, and a well-behaved time response. However, its input-referred noise current,  $I_{n,in}$ , is relatively high. This is because  $\overline{I_{n,in}^2}$  per unit bandwidth at low frequencies is given by

$$\overline{I_{n,in}^2} = \overline{I_{n,M2}^2} + \overline{I_{n,RD}^2} \quad (1)$$

$$= 4kT(\gamma g_{m2} + \frac{1}{R_D}), \quad (2)$$

where  $\gamma$  denotes the excess noise coefficient of  $M_2$  ( $\gamma \approx 2.5$  in  $0.25\text{-}\mu\text{m}$  technology). Interestingly, the noise currents of  $M_2$  and  $R_D$  are directly referred to the input with a unity factor. Furthermore, for a given supply voltage,  $I_{n,M2}$  and  $I_{n,RD}$  trade with *each other* because the minimum drain-source voltage of  $M_2$  plus the voltage drop across  $R_D$  cannot exceed  $V_{DD}$ . In other words,  $g_{m2}$  and  $1/R_D$  are inevitably large. It can also be shown that the noise contributed by  $M_1$  and  $R_D$  rises as the frequency increases and the photodiode capacitance,  $C_D$ , shunts the input [2].

A TIA configuration achieving more relaxed noise-headroom trade-offs is the shunt-shunt feedback topology. Shown in Fig. 3(a) as feedback around a voltage amplifier  $A_1$ , the circuit exhibits a  $-3\text{-dB}$  bandwidth of  $(2\pi)A_1/(R_F C_D)$  (if the poles of  $A_1$  are neglected) and an input-referred noise current per unit bandwidth equal to

$$\overline{I_{n,in}^2} = \frac{4kT}{R_F} + \frac{V_{n,A1}^2}{R_F^2}, \quad (3)$$

where  $V_{n,A1}$  denotes the input-referred noise voltage of  $A_1$ .<sup>1</sup> The key point here is that  $R_F$  does not carry significant dc current and can therefore be maximized so as to reduce both terms in (3). This is in contrast to the behavior represented by (2).

Actual implementations of the feedback TIA suffer from voltage headroom, stability, and overshoot problems. Considering the example shown in Fig. 3(b), we recognize that

<sup>1</sup>The input-referred noise current of  $A_1$  is neglected for simplicity.

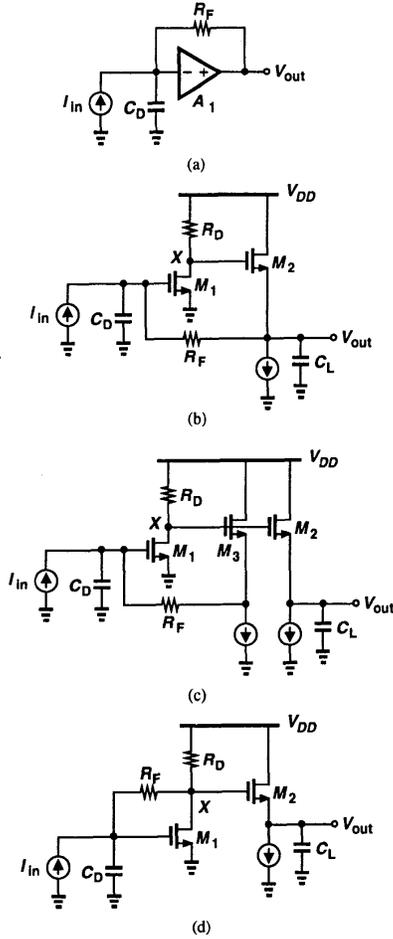


Fig. 3. Feedback TIA and its realizations.

$V_{GS1} + V_{GS2}$  significantly constrains the dc drop across  $R_D$ , thereby limiting the open-loop gain and raising the noise contributed by  $R_D$  and  $M_2$ . Furthermore, the three poles at the input node, the drain of  $M_1$ , and the output node degrade the phase margin and hence the step response. Figure 3(c) suggests a modification that isolates the feedback path from the input capacitance of the subsequent stage. Finally, Fig. 3(d) eliminates the source follower from the feedback loop to allow a greater drop across  $R_D$  [3].

It is possible to choose the pole at node  $X$  in Figs. 3(c) or (d) so as to *increase* the bandwidth of the TIA. In fact, if the magnitude of this pole is equal to  $2A_1/(R_F C_D)$ , the TIA exhibits a slightly underdamped step response but a bandwidth of  $(2\pi)\sqrt{2}A_1/(R_F C_D)$ , i.e., 40% greater than that for an ideal core amplifier.

