CMOS Transceiver with Baud Rate Clock Recovery for Optical Interconnects

Azita Emami-Neyestanak, Samuel Palermo, Hae-Chang Lee and Mark Horowitz

Computer Systems Laboratory, Stanford University
Stanford, CA 94305

Abstract

An efficient baud rate clock and data recovery architecture is applied to a double sampling/integrating front-end receiver for optical interconnects. Receiver performance is analyzed and projected for future technologies. This front-end allows use of a 1:5 demux architecture to achieve 5Gb/s in a 0.25μm CMOS process. A 5:1 multiplexing transmitter is used to drive VCSELs for optical transmission. The transceiver chip consumes 145mW per link at 5Gb/s with a 2.5V supply.
Keywords: optical interconnects, clock and data recovery, integrating receiver, I/O, double sampling, baud rate, VCSEL

Introduction

Many researchers have proposed using a large number of optical beams in parallel to achieve very high data rates for short-haul chip-to-chip communication. Bonding 2D arrays of hundreds of photo-detectors and VCSELs to silicon substrates has been demonstrated and scaling to thousands of devices is possible [1]. Such a system requires receiver and transmitter circuitry that is very small and has low power consumption. In our previous work [2], we showed that a double sampling/integrating front-end could be an excellent candidate for parallel optical interconnects due to its low power and area consumption. In this work we increase the performance of the receiver by 2.5x by using a higher multiplexing factor, and create a complete link by applying an efficient, novel baud rate clock and data recovery (CDR) scheme to the previous architecture. Our 3x3 array transceiver test-chip with VCSEL drivers for optical transmission is shown in Fig. 1.

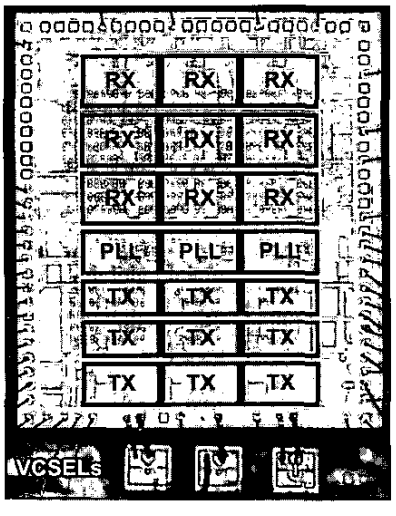

Fig. 1 Parallel optical transceiver test-chip

This transceiver uses time division de/multiplexing of 5 to support a bit-time less than 2 FO4 inverter delays. The low bit time is possible since this architecture eliminates the need for a transimpedance amplifier that runs at the bit rate. To save power, each column of three transmitters shares one PLL that generates clock phases for the multiplexing. The very last row of transmitters are wire-bonded to three VCSELs and the rest of the array is designed to have flip-chip bonded multiple quantum well (MQW) p-i-n diodes that can be used both as optical modulators for transmission and as photo-detectors at the receivers (the MQW modulator transmitters are not discussed in this paper). The under 2 FO4 bit-time yields 5Gb/s per link in our 0.25μm technology, and we show how these techniques scale to future technologies. Greater than 10Gb/s is achievable in a 0.13μm CMOS technology.

Optical Receiver

In the double sampling/integrating front-end described in [2], the optically generated current is integrated onto the parasitic capacitor of the input node and voltage-samples at the end of two consecutive bit-times are compared for data recovery. The input node is effectively AC-coupled using a negative feedback loop that subtracts a DC current equal to the average optical current. In this new design we achieve higher data rates by using five sets of samplers and five clock phases to build a 1:5 demultiplexing front-end that operates off an 8-10 FO4 clock. The multiplexing is possible since the receiver avoids a transimpedance amplifier that must run at the bit rate. The basic input receiver is shown in Fig. 2.

