

Efficient Optical Communications Using Multi-Bit Differential Signaling

Donald M. Chiarulli¹, Steven P. Levitan², Samuel J. Dickerson², Jason D. Bakos³, Joel Martin²

¹Department of Computer Science, Room 6135 Sennott Square Bldg, University of Pittsburgh, Pittsburgh PA 15260

²Department of Electrical and Computer Engineering, Room 348, Benedum Hall, University of Pittsburgh, Pittsburgh, PA 15261

³Department of Computer Science and Engineering, 301 Main St., University of South Carolina, Columbia, SC, 29208.

ABSTRACT

We present an alternative signaling method for multi-channel fiber ribbon based optical links. The method is based on a hybrid of differential signaling and single-ended channels. Channels are grouped into code blocks of n -bits. Each code word transmitted in the block is restricted to conform to an n choose m rule. Electrical drivers steer current between m active VCSELS with no dummy loads. A virtual reference is synthesized from the received signals and used for differential discrimination. This signaling method approaches the signal-to-noise characteristics of fully differential signaling but can be implemented with significantly lower channel overhead, giving as much as a 33% reduction in fiber count and a 44% reduction in power. Further, code utilization rates on these links can be as low as 51%, leaving substantial code space available for ECC or channel management functions. In this paper, we describe the signaling method and present a prototype transceiver chip. The transceiver is implemented in 0.25 μ m UTSi Silicon-on-Sapphire technology with flip-chip bonded VCSEL and photodetector arrays. The design demonstrates a pin-compatible alternative to the POP4-MSA transceiver standard with 125% greater data throughput and 25% better power efficiency.

Keywords: Fiber optics links and subsystems; Optical communications; Balanced encoding

INTRODUCTION

Electrical links have long embraced the use of differential signaling for low noise and high bandwidth transmission. Most optical transceivers use differential circuitry in their driver and receiver circuits for the same reason. However, constraints on fiber cost, channel count, and packaging have limited virtually all optical channels to single ended signaling through the fibers themselves. The costs of these constraints are typically balanced against inefficient power dissipation in the driver circuits, and lower signal-to-noise tolerance in the receivers.

In this paper, we present a new encoding method for short-range, parallel optical communication links that provides most of the performance benefits of differential signaling but is far less constrained by channel count and packaging issues. In this method, we replace the 2-channel-per-bit format of differential signaling with an n -channel system in which all n -bit code words must conform to an n choose m (written nC_m) rule for n even and $m=n/2$. The result is a spatially balanced code, with an equal number of zero and one bits in each code word, across a multi-channel fiber ribbon. We refer to this signaling method as multi-bit differential signaling (MBDS). Traditional differential signaling is actually a degenerate case of a 2C1 MBDS code. However, by generalizing to wider channels, it is possible to design links with considerably higher code density, better power efficiency, and lower channel count than equivalent links built from pair-wise differential channels.

These characteristics make MBDS an attractive approach for improving the performance of parallel optical (single ended) channels with a minimal commitment of additional fibers. In this paper, we present an example design for a 4-bit transceiver using 6C3 MBDS encoding. This configuration is particularly convenient as an alternative configuration for the POP4 MSA standard package. In this case, the same 12 fiber ribbon used in current POP4 package can be configured with two 6C3 MBDS links with 125% greater data throughput using 75% of the electrical and optical power

of conventional 4-bit POP4 transceivers. In addition, the receivers in this link have the additional benefit of common mode noise rejection inherent in the differential encoding style of MBDS.

We begin with a brief discussion of the key electrical and optical characteristics of typical single ended VCSEL drivers and receivers. This is followed by a description of a prototype transceiver chip implemented in .25um UTcmos, silicon on sapphire technology, and simulation data for a prototype link.