### B. Limiters

The voltage swing produced by TIAs at the minimum light level is usually inadequate to drive the CDR circuit, necessitating further amplification. Used to boost the binary swings, limiters typically consist of a cascade of differential pairs with enough bandwidth and a relatively linear phase response so as

to amplify the signal with negligible ISI. The high small-signal gain requires low-frequency negative feedback to prohibit the offset voltages of the differential pairs from saturating the latter stages.

Interestingly, limiter design must cope with difficulties at both the low corner and the high corner of the passband. Consider the limiter topology shown in Fig. 4, where the feedback network suppresses the offset of the last three stages. Since

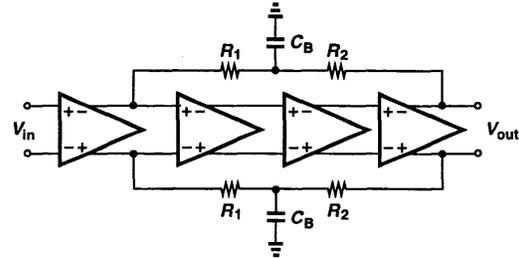


Fig. 4. Simple limiter.

some optical standards require that the low end of the passband fall around a few tens of kilohertz, the values of  $R_2$  and  $C_B$  must be very large. More specifically, with a small-signal gain of  $A$  per stage, the low corner frequency is given by  $(A^3 + 1)/(2\pi R_2 C_B)$ , demanding an  $R_2 C_B$  product on the order of 1 ms if  $A$  is around 5. For this reason, the capacitors are usually placed off chip, raising the number of package pins and also the possibility of crosstalk from other bond wires. New circuit topologies may resolve this issue.

At the upper end of the passband, high-speed amplification techniques must provide a well-behaved magnitude and phase response for both small and large signals. Shown in Fig. 5, configurations such as the Cherry-Hooper amplifier [4, 1] and the Gilbert gain cell [5] have been used but their utility becomes more limited as the supply voltage falls. In particular, the voltage drops across  $R_{1,2}$  and  $R_{3,4}$  in Fig. 5(a) and the cascode in Fig. 5(b) both constrain the voltage headroom and mandate level-shift circuits between the stages.

An attractive solution for low-voltage broadband amplifiers is inductive peaking. Owing to the extensive work on monolithic inductors in RF design, this method can now be realized with accurate modeling and prediction of the performance in optical communication circuits as well. Interestingly, inductor quality factors ( $Q$ 's) as low as 3 to 4 prove adequate for increasing the bandwidth, allowing the use of simple, compact spiral structures.

Figure 6(a) shows a limiting stage incorporating inductive peaking. It can be shown that an ideal inductor increases the bandwidth by approximately 82% if a 7.5% overshoot in the step response is acceptable. With the finite  $Q$  and parasitic capacitance of the inductors included, the enhancement is around 50%, still quite a significant factor.

An interesting difficulty in modeling the inductors in the above circuit arises from the narrowband nature of the definition of the  $Q$ , an issue rarely encountered in RF design. Figure 6(b) depicts a rough model where  $R_P/(L_P\omega)$  yields the correct  $Q$  at about 3/4 of the  $-3$ -dB bandwidth. The ap-

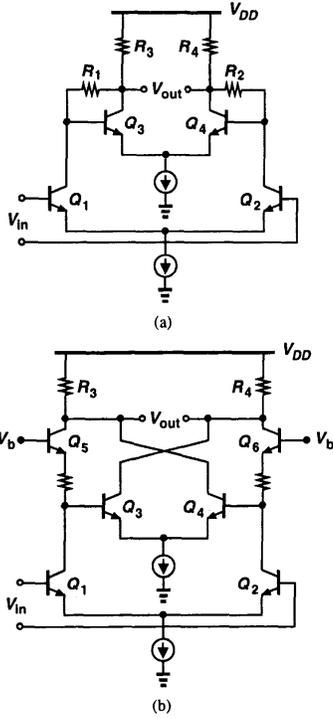


Fig. 5. (a) Cherry-Hooper amplifier, (b) Gilbert gain cell.

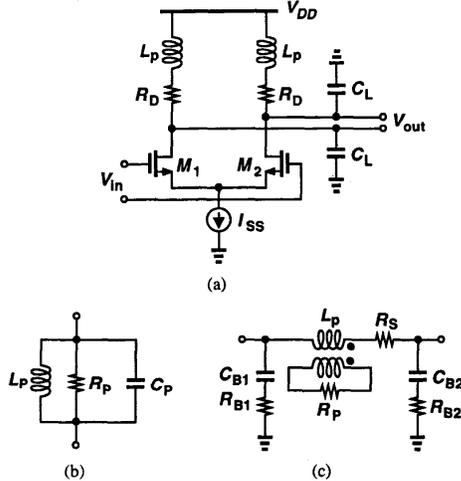


Fig. 6. (a) Inductive peaking, (b) simple inductor model, (c) more complete inductor model.

proximation is reasonable because the inductor manifests itself only near the high end of the band. Alternatively, a more complete model such as that in Fig. 6(c) can be used. Here,  $R_S$  denotes the effective series resistance,  $R_{B1}$  and  $R_{B2}$  represent the resistance seen by the electric coupling to the substrate,  $R_P$  models the resistance seen by the magnetic coupling to the substrate, and the capacitors approximate the parasitic capacitances. While the values of some of the components in this model do vary with frequency, the overall model can be fitted to measured data over a broader range than the parallel tank of

Fig. 6(b) can.

