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.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

Bell & Howell Information and Learning
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA

UMI®
800-521-0600

Packaging of surface active optoelectronic device arrays

David Kabal

Department of Electrical Engineering
McGill University
Montréal, Québec
Canada
October, 1997

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

Abstract

Requirements of electronic systems are increasing, as the applications for these systems become more complex and require more bandwidth. Optical interconnection can help satisfy these system requirements. In order for an optical interconnect to be successful, system design must satisfy new demands on the packaging of the optically interconnected devices, different from those encountered in packaging electronic interconnects.

This thesis addresses the issues faced in designing the packaging for a two-dimensional array of surface active optoelectronic devices, within the framework of a free-space optically interconnected backplane. Two demonstrator systems' packaging are presented, as well as a general discussion of the issues and constraints involved in the design, assembly and testing of this packaging. These issues include connectivity, bandwidth, thermal management, optomechanical interfacing and space constraints.

The particular optoelectronic devices examined are Quantum Confined Stark Effect (QCSE) modulators and detectors. The impact of temperature sensitivity on these particular devices is examined, and a technique for optimizing their performance at a specific temperature is presented. As a thermal diagnostic aid, as well as a thermal management tool, the design and characterization of an array of temperature sensors integrated into an optoelectronic device array is presented.

Résumé

En raison de la croissance de leur complexité et de leur besoins en largeur de bande, les exigences des systèmes électroniques deviennent de plus en plus sévères. Les interconnexions optiques promettent de répondre à ces besoins. Afin de réaliser une interconnexion optique, la conception du système doit répondre aux nouvelles contraintes de packaging de dispositifs reliés par l'optique qui sont différentes de celles pour le packaging purement électronique.

Cette thèse aborde les défis à relever au cours de la conception de packaging pour des matrices deux dimensionnelles de composantes optoélectroniques à surface active dans le cadre d'un backplane interconnecté à optique libre. Le packaging pour deux systèmes de démonstration est présenté. Il en suit une discussion sur les considérations et contraintes associées à la conception, au montage et à la vérification de ce packaging. Les considérations comprennent la connectivité, largeur de bande, gestion thermique, interface optomécanique et contraintes de volume.

Les dispositifs optoélectroniques étudiés sont des détecteurs et modulateurs à effet quantique confiné Stark. L'effet sur la performance du système de la sensibilité aux changements de température de ces dispositifs est étudié. Une technique pour optimiser leur performance à une température donnée est aussi mise de l'avant. La conception et l'étude des caractéristiques d'une matrice de senseurs de température intégrée dans une matrice de dispositifs optoélectroniques pour fins de diagnostic thermique est présentée.

Acknowledgements

I would like to thank the following people, without whom this research would not be possible;

Firstly, I would like to acknowledge Prof. David Plant, my supervisor. Without his invulnerable enthusiasm, invaluable guidance and irrefutably relentless search of funding this research could never have been realized.

Most of the packaging research for the Phase III demonstrator packaging system was conducted as part of a two-person team. Mike Ayliffe handled the optical issues of the packaging as I worked on the electrical, thermal and other aspects of the design. The inevitable area of overlap occured at the optomechanics, where the abstract design met reality. We are jointly reponsibly for this work, and our collaboration was essential to the realization of our designs. A more detailed description of the optomechanical and optical characterization of this work will be published in

due time. It is not simply for our collaboration that I must thank Mike Ayliffe, but also for the almost fantastic ease which changed a half-spoken idea into a deliverable part.

I would also like to thank David Rolston for designing the core and substance of both devices which this thesis talks about: the CMOS-QCSE Workshop Chip and the BiCMOS-QCSE Phase III chip. His advice, inspiration and friendship were invaluable to this thesis.

Mike Venditti is the gracious soul who put up with my relentless nagging and made sure I had all the data I needed on QCSE temperature dependence.

I also thank the following people for their experimentation, simulation, advice and inspiration: Pritha Khurana, Madeleine Mony, Prof. Frank Tooley, and Eric Bernier.

This research was partially funded by the Canadian Institute for Telecommunications Research (CITR) and NorTel.

Table of contents

Chapter 1: Introduction

1.1 Motivation

Requirements of electronic systems are increasing, as the applications for these systems become more complex and require more bandwidth. Electronic systems are now being used in more ways than ever before, and a premium is now placed on the networks that must interconnect these systems. Telecommunications networks must service the bandwidth and connectivity requirements of these electronic systems, evolving to satisfy the increasing needs of a information driven society.

The new performance parameters have changed the purview of the system designer so that the connectivity and speed *between* discrete components and subsystems is becoming a critical area of research, in addition to the performance of these individual components and subsystems. The interconnectivity between these elements is provided by their packaging, and a general packaging hierarchy distinguishes the interconnection levels. The physical layer of a system is composed

of a hierarchy of connectivities: die to die, module to module, board to board and finally subsystem to subsystem. In addition to the interconnections, the physical packaging layer must implement the electrical support and the mechanical structure, including thermal management and the mechanical base upon which a system is built.

As the architectural requirements of the interconnection in a system are expanded, the physical size must remain the same in order to maintain several performance metrics, including speed, latency, reliability, and others. Preservation of compact size ensures new system requirements, as the interconnection density grows higher and more electronics are packed even denser, running at higher speeds and dissipating more power. These requirements will enforce constraints at the backplane level of interconnection (board-to-board) [1], where the highest interconnection density of a system is acheived. The current technological solution, the electronic backplane, is failing to keep pace with the increasing demands placed upon it by increasingly high performance parallel computing and switching systems.

For these reasons, it is apparent that data communication through a photonic interconnect layer can be advantageous, by increasing the overall speed and connectivity of a backplane [2, 3]. For long, medium and short haul telecommunications networks, the jump to photonic interconnects has already been made, and installed fiber optic networks are in use. This jump has yielded tremondous gains in the field of optoelectronic packaging, where the tight

integration of optoelectronic, electronic and optical elements is performed, yielding integrated, commercial packaging solutions that are reliable [1]. These packaging solutions are maturing for optoelectronic/fibre or optoelectronic/lightguide integration.

1.1.1 Towards a photonic backplane

At the backplane level, or very short distance communication, photonic interconnects are an emerging technology, and research is starting to demonstrate the potential of this technology towards commercial development. Although guided wave structures can provide nearly lossless transmission for large distances, their applicability to backplanes is limited. These applications are limited by the inherent constraint of fibre and cladding size, limiting the total number of interconnections, and because of their point-to-point nature.

For an optical backplane, another approach has great potential: optically interconnected two-dimensional arrays of optoelectronic devices. The impetus for this development is the potential for bandwith, and particularly, for the large number of channels that the volumetric interconnect that a Free-Space Optically Interconnected (FSOI) Backplane can provide. Current hybrid optoelectronic devices have demonstrated these large number of optical data links operating at high speed [4]. With the availability of devices capable of supporting a volumetric backplane, the effective integration of such devices into an operational interconnect becomes the challenge [5].

Development of the optoelectronic devices that could support a photonic interconnect fabric continues to produce more advanced devices supporting more optical and electrical inputs and outputs, and their speed continues to increase. Hybrid integration of two-dimensional arrays of Quantum Confined Stark Effect modulators and detectors onto CMOS circuitry is a reality, and reliable VCSEL and detector arrays integrated on CMOS circuitry is in progress, with promising applications forthcoming.

The last hurdle in this development cycle is the effective integration of such a device into a system. Such integration will require the successful satisfaction of three fundamental requirements: optomechanical interfacing to the photonic interconnection fabric, electronic interfacing to the photonic interconnect fabric, and optical interconnect design. This work will present one aspect of the interfacing, the electronic and to a smaller extent the optomechanical interfacing required for the construction of an optical backplane.

1.1.2 Optoelectronic packaging

Current electronic packaging has advanced to the point where large chips can be interconnected in a hybrid manner, where the delineation between board and chip-level integration and interconnection is blurred. In such a packaging system, chips may be wirebonded or connected using tape bonding to either a module (such as an Multi-Chip-Module), a board (using chip-on-board packaging) or a chip carrier (such as a Ball-Grid-Array, Flatpack, Leaded or Leadless Chip Carrier). The clear

hierarchy of packaging electronic devices has been broken by necessity, as the overhead of the hierarchy proved limiting to either bandwidth or connectivity.

Borrowing from the field of electronic packaging, optoelectronic packaging has advanced steadily [6]. In this field, a smaller subset of practitioners work on the specific problem of packaging two-dimensional arrays of surface-active devices, to interface them to free-space optical interconnects and switching fabrics. Successful examples of integration for two-dimensional device arrays are starting to make the conceptual benefits of free-space optical interconnect into real demonstrators. Starting with standard off-the-shelf electronic packaging used as a basis for optoelectronic packaging, the demands of faster, more dense interconnects have brought more advanced electronic packaging, and packaging unique to optoelectronics, to the field.

1.1.3 Packaging of optoelectronics in optical backplanes

This thesis deals with the constraints, challenges and design methodology for packaging two-dimensional arrays of optoelectronic devices. Although the demonstrators presented in this thesis use Quantum Confined Stark Effect (QCSE) modulators and detectors, the approaches used can be applied to emitter based optical interconnections, and many of the goals are shared between the available optoelectronic technologies. These shared goals include high optical and electrical connectivity, thermal stability, and good optomechanical interfacing techniques.

1.2 Thesis organisation

This thesis will start with a general statement of the considerations and issues that are vital to the problem of packaging a two-dimensional optoelectronic device array, and discuss the specific application of these issues to the devices under consideration, MQW diodes forming Quantum Confined Stark Effect (QCSE) modulators and P-I-N detectors. Along with the general solutions that have been applied and are currently being applied by researchers in this field, chapter 2 discusses a specific procedure for the optimization of temperature performance for QCSE modulators and detectors. Chapter 3 refines these techniques by application, using a specific demonstrator system to reveal the research and progress in this area. This system was completed in 1996. Chapter 4 is the discussion and application of even more restrictive constraints to a more advanced packaging system that has been completed in 1997. The system demonstrator that this packaging system is integrated into will not be complete until early 1998. Chapter 5 deals with the design and characterization of an array of temperature sensors integrated onto the optoelectronic device whose packaging is discussed in chapter 4. Chapter 6 summarizes the designs, constraints and approaches to the packaging systems for two-dimensional arrays of optoelectronic devices and discusses their applicability in the future.

1.3 Contributions

The system demonstrators in this thesis were realized by the combined efforts of a team, the Photonic Systems Group at McGill University. This section tries to quantify the relative contributions of the participants, as well as giving credit where it is due.

In the temperature optimization technique presented in section 2.4.3, the (representative) data presented is the exclusive work of Michael Venditti, who set up the power/photocurrent/voltage/temperature (PIVT) characterization system, developed the computer software necessary to do the scans the author required, and performed the measurements on the optoelectronic devices.

For the system demonstrator packaging presented in chapter 3, the design, assembly and testing of the electrical and thermal aspects, as well as optomechanical characterization, were performed by the author. This included printed circuit design, chip attach assembly, gold-plating of heat spreader, wirebonding, thermal characterization, assembly of package and connector selection. All boards used to test the bandwidth (i.e., connected to the daughterboard) and electrical connectivity were designed by the author with the help of Michael Venditti. The assembly and wirebonding rig was designed by the author, David Rolston and Guillaume Boisset. The mechanics of the three point attach and optical constraints were supplied by Guillaume Boisset. The optoelectronic device was designed by David Rolston.

For the system demonstrator packaging presented in chapter 4, the packaging "team" consisted of the author and Michael Ayliffe. Most of this work is through our

combined design, characterization and testing efforts. Exceptionally, the flexible printed circuit was designed by the author, as well as the printed circuit test board. Michael Ayliffe's future writing will concentrate on the optomechanical interfacing and characterization of the package module, a subject which is notably and unavoidably absent from this work. The assembly and testing of this package module to date were performed by the author, Pritha Khurana and Michael Ayliffe.

The design, layout, and testing for the temperature sensors were performed by the author and Michael Ayliffe.

1.4 References

[1] Frank A.P. Tooley, "Challenges in Optically Interconnecting Electronics," *IEEE Journal of Selected Topics in Quantum Electronics*, April 1996, Vol. 2, No. 1, pp3–13

[2] H. Scott Hinton, *An Introduction to Photonic Switching Fabrics*, Plenum Press, 1993

[3] S. Yu and R. Forrest, "Implementations of smart pixels for optoelectronic processors and interconnection systems II: SEED-based technology and comparison with optoelectronic gates," *Journal of Lightwave Technology*, Vol. 11, XX, pp1670–1680

[4] Anthony L. Lentine, Keith W. Goossen, J.A. Walker, Leo M.F. Chirovsky, L. Arthur D'Asaro, S.P. Hui, B. J. Tseng, R. E. Leibenguth, J.E. Cunningham, W. Y. Jan, Jen-

Ming Kuo, D. W. Dahringer, D. P. Kossives, D.D. Bacon, Gabriela Livescu, R. L. Morrison, Robert A. Novotny, and D.B. Buchholz, "High-Speed Optoelectronic VLSI Switching Chip with >4000 Optical I/O Based on Flip_Chip Bonding of MQW Modulators and Detectors to Silicon CMOS," *IEEE Journal of Selected Topics in Quantum Electronics*, April 1996, Vol. 2, No. 1, pp77–83

[5] David V. Plant, "Constructing a Free-Space Optical Backplane: Challenges and Choices," Technical Digest of *Optics in Computing*, 1997, Vol. 8, pp156–158

[6] Alan R. Mickelson, Nagesh R. Basavanhally and Yung-Cheng Lee, editors, *Optoelectronic packaging*, John Wiley and Sons, Inc., 1997

Chapter 2: QCSE device packaging

2.1 Introduction

Providing an integrated packaging solution for a two-dimensional array of surface-active optical devices presents a series of constraints that complicate the employment of standard electronic packaging in systems applications. Through the construction of demonstrator systems, the Photonic Systems Group at McGill University (PSG McGill) has designed, fabricated, and implemented array packaging which uniquely addresses the critical issues associated with successfully integrating two-dimensional optoelectronic device arrays into systems. These system design considerations are: electrical bandwidth, electrical connectivity, thermal management, and optomechanical interfacing.

An overview of the packaging techniques that have been applied to system demonstrators by PSG McGill is presented in Table 2.1. In this chapter, an overview of the constraints inherent to this type of system are presented. In subsequent,

Table 2.1: PSG McGill Packaging

Packaging [ref.]	Year	Optoelectronic device technology	Chip I/O		Aggregate optical BW	TE Cooler	Thermal resistance θ_{jc}	Power
			Electrical	Optical				
QFP [1]	1994	FET-QCSE	40	32	2.5 GB/s		30 °C/W	~0.5 W
PGA [2]	1995	VCSEL/MSM	68	32	300 MB/s	✔	55 °C/W	~0.8 W
COB I [3]	1995	CMOS-QCSE	22	16	8 MB/s		~50 °C/W	~0.15 W
COB II [4]	1996	CMOS-QCSE	44	64	16 MB/s	✔	7 °C/W	0.5 W
COB III	1997	BiCMOS-QCSE	189	1024	81.9 GB/s	✔	~3 °C/W	~5 W

chapters, this overview is refined by a discussion of the packaging techniques that were used in two separate systems (in table 2.1, COB II and III) that were implemented by the Photonic Systems Group. This packaging was designed, modelled, fabricated and demonstrated for board-to-board and backplane optical interconnect applications.

2.1.1 Optoelectronic devices

When an optical communication system uses Quantum Confined Stark Effect (QCSE) modulators in reflection mode for transmitting and QCSE detectors for receiving light signals, a unique set of design constraints are imposed on the packaging for that system. In such a system, optical power supply (light) sources are incident on the surface active array. The surface active array introduces a modulation on the power supply beams by altering the absorption of light in the QCSE modulator via an electrical signal applied by transmitter circuitry. This reflected optical signal is directed onto the subsequent array of detectors, which are P-I-N (*P*-doped region, *I*ntrinsic region, *N*-doped region) diodes. The detector translates the received optical signal into a photocurrent, which is amplified and thresholded back to digital (voltage) logic levels by receiver circuitry.