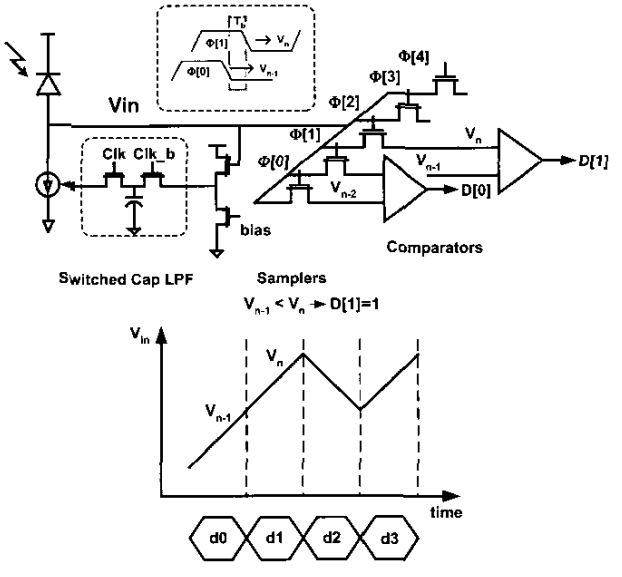

Fig. 2 Double sampling/integrating front-end receiver

 0-7803-8287-0/04/ $20.00 © 2004 IEEE

A. *Double Sampling/Integrating Front-End Analysis*

In this front-end, the data rate is limited by the bandwidth of the samplers $1/C_sR_s$, where the sampling capacitor C_s is mainly the input capacitance of the comparators and R_s is the ON resistance of the NMOS sampler in Fig. 3.

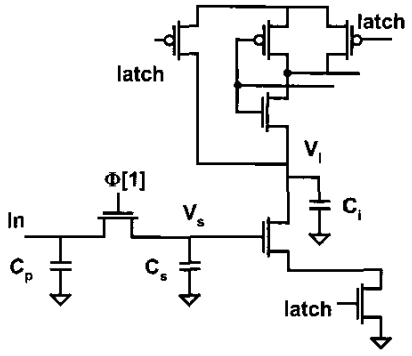

Fig. 3 Sampler and half circuit of the StrongArm latch comparator

The required optical power depends on the photodiode responsivity R, the total parasitic capacitance C_p at the input node, and the minimum required input swing, ΔV_b. $P_{op} = R^{-1}\Delta V_b(C_p + n/2\, C_s)f$ where f is the bit rate and n is the multiplexing factor. The minimum voltage swing per bit required for a certain BER of the integrating receiver is set by the voltage noise and offset of the front-end, $\Delta V_b = \pm(\sqrt{SNR}.\sigma_n + V_{off})$. The dominant source of offset V_{off} is the residual offset of the comparators after being digitally corrected [3], which in our design is roughly 2.5mV. The two main noise sources are the thermal noise of the sampler /comparator and the sampled voltage uncertainty due to clock jitter: $\sigma_n^2 = (KT/C_s) + (A^{-2}\, KT/C_i) + (\sigma_j \Delta V_b / T_b)^2$ where A is the voltage gain from V_s to V_i, σ_j is the RMS jitter and T_b is the bit period. A is between 1-3 depending on the common mode voltage with small dependence on transistor sizes and capacitances.

The electrical power consumption in this first stage (Fig. 3) can be approximated by $P_e = 3C_iVdd^2 f$ since the relative sizes of the devices are set by timing constraints. Therefore, transistor and capacitor sizing of the sampler and clocked comparator are very important and set the sensitivity (required optical power), electrical power and the bandwidth of the receiver. Fig. 4 shows how required optical energy per bit changes as a function of C_s for different values of C_i and C_p. Increasing C_s up to the point that the total input capacitance is not increased significantly can help to reduce the optical energy by decreasing the KT/C noise. For a 200fF C_p, the parasitics of our flip-chip bonded detectors, and multiplexing factor of 5, the optimum value of C_s is around 15fF. Although larger values of C_i can reduce the noise, smaller C_i is preferred for lower electrical power. Thus our test-chip comparators are sized for a C_s =15fF in 0.25µm CMOS technology, resulting in about 250fF total capacitance at the input node. As shown in Fig. 4 the optimum value of C_s is about 2x smaller if C_p is reduced to 50fF.

