

A 90 nm CMOS 16 Gb/s Transceiver for Optical Interconnects

Samuel Palermo, *Member, IEEE*, Azita Emami-Neyestanak, *Member, IEEE*, and Mark Horowitz, *Fellow, IEEE*

Abstract—Interconnect architectures which leverage high-bandwidth optical channels offer a promising solution to address the increasing chip-to-chip I/O bandwidth demands. This paper describes a dense, high-speed, and low-power CMOS optical interconnect transceiver architecture. Vertical-cavity surface-emitting laser (VCSEL) data rate is extended for a given average current and corresponding reliability level with a four-tap current summing FIR transmitter. A low-voltage integrating and double-sampling optical receiver front-end provides adequate sensitivity in a power efficient manner by avoiding linear high-gain elements common in conventional transimpedance-amplifier (TIA) receivers. Clock recovery is performed with a dual-loop architecture which employs baud-rate phase detection and feedback interpolation to achieve reduced power consumption, while high-precision phase spacing is ensured at both the transmitter and receiver through adjustable delay clock buffers. A prototype chip fabricated in 1 V 90 nm CMOS achieves 16 Gb/s operation while consuming 129 mW and occupying 0.105 mm².

Index Terms—Clock and data recovery, equalization, laser driver, optical interconnects, optical receiver, serial transceiver, VCSEL.

I. INTRODUCTION

INTEGRATED circuit scaling has enabled a huge growth in processing power which necessitates a corresponding increase in inter-chip communication bandwidth [1]. This trend is expected to continue, requiring both an increase in the per-pin data rate and the I/O number, as shown in the current ITRS roadmap (Fig. 1). While high-performance I/O circuitry can leverage the technology improvements that enable increased core performance, unfortunately the bandwidth of the electrical channels used for inter-chip communication has not scaled in the same manner. Thus, rather than being technology limited, current high-speed I/O link designs are becoming channel limited. In order to continue scaling data rates, link designers implement sophisticated equalization circuitry to compensate for the frequency dependent loss of the bandlimited channels [3]–[5]. With this additional complexity comes both power and area costs, which will make it difficult to achieve the roadmap targets in a realistic power budget.

Manuscript received October 11, 2007; revised January 17, 2008. This work was supported by MARCO-IFC. Chip fabrication was provided by CMP and STMicroelectronics.

S. Palermo was with the Department of Electrical Engineering, Stanford University, CA 94305 USA. He is now with Intel Corporation, Hillsboro, OR 97124 USA (e-mail: spalermo@vlsi.stanford.edu).

A. Emami-Neyestanak is with the Department of Electrical Engineering, California Institute of Technology, Pasadena, CA 91125 USA.

M. Horowitz is with the Department of Electrical Engineering, Stanford University, Stanford, CA 94305 USA.

Digital Object Identifier 10.1109/JSSC.2008.920330




Fig. 1. I/O scaling projections [2].

A promising solution to this I/O bandwidth problem is the use of optical inter-chip communication links. The negligible frequency dependent loss of optical channels provides the potential for optical link designs to fully leverage increased data rates provided through CMOS technology scaling without excessive equalization complexity. Optics also allows very high information density in both free space systems [6]–[8], with the ability to focus short wavelength optical beams into small areas without the crosstalk issues of electrical links, and in fiber based systems, with the added dimension of wavelength division multiplexing (WDM) [9].

In order for optical interconnects to become viable alternatives to established electrical links, they must be low cost and have competitive energy (mW/(Gb/s)) and area efficiency metrics. While significant work has been done on optical transceivers (Table I), many of these designs are implemented in processes that are more expensive than standard CMOS and/or are not competitive in energy efficiency to electrical link solutions for short distances. Also, these optical transceivers often neglect the power required for data (de)serialization and clock generation/recovery, leading to an incomplete comparison against electrical link systems. The required improvements in cost, area, and energy efficiency motivate an increased level of integration, combining the optical front-ends with the data serialization and clocking circuitry.

This paper describes a dense, low-power, full optical transceiver cell developed in 90 nm CMOS which is capable of 16 Gb/s operation and achieves an energy efficiency of 8.1 mW/(Gb/s). Section II outlines the transceiver architecture, which includes optical front-end circuitry that address key

TABLE I
OPTICAL TRANSCIVER PERFORMANCE COMPARISON

Channel Data Rate (Gb/s)	Energy Efficiency (mW/Gb/s)	Serialization/Retiming	Technology	Reference
0.5	158	Yes	0.5 μ m SOS CMOS	[10]
1	70.3	No	GaAs MESFET	[11]
8.5	417	No	SiGe BiCMOS	[12]
10	87.5	No	0.13 μ m SOI CMOS	[13]
10	19.2	No	SiGe BiCMOS	[14]
10	5	No	0.13 μ m CMOS	[15]
10	2.5	No	80nm CMOS	[16]
16	8.1	Yes	90nm CMOS	This work

issues associated with vertical cavity surface-emitting laser (VCSEL) bandwidth and reliability tradeoffs and achieving adequate receiver sensitivity in low-voltage CMOS. A presentation of a four-tap current summing FIR transmitter which extends VCSEL data rate for a given average current and corresponding reliability level follows in Section III. Section IV discusses an integrating and double-sampling optical receiver architecture [17] which enables low-voltage operation suitable for modern and future CMOS technologies. A description of the clock generation and recovery circuitry which produces low-noise clocks with the high-precision phase spacing required by the time-division multiplexing architecture is given in Section V. Section VI details the full transceiver experimental results, and Section VII summarizes the work with a comparison to state-of-the-art electrical links.