In both of the systems whose packaging is presented in the following chapters, dual rail optical signals were used to represent a single bit of data. With a representative reflectivity difference between high and low states of less than 50% [5], it is desirable to send complementary data, so that the *difference* between two reflected signals can be thresholded by receiver logic. Thus, for these systems, two power

supply beams are reflected by two modulators coupled by their driver logic to produce complementary outputs, and the resulting two optical signals are incident on two detectors where the difference between the resulting photocurrents can be turned into a digital logic level. This arrangement is shown in figure 2.1.

Figure 2.1: Dual rail optical interconnect

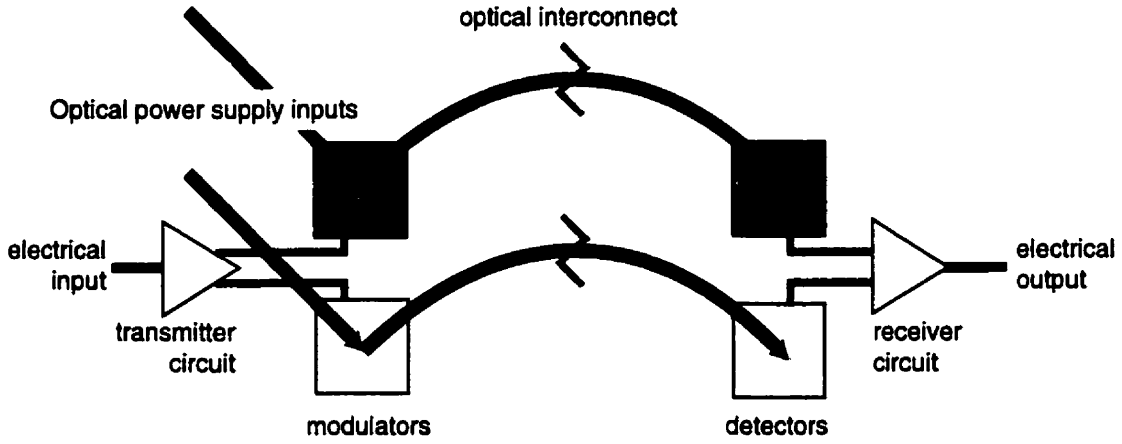

2.1.2 Optical backplane

For an optical interconnect, an array of these QCSE devices can be hybridly integrated onto the surface of a standard CMOS or BiCMOS chip, using a flip-chipping technique that connects the Silicon substrate of the CMOS or BiCMOS device to the Gallium Arsenide substrate of the QCSE devices, at specific bump-bonding points.

The hybrid device thus integrated can be used as an optical communication node, thereby exploiting the two-dimensional, volumetric nature of the signals as depicted in figure 2.2. This figure ignores the details of the optical interconnection scheme used, but shows a representative application of such an interconnect: to connect printed circuit boards. The details of this optical interconnect and its implications on

Figure 2.2: Optical backplane

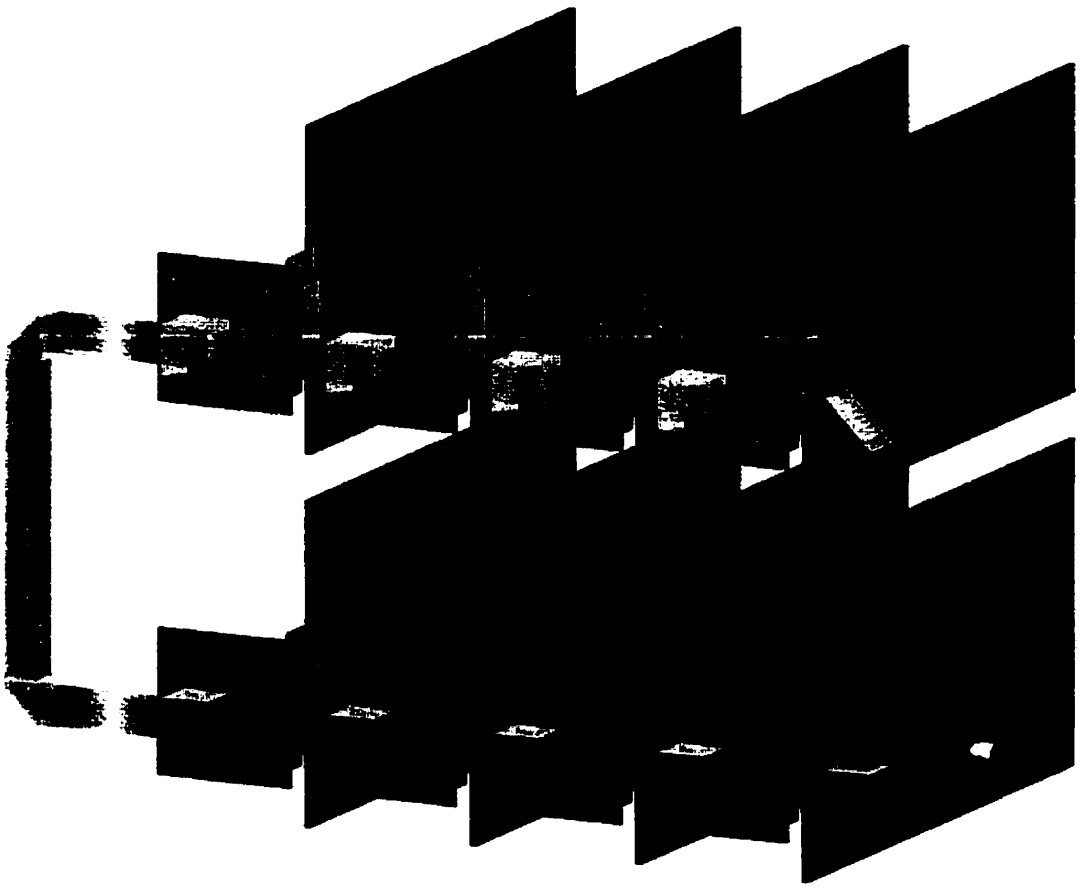

the constraints of the packaging system are discussed in section 2.3. The remaining constraints on the packaging fall from the device's thermal dependence, in section 2.4 and the specifics of the electrical connectivity and signal integrity, in the following section (2.2).

2.2 Electrical constraints

For the system demonstrators in development by the Photonic Systems Group of McGill University, optical nodes in the backplane support the aggregrate bandwidth of the entire backplane, and the individual optoelectronic device arrays are required

to inject or extract data from this backplane, in what has been referred to as a firehose architecture [6]. This architecture is illustrated in figure 2.3.

Figure 2.3: Firehose architecture

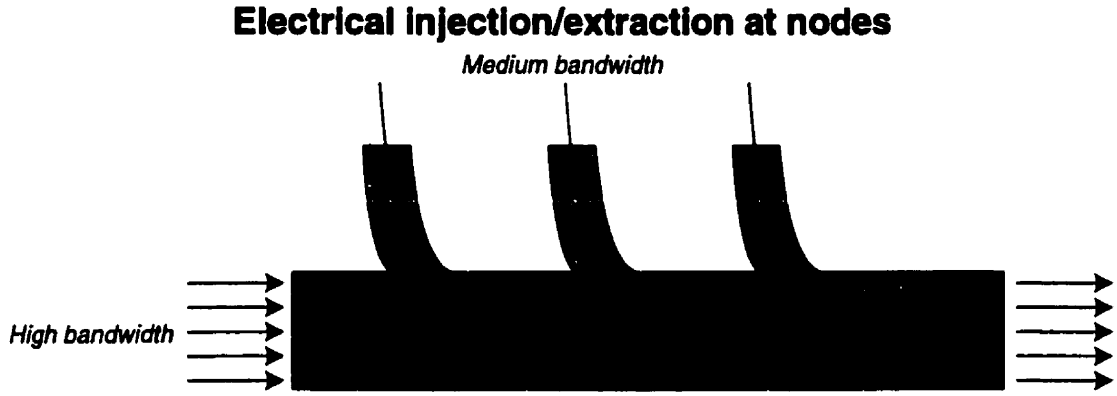

With the firehose architecture, each optoelectronic device is required to support only a medium electronic bandwidth but the entire, high optical bandwidth. This medium electronic bandwidth can still be substantial, and consists, in the Photonic Systems Group incarnation of the optical backplane, of many electrical inputs and outputs, running at speeds of up to 100 Mbit/s. Electrical signal integrity is a serious concern at this speed, and the packaging system must directly address this issue, with impedance control, high-speed, dense connectorization, and adequate termination of electrical signals.

A standard rule-of-thumb that is used to separate low-speed from high-speed packaging problems in the electronics world is the *l*/6 rule [7]. This rule indicates the demarcation between the transmission line domain and lumped model domain. For example, a 100 MHz digital signal with a rise time or fall time of 1000 ps (1/10

the cycle width) travelling along a polyimide board with a delay per unit length of 180 ps/in (7 ps/mm) has a rising edge length of:

$$l = \frac{T_R}{D} = \frac{1000\ \text{ps}}{\left(\frac{7\ \text{ps}}{1\ \text{mm}}\right)} = 143\ \text{mm} \qquad (2.1)$$

where D is the delay per unit length, and T_R is the rise time. For this example, the demarcation between lumped and transmission line models is at 24 mm ($l/6$). Any trace longer than 24 mm should be treated as if it were a transmission line. For most of the examples presented here, the circuits are significantly longer than the validity that a lumped model permits, requiring a transmission line treatment. In this treatment, termination is necessary to prevent reflections. In the packaging systems presented, the inputs to the optoelectronic device are CMOS inputs, presenting a high impedance to the transmission line, necessitating a termination to prevent an amplitude doubling reflection. Various termination schemes are used, but dominant categories are active Schottky (diode) termination and passive (resistor or resistor and capacitor) termination. Both are present in the Phase III system presented in chapter 4.

2.3 Optical and optomechanical constraints

For a packaging system involving an optoelectronic device to be successful, it must interface the optoelectronic device effectively to the optical interconnect. The optical interconnect design, and restrictions on focal length, will generally keep the optoelectronic device in close proximity to optics, forcing tight constraints for

volumes and footprints in the packaging system. The interface itself is clearly also an issue, with the ability to align the optoelectronic device to the optical interconnect a primary concern.

The specific constraints that optomechanical and optical issues form will be introduced in each specific section as they apply to a system demonstrator in chapters 3 and 4.

2.4 Thermal issues

This section deals with the thermal issues inherent in a system that uses QCSE modulators and detectors. These devices are temperature sensitive, and a lack of control over each device's temperature could affect system performance. A stepwise optimization procedure to optimize the performance of modulators and detectors is introduced in section 2.4.3.

2.4.3 Temperature optimization of device performance

Introduction

For a photonic interconnect that utilizes QCSE modulators and detectors in the optical data link, specific optimizations can be performed with regard to the temperature sensitivity. The following section will deal with the temperature dependencies of these devices and suggest an optimization technique.

To introduce the optimization procedure, representative graphs measured from real QCSE modulators and detectors are used. For the data used, the detectors and modulators are identical devices. In a system, these devices would be operated

under different electrical and optical conditions to control their mode of operation. The optimization technique presented here does not preclude the use of separately designed and fabricated devices for modulators and detectors [5].

The modulators and detectors were both Multiple Quantum Well (MQW) diodes created in Gallium Arsenide by researchers at Nortel Technology (wafer growth) and École Polytechnique de Montréal (processing), and measured using an optical/ electrical setup with wavelength, voltage and temperature sweeping capabilities. At each voltage, temperature and wavelength, the output photocurrent of the diode could be measured.

The setup consisted of a monochromator (*Oriel Instruments* MS 257), a DSP lock-in amplifier (*Stanford Research Systems* SR830), temperature controller (*ILX* LDC-3722B), voltage source (*Keithley* 230), optical chopper and microprobes. This setup allowed very small photocurrents to be sampled by probing the active device with microprobes. The photocurrents are isolated from background noise using the optical chopper and lock-in amplifier.

The GaAs devices are tested prior to their hybrid integration onto a CMOS or BiCMOS chip. To do this, the devices are temporarily mounted in a Pin-Grid-Array (PGA) chip carrier cavity alongside a chip thermistor (*Betatherm* 10K3CG2). To control the temperature, a Peltier device (*Marlow Industries* MI1060T) is operated in reverse mode to raise the temperature of the cavity, under the control of the temperature controller. The PGA is mounted on a custom-designed mounting printed circuit, and the monochromator delivers a given wavelength so that the

resulting photocurrent can be measured. The measurement system is under the control of a single Macintosh computer running (*National Instruments*) LabView software.

The wavelength dependency of the quantum efficiency, η, of a detector is an important metric, and the temperature shift of this efficiency versus wavelength is an oft-quoted number; approximately 0.3 nm/°C [5]. In a communications system, specific optimization is needed in order to maximize performance for modulators and detectors at a given wavelength, the system's wavelength, around which all the optical elements are designed. The behaviour of the exciton curve for such a device at wavelengths *other than the system wavelength* may not be relevant, in this case.

For detectors, the absorption should be maximized for this wavelength, at the operating temperature, for a given bias voltage. For modulators driven by circuit logic, the reflectivity difference should be maximized at this wavelength, at the operating temperature, using two parameters: the bias and the available logic swing.

These dependencies can be summarized as:

$$absorption = 1 - R_{detector}(\lambda, T, V_{bias}) \tag{2.2}$$

where R is the reflectivity,λ is the wavelength, T is the temperature, V_{bias} is the reverse bias voltage.

$$reflectivity\ difference = \Delta R_{modulator}(\lambda, T, V_{bias}, \Delta V) \tag{2.3}$$

where λ is the wavelength, T is the temperature, V_{bias} is the reverse prebias voltage and ΔV is the available logic swing.