1.1 Background

Figure 1 shows an electrical model for a conventional single ended VCSEL driver circuit. The driver circuit is designed to operate with a constant supply current to minimize switching noise and crosstalk through the supply rail. A transition from one code state to the other does not require the turning on or off supply current. Instead modulation current is switched between the VCSEL and an equivalent dummy load on the opposite side of the differential driver. The weakness of this design is in the conversion from differential to single ended signaling as power is wasted in the dummy load.

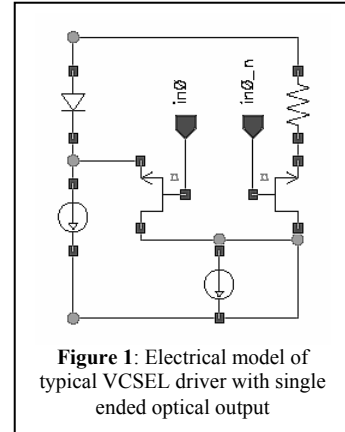


Figure 1: Electrical model of typical VCSEL driver with single ended optical output

Figure 2 shows an electrical model for a conventional receiver [1]. In this case, the conversion from a single ended optical signal to a differential electronic signal happens at the output of the transimpedance amplifier (TIA) where a fixed voltage is applied to the second input of the differential amplifier. This voltage is critical to the proper operation of the receiver. However, since it is a fixed reference, it provides none of the common mode noise rejection of a differential channel. In an MBDS receiver we derive this reference voltage directly from the ensemble of inputs over the balanced code. Thus, the reference signal carries both the common mode voltage and any common mode noise.

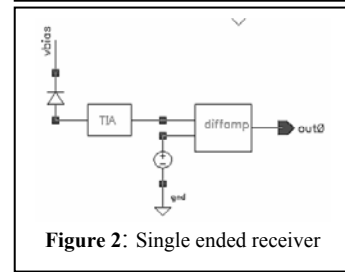


Figure 2: Single ended receiver

In a parallel optical link an array of these driver and receiver circuits is interfaced to VCSEL and detector arrays and packaged with a fiber ribbon connector. In a typical POP4 transceiver example, four of the twelve fibers in the ribbon are input, four are output, and four are dark. In our alternative version, we implement two 6C3 MBDS links using all 12 fibers to support a bidirectional pair of links.

1.2 nCm encoding

In a nCm MBDS link even though there are n-fibers and 2ⁿ available code words, the key assumption is that only a subset of these code words will ever be transmitted. [2, 3] Further, these code words have the property that each has exactly m 1-bits and exactly (n-m) 0-bits. Consider the set X such that X_{nm} = {x_{nm} : x ∈ nCm}. The size of X_{nm}, the number of code words in an arbitrary nCm code set, is

$$\phi\{X_{nm}\} = \frac{n!}{(n-m)!m!}$$

For example a 4C2 code has 6 elements X₄₂={0011,0101,0110,1001,1010,1100}. For a given n-bit code, the number of nCm code words is maximal when m=n/2. For balanced code we restrict ourselves to the maximal case and even values of n.

| Channel encoding | MBDS Code words | Effective bit width | % channel overhead | Excess codes | % Code Utilization | % Power Utilization |
|------------------|-----------------|---------------------|--------------------|--------------|--------------------|---------------------|
| 2C1 | 2 | 1 | 100% | 0 | 100% | 100% |
| 4C2 | 6 | 2 | 100% | 2 | 66% | 100% |
| 6C3 | 20 | 4 | 50% | 4 | 80% | 75% |
| 8C4 | 70 | 6 | 33% | 6 | 91% | 66% |
| 10C5 | 252 | 7 | 42% | 124 | 51% | 72% |
| 12C6 | 924 | 9 | 33% | 412 | 55% | 66% |

Table 1: Comparison of MBDS channels to equivalent parallel optical links based on single ended channels

