Design Issues for Through-Wafer Optoelectronic Multicomputer Interconnects

P. May, S. T. Wilkinson, N. M. Jokerst, D. S. Wills,
M. Lee, O. Vendier, S. W. Bond, Z. Hou, G. Dagnall, M. A. Brooke, A. Brown
School of Electrical and Computer Engineering
Microelectronics Research Center
Georgia Institute of Technology
Atlanta, Georgia 30332-0250

Abstract

This paper presents several design issues associated with the implementation of a three dimensional optically interconnected parallel processing system. A technique for improving bit error rate in low power multistage networks is presented. Error detection codes are transmitted along with message data to guarantee the integrity of the data during each optical hop. To realize three dimensional through-silicon wafer interconnect, thin film emitters and detectors operating at a wavelength of 1.3 μm (to which silicon is transparent) will be bonded to the silicon circuitry. A transfer diaphragm process is used to realize this integration [1]; this process has been used to demonstrate the basic concept: a single silicon circuit has been integrated with both a thin film emitter and detector operating at 1.3 μm wavelength [2]. In this paper, we will utilize one possible integration scenario to illustrate the trade-offs associated with a system of this type, which includes device design, circuit design, and issues which include manufacturability, alignment tolerance, crosstalk, and power dissipation.

1. Introduction

Optoelectronics offers significant benefits in the implementation of multicomputer interconnection networks including three-dimensional interconnect, high interchip bandwidth, and extremely compact packaging. Multistage optoelectronic networks offer interconnect generality and scalability that is not possible in a single stage networks. To realize their full potential, optoelectronic networks must be designed to meet the system requirements of current and future parallel processing systems.

This paper presents several design issues associated with the implementation of a through-wafer three dimensional multicomputer interconnection network. The next section presents a technique for incorporation of error detection codes in short distance, low power transmission to lower the bit error rate. Section three addresses device parameters required to achieve a 100 Mbit/sec optoelectronic link.

2. Error detection codes for improved bit error rate

Wire-based multistage interconnection networks have been built for several decades. The MIT J-Machine [3] is a good example, providing a high throughput, low latency implementation of a three-dimensional mesh with 4096 nodes. Because of the lower speed of the wires (compared with optoelectronics) and the high drive current of the output pads, transmission errors are rare. Even in long distance telecommunications application, end-to-end error detection schemes are used in which transmission errors are detected only at the message destination.

For optoelectronic interconnection networks employed in dense three dimensional systems (e.g., the Georgia Tech Pica architecture [4]), a direct correlation exists between link power (emitter driver, emitter, and detector amplifier power) and improved bit error rates. Increasing link power drives down the error rate, but complicates power management in a compact system.

This section describes a *hop-to-hop* approach to reduce bit error rates where hardware validates the integrity of the message at each optoelectronic hop between the source and the destination. It begins with an overview of direct multistage optoelectronic interconnection network. Then a description of cyclic redundancy codes or *CRC*s for error detection is presented. Finally, the link hardware required for hop-to-hop error detection and correction is described.

2.1 Network architecture

A typical multistage optoelectronic interconnection network is shown below in Figure 1.

Each stage transition or *hop* includes output buffering, optical transmission and reception, and input buffering, and switch routing. In a wormhole routed network [5, 6], a message is routed along the transmission path in pipelined technique. This means that each transmitted segment of the message (called a flow control digit or *flit*) is routed in a chain from the head of the message to the

0-8186-7101-7/95 $4.00 © 1995 IEEE

tail. Because buffering exists between each hop of the message, an opportunity exists for hop-to-hop error detection by validating the flit at the receiver. If a transmission error occurred, a duplicate copy of the flit can be requested from the transmitter's buffer. (Transmissions within a router are assumed to be error free.) This approach requires a technique for validating a value's integrity. Cyclic redundancy codes (CRCs) fill this need.

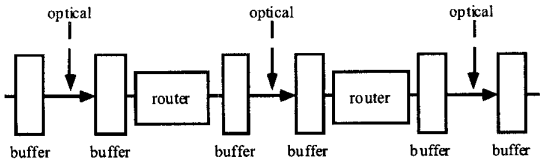

Figure 1: A multistage optoelectronic interconnection network.

2.2 Cyclic redundancy codes

CRCs are used to protect data transmitted on noisy links by adding additional information to the data stream. This information is added in such a way that the receiving end of the link can determine whether errors have occurred in the transmission. The easiest way to think of CRCs are as remainders of a long division process involving the data to be transmitted. To start with, data to be transmitted is represented by a polynomial D(x), with each of its coefficients being either a 0 or a 1. These 0's and 1's correspond exactly to the bits in the data word. Addition and subtraction performed on this polynomial are done module base two term by term. Another polynomial called the generator polynomial or G(x) is selected based on its ability to detect errors. The data that the system transmits is the message polynomial M(x). M(x) is D(x) times the greatest power of x in G(x) (represented by x^g) plus the remainder of $x^gD(x)/G(x)$ which is represented by R(x).

$$M(x) = xgD(x) + \text{remainder of } [xgD(x) /G(x)]$$

On the receive end of the link, M(x) is divided by G(x), and since M(x) is created by adding the remainder and the dividend of a previous division by G(x), the remainder will be 0. If there are bit errors present, the remainder should be non-zero.