In a system, the wavelength is fixed (λ = constant), and frequently the available logic swing is fixed (ΔV = constant). If the absorption and the reflectivity difference are independently optimized (maximized) to find the ideal V_{bias} for each, at any given temperature, the dependencies become:

$$1-R_{detector}(T) \tag{2.4}$$

$$\Delta R_{modulator}(T) \tag{2.5}$$

In the following graphs illustrating this optimization, ΔV was set to the logic swing of our system, 5 V, λ was set to the operating wavelength of our system, 852.0 nm. The resulting graphs are shown in figures 2.4 and 2.5. For these values, the

Figure 2.4: Modulator reflectivity difference

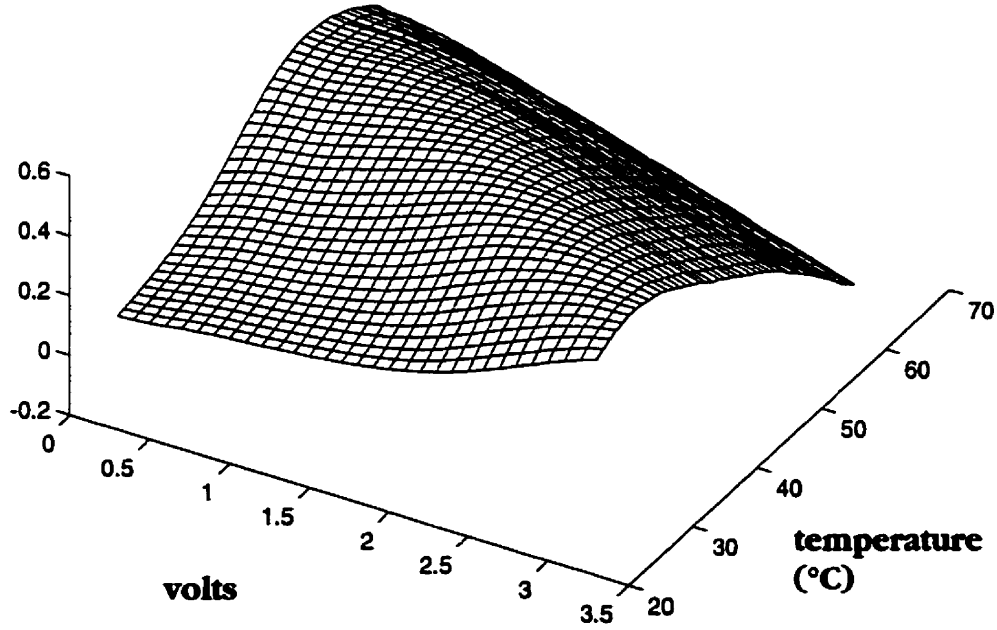

Figure 2.5: Detector (1-*R*)

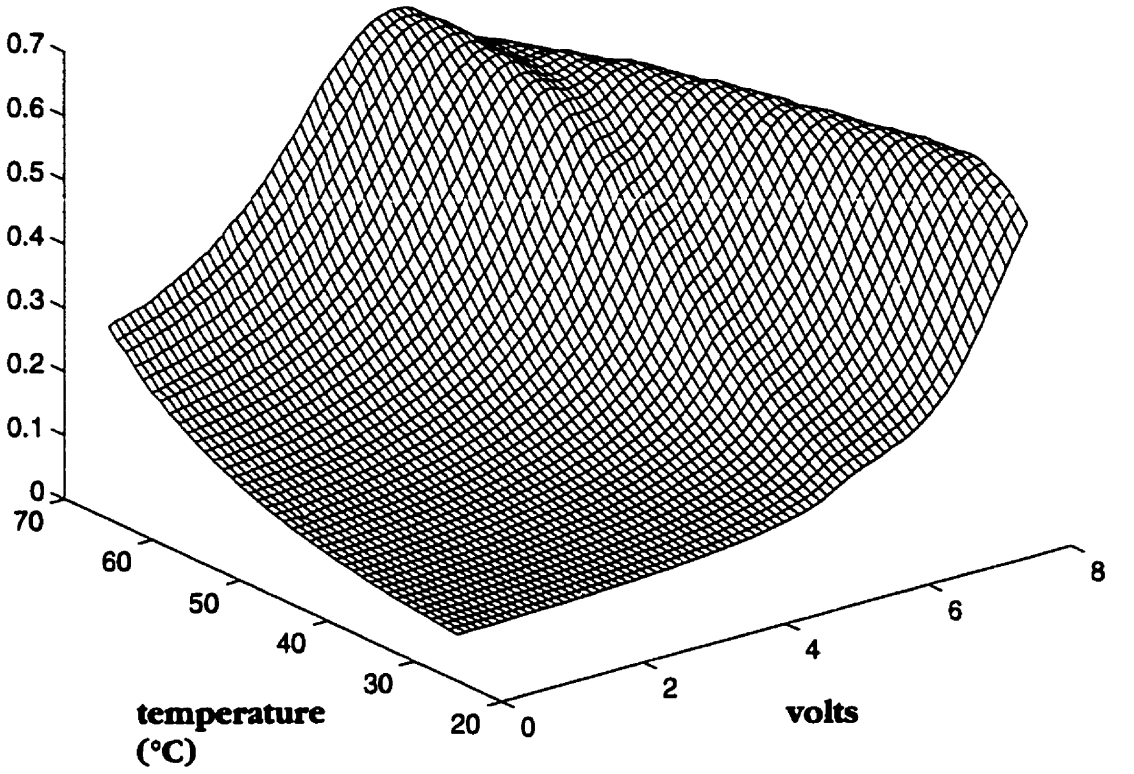

reflectance of the back mirror of the QCSE device (gold) is assumed to be 91.2%, so that the difference between input and reflected optical power can be used to develop a reflectance for the device.

Using an existing metric [5], equations 2.4 and 2.5 can be combined by multiplication, with the critical assumption that modulator and detector must operate at the same temperature. This assumption is reasonable given a standard configuration where modulators and detectors are coplanar in an optical

communication device on the same underlying electronics. The resulting equation is:

$$combined\ efficiency = \eta_B(T) = (1\text{-}R_{detector}(T)) \bullet \Delta R_{modulator}(T) \qquad (2.1)$$

If both the modulators and detectors performance are maximized at each temperature by picking the optimum prebias and bias voltage, the resulting graph is shown in figure 2.6. The peak of the η_B defines the optimal operating temperature

Figure 2.6: Combined efficiency versus temperature (maximized)

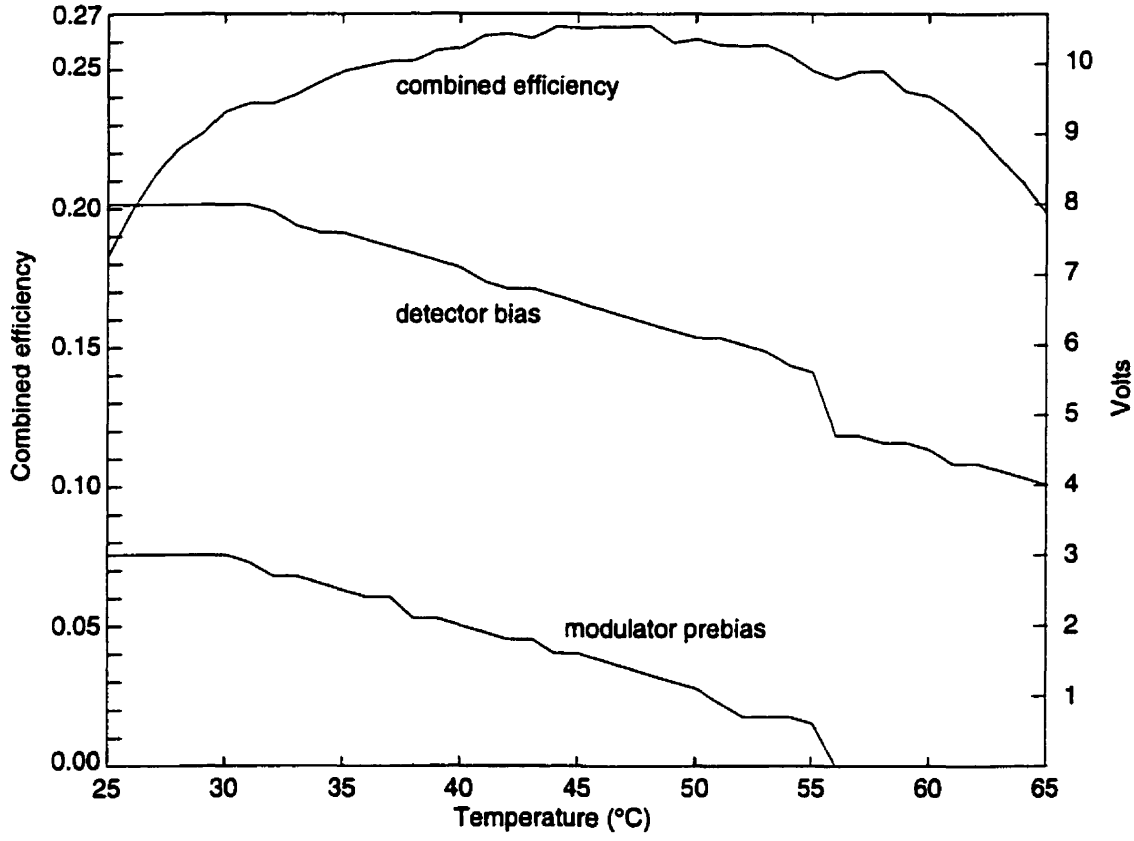

of a system which uses these modulators and detectors, where a read beam is incident on a modulator, reflected, and the resulting modulated beam is incident on a detector. This graph indicates the temperature dependency of performance for

this optical data link. Additionally, the optimized values for the voltage for prebias for modulators and bias for detectors at any given temperature can be found on the scale on the right side of this graph.

Optimization procedure

The steps used to produce the graph shown in figure 2.6 can be summarized as follows:

1. At the operating wavelength of the system, using the optical devices, measure the photocurrent produced by a quantum well diode (to be used as either a modulator or detector) across a range of voltages and temperatures. The range of voltages should include both the prebias value for the modulator, $V_{prebias}$, and $V_{prebias}+\Delta V$, the maximum voltage that the system is expected to use. For separately fabricated modulators and detectors, two measurements must be taken
2. From this data, it is possible to derive the graphs of: a) absorption versus voltage and temperature and b) reflectivity difference versus voltage and temperature. The reflectivity difference graph assumes a system logic swing, ΔV.
3. Maximizing both absorption and reflectivity difference for a given temperature and multiplying these results yields the graph of combined efficiency versus temperature.

From this optimization procedure, it is possible to determine the optimal operating temperature of the device, and the resulting V_{bias} and $V_{prebias}$ at this temperature for detectors and modulators, respectively. Using thermal management techniques

with a feedback loop for temperature control it is possible to set the device temperature accurately, as will be shown in the following chapters.

If it is desirable to operate the devices at a temperature other than this optimal temperature, then the resulting V_{bias} and $V_{prebias}$ for another temperature can be determined from the optimization procedure. This is useful if, for example, thermal stabilization is possible using available packaging techniques but thermal control to an *optimal* temperature is not possible.

2.4.4 Thermal management

Once it has been established what the desired operating conditions of the device are, packaging must enforce these conditions. Accurate thermal modelling using finite element analysis (FEA) and computational fluid dynamics (CFD) software can help predict the thermal properties of a proposed package, and determine what elements this thermal management must incorporate.

Thermoelectric coolers (figure 2.7) using the Peltier effect have been used in the

Figure 2.7: Thermoelectric cooler

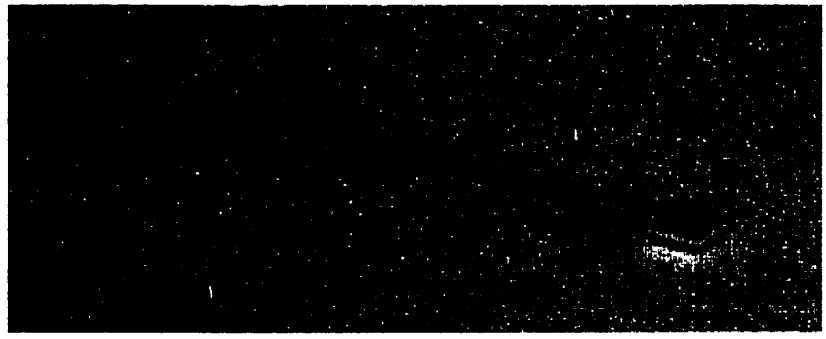

optoelectronics industry to control both CCD and laser diode temperatures [for

example, 6], and more recently these temperature control elements have been used by researchers for two-dimensional arrays of surface-active devices [9]. A thermoelectric cooler uses thermocouples to set up a heat pump between its two surfaces [10], creating a hot and cold side. In general, this is used to cool a device in thermally conductive contact with the cold side that would, under natural convection, operate at a higher temperature. With current flowing in the reverse direction, the symmetric thermoelectric cooler pumps heat towards the device. With feedback from temperature sensors, this bidirectional nature can be used to accurately control the device temperature. Many commercial instruments are available to complete the control loop. These instruments use a single temperature sensor input with an internal control loop to set the output current magnitude and direction for a thermoelectric cooler [11]. The IDT 5910 temperature controller from ILX is a representative instrument that can accept temperature sensor input from a thermistor, a current output temperature sensor (the AD590) or a voltage output temperature sensor (the LM335).

For packaging systems that will be presented in subsequent chapters, this temperature controller (or a similar device) was used in a control loop with a thermoelectric cooler. In chapter 3, the CMOS-QCSE packaging used a sensor input from a thermistor chip mounted on a surface in immediate contact with the device side of the thermoelectric cooler. In chapter 4, BiCMOS-QCSE packaging used a simulated LM335 input that was generated by a computer scanning and then averaging an array of temperature sensors integrated into the electronics on the top side of the surface-active device array.

A serious concern in the design of such a temperature management system is the uniformity of device temperatures across a large array. It can be expected that convective and conductive heat transfer, dependent on package design, as well as nonuniform electrical heat dissipation will introduce nonuniformities in temperature on the device surface. It is not feasible to individually control the device temperatures in an array in the same active manner that a thermoelectric cooler can control the entire die temperature. Without control of the individual device temperatures, a nonoptimized package design will generate varying (nonoptimal) performance of devices across the surface, dependent on location.

Using thermal analysis software coupled with experimentation, it is possible to design the package structure so that these nonuniformities are made small and perhaps negligible. Such a method has been used in the past for electronic devices by James Sweet [12], where an array of temperature sensors was designed onto a test chip and incorporated into a mockup of the packaging system, so that thermal nonuniformities could be predicted and the thermal model in FEA software validated.

For optoelectronic devices, Nagesh Basavanhally has used FEA software to model a device array and to predict package design thermal performance with respect to temperature nonuniformities [9]. With software, the machining of the package design was optimized to minimize these nonuniformities.

For the system presented in chapter 3, no regard was paid to nonuniformities caused by package design. The package was simply athermalized, with a heat

spreader (pedestal) larger than the die assumed to produce temperature unformity over the 2 x 2 mm die. This design assumption was not verifiable using laboratory equipment that was available.

This assumption was removed in the system presented in chapter 4. This packaging system used a combination of CFD thermal analysis and an array of temperature sensors to verify the minimization of package temperature nonuniformities. In this system, an athermalized pedestal was machined in the mount, connected to the rest of the mount only by narrow, thermally resistive joints. This pedestal shared its dimensions with those of the die (8 x 8 mm). The success of this technique, using software to predict thermal performance, allowing this prediction to be corrected by thermal sensors, and then redesigning the package to compensate, compels the author to propose that it be continued in future systems, as an iterative technique to optimize package design.

2.5 Manufacturing and reliability

A packaging approach must satisfy additional constraints in order to make the transition between a realizable package and a reliable, manufacturable package. This thesis deals exclusively with the task and theory of realizing a packaging approach. The issues of reliability and manufacturability, although important, are not addressed here.

2.6 Conclusions

The design issues inherent to a system using Quantum Confined Stark Effect modulators and detectors for communication in an optical backplane have been presented, indicating the specific design issues of thermal management, electrical connectivity and the particular optomechanical constraints of an interconnect architecture. Specific application of a temperature optimization technique has been developped, allowing the optimum temperature of a QCSE device to be determined. Subsequent chapters will demonstrate and address these design constraints for demonstrator systems.

2.7 References

[1] David V. Plant, Alain Z. Shang, Marcos R. Otazo, David R. Rolston, Brian Robertson, H. Scott Hinton, "Design, Modelling and Characterization of FET-SEED Smart Pixel Transceiver Arrays for Optical Backplanes," *IEEE Journal of Quantum Electronics*, August 1996, Vol. 32, No. 8, pp1391–1398.

[2] David V. Plant, Brian Robertson, H.Scott Hinton, W.M. Robertson, M.H. Ayliffe, Guillaume C. Boisset, Wayne S. Hsiao, David Kabal, Nam-Hyong Kim, Yongsheng S. Liu, Marcos R. Otazo, and Alain Z. Shang, "A 4x4 VCSEL/MSM Optical Backplane Demonstrator System," *Applied Optics*, 1996

[3] David R. Rolston, David V. Plant, Ted H. Szymanski, H. Scott Hinton, Wayne S. Hsiao, Michael H. Ayliffe, David Kabal, Michael B. Venditti, P. Desai, Ashok V. Krishnamoorthy, Keith W. Goossen, J.A. Walker, B. Tseng, S.P. Hui, J.E. Cun-

ningham, and W.Y. Jan, "A Hybrid-SEED Smart Pixel Array for a Four-Stage Intelligent Optical Backplane Demonstrator," *IEEE Journal of Selected Topics in Quantum Electronics*, April 1996, Vol. 2, No. 1, pp97–105.

[4] D. Kabal, G. Boisset, D.R. Rolston, D.V. Plant, "Packaging of two-dimensional smart pixel arrays," Presented at IEEE Summer Topical Meeting on Smart Pixels, Keystone, Colorado, Aug. 1996.