Table 1 compares the code capacity and channel requirements of various MBDS channels to equivalent single ended parallel optical links. The channel encoding is listed in column one on the left and the corresponding number of MBDS code words is listed in column two. When used to transport a binary data stream, each of these

code words must be mapped to one of the 2^k binary data values in the k-bit incoming and outgoing data. Thus, the *effective bit width* listed in column three defines the value of k for a particular link type. It is the largest power of two less than the size of the MBDS code set. This also corresponds to the number of fibers in an equivalent single ended link. The *channel overhead percentage*, listed in column four, is based on the ratio of the fiber count for the equivalent single ended link to fiber count in an MBDS link of equal effective bit width. For example, the fully differential (2C1) channel in row one requires 100% channel overhead since it doubles the number of channels. On the other hand an 8C4 channel, using 8 fibers to transmit 6 bits, corresponds to a 33% channel overhead.

The number code words left over after mapping all of the k-bit data words into the MBDS code set are listed as *excess codes* in column five. These code words are available for error checking, protocol support, or other link management functions.

The number of excess codes available varies significantly between different channel encodings. However, the effectiveness of these codes for ECC can be enhanced by encoding the ECC over multiple code words in temporal or spatial sequences. In other research [4], we have used this technique to implement multi-bit ECC over MBDS encodings with no additional channel requirements other than the MBDS channel overhead in a 12-bit link composed of three 4C2 code words.

PROTOTYPE TRANSCIEVER IMPLEMENTATION

2.1. Design Overview

Figure 3 is an annotated image of MBDS optical transceiver test chip. The chip is implemented in .25um UTSi Silicon-on-Sapphire CMOS technology. [5] The UTSi substrate is both fully insulating and optically transparent making it an excellent host for optoelectronic applications. [6] In this figure the transceiver is shown from the bottom surface with the inset in the lower right showing a top view. Three 1x4 VCSEL and three corresponding 1x4 detector arrays are flip-chip bonded to the top surface. When viewed from the bottom each of the VCSEL and detector sites are clearly visible through input/output windows designed into the driver and receiver circuitry. This allows for input and output of the transceivers through the bottom surface where these signals are coupled to fibers.