The high gain provided by several stages in a limiter may lead to oscillation or at least considerable peaking and ISI. Illustrated in Fig. 7, this phenomenon occurs if the mismatches

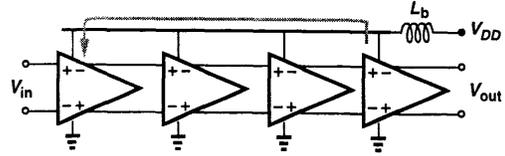


Fig. 7. Instability resulting from feedback through supply line.

in the differential stages create both substantial current switching from the supply and a finite supply rejection, allowing a component to travel from the output stage through the supply and back to the input stage. With a finite bond wire inductance,  $L_b$ , the gain around the loop may exceed unity, leading to high-frequency oscillation. The issue of course becomes much more severe if a single-ended TIA shares the same supply lines with the limiter. For this reason, separate supply lines, careful bypassing, symmetric layout, and accurate package modeling are essential.

#### IV. CLOCK AND DATA RECOVERY CIRCUITS

CDR circuits are among the most challenging building blocks in optical transceivers. In addition to high speed, the properties of random non-return-to-zero (NRZ) data and the stringent jitter and loop bandwidth specifications of optical standards create many design and simulation issues at both circuit and architecture levels.

Figure 8(a) illustrates a generic CDR architecture, where a voltage-controlled oscillator (VCO) is phase-locked to the

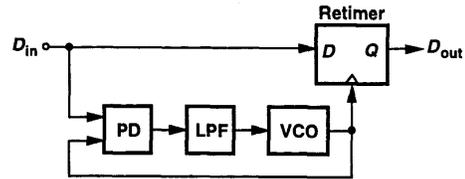


Fig. 8. Generic CDR circuit.

input data by means of a phase detector (PD) and a low-pass filter (LPF). The recovered clock generated by the VCO then samples and retimes the data, thereby reducing the jitter and ISI and providing timing coherency with subsequent operations.

##### A. Phase Detectors

Since the spectrum of NRZ data contains no impulse at a frequency equal to the bit rate, the phase detector in Fig. 8 must generate such a component. This can be accomplished by either explicit edge detection and full-wave rectification or a PD topology in which the edges of data sample the VCO output. In fact, a simple master-slave D flipflop can serve as an NRZ phase detector if its clock input is driven by the data stream and its D input senses the VCO output [Fig. 9(a)]. Called a "bang-bang" phase detector, this topology exhibits a very nonlinear characteristic [Fig. 9(b)], applying large swings

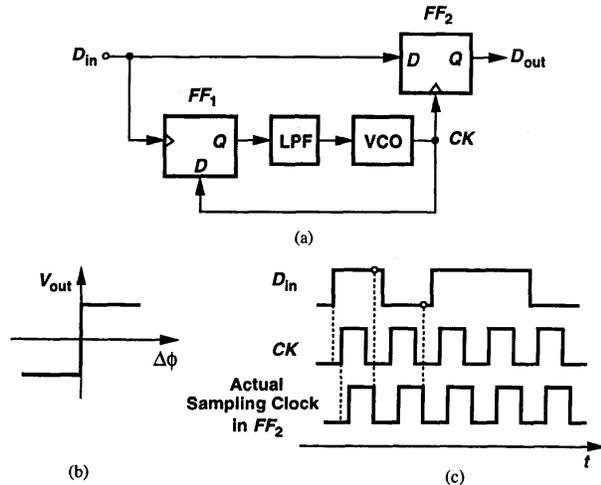


Fig. 9. (a) CDR circuit using a D flipflop phase detector, (b) PD characteristic, (c) addition of skews in  $FF_1$  and  $FF_2$ .

to the loop filter and possibly introducing substantial ripple on the oscillator control line.

A critical drawback of this CDR architecture at high speeds results from the skews in  $FF_1$  and  $FF_2$ . Since typical flipflops suffer from unequal data-to-output and clock-to-output delays, the loop locks such that the recovered clock and the input data sustain a finite, systematic phase offset, compensating for the delay difference. Illustrated in Fig. 9(c), the skews of  $FF_1$  and  $FF_2$  add, resulting in a significant deviation of the clock edge from the middle of the data bits.

