

# An 84-mW 4-Gb/s Clock and Data Recovery Circuit for Serial Link Applications

M.-J. Edward Lee<sup>1,3</sup>, William J. Dally<sup>1,3</sup>, John W. Poulton<sup>2,3</sup>, Patrick Chiang<sup>1</sup>, and Stephen F. Greenwood<sup>1</sup>

<sup>1</sup>Stanford University  
353 Serra Mall #212  
Stanford, CA 94305, U.S.A.

<sup>2</sup>UNC at Chapel Hill  
CB #3175 Sitterson Hall  
Chapel Hill, NC 27599-3175, USA

<sup>3</sup>Velio Communications, Inc.  
2249 Zanker Rd.  
San Jose, CA 95131, USA

Tel: (650) 725-8086, E-mail: elee@cva.stanford.edu

## Abstract

A 4Gb/s serial link tracking clock and data recovery (CDR) circuit fabricated in 0.24 $\mu$ m CMOS technology dissipates 84mW and occupies 0.3mm<sup>2</sup>. The input signal is 2 $\times$  oversampled by 8 offset-cancelled receive amplifiers per receive clock cycle. The samples are processed by a phase controller to position the receive clocks at the center and the edge of the data eye using a semi-digital dual delay-locked loop (DLL) [3]. The quiet-supply p-p jitter of the receive clock is 39ps with 0.33ps/mV supply sensitivity. It allows for plesiochronous clocking with a frequency tolerance of  $\pm 400$ ppm. The worst case phase resolution is < 20ps.

## Introduction

Recent developments in high speed serial links have enabled multi-Gb/s bandwidth through a single pin [1]. For applications such as telecommunication switches and CPU-memory interfaces, hundreds, not just a special few, of the pins on a single chip must operate at Gb/s speed to meet the exponential increase in bandwidth demand. Characteristics that must be met for such high level of integration include low power consumption, small area, and noise immunity. In [2], we introduced a set of techniques targeted at reducing the area, power, and noise sensitivity of high speed I/O circuits. In this paper, we describe a CDR circuit with the same goals of being amenable to large scale integration while capable of operating in excess of 4Gb/s.

## System Architecture

Fig. 1 shows the CDR architecture adapted from [3]. A low-power, noise insensitive core DLL generates 8 clock phases. The DLL has an inverter based delay line regulated by a linear regulator as described in [2]. The absolute phase positions of the 8 clock phases are simultaneously adjusted by 4 differential timing verniers, each composed of two phase muxes and one phase interpolator sequenced by a central phase controller. Each timing vernier selects two adjacent phases using the phase muxes and interpolates between them using the phase interpolator to create 8 finer phase steps. Both the phase mux and phase interpolator are thermometer coded. This architecture is capable of infinite phase rotation and allows for plesiochronous clocking between the transmitter and the receiver. The 8 phases generated by the timing verniers are used to perform 2 $\times$  oversampling on the incoming data with 8 offset-cancelled receive amplifiers described in [2]. The bit rate is 4 $\times$  the receive clock rate. Before feeding into the phase controller, the resulting data samples are further demultiplexed to half the receive clock rate to relax the frequency requirement of the digital logic.

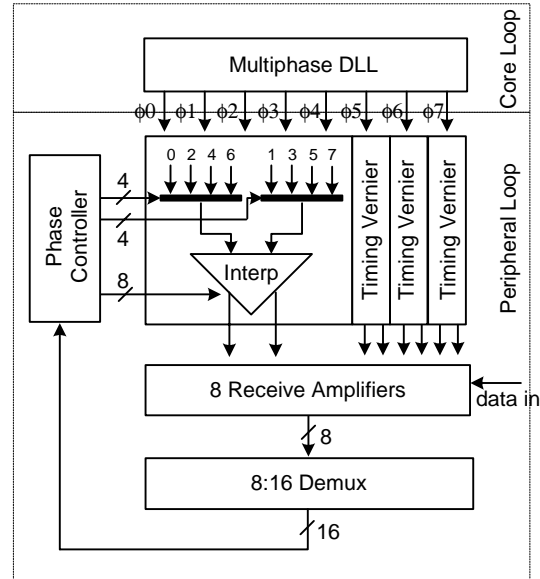


Fig. 1. CDR architecture.

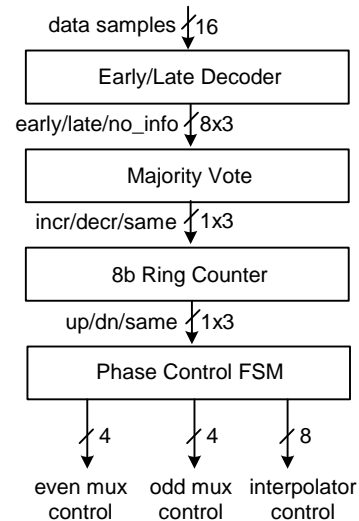


Fig. 2. Phase controller architecture.

