

Performance-based adaptive power optimization for digital optical interconnects

Xiaoqing Wang, Fouad Kiamilev, George C. Papen, Jeremy Ekman, Ping Gui, Michael J. McFadden, Joseph C. Deroba, Michael W. Haney, and Charles Kuznia

Optical links are traditionally set to transmit maximum power for worst-case loss and consequently to dissipate more power than is required. We describe a technique to minimize power consumption based on the measured bit-error rate (BER) of the link. This technique uses a novel power-negotiation algorithm that optimizes the link power setting to achieve minimum power dissipation for a target BER. A 0.5 μm complementary metal-oxide semiconductor optical transceiver chip was fabricated, and a free-space optical interconnect system was built for validation. The results showed that the algorithm was able to find the optimum power settings for the VCSELs for a target BER and to account for dynamic changes such as variation in the optical loss in the system. © 2005 Optical Society of America

OCIS codes: 200.4650, 250.7260, 250.3140.

1. Introduction

Optical interconnects are potential candidates for high bandwidth, very short-range (VSR) communication.^{1–4} Large vertical-cavity surface-emitting laser (VCSEL) and photodetector arrays can now be reliably produced and integrated with electronics. Aggregate bandwidths of the order of terabits per second per square centimeter between chips have been reported.⁵

The high density of very large-scale optoelectronic (OE) integration permits great flexibility in increasing aggregate bandwidth. The scaling of device arrays, decreasing physical dimensions, increasing speed, and tighter integration lead to higher power densities and a nonuniform thermal environment. Fluctuations in temperature and in the manufacturing process introduce variations in threshold current and slope efficiency of the VCSEL.

Current research suggests reducing the power dissipation by using zero-biased, low–ultralow threshold laser devices.⁶ However, below-threshold biasing leads

to a significant turn-on jitter, and it may not be power efficient for high-speed digital communications.⁷ Kibar *et al.* propose a power-optimization method that takes the link design constraints as input and iterates over the design variables to find the optimum link design for a given bit rate.⁸ Design constraints may include optical system efficiency, transceiver characteristics, and system fan-out, while design variables may be the number of stages in the transceiver, the feedback resistor, and the bias voltage. This technique emphasizes circuit-level power optimization in the design stage for a system with given parameters. Here we report on an alternative approach that utilizes an adaptive power-optimization methodology achieved by integrating power management firmware with the optical transceiver.

VSR optical links are usually designed to compensate for the worst-case power requirement for a given link configuration. Some advanced commercial VSR links use various monitoring techniques to detect the effects of device aging and adjust the laser driving strength accordingly.^{9,10} However, such a compensation scheme strives to maintain constant optical output power at the transmitter, without regard to optical link length, optical receiver performance, and optical channel losses. Consequently, this scheme usually results in unnecessarily high power consumption over short optical links. In contrast, wireless links use sophisticated algorithms and circuits in both transmitters and receivers to reduce power consumption. For example, in the Bluetooth wireless standard, the transceiver power is dynamically

X. Wang (xiaoqing.j.wang@intel.com), F. Kiamilev, J. Ekman, P. Gui, M. J. McFadden, J. C. Deroba, and M. W. Haney are with the Department of Electrical and Computer Engineering, University of Delaware, Newark, Delaware 19716. G. C. Papen is with the Department of Electrical and Computer Engineering, University of California at San Diego, La Jolla, California 92093. C. Kuznia is with Peregrine Semiconductor, San Diego, California 92121.

Received 7 October 2004; revised manuscript received 22 March 2005; accepted 14 June 2005.

0003-6935/05/296240-13\$15.00/0

© 2005 Optical Society of America

adjusted, depending on the link data rate and the signal attenuation associated with the communication channel.¹¹

We present protocols and circuits with which to achieve power-saving and optimization characteristics for VSR optical data links in much the same way as is done in wireless links. Our approach focuses on optimizing the power consumption of optical links based on the bit-error rate (BER) of each optical link by use of a built-in power-negotiation algorithm. Theoretical analysis of the optical communication channels reveals the dependence of the channel BER on the modulation and bias current settings of the VCSELs. Based on this dependence, we develop an algorithm with which to adjust the power settings of the optical transmitter according to the BER feedback from the receiver. The setup is shown in Fig. 1. This algorithm converges to an optimal setting for each transmitter that results in the minimum power consumption for a given target BER.

This algorithm can be utilized not only when the system is first powered on but also during normal system operation if the link is idle or a change has been made. In addition, for a system that has built-in error-detection capability, such an algorithm can be run in the background even during real data transmission, and the accumulated BER can be used as the feedback to optimize link power consumption. Therefore it provides the system with a pseudo-real-time power-optimization capability. Here it is assumed that such capability is available or can be added to the system. For a system that lacks error-detection functions, the power-saving algorithm can be run when the system is taken off line, and the resultant optimized transmitter setting can be saved for use in actual data transmission. In such a system, some control mechanism, which can be a simple-state machine, can be used to determine the system's operating mode and ensure a smooth transition among various operating modes. The triggering events could be the idle state of the channel or any detected change in system performance. Additionally, the power-optimization algorithm can be run periodically if needed. The frequency of rerunning the algorithm depends on the requirements of each system, such as the stability of the operating environment and the characteristics of live traffic. Such specification usually needs to be made during the design stage and may be adjusted accordingly over time.

Many commercial optical receiver modules include an




Fig. 1. System setup for the power-negotiation algorithm: TX, transmitter; RX, receiver; PRBS, pseudorandom bit sequence.

electrical output that monitors average optical power. This output is typically used to set up the initial link operation and make coarse adjustments. Achieving a specific BER on that link still requires performing a BER measurement. The power-optimization algorithm presented here is based on BER measurement.

One of the major advantages of this technique is that the link BER accounts for all the factors that affect link performance, including OE device performance and electrical transceiver design as well as the attenuation associated with the noisy communication channel. Therefore a BER-based optimization approach provides the system with a simple control mechanism with which to achieve both low-power and high-performance characteristics.

Meanwhile, one of the challenges of implementing this algorithm is the long test time required for a high-confidence BER measurement for a multigigabit-per-second link with a low target BER ($<10^{-12}$) by use of the conventional test method. For example, it takes more than 28 h to perform a BER measurement that allows 256 errors to occur for a target BER of 10^{-12} on a link running at 2.5 Gbits/s. Fortunately, various accelerated BER extrapolative techniques to speed up the measurement process and reduce the test time to minutes have been proposed.¹² Recent research has demonstrated the possibility of integrating such accelerated BER measurement capability into optical transceivers.¹³ Alternatively, in some applications, various coding and decoding schemes can be used to improve the final BER performance while they reduce the actual BER required to a much higher range, such as 10^{-9} and greater.¹⁴ Incorporation of the accelerated BER measurement circuits and other control mechanisms inevitably increases design complexity and eats into the system's power budget, however. In addition, for a system that has a large number of channels, the total processing time for performing such a power-optimization algorithm can be significant. Those are the trade-offs that need to be well considered and carefully balanced during system design. Nevertheless, we believe that, for most VSR optical systems, the benefits of such a BER-based power-optimization technique can well offset these design overheads, and an overall power saving together with high performance can be achieved.