For our test-chip the input noise amplitude was measured by gradually increasing the offset and looking at the error rate. For a 10^{-10} error rate, the measured noise amplitude is ±6mV. The calculated sampler's KT/C noise plus comparator's input referred noise for this error rate is about ±4.8mV (0.8mV total RMS). In order to achieve a BER better than 10^{-10}, ΔV_b needs to be $\pm(6\sigma_n + V_{off}) \approx \pm 9$mV, which corresponds to 4.5fJ optical energy per bit and 22.5µW of optical power for 5Gb/s data rate. For this sizing, the electrical power consumption is 0.5mW for the first stage comparator, 0.4mW in the second stage sense amplifier and RS latch, and 0.5mW in the clock buffers (about half of it in wires) at 1GHz clock. Samplers' power is in the order of µW and negligible. This gives a total of 7mW power for the five samplers/comparators needed to support a 5 Gb/s data rate.

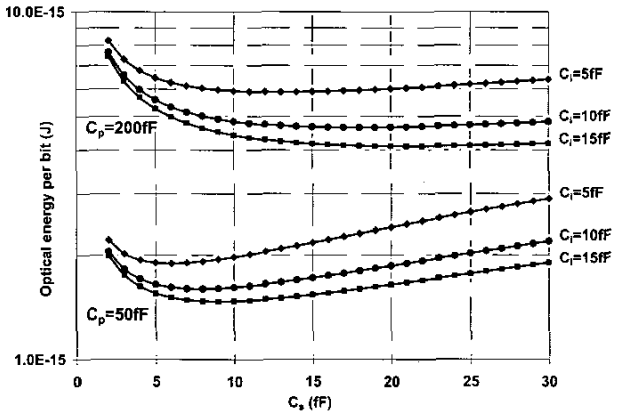

Fig. 4 Required optical energy per bit versus C_s, C_i and C_p assuming R=0.5A/W, SNR=36 (BER =10^{-10}), and A=1

B. *Scaling*

Technology scaling has two main effects on the performance of the front-end receiver: the lower operating voltages decrease the required power and the higher performance devices allow higher bit rate operation. Since the size of the input circuitry is set by KT/C noise issues, scaling does not directly change the sizes of the input devices (Fig 4). Thus, the optical energy per bit will not scale directly, and the optical power will need to increase as the bit rate scales. Constant optical power is possible if the photodiode parasitics can be scaled with technology.

If the feature sizes and supply voltage scale by α, and data rate by α^{-1}, with constant C_i, electrical power in the first stage of receiver approximately scales as α. The power consumption in the following stages (second sense amplifier and SR latch) and clock wires scale with α^2. Note that for the 0.25µm technology test-chip most of the overall power is dissipated in the clocking circuits (PLL) that also scale as α^2. The total test-chip receiver power is 75mW and will scale to about 20mW in a 10Gb/s, 0.13µm design. As technology continues to scale, supply scaling will slow when Vdd reaches around 1V. When this occurs, front-end power will increase linearly with data rate.

C. Baud Rate CDR

An interesting problem in a clocked integrating front-end is to recover the clock from the incoming data. We can apply the standard 2x oversampled technique for clock and data recovery (CDR) to our double sampling/integrating front-end by duplicating our samplers/comparators and clocking the second set with a clock shifted by half a bit period in a bang-bang control loop. The control loop adjusts the clock phase by trying to equalize the consecutive middle samples (Vm_n and Vm_{n-1} in Fig. 5a) at any transition. While this technique has the advantage of having phase correction at any data transition, it requires extra sets of samplers and clock phases that add to the area, power consumption and design complexity.