The above skew phenomenon manifests itself in most CDR topologies based on the generic architecture of Fig. 8. At high speeds, therefore, it is desirable to retime the data inside the phase detector. Two such PDs are the Alexander [6] and Hogge [7] topologies. Depicted in Fig. 10(a), the Alexander PD employs four flipflops to realize a three-point sampling scheme and two XOR gates to compare the samples. As illustrated in Fig. 10(b),  $FF_1$  and  $FF_2$  sample the data bits on two consecutive rising (falling) edges of the clock and  $FF_3$  samples the data on the falling (rising) edge of the clock.  $FF_4$  aligns this sample with that at A. As a consequence, if  $CK$  leads  $D_{in}$ , then  $A \neq C = B$ , generating a low  $V_{out1}$  and a high  $V_{out2}$ . Conversely, if  $CK$  lags  $D_{in}$ , then  $A = C \neq B$ , forcing a high level at  $V_{out1}$  and a low level at  $V_{out2}$ . A PLL utilizing this PD therefore locks such that the falling edge of  $CK$  wanders in the close vicinity of the data edges. Under this condition, the output of  $FF_2$  is an optimally-retimed version of the input, obviating the need for an explicit retimer and its associated skew.

Note that in the absence of data transitions,  $A = C = B$  and  $V_{out1} = V_{out2}$ , i.e., the circuit produces no new output, a behavior similar to that of  $FF_1$  in Fig. 9. For this reason, the Alexander PD contains edge detection as well, allowing phase locking to NRZ data.

The Alexander topology is a bang-bang PD: a small phase difference between data and clock edges drives  $V_{out1}$  and  $V_{out2}$

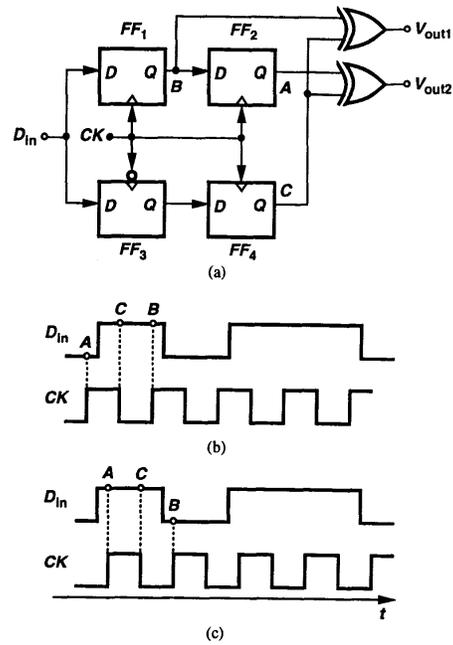


Fig. 10. (a) Alexander phase detector, (b) waveforms for when  $CK$  leads  $D_{in}$ , (c) waveforms for when  $CK$  lags  $D_{in}$ .

to full logical levels. A better-behaved PD exhibiting a linear characteristic as well as producing a retimed output is the Hogge configuration. Depicted in Fig. 11(a), the circuit in-

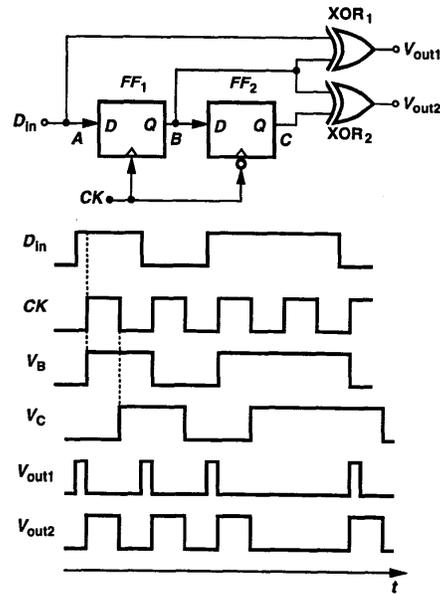


Fig. 11. Hogge phase detector and its waveforms.

corporates two D flipflops driven by opposite phases of the clock and two XOR gates.  $FF_1$  samples the data on the rising (falling) edge of  $CK$  and  $XOR_1$  compares the phases of  $D_{in}$  and  $V_B$ , generating a pulse whose width is equal to half a clock period if the clock samples the midpoint of the data bits.