II. TRANSCIVER ARCHITECTURE

The optical interconnect transceiver architecture is shown in Fig. 2 [18]. In order to enable short bit periods without consuming excessive area and power in clock generation and distribution, multiple clock phases are employed to create a multiplexing architecture at both the transmitter and receiver. At the transmitter side, a supply-regulated ring oscillator is used in the frequency synthesis phase-locked loop (PLL) [19] to provide five sets of complementary clock phases spaced a bit period apart which switch a five-to-one multiplexer. This allows a 16 Gb/s serial data stream to be produced with only 3.2 GHz clock phases. The multiplexer serial output is buffered by the VCSEL driver output stage [20], which consists of a four-tap current-mode FIR filter that equalizes the VCSEL response at high data rates. At the receiver side, a low-voltage integrating and double-sampling front-end performs data demultiplexing directly at the input node using five uniform clock phases from the clock and data recovery (CDR) system. Clock recovery is performed with a dual-loop architecture which employs baud-rate phase detection and feedback interpolation to achieve reduced power consumption. High-precision phase spacing is ensured at both the transmitter and receiver through adjustable delay clock buffers applied independently on a per-phase basis that compensates for circuit and interconnect mismatches.

III. VCSEL TRANSMITTER

Total VCSEL bandwidth is limited by a combination of electrical parasitics and the electron-photon interaction dynamics. The laser diode's dominant electrical time constant comes from

the bias-dependent junction RC , with the dominant junction capacitor value typically between 0.5–1 pF for 10 Gb/s class 850 nm VCSELs [21], [22]. In addition to the bias-dependent junction resistance, there is also significant series resistance due to the large number of distributed Bragg reflector (DBR) mirrors used for high reflectivity, with a total device series resistance typically between 50 to 150 Ω .

VCSEL optical bandwidth is regulated by two coupled differential equations which describe the electron-photon interaction [23]. Derived from these rate equations, the VCSEL relaxation oscillation frequency ω_R , which is proportional to the effective bandwidth, is directly proportional to the square root of the injected current above the threshold current I_{TH}

$$\omega_R \propto \sqrt{I - I_{TH}}. \quad (1)$$

Combining an electrical parasitic model with the optical rate-equation model yields the total frequency response of a 10 Gb/s class VCSEL, shown in Fig. 3 [22].

Output power saturation due to self-heating [24] and also device lifetime concerns [25] restrict excessive increase of VCSEL average current levels to achieve higher bandwidth. VCSEL reliability potentially poses a series impediment to very high-speed modulation, as the mean time to failure (MTTF) is

$$\text{MTTF} = \frac{A}{j^2} e^{\left(\frac{E_A}{k}\right)} \left(\frac{1}{T_j} - \frac{1}{373}\right) \quad (2)$$

where A is a proportionality constant dependent on the type of interconnect, j is device current density, E_A is the activation energy (typically 0.7 eV), and T_j is the junction temperature [26].

The conflicting dependencies of VCSEL bandwidth and reliability on device current yield the following steep tradeoff:

$$\text{MTTF} \propto \frac{1}{(\text{VCSEL Bandwidth})^4}. \quad (3)$$

Thus, in order to ease this tradeoff, an equalizing FIR output stage is used to extend the data rate for a given average current. While the VCSEL's varying frequency response with current limits the performance of a linear equalizer for large signal modulation, the frequency response variations diminish with increasing average current due to the square root relationship and a linear equalizer is effective in canceling intersymbol interference (ISI).

Fig. 4(a) shows the VCSEL transmitter with a four-tap equalizer consisting of one pre-cursor, one main, and two post-cursor




Fig. 2. Optical transceiver architecture.




Fig. 3. Modeled 10 Gb/s class VCSEL frequency response [22].

taps implemented by summing current sources at the output node. Five parallel data bits, $D[4:0]$, are routed to the taps, where they are shifted one bit time with respect to the clock phases to implement the necessary filter delays. At each tap, a pseudo-differential multiplexer serializes the five parallel input bits and drives a differential output stage which steers current between the VCSEL and dummy diode-connected thick-oxide nMOS devices that are connected to a separate 2.8 V LV_{dd} supply. This higher supply is necessary to support the 1.5 V VCSEL knee voltage. A static DC current source, I_{DC} , is also

used to bias the VCSEL above the threshold current to insure adequate bandwidth. This bias current and the leakage current from the tap driver transistors, I_{leak} , provide sufficient voltage drop across the VCSEL and dummy load to prevent excessive voltage stress on the output stage transistors.

As shown in Fig. 4(b), at each tap the five two-transistor multiplexing segments are switched with pairs of complementary clock phases spaced a bit time apart in order to form a current pulse that defines the data bit. Tunable delay predrivers, which compensate for clock static phase offsets and duty cycle errors, qualify the clocks with the data and provide buffering to drive the multiplexing segments. Eight-bit current mirror DACs bias the output stages to the desired current value. Because of the smaller current requirements of the pre/post-cursor taps, their muxes and output stages are set to one-fourth the size of the main tap to save power.

Fig. 5 shows measured optical eye diagrams at 16 Gb/s from a 10 Gb/s-class commercial VCSEL with an average current of 6.2 mA and a 3 dB extinction ratio. The four-tap equalization improves vertical eye opening by 45% while maintaining the same average operating current, and thus the same level of VCSEL reliability. While optimizing the equalizer tap values for maximum vertical eye opening resulted in overall improved link margin, the symbol-spaced equalization does introduce slightly more jitter ($\sim 8\%$ UI). It is possible to reduce this jitter at the expense of vertical eye opening improvement by adjusting the symbol-spaced tap values to co-optimize for both horizontal and vertical eye opening. While further improvement is possible by altering the architecture to include half-symbol-spaced taps




Fig. 4. VCSEL transmitter. (a) Four-tap equalizer. (b) Tap multiplexer and output stage schematic.

dedicated to canceling edge ISI, this was deemed not worthy of the additional equalization complexity and power consumption.