An optical transceiver chip was designed and fabricated by use of a $0.5\ \mu\text{m}$ complementary metal-oxide semiconductor (CMOS) silicon-on-sapphire process. The chip was integrated with either a VCSEL or photodetector chips by flip-chip bonding. The hybrid chips communicated with a field-programmable gate array (FPGA) in which the power negotiation algorithm was implemented. A free-space optical interconnect system was built, and the power-negotiation algorithm was tested with a 1.8 Gbit/s optical link. Our test results showed that the algorithm was able to find the optimum modulation and bias settings for the VCSELs to achieve a target BER, despite changes such as variation in the optical loss in the system.

To our knowledge, this is the first demonstration

of adaptive power optimization based on a link BER for digital optical communication. The proposed approach provides the capability of link optimization and thus facilitates efficient thermal management for densely integrated systems. The utilization of digitally tunable drivers and receivers offers individual control of OE devices, which provides an efficient way to compensate for parameter variations caused by device and thermal environment nonuniformity. The benefit of this algorithm can also be extended to OE packaging to relax optical alignment requirements and permit dynamic variation within the optical system. For example, the system can compensate for slight optical misalignment by adjusting the transmitted power while it maintains its performance.

This paper is organized as follows: In Section 2 we analyze the effect of operating conditions on BER. This analysis is followed in Section 3 by details of our power-saving algorithm and its implementation. The custom transceiver design, chip layout, and integration of optical system are introduced in Section 4, and test results are presented in Section 5.

2. Theoretical Background

For a single-channel, on-off keyed, intensity-modulated link, error-causing noise may include thermal noise, shot noise, and laser relative intensity noise. The sum of all noise sources is assumed to have a Gaussian distribution. An estimate of error probability P_e or the BER with decision threshold I_D is given by¹⁵

$$P_e = \frac{1}{4} \left[\operatorname{erfc} \left(\frac{I_1 - I_D}{\sqrt{2}\sigma_1} \right) + \operatorname{erfc} \left(\frac{I_D - I_0}{\sqrt{2}\sigma_0} \right) \right], \quad (1)$$

where I_1 and I_0 are the induced photocurrents from the detector and σ_1 and σ_0 are the root-mean-square noise values when bit 1 and bit 0, respectively, are received. In Eq. (1), erfc is the complementary error function, defined as¹⁶

$$\operatorname{erfc}(x) = \frac{2}{\sqrt{\pi}} \int_x^\infty \exp(-u^2) du. \quad (2)$$

For most VSR optical systems of interest, the channel is nearly bisymmetric and the probability that bit 1 or bit 0 will be transmitted is assumed the same. The shot noise and the laser's relative intensity noise are small because of the highly sensitive receiver used and the small VCSEL driving current (a few milliamperes). Thus it is reasonable to assume that the thermal noise dominates and limits the receiver's performance, which is a good approximation for most systems that use p-i-n detectors.¹⁵ Therefore the channel noises (σ_1 and σ_0) are independent of the data transmitted (I_1 and I_0), and we have $\sigma_1 = \sigma_0 = \sigma$. For the optimum decision threshold $I_D = (\sigma_1 I_0 + \sigma_0 I_1) / (\sigma_1 + \sigma_0) = (I_1 + I_0) / 2$, which corresponds to setting the decision threshold midway between I_1 and I_0 to

achieve the minimum BER, Eq. (1) can be simplified to

$$P_e = 0.5 \operatorname{erfc}(Q/\sqrt{2}), \quad (3)$$

where $Q = (I_1 - I_0) / (\sigma_1 + \sigma_0) = (I_1 - I_0) / 2\sigma$ is the Q factor at the decision circuit.

As the induced photocurrent (I_1 or I_0) is proportional to the optical power incident upon the detector, and the link's optical power is determined by the driving current for the VCSEL (I_{on} for data 1 and I_b for data 0) with direct intensity modulation, Eq. (3) can be written as

$$P_e = 0.5 \operatorname{erfc} \left[\frac{\gamma(I_{on} - I_b)}{\sqrt{2}} \right], \quad (4)$$

where

$$\gamma = \frac{\eta\alpha R}{\sigma_0 + \sigma_1} = \frac{\eta\alpha R}{2\sigma}. \quad (5)$$

In Eq. (5), η , α , and R are the quantum efficiency of the laser, the link loss from the laser to the detector, and the responsivity of the photodetector, respectively.

Equation (4) is true only for a VCSEL biased above its threshold current. When a VCSEL is biased below its threshold current (I_{th}), a significant turn-on delay occurs at the rising edge of the optical pulse.^{7,17-19} This delay is a result of the time that it takes for the photon population to build up in the laser cavity after the carrier density has been decaying during its off period, i.e., when it is biased below threshold. Consequently the turn-on delay varies from pulse to pulse, depending on the data rate and on the total off period before the pulse, which results in pattern-dependent jitter. An approximate form for the worst-case error rate below threshold is²⁰

$$P_e = \frac{1}{2} \operatorname{erfc} \left[\frac{\gamma(I_{on} - I_{th})h(Bt_{d,max})}{\sqrt{2}} \right], \quad (6)$$

where the bit rate, B , is usually in the range of a few gigabits per second for practical optical links. $t_{d,max}$ is the maximum turn-on delay at the rising edge of the optical pulse and can be evaluated by^{7,21}

$$t_{d,max} = \tau \ln \left(\frac{I_{on} - I_b}{I_{on} - I_{th}} \right) = \tau F, \quad (7)$$

where τ is the carrier recombination time and typically has a value of the order of 1 ns for a VCSEL. In Eq. (6), function $h(Bt)$ describes the pulse shape as a result of filtering and equalization at the receiver, which in many cases can be approximated by a raised cosine function:

$$h(Bt) = \frac{1}{2} [1 + \cos(\pi Bt)]. \quad (8)$$

Substituting Eqs. (7) and (8) into (6), we obtain

$$P_e = \frac{1}{2} \operatorname{erfc} \left[\frac{\gamma(I_{\text{on}} - I_{\text{th}}) \cos^2(\pi B\tau F/2)}{\sqrt{2}} \right]. \quad (9)$$

In this expression, τF represents the maximum time delay between the injected current pulse and the output light pulse for below-threshold conditions. The cosine term represents the reduction in the detected signal that is due to this delay. This reduction is caused by not sampling at the optimal location within the received eye where we assume that the optimal sampling time is determined for a received raised cosine pulse shape by use of the above-threshold conditions. In most systems of interest, delay τF is typically small, and the error rate is dominated by the term $I_{\text{on}} - I_{\text{th}}$.

