INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps.

ProQuest Information and Learning
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA
800-521-0600

A 1Gb/s, 9x1 optical receiver array, for an adaptive redundant free-space interconnect system

Madeleine Mony

Department of Electrical and Computer Engineering
McGill University
Montreal, Canada
July, 2000

A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements of the degree of Master of Engineering

Abstract

Free-space optical links offer flexible, high-density, and high-bit rate interconnections between circuit boards. One of the main challenges in realizing free-space interconnects is achieving misalignment tolerance. This thesis presents the design of a receiver used in a system that makes use of spatial redundancy to increase misalignment tolerance. The system consists of four 1Gb/s free-space optical links between two printed circuit boards. It has ±1mm of lateral misalignment tolerance and ±1 degree of tilt misalignment tolerance with the spacing between the boards varying from 5cm to 21.5cm. This makes the system ideal for a rack-based interconnection.

The receiver for this system is a rapidly re-configurable array that accepts nine low-amplitude, high-speed photocurrents, selects one of them, and then outputs that signal as a digital differential positive emitter coupled logic (PECL) signal. The selection of which channel to amplify is based on received power, and is performed off-chip. The design, packaging, and testing of the receiver are presented.

Sommaire

Les interconnections optiques en espace libre offrent des canaux de communication flexibles, à haute-vitesse ainsi qu'une haute densité de canaux. L'obtention d'une bonne tolérance aux mésalignements constitue un des défis fondamentaux dans l'atteinte de ce type d'interconnections. Cette thèse présente la conception d'un circuit récepteur, utilisé au sein d'un système qui exploite une redondance spatiale pour procurer au système en question un plus haut niveau de tolérance au mésalignments. Le système est composé de quatre liens optiques en espace libre, chacun fonctionnant à un gigabit par seconde (1Gb/s), entre deux circuits imprimés. Le système fonctionne avec un mésalignment latéral de ±1mm et angulaire de ±1°. La distance entre les deux circuits imprimes peux varier de 5 à 21.5 cm. Cette distance rend le système idéal pour l'interconnection de cartes à l'intérieur d'un panier.

Le circuit constituant le récepteur pour ce système est une matrice rapidement re-configurable, qui accepte neuf signaux à haute vitesse et basse amplitude, en sélectionne un, puis livre ce signal sous forme numérique différentielle, émetteur positif à logique couplée (PECL). La sélection du canal à amplifier se base sur la puissance reçue au niveau des photodetecteurs et se fait à l'extérieur du circuit récepteur. La conception du circuit intégré, son utilisation sur un circuit imprimé, ainsi que sa caractérisation sont présentés.

Acknowledgements

I would like to thank my supervisor, Prof. David Plant, for all of his support, guidance, and encouragement throughout my degree, and for teaching me what it means to be an experimentalist.

The original idea for the system came from Dominic Goodwill at Nortel Networks, and I would like to thank him for all of the guidance he provided on the project. Petre Popescu from Nortel Networks also offered invaluable advice on many aspects of the receiver design, and I am very grateful to him.

This thesis covers only a part of a system developed at McGill University in the Photonic Systems Group. I would like to acknowledge Prof. Andrew Kirk, Eric Bernier, Eric Bisaillon, Daniel F. Brosseau, Mitch Salzberg, and Tsuyoshi Yamamoto for all of their very hard work in getting this system working. A special thanks to Eric Bisaillon for all of his help and support at the conferences where we presented this work, and all throughout the project.

I would also like to thank all of the many other students in the Photonic Systems Group, past and present, who all helped by providing a wonderful atmosphere to work in. No other research group can possibly have as many cake parties!

I would finally like to thank my parents, Georges and Mary, my sisters Thérèse, Ann, and Emilie, for their loving support through all of my studies. I would especially like to thank my father who read and commented my entire thesis and helped me analyze some of the trickier results.

This work was supported by a contract from Nortel Networks. The silicon chip, which was designed in NT25 technology, was gracefully provided by the Canadian Microelectronics Corporation (CMC). In addition, NSERC and FCAR post-graduate fellowships are acknowledged.

TABLE OF CONTENTS

LIST OF FIGURES

Chapter 1. Introduction

1.1 Motivation

The bandwidth requirements for board-to-board interconnections are constantly increasing. Optical interconnects have been proposed as a way to increase that bandwidth, while alleviating the problems caused by capacitive loading and electromagnetic crosstalk [1-4]. Several board-to-board optical interconnects have been designed, using such technologies as fiber ribbon [7, 8], two-dimensional arrays of optical fibers [9], and flexible image guides [10]. However, in certain applications free-space interconnections are more desirable. Free-space optical interconnects offer not only high bit rate and high density connections, but also high interconnection flexibility. For example, in a rack based system, a free-space connection between boards allows for the dynamic insertion and removal of boards. Examples of such systems are shown below in Figure 1.

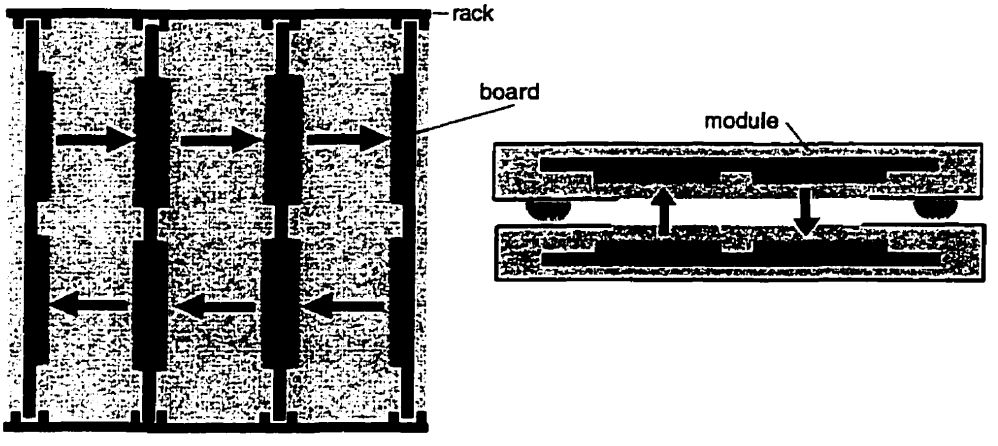

Figure 1. Board-to-board optical interconnections

One of the main challenges in realizing free-space optical interconnects is achieving misalignment tolerance. For example, in a free-space system built by E. Strzelecka, et al. [13], an array of lasers is interconnected to an array of photodetectors, and the lateral alignment tolerance between the two is ±50μm. We are exploring a novel way of relaxing these tolerances by using spatial redundancy [3, 11]. Specifically, an attractive option is to use two-dimensional arrays of lasers and detectors which provides many possible optical channels, and where only the one with the highest received power is used to transmit data. This concept was originally proposed by Tewksbury, et al [5], and then again by D. J. Goodwill, et al [11]. In our system, we interconnect a 3×3 Vertical Cavity Surface Emitting Laser (VCSEL) array to a 3×3 array of PIN photodetectors, creating a total of 81 possible optical channels between them. Before the optical data is transmitted there is an FPGA controlled hunt algorithm that finds the strongest channel, based on received power. The chosen channel is used to transmit 1Gb/s data [14, 15].

Using the combination of redundancy and a custom optical interconnect design [16], simulations indicated that we could achieve a large misalignment tolerance between boards. In fact, we are able to handle lateral misalignments of ±1mm and angular misalignments of ±1 degree, with the distance between the boards varying continuously from 5cm to 21.5cm. An application of this design can be, for example, to insert the boards into an electrical rack where they are aligned and spaced to within the required specifications defined by the maximum misalignment that can be handled by the opto-electronics and the optics.

1.2 Thesis Organization

This thesis presents the design and testing of a receiver array used in a system that exploits redundancy to increase misalignment tolerance. The receiver chip is a 9×1 receiver array that is connected to a 3x3 array of photodetectors. A multi-chip module packaging scheme is used, in which the two dies are wirebonded to each other inside a quad flat-pack chip carrier. The receiver array was designed and manufactured in NT25 technology [17], which is Nortel's high-speed silicon bipolar process. This process was made available through the Canadian Microelectronic Corporation (CMC) foundry service [18].

The receiver array is designed to receive nine input photocurrents. These currents are amplified by nine transimpedance amplifiers (TIAs). One of the signals is then selected through a pyramid of four 3-to-1 analog switches. This signal is sent off-chip for power monitoring purposes. Once all 9 signals have been scanned via a hunt algorithm, the strongest one is selected, amplified, and sent off-chip as a digital positive emitter coupled logic (PECL) output signal. The circuitry that compares the signal strengths from all the receivers is off-chip, and the control for the switch pyramid is handled by an FPGA. The analog 3-to-1 switches can be switched at 50MHz, allowing for a rapidly re-configurable receiver array.

This thesis is organized into the following sections: Chapter 2 will give an overview of the system for which the receiver was designed, and in Chapter 3 the receiver architecture will be described. Chapter 4 discusses the simulation results

of the designed architecture. Chapter 5 describes the receiver module packaging scheme and Chapter 6 describes the printed circuit board design for the system. Chapter 7 explains the test methods that were used and presents the experimental results. Chapter 8 presents the conclusions.

1.3 References

1. M. Feldman, et al. "Comparison between optical and electrical interconnects based on power and speed considerations", *Applied Optics*, Vol. 27, No. 9, pp. 1742-1751, 1988.

2. A. Krishnamoorthy, et al., "Scaling optoelectronic-VLSI circuits into the 21st century: a technology roadmap", *IEEE Journal of Selected Topics in Quantum Electronics*, Vol. 2, No. 1, pp.55-76, 1996.

3. F. Tooley, "Challenges in optically interconnecting electronics", *IEEE Journal of Selected Topics in Quantum Electronics*, Vol. 2, No. 1, pp.3-13, 1996.

4. H. S. Hinton, et al. "Free-space digital optical systems", *Proceedings of the IEEE*, Vol.82, No. 11, pp. 1632-1649, 1994.

5. Tewksbury, et al. "Toward cointegration of optical interconnections within silicon microelectronics systems", Journal of Parallel and Distributed Computing, Vol. 17, pp. 188-199, 1993.

6. Semiconductor Industry Association, *International Technology Roadmap for Semiconductors: 1999.*

7. M. Takaya, et al., "Design and performance of a modular jack type Mini-MT connector for high-density optical duplex connection", *IEEE Photonics Technology Letters*, Vol. 11, No. 11, pp. 1146-1148, 1999.

8. A. Van Hove, et al., "Direct MT(TM)-compatible connectorisation of VCSEL and LED-arrays to plastic optical fiber ribbon for low cost parallel datalinks", Electronic Components and Technology Conference Proceedings (49th), pp. 1096-1102, 1999.

9. C. V. Cryan, "Two-dimensional multimode fiber array for optical interconnects", *Electronics Letters*, Vol. 34, pp. 586-587, 1998.

10. T. Maj, et al., "Interconnection of a two-dimensional array of vertical-cavity surface-emitting lasers to a receiver array by means of a fiber image guide", Applied Optics, Vol. 39, No. 5, pp. 683-689, 2000.

11. D. J. Goodwill et al., "Free space optical interconnect at 1.25Gb/s/channel using adaptive alignment," OFC/IOOC 1999, San Diego, USA, paper WM22, 1999.

12. D. A. Louderback, et al., "Modulation and free-space link characteristics of monolithically integrated vertical-cavity lasers and photodetectors with microlenses", *IEEE Journal of Selected Topics in Quantum Electronics*, Vol. 5, No. 2, pp.157-165, 1999.

13. E. Strzelecka, et al. "Parallel free-space optical interconnect based on arrays of vertical-cavity lasers and detectors with monolithic microlenses", *Applied Optics*, Vol. 37, No. 14, pp.2811-2821, 1998.

14. D. V. Plant, et al., "A 5 Gb/s, 2 channel bi-directional adaptive redundant FSOI demonstrator system", in *Optics in Computing, 2000*, Roger A. Lessard, Tigran Galstian, Editors, SPIE Vol. 4089, pp. 465-472, 2000.

15. M. Mony, et al. "1Gb/s, 9x1 optical receiver array, for an adaptive redundant interconnect system, using NT25 technology", in *Optics in Computing, 2000*, Roger A. Lessard, Tigran Galstian, Editors, SPIE Vol. 4089, pp. 1010-1020, 2000.

16. E. Bisaillon, et al., "Optical link for an adaptive redundant free space interconnect", in *Optics in Computing, 2000*, Roger A. Lessard, Tigran Galstian, Editors, SPIE Vol. 4089, pp. 999-1009, 2000.

17. "NT25: Nortel's High-Speed Bipolar Process", White Paper, Nortel Semiconductors.

18. Canadian Microelectronic Corporation (CMC), http://www.cmc.ca

19. T. Vanisri, et al., "Integrated high frequency low-noise current-mode optical transimpedance preamplifiers: theory and practice", *IEEE Journal of Solid-State Circuits*, Vol. 30, No. 6, pp.677-685, 1995.

Chapter 2. Description of system

2.1 Introduction

In order to increase data throughput between two independent printed circuit boards, we designed a system consisting of two bi-directional optical links. As described in the previous chapter, this system is designed to interconnect boards within a rack, or for a stackable module system. Therefore the data links between the boards are free-space optical links and there is no electrical communication between the boards [1]. Such free-space systems usually require very high alignment tolerances between the printed circuit boards. In our system, using the concept of redundancy which is described further in this chapter, and using an optimized optical design, we have greatly increased the misalignment tolerance between the boards to ±1mm of lateral misalignment tolerance and ±1 degree of tilt misalignment tolerance [1,2]. The separation between the boards can vary from 5cm to 21.5cm. These mechanical tolerances allow the boards to be inserted and removed from a rack without any additional alignment requirements.

2.2 System

Our system consists of two printed circuit boards defined as board A and board B. Referring to Figure 2, there are four optical links between the boards, each transmitting data between 100Mb/s and 1Gb/s. This bandwidth for the

transceivers was chosen to be compatible with an 8B/10B encoding scheme used for Gigabit Ethernet.

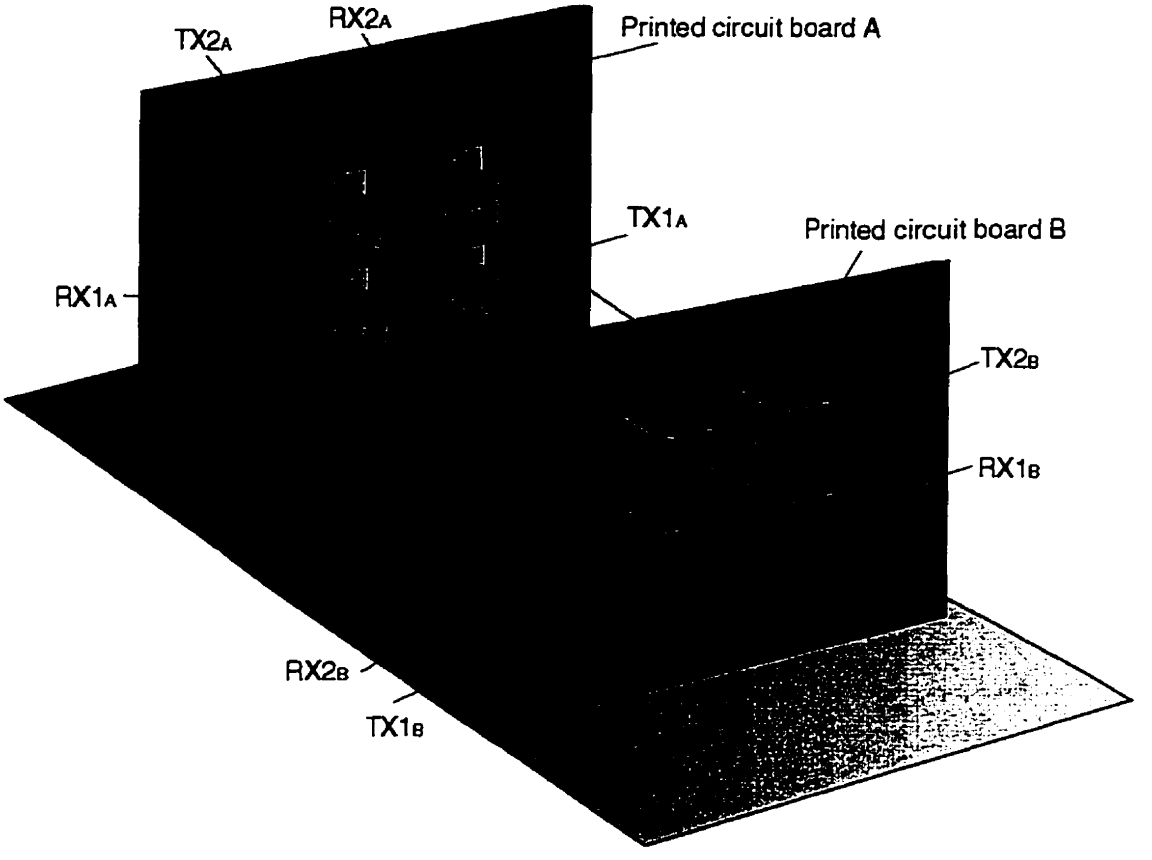

Figure 2. Two bi-directional links between two circuit boards

On each board, there are two transmitter modules ($TX1_{A,B}$ and $TX2_{A,B}$) and two receiver modules ($RX1_{A,B}$ and $RX2_{A,B}$). The first bi-directional link is made up of $TX1_A$ on board A transmitting data to $RX1_B$ on board B, and $TX1_B$ on board B sending data to $RX1_A$ on board A. The second bi-directional link is made up of $TX2_A$ on board A transmitting data to $RX2_B$ on board B, and $TX2_B$ on board B sending data to $RX2_A$ on board A. For the optical links to be established, the boards must face each other. The two bi-directional links are independent, and

having a second bi-directional link between the boards doubles the data throughput rather then providing redundancy. Between the boards there is no electrical contact except for power, so all communication or feedback between the boards is performed optically.

2.3 Redundancy

The transmitter module is made up of a transmitter chip, a 3×3 array of VCSELs, a microlens array and a bulk lens. The receiver module is made up of a receiver chip, a 3×3 array of PIN photodetectors (PDs), and a bulk lens. The operation of the system is as follows: one of the lasers in the 3×3 VCSEL array emits light which is detected by one of the photodetectors in the 3×3 detector array. Given that there are 9 lasers and 9 detectors that can be used to establish one optical link, there are 81 possible combinations (9x9). The VCSEL-PD pairing which is selected will be the one that has the lowest link loss and therefore produces the highest average optical power at the selected PD from the selected VCSEL.

The VCSEL spots are defocused at the photodetector plane to about twice the photodetector active area size to ensure smooth power coupling between the two arrays as the system gets misaligned. Also, the VCSEL array is rotated by 45 degrees with respect to the optical axis to give better power coupling [3]. Figure 3 illustrates this concept by showing the VCSEL spots at the photodetector plane.

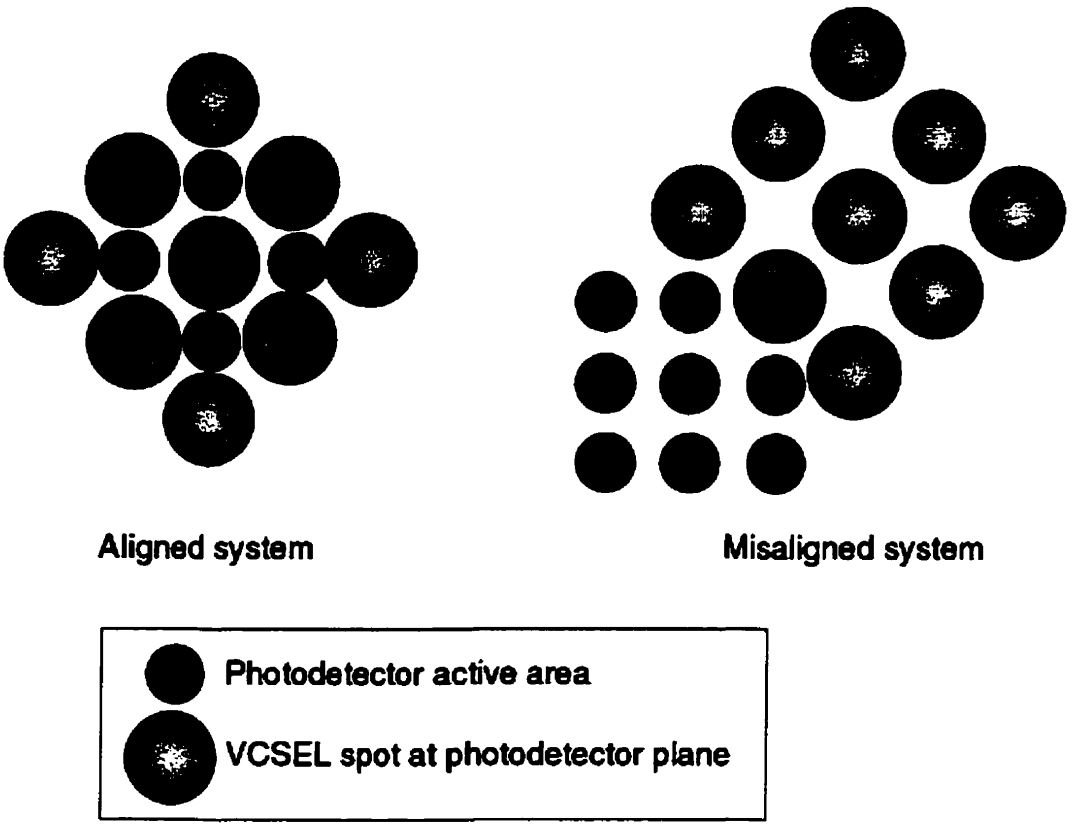

Figure 3. Optoelectronic redundancy leading to misalignment tolerance

In the aligned system shown of the left-hand side of Figure 3, there are 5 possible VCSEL-photodetector pairs that could be used to transmit data. In the misaligned system, there is one possible VCSEL-photodetector pair to be used as the optical channel.

Before the optical data is transmitted, there is an FPGA controlled hunt algorithm that finds the best channel by testing all 81 possibilities for each of the 4 links. The best channel is defined as the one with the highest received power, and is the one used to transmit the 1Gbit/s data across the link.

This method of choosing the best channel out of many for an optical link is easily scalable to much larger arrays of VCSELs and PDs. The larger the array,

the larger the misalignment tolerance in between the two circuit boards. With this in mind, the receiver array architecture was designed to be easily scalable. In addition, more than one channel could be used per array of VCSELs and PDs, if the circuitry and optics were appropriately designed. This would involve, for example, separating the VCSEL and PD arrays into clusters or subsets, where a single VCSEL-photodetector pair would be used per cluster to transmit data.