Errors can also be represented by a polynomial E(x). The received data with errors is now represented by E(x) + M(x). In order for errors to be detected on the receive end, [E(x) + M(x)]/G(x) must produce a remainder. [E(x) + M(x)]/G(x) = M(x)/G(x) + E(x)/G(x) which means that E(x)/G(x) must produce a remainder for the errors to be detected (since we know that M(x)/G(x) produces no remainder). Picking the generator is critical since the type of errors that can be detected are dependent on its properties. Any E(x) with G(x) as a factor will produce a zero remainder in the presence of errors.

Generators with (x+1) as a factor can detect any odd number of errors [7]. To detect two bit errors, G(x) must not be evenly divisible by any possible error polynomial E(x) that can create two bit errors. That means that E(x)/G(x) must have a non zero remainder for all $x^n + x^m$ where n and m are the powers of x in D(x) corresponding to the bits in error. In other words, any possible spacing of two bits in the message must not be evenly divisible by G(x). If G(x) is not divisible by x (which we will assume in all cases), then G(x) must not evenly divide $x^{n-m} + 1$ where n-m is the in the range from 1 to the maximum spacing of bits in the message (e.g., n-m is in the range 1 to 15 for 16 bit data). To restate this, if $(x^{n-m} + 1)/G(x)$ produces a remainder for all n-m from 1 to the message length minus 1, then G(x) will detect all possible double bit errors.

The following example illustrates the way data is converted to a polynomial using a 16 bit data word:

$$D = 1011010010101001$$
$$D(x) = x^{15}+x^{13}+x^{12}+x^{10}+x^7+x^5+x^3+1$$

To add two of these polynomials, the coefficients of the terms with the same powers of x are added modulo base two. It is important here to note that since the additions are being performed modulo base two, they are identical to subtraction. So, adding two polynomials created in the above manner produces the same result as subtracting them. This will be important when long divisions are considered. The following example shows two data words D and C, and the addition of the two associated polynomials:

$$\begin{array}{r} D = 1011010010101001 \\ \underline{+\ C = 0011011101101011} \\ 1000001111000010 \end{array}$$

$$\begin{array}{rl} D(x) = & x^{15}+x^{13}+x^{12}+x^{10}+\quad x^7+\quad x^5+x^3+\ 1 \\ \underline{+\ C(x) = \quad x^{13}+x^{12}+x^{10}+x^9+x^8+\quad x^6+x^5+x^3+x+1} & \\ x^{15} \quad +x^9+x^8+x^7+x^6+\quad x \end{array}$$

In this example, the polynomials are added(subtracted) the same way as polynomials with base 10 coefficients, the only difference being that the coefficients in this case can only have the values 0 or 1. When the two polynomials are represented in their original binary form, the operation for addition or subtraction is simply a bitwise XOR.

Multiplying a data polynomial by a x^g is simply a left shift of the bits by the amount g. This is illustrated in the following example using both the data word D and its polynomial D(x) with g = 5.

$$D = 1011010010101001$$
$$D(x) = x^{15}+x^{13}+x^{12}+x^{10}+x^7+x^5+x^3+1$$
$$x^5D(x) = x^{20}+x^{18}+x^{17}+x^{15}+x^{12}+x^{10}+x^8+x^5$$
$$10110100101010100100000$$

The data D in the above example is shifted left by 5, the power of x by which D(x) is multiplied.

2.3 Hardware implementation

The division described above can be carried out using specialized hardware which performs these divisions efficiently on serial data streams [8, 9]. An example of one such circuit that implements the G(x) used in the above examples is shown in Figure 2.

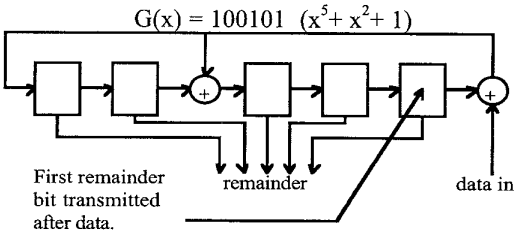

Figure 2: Hardware implementation of example CRC.

D(x) is fed into the input of this circuit and after the last bit has been clocked in, the remainder (R(x)) is held in the circuit storage elements. This remainder is appended immediately following the data, creating M(x) to transmit. In this implementation, the data D(x) does not need to be shifted by g bits (multiplying by x^g) as this CRC circuit performs that function as well. The receiver uses identical circuitry. After M(x) has been shifted in on the receive end, the storage elements of the circuit should be all 0's. If any of the values in the circuit are 1, an error has occurred.