[5] David T. Neilson, "Optimization and Tolerance Analysis of QCSE Modulators and Detectors," *IEEE Journal of Quantum Electronics*, July 1997, Vol. 33, No. 7, pp1094–1103

[6] Ashok V. Krishnamoorthy, D.A.B. Miller, "Firehose Architectures for Free-Space Optically-Interconnected VLSI Circuits," *Journal of Parallel and Distributed Computing*, November 1996

[7] Howard W. Johnson, Martin Graham, *High-Speed Digital Design: A Handbook of Black Magic*, Prentice-Hall, 1993

[8] Alan R. Mickelson, Nagesh R. Basavanhally and Yung-Cheng Lee, editors, *Optoelectronic packaging*, John Wiley and Sons, Inc., 1997

[9] Nagesh R. Basavanhally, Michael F. Brady, D. Bruce Buchholz, "Optoelectronic packaging of two-dimensional surface active devices," *IEEE Transaction on Components, Packaging and Manufacturing Technology — Part B*, Feb. 1996, Vol. 19, No. 1, pp107–115.

[10] *Frigichip Thermoelectric Cooling Devices*, Thermoelectric handbook, Melcor Thermoelectrics, 1995

[11]LDT 5910B Instruction manual, Thermoelectric Temperature Controller, ILX Lightwave, June 20, 1997

[12]James N. Sweet, David W. Peterson, Dahwey Chu, Bruce L. Bainbridge, Richard A. Gassman, Cathy A. Reber, "Analysis and measurement of Thermal Resistance in a 3-Dimensional Silicon Multichip Module Populated with Assembly Test Chips," *Proceedings of Ninth Annual IEEE SEMI-THERM Symposium*, November 1993.

Chapter 3: CMOS-QCSE packaging

3.1 Introduction

This chapter describes the design issues and constraints that formed the basis of a four board optical interconnect implemented by the Photonic Systems Group at McGill University. The packaging for this system had to take into account various constraints arising from the use of a two-dimensional array of surface-active devices: the QCSE modulators and detectors. The optoelectronic chip that was packaged in this system was a CMOS-QCSE device that required a GaAs substrate containing the QCSE devices, which was flip-chipped onto a CMOS die resulting in a hybridized chip of dimensions, 2 mm x 2 mm.

3.2 Chip-on-board packaging

When the use of an explicit chip-carrier becomes impossible due to system constraints, an alternate approach may be attempted. If a system has thermal

requirements, I/O counts and volume constraints that prohibit the use of standard packages, the optoelectronic device array may be mounted directly on the board, employing the alternative approach of chip-on-board (COB) or MCM-L mounting [1]. To accomplish this packaging task for a CMOS-QCSE optoelectronic device array, the following packaging hierarchy and constraints were proposed:

1. The 2-D optoelectronic device array should reside on a small printed circuit *daughterboard* (mount) that must satisfy optical design constraints. These constraints limits the maximum board area to be used by the packaging to 65.00 by 46.05 mm. The first optical surface needs to be 800 μm from the die surface. Via thermal management, the chip's surface should be maintained at a constant temperature to optimize the responsivity of the optoelectronics.
2. This daughterboard should be connected to the next level of hierarchy by a high-speed connection that provides at least 40 I/O pins at 100 Mb/s and can mechanically isolate the daughterboard from other electrical components.
3. A motherboard should hold the processing electronics for the optoelectronic device array, as well as connectivity to the outside world and any application-specific electronics.

The fabricated daughterboard for COB II is shown in figure 3.2 An illustration of this packaging is shown in figure 3.3. The dimensions of the board are 63.88 mm x 41.80 mm. Including a thermoelectric cooler and heatsink, the thickness of the packaging is 40 mm. These dimensions satisfy the optical design constraints imposed by the system. This board follows an impedance controlled 4-layer FR-4

Figure 3.1: Packaging solutions

a) COB I

b) COB II

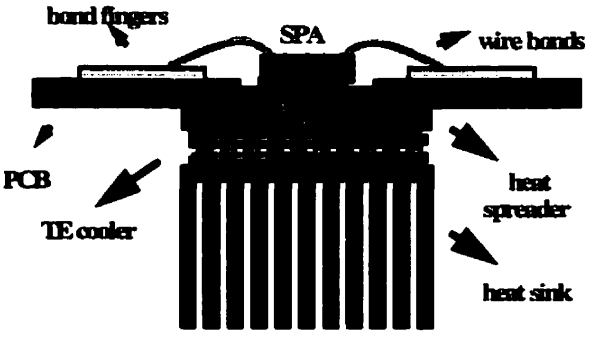

c) COB III

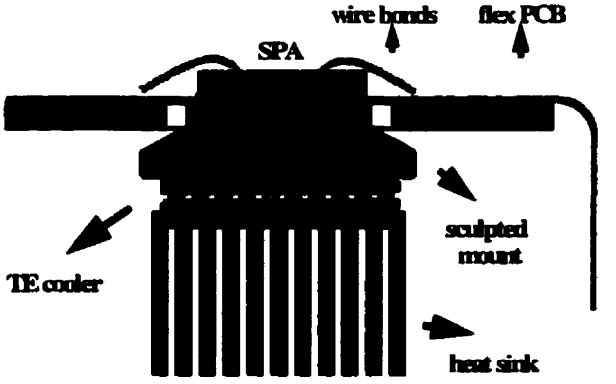

microstrip design with a characteristic impedance of 50 ¾ (figure 3.4). The impedance is set by fixing the dielectric thickness, h, (0.006"), trace width, w, (0.008") and trace thickness, t, (0.0015"):

$$Z_0 = \frac{87}{\sqrt{\varepsilon_R + 1.41}} \ln\left(\frac{5.98h}{0.8w + t}\right) \text{ (from [2])} \qquad (3.1)$$

Figure 3.2: Completed *daughterboard* chip-on-board packaging

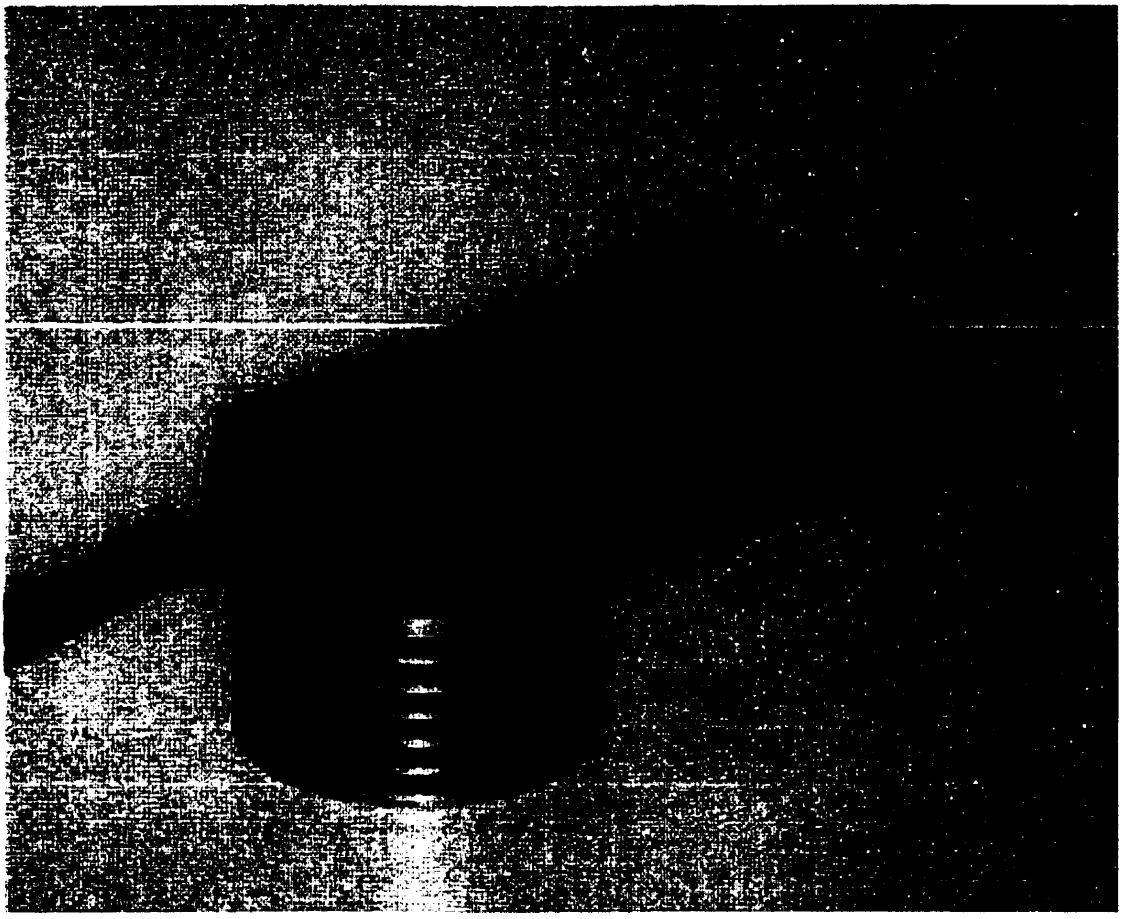

Figure 3.4: Microstrip stackup

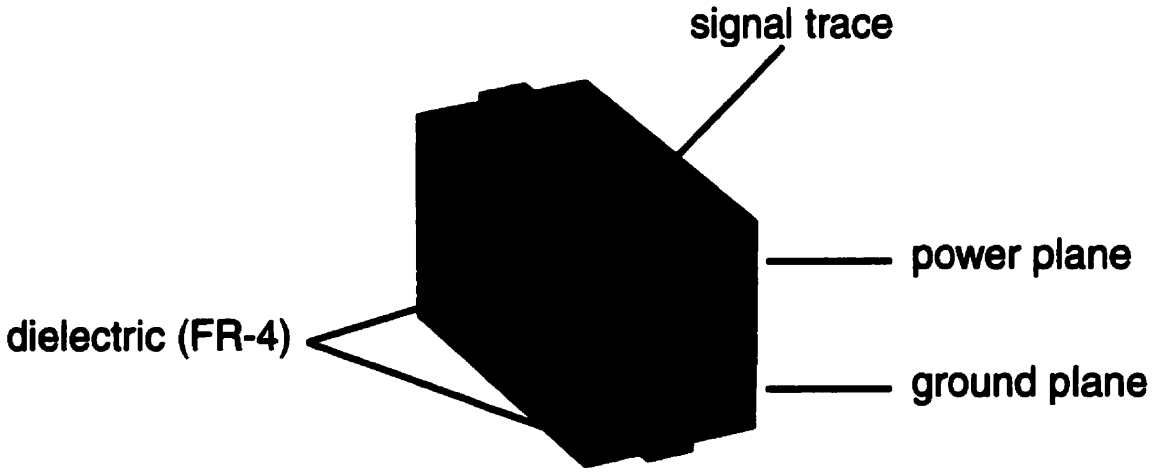

Figure 3.3: Daughterboard packaging assembly

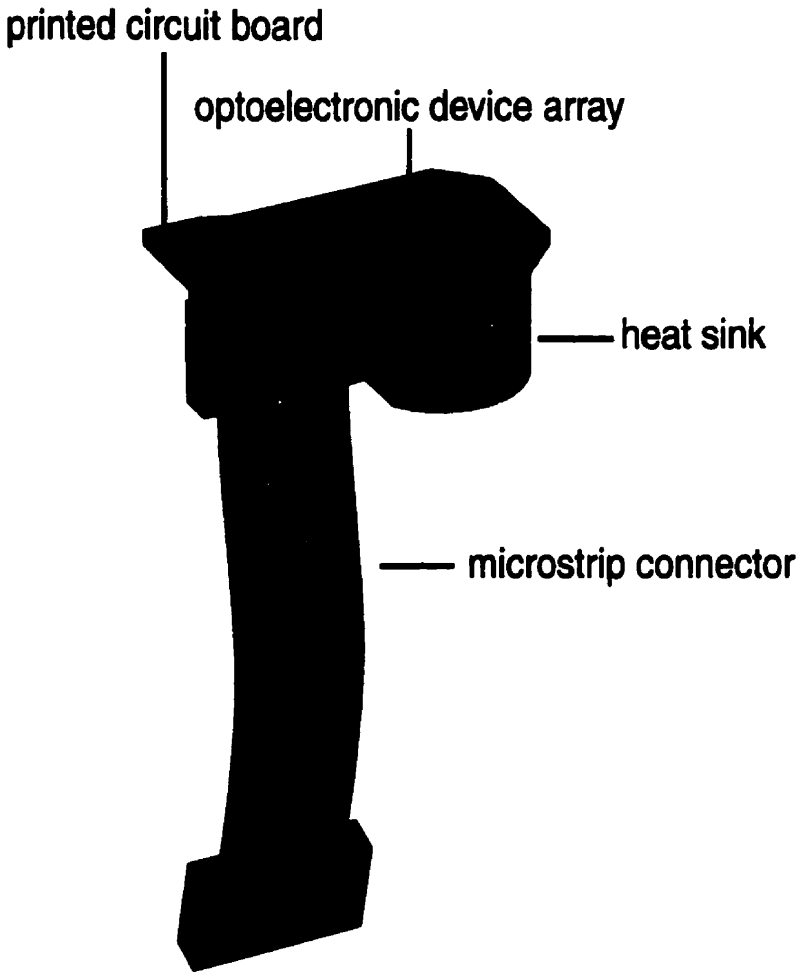

Inner layers in the daughterboard are power planes, while outer layers are reserved for signal lines.

The total board thickness is 0.092", allowing for a more rigid board than the standard 0.062", which was desirable for optical stability in alignment. The CMOS-QCSE optoelectronic device array was wirebonded to 44 bond fingers on the immersion-gold-plated daughterboard.