The maximum data rate (minimum 80% vertical eye opening) versus average VCSEL current with and without equalization is shown in Fig. 6. At 14 Gbps, equalization allows the VCSEL to run at 35% less average current, which due to the fourth-order power dependence results in a potential 138% increase in VCSEL lifetime. The four-tap equalization extends the maximum data rate from 14 to 18 Gbps before exceeding driver current levels.

IV. OPTICAL RECEIVER

In traditional optical receiver front-ends, a transimpedance amplifier (TIA) converts the photocurrent into a voltage and is followed by limiting amplifier stages which provide amplification to levels sufficient to drive a high-speed latch for data recovery. Excellent sensitivity and high bandwidth can be achieved by TIAs that use a negative feedback amplifier to reduce the input time constant [11], [13], [27]. Unfortunately,

while process scaling has been beneficial to digital circuitry, it has adversely affected analog parameters such as output resistance which is critical to amplifier gain. Another issue arises from the inherent transimpedance limit [28], which requires the gain-bandwidth of the internal amplifiers used in TIAs to increase as a quadratic function of the required bandwidth in order to maintain the same effective transimpedance gain. While the use of peaking inductors can allow bandwidth extension for a given power consumption [27], [28], these high-area passives lead to increased chip costs. These scaling trends have reduced TIA efficiency, thereby requiring an increasing number of limiting amplifier stages in the receiver front-end to achieve a given sensitivity and leading to excessive power and area consumption.

A receiver front-end architecture that eliminates linear high-gain elements, and thus is less sensitive to the reduced gain in modern processes, is the integrating and double-sampling front-end developed by Emami [17]. The absence of high-gain amplifiers allows for savings in both power and area and makes the integrating and double-sampling architecture advantageous




Fig. 5. 16 Gb/s optical eye diagrams from four-tap VCSEL TX.




Fig. 6. VCSEL transmitter maximum data rate versus average current.

for chip-to-chip optical interconnect systems where retiming is also performed at the receiver.

The integrating and double-sampling receiver front-end, shown in Fig. 7, demultiplexes the incoming data stream with five parallel segments that include a pair of input samplers, a buffer, and a sense amplifier. Two current sources at the receiver input node, the photodiode current and a current source that is feedback biased to the average photodiode current, supply and deplete charge from the receiver input capacitance, respectively. For data encoded to ensure DC balance, the input voltage

will integrate up or down due to the mismatch in these currents. A differential voltage, ΔV_b , that represents the polarity of the received bit is developed by sampling the input voltage at the beginning and end of a bit period defined by the rising edges of the synchronized sampling clocks $\Phi[n]$ and $\Phi[n + 1]$ that are spaced a bit-period, T_b , apart. This differential voltage is buffered and applied to the inputs of an offset-corrected sense amplifier [29] which is used to regenerate the signal to CMOS levels.

The use of multiple receiver segments clocked with multiple sampling phases spaced a bit period apart allows for demultiplexing of the serial data stream directly at the input node. Input demultiplexing provides an increase in the achievable data rate by reducing the receiver clocks frequency and the individual receiver segments bandwidth by the demultiplexing factor. While one receiver segment is in sampling mode, the sense amplifiers in the other receiver segments have time to resolve the data and pre-charge, allowing for continuous data resolution. As in the transmitter, a demultiplexing factor of five is used.

While in a previous implementation [17] ΔV_b was applied directly to the sense amplifier for data regeneration, the reduced supply voltage that comes with modern CMOS technologies causes the integrating input to exceed the sense-amp input range. In order to fix the sense amplifier common-mode input level and buffer the sensitive sample nodes from kickback charge, a differential buffer is inserted between the samplers and the sense-amp. The power penalty of the additional buffer is quite small ($250 \mu\text{W}$ per segment), as buffer gain is low to avoid sense amplifier offset saturation and bandwidth requirements are relaxed due to input demultiplexing.

Due to the front-end's integrating nature, the receiver sensitivity is a strong function of the bit period, total input capacitance C_{in} , and photodiode responsivity, ρ . The receiver sensitivity can be expressed as

$$P_{avg} = \frac{I_b}{\rho} = \frac{\Delta V_b C_{in}}{\rho T_b} \quad (4)$$

where P_{avg} is the minimum average optical power that generates the integrating current per bit I_b sufficient for a given bit error rate (BER).

The input capacitance consists of

$$C_{in} = C_{pd} + C_{int} + 2nC_s \quad (5)$$

where C_{pd} is the photodetector capacitance, C_{int} is the input interconnect capacitance, n is the demultiplexing factor (5), and C_s is the total hold capacitance for each sampler. Note that while only half the samplers are active at one time, (5) includes the factor of $2C_s$ which accounts for the equal number of phase samplers required for the clock recovery system discussed in Section V.

The required ΔV_b is set by input referring the sum of the residual sense amplifier offset after correction, V_{offset} , and the voltage necessary for the sense amplifier to correctly resolve at a given data rate, V_{min} . In addition, a minimum signal-to-noise ratio (SNR) must be maintained in order to achieve a given BER and the interference associated with the average current




Fig. 7. Integrating and double-sampling receiver front-end.

variation, $V_{\Delta I_{avg}}$, must be accounted. Combining these terms results in a total minimum voltage swing per bit of

$$\Delta V_b = \sqrt{\text{SNR}}\sigma_n + V_{\text{offset}} + V_{\Delta I_{avg}} + V_{\text{min}} \quad (6)$$

where σ_n^2 is the total input voltage noise variance which is computed by input referring the receiver segment circuit noise and the effective clock jitter noise.