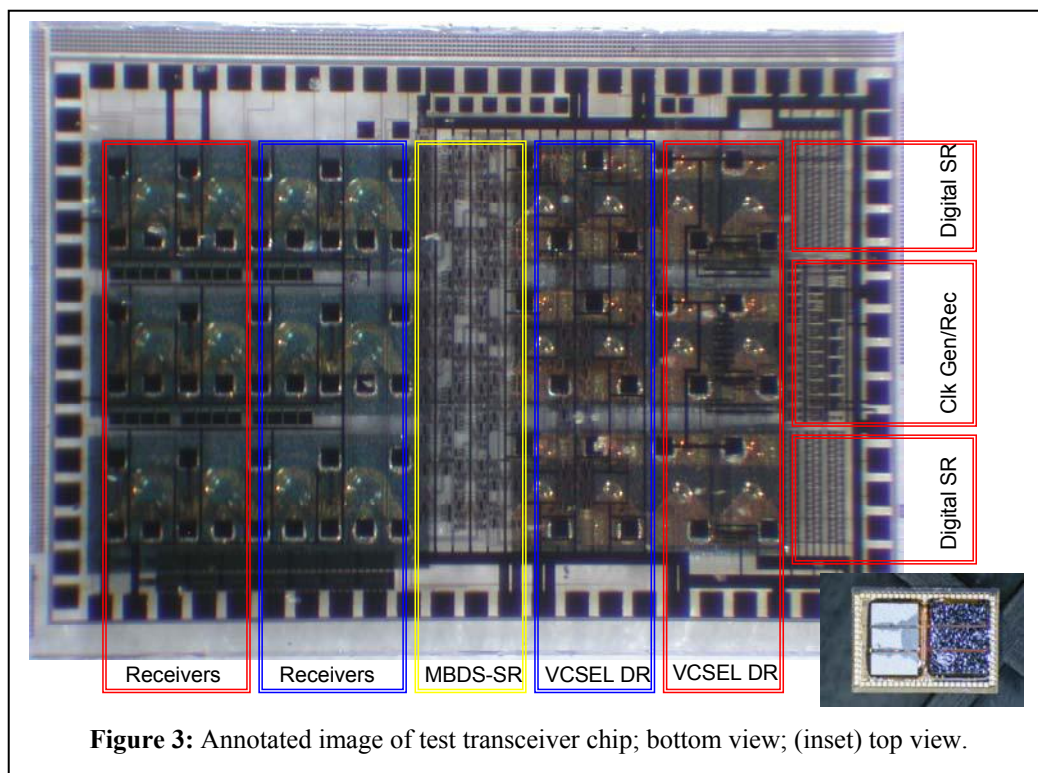
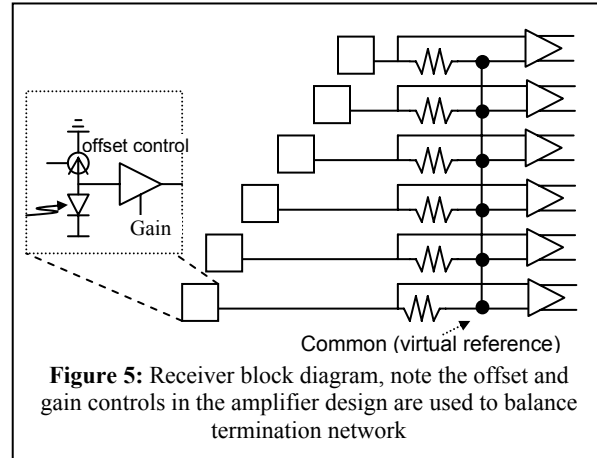
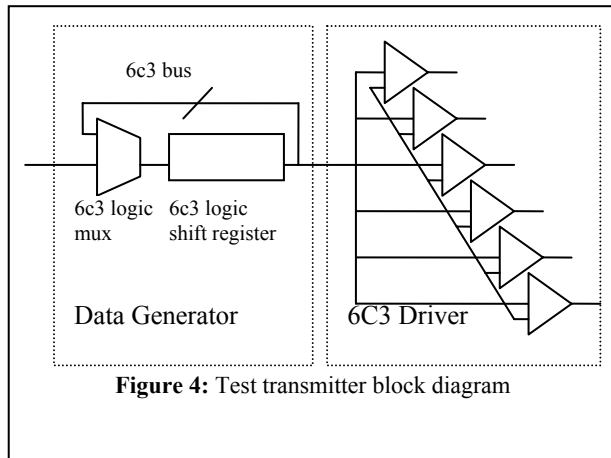


Figure 3: Annotated image of test transceiver chip; bottom view; (inset) top view.

The test transceiver chip implements two versions of the driver circuit, shown outlined in red and blue on the right side of figure 3, and two versions of the receiver circuit, shown outlined in red and blue on the left of the figure. Both versions are based on the same basic design shown in the block diagrams of figure 4 and figure 5, respectively. In the driver case, this design consists of a data generator based on a programmable, 24-bit, circular shift register with output connected to a 6C3 driver. The difference between the two versions is in the implementation of the data generator. The first version is based on conventional current mode logic, labeled “Digital SR” in figure 3. The second uses a unique MBDS implementation of current mode logic to implement the shift register and multiplexer in the data generator. This logic is outlined in yellow and labeled MBDS-SR in figure 3.