Positional tolerances were mandated by optical design, necessitating a custom assembly procedure using a custom alignment, wire-bonding and gluing rig. This rig is shown in figure 3.5. Application of this procedure to multiple daughterboards

Figure 3.5: Custom alignment, wire-bonding and gluing rig

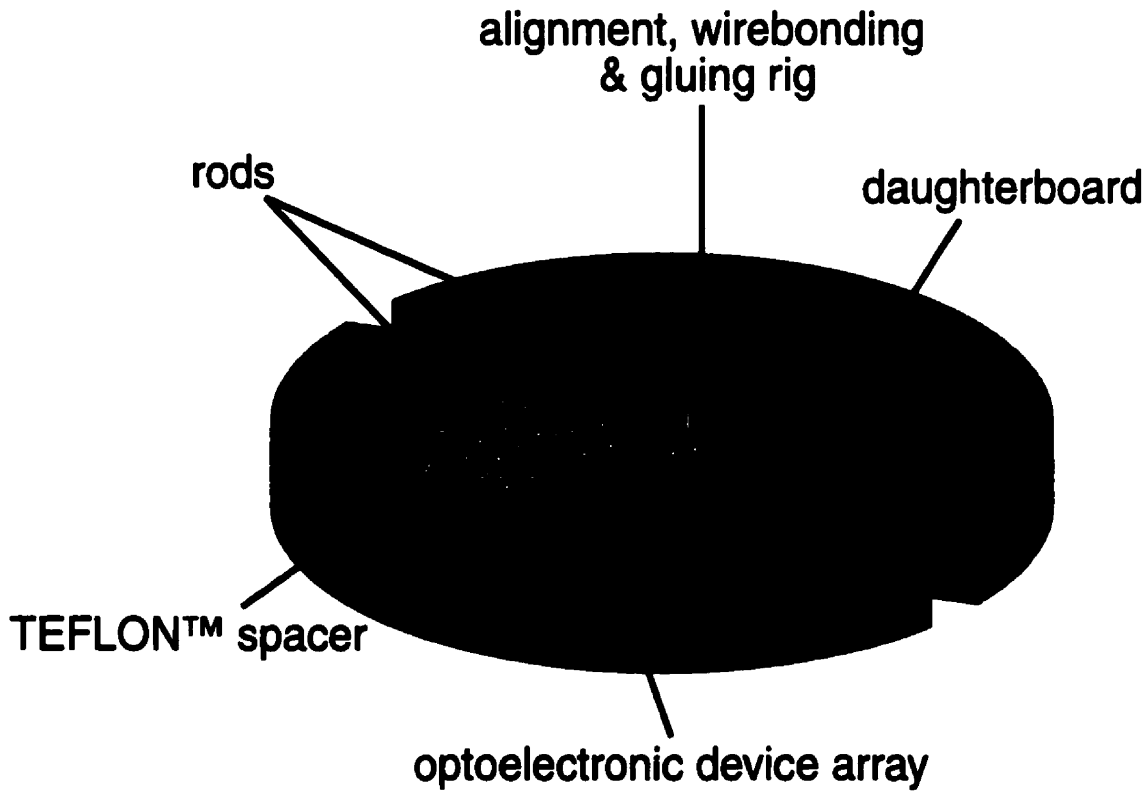

using the rig consistently acheived a lateral alignment tolerance of less than 150 μm (in chip's plane) to a 50 μm measurement resolution. The lateral alignment tolerance data are shown in figure 3.6 for the six fabricated circuit boards, where an HP ScanJet IIc scanner with a input resolution of 600 dots/per/inch was used to measure the alignment precision relative to the printed circuit board. Using microscope depth travel, with a repeatability of 10 μm on the chip corners, the

Figure 3.6: Lateral misalignment for six assembled daughterboards

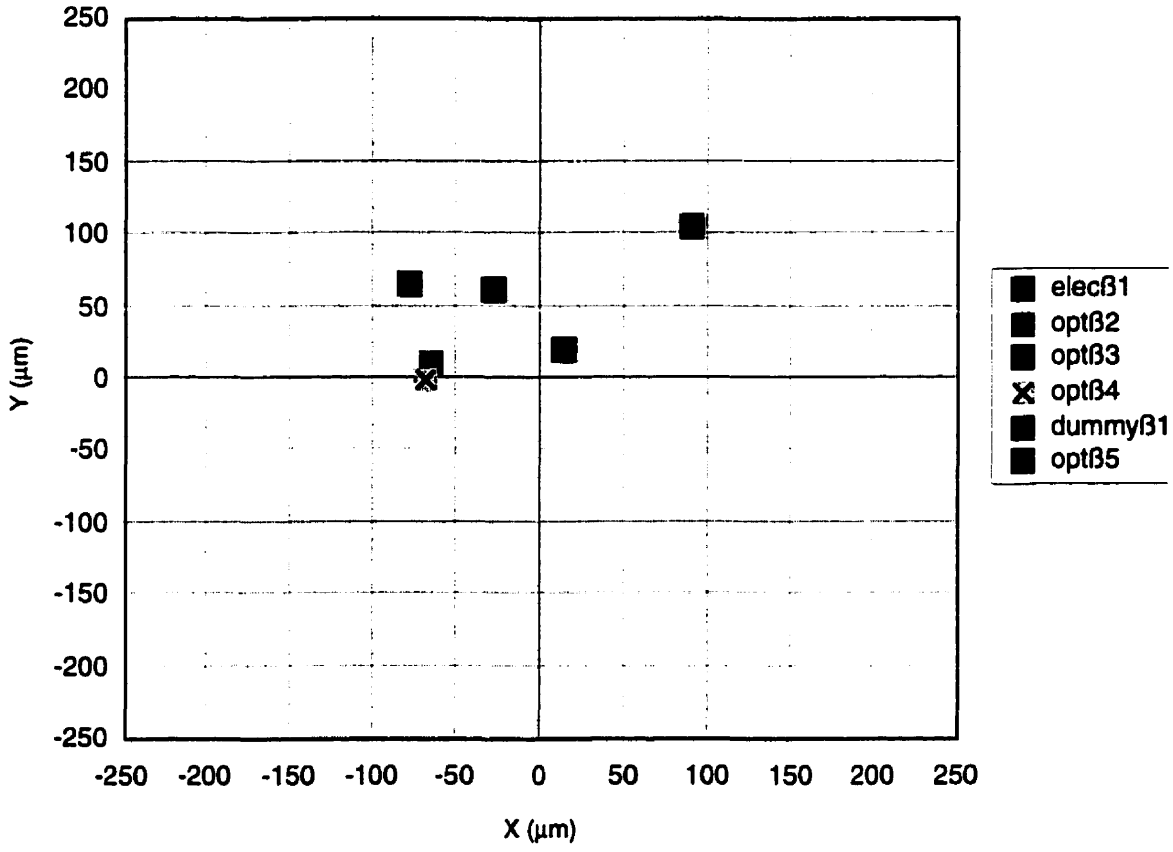

packaging approach generated a die tilt of less than 2.5° (board plane to chip surface plane), these data are shown in figure 3.7.

This daughterboard supports a single high-speed microstrip connector for connection to the motherboard via a flexible, impedance-controlled (50 ¾) ribbon. In this way, optical alignment of the photonic interconnect fabric is decoupled from the support electronics, while a reasonably high connectivity (40 pins) is maintained between motherboard and daughterboard at a bandwidth of 1 GHz per connection [3]. The measured -3dB frequency of the entire link between motherboard and daughterboard, including the ribbon cable, was 461 MHz (figure

Figure 3.7: Tilt measurement

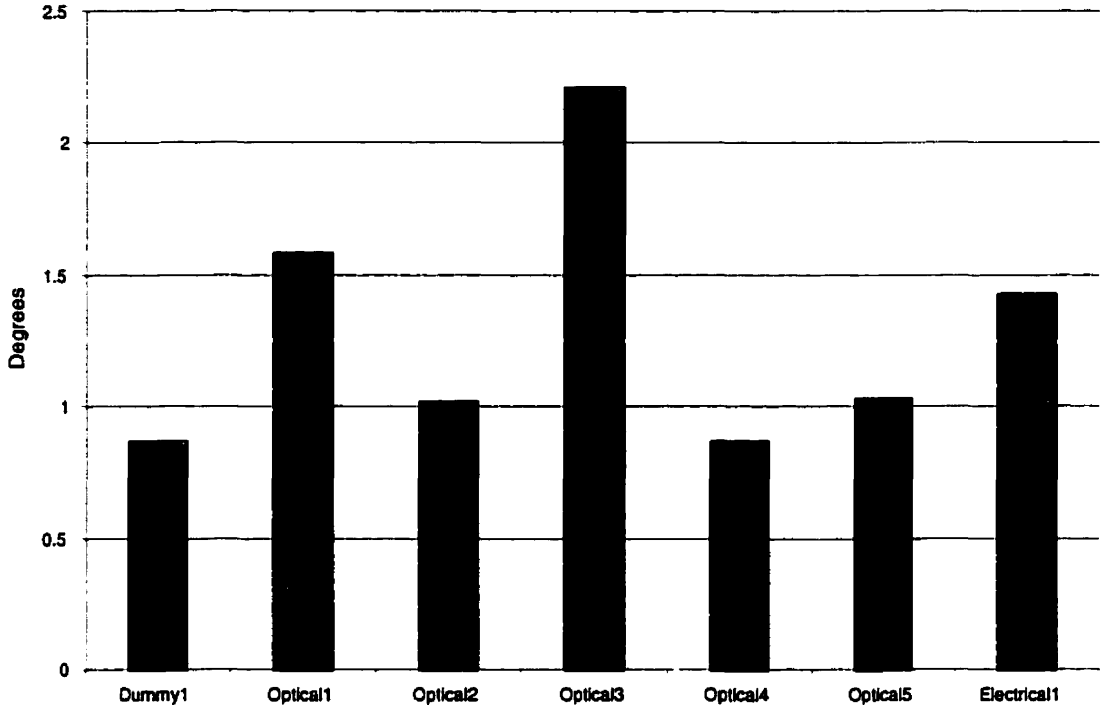

3.8). This data was measured using a network analyser. This value indicates that our system was not package bandwidth limited.

In COB I (Fig. 3.1a), the chip was directly mounted on the daughterboard. To manage the higher heat dissipation of the CMOS-QCSE die in COB II, the optoelectronic device array is mounted on a gold-plated, ground-connected heat spreader that moves heat via conduction from the chip to a thermoelectric cooler mounted on the back of the board (figure 3.1b and figure 3.9). The designed peak responsivity of the CMOS-QCSE modulators occurs at 25°C with the available bias voltages (no temperature dependent optimization was performed), and the thermoelectric cooler was able to maintain the die's center at this temperature to within 0.25°C with the die dissipating up to 0.5W. This measurement does not take

Figure 3.8: Network analyser results for package bandwidth

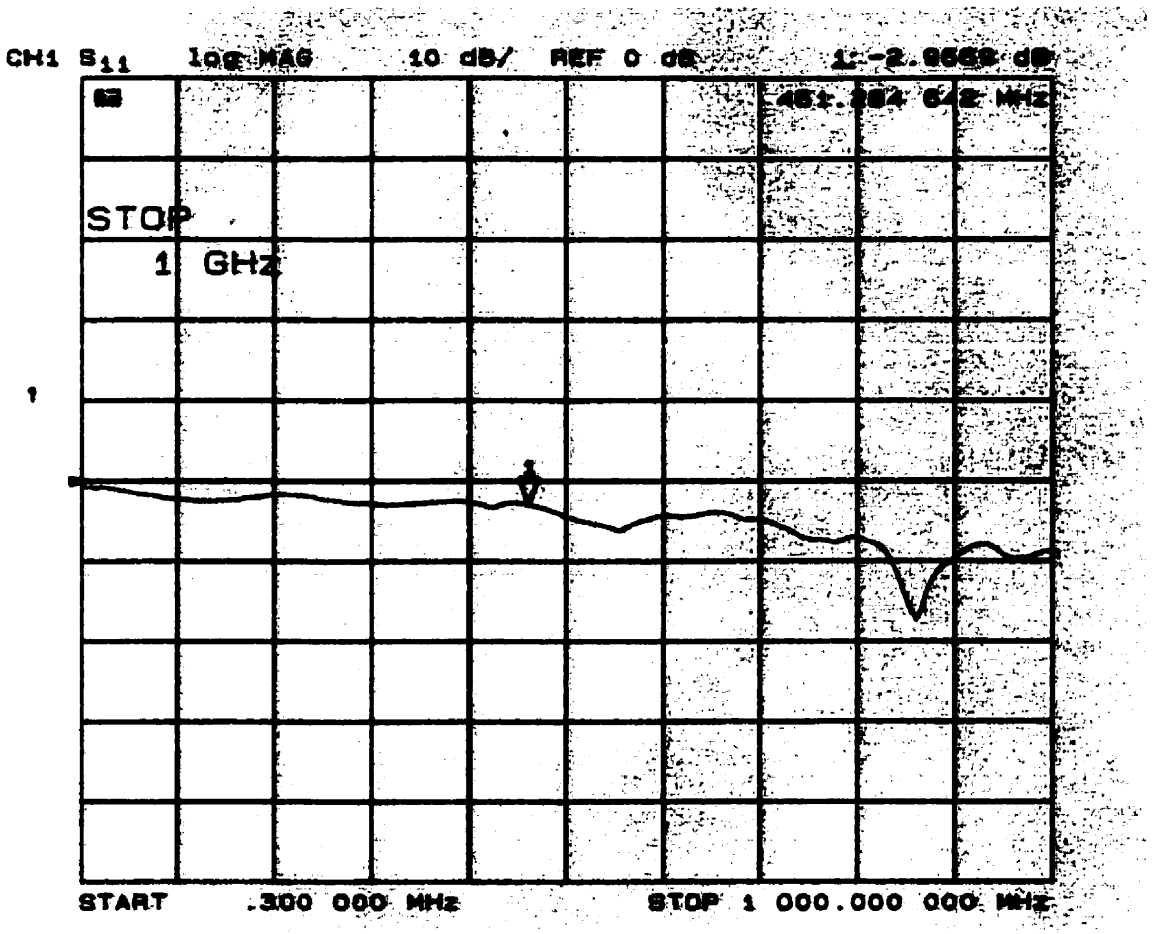

into account die temperature nonuniformity as these data were not available for the CMOS-QCSE die.

Using a dedicated thermal characterization daughterboard, the thermal resistance of the packaging was accurately measured. The thermal resistance from junction to case, θ_{jc}, is defined as:

$$\theta_{jc} = \frac{T_j - T_c}{P_{in}} \tag{3.1}$$

Where T_j and T_c are the temperatures of the junction (chip surface) and case (for this package, the back of the heat spreader), respectively, and P_{in} is the power

Figure 3.9: Thermal management

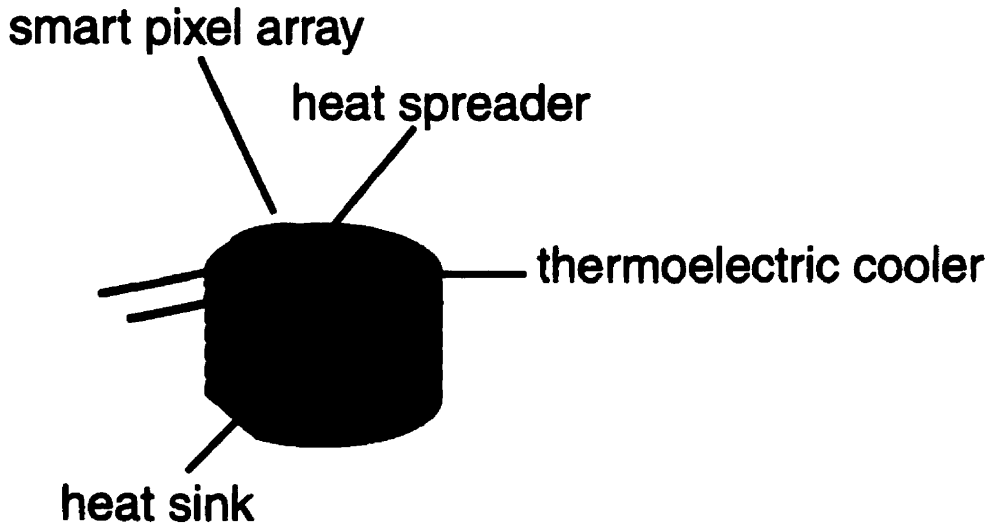

being dissipated on the chip surface. For the measurement of thermal resistance, a thermistor was mounted on the back of the heat spreader and on the top of an electrical-only version of the CMOS-QCSE die. The CMOS inputs were left to float thereby forcing a high power dissipation in the die. With this setup, the thermal resistance of the package was measured to be 7 ± 2 °C/W.

The thermal model assumed a conduction-dominated thermal resistance, with the resistance of a conductive path being defined as:

$$R = \frac{l}{kA} \tag{3.2}$$

with A being the cross-sectional area (m^2), l the length of the path (m), and k the thermal conductivity of the material (m • °C / W). The predicted values using this first order model are tabulated below:

Table 3.1: Modelled thermal resistances

Component	Thermal resistance
Silicon die	0.8 °C/W
Epoxy	2.0 ° C/W [4]
Heat spreader	0.2 °C/W

Using these values, the thermal resistance was 3°C / W, which agrees quite closely with the measured value.

It is clear, even in the coarse view offered by this simple model, that the thermal resistance of the epoxy used to glue the chip to the heat spreader dominates the thermal resistance of this packaging approach.

3.3 Conclusions

The packaging approach presented here can be expanded and modified for future, increasingly demanding applications.

Results presented here indicate that the packaging techniques used for the CMOS-QCSE packaging (COB I & II) can be scaled to future systems employing larger and faster 2-D optoelectronic device arrays that will run hotter. Using hybridized QCSE on CMOS or BiCMOS technology, this future package could support a 1 cm x 1 cm

chip dissipating 5 W. This scaleability of approach will be demonstrated in the analysis of the Phase III packaging, which is in the following chapter.

3.4 References

[1] R.R. Johnson, "Multichip modules: next-generation packages," *IEEE Spectrum*, March 1990, pp34–36, 46, 48.