Whereas the expression for the error rate is approximate, it models the fact that, for below-threshold conditions, the error rate increases because of pattern-dependent jitter. Therefore, one can determine a complete representation of the BER as a function of transmitter power setting by combining Eq. (4) for biasing above-threshold current with Eq. (9) for biasing below-threshold current.

To further illustrate the magnitude of the parameters described above, we take a 1.5 Gbit/s free-space optical link as an example; a VCSEL and a photodetector are used, and we assume that $\eta = 0.5$ mW/mA, $I_{\text{th}} = 1$ mA, and $\tau = 1$ ns for the VCSEL; $\alpha = 0.1$ for link loss; and $R = 0.4$ mA/mW for the detector. We choose a Q of 6 ($P_e = 10^{-9}$) and assume that $I_1 = 60$ μ A and $I_0 = 0$ μ A at the receiver. From the expression of Q , we get $(\sigma_1 + \sigma_0) = 5$ μ A and $\gamma = 4$ mA $^{-1}$. Thus, for $B\tau = 1.5$, $I_{\text{th}} = 1$ mA, and $\gamma = 4$ mA $^{-1}$, we plot Eqs. (4) and (9) in three dimensions, as shown in Fig. 2.

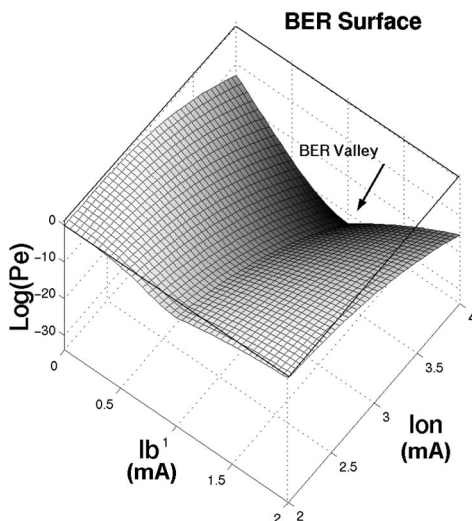


Fig. 2. BER surface as a function of transmitter power setting (I_{on} and I_b), showing the BER valley floor where the minimum BER exists for an optical link.




Fig. 3. BER contour, showing the dependency of the channel BER on the power setting (I_{on} and I_b) and the BER valley floor that corresponds to the optimum power setting. $B\tau$ is unitless, I_{th} is in milliamperes, and γ is in inverse milliamperes.

Figure 2 shows how the BER changes as the power settings are varied. The X and Y axes represent values of I_{on} and I_b , respectively, whereas the Z axis varies with $\log(P_e)$. With a fixed I_{on} , the minimum BER appears at the point when I_b is equal to I_{th} . Varying I_b from I_{th} in either direction will result in degrading the BER. For a given I_b , however, the BER decreases with increased I_{on} . In other words, a BER valley exists that is parallel to the I_{on} axis at $I_b = I_{\text{th}}$, where the lowest BER is located. Therefore this BER valley indicates that there is an optimal power setting that results in a minimum BER.

To further illustrate the dependence of BER on power setting, in Fig. 3 we have plotted the BER contours with BERs of 10^{-6} to 10^{-15} . Each curve corresponds to a certain BER, and the X and Y coordinates of any point on the curve represent the power setting (I_b and I_{on}) that results in the value of the BER. For a given BER the minimum value of I_{on} occurs when the laser is biased at threshold. The average power consumption of the laser can be approximated by

$$P_{\text{elec}} \approx 0.5 V_{\text{on}}(I_{\text{on}} + I_b), \quad (10)$$

where V_{on} is the turn-on voltage of the laser and we assume that data 1 and 0 are transmitted with equal probability. Therefore, to minimize the power consumption for a certain BER value, we need only to find the power setting along the BER contour that results in the minimum sum of I_b and I_{on} . In most cases, because I_{on} increases on both sides of threshold biasing, it can be seen that the power setting at the BER valley floor leads to the minimum sum of I_b and I_{on} , and thus to the minimum power dissipation.

The above analysis illustrates the BER as a function of I_{on} and I_b , assuming that parameters $B\tau$, γ , and I_{th} , are fixed. In the following paragraphs we investigate how these three parameters individually affect the BER contours.




Fig. 4. Effect of $B\tau$ on BER contour: $B\tau$ affects only the slopes of contours for below-threshold biasing. Units are the same as for Fig. 3.

First we look at the effect on BER contours of changing $B\tau$. Figure 4 shows a comparison between two plots with $B\tau = 1.5$ and $B\tau = 3$. The value of $B\tau$ has no effect on the curves for above-threshold biasing because they overlap exactly for both values of $B\tau$. It affects the slopes of curves only for below-threshold biasing. As $B\tau$ increases, the slope becomes greater. Therefore, for a reasonably large $B\tau$, threshold biasing remains the optimum power setting and does not vary significantly with changing $B\tau$.

However, at small $B\tau$, typically as $B\tau$ drops below 1.2, the minimum power setting will no longer be at the threshold biasing point. Instead, zero bias is the lowest-power solution.²⁰ Generally, for semiconductor lasers, the rising speed of a pulse is closely related to the relaxation oscillation frequency of the laser, which is approximately proportional to the square root of the current density above threshold. Therefore, for very high data transmission (i.e., 10 Gbits/s), the VCSEL is usually biased at well above the threshold current to boost the pulse edge rate¹⁷; this operating regime is not considered here. For the remainder of this paper it is assumed that the value of $B\tau$ is greater than 1.2, unless an exception is explicitly stated. This assumption is justified because, for values of τ of the order of 1 ns, this corresponds to a signaling rate of approximately 1.2 Gbits/s. This rate is the approximate signaling rate of current gigabit-per-second physical layer standards such as Fiber Channel and Gigabit Ethernet that use 8B10B line coding. Thus the use of the BER valley will produce minimum power for existing high-speed interconnect standards.

The scaling factor, γ , takes into account the quantum efficiency of the laser, the optical loss, and the responsivity of the photodetector. As shown in Fig. 5, the BER contours for $\gamma = 3.7 \text{ mA}^{-1}$ shift vertically to higher values of I_{on} compared with the contours for $\gamma = 4 \text{ mA}^{-1}$, indicating that, with less noise, the same error rate can be obtained with a lower current. There is no shifting horizontally along the I_b axis, because I_{th} is a constant and all slopes remain the same. For a specific value of BER, the optimum power setting




Fig. 5. Effect of γ on BER contour: Changing γ causes the contour to shift vertically, and a lower I_{on} is required for the link with a higher γ to yield the same error rate. Units are the same as for Fig. 3.