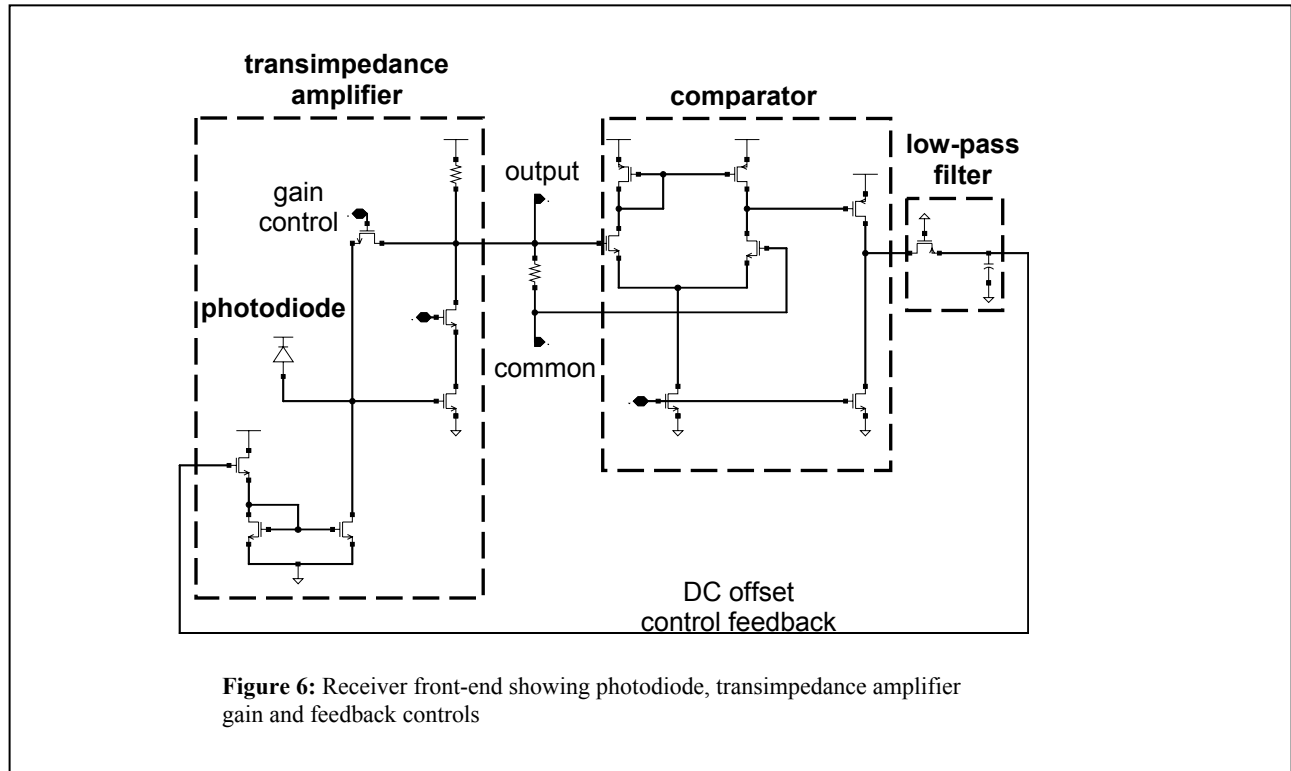
The basic MBDS receiver design is shown in figure 5. In this receiver, each channel is detected independently and the photodetector current is converted to voltage in a separate transimpedance amplifier. Each of these signals is applied to the “star” resistor network shown in figure 5 such that the center of the star creates a common mode reference voltage. Since at any time, MBDS coding requires that there be exactly the same number of zero and one bits, the common mode reference voltage is centered at the exact average of the zero and one state, and carries common mode noise for all signals. Differential limiting amplifiers use the difference between this common mode signal and the individual transimpedance amplifiers to generate the final receiver output.

Two versions of the receiver have been implemented on the test chip. The versions use alternative approaches to the handling of a key issue for the proper operation of the receiver, maintenance of balanced signals across the star resistor network that generates the common mode reference. Channel-to-channel variations in the signal offset and amplitude can be introduced by several mechanisms ranging from thermally induced differences in individual VCSEL performance to differences in insertion loss caused by mechanical tolerances. These variations are compensated by adjustments to offset and gain controls in the transimpedance amplifiers. In the first version of the receiver, these controls are driven externally. In the second, we have implemented automatic compensation circuitry. This circuitry is described in the next section.