The integrating front-end allows us to create efficient baud rate CDR based only on data samples with reduced complexity and power consumption, Fig. 5b. Instead of comparing each sample V_n with a one-bit older sample V_{n-1} as done for data recovery, each data sample is compared with its two-bit older sample V_{n-2} for phase recovery (the P comparators in Fig. 5c).

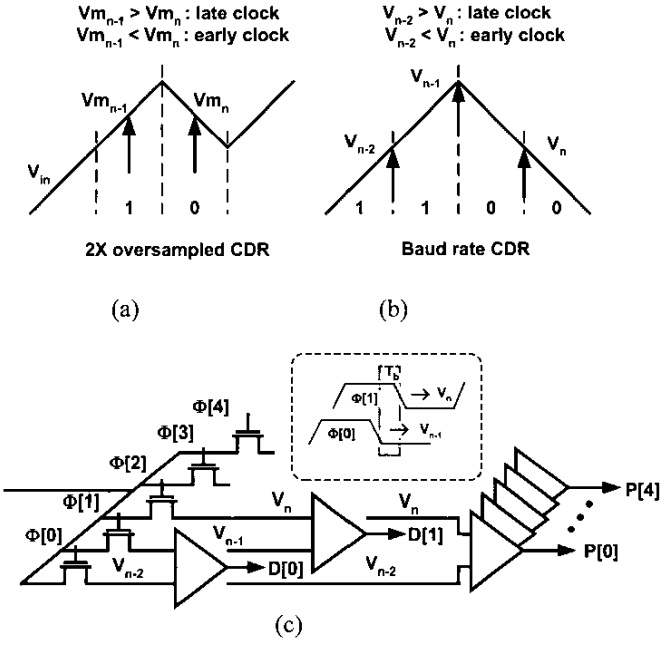

Fig. 5 CDR for double sampling/integrating receiver

The error information for the CDR loop is the difference in these two samples and the 4 bit pattern that corresponds to samples V_{n-3} to V_{n+1}. The valid patterns for phase corrections are those that give equal V_n and V_{n-2} samples when the clock is synchronized with the incoming data. "0011" and "1100" are patterns that have complete early/late phase information. Most other patterns have conditional phase information, e.g. 1101 only gives robust results when the input leads the clock. The overall phase correction probability in this CDR is 0.25 for random data, while it is 0.5 for a normal 2x oversampled system. A pattern and phase detector block can generate the up/dn command for a bang-bang phase correction loop based on this technique. The graph in Fig. 6 compares the performance of the two techniques by looking at the probability of these up/dn commands versus phase misalignment in the presence of noise and offset. Reduction in the effective gain is the main trade-off for using baud rate clock recovery.

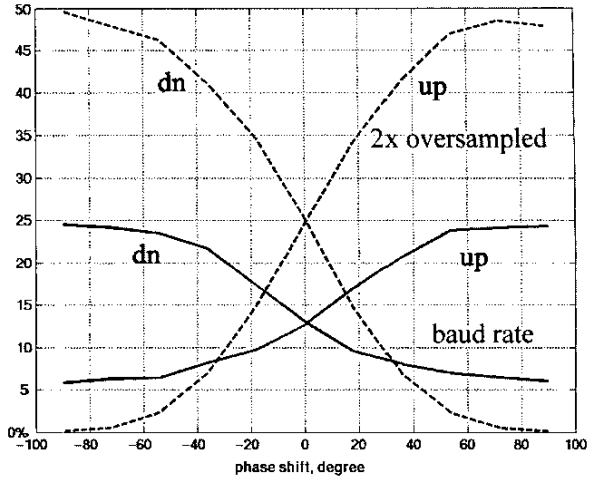

Fig. 6 Phase detector performance for 2x and baud rate CDR

In the test chip to save power we used the VCO PLL design described in [4]. The control voltage from the transmitter's voltage-controlled ring oscillator is used to set the coarse frequency level of a similar VCO at the receiver. The phase correction signals from the CDR then drive the fine control loop, as shown in Fig. 7a. Our testing results showed that for achieving acceptable jitter numbers and BER the required voltage swing per bit is higher than what we expected. The 1.0GHz recovered clock with 4.8ps RMS jitter shown in Fig. 7b corresponds to 5Gb/s data rate with about 40mV voltage swing per bit.