(I_{on} only) is also moved upward, while the optimum I_b remains the same.

Figure 6 shows the effects of different VCSEL threshold currents. All the BER contours shift both horizontally and vertically for an increased I_{th} . The new BER valley is aligned with the new I_{th} , and the contours for above-threshold biasing remain the same. Therefore, the optimum setting for both I_b and I_{on} needs to be changed if the threshold current changes. If more than one parameter changes, the BER contours are the superimposition of shifting caused by each individual parameter.

In summary, for any given target BER, the power setting in the BER valley floor results in minimum power consumption. If we can develop an algorithm to locate the valley floor and use the optimum setting for the transmitter, it will lead to the lowest power dissipation for this link. In addition, the algorithm needs to be able to recover the optimum power setting by reexamining the BER information when link performance is affected by changed parameters. The detailed algorithm is given in Section 3.




Fig. 6. Effect of I_{th} on BER contour: Changing I_{th} causes the contour to shift in both horizontal and vertical directions. Increasing I_{th} results in increasing of the optimum I_{on} and I_b . Units are the same as for Fig. 3.

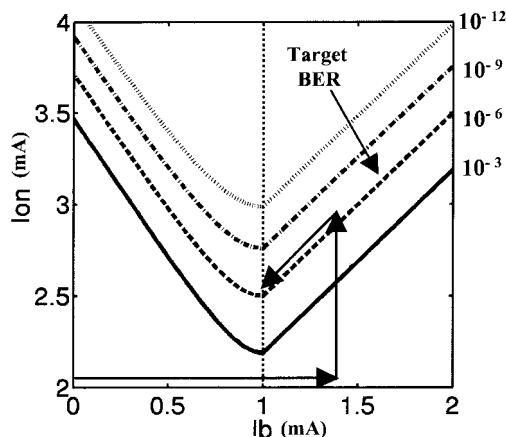


Fig. 7. Steps show the power-on optimization algorithm converging to the optimum power setting when the system is powered up.

3. Power-Negotiation Algorithm and Implementation

The complete power-negotiation algorithm is composed of two parts: a BER tester (BERT) and a decision algorithm. The BERT provides the BER information on the optical link. The decision algorithm adjusts the power (I_{on} and I_b) based on the channel BER information by following the analysis in Section 2.

The decision algorithm can be further divided into two subcategories: a power-on optimization algorithm, which can be used to set the initial optimum setting when the system is powered up, and a real-time correction algorithm, which can be used to fine-tune the setup when changes have been made that cause the BER to deviate from the target value.

A. Power-On Optimization Algorithm

The goal of the power-on optimization algorithm is to find the optimum power setting before transmission of the real data. Following the steps below will lead to the optimal setting, as illustrated in Fig. 7:

- (1) Set bias current I_b above the threshold current suggested by the device's manufacturer and set I_{on} at the low end of the operational range.
- (2) Keep I_b fixed and keep increasing I_{on} until the target error rate is achieved.
- (3) Decrease I_b and I_{on} together until the BER starts to be greater than the target BER.

For $B\tau > 1.2$, the algorithm converges to the power setting corresponding to the BER valley floor. In the case of lower data rate systems for which $B\tau < 1.2$, it may converge to a power setting that has I_b below threshold, resulting in an average power consumption that is slightly higher than the one biased at zero current.

B. Real-Time Correction Algorithm

The goal of the real-time correction algorithm is to correct the power setting to account for system parameter deviations between the data transmission periods. The power-on optimization algorithm could be rerun to yield the optimum power setting in this




Fig. 8. Steps show the real-time correction algorithm fine-tuning the power setting during data transmission. Dashed curves, original system setup; solid curves, the new BER contour after the system's operating parameters have been changed.

case. However, a fast real-time correction algorithm is needed to reduce the number of iterations and speed up the process of fine-tuning the power setting.

Various factors, such as temperature variation, attenuation changes in the optical path, and optical misalignment, can lead to operating-parameter changes. Each of these factors may cause one or more parameters to change. For example, increasing temperature will lead to an increased I_{th} and a decreased γ . Changing attenuation in the optical path affects only γ . However, when the link BER has changed in a real system, the cause of the change may be unknown. Regardless, one thing remains the same: The optimum power setting is still located at the BER valley floor. Therefore the goal remains to find the new BER valley floor and then adjust I_{on} . The following procedure describes the steps of this algorithm as illustrated in Fig. 8:

- (1) Fix current I_{on} ; slightly change I_b toward the direction of decreasing BER until the BER begins to increase again. For changing attenuation (only γ changes), this will result in the same I_b because the BER valley floor has not been affected.
- (2) Fix I_b and change I_{on} toward the target BER.

Figure 9 shows a flow chart of the negotiation algorithm's implementation. It starts with a transmitter sending a pseudorandom bit sequence (PRBS) over the optical link to a receiver. The receiver checks incoming data and records the number of errors. Once the data transmission is done, the error information is sent back to the transmitter and the decision algorithm makes the proper adjustments to the power setting for the VCSEL. Therefore each iteration involves one BER collection and one power-setting adjustment. This iteration repeats until the optimum setting has been found.

4. Hardware Implementation

A CMOS chip with four optical transceivers was fabricated, and OE devices were flip-chip bonded on top

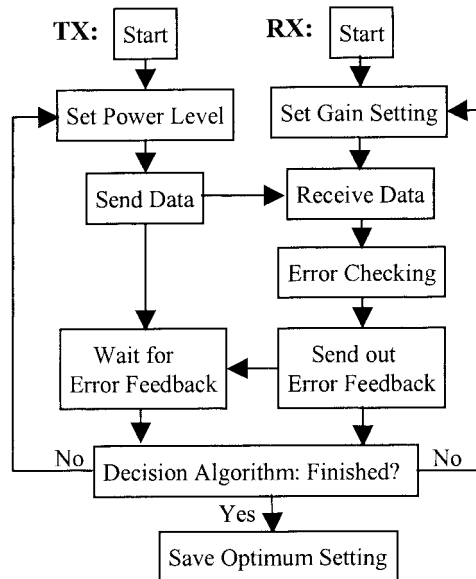


Fig. 9. Power-negotiation algorithm flow chart showing details of the implementation of both the power-on optimization algorithm and the real-time correction algorithm.

of the CMOS chip. The hybrid chip was then wire-bonded to a carrier board that was eventually integrated onto a motherboard to form a FPGA-based demonstration system. Four unidirectional optical links were established between two chips by use of free-space optics, and the power-negotiation algorithm was implemented in the FPGA to provide maximum flexibility in its development and testing.