2.4 Conclusion

We designed and implemented a free-space board-to-board interconnect. There are two bi-directional links in between the boards, leading to an aggregate bit rate of 4Gb/s. Using the concept of optoelectronic redundancy and appropriately designed custom optics, the system can handle lateral misalignments of ±1mm and angular misalignments of ±1 degree, with the distance between the boards varying continuously from 5cm to 21.5cm.

2.5 References

1. D. V. Plant, et al., "A 5 Gb/s, 2 channel bi-directional adaptive redundant FSOI demonstrator system", in *Optics in Computing, 2000*, Roger A. Lessard, Tigran Galstian, Editors, SPIE Vol. 4089, pp. 465-472, 2000.

2. D. J. Goodwill et al., "Free space optical interconnect at 1.25Gb/s/channel using adaptive alignment," OFC/IOOC 1999, San Diego, USA, paper WM22, 1999.

3. E. Bisaillon, et al., "Optical link for an adaptive redundant free space interconnect", in *Optics in Computing, 2000*, Roger A. Lessard, Tigran Galstian, Editors, SPIE Vol. 4089, pp. 999-1009, 2000.

Chapter 3. Description of receiver architecture

3.1 Introduction

A receiver design for the system described in the previous chapter must have specific characteristics. It must accept current swings as input, and output digital PECL signals. It must have a very high sensitivity and a large dynamic range in order to handle the misalignment tolerance of the link that leads to a wide range of input optical powers. It must also be scalable since the design for the optical link is easily scalable to larger arrays of VCSELs and photodetectors [2], leading to increased redundancy and hence misalignment tolerance. Also, it must have a physical layout that makes it easily wirebondable to a photodetector array die.

The receiver chip was designed in NT25 technology [2], which features bipolar transistors with an f_T of 25GHz. This process was made available through the Canadian Microelectronic Corporation (CMC) foundry service [3]. We chose to use a bipolar process rather than CMOS to increase our chances of successfully building transceivers that function at gigabit data rates. The receiver accepts as input nine analog input current swings from the 3×3 array of photodetectors. To make the receiver scalable to a larger number of inputs, a system of switches was used to select one of the nine inputs and amplify it to digital PECL levels. This one amplified differential signal is the electrical output of the chip.

The receiver was designed to operate with a bandwidth from 100Mb/s to 1Gb/s, and accept input current swings ranging from 2.5μA to 25μA. Simulations

of the circuit performance, shown in Chapter 4, indicate that these targets were reached. Experimental results are shown in Chapter 7, and though all of the targets were not reached, the results are comparable to the simulation results.

3.2 Receiver architecture

The receiver architecture is shown in Figure 4. The nine photocurrents from the nine photodetectors are input to the chip through wirebonds. Each signal is independently amplified through transimpedance amplifiers and then through a single-ended gain stage, as shown on the left-hand side of the figure.

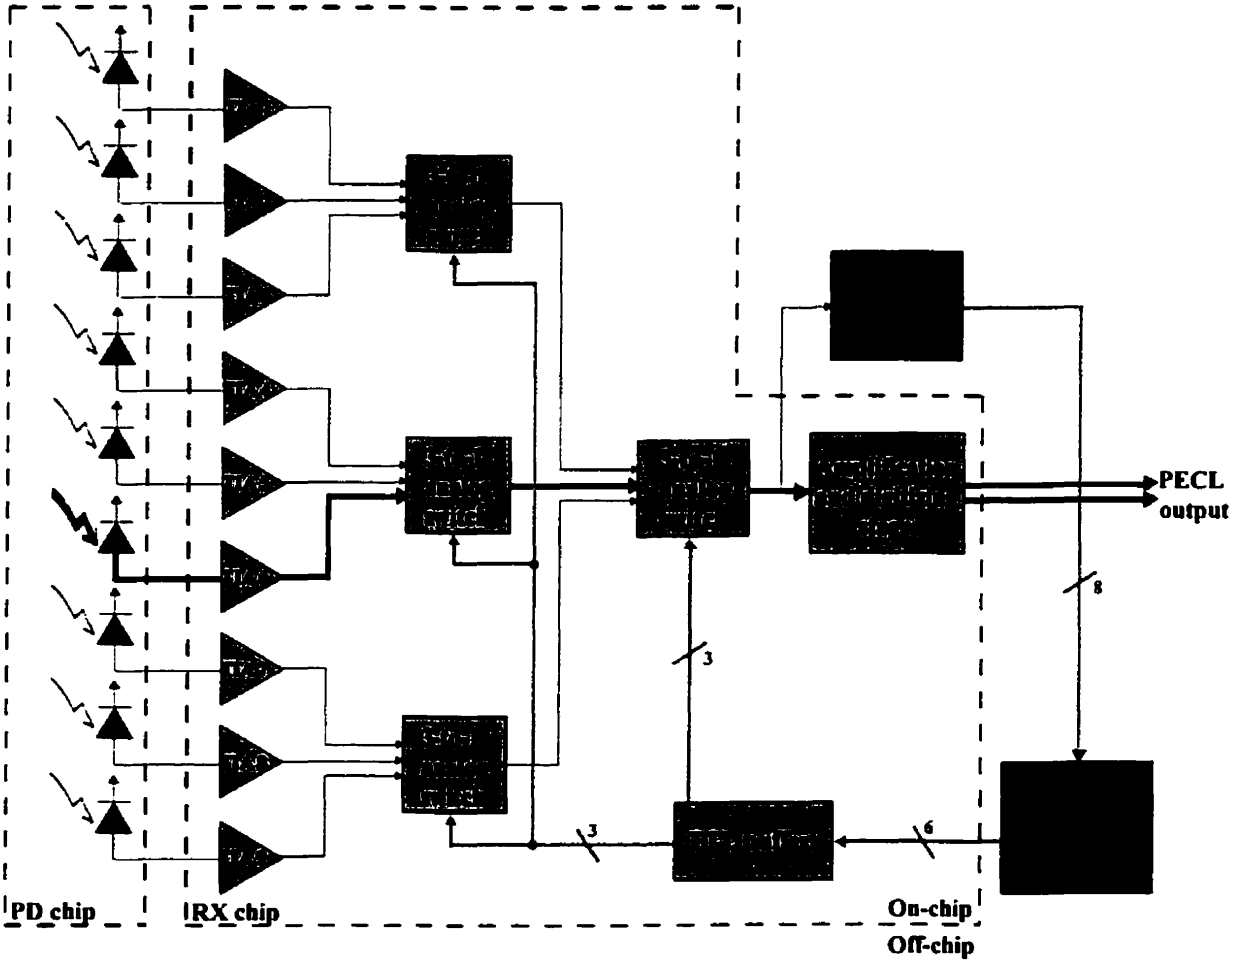

Figure 4. Receiver chip architecture

The amplified single-ended signals then go though a pyramid of switches in which only one of the nine signals is selected. The one selected is the one receiving the most optical power. This selected signal then goes to a final amplification stage and digitizing stage before going through an output buffer and off-chip as a digital PECL signal. In Figure 4, the 6^{th} photodetector is assumed to be the one receiving the largest signal, thus this signal is fully amplified and directed to the output.

3.2.1 Photodiode array

The nine input currents to the receiver come from a 3x3 array of photodiodes. A photo of the photodiode array is shown in Figure 5. As can be seen in the figure, the photodiode array is a 5x5 array but only a 3x3 subset is used for this system.

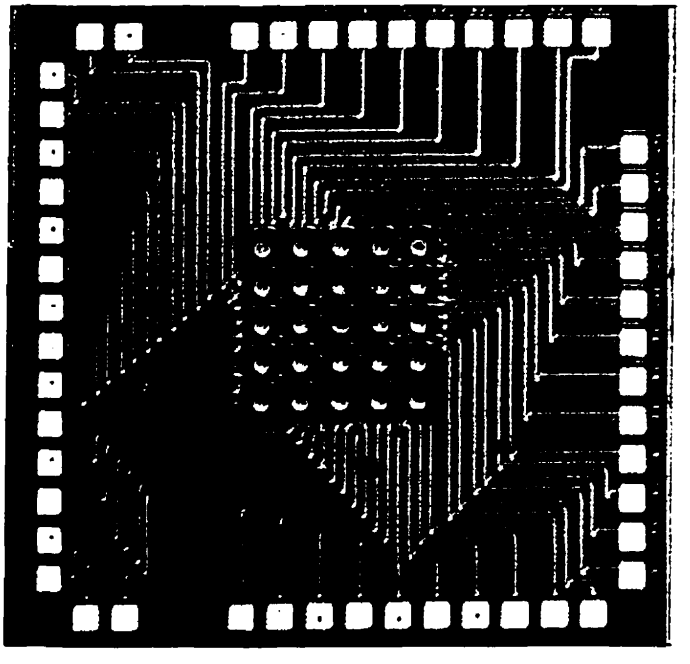

Figure 5. Photodiode array

The receiver die supplies 4V to the anode of the photodiodes in order to reverse bias them by 2.5V, since the cathode is held at 1.5V. This is illustrated in Figure 6.

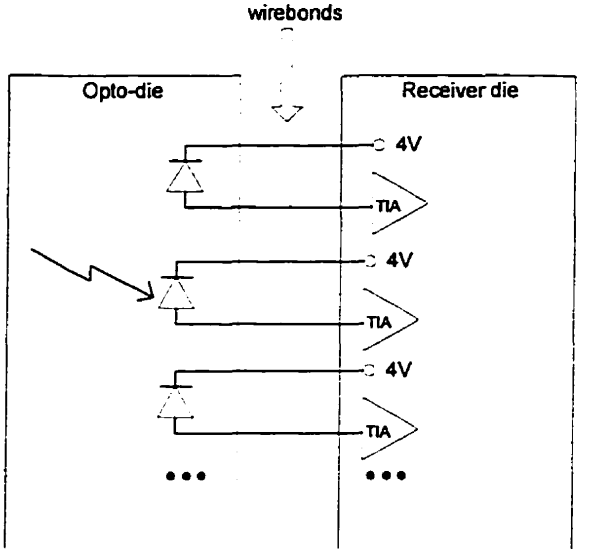

Figure 6. Photodiode array and TIA configuration

The photodiodes are top-illuminated InP PIN photodiodes, and have a responsivity of 0.6A/W at 960nm and at - 2.5V bias [4]. According to simulations of the optical link [5], they will receive between -15dBm and -21dBm of average optical power. This leads to photocurrents ranging from 4.5μA to 18μA. As a safety margin, the receiver was designed to have a larger dynamic range, from 2.5μA to 25μA. The dark current at –2.5V reverse bias is 50nA and is therefore negligible.

The rise and fall times (10%-90%) for the photodetectors at –2.5V bias are approximately 150ps [4]. This leads to possible data rates of 2.2Gb/s, if the RC

time constant of the receiver is not too high. The capacitance of the photodiodes at –2.5V bias has a maximum value of 0.7pF, and a typical value of 0.5pF [4].

The diameter of the active area of the photodetectors is 70μm, and they are on a 125μm center-to-center pitch. The die thickness is 150μm, and the die size 2.35mm x 2.35mm. The bond pads are gold with a thickness of 1.2μm, and are 90μmx90μm. The bond pad pitch is 125μm. A schematic of the photodetector die layout is shown in section 5.3, Figure 36. The photodiode model used for circuit simulations is discussed in section 4.2.

Because of the large diameter of the photodetectors, their capacitance is considerable. This limits the input resistance of the TIA, since the rise time will be too large if the resistance is too high. This is discussed further in the next section.

3.2.2 Transimpedance amplifiers

The photocurrents arriving through wirebonds from the photodetectors are amplified by nine transimpedance amplifiers (TIAs) through a 1.2kΩ transimpedance resistance. A transimpedance amplifier was chosen as a pre-amplifier, rather than choosing a high-impedance amplifier, because its lower input impedance allows for a lower time constant, which translates to a higher frequency response [8]. The trade-off is a reduced sensitivity of the receiver.

At the output of the TIA there are single-ended gain stages with a combined gain of 20dB. For an input current swing from the photodetectors of 2.5μA to 25μA, this leads to a single-ended voltage swing ranging from 30mV to

300mV. Section 4.3 shows the simulation results for this stage. Figure 7 shows the circuitry for the TIA followed by the single-ended gain stages.

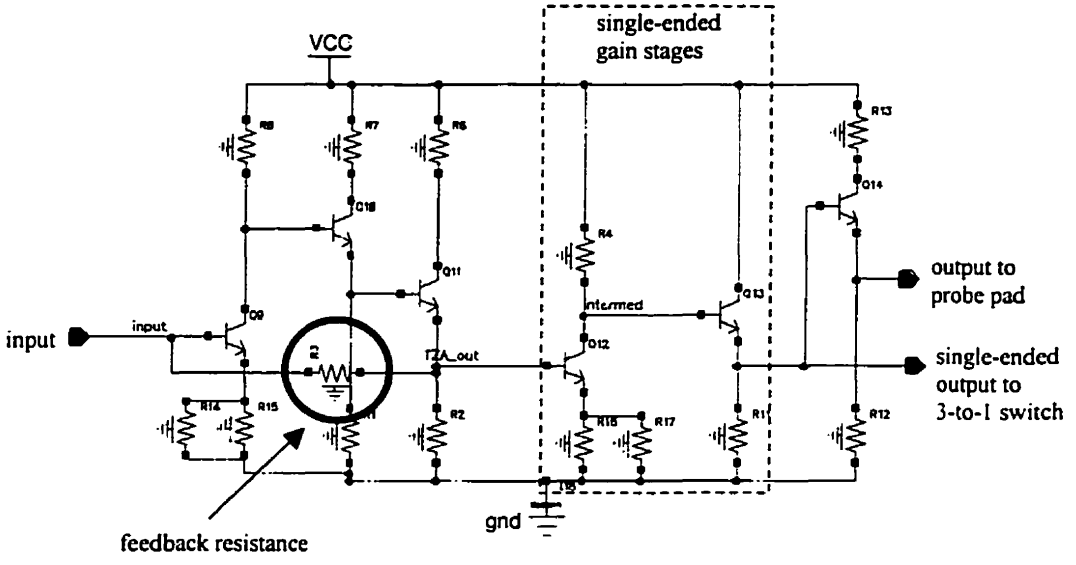

Figure 7. TIA and single-ended gain stages circuitry

The choice of the size of the feedback resistor was crucial to this circuit design. The bandwidth of a transimpedance receiver can be approximated to:

$$BW = \frac{A}{2\pi R_f (C_t + AC_f)} \qquad [8]$$

where A is the amplifier gain (A>>1), C_f is the stray capacitance of the feedback circuit, C_t is the total input capacitance including the contribution from the photodiode, and R_f is the feedback resistance. So, decreasing R_f leads to a larger bandwidth for the TIA. However, the output voltage of the TIA is given by:

$$V_{out} = i_{PD} R_f \qquad [8]$$

where i_{PD} is the photocurrent. Therefore, increasing R_f leads to a larger gain. Through simulations, the best compromise between bandwidth and gain was found when $R_f = 1.2k\Omega$.

After the single-ended gain stages, the signal is sent on to the next stage (3-to-1 analog switch), and to the base of another transistor, labeled Q14. This transistor is in emitter follower configuration with emitter degeneration [11]. The resistance in series with the emitter increases the input resistance seen by the signal. This helps reduce the transistor's loading effect on the signal. A probe pad is attached to the collector of this transistor. This pad can be probed to monitor the output of the TIA stage, to aid with the testing of the receiver. Because of the size of the current that may be drawn while probing, this transistor was made 12 times larger than the minimum size transistor for the NT25 technology [2]. All other transistors in this figure are minimum size.

3.2.3 3-to-1 analog switches

Though nine photocurrents are sent to the receiver array, only one signal is selected and amplified to digital PECL levels before being sent off-chip. After the TIAs and the single-ended gain stages, there is a pyramid of analog switches that selects one of the 9 single-ended voltage swings sent to the pyramid, as shown in Figure 4. The pyramid is composed of two levels of 3-to-1 analog switches, which do not provide any gain or loss. The nine voltage swings are sent in groups of three to the first layer of three switches. After that first layer of switches, there are only three signals left. There is then a second layer to the pyramid, consisting of only one 3-to-1 analog switch. After the last switch, the

nine original signals have been reduced down to one, and the amplitude of the output is the same as the amplitude of the selected input, since the 3-to-1 analog switches provide 0dB gain. Simulation results for this stage are shown in section 4.4.

Each 3-to-1 analog switch is controlled by three select lines. When one of the three lines is low (2.7V) and the other two are high (3.3V), the line that is low selects the corresponding input and passes it to the output. Figure 8 shows the circuitry for the analog switch.

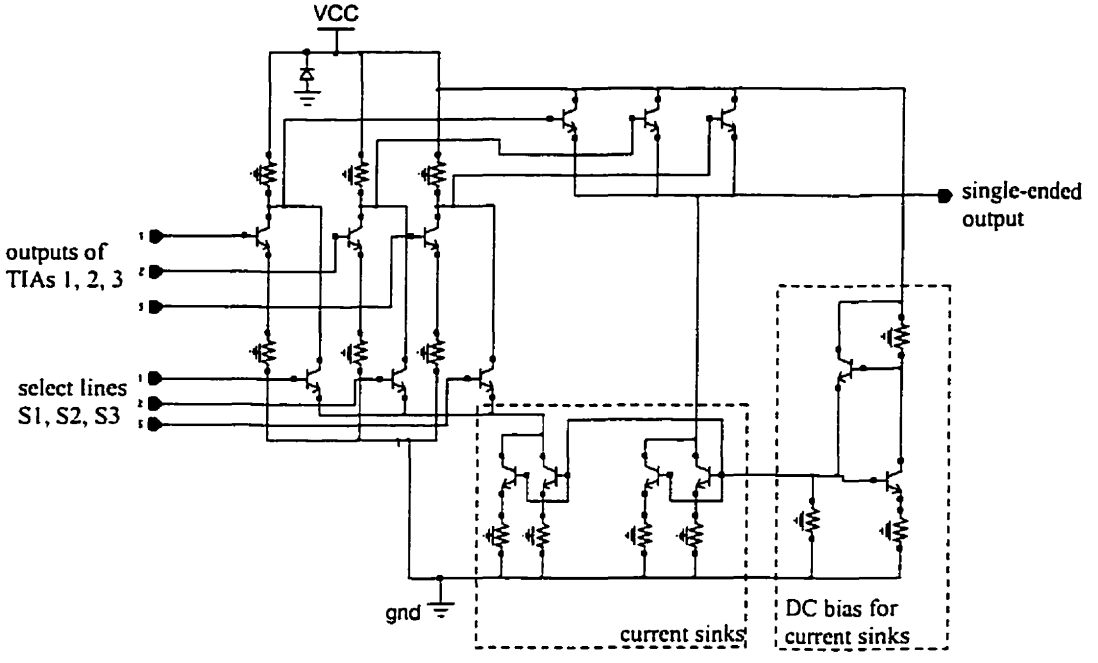

Figure 8. 3 to 1 analog switch circuitry

The three switches that are in the first layer of the pyramid share the same three control lines: S1, S2, S3. If the second select line (S2) is low and the other two are high, inputs in2, in5, and in8 are all selected (these inputs are indicated in Figure 4). The last switch in the second layer has its own independent select

lines: S4, S5, S6. Table 1 shows the select line configurations necessary to select the nine photodiode inputs, with "0" representing the low voltage (2.7V) and 1 representing the high voltage (3.3V).

Input	S1	S2	S3	S4	S5	S6
1	0	1	1	0	1	1
2	1	0	1	0	1	1
3	1	1	0	0	1	1
4	0	1	1	1	0	1
5	1	0	1	1	0	1
6	1	1	0	1	0	1
7	0	1	1	1	1	0
8	1	0	1	1	1	0
9	1	1	0	1	1	0

Table 1 . Select line configuration to select input signal

The voltages for the six select lines are controlled by off-chip circuitry, and have TTL voltage swings when they arrive on-chip. In total, there are six TTL buffers that convert the TTL swings to on-chip 2.7V to 3.3V swings [7]. The maximum switching speed of the buffers is 50Mhz and this determines the maximum switching speed of the pyramid.

Rather than select only one of the signals, a design was considered where all 9 received signals would be "ORed" together. The main advantage with this architecture would be that no channel selection would be necessary, and if the VCSEL spot shifted from one photodiode to the other no data would be lost. However, this design is not scalable to larger arrays of photodetectors where two or more independent optical data links might be established on the same

photodiode array. Therefore a switching mechanism was chosen instead, even thought it meant an increase in the design complexity.

At the output of the switch pyramid, the signal is sent to two different places simultaneously, as shown in Figure 4. Part of the signal remains on-chip and is sent to a last amplification and digitizing stage, which is described in the next section. The remaining part of the signal is sent off-chip for power monitoring purposes. The first element of the power monitoring circuitry is a high-speed buffer placed immediately off-chip, right next to the carrier to minimize trace length. This high-speed buffer limits the loading effect of the off-chip circuitry onto the pyramid circuitry. The off-chip circuitry monitors the power of the output signal of the pyramid, and is described in section 6.3.

3.2.4 Final amplification and digitization stage

The final amplification stage, which is shown in Figure 9, takes the single-ended output signal of the switch pyramid, and amplifies it to digital PECL levels. The circuit consists of two cascaded differential amplifiers. The first differential pair converts the single ended-signal into a differential signal with the aid of a lowpass circuit. The single ended-signal is sent to the left-hand side of the differential pair. This same signal is sent through an RC lowpass circuit to the right-hand side of the differential pair. This sets the DC reference voltage for the differential pair. This design was chosen so that the receiver would be independent of any DC bias on the received signal. The simulations for this circuit are shown in section 4.5.

The second differential pair amplifies the differential signal to digital PECL voltage swings. In between the two stages there is an emitter follower buffer, one for each of the two differential lines.