In our application, 32 bit data words (flits) are being transmitted across our interprocessor communication links. The goal of the CRC in our system is to detect any single or double bit errors in any given flit. Given the criteria for error detection identified in the previous section, there are several possible choices for CRC. The CRC generator polynomial was chosen for our application is:

$$G(x) = x^7 + x^5 + x^3 + 1$$

The circuit for calculating the above CRC is shown in Figure 3.

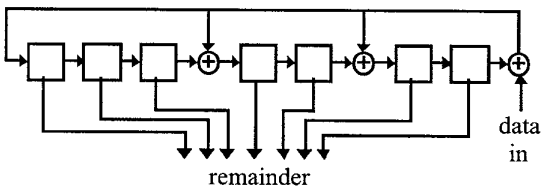

Figure 3: $G(x) = x^7 + x^5 + x^3 + 1$

The CRC hardware implement consists of two halves, a transmit CRC circuit and a receive CRC circuit. These circuits are largely the same, utilizing many of the same functional components. The transmit circuit will calculate a 7 bit CRC for every 32 bits of data fed to it. This data/CRC will be transmitted in the following format:

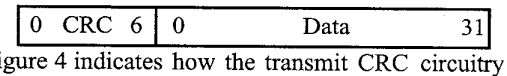

Figure 4 indicates how the transmit CRC circuitry and receive CRC circuitry will fit into the overall communication system.

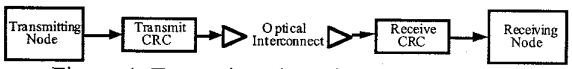

Figure 4: Transmit and receive CRC circuitry.

The receive CRC circuitry will perform a CRC calculation on the transmitted data and check that the remainder of the division that the CRC represents is zero. If the remainder is zero, the data has been received correctly (error free), and the transmitting node receives an 'ACK' message from the receiving node. If, on the other hand, the remainder is not zero, the receiving node will send a 'NACK/Retransmit' message to the transmitting node. This will indicate to the transmitting node that the message is to be retransmitted. There is another form of NACK which is the 'NACK/Blocked' message indicating that the head of the message has blocked. No retransmission occurs in this case. The transmit CRC and receive CRC circuits are described at the block diagram level in the next sections.

The block diagram of the Transmit CRC circuit is shown in Figure 5.

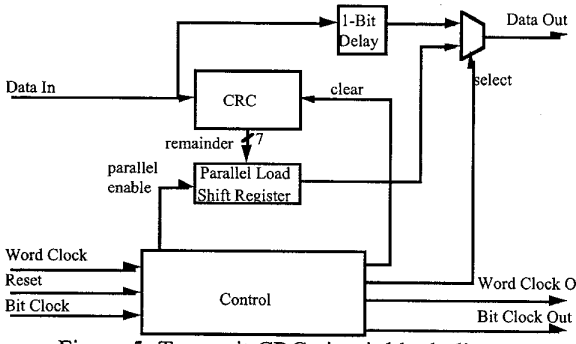

Figure 5: Transmit CRC circuit block diagram.

The transmit circuit receives data serially over the *Data In* line. The timing for this data is provided by the *Bit Clock* and *Word Clock* lines. As the data enters the circuit, each bit is immediately passed to the CRC circuit where the remainder calculation begins. At the same time, the input data is bypassed around the CRC circuit and delayed by one bit time. The multiplexer feeding the *Data Out* line is switched so that the *Data Out* line is the one bit delayed copy of *Data In*. Since the message M(x) is x^g D(x) + R(x), the first 32 bits out of the circuit must be the first 32 bits into the circuit. After the arrival of the 32nd data bit, the CRC calculation will be complete (after the 32nd bit has been clocked into the circuit). At this time, the *Control* circuitry asserts the parallel enable line to the *Parallel Load Shift Register*, latching the value of the remainder calculation. The select line to the multiplexer is switched immediately following the remainder transfer, and the remainder is shifted out. This sequence creates a

data stream that consists of 32 data bits followed by 7 CRC remainder bits on the *Data Out* line.

The block diagram of the Receive CRC circuit is shown in Figure 6.

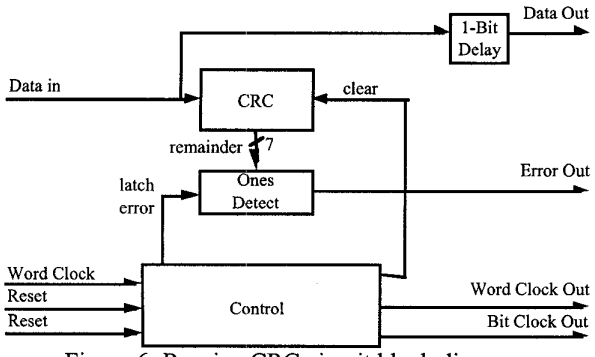

Figure 6: Receive CRC circuit block diagram.