A. Complementary Metal-Oxide Semiconductor Integrated Circuit Design

1. Vertical-Cavity Surface-Emitting Laser Driver Design

Inasmuch as the light output of a VCSEL is linearly proportional to the VCSEL's driving current above threshold, the goal of the driver design is to provide accurate tunable modulation and bias current for the VCSELs. They are obtained by use of two four-bit current-mode digital-to-analog converters (DACs) for both modulation and bias current control. The power settings are loaded and stored in register-based storage cells to provide independent controls for each VCSEL, which allows great flexibility in compensat-




Fig. 10. Schematic of a current-mode digital-to-analog converter.

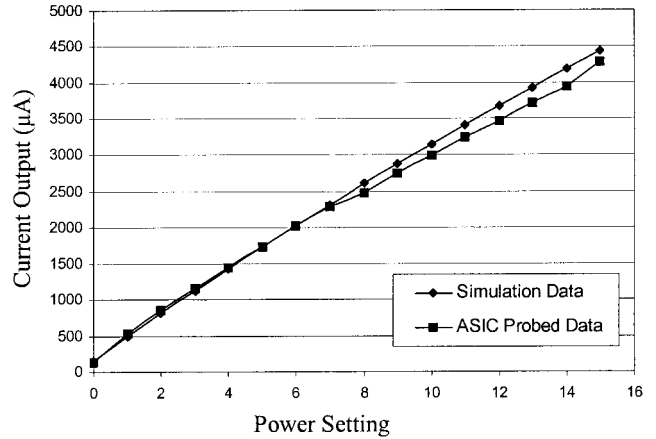


Fig. 11. A close match between simulation and measured data can be observed for DAC output current during a sweep of the digital power setting. ASIC, application-specific integrated circuit.

ing for potential device nonuniformity and parameter variation across a large OE array.

Figure 10 shows a schematic of a 4 bit current-mode DAC. It is based on a current mirror structure, and the transistors in each branch are sized proportionally and binary weighted. The reference current input, I_{ref} , is the step size for the current setting and can be adjusted externally to meet the requirements of a specific system.

Figure 11 shows a comparison of the current output of the DAC for simulation and probed data while the digital settings are swept from 0000 to 1111 with a current step size of 0.3 mA. Good linearity of current stepping is observed, as well as a close correlation between the simulation and the measurements.

2. Optical Receiver Design

The receiver is designed in such a way that it complements the operation of the transmitter. It has a large dynamic range with the ability to trade between gain and bandwidth during operation. It is composed of two DACs, two preamplifiers, a three-stage differential postamplifier, and a current-mode logic (CML) driver, as shown in Fig. 12. The purpose of using two preamplifiers is to achieve a better matching between

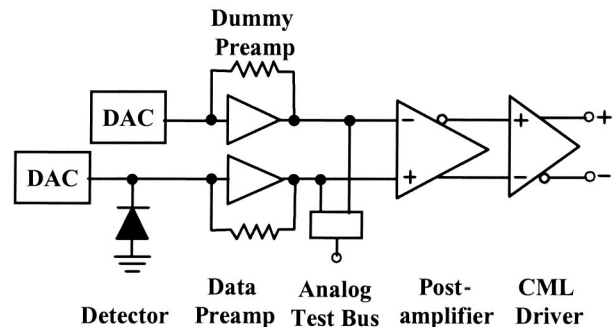


Fig. 12. Building blocks of the optical receiver design: two DACs, two preamplifiers, a three-stage fully differential postamplifier, and a CML driver.

the data and the threshold voltage channel. The use of DACs in both channels allows the dc performance of the preamplifiers to be tested without the need for integration with OE devices. In addition, during the receiver's normal operation the DAC in the dummy preamplifier is used to source a dc current to generate a threshold voltage for the decision circuit for data comparison. The analog test unit consists of a transmission gate with a 50 k Ω resistor connected in series and is designed to monitor the analog output of the preamplifiers. During the data transmission, the test bus senses the average voltage at the data channel and provides a reference voltage at which the optimum threshold voltage should be set in the dummy channel for the decision circuits. It also provides a simple and efficient way to gauge the optical power received by the detector and facilitates optical alignment during system integration. The decision circuits and postamplifiers consist of three stages of fully differential amplifiers for data recovery and waveform shaping. Finally, the CML driver is used to send the data off chip electrically. It has a 100 Ω differential output impedance and can deliver a 400 mV differential peak-to-peak output.

Figure 13 shows schematics of the transimpedance section of the preamplifier. Three digital inputs (Gmin, Gmed, and Gmax) are for gain control, which, when they are enabled, give 1, 2.5, and 10 k Ω transimpedance gain, respectively. It also has a power-control transistor, which can digitally power down the amplifier to reduce power consumption when the link is idle.

The use of the DACs and the analog test bus to enhance the receiver's testability exhibits negligible penalty on the overall noise, speed, and power consumption of the receiver. Each DAC adds a mere 6.7 fA/ $\sqrt{\text{Hz}}$ noise and less than 10 fF parasitic capacitance to the preamplifier input, which can be largely ignored in comparison with an input referred current noise of 4.1 pA/ $\sqrt{\text{Hz}}$ of the preamplifier and the




Fig. 13. Schematic of a transimpedance amplifier design, showing the power-down transistor and three digital gain-control transistors suitable for a wide dynamic range of current input.

~ 400 fF capacitance from the photodiode and packaging. The receiver has a sensitivity of 40 μA when gain setting Gmed is enabled, which leads to power consumption for the DAC that is of the order of 0.1 mW. By contrast, the complete receiver consumes an average power of 30 mW during normal operation. Each DAC takes an area of 60 $\mu\text{m} \times 18 \mu\text{m}$ in a 0.5 μm process.

3. Complementary Metal-Oxide Semiconductor Technology and Chip Layout

Our chip has been designed and fabricated by use of Peregrine Semiconductor Corporation's 0.5 μm , ultrathin, silicon-on-sapphire (UTSi/SOS) process.²² UTSi is a newly developed silicon-on-insulator technology in which devices are made in a very thin silicon layer (100 nm) on top of an insulating sapphire substrate. Compared with bulk silicon technology, it offers a number of advantages for mixed-signal circuit designs, including low power consumption, reduced parasitic capacitance, and minimum cross talk. The sapphire substrate is also a good thermal conductor, and the thermal expansion coefficient of gallium arsenide (GaAs) more closely matches sapphire than bulk silicon does. In addition, the optically transparent substrate facilitates the flip-chip attachment of the top-emitting OE devices, which makes it an ideal candidate for OE integrated circuit (IC) designs.