2.2. Receiver Design

Multi-bit differential signaling requires that the DC offset voltage and signal amplitudes at the input to the termination network be balanced across all of the input bits. When this is the case, the AC portion of each channel’s signal is centered on the common reference voltage. To prevent an imbalance from introducing noise into the common mode signal, automatic DC offset controls are needed for this receiver architecture.

Figure 6 is a schematic of the first stage of the receiver circuit which consists of a transimpedance stage, shown on the left, with the DC offset control for the photodetector input and a comparator stage, shown to the right. The transimpedance amplifier is a resistively loaded cascode amplifier with variable feedback impedance. The feedback is implemented as transistor with a controllable gate voltage. Because the transistor’s impedance varies with the gate voltage, this input can be used as a gain control. Since the TIA loads and layouts are well matched across all channels,

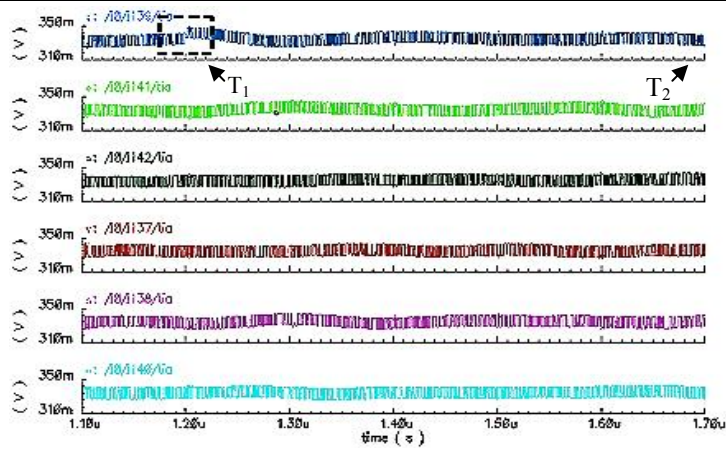


any differences in the DC offset observed at the TIA output will primarily be due to variations in the incoming signal. To compensate for these differences, a current mirror connected to the TIA input node is used to control the photodiode's DC current.

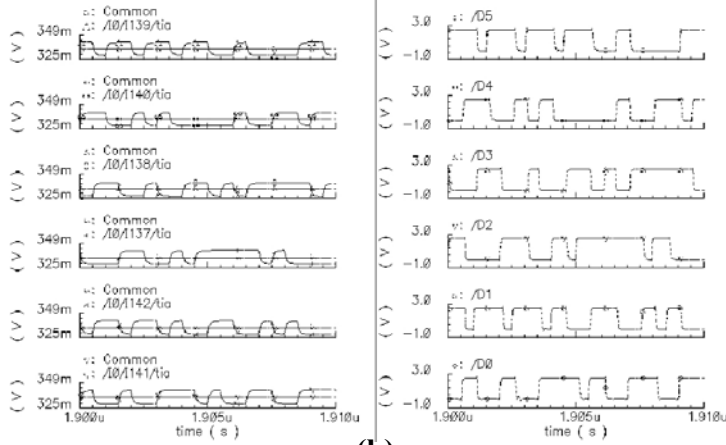
The offset is automatically controlled via a negative feedback loop. Each channel has a high-DC gain comparator with the common reference and TIA output as its two inputs. Source followers are used for buffering so that the comparator does not load down the TIA stage (buffers have been omitted from figure 6 for clarity). The comparator output is low-pass filtered, resulting in a signal that is proportional to the difference between common and the average value of the TIA input and is then fed back to the photodiode current mirror as the controlling voltage for the DC offset control input.

The final stage of the receiver is the limiting amplifier, not shown in figure 6, which amplifies the difference between the small-signal TIA output and common to a CMOS compatible signal level. The limiting amplifier is composed of a cascade of nine differential gain stages, in an open-loop configuration. Each stage is designed to have a low gain and high bandwidth so that an overall high-gain and large bandwidth can be achieved. In addition to gain, the differential stages provide common-mode noise rejection.