Contributing to the input referred circuit noise are the sense amplifier, buffer, and samplers in the receiver segments. The sense amplifier is modeled as a sampler with gain and has an input referred voltage noise variance of

$$\sigma_{sa}^2 = \frac{2kT}{A_{vsa}^2 C_A} \quad (7)$$

Here C_A is the internal sense amplifier node capacitance which is set to approximately 40 fF in order to obtain sufficient offset correction range. The sense amplifier gain, A_{vsa} , is estimated to be equal to near unity for the 0.9 V common-mode input level set by the buffer output, resulting in a sense amplifier voltage noise sigma of 0.45 mV_{rms}. Buffer input referred voltage noise variance is equal to

$$\sigma_{buf}^2 = 8kT \left(\frac{\gamma}{g_m} + \frac{1}{g_m^2 R_D} \right) N_{BW} \quad (8)$$

where γ and g_m are the input nMOS excess noise coefficient and transconductance, R_D is the resistor load, and N_{BW} is the

noise bandwidth. A 250 μA tail current provides sufficient transistor transconductance to achieve a buffer voltage noise sigma of 1.03 mV_{rms} and a bandwidth of 14 GHz. Sampler voltage noise variance is equal to

$$\sigma_s^2 = \frac{2kT}{C_s} \quad (9)$$

where the factor of two is due to the receiver segments' double-samplers which generate the differential input voltage to the buffer. Here C_s is approximately 10 fF, with 55% due to the buffer input capacitance and 45% due to sampler and interconnect capacitance. This results in an input sampler voltage noise sigma of 0.92 mV_{rms}.

Clock jitter also has an impact on the receiver sensitivity because any deviations from the ideal sampling time results in a reduced double-sampled differential voltage. This timing inaccuracy is mapped into an effective voltage noise on the integrated input signal with a variance of

$$\sigma_{clk}^2 = \left(\frac{\sigma_j}{T_b} \right)^2 \Delta V_b^2 \quad (10)$$

which, using the measured clock jitter, is estimated at 0.65 mV_{rms}. Combining the input referred circuit noise and effective clock jitter noise

$$\sigma_n = \sqrt{\sigma_{sa}^2 + \sigma_{buf}^2 + \sigma_s^2 + \sigma_{clk}^2} \quad (11)$$

results in a total input noise sigma of 1.59 mV_{rms}.




Fig. 8. Sense amplifier with capacitive offset correction.

In order for the receiver to achieve adequate sensitivity, it is essential to minimize the sense amplifier input-referred offset caused by device and capacitive mismatches. While the input-referred offset can be compensated by increasing the total area of the sense amplifier [30], this reduces sensitivity by increasing input capacitance and also results in higher power consumption. Thus, in order to minimize the input-referred offset while still using relatively small devices, a capacitive trimming offset correction technique is used [31]. As shown in Fig. 8, digitally adjustable pMOS capacitors attached to internal nodes *A* and *B* cause the two nodes to discharge at different rates and modify the effective input voltage, ΔV_{AB} , to the positive-feedback stage. Using this technique, an offset correction range of ± 70 mV with a residual V_{offset} of 1.15 mV is achieved. The fixed input common-mode voltage provided by the segment buffers eliminates variability in the offset correction magnitude as the input signal integrates over the input voltage range.

The average current variation is limited to less than 5% with frequency content corresponding to 8B/10B encoded data. Assuming that V_{min} is made negligible with adequate sense amplifier regeneration time, a $\Delta V_b = 11.9$ mV is required for a BER = 10^{-10} (SNR = 40.4), which results in an estimated receiver sensitivity of -9.8 dBm at 10 Gb/s with a total input capacitance of 440 fF and a photodetector responsivity of 0.5 A/W.

A wide input voltage range is necessary to maintain adequate receiver dynamic range. Improvements in the dynamic range relative to the original implementation [17] are enabled through the use of pMOS input samplers and by the additional buffers fixing the sense amplifier input voltage independent of the input and thus eliminating offset correction variability. The maximum receiver input voltage is limited to approximately 1.1 V due to incomplete sampler turn-off and excessive leakage corrupting the sampled value, while the input voltage can drop to 0.6 V before the segment buffers drop into low-bandwidth regions.

V. CLOCK RECOVERY AND PER-PHASE ADJUSTMENT

A conventional dual-loop CDR [32], with a frequency synthesis loop and a secondary phase interpolating loop, can achieve high performance from the flexibility to optimize both the frequency synthesis loop bandwidth to filter VCO jitter and the phase loop bandwidth to reduce jitter transfer from the noisy input signal. However, using a straight dual-loop CDR in an




Fig. 9. Dual-loop CDR with feedback interpolation.

input demultiplexing receiver can be costly in terms of area and power, as the number of phase muxes and interpolators equals the demultiplexing factor. In this receiver implementation, five phase muxes and interpolators are required.

A more power-efficient CDR architecture is inspired by the work of Larsson [33], who proposed placing an interpolator in the feedback divide path of a PLL in order to filter large output phase jumps that occur with the switching of the interpolator phase positions. When this concept is extended to the input demultiplexing receiver, as shown in Fig. 9 [18], the phase position of all the VCO output clocks are simultaneously adjusted with only one phase-mux/interpolator pair allowing for significant power and area savings. An additional advantage of this architecture is that the clock paths from the VCO to the input data and phase samplers are now minimized, resulting in reduced jitter accumulation. Also, the static clock paths allows for any VCO and clock distribution phase errors to be tuned out with a low-bandwidth control loop.