The receive circuitry receives M(x) (data bits plus CRC) on the *Data In* line with timing provided by *Bit Clock* and *Word Clock* as in the transmit circuitry. The CRC calculation is performed on the data + remainder and the result is latched into the *Ones Detect* block by the *Control* circuitry after the 39th bit has arrived. The result of the CRC calculation is reflected in the *Error Out* signal. An asserted *Error Out* indicates that there was a non-zero remainder, and there is an error present in the current data. The receiving node will act on the value of the *Error Out* line, generating either an 'ACK' or 'NACK/Retransmit' response. While the CRC calculation is being performed, the received bits are being sent out on the *Data Out* line. The CRC remainder bits are also sent out on the *Data Out* line, but the *Word Clock Out* signal indicates to the downstream circuitry which bits are the 32 data bits.

This CRC circuit has been designed at the gate and transistor level, and a layout of the circuit is also complete. A test circuit is being fabricated using MOSIS CMOS 2 μm technology. This circuit will be tested to determine how it will change the characteristics and overall performance of the optical interconnection network. To realize three dimensional through-silicon wafer interconnect, thin film emitters and detectors operating at a wavelength of 1.3 μm (to which silicon is transparent) will be bonded to the silicon circuitry. A transfer diaphragm process is used to realize this integration [1]. This process has been used to demonstrate the basic concept: a single silicon circuit has been integrated with both a thin film emitter and detector operating at 1.3 μm wavelength [2, 10]. In this paper, we will utilize one possible integration scenario to illustrate the trade-offs associated with a system of this type, which includes device design, circuit design, and issues which include manufacturability, alignment tolerance, crosstalk, and power dissipation.

3. Optoelectronic device design tradeoffs

The example which will be examined herein is the most manufacturable, but engages severe tradeoffs in performance: we will assume, for simplicity of integration, that the emitters and detectors are composed of the same structure. The use of this example does not imply that this integration technology nor this system is limited to this example scenario; on the contrary, the thin film integration used for the physical realization of this project enables the independent optimization of the emitters and detectors bonded to the silicon circuit. From the manufacturing perspective, however, the simplest realization is that in which the emitters and detectors are the same structure, and are bonded simultaneously to the silicon circuit. It is this scenario which will be theoretically examined herein to highlight the issues which are inherent in the design of one particular optoelectronic through-silicon wafer interconnect.

For 1.3 μm device operation, both InGaAsP and InGaAlAs can be grown with an energy bandgap of 1.3 μm and lattice-matched to InP, and both can be processed into LEDs and detectors using the same processing. Both materials have almost the same light absorption coefficient value at 1.3 μm, which directly determines the internal efficiency in the detectors and emitters. However, from the fundamental material growth perspective, there are several tradeoffs between the two compounds that must be considered before using one of these compounds for the optical devices. While InGaAsP and InGaAlAs are both quaternary compounds, InGaAlAs contains three column III elements (indium, gallium, and aluminum) and one column V element (arsenic) while InGaAsP contains two column III elements and two column V elements. For the temperature regions in which the two compounds are grown, the column III elements have sticking coefficients of approximately 1 while the sticking coefficients of the column V elements are much less than 1. Because of this, the two column V elements in InGaAsP add an additional degree of complexity in the growth of this quaternary that is not present in InGaAlAs. In addition, the conduction band offset, $\Delta Ec/\Delta Eg$, is greater in InGaAlAs than for InGaAsP for similar heterojunctions, which provides greater confinement of electrons in active regions. These tradeoffs in InGaAsP are offset by the tradeoffs in InGaAlAs such as the problem that aluminum in the InGaAlAs compound will easily react with any residual oxygen inside the growth chamber of the molecular beam epitaxy machine to form nonradiative centers. Also, InGaAlAs has to be grown at higher temperatures than the InGaAsP material to achieve the same high material quality. The modelling in this paper was performed assuming InGaAsP devices due to the amount of data available on the optical properties of InGaAsP and the relatively sparse data on InGaAlAs,.

The assumptions that are made for this example are that the through-wafer optical interconnections operate at 100 Mbps, that the silicon circuitry will be 0.8 μm MOSIS

foundry CMOS, and that the circuitry will be limited to the well-characterized circuits which have already been demonstrated for this type of interconnect, namely, a five stage transconductance amplifier for the detector receiver [11], and a high current CMOS emitter driver [12]. To achieve 100 Mbps operation, the receiver design dictates that the detector capacitance must be less than or equal to 1 pF; we will treat the cases from 0.1 to 1 pF, as dictated by

$$f = 1/(2\pi RC),$$

where f is the operating frequency of the circuit, R is in the input resistance, and C is the input capacitance plus detector capacitance, the sum of which is generally dominated by the detector capacitance. Thus R can be determined for this circuit. The power dissipation of this circuit is set by the capacitance of the detector, as given by

$$P = V \cdot \frac{(2\pi fC)^2}{k_n\left(\frac{W}{L}\right)}$$

Where V is the rail-to-rail power supply, f the frequency bandwidth, k_n the MOSFET transistor parameter, and W/L the gate width and length of the input transistor. This is also illustrated graphically in Figure 7.