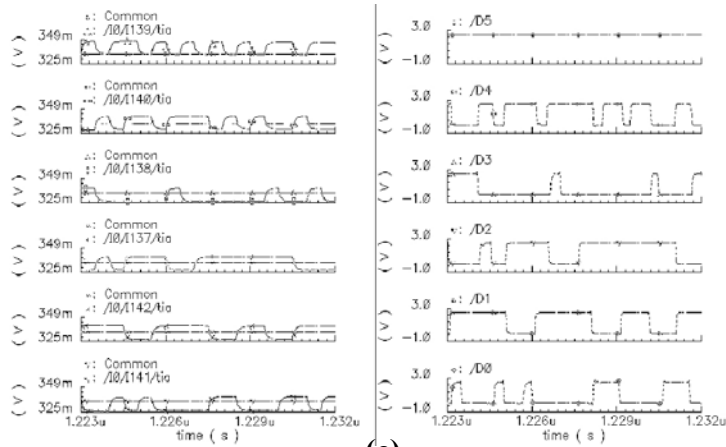
Figure 7 shows a simulation that demonstrates the automatic DC compensation circuit. Figure 7(a) shows the output of all 6 channels operating in parallel. At time $t = 1.22 \mu\text{s}$, an additional DC offset is introduced into the photodetector current of channel 5, as indicated by the boxed region of the waveform. This results in a perturbation of the DC output voltage of the TIA at time T_1 . Figure 7(b) shows how that offset makes the common reference (shown as the solid line superimposed over the TIA output traces on the left hand side of the figure) unbalanced among the 6 channels. This results in bit errors on channels 3 and 5, as shown in the traces of the right hand side of figure 7(b) for the output of the limiting amplifier. Figure 7(c) shows the same data stream several hundred nanoseconds later. After time T_2 , the feedback circuitry has compensated for the new DC offset and the common mode voltage is restored. As shown in the figure, the data output at the limiting amplifier is no longer in error.



(a)



(b)



(c)

Figure 7: (a) DC Offset introduced on channel 5 (b) Bit errors on channels 3 and 5 due to offset (c) Channels no longer in error once feedback reaches steady state

LINK SIMULATION RESULTS

To demonstrate 40Gbs operation of the transceiver, we have performed simulations based on the link model shown in figure 8. The simulation was performed using the Cadence *Spectre* simulator. All circuit models include full extracted parasitic components from the chip layout, generated using the *Diva* circuit extractor. The link was simulated with a bit (code) period of 106.5ps. Since there are 20 available code words in a 6C3 MBDS code set, the effective throughput of the link is 40Gbs.