Finally an output buffer, capable of driving a 50Ω line, was included in the design. The overall receiver output is differential and has a voltage swing from 3.1V to 3.9V, giving a differential swing of 1.6V

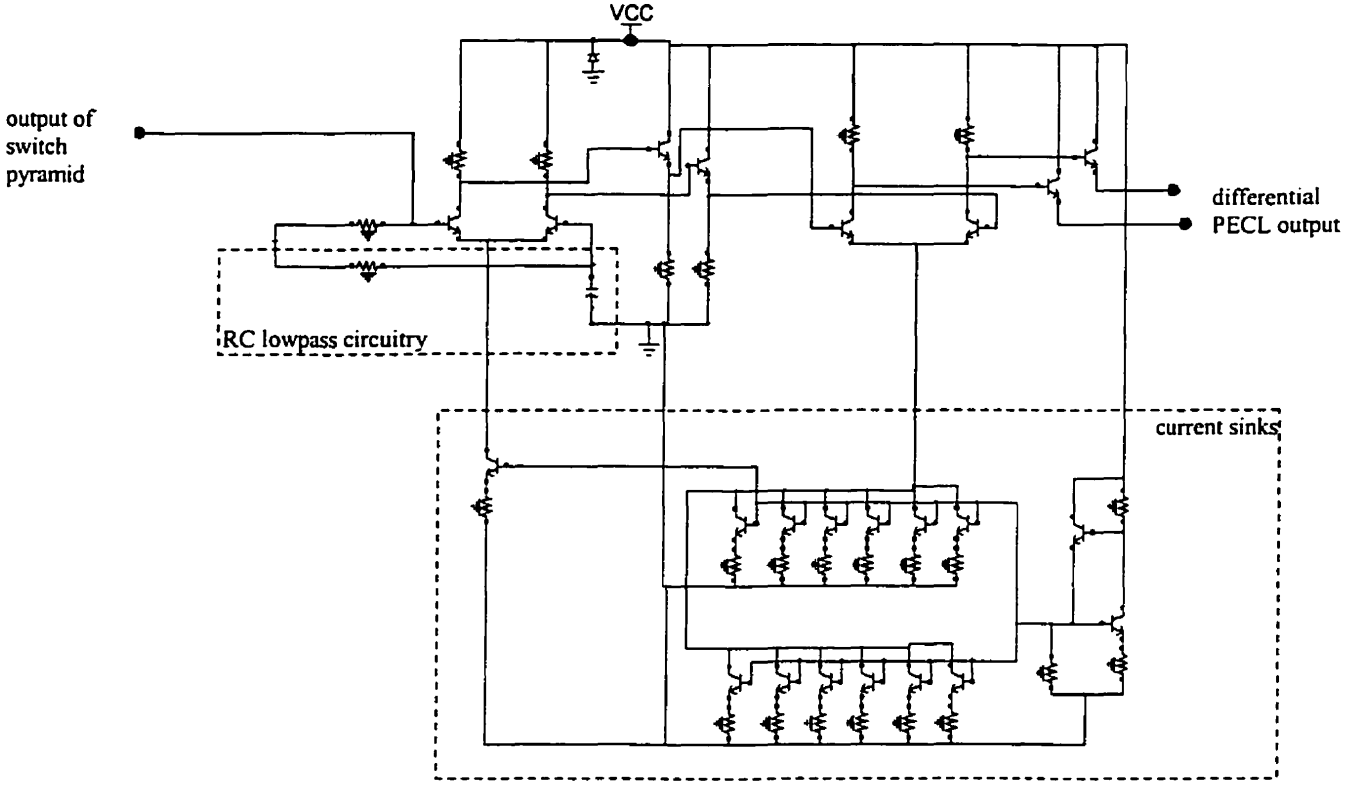

Figure 9. Amplification and digitization circuitry

3.2.5 Alternative final stage

In the previous section, the input signal to the last stage is the signal at the output of the pyramid of switches and is DC coupled to the final amplification

stage. An alternative stage was also designed which must be AC coupled to the output of the pyramid of switches.

The output signal from the pyramid is sent to a carrier pin, which is connected to a 120nf capacitor. The other terminal of the capacitor is connected to the adjacent carrier pin, which is wirebonded to the input of the alternative last stage. This circuitry is shown in Figure 10. The main difference is that the DC component of the signal will be removed by the off-chip capacitor and the DC bias will instead be provided by additional circuitry. This alternative last stage was designed in part to have more control over the lower frequency cutoff of the receiver (this is discussed further in section 4.6). By replacing the on-board capacitor with a capacitor with a different value, the bandwidth of the receiver could be changed. It was also implemented on-chip in the event that this design would be more immune to noise than the one described in the previous section. However, this alternative circuit was found to be noisier once packaged and tested, and therefore was not used for the system implementation.

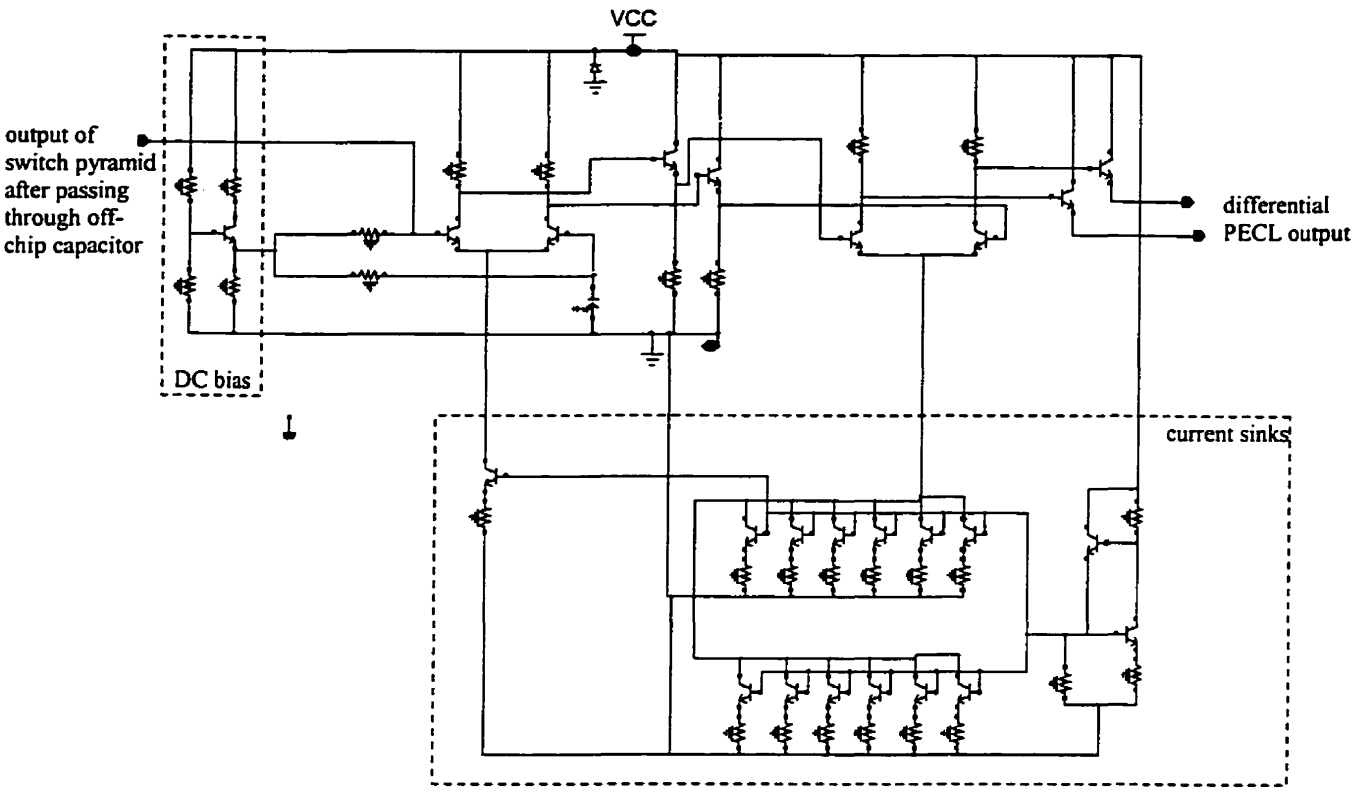

Figure 10. Alternative last stage circuitry

3.3 Test circuitry

Several different test circuits have been included on this chip and are completely independent from the rest of the circuitry on the chip and from each other. The intended purpose of this test circuitry is to independently characterize each of the blocks described above. Once the chip is in its final package, this test circuitry is no longer accessible.

3.4 Layout

The layout of the chip respected both the high-speed constraints and the physical packaging constraints. These constraints included maintaining the signal-integrity of the high-speed signals, minimizing crosstalk, and minimizing

the length of the wirebonds to the photodiode array. The layout of the chip is as shown in Figure 11. The different sections of the chip are labeled in Figure 12.

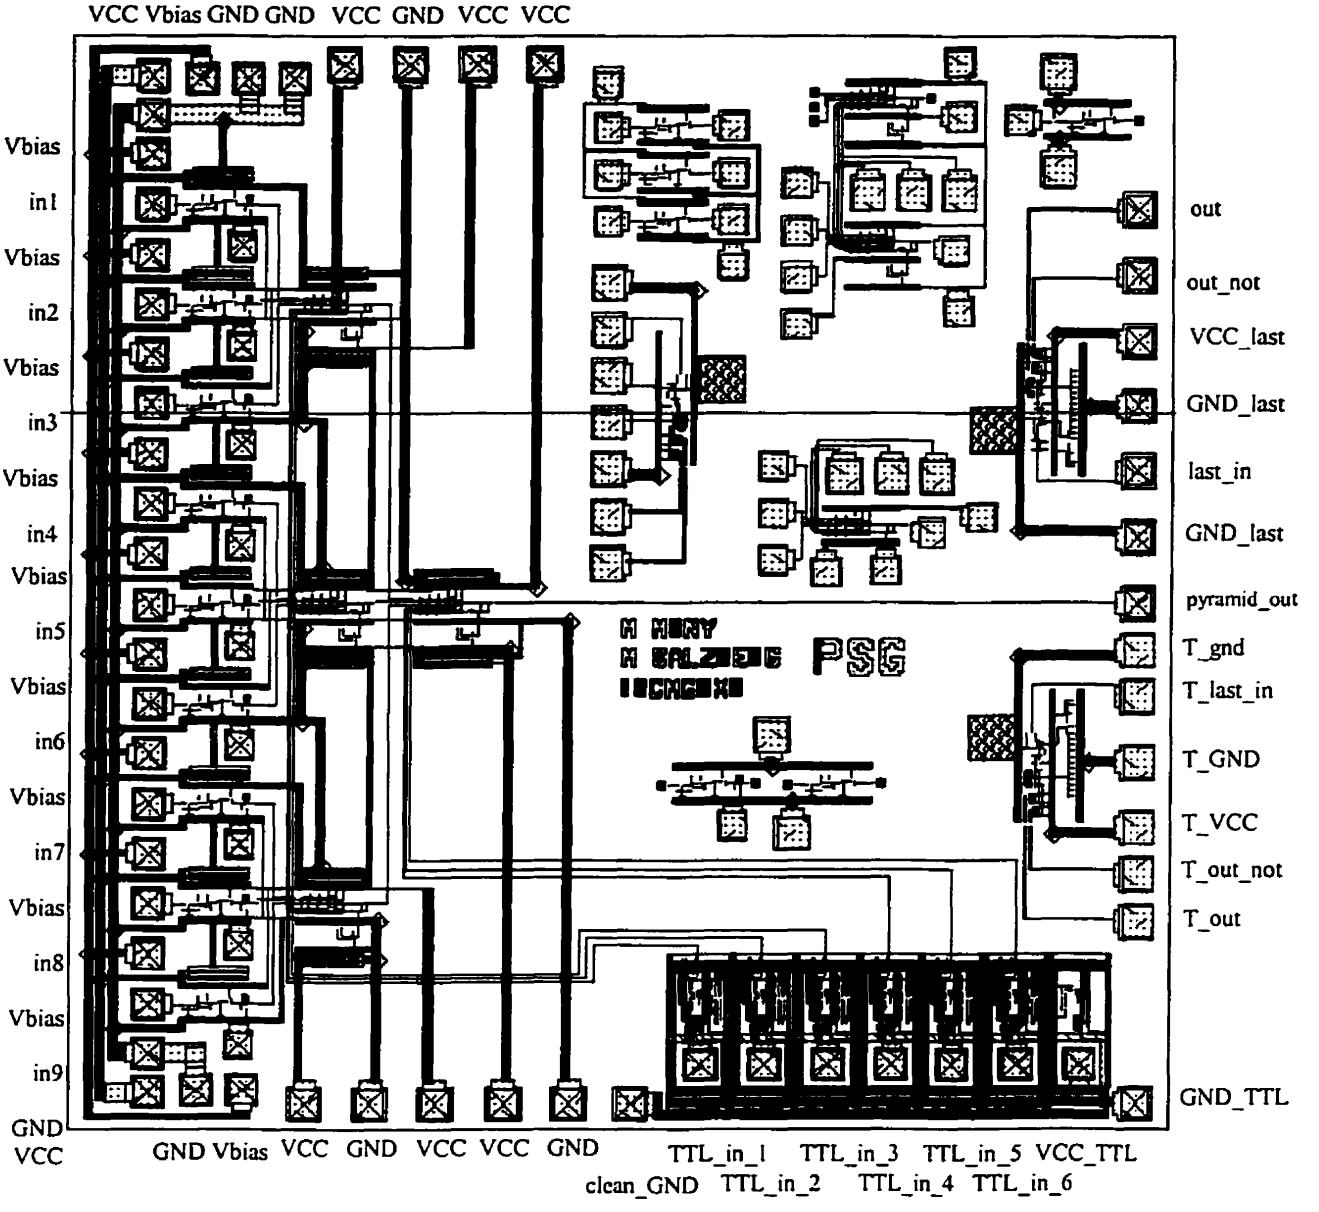

Figure 11. Layout of receiver chip

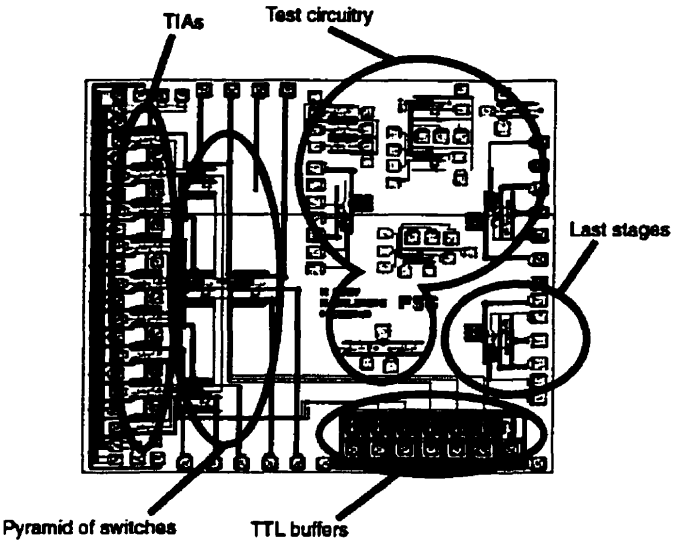

Figure 12. Layout with different sections identified

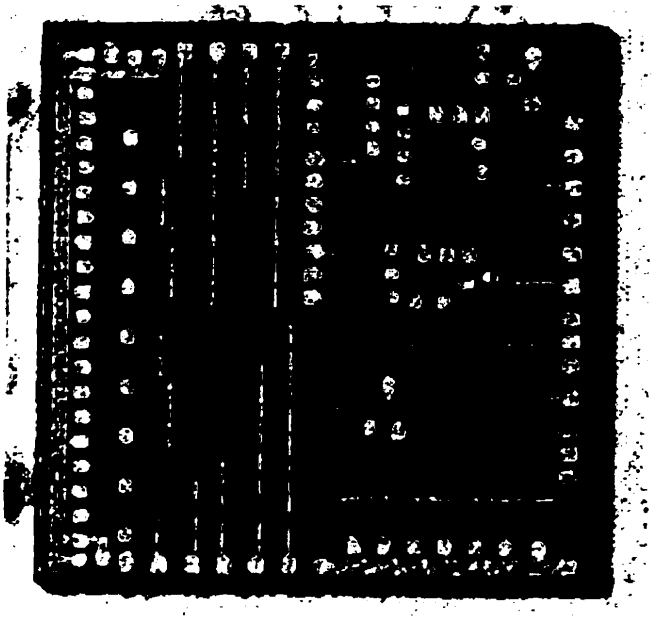

Figure 13. Photo of receiver chip

In order to wirebond the receiver chip to the photodiode chip, all of the bond pads that connect to that chip need to be on one side of the die, in this case the left side. Therefore, all of the TIAs are on the left-hand side of the layout.

Each photodiode needs to be wirebonded to a pad that is attached to the reverse biasing voltage as described in section 3.2.1. To route this bias voltage (called Vbias) to all the wirebond pads, a long strip of metal was used, as is shown schematically in Figure 14. Each TIA also needs to be connected to a common ground, which is also routed in a long strip of metal. The same can be said for the VCC 5V supply for each TIA. By laying them out as shown in Figure 14, with the bond pads on the *inside* of the metal traces, a capacitance is formed in between all the power levels. This decoupling capacitance acts as a low-pass filter and helps eliminate any high-frequency components on the DC power lines [10]. The accumulated capacitance from all of the trace crossings is on the order of 300fF, and through simulations was found to give a –3dB cutoff frequency around 900MHz. This helps to remove some of the very high frequency noise on the DC lines of the chip.

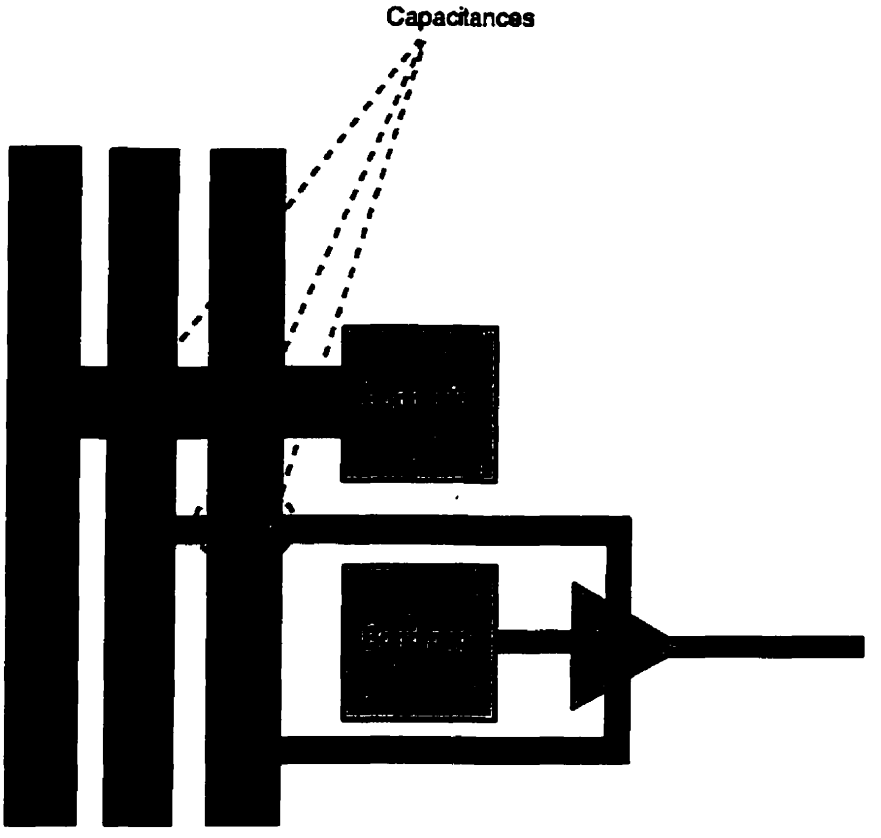

Figure 14. Layout configuration causing decoupling capacitances

Other methods were also used to help maintain signal integrity. In between every two neighboring TIAs, there are two large back-to-back reverse biased p-n diodes. Any charge that may be generated by one of the TIAs will not be able to migrate to a neighboring TIA, but will be swept up by the reverse biased diodes. This helps reduce crosstalk noise in between the receiver channels.

Also, in order to minimize the effect of manufacturing imprecision, all of the resistors in the chip were laid out in the same direction (vertical) and have the same width. So any variation of the width of the resistors will affect them all equally, thus reducing the impact on the circuit performance. This is especially true for the differential pairs, so that each side of the pair remains balanced.

Finally, multiple ground and VCC bond pads were distributed around the chip. This minimizes the chance that part of the circuitry will not be powered by a full 5V, since none of the traces from the bond pads to the circuits are very long. So, there is a minimal resistive voltage drop across the metal routing the power to different parts of the chip.

3.5 Conclusion

The receiver chip was designed to be a 9x1 high-speed receiver array, capable of receiving low amplitude current swings and converting them to digital PECL signals that are sent off-chip. The design was based on scalability, so that the same circuitry could be applied to larger arrays of photodetectors for increased redundancy.

3.6 References

1. D. J. Goodwill et al., "Free space optical interconnect at 1.25Gb/s/channel using adaptive alignment," OFC/IOOC 1999, San Diego, USA, paper WM22, 1999.
2. "NT25: Nortel's High-Speed Bipolar Process", White Paper, Nortel Semiconductors.
3. Canadian Microelectronic Corporation (CMC), http://www.cmc.ca
4. Opto Speed SA, http://www.optospeed.com
5. E. Bisaillon, et al., "Optical link for an adaptive redundant free space interconnect", in *Optics in Computing, 2000*, Roger A. Lessard, Tigran Galstian, Editors, SPIE Vol. 4089, pp. 999-1009, 2000.

6. M. Mony, et al. "1Gb/s, 9x1 optical receiver array, for an adaptive redundant interconnect system, using NT25 technology", in *Optics in Computing, 2000*, Roger A. Lessard, Tigran Galstian, Editors, SPIE Vol. 4089, pp. 1010-1020, 2000.

7. E. Bernier, et al., "1Gb/s transmitter for an adaptive redundant free-space interconnect", in *Optics in Computing, 2000*, Roger A. Lessard, Tigran Galstian, Editors, SPIE Vol. 4089, pp. 772-783, 2000.

8. Robert J. Hoss, "Fiber Optic Communications, Design Handbook", Prentice Hall, 1990.

9. G. P. Agrawal, "Fiber-optic communication systems", Wiley, 1997.

10. J. M. Rabaey, "Digital integrated circuits, a design perspective", Prentice Hall, 1996.

11. P. Gray, R. Meyer, "Analysis and design of analog integrated circuits", 3rd edition, John Wiley & Sons, Inc., 1993.

12. A. Sedra, K. Smith, "Microelectronic circuits", Harcourt Brace College Publishers, 1991.

13. B. Streetman, S. Banerjee, "Solid State Electronic Devices", 5th edition, Prentice Hall, 2000.

14. G. Neudeck, "The bipolar junction transistor", Vol.3 , 2nd edition, Addison-Wesley, 1989.

15. M. Neuhauser, et al., "13Gbit/s Si bipolar preamplifier for optical front ends", *Electronic Letters*, Vol.29, No. 5, pp. 492-493, 1993.

16. M. Moller, et al., "13 Gb/s Si-bipolar AGC amplifier IC with high gain and wide dynamic range for optical-fiber receivers", *IEEE Journal of Solid-State Circuits*, Vol. 29, No. 7, pp.815-821, 1994.

17. M. Neuhauser, et al., "Low-noise, high-gain Si-bipolar preamplifiers for 10Gb/s optical-fiber links, design and realization", *IEEE Journal of Solid-State Circuits*, Vol. 31, No. 1, pp.24-29, 1996.

18. A. Felder, et al., "46 Gb/s DEMUX, 50Gb/s MUX, and 30 GHz static frequency divider in silicon bipolar technology", *IEEE Journal of Solid-State Circuits*, Vol. 31, No. 4, pp.481-486, 1996.