$FF_2$  samples  $V_B$  on the falling (rising) edge of  $CK$ . Thus, as shown in Fig. 11(b), the waveform at node  $C$  is identical to that at  $B$  except for a half clock period delay. As a result,  $XOR_2$  produces a pulse with a width equal to half the clock period for every data transition, serving as reference. In other words, if the data and clock are aligned, each data transition creates pulses of equal width at  $V_{out1}$  and  $V_{out2}$ , yielding a zero average for  $V_{out1} - V_{out2}$ . As the phase difference between the data and the clock increases, the pulses at  $V_{out1}$  change their width accordingly, increasing the average value of  $|V_{out1} - V_{out2}|$  linearly. Under locked condition, the output of  $FF_2$  provides the retimed data with no systematic skew.

While exhibiting a linear characteristic, the Hogge PD entails a number of issues. First, the finite delay through  $FF_1$  introduces a phase offset in the locked condition [7], degrading the clock phase margin. Second, as shown in Fig. 12, the retiming delay through  $FF_2$  leads to a half-period skew between

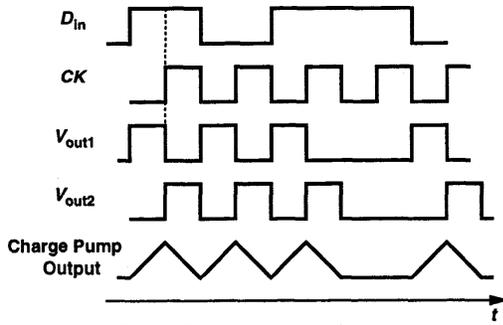


Fig. 12. Problem of triwave in Hogge PD.

the pulses at  $V_{out1}$  and those at  $V_{out2}$ . Consequently, even in lock, a charge pump (CP) and loop filter driven by  $V_{out1}$  and  $V_{out2}$  produce a positive ramp while  $V_{out1}$  is high and a negative ramp while  $V_{out2}$  is high. The control line of the VCO therefore experiences a “triwave” with a positive net area [8, 9], disturbing the VCO on every data transition. Modifications proposed in [8, 9] alleviate these issues.

### B. Frequency Acquisition

Most optical standards specify a very narrow jitter transfer bandwidth for CDR circuits - typically less than 1% of the operating speed. For this reason, clock recovery architectures complying with these specifications and incorporating only a phase detector display a capture range of less than a few percent. On the other hand, the center frequency of the VCO may vary by a large amount with process and temperature - as much as a factor of two for CMOS ring oscillators. CDR circuits must therefore employ a means of frequency acquisition to drive the VCO frequency close to the desired value before phase locking takes over. Such means typically require two loops around the VCO.

A common approach to performing frequency acquisition involves locking the VCO to an external reference through a simple PLL. Illustrated in Fig. 13, the concept employs two loops sharing the VCO; Loop I locks the VCO to  $f_{REF}$  and, when the lock detector decides that  $f_{out}/N$  is sufficiently close to  $f_{REF}$ , Loop II is enabled to phase-lock the VCO to  $D_{in}$ .

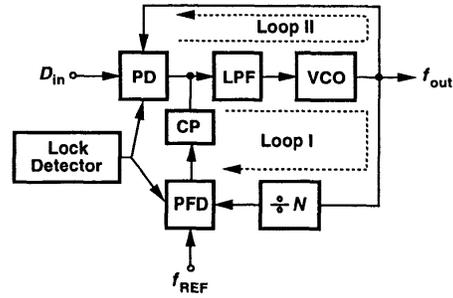


Fig. 13. Dual-loop CDR architecture including frequency acquisition.

During data reception, the lock detector continues to monitor  $f_{out}/N$ , enabling Loop I if unexpected noise drives Loop II out of lock.

Figure 14 shows a variant of the above approach [10]. Here,  $VCO_2$  is permanently locked to  $f_{REF}$  while  $VCO_1$ , a replica

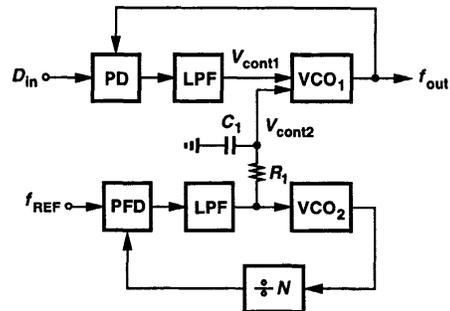


Fig. 14. Frequency acquisition by means of replica VCO.

of  $VCO_2$ , is driven by the CDR loop as well as sensing the control voltage of  $VCO_2$ . The idea is that  $V_{cont2}$  serves as a coarse control for  $VCO_1$ , bringing  $f_{out}$  close to the desired value, and  $V_{cont1}$  provides a fine control, locking  $VCO_1$  to  $D_{in}$ .