One issue with this feedback interpolation architecture is that now the frequency synthesis and phase tracking loops are coupled and care must be taken in setting the two loop bandwidths in order to ensure system stability. Whenever the phase recovery loop state machine updates the interpolator settings, the time for the update to be seen by the phase detector is dominated by the PLL frequency synthesis loop settling time. Thus, the bandwidth of the phase recovery loop must be much less than the frequency synthesis loop to avoid excessive dithering in the receiver clocks. Interestingly, this coincides with the filtering required for VCO noise and input jitter transfer suppression. The frequency synthesis loop bandwidth is set relatively high at 1/20th the input reference clock frequency to filter phase noise from the ring oscillator and allow the PLL to track the CDR updates, while the secondary phase loop update rate is set roughly an order of magnitude lower to suppress input jitter transfer. While a low phase update rate can reduce the CDR frequency tracking range, a potential solution to this is to modify the phase tracking loop to a second-order loop [34] to allow for higher ppm differences between transmit and receive clocks.

The integrating front-end allows for the efficient implementation of baud-rate phase detection [35]. In order to minimize




Fig. 10. Input voltage waveform with baud-rate phase detection [35].

timing offsets, a phase detector consisting of the main data receiver segments and identical phase receiver segments is implemented (Fig. 2). The baud-rate technique uses the same data detection samples for phase detection, with a digital phase signal $Ph[n]$ produced by comparing samples separated by two bit periods, V_n and V_{n+2} . As shown in Fig. 10, valid phase information is extracted for certain four-bit patterns that contain a middle transition and a maximum of one additional transition. The main advantage of baud-rate phase detection is that no quadrature ($1/2$ UI) phases are required. This saves power and area by reducing the number of distributed clock phases by a factor of two when compared to conventional $2\times$ oversampling phase detection. Also, because the same samples are used for both data and phase detection, this architecture is less sensitive to clock phase errors. The primary disadvantage is that it reduces the net update rate to 18.75% for random data due to incomplete phase information with some data patterns.

CDR performance is verified in Fig. 11, which shows receiver clock waveforms at 3.2 GHz, corresponding to a 16 Gb/s data rate. When CDR tracking is disabled, the output jitter is only a function of the frequency synthesis PLL which has 1.74 ps_{rms} jitter. When the CDR is activated to lock onto incoming data, the clock jitter increases only marginally to 1.90 ps_{rms}, implying that the CDR provides sufficient filtering of input noise.


Guaranteeing precise clock phase spacing at the critical points of transmitter multiplexing and receiver demultiplexing is required to ensure adequate link timing margins. Achieving this accuracy is nontrivial due to static phase errors that form in the clock generation and distribution circuitry from both systematic loading imbalances and random mismatches in the VCO, distribution buffers, and interconnect. In this design,

clock phase correction is achieved through adjustable delay buffers with digitally controlled capacitive loads, shown in Fig. 12. As the tuning switches are activated, longer buffer delays occur due to the increased node capacitance. A mixture of both nMOS and pMOS switched capacitors is used to provide uniform rising and falling-edge delay adjustment. An example of the per-phase clock correction performance is shown with the measured phase offsets of the five 3.2 GHz receiver clocks in Fig. 13. The uncorrected clocks have phase errors that exceed 10% of the 16 Gb/s UI. These phase errors are reduced to within 2% UI when the per-phase correction is enabled.


VI. EXPERIMENTAL RESULTS

The optical transceiver was fabricated in a 90 nm standard CMOS process. Both the 850 nm VCSEL and photodetector are attached with short wirebonds, as shown in Fig. 14. The VCSEL output beam is free-space imaged to the receiver board and focused on a photodiode via a system of lenses.

Proper operation of the low-voltage integrating and double-sampling receiver is verified by observing the receiver input integrating node response to a 10 Gb/s 20 bit repeating data pattern obtained with on-die subsamplers, shown in Fig. 15. Receiver sensitivity, plotted in Fig. 16, was measured for both 8B/10B data patterns and also longer runlength data with a maximum variance of 10 bits in order to further stress the integrating receiver. Due to the integrating nature of the front-end, the required optical power increases roughly linearly from 5 to 14 Gb/s, with a sensitivity of -9.6 dBm at 10 Gb/s for a BER of 10^{-10} . At higher data rates, the required optical power increases at a greater rate primarily due to increased ISI from reflections associated with the photodiode wirebond



(a)



(b)

Fig. 11. Clock jitter performance. (a) Frequency synthesis PLL. (b) CDR recovered clock.




Fig. 12. Adjustable delay clock buffer.

connection. A sensitivity of -5.4 dBm is achieved at the maximum data rate of 16 Gb/s. When the 4.8 dB power penalty from the finite transmit extinction ratio is subtracted from the maximum 3.1 dBm average transmit power, this results in a margin of 7.9 dB at 10 Gb/s and 3.7 dB at 16 Gb/s to account for additional link losses and noise sources. It is worth noting that with a more integrated approach, such as flip-chip bonding the photodiodes, superior sensitivity numbers could be achieved due to the minimization of the inductive bondwire parasitics that degrade the ideally capacitive receiver input impedance. Using the measured receiver sensitivity, the integrating receiver




Fig. 13. Receiver clock phase correction performance.




Fig. 14. Micrograph of optical transceiver with bonded VCSEL and optical receiver with bonded photodiode.

can potentially handle runlengths of up to 40 bits at 10 Gb/s and 24 bits at 16 Gb/s. In 8B/10B data systems, the receiver has an estimated dynamic range of 8.2 dB at 10 Gb/s and 6.1 dB at 16 Gb/s.