19. T. Vanisri, et al., "Integrated high frequency low-noise current-mode optical transimpedance preamplifiers: theory and practice", *IEEE Journal of Solid-State Circuits*, Vol. 30, No. 6, pp.677-685, 1995.

20. G. Nasserbakht, et al., "A monolithic GaAs-on-Si receiver front end for optical interconnect systems", *IEEE Journal of Solid-State Circuits*, Vol. 28, No. 6, pp.622-630, 1993.

21. R. Reimann, et al., "Bipolar high-gain limiting amplifier IC for optical-fiber receivers operating up to 4Gbit/s", *IEEE Journal of Solid-State Circuits*, Vol. 22, No. 4, pp.504-511, 1987.

Chapter 4. Simulation results

4.1 Introduction

To aid in the design of the receiver, several simulations were performed. Each stage of the receiver was tested and characterized separately, then all the stages were combined and the overall performance was evaluated. Once the layout was completed, simulations including all parasitic capacitances were performed. All simulations were performed in the Cadence environment, using HSpice. This chapter will present some of the more important simulation results for each of the stages of the receiver.

4.2 Photodiode model

The photodiode model was based on specifications provided by the manufacturer [1]. The following figure shows the photodiode model used.

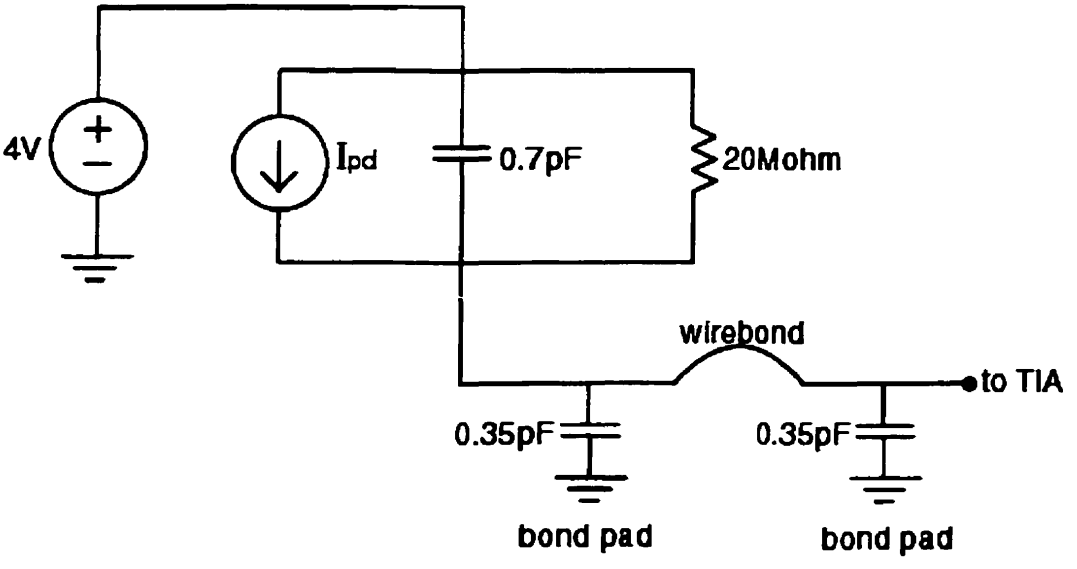

Figure 15. Photodetector model

Each bondpad was approximated by a 350fF capacitor. The wirebond was simulated with a wirebond model, with a wire length of 1mm, and a wire thickness of 25μm. The wirebond could also have been approximated by a 1nH inductance [2]. These dimensions for the wirebond are based on the actual packaging used for the chip.

The responsivity of the photodiodes at 980nm was specified to be 0.65A/W. For simulations, a conservative estimate of 0.5A/W was used. With an estimated –23dBm to –13dBm input optical power to the photodetectors (or 5μW to 50μW) [3], this leads to a high (digital 1) value for the photocurrents ranging from 2.5μA to 25μA. The contrast ratio from the VCSELs was estimated to be 10:1 [4], so the low (digital 0) current value ranged from 0.25μA to 2.5μA. The rise and fall times (10%-90%) were specified to be 150ps, or 19% of a bit period for each, for a bit rate of 1.25Gb/s.

Using the model of the photodiode as shown in Figure 15 for the simulations, the input photocurrents to the receiver chip, and hence to a TIA, are shown in Figure 16 and in Figure 17. Three different input current swings were chosen for the tests: 2.5μA, 13.5μA, and 25μA. These correspond to –23dBm, -15.7dBm, and –13dBm of input optical power. Since the receiver is specified to function within that input power range, the full dynamic range of the receiver was tested during each of the simulations. The following simulation results for all of the stages of the receiver are shown for bit rates of 100Mb/s and of 1.25Gb/s, since that is the bandwidth within which the receiver is specified to operate.

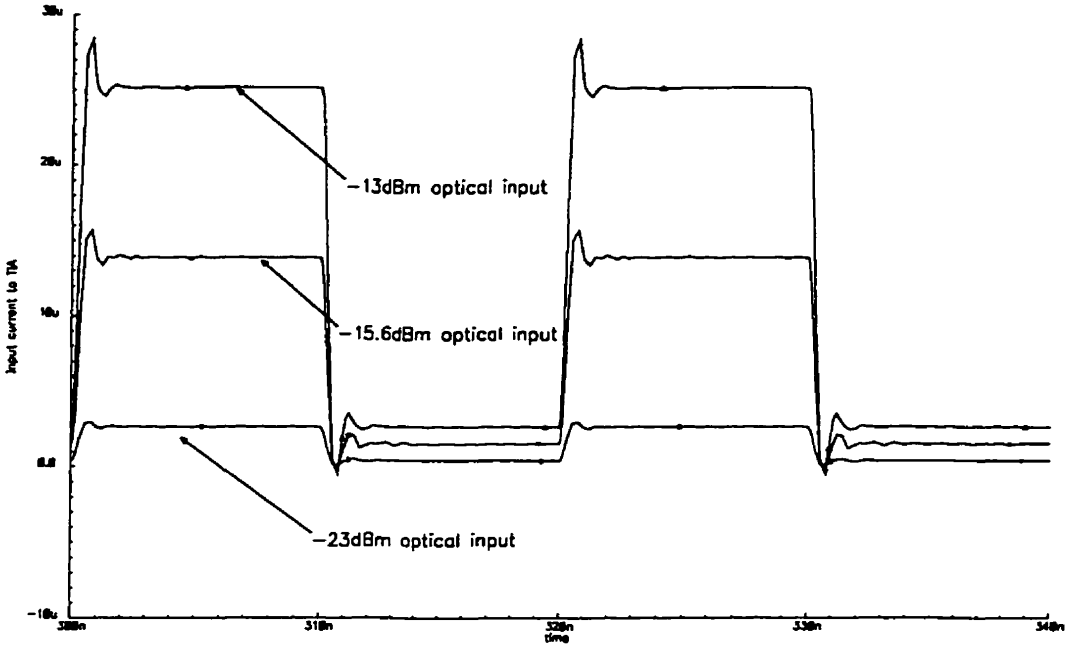

Figure 16. Simulation results at 100Mb/s for 3 different power levels onto PD using PD model in Figure 15 (PD loaded with TIA).

In Figure 16, the effect of the 10:1 contrast ratio can be clearly seen since the low (digital 0) currents are not 0A.

Figure 17 shows the currents generated in the photodetector at the same three input optical power levels as in Figure 16, but at a bit rate of 1.25Gb/s. For this simulation (and as for the one in Figure 16), the photodiode was loaded by the TIA. At a bit rate of 1.25Gb/s, the capacitance of the photodetector and the feedback resistance of the TIA have a considerable impact on the shape of the input waveform. Reducing the feedback resistance of the TIA can reduce the rise and fall times of the input signal (as well as the output signal), but at the expense of gain [7]. This was discussed in section 3.2.2. To illustrate this further, instead of loading the photodiode with the TIA stage, it was loaded with a 50Ω resistor.

The results of that simulation at 1.25Gb/s can be seen in Figure 18. The rise and fall times are much shorter due to the smaller time constant.

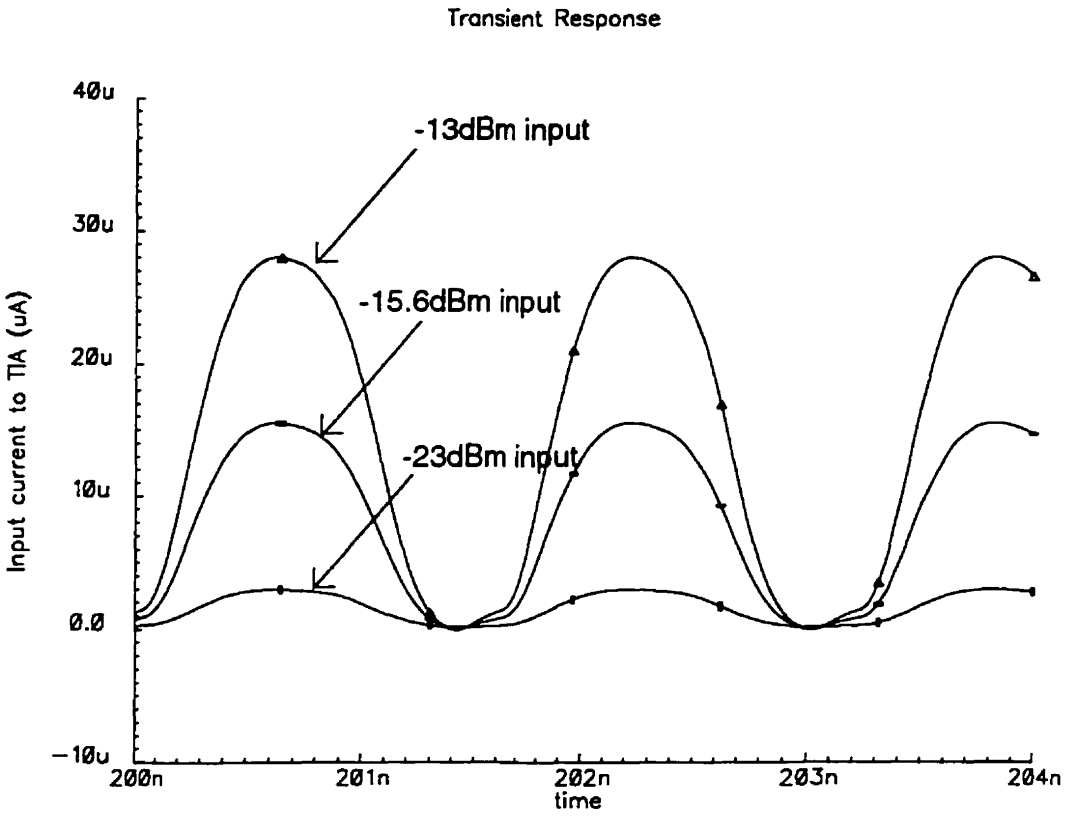

Figure 17. Simulation results at 1.25Gb/s for 3 different power levels onto PD using PD model in Figure 15 (PD loaded with TIA).

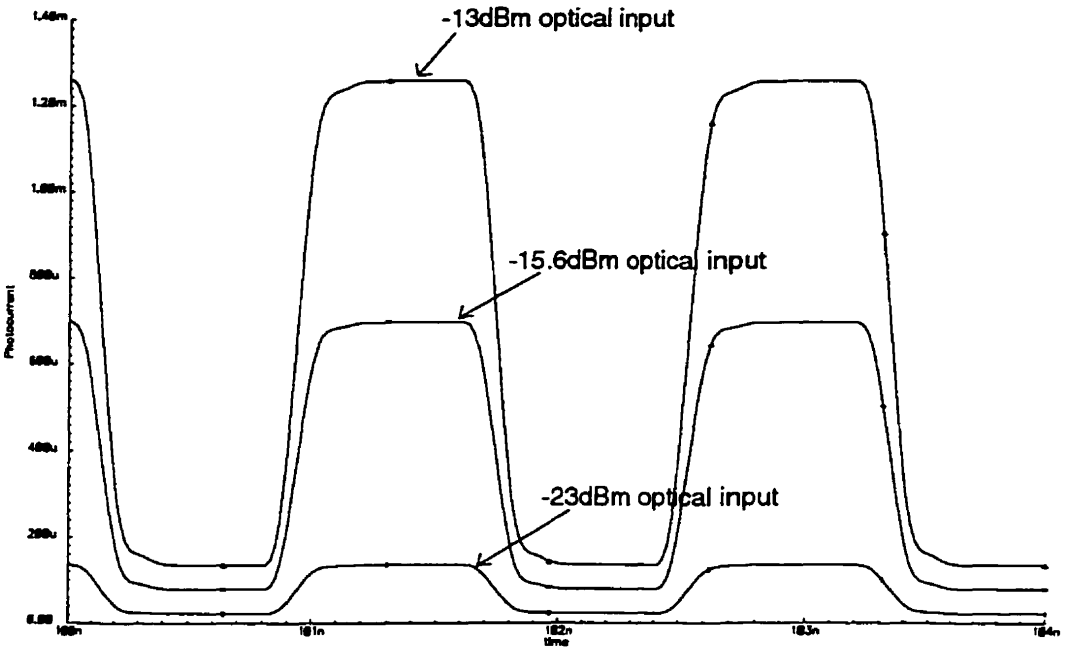

Figure 18. Simulation results at 1.25Gb/s for 3 different power levels onto PD using PD model in Figure 15 (PD loaded with 50Ω load).

4.3 TIA performance

Figure 19 and Figure 20 show the output of the TIA for the 3 different input current swings (2.5μA, 13.75μA, 25μA) for the two bit rates of 100Mb/s and 1.25Gb/s. From these simulation results, the gain of the TIA was found to be 10500 V/A at 100Mb/s, and 12700 V/A at 1.25Gb/s. This gain was suitable, since it was sufficient for the last stage of the receiver to amplify the signal to digital PECL levels, as will be shown in section 4.5.

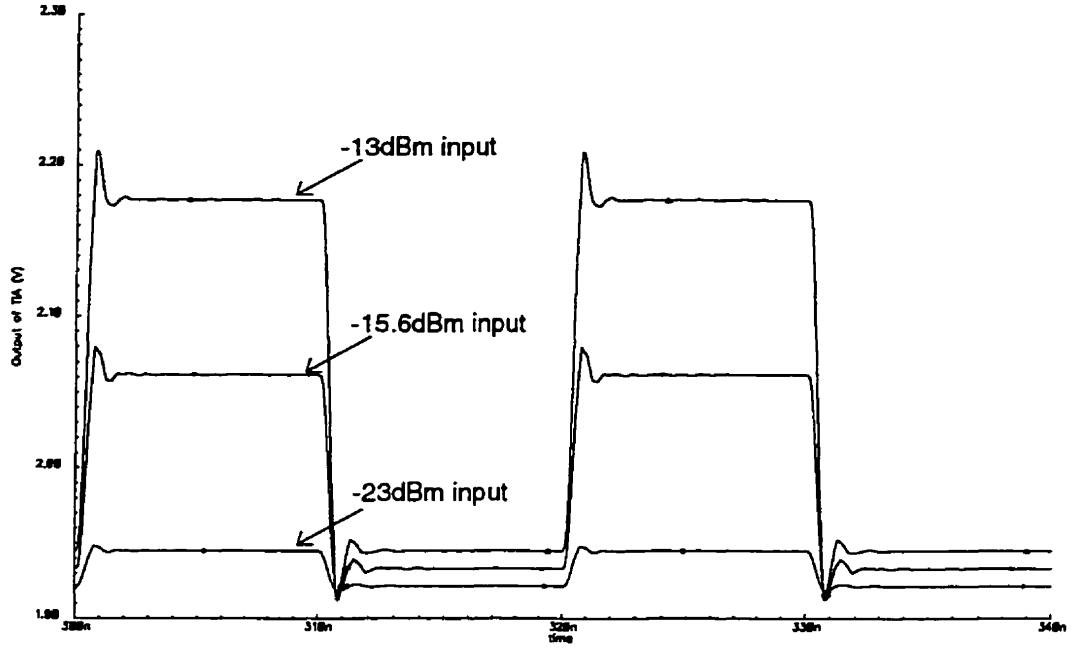

Figure 19. Output voltage swing from TIA at 100Mb/s

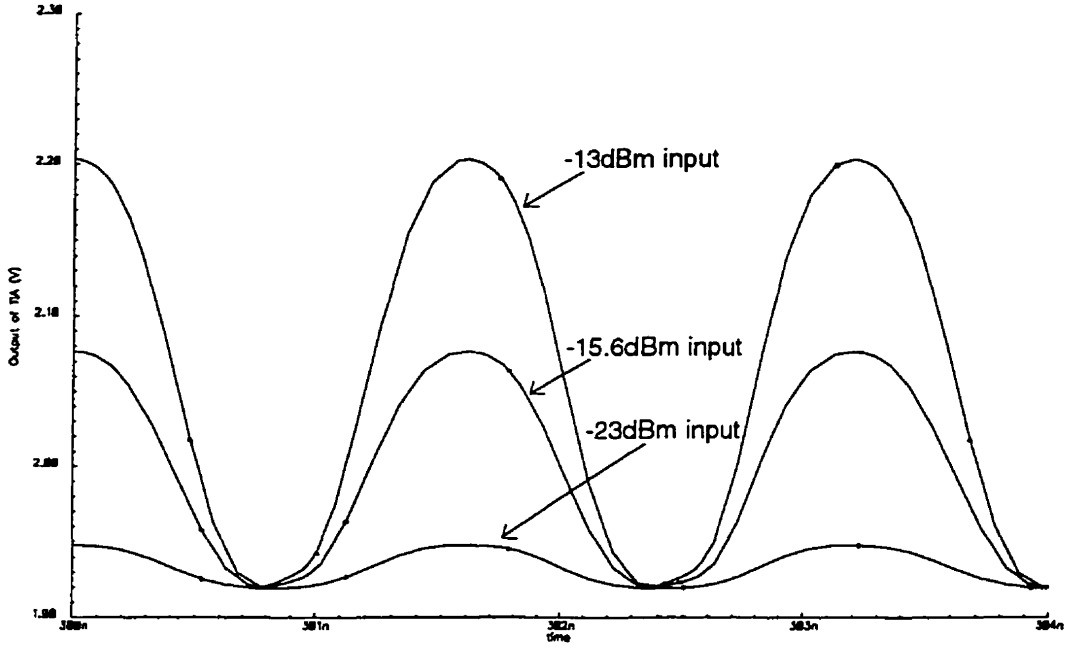

Figure 20. Output voltage swing from TIA at 1.25Gb/s

As discussed in the previous chapter, the choice of feedback resistance was critical in determining the performance of the TIA. Increasing the resistance

increased gain but decreased the bandwidth of the TIA. Through multiple simulations, the resistance that offered the beast compromise between the two was found to be 1.2kΩ. The bandwidth performance of the TIA will be presented in section 4.6.

4.4 3-to-1 analog switch performance

The output of the TIA in the previous section was fed into the pyramid of analog 3-to-1 switches, as shown in Figure 4. For this simulation, the signal was received at TIA number 5, so the control lines to the pyramid were setup to select the output from that TIA. The outputs of the first analog switch in the pyramid are shown in Figure 21 (at a bit rate of 100Mb/s) and in Figure 22 (at a bit rate of 1.25Gb/s). The outputs of the second switch in the pyramid are shown in Figure 23 (bit rate of 100Mb/s) and in Figure 24 (bit rate of 1.25Gb/s). From these results, it was found that the gain of the first 3-to-1 switches was 0.4dB, and the gain of the second switch was about 0dB. This is as desired since the switches are only designed to provide routing and selecting capabilities, not gain.

Crosstalk was simulated by having a small voltage swing at one of the inputs to the switch, and large voltage swings at the other two inputs. The small input swing was 30mV, which corresponds to an optical input of –23dBm, and the large voltage swings were 400mV, which is larger than any expected voltage swing. The control lines were then set to select the signal with the small amplitude. The output of the switch was monitored to see how much effect the large voltage swings had on the output. The effect of crosstalk was at most 1mV on the output of the 3-to-1 switch, which was deemed acceptable.

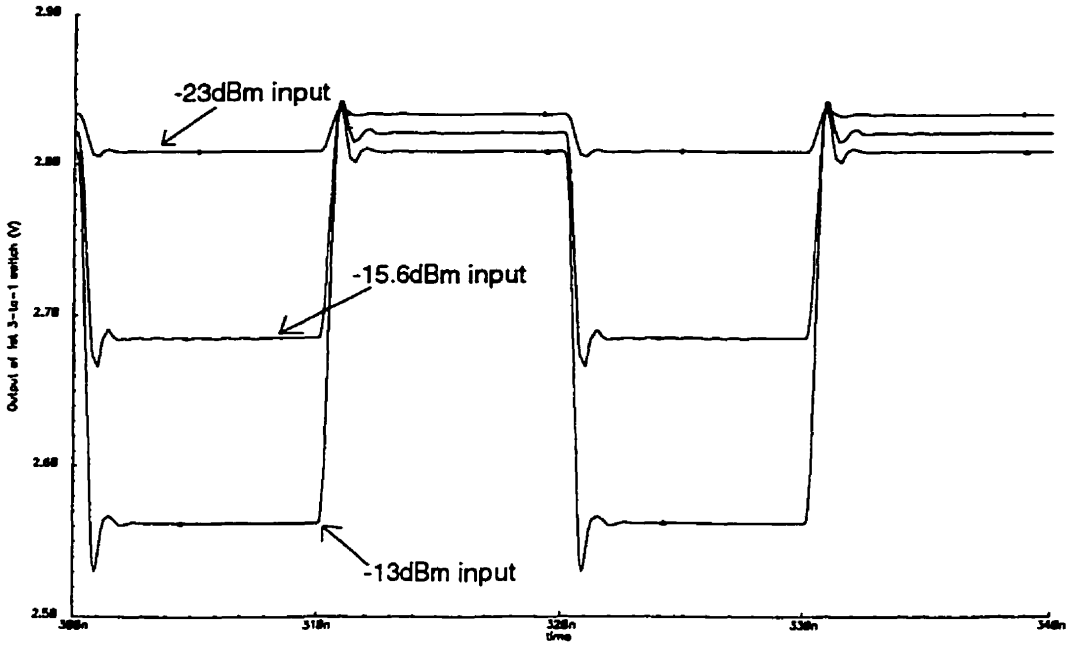

Figure 21. Output voltage swing from 1st 3-to-1 analog switch, at 100Mb/s

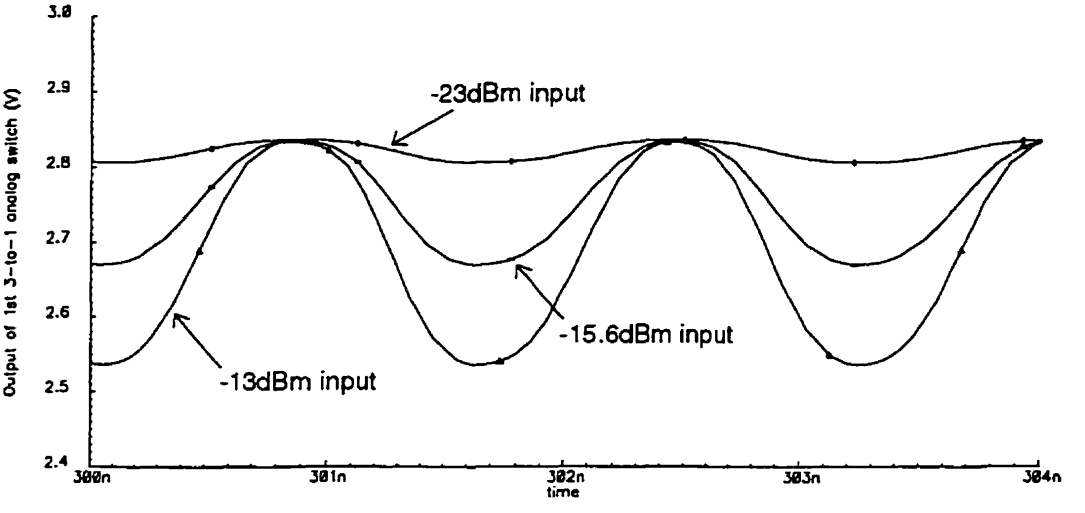

Figure 22. Output voltage swing from first 3-to-1 analog switch, at 1.25Gb/s

The following 2 figures shown the output of the pyramid of switches, after both layers of analog switches, at 100Mb/s and at 1.25Gb/s.