An important advantage of the architecture in Fig. 14 over that in Fig. 13 relates to the sensitivity (gain,  $K_{VCO}$ ) of the oscillators. Decomposing the control into coarse and fine lines, the circuit of Fig. 14 exhibits much less VCO gain in the CDR loop, thereby tolerating more noise on  $V_{cont1}$  for a given amount of output jitter. The coarse control,  $V_{cont2}$ , is heavily filtered by  $R_1$  and  $C_1$  to present a low ripple to  $VCO_1$ .

The architecture of Fig. 14 nonetheless suffers from two disadvantages with respect to that of Fig. 13. First, the mismatch between the two VCO frequencies must be less than the capture range of the CDR loop so as to guarantee lock to  $D_{in}$ . This issue may prove difficult in CMOS implementations. Second, operating at slightly different frequencies, the two VCOs may pull each other through substrate coupling, thus corrupting the recovered clock and data.

Frequency acquisition with no external references is also utilized. “Referenceless” architectures detect the bit rate of the input random data and drive the VCO frequency toward this value [11]. A common approach is to process the data and the VCO output so as to generate two beat waveforms that bear a phase difference of  $+90^\circ$  or  $-90^\circ$  depending on the polarity

of the frequency difference. One beat signal subsequently samples the other, providing a dc value that can be applied to the VCO. Called the “quadricorrelator” [12], this architecture has been realized in analog and digital forms. Shown in Fig. 15 is a digital realization [13] employing D flipflops that sample

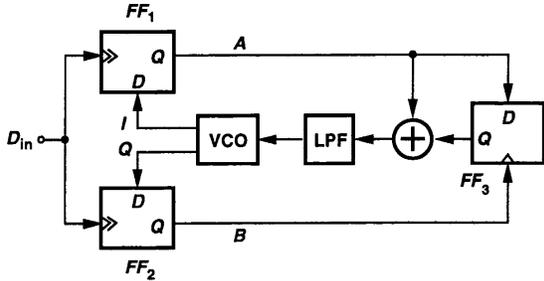


Fig. 15. CDR circuit including frequency detector.

on both the rising and falling edges (hence the name “double-edge-triggered flipflop”). On each transition of the random data,  $FF_1$  and  $FF_2$  sample the quadrature phases of the VCO, producing beat signals at  $A$  and  $B$ .  $FF_3$  then samples one by the other, generating a dc component that represents the polarity of the frequency error. Note that  $FF_1$  also serves as a bang-bang PD, locking the VCO to  $D_{in}$  after frequency acquisition is completed.

CDR loops employing frequency detectors that operate with random data exhibit only a moderate capture range, typically on the order of  $\pm 10\%$  of the center frequency. This limitation can be explained with the aid of the characteristic plotted in Fig. 16 for the frequency detector of Fig. 15. We note that for a

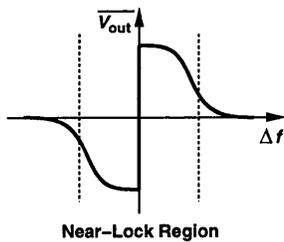


Fig. 16. Characteristic of the frequency detector in Fig. 15.

large difference between the data rate and the VCO frequency, the average output is close to zero, carrying little information. Figure 17(a) shows a part of a dual-loop architecture that substantially increases the capture range [14]. The frequency detector used here is similar to that in Fig. 15. Here, a counter controlling the capacitor array sets the VCO frequency to the lowest value. Under this condition,  $\Delta f$  is very negative and  $V_{FD}$  is close to zero. Thus, the two comparators generate logical zeros, the output of the OR gate remains low, and the counter continues to (slowly) count up until  $V_{FD}$  drops below  $V_L$ . This is an indication that  $V_{FD}$  has reached a reliable level. Now the two flipflops begin to save each state before the next count is carried out. The counter still continues to count until  $\Delta f$  crosses zero and  $V_{FD}$  jumps from negative to positive. The two flipflops then record this change, disabling the counter and enabling the CDR loop (not shown in Fig. 17).

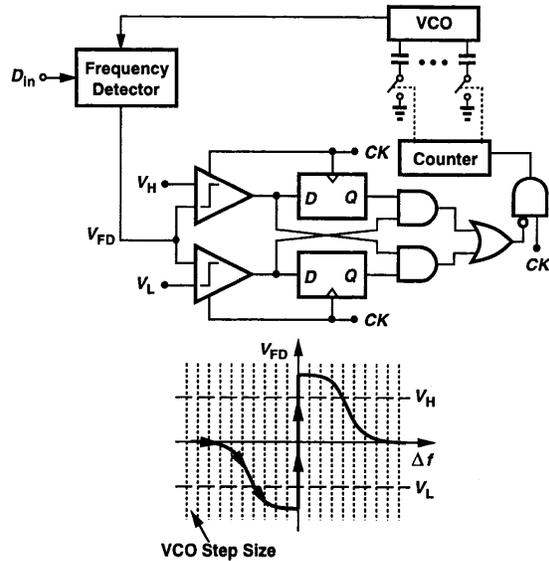


Fig. 17. CDR architecture with wide capture range.