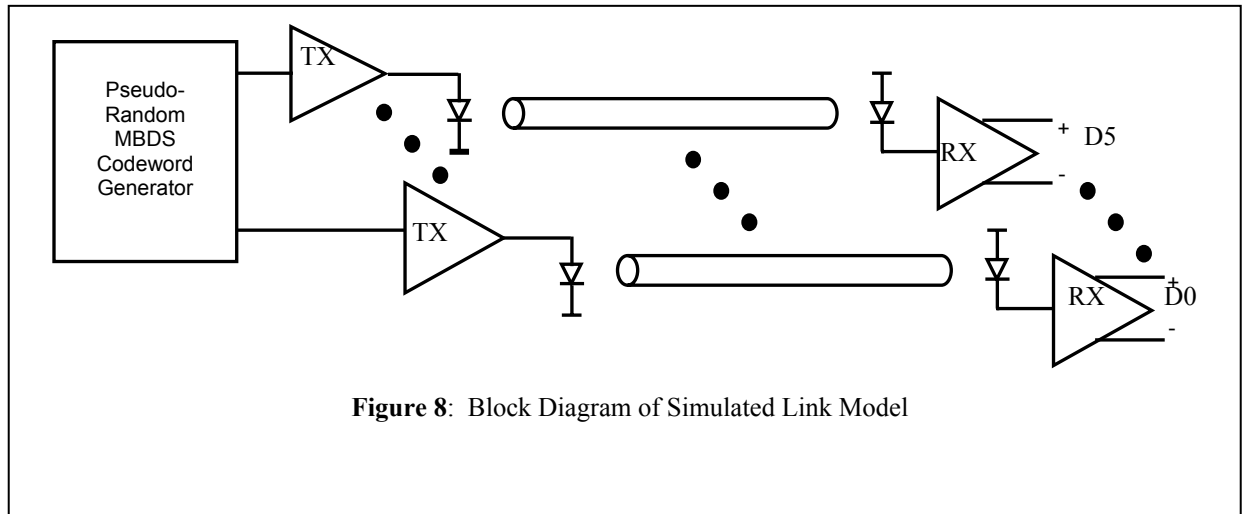


Figure 8: Block Diagram of Simulated Link Model

The simulated link uses the MBDS-logic data generator, this generator was pre-loaded with a pseudo random sequence of MBDS codes. Output from the code generators stimulates the driver circuitry. The VCSEL model at the driver output is an RC network of 50 ohms and 400fF in order to generate an electrical load corresponding to the commercial VCSEL arrays used in the prototype. Simulation models for the conversion efficiency of the VCSELs, fiber loss, and detector responsivity were lumped into a single attenuator. The attenuated current then drives a capacitive load of 300fF to model the photodetector capacitance and is used as input to each of the TIA circuit models.

Eye diagrams for each channel are shown in figure 9. Figure 9(a), on the left, shows modulation current output from the driver as input to the VCSEL model. Figure 9(b), on the right, is the differential output from each of the TIAs. The peak-to-peak differential voltage is 15mV at this bandwidth. Both eyes are well open and demonstrate that the link will operate with acceptable bit error rates at this bandwidth.

SUMMARY

We have demonstrated an alternative signaling method for multi-channel fiber ribbon based optical links using multi-bit differential signaling (MBDS). We have outlined the coding system and shown driver and receiver circuits capable of operation with an effective throughput of 40Gbs through a POP4-MSA compatible transceiver. This transceiver operates using 75% of the optical and electrical power required in a conventional design.

REFERENCES

1. A. Apsel, Z. Fu, A. Andreou, "A 2.5 mW SOS CMOS Optical Receiver for Chip-to-Chip Interconnect," *IEEE Journal of Lightwave Technology*, Volume 22, pp 2149-2157, September 2004.
2. Donald M. Chiarulli, Steven P. Levitan, Jason Bakos, "Area, power, and pin efficient bus structures using multi-bit-differential signaling," *VLSI Circuits and Systems II, SPIE Symposium on Microtechnologies for the New Millennium 2005*, 5837-4, Sevilla, Spain, May 9-11, 2005.
3. Donald M. Chiarulli, Jason D. Bakos, Joel Martin, Steven P. Levitan, "Area, power, and pin efficient bus transceiver using multi-bit-differential signaling," *High Quality Life and Consumer Technology, ISCAS 2005*, (poster), Kobe, Japan, May 23-26, 2005.

4. Jason D. Bakos, Donald M. Chiarulli, Steven P. Levitan, "Lightweight Error Correction Codes for System-Level Interconnects," *IEEE Transactions on Communications*, (submitted November 2005).
5. <http://www.peregrine-semi.com/>
6. D. Chiarulli, S. Levitan, J. Bakos, C. Kuznia, "Active Substrates For Optoelectronic Interconnect," 3385 Session: INV-13: Heterogeneous Systems, *IEEE International Symposium on Circuits and Systems (ISCAS 2004)*, Volume 5, pp. V-592 – V-595, Vancouver, Canada, May 23-26, 2004.