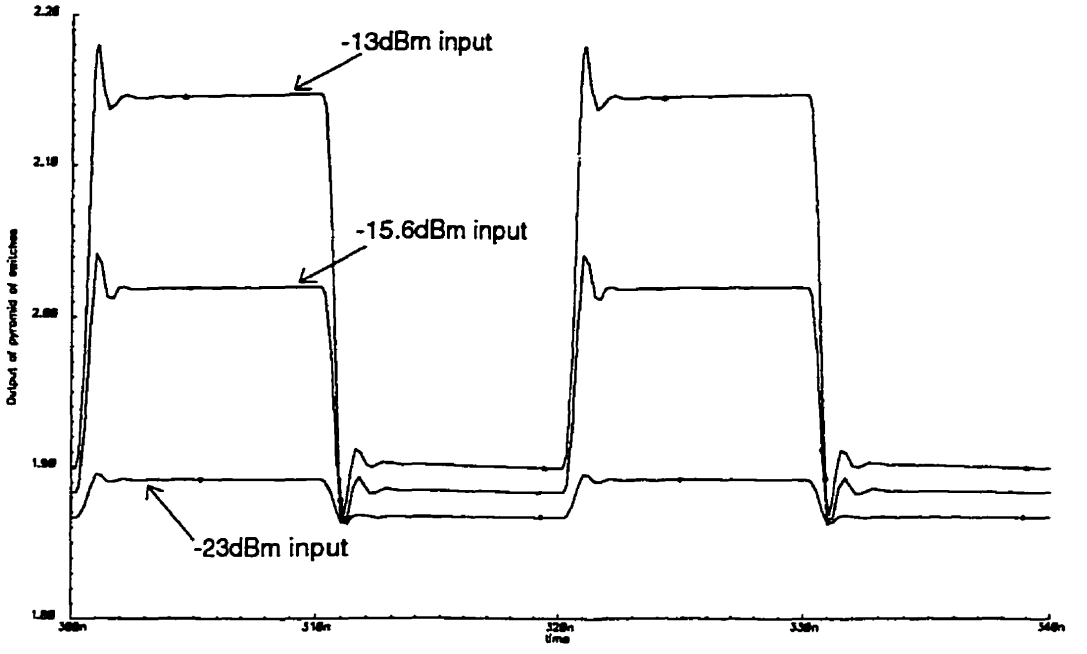

Figure 23. Output of pyramid of analog switches at 100Mb/s

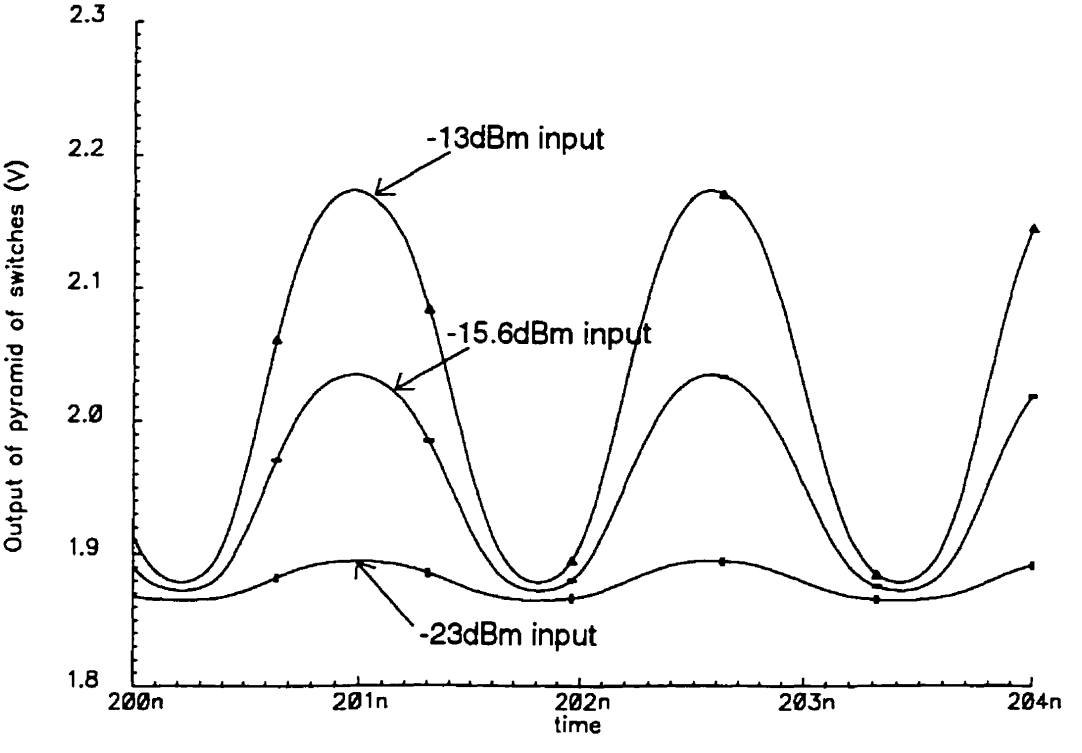

Figure 24. Output of pyramid of analog switches at 1.25Gb/s

4.5 Final amplification and digitization stage performance

The output of the pyramid of switches is fed into the last stages of the receiver. As described in section 3.2.4, the last stages consist of two cascaded differential amplifiers, followed by an output buffer. The input to the first differential pair is shown in Figure 25 for a bit rate of 1.25Gb/s.

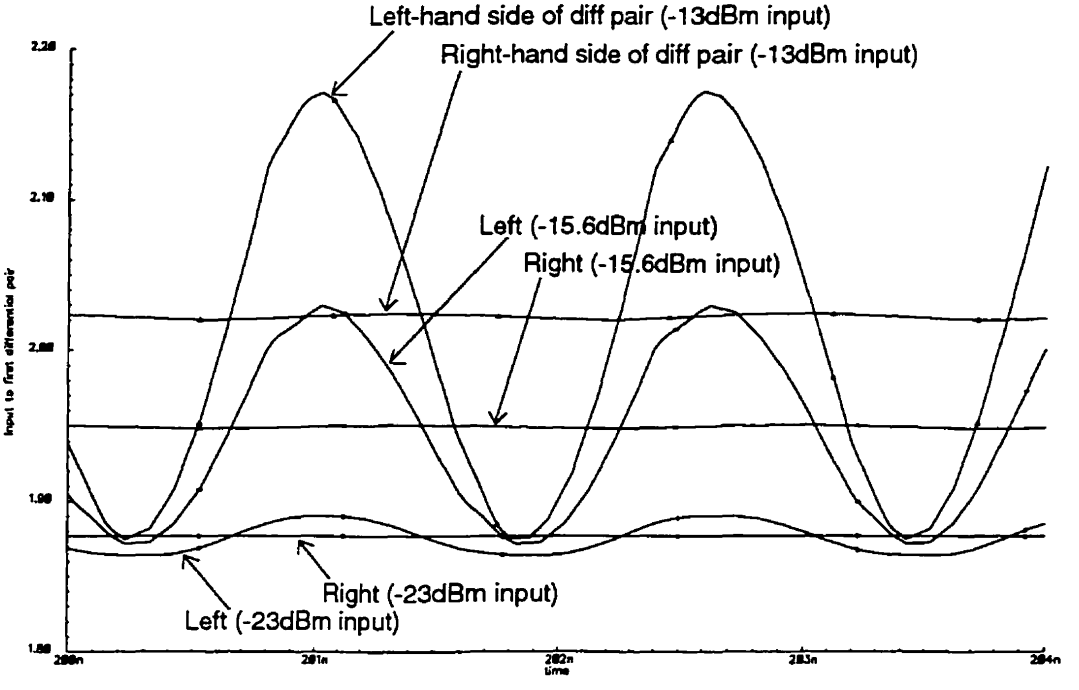

Figure 25. Input to first differential pair of last stages at 1.25Gb/s

The right-hand side of the differential pair is a lowpassed version of the left-hand signal. This lowpass resistor-capacitor combination can be seen in Figure 9. By using the input signal itself as a voltage reference, the circuit becomes independent of any DC bias on the input signal. However, the receiver no longer functions at lower frequencies. Figure 26 shows same inputs to the differential pair (for a −13dBm optical input) but at 100Mb/s. As can be seen in the figure,

the right-hand side signal is beginning to charge up the capacitor but is still functional as a DC reference for the differential pair. The AC response of the receiver is further analyzed in section 4.6.

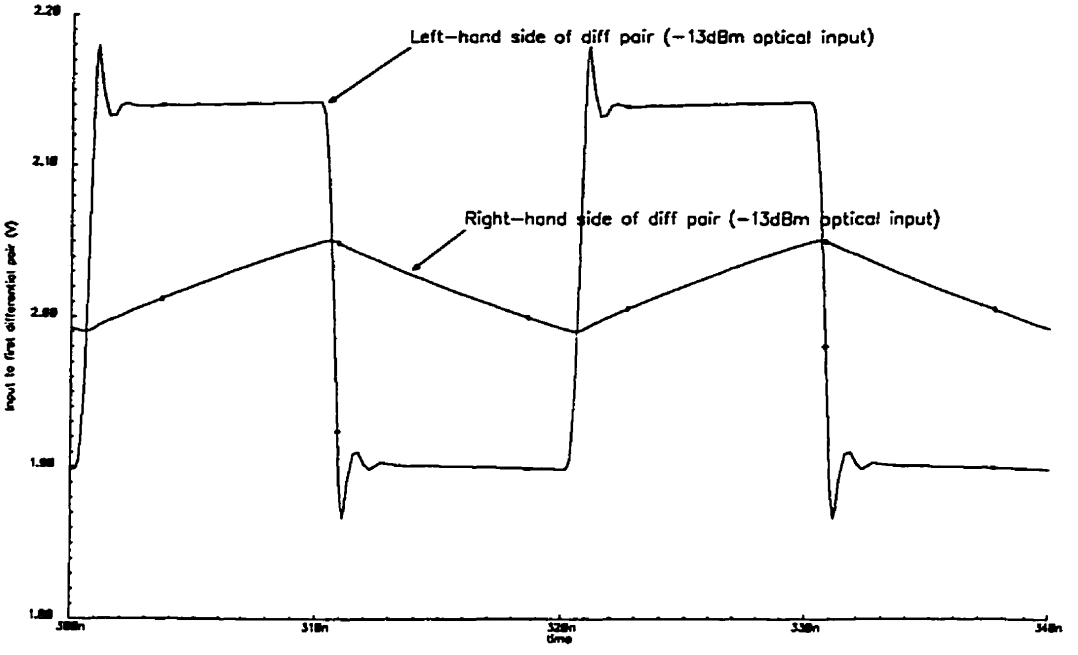

Figure 26. Input to first differential pair of last stages at 100Mb/s

Figure 27 and Figure 28 show the PECL output of these last stages, after the cascade of both differential pairs, and of the output buffer. In these figures, we can see that the smallest input power to the receiver (-23dBm, or 5μW) is just barely being amplified to PECL levels, whereas the other two input power levels produce virtually indistinguishable signals. This highlights the tradeoff in the initial choice of feedback resistance for the TIA. The resistance was made as small as possible to increase the bandwidth of the receiver, but is just big enough to digitize the output signals over the entire dynamic range of the receiver.

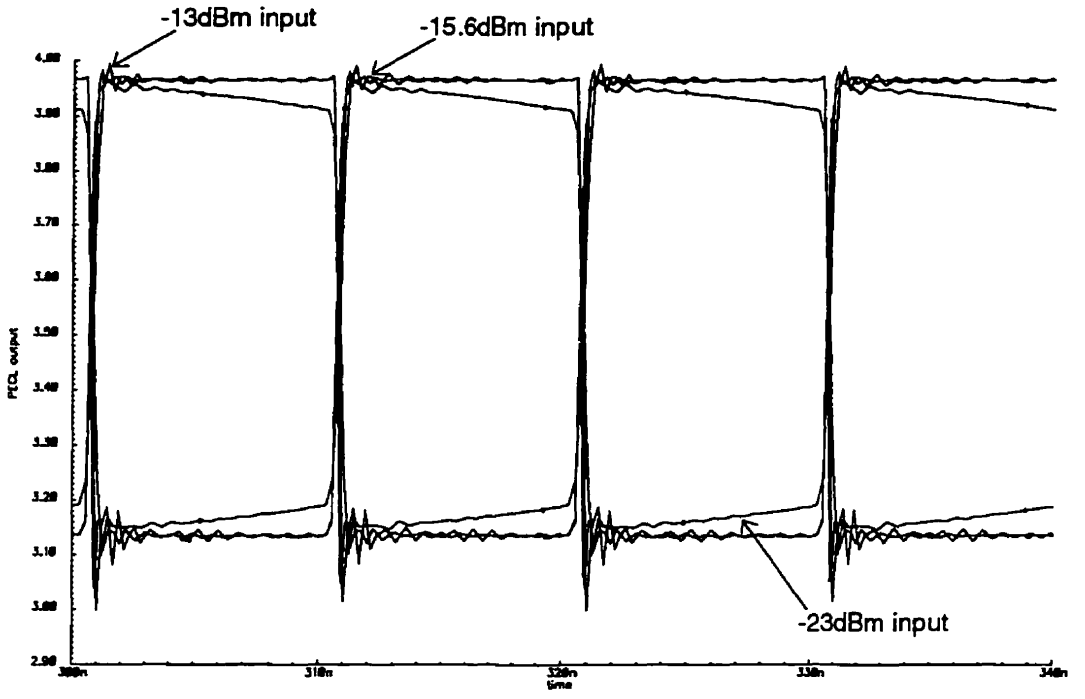

Figure 27. PECL output of receiver at 100Mb/s

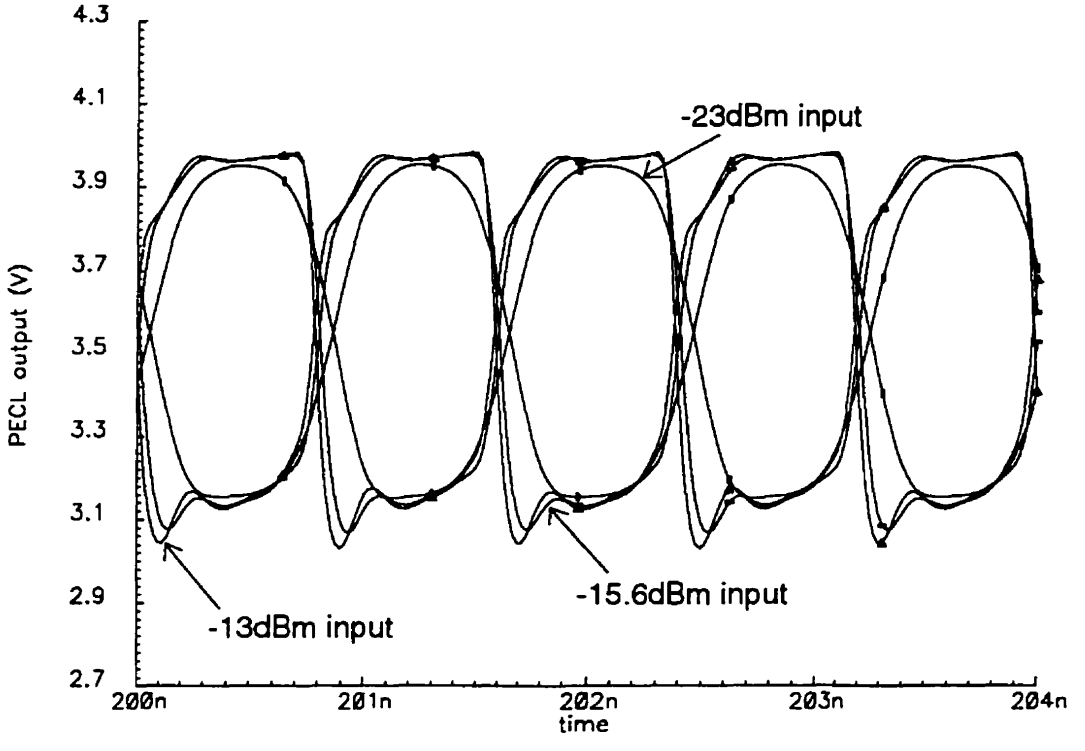

Figure 28. PECL output of receiver at 1.25Gb/s

4.6 Bandwidth performance

An AC response simulation of the chip was performed in order to monitor the performance of the different stages of the receiver at different frequencies. Figure 29 shows the results of the simulation for the first stages of the receiver. The –3dB point of the combined stages is at 1.3GHz.

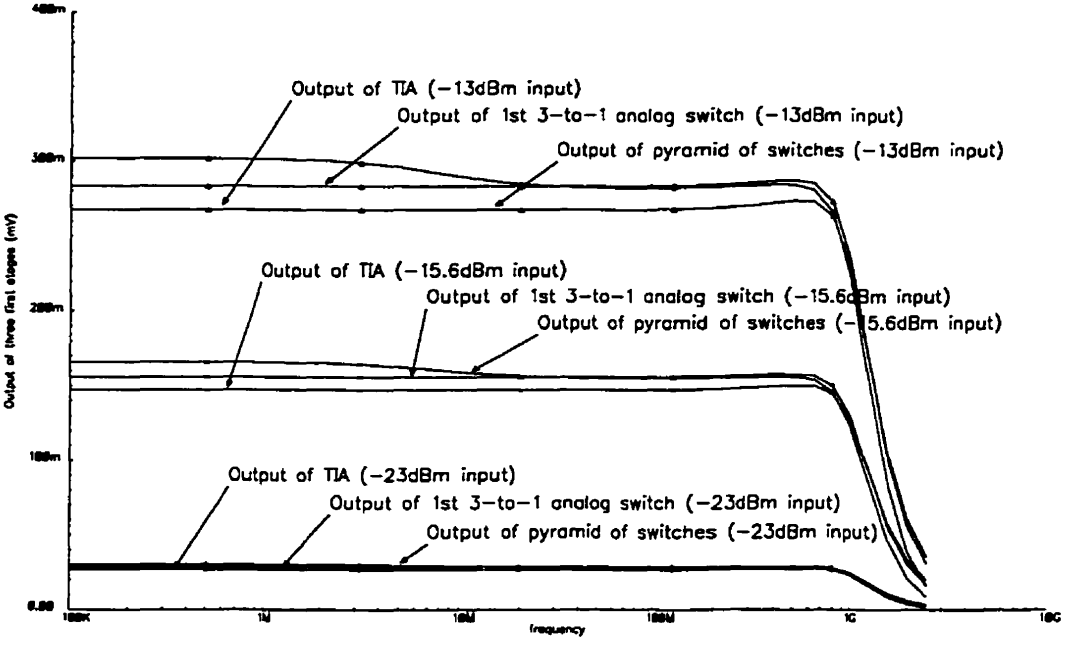

Figure 29. AC response of different stages of the receiver

The last stages of the receiver do not operate under small signal conditions, and therefore the AC response simulation is not entirely valid since it assumes all transistors are operating in the linear mode. Nevertheless, the results are included in Figure 30. The two curves in the graph represent the output of the two differential pairs. The voltage swing at the output (the y-axis) is meaningless since the transistors saturate. However, the tendency to operate within a range of frequencies is clearly seen. The lower frequency –3dB point is at 4.4MHz. This

tendency is confirmed by the experimental results presented in Chapter 7. Increasing the value of the capacitor in the lowpass filter before the first differential pair would have lowered the cutoff frequency even further. However, 4.4MHz was sufficient since the required minimal bit rate was 100Mb/s, and the capacitor on-chip was already 145μm x 145μm, for a capacitance of 20pF. This is discussed further in Chapter 7.

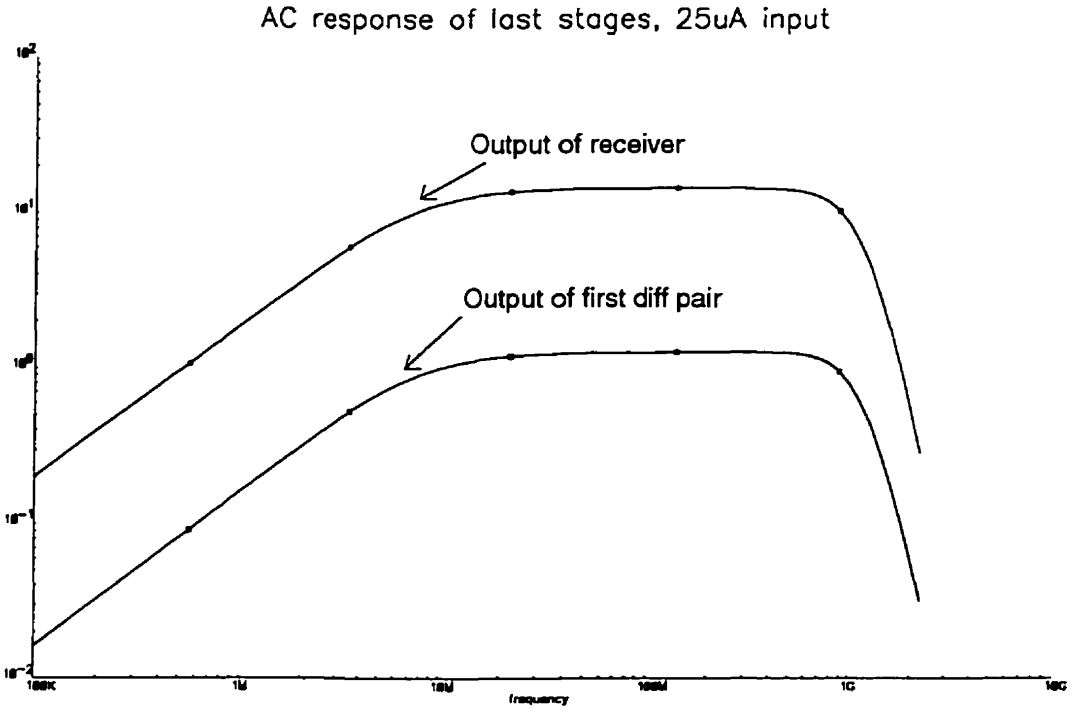

Figure 30. AC response of last stages

4.7 Conclusion

Every stage of the receiver design was simulated, as well as the entire receiver. These simulations showed that the receiver can handle optical inputs ranging from –23dBm to –13dBm, with a bandwidth from 4.4Mb/s to 1.25Gb/s, at

an input contrast ratio of 10:1 dB. The experimental results are discussed in Chapter 7 and are compared to these simulation results.