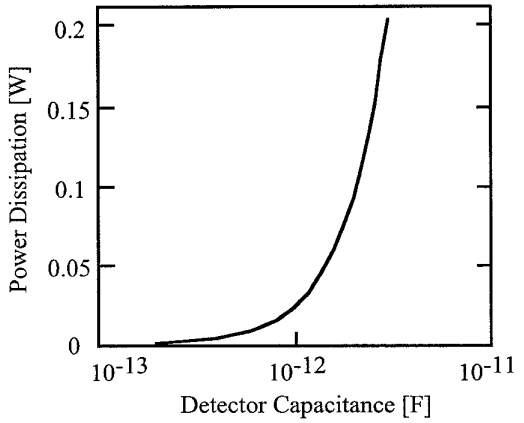

Figure 7: Detector capacitance versus power dissipation.

To determine the noise for this circuit, $\langle i \rangle_{total}^2$,

$$\langle i \rangle_{TOTAL}^2 = 2qI_{da}BI_2 + 2qI_{gate}BI_2 + \frac{8kTg_m}{3}BI_2 + \frac{8kT}{3g_m}(2\pi C_T)^2 I_3 B^3$$

is used, where I_{da} is the dark current of the photodetector, I_{gate} the leakage current of the input transistor, g_m the transconductance of the transistor, C_T the total input capacitance, q the electron charge, B the operating bandwidth, k the Boltzmann constant, T the absolute temperature, I_2 and I_3 are weighting functions which are dependent only on the input optical pulse shape to the receiver and equalized output pulse shape. For non-return-to-zero coding, I_2=0.55 and I_3=0.085 [13]. Using SONET specifications, 1 μA of current from the detector is the minimum signal, and a signal to noise ratio of 10 is necessary. Thus, noise must be limited to 0.1 μA.

The detectors used for integration, and in this example, are double heterostructures with a highly doped active region. The capacitance for a P-i-N is given by

$$C = \frac{\varepsilon A}{d}$$

where ε is the semiconductor permittivity, A is the detector cross-sectional area, and d is the depletion region thickness. For a P-i-N detector, the i region is generally fully depleted, and the thickness of the i region is used for d. Figure 8 illustrates the trade-offs between A and d to achieve the specified 0.1 to 1 pF capacitance. For a p-doped active region, the analysis is slightly more complicated, but the trends are the same. In order to achieve 100 Mbps, the doping in the gain region of the InGaAsP emitters must be on the order of 10^{19} cm^{-3} [14], which will limit the efficiency and size of the detectors since the same material is to be used for both devices. The structure used for modelling of both the emitter and detector is a P-p^+-N InGaAsP structure with doping of $1x10^{17}$ cm^{-3} in the cladding regions and doping of $2x10^{19}$ cm^{-3} in the 0.5 μm thick active absorbing region. The depletion region formed in a structure having this dopant level (assuming voltages possible in standard silicon circuitry) is 0.11 μm [15], so 0.11 μm is used as the depletion region thickness in the detector calculations.

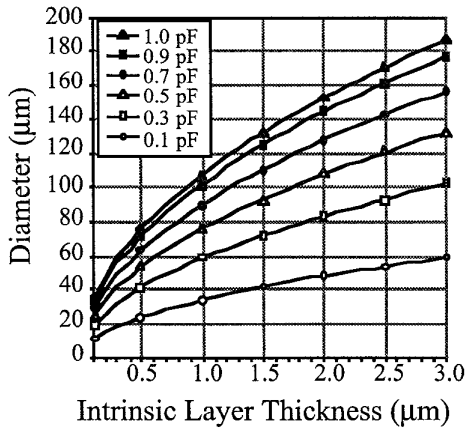

Figure 8: Detector area versus intrinsic region thickness for a variety of capacitances.

For the remainder of our example, a 30 μm diameter detector with a 20 μm detection diameter will be used for calculations because this is the largest detector size which will achieve 100 Mbits/s. Assuming one pass of the incident light through the photodetector, this design will result in an efficiency of η = 6.7%, as given by:

$$\eta = (1 - R_f)(1 - e^{-\alpha d})$$

where R_f is the front mirror reflectivity, which is the semiconductor/air Fresnel reflectivity of 30%, α is the absorption coefficient, which is approximately 0.2 μm^{-1} for a doping of $2x10^{19}$ cm^{-3} [16], and l is the thickness of the absorbing layer (photogenerated carriers created within one depletion length of the depletion region are collected before recombination, the mobility for material doped to this level is approximately $1.5x10^3$ $cm^2/(V \cdot s)$ and the carrier lifetime is 0.45 ns, which results in a diffusion length greater than 0.5 μm). The associated detector responsivity (R) is approximately 0.070 A/W, as calculated from:

$$R = \frac{\eta q}{h\nu}$$

where h is Planck's constant, and ν is the photon frequency.