In an area of 4.4 mm \times 1.5 mm, four single-ended optical transceivers are implemented with 64 perimeter pads and 16 pairs of flip-chip bonding pads in the center for OE array integration. Although the chip is designed as a full transceiver, because of physical constraints only two 1 \times 4 OE arrays—either VCSEL or photodetector arrays—are attached to a given chip. As a result, the system demonstration uses two chips and the optical links between the two are unidirectional. Because the OE arrays have a pitch of 250 μm while our optical transceivers are spaced at 500 μm , only half of the OE devices in two 1 \times 4 arrays are actually connected to the transceivers and are active during normal operation. A micrograph of the chip is shown in Fig. 14. The driver and receiver circuits are 650 $\mu\text{m} \times 75 \mu\text{m}$ and 740 $\mu\text{m} \times 77 \mu\text{m}$, respectively.

The chip was manually laid out by use of a custom layout tool from Cadence Design Systems, and transistor-level simulation was performed with Spectre simulator. Special care was taken at the layout level to ensure signal integrity. Power and ground




Fig. 14. Micrograph of a 1 \times 4 CMOS optical transceiver IC.

rails were separated for analog and digital circuits, and abundant decoupling capacitors were added to minimize the switching noise and the cross talk through the power-ground network. An interleaving scheme, a dual-gate structure, and dummy gates were used for better matching for the preamplifiers and the differential circuits. Also, on-chip electrostatic discharge circuitry was included for both the perimeter and the flip-chip bonding pads.²³

4. Design Overhead for Integrated Power-Negotiation Algorithm

Because of the limited chip area available in our fabrication, the BERT and the decision algorithm are not integrated into our optical transceiver chip. In practice, it is possible to incorporate them into an optical transceiver chip. For example, parallel electrical transceivers that include on-chip BERTs are commercially available.²⁴ Typically, a BERT consists of a PRBS generator and a verifier. For a 15 bit sequence, it could take approximately $300\ \mu\text{m} \times 400\ \mu\text{m}$ in a $0.5\ \mu\text{m}$ process and consume 140 mW of power at 1.5 Gbits/s, based on our simulation. The cost of the decision algorithm is of the order of a few hundred gates. In addition, one power-negotiation algorithm implementation can be shared among many channels to further reduce the design overhead.

As described in Section 3, this power-negotiation algorithm requires a bidirectional path: a forward data path and an error feedback path. However, the error feedback does not require high-speed links, and one feedback link can be shared among multiple receivers. In our case, an electrical feedback path was used to pass the error information.

B. System Integration

1. Integration of a Complementary Metal-Oxide Semiconductor Chip with an Optoelectronic Device

The OE devices used are an Emcore high-performance 1×4 VCSEL and 1×4 GaAs p-i-n photodiode arrays on a $250\ \mu\text{m}$ pitch. The VCSEL has a nominal wavelength of 850 nm, a typical slope efficiency of 0.45 mW/mA, and a differential resistance of $50\ \Omega$. The p-i-n photodiodes have a typical responsivity of 0.5 A/W and a capacitance of 0.4 pF. They are both rated for 3.125 Gbit/s applications.²⁵

Because flip-chip bonding offers low parasitic capacitance and high integration density, it was chosen for the integration of the CMOS chip and the OE devices, which was done by the Peregrine Semiconductor Corporation. Figure 15 shows the CMOS chip with two photodetector arrays attached. As both VCSEL and photodetector arrays are top-emitting devices, the laser beams emit through the sapphire substrate of the CMOS chip during their operation.

2. Chip-on-Board Packaging

A $5\ \text{cm} \times 3.8\ \text{cm}$ carrier printed-circuit board (PCB) was designed and fabricated. The hybrid CMOS chip with the OE arrays attached was placed in the center of the carrier PCB and wire-bonded to it. Directly



Fig. 15. Micrograph of the chip-on-board packaging, showing that two photodetector arrays are attached to the transceiver chip by the flip-chip process and the hybridized CMOS chip is wire-bonded to the carrier board.

underneath the hybrid CMOS chip, a $4.5\ \text{mm} \times 1.2\ \text{mm}$ opening was made on the carrier PCB for optical access.

3. Motherboard

A 10 layer, $21.6\ \text{cm} \times 21.6\ \text{cm}$ FR4 motherboard was designed and fabricated for final system integration. Two carrier boards were connected perpendicularly to the motherboard simultaneously by Mictor impedance-controlled connectors. Right-angle metal clamps were used to stabilize the attachment. The whole setup was assembled on top of an optical breadboard, which provided good vibration isolation and enhanced system stability. The complete demonstration system is shown in Fig. 16.

A high-performance FPGA (XC2VP7) from the Xilinx Virtex-II Pro family was used as a central control unit. It has eight full-duplex multigigabit serial transceiver cores (RocketIO, Xilinx, Inc.) that include an embedded serializer-deserializer, clock-data-recovery circuits, and the options of an 8B10B and cyclic redundancy check (CRC) encoder and decoder.




Fig. 16. Overview of the demonstration system, showing the integration of two carrier boards, a Virtex II Pro FPGA, and the free-space optics mounted on micropositioning stages.




Fig. 17. Block diagram of the test setup overview. The power-negotiation algorithm is implemented in a FPGA and communicates with the optical transceiver ICs through the embedded RocketIO serial transceivers at a data rate of 1.8 Gbits/s.

The CML input–output connections are used for high-speed data transmission and also provide programmable impedance termination, a pre-emphasis level, and differential output swing. All these features provide great flexibility in system testing.

4. Free-Space Optical Interconnection

A free-space optical interconnection system was used to establish communication links between two chips. Two seven-element Universe Kogaku f -1.2 lenses were used to link a pair of carrier PCBs, separated by 89 mm. A CCD camera permitted direct observation of both the VCSEL and detector arrays to facilitate alignment. The voltage drop of the receiver preamplifier output through the analog test bus was monitored during active alignment, when the VCSELs were turned on. Owing to the large disparity between the resolution of the incident spots (full width at half-maximum, $\approx 50 \mu\text{m}$) and to the separation of the detectors ($500 \mu\text{m}$), there was no detectable optical cross talk between adjacent channels. In addition, with only four channels and little scattering, there was not a significant level of ambient light.

5. Experimental Results

A. Test Setup Overview and Initial System Testing

The goal of this test was to see how well the power-negotiation algorithm works in a practical serial data transmission system, in which serializer–deserializer and clock-data-recovery circuits are normally used. Here, we used the FPGA with integrated serial transceivers to fulfill such test needs. All experiments were carried out based on the custom-built test bed described above, and a block diagram of the system is shown in Fig. 17.