4.8 References

1. Opto Speed SA, http://www.optospeed.com
2. H. Johnson, M. Graham, "High-speed digital design: a handbook of black magic", Prentice Hall, 1993.
3. E. Bisaillon, et al., "Optical link for an adaptive redundant free space interconnect", in *Optics in Computing, 2000*, Roger A. Lessard, Tigran Galstian, Editors, SPIE Vol. 4089, pp. 999-1009, 2000.
4. E. Bernier, et al., "1Gb/s transmitter for an adaptive redundant free-space interconnect", in *Optics in Computing, 2000*, Roger A. Lessard, Tigran Galstian, Editors, SPIE Vol. 4089, pp. 772-783, 2000.
5. J. M. Rabaey, "Digital integrated circuits, a design perspective", Prentice Hall, 1996.
6. G. P. Agrawal, "Fiber-optic communication systems", Wiley, 1997.
7. Robert J. Hoss, "Fiber Optic Communications, Design Handbook", Prentice Hall, 1990.
8. P. Gray, R. Meyer, "Analysis and design of analog integrated circuits", 3rd edition, John Wiley & Sons, Inc., 1993.
9. A. Sedra, K. Smith, "Microelectronic circuits", Harcourt Brace College Publishers, 1991.
10. B. Streetman, S. Banerjee, "Solid State Electronic Devices", 5th edition, Prentice Hall, 2000.
11. G. Neudeck, "The bipolar junction transistor", Vol.3 , 2nd edition, Addison-Wesley, 1989.

Chapter 5. Receiver module packaging

5.1 Introduction

When packaging the receiver, several factors need to be taken into consideration. Mainly, the signal integrity of the high-speed signals needs to be maintained, and the optical signal must reach the photodetector plane with the least amount of loss or distortion [1]. The packaging scheme chosen allows for a passive alignment of the components.

5.2 RX module

There are two identical receiver modules per printed circuit board (RX1 and RX2), as shown in Figure 2. Each receiver module receives a 960nm optical data signal and outputs a 1Gb/s differential PECL signal. The receiver module consists of a ceramic carrier, a receiver die, a PIN photodiode array die, and a bulk lens held in place by an optomechanical piece made out of aluminum. These components are shown in Figure 32.

The photodiode array consists of a 5×5 array of InP PIN photodetectors, of which we are using a 3×3 subset as will be shown later in Figure 36. The photodiodes have a room temperature responsivity of 0.6A/W at 960nm, at –2.5V bias [2]. Their rise and fall times (10%-90%) are 150ps, which makes them well suited for use in a 1Gb/s receiver. The receiver chip and photodiode array chip are wirebonded together inside the 100-pin Cerquad flat-pack chip carrier, as shown in Figure 31. The 3dB bandwidth of the unloaded ceramic carrier was

experimentally measured to be 2.1GHz [3]. The optomechanical piece is glued to the corners of this carrier, holding the lens for the photodetectors at a distance of 12mm from the PD die surface. The final module can be seen in Figure 33.

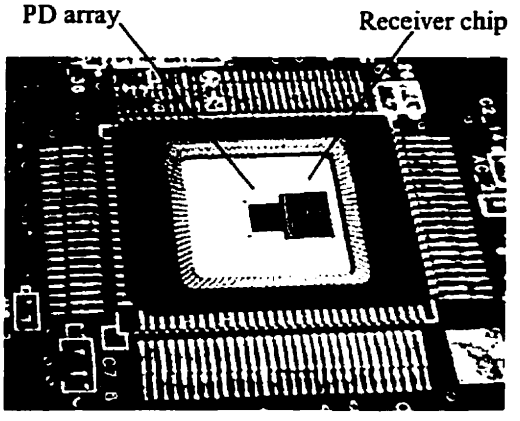

Figure 31. Carrier without optomechanics

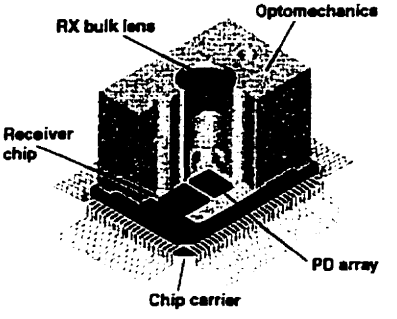

Figure 32. Schematic of receiver module

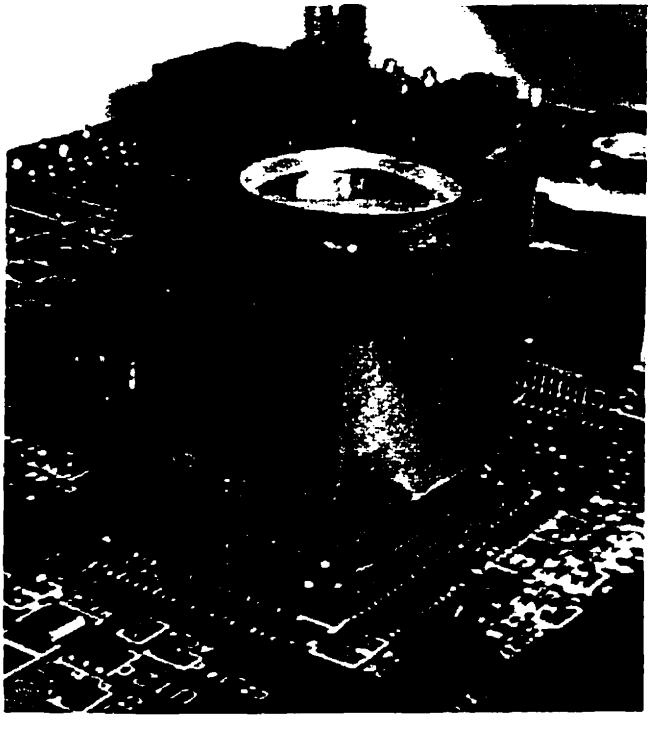

Figure 33. Photo of completed receiver module

The placement of the die containing the array of InP photodetectors is important so that the packaging of the die doesn't consume too much of the misalignment tolerance budget [1]. Ideally, the optical center of the 3×3 array of photodetectors is at the center of the carrier, which centers it on the optical axis. Alignment markers were laser drilled into the cavity of the carrier with a precision of ±40μm with respect to the datum point of the carrier. There is one marker for each corner of the photodetector chip, as shown in Figure 34.

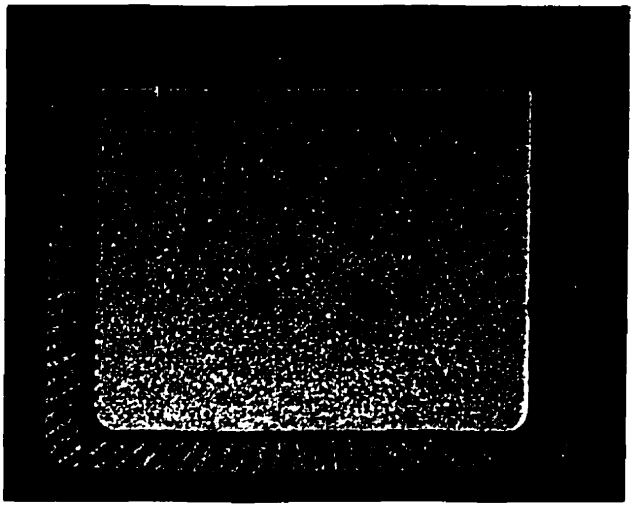

Figure 34. Photo of alignment markers in carrier

The die was then placed by hand in the carrier with the aid of the alignment markers, for a total precision of ±50μm with respect to the datum point, and glued into place with conductive epoxy. The conductive epoxy used was chosen not only for its electrical and thermal conductivity, but also for its even bond line. The epoxy contains several metal balls of 25μm in diameter, with a density of approximately a dozen balls under the photodetector chip that is 2.2mm×2.2mm in surface area. These balls guarantee that the photodiode chip

will be parallel to the ceramic carrier to within ± 0.07°. The lateral and angular placement tolerances are well within the required tolerances found through simulations of the optical link [1].

The ceramic carrier is a gull-wing surface mount component that was placed on the printed circuit board through a reflow process. This gave a placement precision of about ±45μm, which consumes a minimal amount of the lateral misalignment budget.

The optomechanical piece was attached to a 6-degree of freedom computer controlled motion stage, and then aligned to the carrier [1]. Once in place, it was glued to the corners of the ceramic carrier. The focal length of the lens, which is held in place by the optomechanical piece, is 15mm. Since the distance between the lens and the photodetectors is approximately 12mm, the incident beam of light will be defocused at the photodetector plane, and will have a diameter of 135μm. The photodetectors have an active area of 70μm in diameter. Since the photodetectors have a center-to-center pitch of 125μm, the light will never fall into a "blind" spot, where none of the photodetectors could detect the light. This concept is illustrated in Figure 35. This means that within the specified misalignment tolerances, there will always be at least one VCSEL-PD pair with sufficient optical power transfer to transmit data.

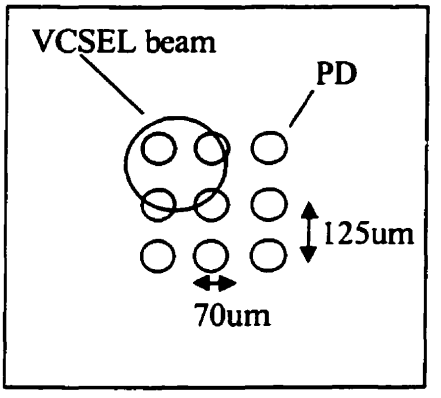

Figure 35. PD-die with light beam from VCSEL landing on photodetector plane

5.3 Photodetector layout

The photodetector die has a 5x5 array of photodiodes, of which a 3x3 subset is being used for this system. Because each photodetector must have both its anode and its cathode wirebonded to the receiver chip, as shown in Figure 6, there are 18 bonds needed in between the two dies. All 18 wirebond pads do not fit on one side of the photodetector chip. So the photodetector layout was picked to be as shown in Figure 36. The pads on the top edge of the die are the ones wirebonded to the receiver chip and they are labeled 1, 2, 35-50.

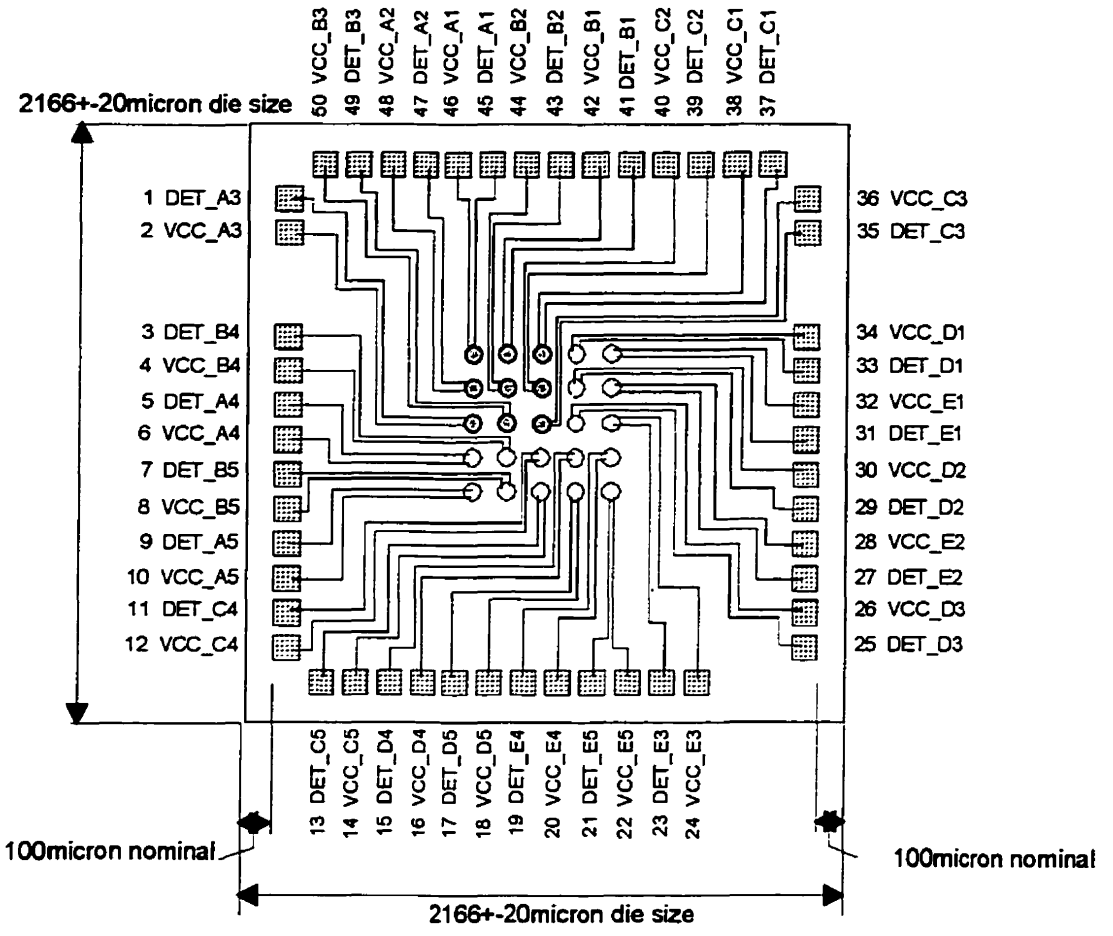

Figure 36. Photodetector die layout

5.4 Conclusion

A receiver module for use in a free space optical link was designed and built. This module is soldered onto a printed circuit board, which is discussed in the next chapter. The module contains the photodiode array, receiver chip, optomechanical piece, and bulk lens. The assembly of all the components was carefully performed in order to consume a minimal amount of the misalignment tolerance budget.

5.5 References

1. E. Bisaillon, et al., "Optical link for an adaptive redundant free space interconnect", in *Optics in Computing, 2000*, Roger A. Lessard, Tigran Galstian, Editors, SPIE Vol. 4089, pp. 999-1009, 2000.

2. Opto Speed SA, http://www.optospeed.com

3. E. Bernier, et al., "1Gb/s transmitter for an adaptive redundant free-space interconnect", in *Optics in Computing, 2000*, Roger A. Lessard, Tigran Galstian, Editors, SPIE Vol. 4089, pp. 772-783, 2000.

4. D. V. Plant, et al., "A 5 Gb/s, 2 channel bi-directional adaptive redundant FSOI demonstrator system", in *Optics in Computing, 2000*, Roger A. Lessard, Tigran Galstian, Editors, SPIE Vol. 4089, pp. 465-472, 2000.

Chapter 6. Printed circuit board

6.1 Introduction

The printed circuit board for the system holds two transmitter modules and two receiver modules for the two bi-directional links, as shown in Figure 2. In addition to the four transceiver modules, it has the biasing circuitry for the transmitter modules, and the power monitoring circuitry for the receiver modules which is described in section 6.3. The algorithm to select the VCSEL-PD pair for each optical link is implemented in an FPGA that handles all the control lines for the transceiver modules. This hunt algorithm is described in Appendix A. A block diagram for the board electronics is shown in Figure 37.

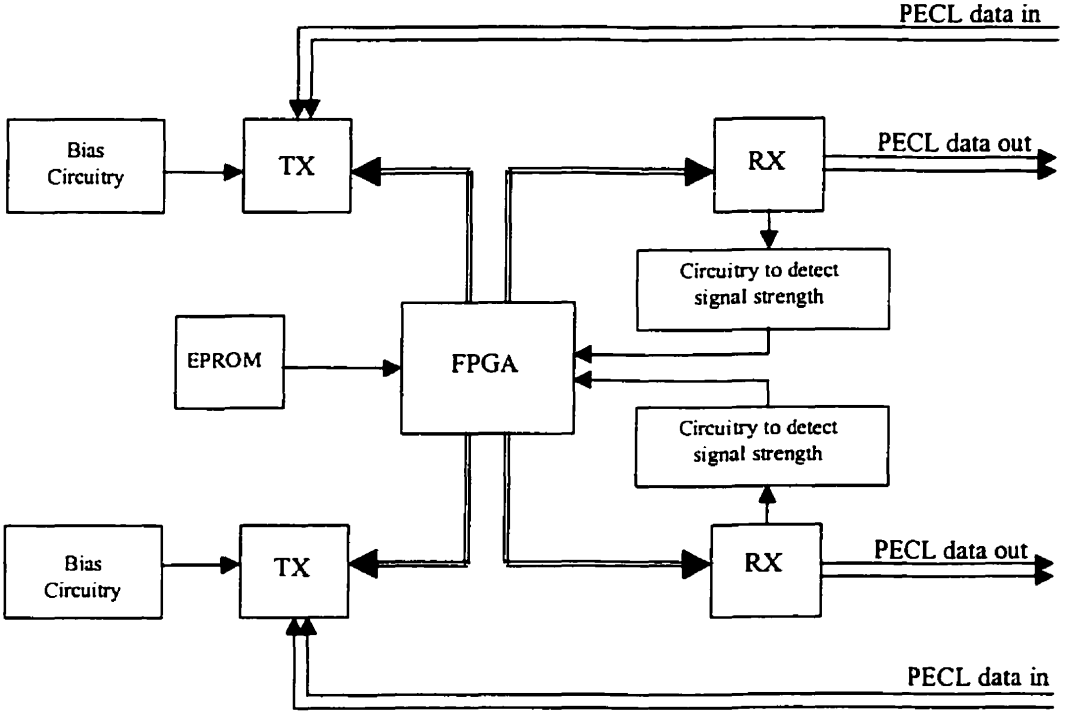

Figure 37. Board block diagram

6.2 Board description

The circuit board is a 16cm x 16cm eight-layer board with a stack-up as shown in Figure 38. To reduce switching noise on the analog lines, the digital and analog ground planes are kept separate. For similar reasons, the power plane for the transmitter dies is kept separate from the power plane supplying the receiver dies. The power plane for the rest of the digital circuitry on the board is also separate. The pairs of differential high-speed lines, used for the PECL output signals from the receiver modules, were designed to have an impedance of 50Ω each and a differential impedance of 100Ω.

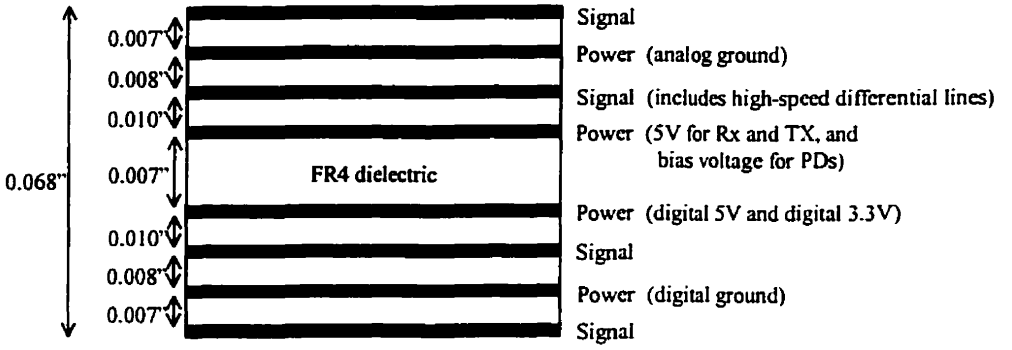

Figure 38. Board stack-up

For our system it is important to minimize board warpage in order to not consume too much of the tilt tolerance in the misalignment budget. If the board is too curved or warped, it may be impossible establish 4 links in between two boards. Two methods were used to eliminate this problem. First, copper fills were used on all of the signal layers so that every layer would have roughly the same amount of copper. Secondly, to reduce the impact of any remaining

warpage, the ceramic carriers were placed very close to each other in the center of the board. A photo of the board is shown in Figure 39.

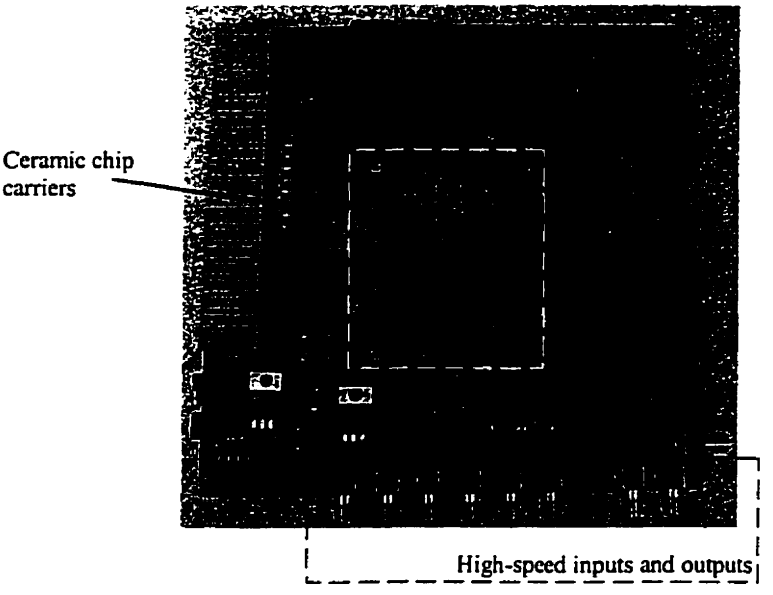

Figure 39. Printed circuit board before placement of dies

6.3 Power monitoring circuitry

In order to obtain the lowest bit error rate for a particular VCSEL-PD pair, the one with the highest optical power transfer must be selected. To determine which photodetector is receiving the strongest signal, the analog single-ended output signal from the pyramid of analog switches is sent off-chip for power monitoring purposes. This was shown schematically in Figure 4.