[2] *MECL System Design Handbook, Fourth Edition,* Motorola Semiconductor Products, 1983

[3] AMP™, Inc., "AMP Micro-Strip Cable Assemblies," Cat. 65069, Dec. 1991.

[4] Data sheet, Epo-Tek H20E, Electrically conductive, Silver Epoxy, Epoxy Technology Inc.

Chapter 4: BiCMOS-QCSE packaging

4.1 Introduction

The Photonic Systems Group Phase III BiCMOS-QCSE chip is the core device for a future optical backplane, the Phase III demonstrator. It is a large, 8 x 8 mm chip with a 16 x 16 array of optoelectronic transceivers. On this surface-active chip there is 1024 QCSE devices, arranged into an 8x8 array of optical clusters. This chip must be integrated into a 4 stage free-space optical interconnect. The following section deals with the design, implementation and characterization of the packaging that was designed for this device, with the expectation that this packaging will be integrated into the completed system.

Figure 4.1: Optical interconnect design

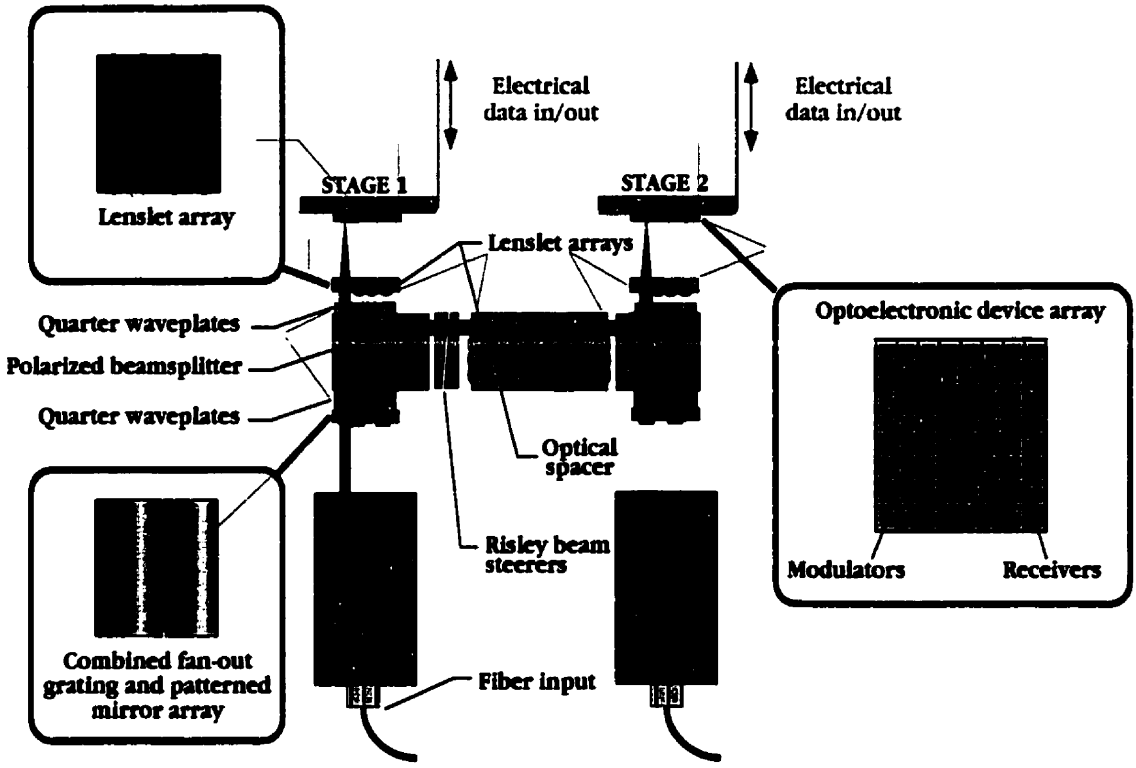

4.2 Optical interconnect

The optical interconnect for the Phase III demonstrator was designed by Brian Robertson. A schematic of the interconnect is shown in figure 4.1. A detailed description of this interconnect and the design constraints and trade-offs that led to its development can be found in [1]. The optical system consists of mini-lenses that provide relay and interconnect functionalities. The optical data links (ODLs) that this design provides are arranged into clusters [2], because of the resulting increase in alignment tolerances and to increase the overall dimensions of the system.

Figure 4.2: Phase III chip

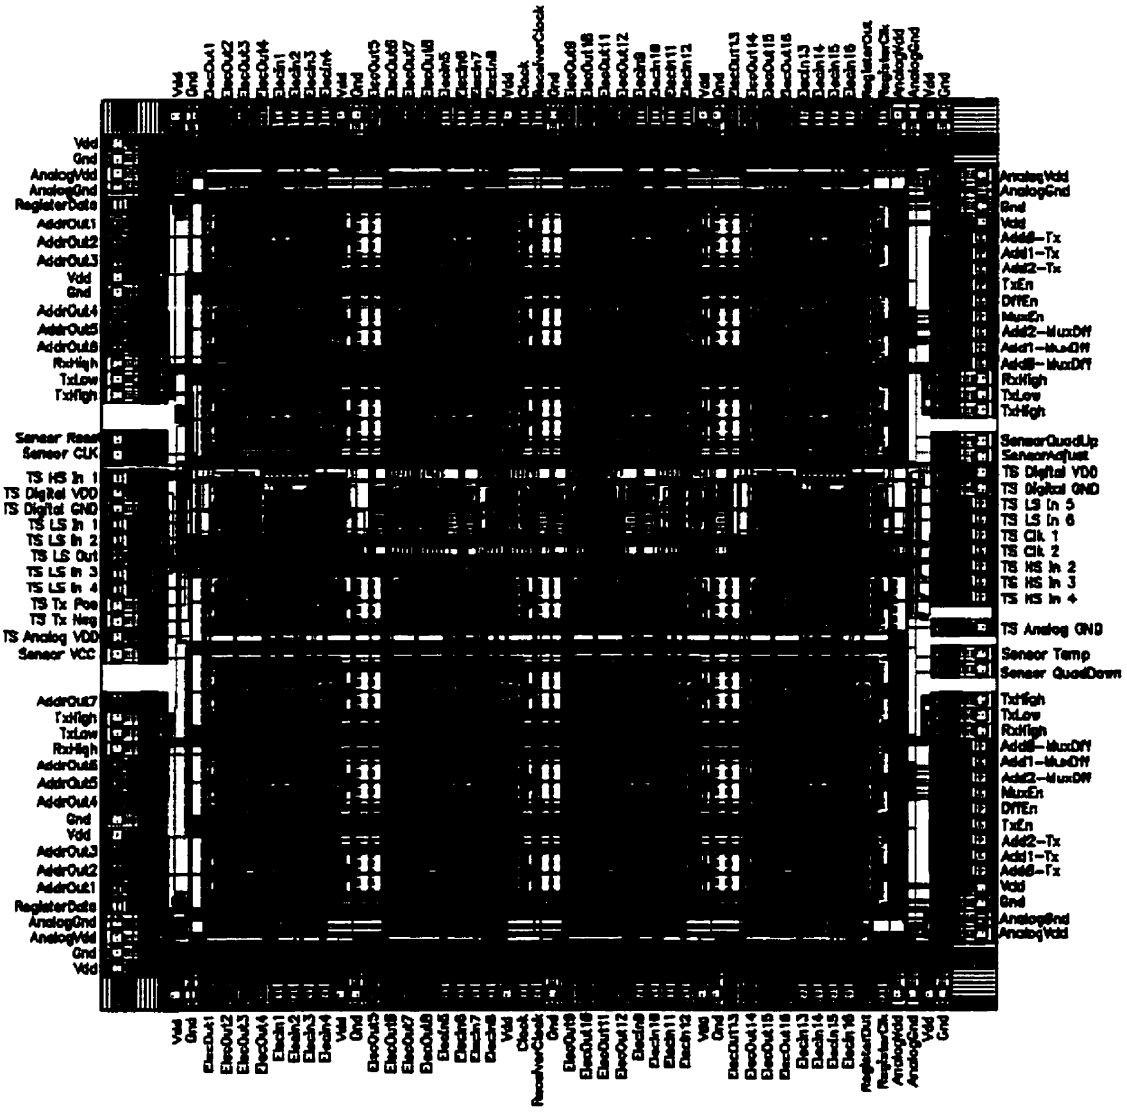

4.3 Optoelectronic device description

The Phase III chip (figure 4.2) is an 8x8 mm, 0.8 µm BiCMOS device. It incorporates almost 0.5 million transistors. The optically active section of the chip is 6.4 mm x 6.4 mm (after it has been flip-chipped with the QCSE devices). The optically active section is divided into an 8x8 array of clusters, each containing 16 QCSE modulators

and 16 detectors. A single cluster is shown in figure 4.3. The Phase III chip has 189

Figure 4.3: Single cluster

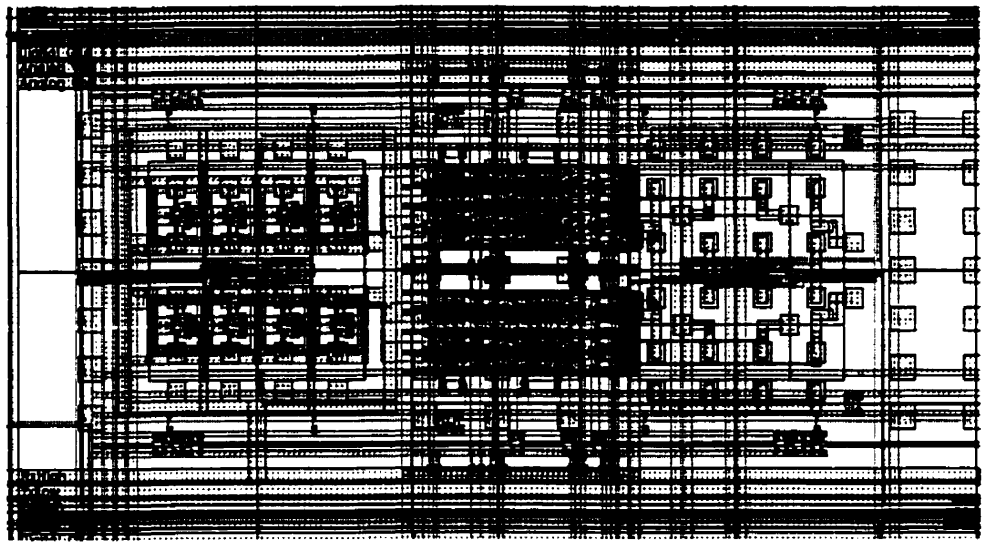

bond pads, consisting of 32 inputs, 32 outputs, 2 clocks and control and power pads.

4.4 Physical design constraints

The optical interconnect for this system is extremely compact, and this compactness leads to a space constraint for the optoelectronic packaging. The total available footprint (area in chip's plane) was 4 cm x 4 cm. There was no inherent volumetric constraint.

Examining standard consumer electronic carriers and packaging techniques, it became apparent that none of these were adequate for this packaging task. The carrier and connectorization that allowed for mechanical decoupling of the support electronics would have to be housed within the 4 cm by 4 cm area. This area was not sufficient to support a standard 25 mil pitch (625 μm) high-density connector,

with the required 139 lines of off-chip connectivity. An alternate approach, moving the connector away from the packaging, proved to be very effective in side-stepping the problem of connector size.

4.5 Electrical design constraints

To connect the Phase III chip to support and application electronics, it was apparent that the electronic packaging and connection technique would have to have the following specifications:

1. 139 signal connections to motherboard
2. mechanical flexibility to allow the packaging to be mechanically isolated from the motherboards, while still remaining electrical contiguous with them
3. termination for 32 high-speed inputs and 32 high-speed outputs (100 MHz) and for the two clocks (320 MHz)
4. impedance control for high-speed inputs and outputs and for the two clocks
5. low crosstalk between adjacent signal lines
6. power supply decoupling to reduce noise

In the packaging description, it will be shown how the design of a *package module* directly addresses each of these specifications.

4.6 Thermal design constraints

The Phase III chip is a large chip with a total expected power consumption of 5 W. Based on preliminary design and later characterization data for the QCSE devices, it

was apparent that this chip needed to operate at a temperature near 50°C to acheive optimal performance. This data was later confirmed by temperature dependent optimization, and the combined efficiency (as presented in the chapter 2) of the modulators and detectors was found to peak at 48°C. With an expected power of 5 W, cooling the chip to 48°C requires active temperature control and stabilization techniques.

4.7 General package module description

The designed *packaging module* is shown in figure 4.4. The packaging satisfies the

Figure 4.4: Photograph of package module

design constraints imposed by the system using a number of novel approaches, as

illustrated in Table 4.1. An exploded view of the packaging, showing the various

Table 4.1: Satisfaction of constraints with package module

Specification	Application
5 W power dissipation	Thermoelectric cooler and temperature sensors provide temperature control. Requires approximately 5 W cooling to dissipate 5 W of power. Thus, heat sink must dissipate 10 W
4 x 4 cm footprint	Dimensions of package module: 4 cm x 4 cm x 6 cm
139 signal connections to motherboard	Flexible printed circuit with 8 mil line pitch connects package module to motherboard, using high-density connector at motherboard end. Bond pads are wirebonded directly to flexible PCB
Mechanical flexibility	Flexible printed circuit is 6 mil thick, and is extremely flexible
Impedance control	50 ¾ impedance controlled connector and microstrip stack-up for flexible printed circuit
Termination for high-speed inputs and outputs and clocks	High speed inputs are terminated using active Schottky diode termination, adjacent to chip. Clocks are terminated with 1/4 W resistors. Outputs are terminated on motherboard.
Low crosstalk	<0.1 V for adjacent lines
Power supply decoupling	Surface-mount capacitors are located adjacent to chip in package

components, is shown in figure 4.5.