The phase controller is shown in Fig. 2. It is clocked at half the frequency of the receive clock (500MHz at 4Gb/s). The 16 data samples generated every cycle first go through a set of early/late decoders. The early/late decoder determines

whether there is a data transition for each bit. If there is, then the edge sample is used to decide whether the receive clocks are too early or too late. Otherwise, there is no timing information contained in that particular bit and the decoder outputs a no\_info. The resulting 8 early/late/no\_info signals are then resolved by a majority vote unit. To avoid loop instability, the summarizing signal generated by the majority vote is low-pass filtered by an 8-bit ring counter, which updates a finite state machine to generate the appropriate phase control signals. The phase interpolators are sequenced all the way to the ends before the phase muxes are switched. This is to ensure that all interpolation weight is on one phase before the other is switched, avoiding potential glitches. This sequencing creates a step which does not have significant effect on the phase position when the mux is switched. Other than decreasing the peripheral loop bandwidth and frequency tolerance, this extra phase step does not affect the performance of the CDR. A more important factor is the maximum phase step, which will be addressed in the next section.

### Circuit Implementation

Fig. 3 and Fig. 4 shows the phase-only comparator [4] and the charge pump used in the core DLL.




Fig. 3. Core DLL phase comparator.




Fig. 4. Core DLL charge pump.

At lock, a mismatch between the up and down charge injection could cause a finite phase offset, which degrades the overall timing budget. An important cause of the charge injection mismatch is channel length modulation within the charge pump. Assuming the phase comparator generates equal up and down pulses at lock, the up and down charge injection difference is given by:

$$\Delta Q = (I_{up} - I_{down}) \times t_s = \Delta I \times t_s \quad (1)$$

where  $\Delta I$  is the difference between up and down current of the charge pump due to channel length modulation, and  $t_s$  is the width of the up and down pulses at steady state. Equation (1) indicates that to minimize the phase offset,  $\Delta I$  and  $t_s$  should be minimized. This phase-only comparator design is capable of output pulses of extremely short duration at steady state due to its simplicity and non-feedback operation. To minimize channel length modulation, a high swing cascode circuit is employed in the charge pump. Simulations indicate <10ps of offset between clk1 and clk2 at lock in all process corners.  $V_{ctrl}$  is the regulated supply voltage for the core DLL delay line and controls the amount of charge pump current. It ensures that the loop bandwidth of the DLL remains a fixed fraction of the operating frequency [2].

Since the timing verniers adjust the phase of the receive clocks in discrete steps (i.e. bang-bang control), it is important to minimize the maximum phase step to avoid excessive dithering at steady state. To do so under a fixed number of phase steps requires a phase interpolator which sweeps phases in linear steps. Fig. 5 shows the current mirror interpolator employed in this design. w0-7 and w0-7b are the interpolator control. Fig. 6 shows the measured phase step as a percentage of the interpolating phase interval. The glitch-guarding step at the boundary, which has negligible effect since the 8 phase steps cover almost 100%, is not included in this plot. The maximum differential nonlinearity (DNL) is 0.16 LSB of the interpolating interval. The phase imbalance among the 8 phases generated by the core DLL further adds to the phase step variation. The measured maximum of the 72 steps (including the 8 glitch guarding steps) in a cycle is 19ps with a receive clock frequency of 1GHz, corresponding to the maximum overall DNL of 0.2 LSB.

Current mirror logic is also used for the phase mux, shown in Fig. 7. Previous phase mux and interpolator designs use either CMOS logic [5] or source-coupled logic [3]. Current mirror logic has significant speed advantage compared to CMOS logic. As a result, it achieves higher frequency and lower jitter. For example, to operate a CMOS mux at the maximum frequency of a current mirror mux, multiple stages have to be used, increasing the jitter sensitivity of the supply. Although CMOS logic does not consume any static current, it was found that the power consumption is comparable at 1GHz since the current mirror logic has less capacitance to switch. Compared to source-coupled logic, current mirror logic is easier to implement since it requires no biasing and is easy to interface with digital circuits (it accepts and generates full swing signal).