A free-space optical system was constructed to establish optical links between the transmitter and receiver chips. The FPGA sources and sinks data through the RocketIO serial transceivers and also configures the optical transceiver chips through low-speed lines. The BERT and the decision algorithm




Fig. 18. Eye diagram of the receiver output at 2 Gbits/s during the link test.

were implemented on the FPGA, and a custom interface was developed to permit communication between the RocketIO transceiver and the test logic. To reflect the actual performance of the data link, the 8B10B and CRC coding and decoding options of the RocketIO were bypassed and disabled during the testing.

Extensive testing was performed to verify the proper functioning of the BERT logic, the quality data transmission between the RocketIO transceivers, the stability of the free-space optic interconnects, and the performance of the optical transceivers.

We verified the error-detection capability by introducing artificial errors to the BERT through VHDL coding and comparing them with the detected error bits. Our test results showed that the correct number of errors was detected and displayed on the LEDs on the motherboard. To verify the integrity of the electrical data transmission, we ran the data links between the RocketIO transceivers in the loop-back mode continuously for several hours, and no errors were detected. To show reliable, extended operation, we turned on four optical channels simultaneously with the I_{on} and I_b of the VCSELs set to 4.8 and 1.2 mA, respectively. The system was run continuously for 8 h, and no errors were observed.

To test how fast our optical transceiver could operate, we used a pattern generator to send a PRBS to the transmitter and monitored the output of the receiver with an oscilloscope. Figure 18 shows an eye diagram of the receiver output at 2 Gbits/s. Although the bandwidth of our transceiver design is >2 Gbits/s, all experiments were conducted at a rate of 1.8 Gbits/s, which was limited by the clock oscillator available at the time of experiment.

B. Power-Negotiation Algorithm Test Results

Both the power-on optimization algorithm and the real-time correction algorithm were implemented. A PRBS pattern of $2^{20} - 1$ was chosen for the BER testing, and a BER of 3×10^{-9} was selected as the target BER. With a system clock of 90 MHz, the optical links were running at a data rate of 1.8 Gbits/s, which corresponds to a $B\tau$ of 1.8.

During the testing of the power-on optimization algorithm, an optical filter with an optical density of 0.2 (equal to a 63% transmission rate) was inserted




Fig. 19. Ideal BER contour and experimental power settings showing the power setting of each iteration the power-on algorithm steps through during testing for a 1.8 Gbit/s optical link. For a target BER of 3×10^{-9} , the algorithm converges to the threshold-biasing point ($I_b = 0.75$ mA) at the BER valley floor.

between the two compound lenses to attenuate the optical power incident upon the detector. To ensure accurate BER measurement, we ran each iteration long enough to allow more than 256 errors to be received. The current step size of the transmitter was 0.25 mA, and the minimum gain setting was selected on the receiver side. The results of each iteration of the algorithm were tracked and monitored by use of displays on the motherboard. The result of this test was that the algorithm converged to a power setting with an I_{on} of 2 mA and an I_b of 0.75 mA at a data rate of 1.8 Gbits/s. Figure 19 shows the BER and the power setting of each iteration that the algorithm went through.

To show that this derived power setting is the desired optimum power setting for the VCSEL, we used different optimum power settings near the derived setting to run the link and we measured the BER for each setting. The results showed that for any other power setting that led to equal or less power consumption of the optical link, a higher BER than the target BER was observed. Thus it is verified that our power-on algorithm is able to find the optimum power setting that results in the minimum power consumption for a given target BER. The algorithm converged to the same setting each time after being rerun several times, which illustrates the stability of the algorithm's implementation with a fixed operating environment. The algorithm was also run on other channels in multiple transceiver chips, and all the results showed its capability of locating the optimum power setting. The low modulation current needed can be explained by the low-loss (short) optical path and the high sensitivity of the receiver.

To verify that the derived optimum power setting corresponds to the BER valley floor where the threshold current is located, a verification of the threshold current of the VCSEL was carried out. An optical powermeter was placed between the two lenses to




Fig. 20. VCSEL optical power output versus current; 0.75 mA of the threshold current is extrapolated.

measure the optical output of the VCSEL under test with other channels turned off. The digital power setting was stepped through, and the optical output power was recorded for each step. Figure 20 shows a plot of the optical output power versus the current driven into the VCSEL. We extrapolated a threshold current of 0.75 mA by fitting a straight line over our data and extending the line to the ordinate, which shows the agreement of the derived optimum I_b and the actual threshold current of the device.

To demonstrate the ability of the real-time correction algorithm to find the optimum power setting when the system's operating environment was changed, we used an optical filter with an optical density of 0.5 (equal to a 32% transmission rate) instead. The real-time correction algorithm was executed, and the same I_b of 0.75 mA was found, together with an I_{on} of 2.75 mA for this setup, which illustrated the ability of the algorithm to locate the optimum power setting under dynamic conditions. The power-on algorithm was also run with this setup, and the same power setting was found, which shows agreement between the power-on and real-time correction algorithms.

To further illustrate that threshold biasing yields a better link performance than below-threshold biasing, we show a comparison of the eye diagrams of the




Fig. 21. Eye diagrams of the optical receiver output at 1.8 Gbits/s with the VCSEL working at a power setting of $I_{on} = 2.75$ mA and with (a) $I_b = 0.75$ mA (threshold biasing) and (b) $I_b = 0.25$ mA (below-threshold biasing). Significant jitter can be observed at the rising edge in (b), which confirms that threshold biasing minimizes data-dependent jitter and results in a better performance for the optical link.

receiver output in Fig. 21. A significant pattern-dependent jitter resulting from the turn-on delay can be observed at the rising edge in Fig. 21(b) when the VCSEL is biased at below-threshold current. This confirms that threshold biasing minimizes data-dependent jitter and results in better performance for the optical link.

The experimental results given above verified the performance of the power-negotiation algorithm for a multigigabit-per-second optical link ($B\tau > 1.2$). To determine how well the algorithm will work in a system at a lower data rate ($B\tau < 1.2$) we carried out separate experiments on a 100 Mbit/s link with a selected target BER value of 10^{-7} . The algorithm resulted in 4.5 mW power consumption for the transmitter, while zero biasing produced a minimum of 4.0 mW to achieve the same BER. Thus, even for a low data rate ($B\tau < 1.2$), the algorithm can produce a nearly optimal power setting that is close to the minimum setting achieved by zero biasing.

C. Discussion

Although our experiments were carried out with dc coupled receivers and a free-space optical interconnect system, the proposed power-negotiation algorithm can be extended to ac coupled receivers and fiber-optic communication systems. In our experimental setup there were no measurable levels of optical cross talk or ambient light caused by scattering. Therefore in this system a Gaussian approximation is applicable and agrees with the noise variations measured for data 1 and data 0 transmission. In other systems, which suffer from non-Gaussian noise sources such as cross talk, high aggregate levels of scattered light, and modal noise in fiber links, extending this technique may require additional theoretical analysis with non-Gaussian noise models, and the performance of the power-negotiation algorithm needs to be reexamined.