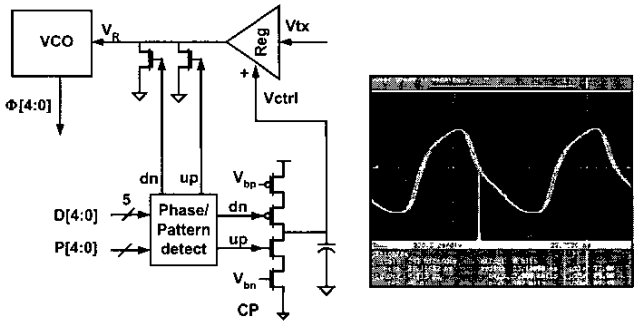

Fig. 7 CDR loop used in our test chip and recovered 1GHz clock

The main reason behind degraded jitter and higher voltage requirements is that the *SNR* for the phase comparators are much lower than the expected $\Delta V_b/\sigma_n$ for data comparators, and as shown in Fig. 6, wrong up/dn decisions are often made for small phase errors. This means that to reduce the clock jitter caused by the input, we need to heavily filter the up/dn phase correction commands before applying any phase correction to the VCO. Therefore, the bandwidth of our CDR loop should be very small, requiring a very stable, low jitter VCO. Unfortunately, this requirement makes our PLL design, with a ring oscillator, less than optimal. Future designs will

either use stable LC oscillators like the ones used in clock synthesizers [5], or use a dual loop architecture [6], where the VCO (which could be a ring oscillator) is locked to a clean reference clock and the filtered CDR drives a digital phase interpolator. While the low loop bandwidth does reduce the jitter in a dual-loop CDR, it also reduces the frequency tracking range of the PLL. Recently a number of researchers have proposed building a second order phase tracking loop [7] to improve the frequency tracking range.

Transmitter VCSEL Driver Design

This link requires a VCSEL current driver running at a bit period as low as 2 FO4 inverter delays. Fig. 8 shows the designed driver stage with a separate 4V supply (LVDD). This higher supply is necessary to support the large 1.6V DC forward voltage drop of the VCSEL. Differential drivers steer modulation current between the VCSEL and a dummy nMOS device to guarantee constant current draw from the LVDD supply to avoid cross-talk/ISI. A bias current is applied to keep the VCSEL lasing to prevent optical turn-on delay. 5:1 multiplexing is implemented directly at the low impedance VCSEL with each differential pair activated during the overlap of adjacent clock phases from the 5-stage VCO in the transmitter PLL. The output pulse-width is controlled with a feedback loop that varies the pre-drive delay using digitally-adjustable capacitive loads [4].