### Experimental Results

A complete custom-designed transceiver, including a 4:1 input-multiplexed 2-tap pre-emphasis transmitter [2] and a 1:4 demultiplexed CDR described in this paper, was fabricated in National Semiconductor's 0.24 $\mu$ m CMOS technology. A 20 bit pseudo-random bit sequence (PRBS) generator and checker were also integrated. Figure 6 shows the photomicrograph of the die, which is 2.6 $\times$ 1.4 mm<sup>2</sup>. The chip is packaged in a 52-pin leaded chip carrier (LDCC) with internal power

planes for impedance control. The active areas of the CDR and the complete transceiver are  $0.3\text{mm}^2$  and  $0.38\text{mm}^2$ , respectively.




Fig. 5. Current mirror phase interpolator.

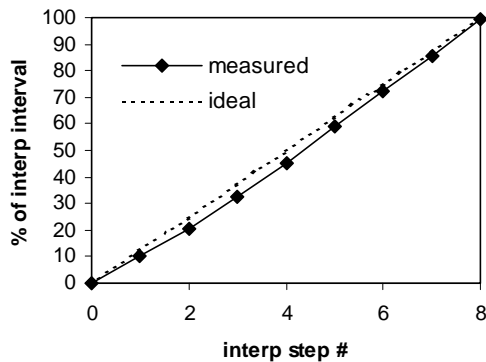


Fig. 6. Phase interpolator steps as % of the interpolating interval.




Fig. 7. Current mirror phase mux.

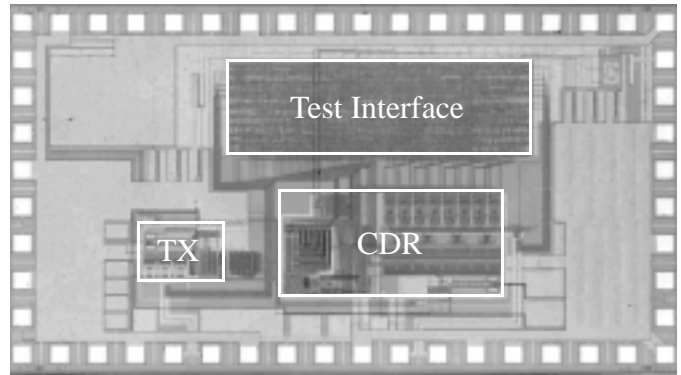


Fig. 8. Test chip die photomicrograph.

Fig. 9 shows the power consumption of the transmitter (including the PRBS encoder), the CDR (including the PRBS decoder), and the total as a function of bit rate at 50mV differential swing. The minimum supply voltage is also indicated at each point. The maximum speed of the link is 5.32Gb/s at 2.5V. At 4Gb/s, the power consumptions of the CDR and the overall transceiver are 84mW and 127mW with a supply voltage of 1.93V. The minimum operating speed of the link is 1Gb/s. Fig. 10 shows the differential eye diagram at the transmitter output. The swing is around 300-mV.

Fig. 11 shows the jitter histogram of the receive clock with a quiet supply. The p-p jitter is 38.9ps. The two peaks result from the dithering of the receive clock at steady state. When this dithering is removed by manually bypassing the timing vernier control, the p-p jitter is 15ps. Fig. 12 shows the jitter histogram of the receive clock with a 200mV 1MHz square wave noise superimposed on the power supply. The noise is generated with a supply shorting transistor [3], which also generates substrate noise. The p-p jitter is 106.7ps. The supply noise sensitivity is 0.34ps/mV. Although not implemented in this chip, a linear regulator described in [2] can be used to regulate the timing verniers to significantly reduce the supply noise sensitivity. With a transmitter differential swing of only 50mV and under 200-mV supply noise, the link operated without error at 4Gb/s for at least one day, corresponding to a BER of at least  $10^{-14}$ . The maximum measured frequency tolerance of the link in plesiochronous mode is  $\pm 400$  ppm. Table 1 summarizes the performance of the test chip.

## Conclusion

In this paper, a 4Gb/s CDR which dissipates 84mW and occupies  $0.3\text{mm}^2$  was described. A complete 4Gb/s transceiver which also includes a transmitter described in [2] dissipates 127mW and occupies  $0.38\text{mm}^2$ . The techniques we have described are amenable to large scale integration since they allow for small, low-power, and noise insensitive high speed serial link designs. Integrating 100 of these transceivers on a chip would achieve an aggregate bandwidth of 400Gb/s both into and out of a chip and requires only  $38\text{mm}^2$  and  $12.7\text{W}^1$ .

1. This is assuming multiple links do not share common blocks. In most cases, timing circuits such as DLLs and clock recovery can be shared, reducing the overall area and power consumption.