Transceiver power consumption versus data rate is shown in Fig. 17. The power consumption scales nearly linearly with the data rate. This is mainly due to the large percentage of CMOS-style circuitry used in both transmitters and the in receiver. Also, as data rates are lowered the integrating receiver sensitivity improves, allowing for reduced transmit power or VCSEL current. At 16 Gb/s, the power is 129 mW or 8.1 mW/Gb/s. The transceiver power breakdown in Fig. 18 shows that 45% of the power is consumed in the receiver and 55% in the transmitter.

Table II summarizes the transceiver performance. The transceiver operates at a data rate of 5 to 16 Gb/s, with a nominal transmit extinction ratio of 3 dB and a maximum average optical launch power of 3.1 dBm. Total transceiver area is 0.105 mm².




Fig. 15. Integrating receiver input node response to a 10 Gb/s 20 bit repeating pattern. Note from the on-die measurement, bits 3 and 13 are somewhat distorted due to periodic noise on the subsamplers supply that is believed to not be present on the input waveform.




Fig. 16. Measured integrating receiver sensitivity versus data rate.




Fig. 17. Optical transceiver power versus data rate.

VII. CONCLUSION

This paper presented a power-efficient optical transceiver architecture which achieves high data rates and addresses issues in reliably driving optical VCSELs and low-voltage optical receiver design. The VCSEL driver eases the tradeoff between VCSEL bandwidth and reliability by employing simple transmitter equalization techniques in order to extend the effective




Fig. 18. Optical transceiver power breakdown at 16 Gb/s.

TABLE II
TRANSCIEVER PERFORMANCE SUMMARY

Technology	90nm CMOS
Supply Voltages	V _{dd} =1V, LV _{dd} =2.8V
Data Rate	5-16Gb/s
ER	3dB
Average Optical Launch Power	3.1dBm
VCSEL I _{avg}	6.2mA
RX C _{in}	440fF
RX Sensitivity, BER=10 ⁻¹⁰	
10Gb/s	-9.6dBm
16Gb/s	-5.4dBm
TX Clock Jitter @ 16Gb/s	1.90ps _{rms} , 14.0ps _{pp}
RX Clock Jitter @ 16Gb/s	1.90ps _{rms} , 15.1ps _{pp}
Area	
TX Output	0.017mm ²
TX PLL	0.013mm ²
RX Front-End	0.025mm ²
RX PLL/CDR	0.050mm ²
Total	0.105mm ²
Power Consumption @ 16Gb/s	
TX Output	48mW
TX PLL	23mW
RX Front-End	23mW
RX PLL/CDR	35mW
Total	129mW (8.1mW/Gb/s)

device bandwidth at a given average current and corresponding reliability level. An improved low-voltage integrating receiver provides adequate sensitivity in a power efficient manner by avoiding the use of linear high-gain elements whose efficiency is degraded with the reduction in both voltage headroom and intrinsic device gain associated with CMOS scaling. Further improvements in power efficiency are realized with a clock recovery system which employs baud-rate phase detection and feedback interpolation. At both the transmitter and receiver, adjustable delay clock buffers are applied independently on a per-phase basis to ensure high-precision phase spacing at the critical (de)multiplexing points.

Fig. 19 compares the energy efficiency and area performance of the optical transceiver with state-of-the-art electrical links. The optical link compares favorably due to the use of only very simple transmitter equalization. Conversely, the majority of the electrical links employ both transmitter equalization and either analog or sophisticated decision feedback equalization at the receiver. While there has been recent work on reducing link power




Fig. 19. Optical versus electrical transceiver performance comparisons. (a) Energy efficiency. (b) Circuit area.

[36], [37], these implementations have focused on moderate data rates over refined channels. In order to meet future system bandwidth demands, this approach will require extremely dense I/O architectures over optimized electrical channels that will ultimately be limited by the chip bump/pad pitch and crosstalk constraints.

The relative performance should scale well for the optical link with improved optical devices. VCSEL technology continues to evolve, with higher bandwidths [38], reduced threshold currents [39], and the development of longer wavelength devices [40] allowing for reduced forward voltages and link budget improvements due to correspondingly less fiber loss and improved photodetector responsivity. In addition, advances made in photodetectors [41], [42] allow for high responsivity at low capacitance, resulting in improved optical receiver sensitivity. In contrast, increased system bandwidth demands even more equalization and/or modulation complexity from electrical links

in order to signal at higher data rates over the bandlimited electrical channels.

ACKNOWLEDGMENT

The authors would like to acknowledge the help and support of D. Patil, B. Nezamfar, P. Chiang, and B. Gupta, CMP and STMicroelectronics for chip fabrication, ULM photonics for VCSELs, Albis Optoelectronics for photodiodes, and MARCO-IFC for funding. In addition, they would like to thank Prof. D. Miller and his research group for testing assistance. S. Palermo thanks Sh. Palermo for constant help and support.