This detector efficiency can be greatly improved. The assumption that the incident light traverses the absorbing region only once is a conservative assumption. These thin film detectors generally exhibit resonant cavity enhanced behavior, since one side of the cavity has an enhanced reflectivity due to the broad area metallic ohmic contact. For detection optimized at a cavity resonance, the efficiency is calculated from the equation:

$$\eta = \frac{\left(1+R_b e^{-\alpha l}\right)\left(1-R_f\right)\left(1-e^{-\alpha l}\right)}{1-2\sqrt{R_b R_f}\cos\left(2\beta l+\delta_b+\delta_f\right)e^{-\alpha l}+R_b R_f e^{-2\alpha l}}$$

where R_b and R_f are the back and front mirror reflectivities, respectively (0.92 for metal ohmic contact mirror and 0.3 for air/semiconductor Fresnel interface), $\beta=2\pi n/\lambda$, and δ_b, and δ_f are the phase changes due to reflection on the back and front mirrors, respectively. For our calculation, we assume that the resonant cavity has been optimized, so that the cosine term is equal to 1. The resonant case efficiency is 44.3% (considerably greater than the single pass efficiency of 6.7%).

For our current design, the detector design is limited since the same semiconductor material is used for both the emitters and the detectors. Although P-p^+-N detectors are used for this example, metal-semiconductor-metal (MSM) detectors are particularly attractive for integration since they have a low capacitance coupled with a large area. MSM detectors, however, have all undoped layers and cannot be used as emitters. The efficiency of MSMs is typically poor due to finger shadowing. However, recently reported inverted MSMs (with fingers on the bottom [17]), which demonstrate much higher efficiency than conventional MSMs are an excellent option for this type of integration.

The light emitting diode output power (L) is directly proportional to the drive current (I) and the emission (active) volume (V), as given by

$$L = kIV$$

where k is a proportionality constant given by

$$k = \frac{\left(\frac{\partial L}{\partial I}\right)}{V}.$$

GaAs-based double heterostructure thin film emitters with a similar structure to that examined herein have been fabricated and characterized [18]. The dimensions of the GaAs LEDs were 250 μm x 250 μm square, with a 0.5 μm thick active region. V was 31250 μm^3, and the measured ($\delta L/\delta I$) = 30 μW/mA. The measured divergence half-angle of these thin film (resonant cavity enhanced) light emitting diodes was 23.7°, with a far-field pattern of $\cos^{11.3}(\theta)$. For ease of simulation, this pattern description was rounded to $\cos^{12}(\theta)$.

The interconnect efficiency is a strong function of misalignment, as shown in Figure 10, where z is the distance along the detector plane, x is the distance along the emitter plane, and o is the offset of the detector from z=0. The measurement of the far-field pattern assumed a point source, which was valid since the source-to-detector distance used during measurement was much greater than the device sizes. For the through-wafer optical interconnect, however, the separation distances are on the order of the device size, so the source is modeled as a continuous distribution of point sources with a $\cos^{12}(\theta)$ distribution. To simplify the calculation, only two dimensions were used, so offsets can only be in a single direction, as illustrated in Figure 9.

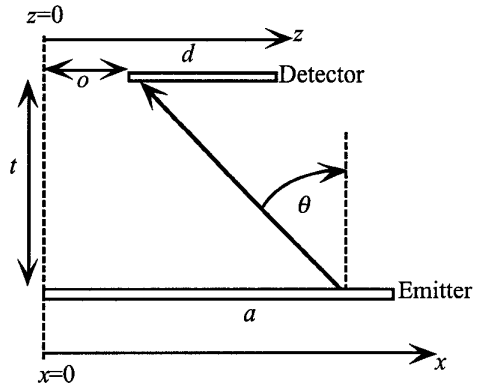

$$\cos(\theta) = \frac{t}{\sqrt{t^2+(z-x)^2}}$$

Figure 9: Two dimensional alignment geometry.

The power density at any point (z) on the detector plane for separations (t) is the integral over the entire source area of the continuous distribution of point sources, as given by:

$$P(z,t) = \int_0^a \frac{t^{12}}{\left((z-x)^2+t^2\right)^6}dx$$

The total power incident onto the detector plane is the integral of the power density along the entire plane, as given by:

$$P_{total}(t) = \int_{-\infty}^{\infty} P(z,t)dz$$

The percentage of total power incident onto the detector is the integral of the power density across the detector aperture divided by the total power incident onto the detector plane, as given by

$$P_{detector} = \frac{\int_{o}^{o+d} P(z,t)dz}{P_{total}(t)} \cdot 100$$

The result is plotted in Figure 10, below, showing percent coupling as a function of offset (o) for a variety of vertical separations (t).