6. Conclusions

An adaptive power-optimization technique for optical links based on channel bit-error-rate performance has been proposed and demonstrated. Based on an analysis of optical links dominated by additive Gaussian noise, a power-negotiation algorithm has been developed with which to determine the optimum power setting for a given target BER. Details of the algorithm and its implementation have been developed. An optical transceiver–demonstration system based on VCSELs and photodetectors has been presented. The power-optimization algorithm was implemented, and the test results proved the concept. By using these techniques it is possible to build simpler, more power-efficient optical links that have the flexibility to account for dynamic changes.

This research was supported and funded by the Defense Applied Research Projects Agency under contract 1120-24596, the OE center and the Consortium for Optical and Optoelectronic Technologies in Computing of George Mason University, Fairfax, Vir-

ginia, and the COOP–Peregrine–University of Southern California workshop and foundry run.

References

1. F. Kiamilev, P. Marchand, A. V. Krishnamoorthy, S. C. Esener, and S. H. Lee, "Performance comparison between optoelectronic and VLSI multistage interconnection networks," *J. Lightwave Technol.* **9**, 1674–1692 (1991).
2. M. W. Haney, M. P. Christensen, P. Milojkovic, J. Ekman, P. Chandramani, R. Rozier, F. Kiamilev, Y. Liu, and M. H. Brenner, "Multichip free-space optical interconnection demonstration based on integrated arrays of vertical cavity surface emitting lasers and photodetectors," *Appl. Opt.* **38**, 6190–6200 (1999).
3. A. Shang and F. Tooley, "Digital optical interconnects for networks and computing systems," *J. Lightwave Technol.* **18**, 2086–2094 (2000).
4. D. V. Plant, M. B. Venditti, E. Laprise, J. Faucher, K. Razavi, M. Chateaufneuf, A. G. Kirk, and J. S. Ahearn, "256-channel bidirectional optical interconnect using VCSELs and photodiodes on CMOS," *J. Lightwave Technol.* **19**, 1093–1103 (2001).
5. M. W. Haney, M. P. Christensen, P. Milojkovic, G. J. Fokken, M. Vickberg, B. K. Gilbert, J. Rieve, J. Ekman, P. Chandramani, and F. Kiamilev, "Description and evaluation of the FAST-Net smart pixel-based optical interconnection prototype," *Proc. IEEE* **88**, 819–828 (2000).
6. D. M. Cutrer and K. Y. Lau, "Ultralow power optical interconnect with zero-biased, ultralow threshold laser—how low a threshold is low enough?" *IEEE Photon. Technol. Lett.* **7**, 4–6 (1995).
7. L. Zei, K. Petermann, R. Jager, and K. J. Ebeling, "Operation range of VCSEL-interconnect links with 'below-threshold' biasing," *J. Lightwave Technol.* **18**, 477–481 (2000).
8. O. Kibar, D. A. Van Blerkom, C. Fan, and S. C. Esener, "Power minimization and technology comparisons for digital free-space optoelectronic interconnections," *J. Lightwave Technol.* **17**, 546–645 (1999).
9. J. M. Steininger and E. J. Swanson, "A 500 Mb/s CMOS optical data link receiver integrated circuit," in *IEEE International Solid-State Circuits Conference Digest of Technical Papers* (Institute of Electrical and Electronics Engineers, 1986), pp. 60–61, 304.
10. R. G. Swartz, B. A. Wooley, A. M. Voshchenkov, V. D. Archer, and G. M. Chin, "A monolithic laser driver for optical communications," in *IEEE International Solid-State Circuits Conference Digest of Technical Papers* (Institute of Electrical and Electronics Engineers, 1982), pp. 162–163.
11. G. Held, *Data over Wireless Networks: Bluetooth, WAP, and Wireless LANS* (McGraw-Hill, 2001), Chap. 6.
12. P. Palacharla, J. Chrostowski, R. Neumann, and R. J. Gallenberger, "Techniques for accelerated measurement of low bit error rates in computer data links," in *Proceedings of IEEE 14th Annual International Phoenix Conference on Computers and Communications* (Institute of Electrical and Electronics Engineers, 1995), pp. 184–190.
13. X. Wang, F. Kiamilev, P. Gui, J. Ekman, G. C. Papen, M. J. McFadden, M. W. Haney, and C. Kuznia, "A 2-Gb/s optical transceiver with accelerated bit-error-ratio test capability," *J. Lightwave Technol.* **22**, 546–645 (2004).
14. V. Tarokh, H. Jafarkhani, and A. R. Calderbank, "Space-time block coding for wireless communications: performance results," *IEEE J. Sel. Areas Commun.* **17**, 451–460 (1999).

15. G. P. Agrawal, *Fiber-Optic Communication Systems* (Wiley, 2002), pp. 155–164.
16. M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables* (Dover, 1972), p. 299.
17. Honeywell International, “Modulating VCSELs,” Honeywell application notes, <http://content.honeywell.com/vcSEL/publications/app.stm>.
18. K. Obermann, S. Kindt, and K. Petermann, “Turn-on jitter in zero-biased single-mode semiconductor lasers,” *IEEE Photon. Technol. Lett.* **8**, 31–33 (1996).
19. M. Bruensteiner and G. C. Papen, “Extraction of VCSEL rate-equation parameters for low-bias system simulation,” *IEEE J. Sel. Top. Quantum Electron.* **5**, 487–494 (1999).
20. L. P. Chen and K. Y. Lau, “Regime where zero-bias is the low-power solution for digitally modulated laser diodes,” *IEEE Photon. Technol. Lett.* **8**, 185–187 (1996).
21. N. K. Dutta, “Power penalty due to timing jitter for laser modulated without prebias,” *Appl. Phys. Lett.* **67**, 3230–3232 (1995).
22. Peregrine Semiconductor Corporation, “UTSi CMOS SOS for mixed signal ICs,” Peregrine Semiconductor Corp. Tech. Doc., http://www.peregrine-semi.com/prd_proinfo.html.
23. A. Hastings, *The Art of Analog Layout* (Prentice-Hall, 2001).
24. Broadcom Corporation, “8-Channel multirate 1.0–3.2-GBPS transceiver,” Product brief, BCM8020, Broadcom Corporation, Irvine, Calif.
25. Emcore Corporation, “PIN array die 1×4 3.125 Gb, S8485-1405” and “Oxide Array Die, 1×4 3.125 Gb (U-lot), 8685-1402,” Emcore data sheets (Emcore Corporation, Somerset, N.J.).