REFERENCES

- [1] B. Landman and R. L. Russo, "On a pin versus block relationship for partitions of logic graphs," *IEEE Trans. Comput.*, vol. C-20, no. 12, pp. 1469–1479, Dec. 1971.
- [2] International Technology Roadmap for Semiconductors 2006 Update. Semiconductor Industry Association (SIA), 2006.
- [3] R. Payne *et al.*, "A 6.25-Gb/s binary transceiver in 0.13- μ m CMOS for serial data transmission across high loss legacy backplane channels," *IEEE J. Solid-State Circuits*, vol. 40, no. 12, pp. 2646–2657, Dec. 2005.
- [4] J. F. Bulzacchelli *et al.*, "A 10-Gb/s 5-tap DFE/4-tap FFE transceiver in 90-nm CMOS technology," *IEEE J. Solid-State Circuits*, vol. 41, no. 12, pp. 2885–2900, Dec. 2006.
- [5] B. S. Leibowitz *et al.*, "A 7.5 Gb/s 10-tap DFE receiver with first tap partial response, spectrally gated adaptation, and 2nd-order data-filtered CDR," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2007, pp. 228–229, 599.
- [6] G. A. Keeler *et al.*, "The benefits of ultrashort optical pulses in optically interconnected systems," *IEEE J. Sel. Topics Quantum Electron.*, vol. 9, no. 2, pp. 477–485, Mar. 2003.
- [7] J. J. Liu *et al.*, "Multichannel ultrathin silicon-on-sapphire optical interconnects," *IEEE J. Sel. Topics Quantum Electron.*, vol. 9, no. 2, pp. 380–386, Mar. 2003.
- [8] D. V. Plant *et al.*, "256-channel bidirectional optical interconnect using VCSELs and photodiodes on CMOS," *J. Lightw. Technol.*, vol. 19, no. 8, pp. 1093–1103, Aug. 2001.
- [9] D. Agarwal and D. A. B. Miller, "Latency in short pulse based optical interconnects," in *IEEE Lasers Electro-Optics Soc. Annu. Meeting (LEOS 2001)*, Nov. 2001, vol. 2, pp. 812–813.
- [10] P. Gui *et al.*, "A source-synchronous double-data-rate parallel optical transceiver IC," *IEEE Trans. Very Large Scale Integrat. (VLSI) Syst.*, vol. 13, no. 7, pp. 833–842, Jul. 2005.
- [11] V. M. Hietala *et al.*, "Two-dimensional 8x8 photoreceiver array and VCSEL drivers for high-throughput optical data links," *IEEE J. Solid-State Circuits*, vol. 36, no. 9, pp. 1297–1302, Sep. 2001.
- [12] L. A. B. Windover *et al.*, "Parallel-optical interconnects >100 Gb/s," *J. Lightw. Technol.*, vol. 22, no. 9, pp. 2055–2063, Sep. 2004.
- [13] A. Narasimha *et al.*, "A fully integrated 4 x 10 Gb/s DWDM optoelectronic transceiver in a standard 0.13 μ m CMOS SOI," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2007, pp. 42–43.
- [14] D. M. Kuchta *et al.*, "120-Gb/s VCSEL-based parallel-optical interconnect and custom 120-Gb/s testing station," *J. Lightw. Technol.*, vol. 22, no. 9, pp. 2200–2212, Sep. 2004.
- [15] L. Schares *et al.*, "Terabus: Terabit/second-class card-level optical interconnect technologies," *IEEE J. Sel. Topics Quantum Electron.*, vol. 12, no. 5, pp. 1032–1044, Sep./Oct. 2006.
- [16] C. Kromer *et al.*, "A 100-mw 4x10 Gb/s transceiver in 80-nm CMOS for high-density optical interconnects," *IEEE J. Solid-State Circuits*, vol. 40, no. 12, pp. 2667–2679, Dec. 2005.
- [17] A. Emami-Neyestanak *et al.*, "A 1.6 Gb/s, 3 mW CMOS receiver for optical communication," in *IEEE Symp. VLSI Circuits Dig.*, Jun. 2002, pp. 84–87.
- [18] S. Palermo, A. Emami-Neyestanak, and M. Horowitz, "A 90 nm CMOS 16 Gb/s transceiver for optical interconnects," in *IEEE Int. Solid-State Circuits Conf. Dig.*, Feb. 2007, pp. 44–45.
- [19] S. Sidropoulos *et al.*, "Adaptive bandwidth DLLs and PLLs using regulated supply CMOS buffers," in *IEEE Symp. VLSI Circuits Dig.*, Jun. 2000, pp. 124–127.
- [20] S. Palermo and M. Horowitz, "High-speed transmitters in 90 nm CMOS for high-density optical interconnects," in *Proc. Eur. Solid-State Circuits Conf. (ESSCIRC 2006)*, Feb. 2006, pp. 508–511.