### C. Half-Rate Architectures

If the data rate is higher than the maximum tolerable speed of phase detectors and VCOs, a half-rate CDR architecture can be used [15, 16]. The idea is to run the VCO at a frequency equal to half of the data rate, thereby relaxing the design of the circuits in the signal path. Half-rate architectures usually demultiplex the data as well.

The principal challenge in half-rate CDR circuits relates to the design of phase and frequency detectors that operate properly while sensing full-rate data and a half-rate clock. Figure 18 depicts an example of a linear half-rate PD [17].

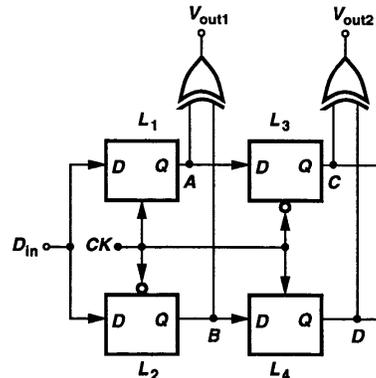


Fig. 18. Half-rate linear PD.

The circuit employs four D latches and two XOR gates. Since latches  $L_1$  and  $L_2$  sample  $D_{in}$  on rising and falling edges of  $CK$ ,  $A \oplus B$  produces a pulse each time a data transition occurs between a rising edge and a falling edge of the half-rate clock. The waveforms at  $C$  and  $D$  are identical except for a phase difference equal to half of the clock period. Thus,  $C \oplus D$  produces a constant-width pulse on every data transition, serving as a reference. Interestingly, if the clock edges are aligned with the middle of the data eye, then  $\overline{V_{out1}}$  is equal to

half of  $\overline{V_{out2}}$ . For this reason, a current mirror arrangement following the XOR gates generates a current proportional to  $2V_{out1} - V_{out2}$ , a quantity that falls to zero when the clock samples the midpoint of data bits. Note that the waveforms at  $C$  and  $D$  are the retimed, demultiplexed versions of the input stream.

Other examples of half-rate phase and frequency detectors are described in [15, 16, 18].

## V. JITTER ISSUES

Optical standards impose severe restrictions on the jitter in the transmitted and recovered data, mandating low-noise VCOs and PLLs in both the transmitter and the receiver. The following sources of jitter can be identified: (1) VCO jitter due to electronic device noise, supply noise, and ripple on the control voltage; (2) VCO pulling due to the coupling of data transitions through the phase detector and the retiming circuit (Fig. 19); (3) jitter in the data itself. The overall jitter therefore

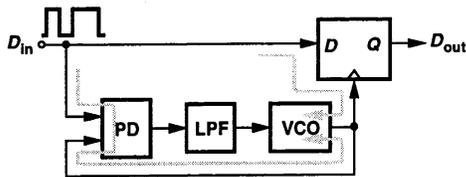


Fig. 19. Coupling of data transitions to the VCO.

depends on not only circuit design but layout and packaging as well.

In design or measurement, it is often necessary to predict the output jitter of a PLL if the electronic noise in the VCO is the dominant source. We describe a simple approach that estimates the closed-loop jitter with reasonable accuracy.

Using simulations or measurements, we first compute the relative phase noise of the free-running VCO due to the sources of white noise. The cycle-to-cycle jitter is then calculated from the phase noise with the aid of the following equation:

$$\Delta T_{cc}^2 \approx \frac{4\pi}{\omega_0^3} S_\phi(\Delta\omega) \Delta\omega^2, \quad (4)$$

where  $\omega_0$  denotes the oscillation frequency and  $S_\phi(\Delta\omega)$  represents the relative phase noise power at an offset frequency of  $\Delta\omega$  [19].

In the next step, we relate the jitter of the PLL to that of the free-running VCO. It has been shown that the closed-loop jitter can be viewed as if the VCO jitter rises with the square-root of time and saturates at a time equal to the inverse of the loop bandwidth [20]. If the loop bandwidth is  $B_L$  hertz, then the VCO produces a total of  $[\omega_0/(2\pi)]/B_L$  cycles in  $1/B_L$  seconds. Thus, the total accumulated jitter due to the VCO is equal to  $\sqrt{[\omega_0/(2\pi)]/B_L} \Delta T_{cc}$ . For example, if  $\omega_0 = 2\pi \times 10$  GHz and  $B_L = 10$  MHz, then  $\Delta T_{cc}$  must be less than 33 fs for the closed-loop jitter to remain below 1 ps. Equation (4) suggests that this can be achieved if the free-running VCO phase noise at 1-MHz offset is below  $-93$  dBc/Hz.