The first element the signal encounters when it leaves the chip carrier is a buffer. This buffer has a bandwidth of 1GHz, and is placed adjacent to the carrier to minimize loading on the receiver circuit. The buffered signal is then sent to an RMS to DC chip. This chip converts the RMS voltage of the input signal into a DC output voltage. The DC level produced by this chip is sent to an analog to

digital converter chip that has 8-bit precision. The output of the A to D is fed directly to an FPGA. With the current choice of RMS to DC chip and A to D converter, we have a resolution of 13mV on the amplitude of the single-ended analog signal coming out of the switch pyramid.

A hunt algorithm implemented on the FPGA controls the selection process, and is described further in Appendix A. The six select lines (S1 to S6) for the pyramid of analog switches are tied to output pins of the FPGA. To determine which photodetector is receiving the highest amount of optical power, the hunt algorithm is able to control the select lines and scan through the 9 channels on the receiver.

Once the channel with the highest received power has been selected, the optical power received by that channel is continuously monitored by the FPGA. If the power level drops beneath a certain threshold, to be determined experimentally, the hunt algorithm immediately stops the data transfer so that no bits are lost. Then the algorithm finds the new best channel. The clocking speed of the FPGA is 20MHz, so the 50MHz switching speed of the pyramid of analog switches does not limit the algorithm.

6.4 Electrical characterization of circuit board

The high-speed lines on the printed circuit board were characterized to ensure that they could transport digital PECL signals up to 1.25Gb/s. Figure 40 shows the time domain reflectrometry (TDR) of one of the high-speed trace on the PCB, which was 7.7 cm long.

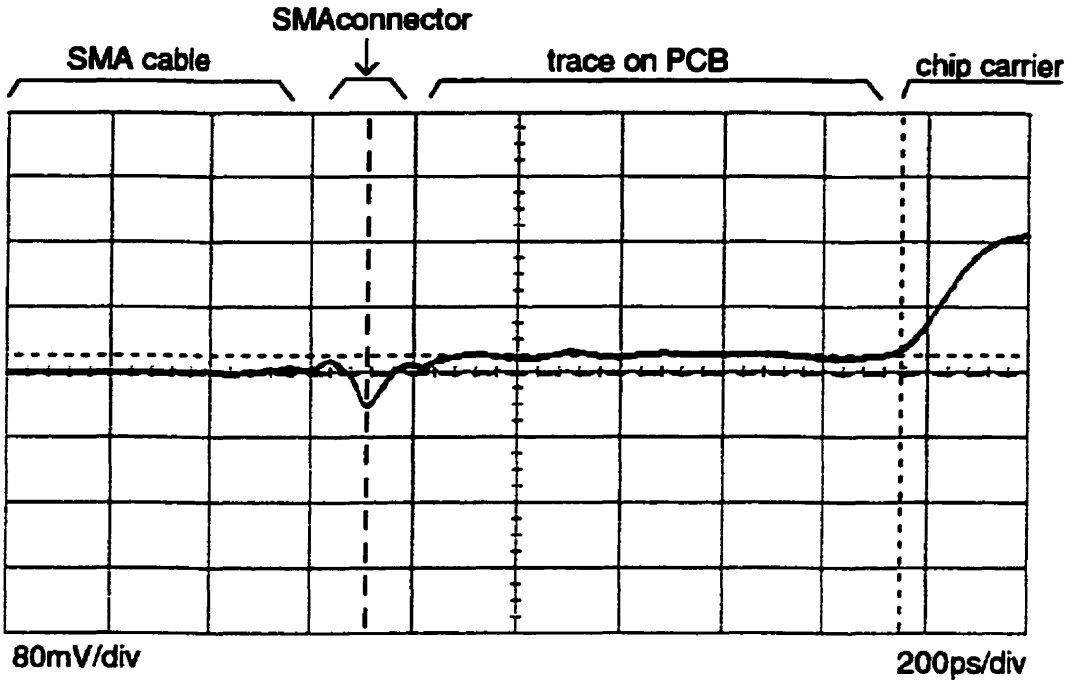

Figure 40. TDR measurement of board

As can be seen in the figure there is a discontinuity when the signal leaves the 50Ω SMA cable and arrives on the circuit board. The amplitude of the wave sent down the line was 200mV, and the amplitude of the reflected wave was 21 ± 1mV. Using the formula for the reflection coefficient of a wave:

$$\rho = \frac{Z_l - Z_o}{Z_l + Z_0} \quad [3]$$

where Z_o is the impedance of the SMA cable and Z_l is the impedance of the trace on the board, the impedance of the high-speed trace was found to be 62 ± 3Ω. This deviation from the designed impedance is to be expected from the manufacturing process. To confirm that this impedance value is acceptable for the high-speed signals on the board, a loop-back test was performed. This test involved sending a high-speed signal down one of the high-speed traces,

connecting that trace through a wirebond to a neighboring trace and returning the signal through the second trace. The eye-diagram generated by this test at 1Gb/s is shown in Figure 41. Since the eye is open and clean, the packaging has a minimal impact on the performance of the receiver.

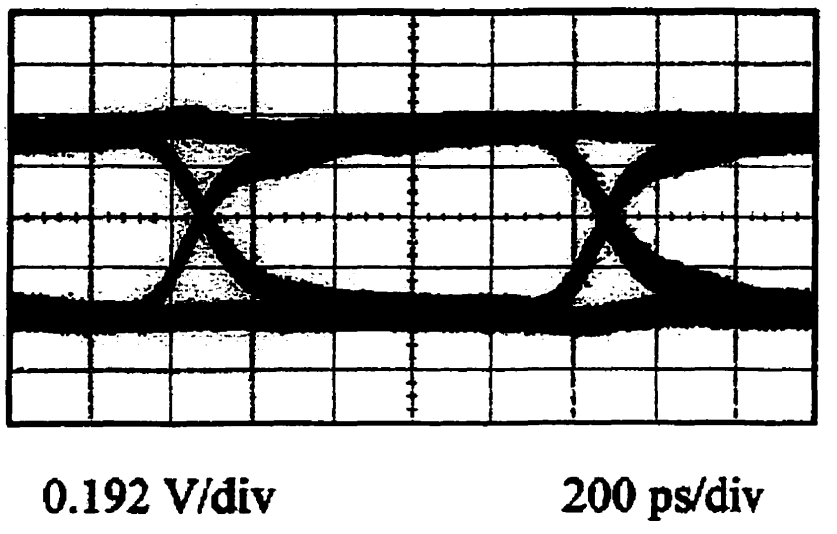

Figure 41. 1Gb/s loop-back test on circuit board

6.5 Conclusion

A printed circuit board was designed and manufactured to hold all four transceiver modules and the associated control circuitry. The high-speed traces on the board were successfully designed to not limit the performance of the transceivers.

6.6 References

1. H. Johnson, M. Graham, "High-speed digital design: a handbook of black magic", Prentice Hall, 1993.

2. Mentor Graphics Corporation, http://www.mentorg.com/

3. D. Griffiths, "Introduction to electrodynamics", Prentice Hall, 1989.

4. S. D. Corcy, A. Yang, "Interconnect characterization using time-domain reflectrometry", *IEEE transactions on microwave theory and techniques*, Vol. 43, No. 9, pp. 2151-2156, 1995.

5. K. K. Mihan, et al., "Interdisciplinary design optimization for high-speed packages and interconnects", *Canadian Conference on Electrical and Computer Engineering*, Vol. 1, pp. 389-395, 1995.

Chapter 7. Experimental results

7.1 Introduction

The receiver chips were packaged and tested to compare the performance to the simulated results. A special test rig was built in order to control the amount of power arriving in the active area of the photodetectors. Several tests were performed, including tests for uniformity, bandwidth, sensitivity, dynamic range, and bit error rate.

7.2 Test rig

For the receiver tests, the amount of power arriving in the active area of the photodiode must be known and controllable. However, to get large misalignment tolerances, the optics for the system were designed to have an unfocused spot arriving at the photodetector plane. This was described in more detail in Chapter 2. Therefore the receiver test were not performed with the optics designed for the system [1]. Instead, before the optomechanical pieces were glued onto the ceramic carriers, a 4-f telecentric relay was used in between the two printed circuit boards. A schematic of the test rig is shown in Figure 42.

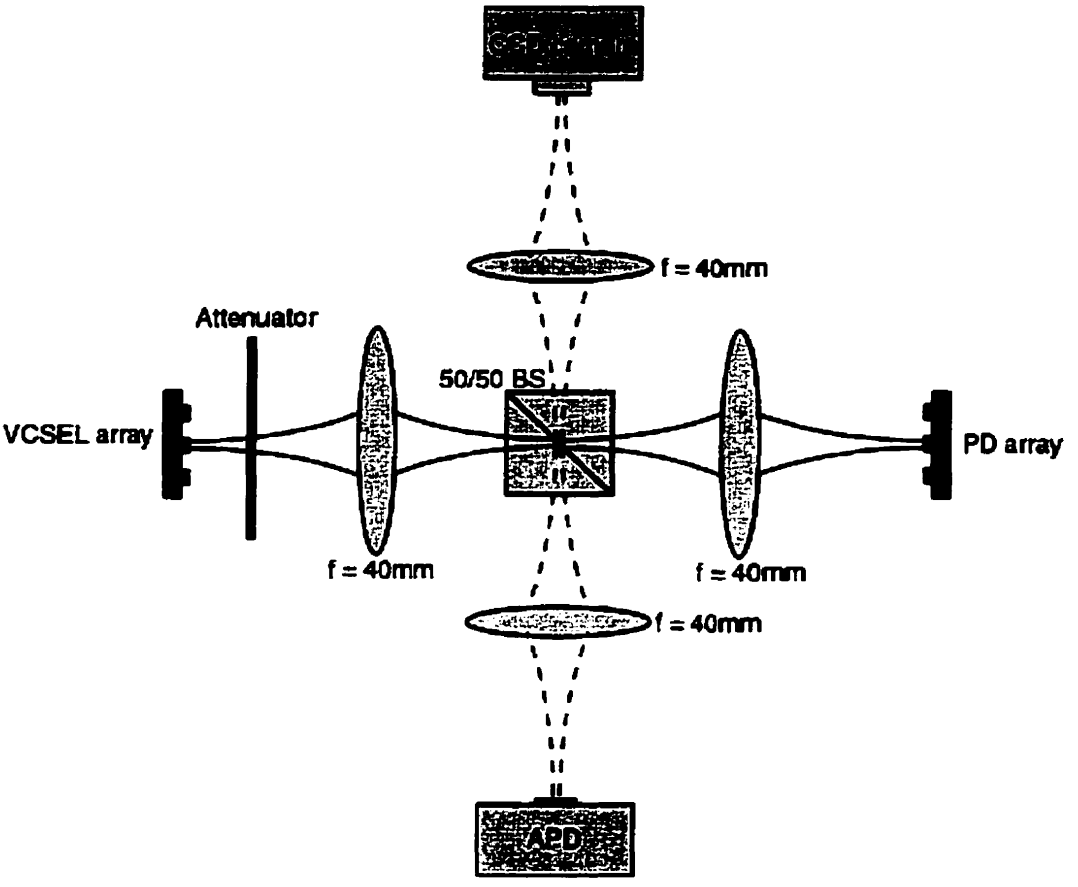

Figure 42. Test rig used to test receiver chip

Using a 50/50 beam splitter placed in between the two lenses, the photodetector plane could be viewed using a CCD camera, and the output of the VCSELs could be monitored using an avalanche photodiode (APD). The light from the VCSEL was focused on the photodiode plane to a spot size under 20μm in diameter. The profile of the spot used for testing is shown in Figure 43. The photodiode was therefore underfilled since the active area of the photodetectors is 70μm. It could then be assumed that all of the power arriving at the photodetector plane was entering the photodiode. Also, a variable attenuator was placed at the

output of the VCSEL and before the first lens, to be able to vary the amount of power entering the photodetector.

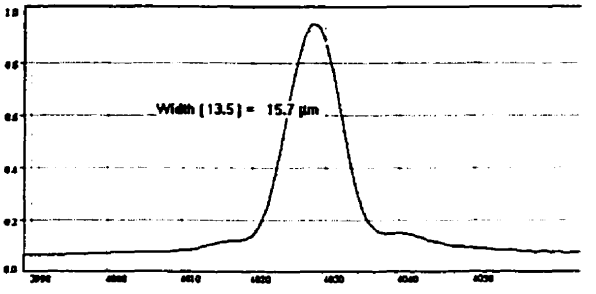

Figure 43. Profile of VCSEL spot on photodetector plane

7.3 Bandwidth tests

The first tests performed were to determine the bandwidth of the receiver. For these tests, the average power to the photodetector was 250±5µW (-6dBm), which was experimentally found to be well above the minimum sensitivity of the receiver, but before any overloading effects could be detected. The outputs shown in Figure 44 are one of the two PECL outputs from the receiver. The VCSEL was modulated with a periodic signal and the illuminated photodetector was photodetector number 6.

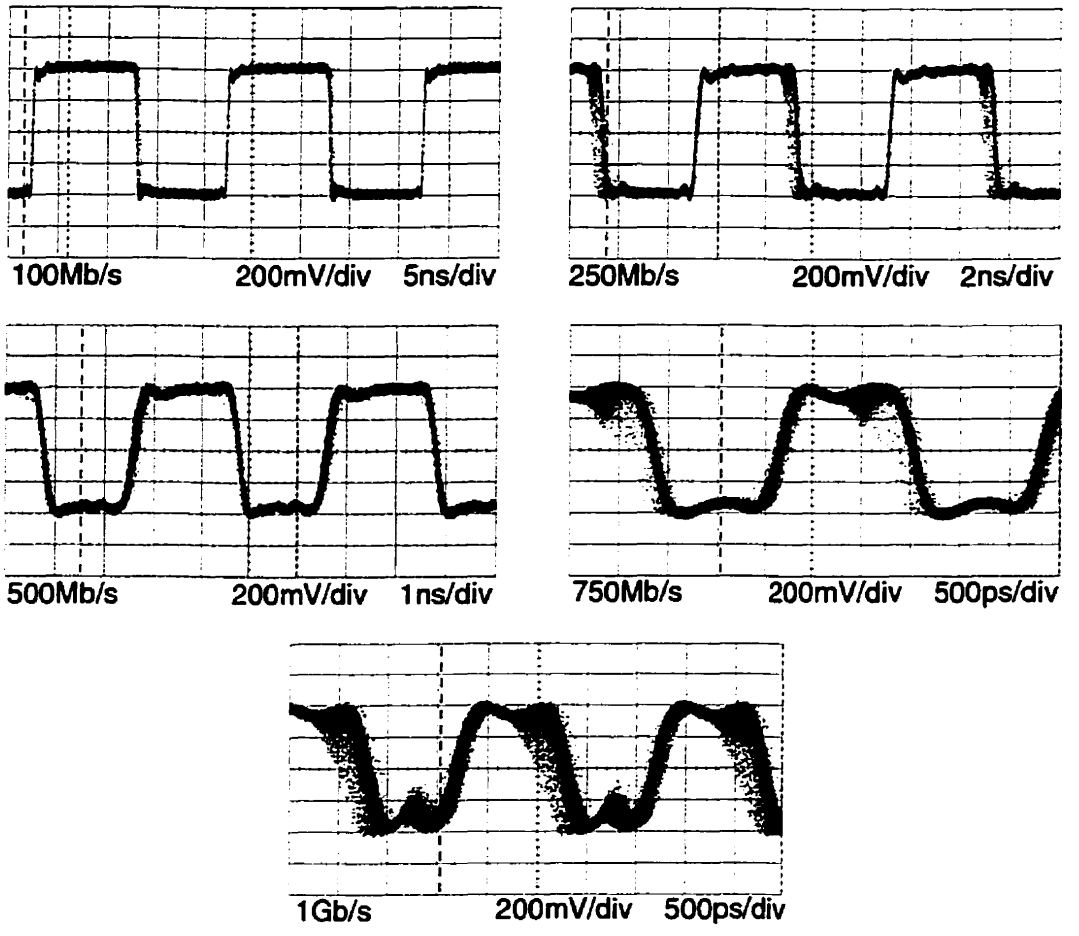

Figure 44. Periodic output from the receiver

In viewing the waveforms in Figure 44, we can see that the performance at 100Mb/s is very close to the simulation results presented in Chapter 4. And as expected, the performance deteriorates as the bit rate increases, but the amplitude of the output signal is still a full 0.8V. However, there is a considerable amount of timing jitter at 1Gb/s, especially on the falling edge of the signal. This is probably due to the last stage of the receiver, where the transistors of the differential amplifiers enter saturation in order to digitize the output of the receiver. On the falling edge of the signal, the transistor is leaving the saturation

region to enter the active region. This may take a longer amount of time than entering the saturation region on the rising edge of the signal.

Eye diagrams were also generated to include more frequency components in the signal and get a more accurate analysis of the bandwidth. For these tests, the VCSEL was modulated using a pseudo-random bit sequence (2^{17}-1). The eye diagrams in Figure 45 show the differential outputs of one channel of the receiver at different bit rates. The two traces at the top of every image are the differential PECL outputs shown at 800mV/div. The trace at the bottom on every image is the measured eye diagram of the VCSEL being used to drive the receiver, shown at 80mV/div. The transmitter data was captured using the APD.

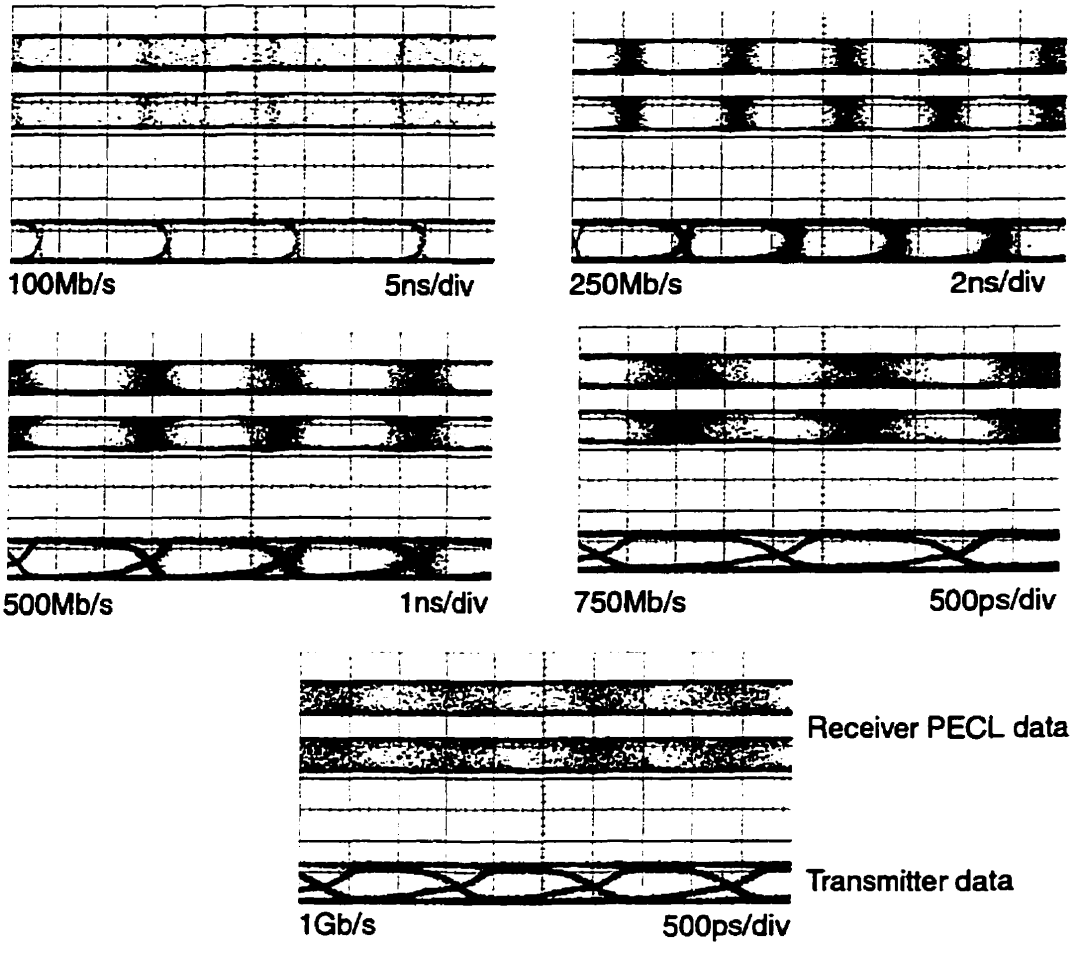

Figure 45. Eye diagrams of output of receiver at different bit rates

From these eye diagrams, it can be seen that the receiver operates with a reasonable eye up to approximately 750Mb/s. Above that bit rate, the signal became unacceptably noisy, where most of the noise seems due to timing jitter. At data rates below 200Mb/s, we found the signal became unacceptably noisy as well. This is to be expected, because the pseudo-random data at 200Mb/s includes many lower frequency components that may go well below 100Mb/s, and 100Mb/s was the low end of the receiver by design. The lowest bit rate that produced acceptable output was at 18Mb/s and is shown in Figure 46. This is

very comparable to the simulation results presented in Chapter 4. Using an encoding scheme such as 8B/10B encoding, the bit rate of the data does not go below 100Mb/s, and therefore the lower frequency cutoff is not a limitation.

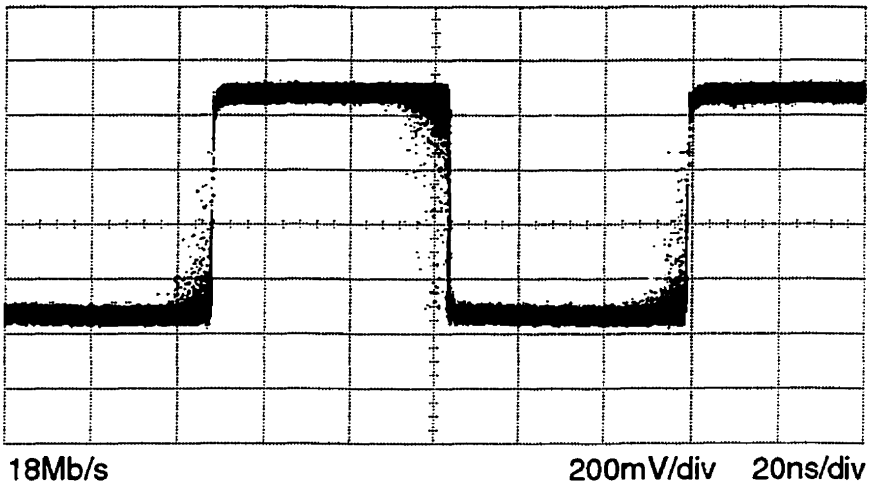

Figure 46. Periodic output from the receiver at 18Mb/s

7.4 Uniformity test

A uniformity test was performed to compare the performance of each channel of the receiver. The eye-diagrams shown in Figure 47 are for a bit rate of 500Mb/s. The top signal of each graph is the input to the receiver viewed with the APD. The bottom two signals are the two PECL outputs from the receiver. The second signal in the graphs is the difference of the two outputs.