- [21] D. Wiedenmann *et al.*, "Design and analysis of single-mode oxidized VCSELs for high-speed optical interconnects," *IEEE J. Sel. Topics Quantum Electron.*, vol. 5, no. 3, pp. 503–511, May 1999.
- [22] D. Bossert *et al.*, "Production of high-speed oxide confined VCSEL arrays for datacom applications," *Proc. SPIE*, vol. 4649, pp. 142–151, Jun. 2002.
- [23] L. A. Coldren and S. W. Corzine, *Diode Lasers and Photonic Integrated Circuits*. New York: Wiley-Interscience, 1995.
- [24] Y. Liu *et al.*, "Numerical investigation of self-heating effects of oxide-confined vertical-cavity surface-emitting lasers," *IEEE J. Quantum Electron.*, vol. 41, no. 1, pp. 15–25, Jan. 2005.
- [25] K. W. Goossen, "Fitting optical interconnects to an electrical world—packaging and reliability issues of arrayed optoelectronic modules," in *IEEE Lasers Electro-Optics Soc. Annu. Meeting (LEOS 2004)*, Nov. 2004, vol. 2, pp. 653–654.
- [26] M. Teitelbaum and K. W. Goossen, "Reliability of direct mesa flip-chip bonded VCSEL's," in *IEEE Lasers Electro-Optics Soc. Annu. Meeting (LEOS 2004)*, Nov. 2004, vol. 1, pp. 326–327.
- [27] C.-F. Liao and S.-I. Liu, "A 40 Gb/s transimpedance-AGC amplifier with 19 dB DR in 90 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2007, pp. 54–55.
- [28] S. S. Mohan *et al.*, "Bandwidth extension in CMOS with optimized on-chip inductors," *IEEE J. Solid-State Circuits*, vol. 35, no. 3, pp. 346–355, Mar. 2000.
- [29] J. Montanaro *et al.*, "A 160-MHz, 32-b, 0.5-W CMOS RISC microprocessor," *IEEE J. Solid-State Circuits*, vol. 31, no. 11, pp. 1703–1714, Nov. 1996.
- [30] M. J. M. Pelgrom, A. C. J. Duijnmaijer, and A. P. G. Welbers, "Matching properties of MOS transistors," *IEEE J. Solid-State Circuits*, vol. 24, no. 5, pp. 1433–1439, Oct. 1989.
- [31] 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, no. 11, pp. 1591–1599, Nov. 2000.
- [32] S. Sidiropoulos and M. Horowitz, "A semidigital dual delay-locked loop," *IEEE J. Solid-State Circuits*, vol. 32, no. 11, pp. 1683–1692, Nov. 1997.
- [33] P. Larsson, "A 2–1600-MHz CMOS clock recovery PLL with low-VDD capability," *IEEE J. Solid-State Circuits*, vol. 34, no. 12, pp. 1951–1960, Dec. 1999.
- [34] H. Lee *et al.*, "Improving CDR performance via estimation," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2006, pp. 1296–1303.
- [35] A. Emami-Neyestanak *et al.*, "CMOS transceiver with baud rate clock recovery for optical interconnects," in *IEEE Symp. VLSI Circuits Dig.*, Jun. 2004, pp. 410–413.
- [36] R. Palmer *et al.*, "A 14mW 6.25Gb/s transceiver in 90nm CMOS for serial chip-to-chip communications," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2007, pp. 440–441.
- [37] G. Balamurugan *et al.*, "A scalable 5–15Gbps, 14–75mW low power I/O transceiver in 65 nm CMOS," in *IEEE Symp. VLSI Circuits Dig.*, Jun. 2007, pp. 270–271.
- [38] N. Suzuki *et al.*, "1.1- μ m-range InGaAs VCSELs for high-speed optical interconnections," *IEEE Photon. Technol. Lett.*, vol. 18, no. 12, pp. 1368–1370, Jun. 2006.
- [39] S. A. Blokhin *et al.*, "Vertical-cavity surface-emitting lasers based on submonolayer InGaAs quantum dots," *IEEE J. Quantum Electron.*, vol. 42, no. 9, pp. 851–858, Sep. 2006.
- [40] M. A. Wistey *et al.*, "GaInNAsSb/GaAs vertical cavity surface emitting lasers at 1534 nm," *Electron. Lett.*, vol. 42, no. 5, pp. 282–283, Mar. 2006.
- [41] M. Yang *et al.*, "A high-speed, high-sensitivity silicon lateral trench photodetector," *IEEE Electron Device Lett.*, vol. 23, no. 7, pp. 395–397, Jul. 2002.
- [42] M. R. Reshotko, D. L. Kencke, and B. Block, "High-speed CMOS compatible photodetectors for optical interconnects," *Proc. SPIE*, vol. 5564, pp. 146–155, Oct. 2004.



Samuel Palermo (S'97–M'07) received the B.S. and M.S. degrees in electrical engineering from Texas A&M University, College Station, in 1997 and 1999, respectively, and the Ph.D. degree in electrical engineering from Stanford University, Stanford, CA, in 2007.

From 1999 to 2000, he was with Texas Instruments, Dallas, TX, where he worked on the design of mixed-signal integrated circuits for high-speed serial data communication. He is currently with Intel Corporation, Hillsboro, OR, working on high-speed optical and electrical I/O architectures. His research interests include high-speed electrical and optical links, clock recovery systems, and techniques for device variability compensation.



Azita Emami-Neyestanak (S'97–M'04) was born in Naen, Iran. She received the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1999 and 2004, respectively. She received the B.S. degree with honors in electrical engineering from Sharif University of Technology, Tehran, Iran, in 1996.

She is currently an Assistant Professor of electrical engineering at the California Institute of Technology, Pasadena, CA. She was with Columbia University, New York, NY, as an Assistant Professor in the Department of Electrical Engineering from July 2006 to August 2007. She also worked as a Research Staff Member at IBM T. J. Watson Research Center, Yorktown Heights, NY, from 2004 to 2006. Her current research areas are VLSI systems, and high-performance mixed-signal integrated circuits, with the focus on high-speed and low-power optical and electrical interconnects, synchronization, and clocking.



Mark Horowitz (S'77–M'78–SM'95–F'00) received the B.S. and M.S. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 1978, and the Ph.D. degree from Stanford University, Stanford, CA, in 1984.

He is the Associate Vice Provost for Graduate Education working on Special Programs and the Yahoo! Founders Professor of the School of Engineering at Stanford University. In addition, he is Chief Scientist at Rambus Inc. His research interests are quite broad and span using EE and CS analysis methods to problems in molecular biology to creating new design methodologies for analog and digital VLSI circuits. He has worked on many processor designs, from early RISC chips, to creating some of the first distributed shared memory multiprocessors, and is currently working on on-chip multiprocessor designs. Recently, he has worked on a number of problems in computational photography. In 1990, he took leave from Stanford to help start Rambus Inc., a company designing high-bandwidth memory interface technology, and has continued work in high-speed I/O at Stanford. His current research includes multiprocessor design, low-power circuits, high-speed links, computational photography, and applying engineering to biology.

Dr. Horowitz has received many awards including a 1985 Presidential Young Investigator Award, the 1993 ISSCC Best Paper Award, the ISCA 2004 Most Influential Paper of 1989, and the 2006 Don Pederson IEEE Technical Field Award. He is a Fellow of IEEE and ACM and is a member of the National Academy of Engineering.