4.7.1 Chip attach

The Phase III chip is attached directly to a custom, sculpted mount (figure 4.6) using *Epoxy Technology* H20E thermally and electrically conductive silver-filled epoxy. Previous measurements (necessary for the die tilt measurements of chapter 3) have determined that the average thickness of this epoxy after setting for

Figure 4.5: Exploded view of packaging

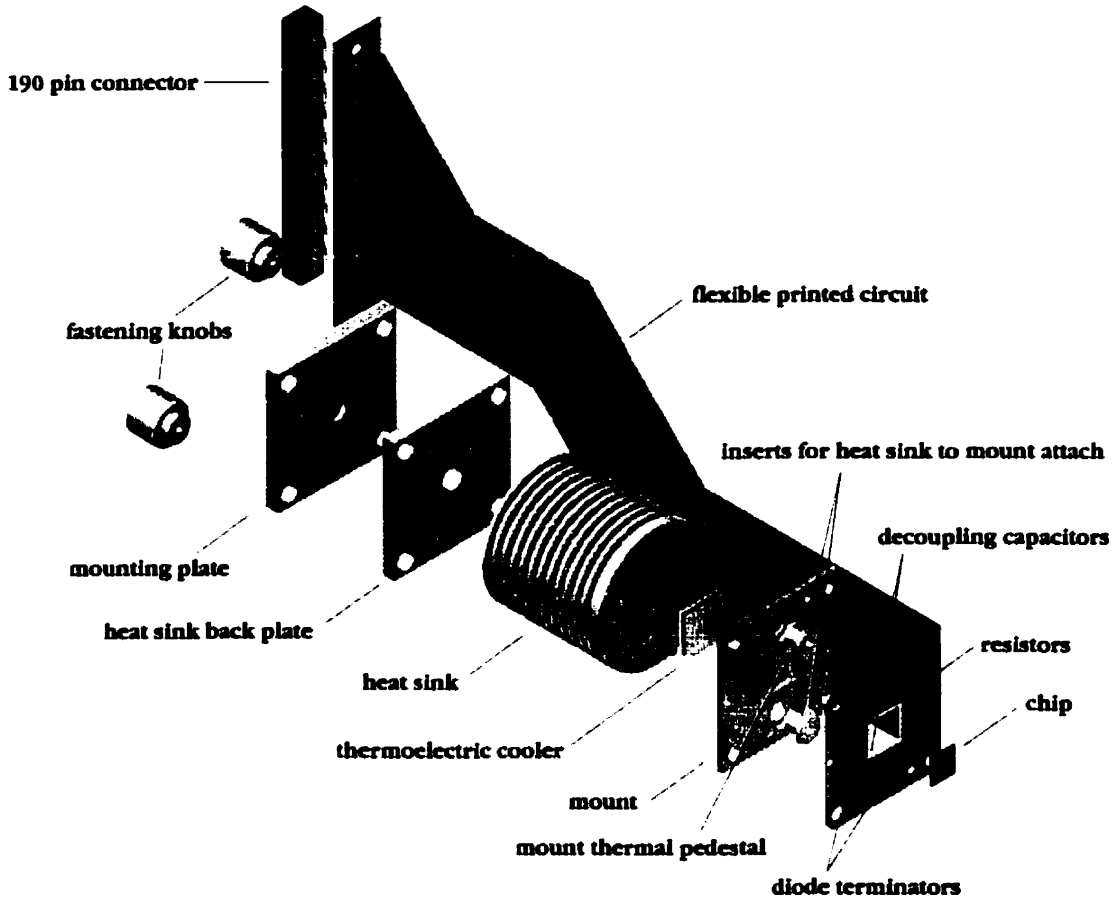

multiple chip-attaches is approximately 50 μm. The epoxy is highly thermally and electrically conductive [3]. The rated thermal resistance for this glue thickness is less than 1 °C/W for the bond area of 8 mm x 8 mm (315 mil x 315 mil). The mount is also used as a ground connection for the rear side of the chip, so the high electrical conductivity of the H20E adhesive was also desireable. The calculated electrical resistance of the die attach is based on the H20E maximum electrical volume resistivity of 0.004 ¾-cm, yielding a maximum total resistance of:

Figure 4.6: Mount

$$R = \frac{r \times l}{A} = \frac{4 \times 10^{-3} \times 50 \times 10^{-4}}{0.8 \times 0.8} = 31\mu\Omega \qquad (4.1)$$

Thus, using H20E represents a good, solid connection to the ground plane supported by the mount.

The chip was placed under a microscrope, and air gaps surrounding the chip were used as alignment features.

4.7.2 Off-chip connectivity

A custom-designed, two-layer flexible printed circuit is glued onto the mount using high temperature thermal set lamination techniques, providing a hard backing for the printed circuit. The chip is wirebonded directly to the gold-on-nickel plated printed circuit, using an aluminum wedge ultrasonic wirebonder. The remaining, highly flexible portion of the flexible PCB provides the required electrical connectivity to a dense, 190 position connector (from AMP) with an integral ground

plane and 50 ¾ impedance control. The flexible printed circuit is designed to allow for the required bends and space constraints imposed by the presence of four 100 x 160 mm (3 U x 160 mm) motherboards., in a 3 U VME chassis. The complete packaging system required for a 4 node optical interconnect is shown in figure 4.9, including application cards.

Flexible printed circuit

The flexible printed circuit was designed in the Mentor Graphics CAD environment, and manufactured by Dynaflex of San Jose, CA. The resulting (flattened) flexible printed circuit is shown in figures 4.7 and 4.8. The overall dimensions of the flexible

Figure 4.7: Close-up of flexible printed circuit, chip region

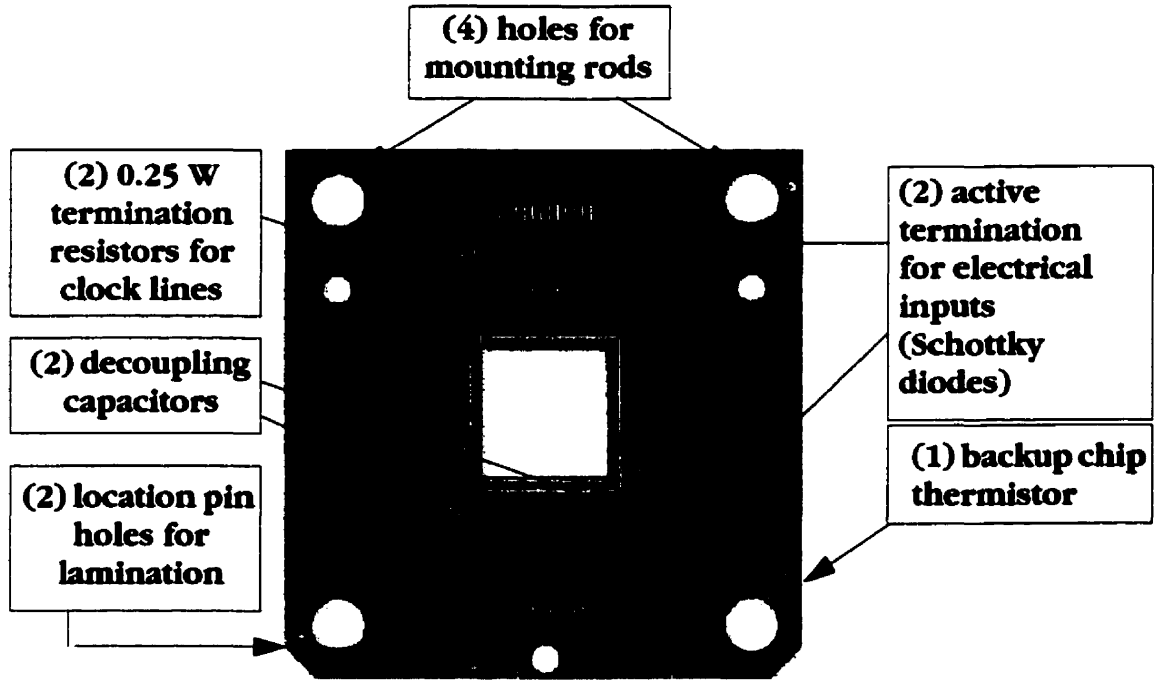

circuit are 260 x 120 mm, when flattened. In the final packaging configuration, the flexible circuit is designed to have either two or three bends, depending on its

Figure 4.8: Flexible printed circuit (approximately 2/3 actual size)

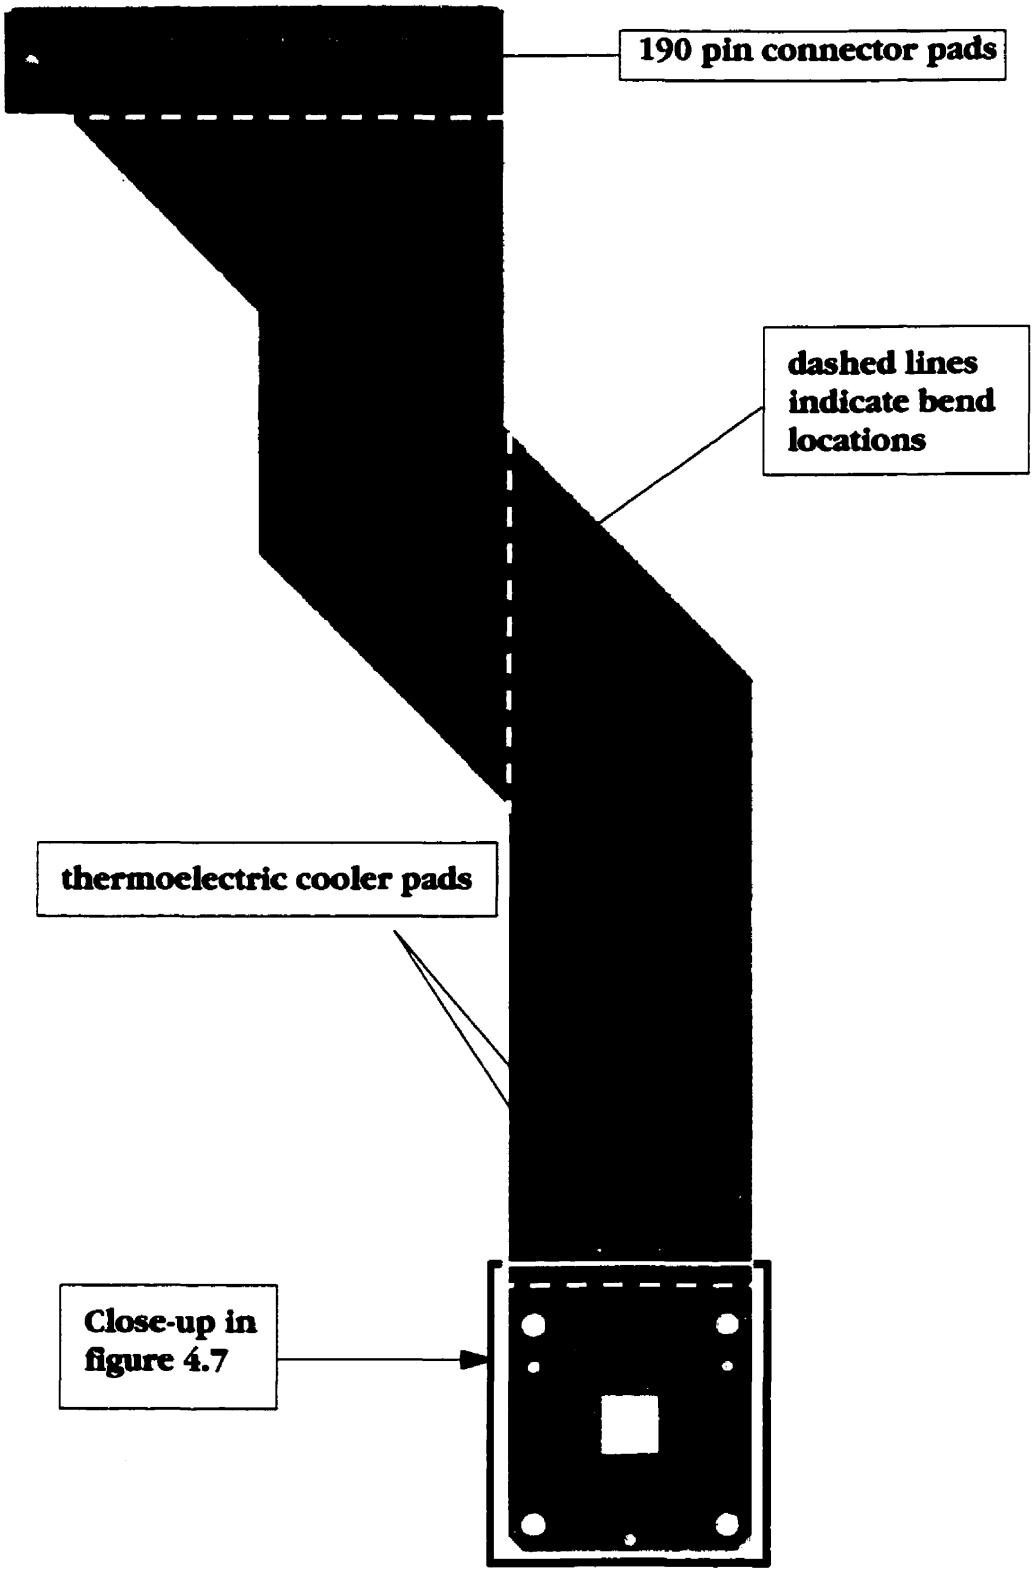

Figure 4.9: Packaging for four motherboard interconnect

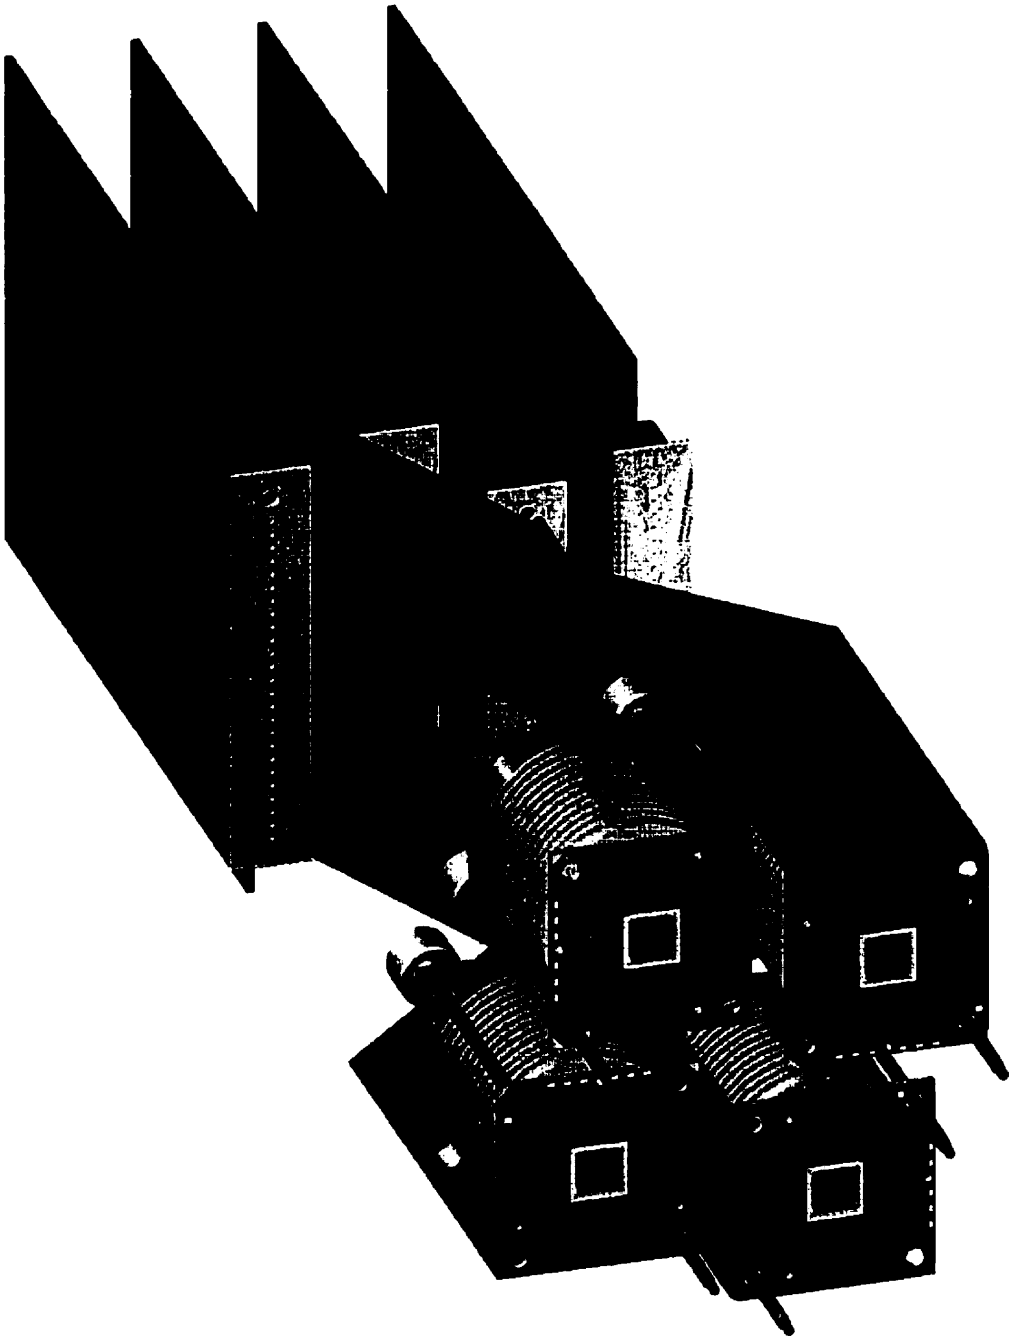

relative position in the system. This is shown in figure 4.9, where the packaging module for nodes 2 and 4 have three bends in the flexible PCB, and the nodes 1

and 3 have two bends in the flexible PCB. This altenating bending was mandated by the optical system, where each packaging module must be rotated 90° clockwise relative to its predecessor, in order to satisfy the optical interconnect's design (this is shown in figure 4.10). To economize on the overall cost of the packaging system, it was decided that a single flexible printed circuit would have to satisfy both of the configurations, necessitating an alternating number of bends.