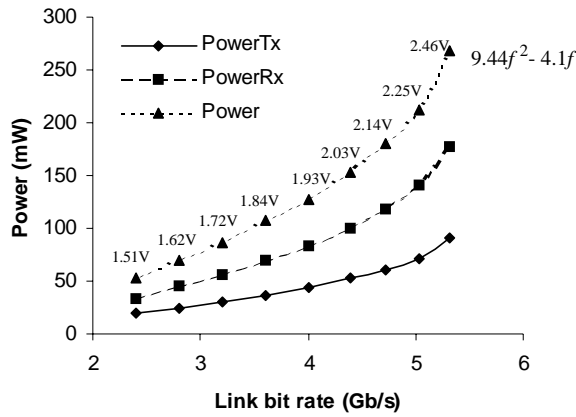


Fig. 9. Test chip power consumption.




Fig. 10. Transmitter differential eye diagram at 4Gb/s. The grid is 100-ps by 100-mV.

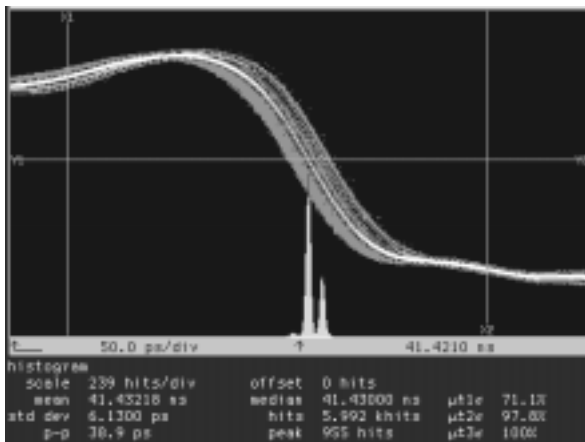


Fig. 11. Jitter histogram of the receive clock at 1GHz with quiet supply.

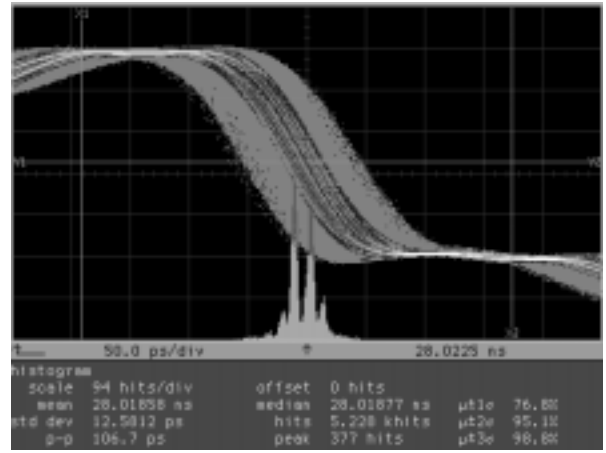


Fig. 12. Jitter histogram of the receive clock at 1GHz with 200mV square wave noise superimposed on the supply.

TABLE I  
TEST CHIP PERFORMANCE SUMMARY

Active area	Transmitter: 0.08mm <sup>2</sup> CDR: 0.3mm
Power at 4Gb/s 50mV diff. swing	Transmitter: 43mW CDR: 84mW
Maximum speed	5.32Gb/s
CDR quiet supply jitter	38.9ps
CDR supply sen. of jitter	0.34ps/mV
BER at 50mV diff. swing	< 10 <sup>-14</sup>
Frequency tolerance	±400ppm

### Acknowledgement

The authors thank Ramin Farjad-Rad and Mark Horowitz for discussions, Dean Liu and Jaeha Kim for CAD assistance, and National Semiconductor for fabricating the prototype chip. This work was supported in part by the Defense Advanced Research Projects Agency under ARPA Order E253 and monitored by the U.S. Army under Contract DABT63-96-C-0039, and in part by Intel Corporation.

### References

- [1] W. J. Dally and J. Poulton, "Transmitter equalization for 4Gb/s signaling," *Proc. Hot Interconnect*, Aug. 1996, pp. 29-39.
- [2] M.-J. E. Lee, W. J. Dally, and P. Chiang, "Low-power area efficient high speed I/O circuit techniques," *IEEE J. Solid-State Circuits*, vol. 35, pp. 1591-1599, Nov. 2000.
- [3] S. Sidiropoulos and M. Horowitz, "A semidigital dual delay-locked loop," *IEEE J. Solid State Circuits*, vol. 32, pp. 1683-1692, Nov. 1997.
- [4] Y. Moon, *et al.*, "An all-analog multiphase delay-locked loop using a replica delay line for wide-range operation and low-jitter performance," *IEEE J. Solid Circuits*, vol. 35, pp. 377-384, Mar. 2000.
- [5] G.-Y. Wei, *et al.*, "A variable-frequency parallel I/O interface with adaptive power-supply regulation," *IEEE J. Solid State Circuits*, vol. 35, pp. 1600-1610, Nov. 2000.