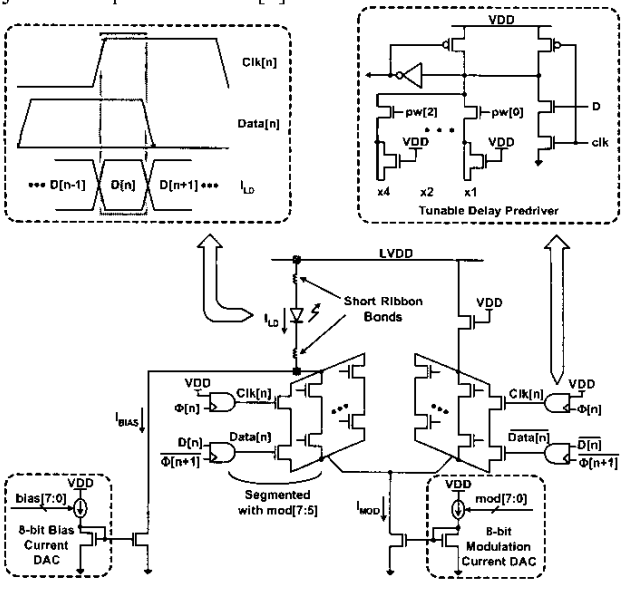

Fig. 8 VCSEL driver

VCSELs are attached with short wire-bonds to the last row of transmitters. The VCSELs used in this experiment have threshold currents of 800μA and slope efficiencies of 0.34mW/mA. Fig. 9 shows the eye diagram from a VCSEL transmitter at the maximum bit rate (5Gb/s).

Conclusion

For practical optical chip I/Os, simple, low-power CMOS interface circuits are needed. The double sampling/integrating receiver looks promising for these applications. By trading off a little sensitivity, it removes the need for gain at the bit rate, allowing use of a 1:5 demux architecture, and achieves less than 2 FO4 bit time, or 5Gb/s in our 0.25μm test-chip. The integrating nature of the input allows one to build a baud rate clock recovery by looking at voltage samples that are separated by two bits. While this phase measurement is noisy and has low gain, it is effective in a low-bandwidth CDR loop. Muxing is also used to create direct drivers for the VCSEL. The DC drop across the laser diode can be easily handled by using an additional supply, and clamping the CMOS to ensure transistors are not overstressed.

With a 2.5V power supply, the receiver consumes 75mW at 5Gb/s. Total power at the transmitter with LVDD=4V is 72mW. Scaling to a 0.13μm technology would provide a 10Gb/s link that dissipates less than 40mW per transceiver.

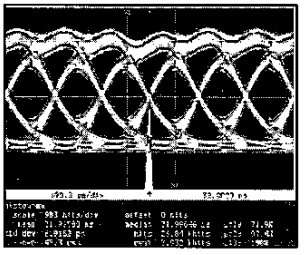

Fig. 9 VCSEL driver 5Gb/s optical eye diagram

Acknowledgements

The authors would like to thank Jaeha Kim, Elad Alon, Aparna Bhatnagar and Tim Drabik for technical discussions, National Semiconductor for fabrication of the test chip, and NSF, TI, DARPA, and the MARCO Interconnect Focus Center for funding.

References

[1] A. L. Lentine et al. "Arrays of Optoelectronic Switching Nodes Comprised of Flip-Chip-Bonded MQW Modulators and Detectors on Silicon CMOS Circuitry," *IEEE Photon. Technol. Lett.*, vol. 8, pp. 221-223, Feb. 1996.

[2] A. Emami-Neyestanak et al. "A 1.6Gb/s, 3mW CMOS Receiver for Optical Communication," in *Proc. IEEE VLSI Circuits Symp.*, pp. 84-87, June 2002.

[3] M. E. Lee et al. "Low-Power Area-Efficient High-Speed I/O Circuit Techniques", *IEEE J. of Solid State Circuits*, vol. 35, pp. 1591-1599, Nov. 2000.

[4] J. Kim and M. Horowitz, "Adaptive Supply Serial Links With Sub-1-V Operation and Per-Pin Clock Recovery," *IEEE J. Solid-State Circuits*, vol. 37, pp. 1403-1413, Nov. 2002.

[5] N. D. Dalt, and C. Sandner, "A Subpicosecond Jitter PLL for Clock Generation in 0.12μm Digital CMOS", *IEEE J. Solid-State Circuits*, vol. 38, pp. 1275-1278, Jul. 2003.

[6] S. Sidiropoulos and M. Horowitz, "A semidigital dual delay-locked loop", *IEEE J. of Solid-State Circuits*, pp 1083-1092, Nov. 1997.

[7] M. E. Lee et al. "A second-order semi-digital clock recovery circuit based on injection locking", *ISSCC Digest of Technical Papers*, pp 1-8, Feb. 9-13, 2003