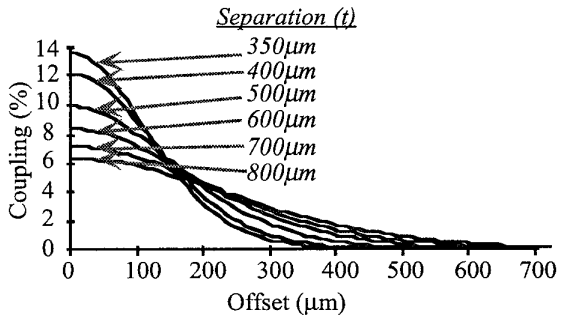

Figure 10: Coupling efficiency versus misalignment at a variety of separations
(detector aperture=20μm, emitter aperture=150μm)

Since the minimum signal for the receiver circuit is 1 μA, and the detector has a responsivity of 0.070 A/W, 14 μW must be incident upon the detector. Assuming that the back surface of the silicon circuitry is polished and an antireflection coating is applied, the primary source of coupling loss between the emitter and detector is misalignment and vertical separation. Assuming that the vertical separation is 350 μm, which is a standard silicon circuit thickness, perfect alignment yields a coupling factor of 13.5%. This implies that an output power of 105 μW is required from the LED. Linearly scaling the drive current from $L=kIV$, results in a 3.5 mA drive current. The power dissipation from the emitter driver (assuming 100% thermal power dissipation from the LED, which is valid, since these LEDs are approximately 1.4% efficient) is:

$$P_{max} = I_{LED} \cdot V_{LED} + \left(I_{Modulation} + I_{bias}\right) \cdot V_{Circuit}$$

where I_{LED} and V_{LED} are the current and voltage across the LED, $I_{modulation}$ and I_{bias} are the currents through the modulation and bias branches of the circuit, and $V_{circuit}$ is the power supply voltage for the cirucit. The maximum values for $I_{modulation}$ and I_{bias} are $^1/_5$ of I_{LED}. $V_{circuit}$ is typically 5V, and V_{LED} has a maximum value of 5V. Therefore, assuming a drive current of 10 mA (engineering from 3.5 mA to 10 mA for a safe margin of operation; similar GaAs LEDs have demonstrated up to 750 μW of output power at 25 mA [18]), the maximum total power dissipation is 55 mW.

The tight tolerance margins on this design can be greatly relieved simply by using two different semiconductor structures for the emitter and detector. The high doping necessary to increase the speed of the LEDs results in a significant loss of efficiency in the detector. A lower efficiency requires more input power from the emitter, which requires a higher output LED and/or extreme alignment tolerances. A higher output LED requires more current (which causes the power dissipation to rise), and perfect alignment is difficult to achieve. However, by using a P-i-N detector with an undoped active region, the characteristic of the entire system can be significantly improved. Assuming a detector structure with a 1.5 μm thick intrisic active region, the efficiency is 57% (compared to 6.7% for the highly doped gain region). This efficiency yields a responsivity of 0.64 A/W, so that only 1.6 μW of incident light is necessary to achieve 1 μA of current (compared to 14 μW). Also, the fully depleted active region thickness means that a much larger detector (95 μm diameter with an 80 μm window) can be used while retaining the 0.5 pF capacitance necessary to achieve 100 Mbit/s.

Even without preparing the back surface of the silicon circuitry, optical signals can be detected through-wafer when an undoped detector is used. With a 30% silicon/air reflectivity (with no anti-reflection coating), and assuming 50% scattering loss at the silicon/air interface (assuming that the back of the silicon integrated circuits are unpolished), and a 13% coupling efficiency (recalculating the data in Figure 10 with an 80 μm detector window, and assuming 200 μm misalignment), 4.5% of the emitted light is coupled into the emitter. Using an engineering safety factor of about 10, 0.5% of the light from the emitter will be incident upon the detector in this through-silicon optical interconnect example. Thus, 400 μW of light must be output from the light emitting diode. Linearly scaling the drive current results in a 13 mA drive current. The power dissipation from the emitter driver (once again, engineering from 13 mA to 15 mA) is 82.5 mW.

The maximum separation of the two through-wafer optically interconnected layers is set by the coupling loss which can be tolerated. For the example used, the separation is 350 μm, which is, on average, the thickness of the silicon circuits from MOSIS. Given the engineering factor of 10, the coupling can be represented as 1.3% rather than 13% so, from recalculating the data in Figure 10 for an 80 μm detector window, the offset (misalignment) can range from 350 μm to 650 μm with a corresponding separation of 350 μm to 800 μm.

The sample integration structure examined herein is not representative of the best nor worst design of a through-silicon wafer optically interconnected system, however, it

is instructive to investigate the trade-offs associated with this type of interconnection system.