## REFERENCES

- [1] H.-M. Rein and M. Moller, "Design Considerations for Very High Speed Si Bipolar ICs Operating up to 50 Gb/s," *IEEE Journal of Solid-State Circuits*, vol. 31, pp. 1076-1090, August 1996.
- [2] B. Razavi, *Design of Integrated Circuits for Optical Communication Systems*, to be published in 2001.
- [3] M. Neuhauser, H.-M. Rein, and H. Wrenz, "Low-Noise, High-Gain Si Bipolar Preamplifiers for 10 G/s Optical Fiber Links - Design and Realization," *IEEE Journal of Solid-State Circuits*, vol. 31, pp. 24-29, Jan. 1996.
- [4] E. M. Cherry and D. E. Hooper, "The Design of Wideband Transistor Feedback Amplifiers," *Proc. IEE*, vol. 110, pp. 375-389, Feb. 1963.
- [5] B. Gilbert, "A New Wideband Amplifier Technique," *IEEE Journal of Solid-State Circuits*, vol. 3, pp. 353-365, Dec. 1968.
- [6] J. D. H. Alexander, "Clock Recovery from Random Binary Data," *Electronics Letters*, vol. 11, pp. 541-542, Oct. 1975.
- [7] C. R. Hogge, "A Self-Correcting Clock Recovery Circuit," *IEEE J. Lightwave Tech.*, vol. 3, pp. 1312-1314, Dec. 1985.
- [8] L. De Vito et al, "A 52MHz and 155MHz Clock Recovery PLL," *ISSCC Dig. of Tech. Papers*, pp. 142-143, Feb. 1991.
- [9] T. H. Lee and J. F. Bulzacchelli, "A 155-MHz Clock Recovery Delay- and Phase-Locked Loop," *IEEE Journal of Solid-State Circuits*, vol. 27, pp. 1736-1746, Dec. 1992.
- [10] J. C. Scheytt, G. Hanke, and U. Langmann, "A 0.155, 0.622, and 2.488 Gb/s Automatic Bit Rate Selecting Clock and Data Recovery IC for Bit Rate Transparent SDH Systems," *ISSCC Dig. of Tech. Papers*, pp. 348-349, Feb. 1999.
- [11] D. G. Messerschmitt, "Frequency Detectors for PLL Acquisition in Timing and Carrier Recovery," *IEEE Trans. Comm*, vol. 27, pp. 1288-1295, Sept. 1979.
- [12] D. Richman, "Color-Carrier Reference Phase Synchronization Accuracy in NTSC Color Television," *Proc. IRE*, vol. 42, pp. 106-133, Ja. 1954.
- [13] A. Pottbacker, U. Langmann, and H. U. Schreiber, "A Si Bipolar Phase and Frequency Detector for Clock Extraction up to 8 Gb/s," *IEEE Journal of Solid-State Circuits*, vol. 27, pp. 1747-1751, Dec. 1992.
- [14] S. Butala and B. Razavi, "A 2.75-Gb/s CMOS Clock Recovery Circuit with Broad Capture Range," to be presented at *ISSCC*, San Francisco, Feb. 2001.
- [15] M. Rau et al, "Clock/Data Recovery PLL Using Half-Frequency Clock," *IEEE Journal of Solid-State Circuits*, vol. 32, pp. 1156-1159, July 1997.
- [16] M. Wurzer et al, "A 40-Gb/s Integrated Clock and Data Recovery Circuit in a 50-GHz  $f_T$  Silicon Bipolar Technology," *IEEE Journal of Solid-State Circuits*, vol. 34, pp. 1320-1324, Sept. 1999.
- [17] J. Savoj and B. Razavi, "A 10-Gb/s CMOS Clock and Data Recovery Circuit," *VLSI Circuits Symp. Dig. Tech. Papers*, pp. 136-139, June 2000.
- [18] J. Savoj and B. Razavi, "A 10-Gb/s CMOS Clock and Data Recovery Circuit with Frequency Detection," to be presented at *ISSCC*, San Francisco, Feb. 2001.
- [19] F. Herzel and B. Razavi, "A Study of Oscillator Jitter Due to Supply and Substrate Noise," *IEEE Trans. Circuits and Systems, Part II*, vol. 46, pp. 56-62, Jan. 1999.
- [20] J. A. McNeill, "Jitter in Ring Oscillators," *IEEE Journal of Solid-State Circuits*, vol. 32, pp. 870-879, June 1997.