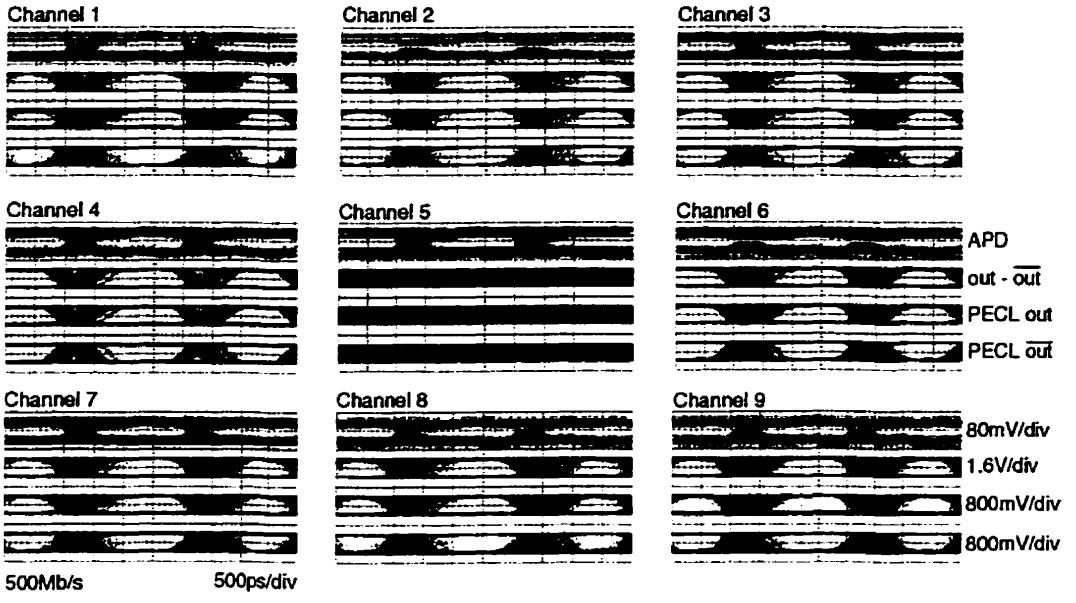

Figure 47. Comparison of the performance of the 9 channels of the receiver

As can be seen in Figure 47, channel 5 does not work, but all the other channels are virtually identical, based on a visual inspection of the measurements. When other receiver dies were tested, other channels sometimes did not work, but channel 5 did. Therefore, the dead channels are not intrinsic to the chip, but are due to manufacturing defects.

7.5 Sensitivity test

A sensitivity test was performed in order to determine the minimum amount of input optical power required for the receiver to operate with an open eye. To perform this test, the receiver was sent pseudo-random data at a certain bit rate, and the variable attenuator (shown in Figure 42) was used to slowly decrease the amount of power arriving at the photodetector. Just before the onset of eye closure, the average power arriving at the photodetector was measured using a

power meter. The lower bound on the sensitivity of the receiver is shown in Figure 48.

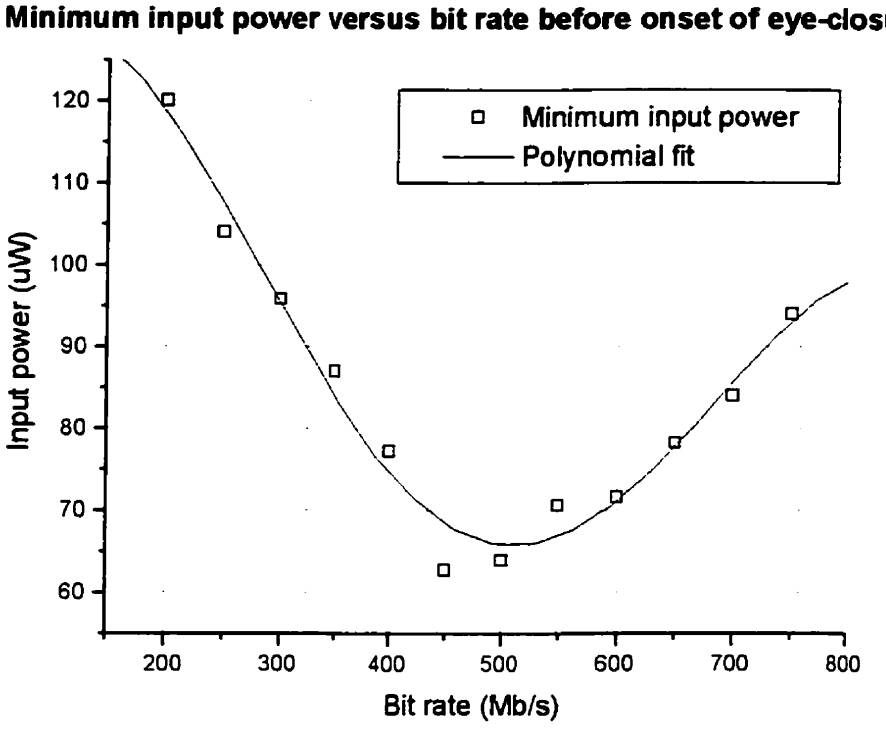

Figure 48. Receiver sensitivity

Referring to Figure 48, the optimal bit rate for the receiver is at 500Mb/s. Because this bit rate requires the least amount of input optical power, it is also the bit rate at that will have the highest tolerance to misalignment. The required power at this bit rate, 65μW (-11.9dBm), is higher than was required in the simulations where the minimum input power was 5μW (-23dBm). However, when tested in the final system, even with a higher required input power the lateral and angular misalignment tolerances between boards were still ±1mm and ±1 degree [1], for a board separation ranging from 5cm to 22cm.

The upper bound for the input power to the receiver could not be determined. In the test rig, the beam splitter was removed and all of the power from the VCSEL was coupled into the receiver. This led to an average input optical power of 371μW. With this power arriving at the photodetector, no performance degradation could be observed. Therefore, the upper limit of the dynamic range of the receiver is above anything the receiver will be required to handle in the final system.

7.6 Bit error rate test

The bit error rate (BER) test was not performed in the test rig previously described. Instead, the final system configuration was used. Two boards, fully populated with their optomechanical pieces, were placed facing each other, as shown schematically in Figure 2, and a photo of the setup is shown in Figure 49. The distance between the boards was 13.2 cm. Both boards were placed on micro-controlled motion stages to control the misalignment between the boards. For this test, the misalignment was minimized.

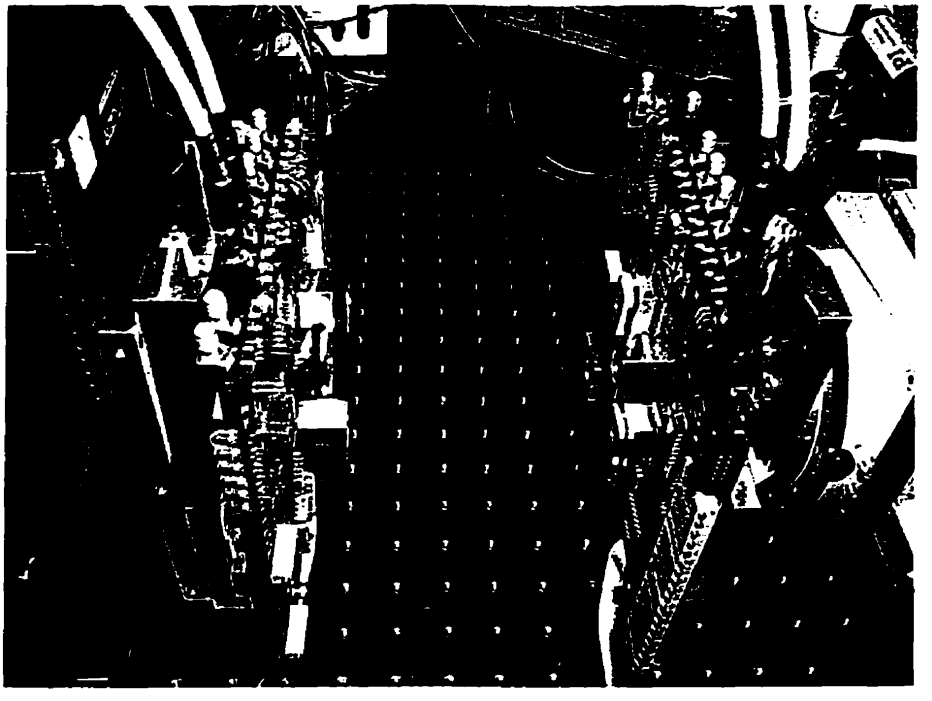

Figure 49. Bit error rate test setup

The transmitter $TX2_A$ on board A was used to establish an optical link to $RX2_B$ on board B. Before performing the BER test, the best VCSEL-PD pair was selected, which for this test was VCSEL number 8 from $TX2_A$ and photodetector

number 6 from $RX2_B$. A pseudo-random bit stream (2^7-1) was used to transmit data. A longer pseudo-random stream, such as 2^{23}-1, was not used, so as to minimize the lower frequency components in the data.

The results of the BER test are shown in Figure 50. Below 100Mb/s and above 730Mb/s, the BER tester was not able to synchronize to the data to give an error rate.

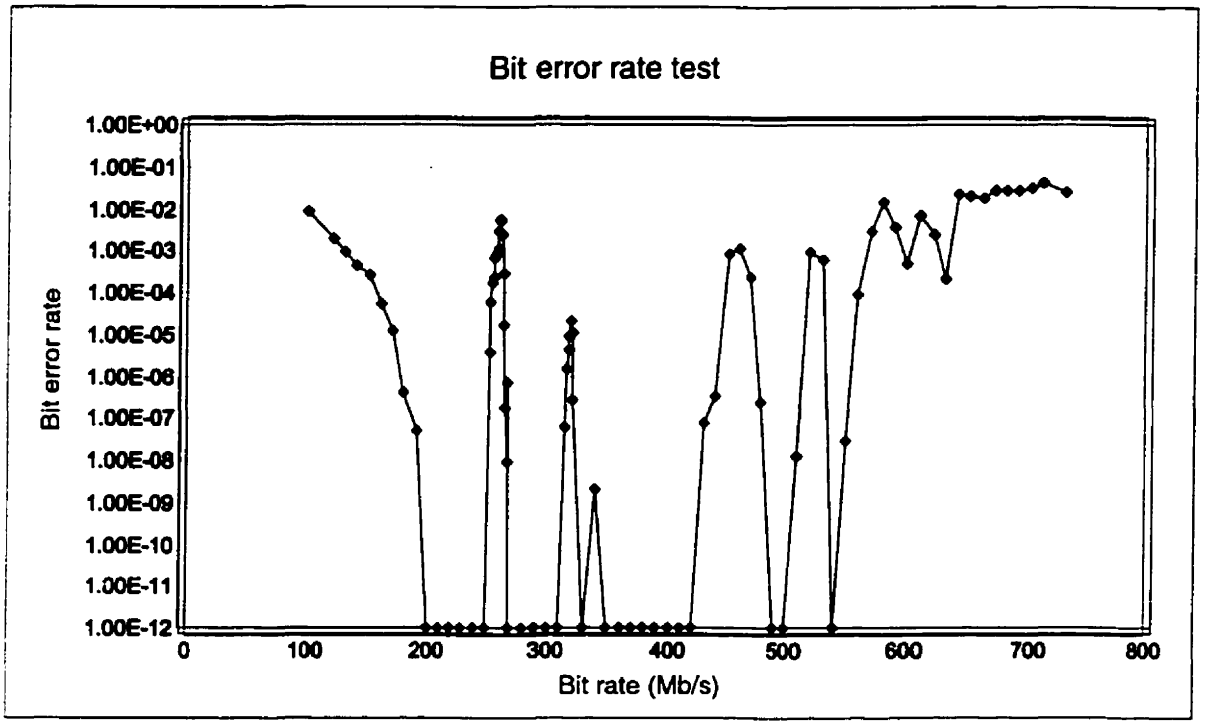

Figure 50. Bit error rate test results

The bit error rate decreased from 100Mb/s to 200Mb/s, and then increased again around 500Mb/s. This is consistent with the bandwidth of the receiver. All of the points in the graph that are shown to have a BER of 10^{-12} actually have a bit rate that is lower but the test was not continued long enough to detect an error. A

BER test was performed at 420Mb/s for 24 hours and no errors were detected, so at that bit rate the BER is $<10^{-13}$.

There are several points in the BER test where the bit error rate is quite high, for instance at 260Mb/s, and at 460Mb/s. However, no visible change or noise in the eye-diagrams could be detected at those frequencies when compared to the surrounding frequencies. To analyze this phenomenon further, a BER test was performed with a periodic signal as the input, instead of a pseudo-random sequence. The results of this test are shown in Figure 51.

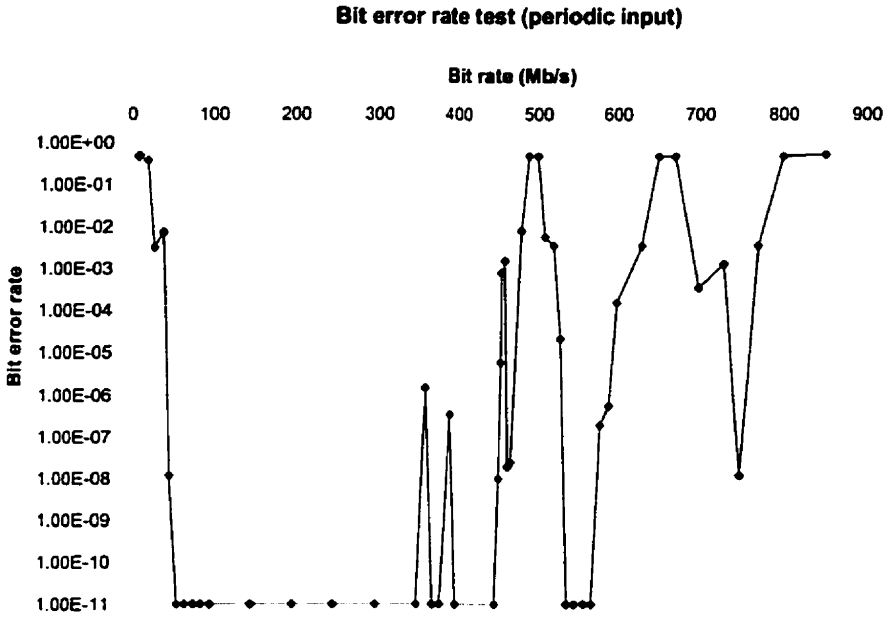

Figure 51. Bit error rate test results (periodic input)

Once again there are bit rates (500Mb/s for example) that generate very high error rates. One hypothesis to explain the high BER "spikes" in the graph is that at those bit rates the clock of the BER tester may have trouble synchronizing with the data and may sample the data at the wrong moment. Referring to Figure

44 the duty cycle of the waveform at 500Mb/s is not 50-50. This may cause timing problems since the duty cycle of the clock is 50-50.

7.7 Conclusions

The receiver was tested in a custom test rig to determine its bandwidth, sensitivity, and dynamic range. The bandwidth of the receiver was found to range from 18Mb/s to 750Mb/s. The receiver was not found to be as sensitive as simulated, but this did not reduce the misalignment tolerances to below specifications for the system. The upper bound of the dynamic range of the receiver could not be measured, but is above any input power the system could provide to the photodetector. The bit error rate of the receiver at certain bit rates is less than 10^{-12}, but at other bit rates it is quite high. The cause of this poor performance at those bit rates is as yet unknown, but may be due to a timing problem with the BER tester.

7.8 References

1. E. Bisaillon, et al., "Optical link for an adaptive redundant free space interconnect", in *Optics in Computing, 2000*, Roger A. Lessard, Tigran Galstian, Editors, SPIE Vol. 4089, pp. 999-1009, 2000.

2. Opto Speed SA, http://www.optospeed.com

3. M. Mony, et al. "1Gb/s, 9x1 optical receiver array, for an adaptive redundant interconnect system, using NT25 technology", in *Optics in Computing, 2000*, Roger A. Lessard, Tigran Galstian, Editors, SPIE Vol. 4089, pp. 1010-1020, 2000.

4. E. Bernier, et al., "1Gb/s transmitter for an adaptive redundant free-space interconnect", in *Optics in Computing, 2000*, Roger A. Lessard, Tigran Galstian, Editors, SPIE Vol. 4089, pp. 772-783, 2000.

5. D. V. Plant, et al., "A 5 Gb/s, 2 channel bi-directional adaptive redundant FSOI demonstrator system", in *Optics in Computing, 2000*, Roger A. Lessard, Tigran Galstian, Editors, SPIE Vol. 4089, pp. 465-472, 2000.

Chapter 8. Conclusion

8.1 Review

A 9×1 receiver array was successfully designed, fabricated, packaged, and tested for use in a self-aligning optical interconnect. This optical interconnect consists of two board interconnected by two bi-directional optical links. Each link makes use of spatial redundancy to increase the misalignment tolerance of the entire system. The system is able to handle lateral misalignments of ±1mm and angular misalignments of ±1 degree, with the distance between the boards varying continuously from 5cm to 21.5cm.

The receiver die is wirebonded to a photodetector die within a ceramic chip carrier. The receiver array accepts nine low-amplitude photocurrent swings, selects one of them, and then outputs that signal as a digital PECL signal. The selection of which channel to amplify is based on received power, and is performed off-chip.

The receiver module consists of the receiver chip, the photodetector chip, the ceramic carrier, and an optomechanical piece glued onto the corners of the ceramic carrier. The optomechanical piece holds a bulk focusing lens at the correct distance from the photodiode plane to defocus the VCSEL spot to approximately twice the diameter of the active area of the photodiodes. Having a defocused spot aids in increasing the misalignment tolerances between the circuit boards.

The receiver module is mounted on a printed circuit board. The high-speed traces leading to the receiver module are impedance controlled and do not limit the high-speed performance of the circuit. The circuit board also contains all of the external control circuitry that selects which photodiode to use to establish an optical channel.

The bandwidth of the receiver was found to range from 18Mb/s to 750Mb/s, which is comparable to the simulation results. The receiver was not found to be as sensitive as simulated, but this did not reduce the misalignment tolerances to below specifications for the system. The bit error rate of the receiver at certain bit rates was less than 10^{-12}.

8.2 Future work

Future work on the receiver would include increasing the bandwidth of the receiver, increasing the sensitivity of the receiver, and improving the bit error rate results for all frequencies within the operating bandwidth of the receiver. One method that could be investigated to increase the bandwidth of the receiver is to flip-chip the photodiode array onto the receiver die, rather than wirebonding the two together.

The transceiver circuitry could be easily extended to handle larger arrays of VCSELs and photodetectors. These larger arrays could be used for increased redundancy, which would further increase the misalignment tolerances between the circuit boards. Alternatively, the larger arrays of VCSELs and photodetectors could be used to increase the data throughput in between the circuit boards. In a larger array, several VCSEL-PD pairs may be aligned at the same time, and could

be used to send independent data streams. In such a system crosstalk would need to be carefully avoided by not choosing adjacent channels to transmit independent data. A third alternative would be to use multiple channels to send the same data, and use the redundancy for error correction.

Appendix A. Hunt Algorithm

The selection of which VCSEL-PD pair to use to transmit the optical data is handled by a hunt algorithm. This algorithm is based on a master/slave scheme and chooses the best VCSEL-PD pair for both channels in the bi-directional link. Since there are no electrical connections in between the boards, all of the feedback between boards must be performed optically. The algorithm is shown schematically in Figure 52.

One board is defined in hardware as the master board, and the other as the slave board. The first step of the hunt algorithm is to synchronize the two boards. The master board modulates, or flashes, all 9 of its VCSELs at once, and the slave board scans the 9 receiver channels one at a time. As soon as the slave board detects a signal on one of the channels, it responds by flashing all of its VCSELs. The master board detects this response signal, and as soon as the slave board stops transmitting its response, both boards enter the next state of the algorithm at roughly the same time, hence synchronized.

The next state involves finding the best VCSEL-PD pair by scanning through all the VCSELs on the master board, and finding which receiver channel gets the highest received power on the slave board. The boards then enter the coordination state where the information about which VCSEL to use is relayed back to the master board through a series of timed flashes of the VCSELs on the slave board.

The boards then swap roles, and the best VCSEL-PD pair is found for the other half of the bi-directional link. Once both VCSEL-PD pairs have been

found, data is transmitted across the two channels. If the received power on one board ever drops below a certain threshold, all data transmission stops, and the hunt algorithm starts again and finds the new best VCSEL-PD pair.

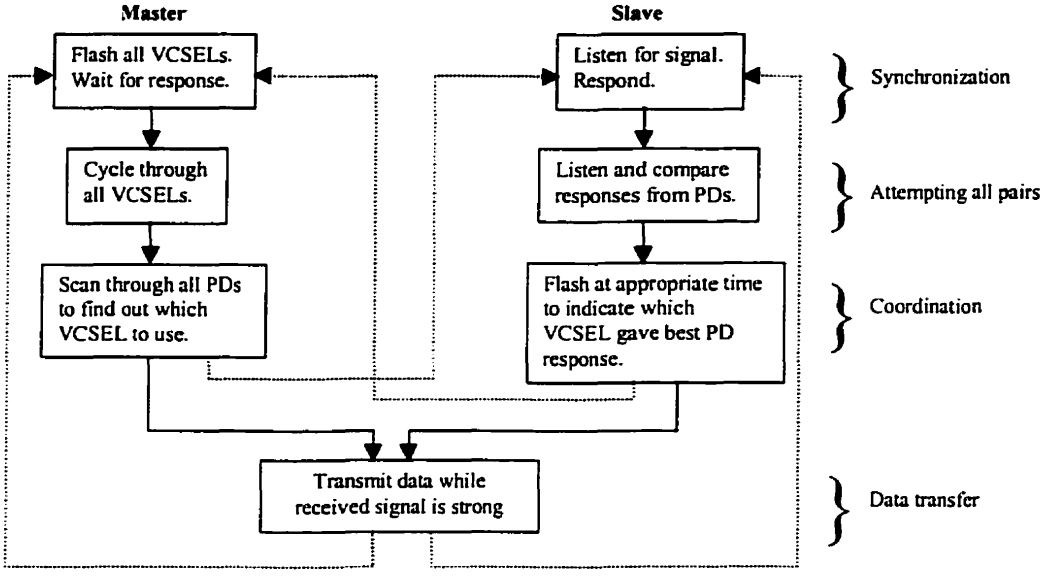

Figure 52. Hunt algorithm