4. Conclusions

We have presented several design issues associated with the implementation of a three dimensional optically interconnected parallel processing system. Two primary concerns, improvement of bit error rate and realization of through-silicon wafer interconnect, have been discussed. Bit error rate improvement is realized by transmitting error detection codes along with message data to guarantee the integrity of the data during each optical hop. Three dimensional through-silicon wafer interconnect (using thin film emitters and detectors operating at a wavelength of 1.3 μm bonded to silicon circuitry) has been analyzed for device efficiency, alignment tolerance, power dissipation, and system throughput. The example uses the same material structure for both the detectors and the emitters for increased manufacturability. However, increased system performance is possible if separate device structures are used.

5. References

[1] C. Camperi-Ginstet, M. Hargis, N. Jokerst, and M. Allen, "Alignable Epitaxial Liftoff of GaAs Material with Selective Deposition Using Polyimide Diaphragms," *IEEE Photon. Tech. Letters*, 3(12):1123-1126, December 1991.

[2] K. H. Calhoun, C. B. Camperi-Ginestet, and N. M. Jokerst, "Vertical Optical Communication Through Stacked Silicon Wafers Using Hybrid Monolithic Thin Film InGaAsP Emitters and Detectors," *IEEE Photon Tech. Letters*, 5(2):254-257, February 1993.

[3] W. J. Dally, J. A. S. Fiske, J. S. Keen, R. A. Lethin, M. D. Noakes, P. Nuth, R. Davison, and G. Fyler, "The Message-Driven Processor: A Multicomputer Processing Node with Efficient Mechanisms," *IEEE Micro*, 12(2):23-39, April 1992.

[4] W. S. Lacy, C. Camperi-Ginestet, B. Buchanan, D. S. Wills, N. M. Jokerst, and M. Brooke, "A Fine-Grain, High Throughput Architecture Using Through-Wafer Optical Interconnect," in *First International Workshop on Massively Parallel Processing Using Optical Interconnections*, pages 27-36, Cancun, Mexico, 26-27 April 1994.

[5] W. J. Dally and C. L. Seitz, "Deadlock-Free Message Rounting in Multiprocessor Interconnected Networks," *IEEE Trans. on Computers*, C-36(5):547-553, May 1987.

[6] William J. Dally, Performance Analysis of *k*-ary *n*-cube Interconnection Networks, *IEEE Transactions on Computers*, C-39(6):775-785, June 1990.

[7] A. S. Tanenbaum, Computer Networks. Prentice-Hall, Inc., N.J., 1981.

[8] T. V. Ramabadran and S. S. Gaitonde, "A Tutorial on CRC Computations," *IEEE Micro*, pp. 62-75, August 1988.

[9] T. R. N. Rao and E. Fujiwara, Error-Control Coding for Computer Systems. Prentice-Hall, Inc., N.J., 1989.

[10] C. Camperi-Ginestet, B. Buchanan, N. Jokerst, M. Brooke, "Bi-Directional Communication Through Stacked Silicon Circuitry Using Integrated Thin Film InP-Based Emitters and Detectors," *Proceedings of the Conference on Lasers and Electro-Optics*, Baltimore, June 1995.

[11] M. Lee, C. Camperi-Ginestet, M. Brooke, N. Jokerst, "Silicon CMOS Optical Receiver with Integrated Compound Semiconductor Thin Film P-i-N Detector," *Proceedings of the IEEE LEOS Summer Topical Meeting on Integrated Optoelectronics*, p. 58-59, Lake Tahoe, 6-8 July 1994.

[12] To be published, Georgia Tech, 1995.

[13] R. H. Smith and S. D. Personick, "Receiver Design for Optical Fiber Communications Systems" in Semiconductor Devices for Optical Communication (H. Kressel, ed,). Spring-Verlag, N.Y. 1980.

[14] K. Ikeda, S. Horiuchi, T. Tanaka, and W. Susaki, "Design Parameters of Frequency Response og GaAs-(Ga,Al)GaAs Double Heterostructure LED's for Optical Communication," *IEEE Trans. Elect. Dev.*, ED-24(7):1001-1005, July 1977.

[15] S. M. Sze, Physics of Semiconductor Devices. John Wiley and Sons, N.Y., 1981.

[16] G. P. Agrawal and N. K. Dutta, Semiconductor Lasers. Van Nostrand Reinhold, N.Y., 1993.

[17] O. Vendier, N. Jokerst, R. Leavitt, "High Efficiency Inverted MSM Photodetectors," *Proceedings of the Conference on Lasers and Electro-Optics (CLEO)*, Baltimore, June 1995.

[18] S. T. Wilkinson, N. M. Jokerst, and R. P. Leavitt "Resonant Cavity Enhanced Thin Film AlGaAs/GaAs/AlGaAs LEDs with Metal Mirrors," submitted to *Applied Optics*, May 1995.