The flexible circuit is 150 μm in thickness (6 mil) and has two (copper) layers. Where the coverlayer does not cover the circuit, it has been plated with nickel and then a protective layer of gold, to allow exposed pads to be wirebonded using an aluminum wedge ultrasonic wirebonder. The top copper layer has high speed traces and the back copper layer is a ground plane to allow for the microstrip board stackup shown in figure 4.11. Using the electromagnetic field solver built into the CAD tools package (Mentor Graphics), the calculated impedance of the microstrip, given nominal material thicknesses and dielectric constants, is 50 ¾. These values are shown in Table 4.2. Using the same field solver, the estimated coupled crosstalk between adjacent lines is 0.1 V, over the 280 mm maximum parallel distance that two signals will couple.

Experimental verification of these electrical field analyses can be found in section 4.8.

Termination

The design constraints imposed by the system and chip design dictated the termination of the signal and clock lines. The clock lines are required to operate at

Figure 4.10: Optical bitmapping of packaging (looking onto devices)

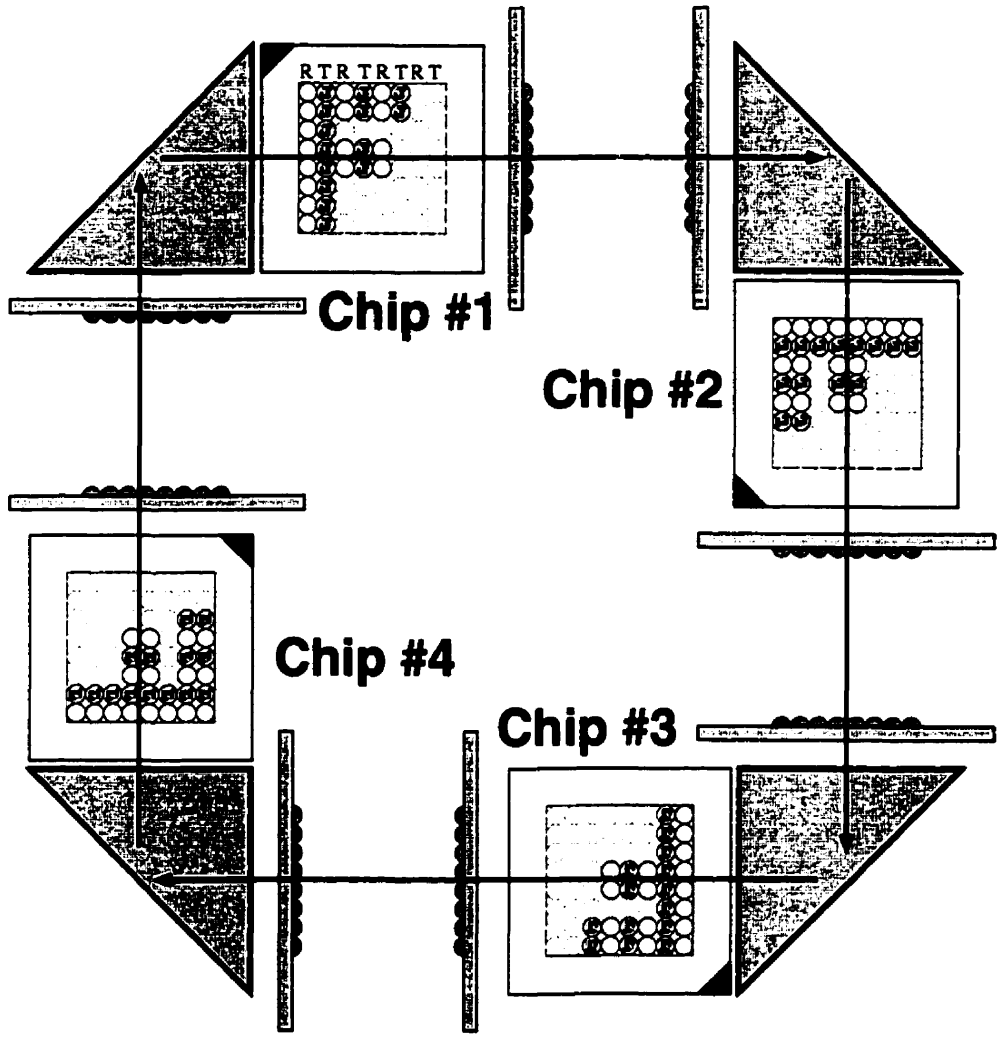

1) **The letter 'F' formed with lenslets indicates the cluster-to-cluster optical mapping.**

2) **The letter 'J' formed with QCSE devices indicate the device-to-device optical mapping.**

a 5 V logic swing at a frequency of 320 MHz. The input pads on the chip are standard high impedance CMOS pads, necessitating external resistor termination. The two clocks, designated by chip design as Clock and ReceiverClock are

Figure 4.11: Microstrip stackup for impedance control in flexible PCB

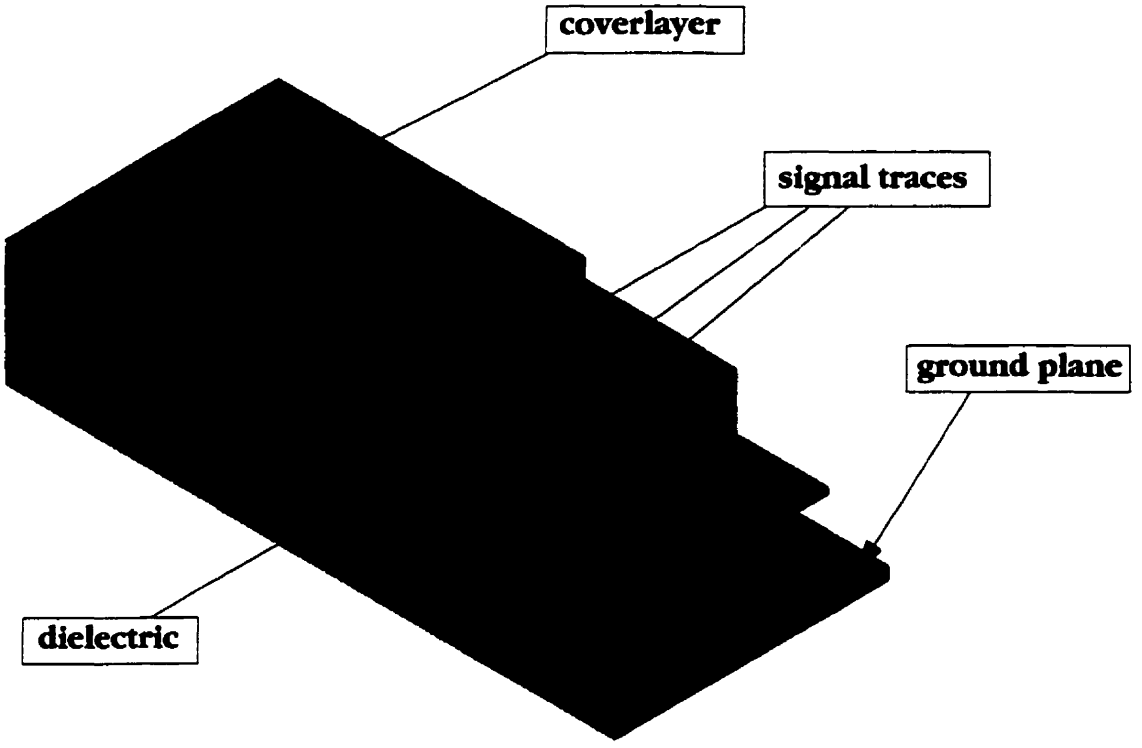

Table 4.2: Microstrip stackup and trace parameters

	Parameter	Symbol	Value	
			S.I. units	Standard units
Stackup parameters for microstrip impedance control	Dielectric thickness	h	50 µm	2 mil
	Coverlayer thickness	—	50 µm	2 mil
	Signal trace thickness	t	17.8 µm	1 oz/ft²
	Signal trace width	W	100 µm	4 mil
	Dielectric relative permittivity	ε_R	3.2	3.2
	Combined relative permittivity of coverlayer and dielectric	ε_R	4.7	4.7

Table 4.2: Microstrip stackup and trace parameters

	Parameter	Symbol	Value	
			S.I. units	Standard units
Trace parameters	(Minimum) Trace and bond pad pitch	—	200 μm	8 mil
	Maximum trace length	—	285.9 mm	11.25 in

terminated with a resistor of the characteristic impedance of the flexible printed circuit. The experimental measurement of this characteristic impedance is presented in section 4.8.5.

Power considerations in the terminating resistors prohibited terminating the resistors to ground, as the clock drivers would have to source and the resistors would have to sink the calculated:

$$P_{PEAK} = \frac{(\Delta V)^2}{R} = \frac{5^2}{50} = 0.5W \tag{4.1}$$

The required resistors would have been prohibitively large to place in close proximity to the load (Phase III chip), as required by good signal integrity design. Instead, the design opted to terminate these clock lines with resistors terminated to an additional power trace, a 2.5 V signal. With this arrangement, the peak and average power were the same:

$$P_{PEAK} = \frac{(\Delta V)^2}{R} = \frac{2.5^2}{50} = 0.125W \tag{4.2}$$

This termination allowed the use of 0.25 W (to overrate the 0.125 W load) surface-mount 1210 package resistors, available in many different resistances. The resistance placement is indicated in figure 4.7.

For the high-speed inputs to the Phase III chip, a termination method was need to allow them to operate at the expected chip performance bit rate of 100 MHz. Resistive termination was prohibitively large in footprint, as was passive filtering using resistors and capacitors in series. Instead, active termination using Schottky diode terminators (the PACDN005 from *California Micro Devices*) was used. This method allows for characteristic impedance independence, as the diode terminators are designed to prevent only excursions from the supply rails (in this case, 0V and 5 V) and the resulting reflection. Simulation of these lines using HSPICE (with and without active diode termination) is shown in figure 4.12. The models used for this simulation were provided by the manufacturer (CMD) and by our chip foundry representative (Canadian Microelectronics Corporation).

Decoupling capacitors

Two decoupling capacitors are located immediately adjacent to the chip on the flexible printed circuit, as shown in figure 4.7. These decoupling capacitors decouple the chip supply to the ground plane.

Thermistor chip

For calibration of the on-chip temperature sensors and as a backup in case of temperature sensor failure, a single pad used to glue a wirebondable chip thermistor (*BetaTherm* 10K3CG2) is located as shown in figure 4.7.

Figure 4.12: HSPICE simulation of square wave signal at load

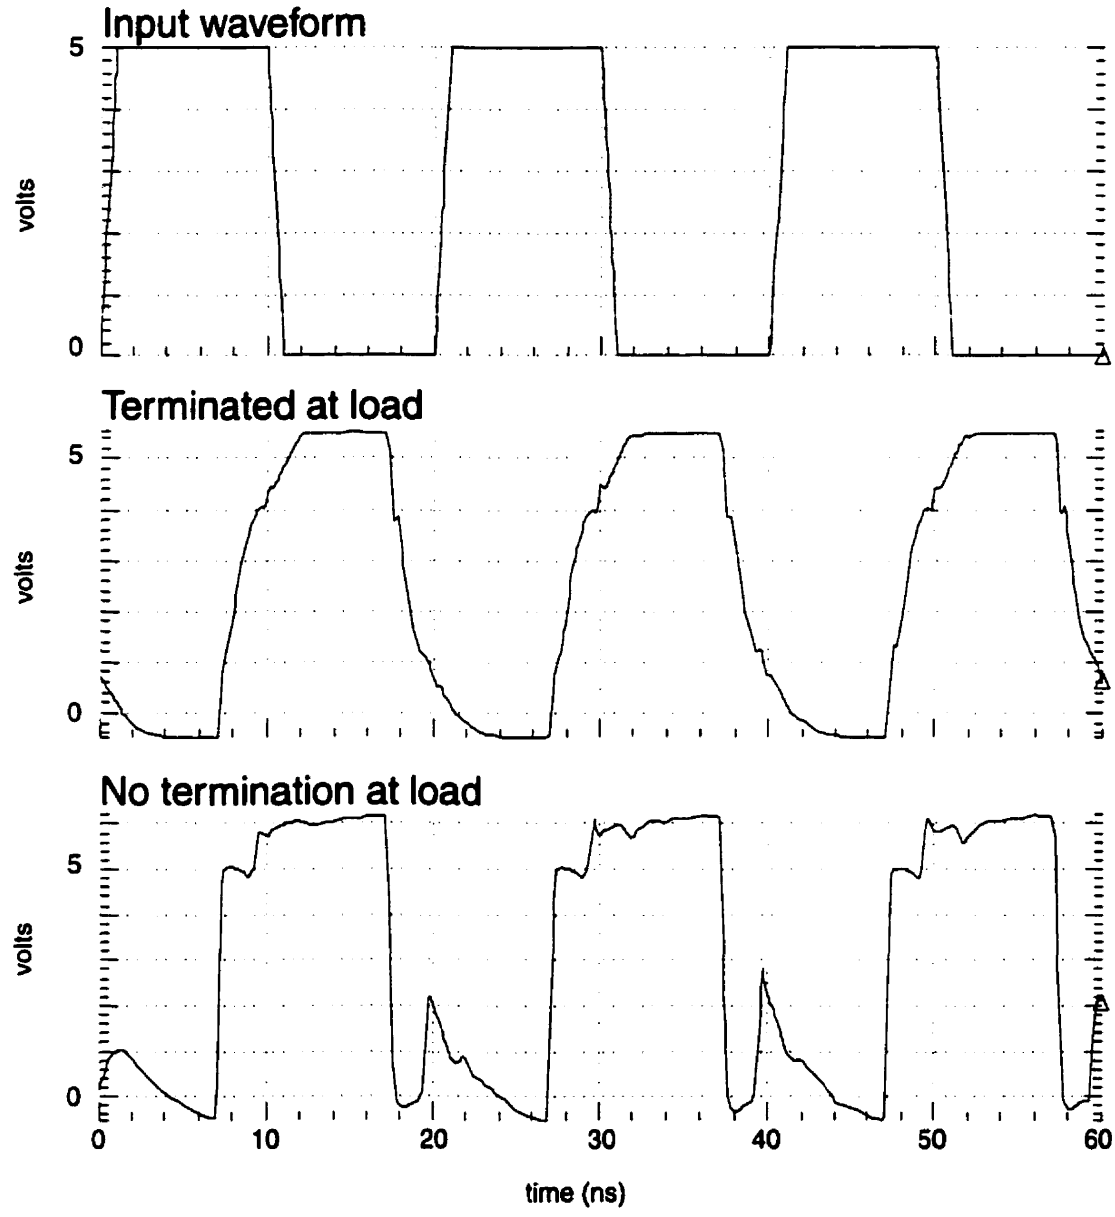

Alignment features

Additional alignment features, used for kinematic mounting of packaging into the optical backplane, are incorporated into the flexible PCB design and chip design. These features include 1206 package surface-mount LEDs (Light Emitting Diodes), their current-limiting 1206 package surface mount resistors, and pads for off-board

wiring. These features and their application and characterization will not be discussed here, although future publications will examine these features in detail.

Optomechanical interfacing

In the design of the system, the first optical surface, a lenslet array, is to be prepackaged in tight 6 degree of freedom alignment to the chip surface. In order to do this, 3 wells are located on the mount surface. A custom-machined spacer, with the lenslet preglued to it, is located using quadrant detectors on the Phase III chip surface, to a precise location above it. In this configuration, three small (<1 mm) pins reside within the wells on the mount surface. At this point, the spacer housing the lenslet can be glued to the package module by filling the well with glue, thereby fixing the pin into the well. This is shown in figure 4.13.

Figure 4.13: Lenslet to chip mounting

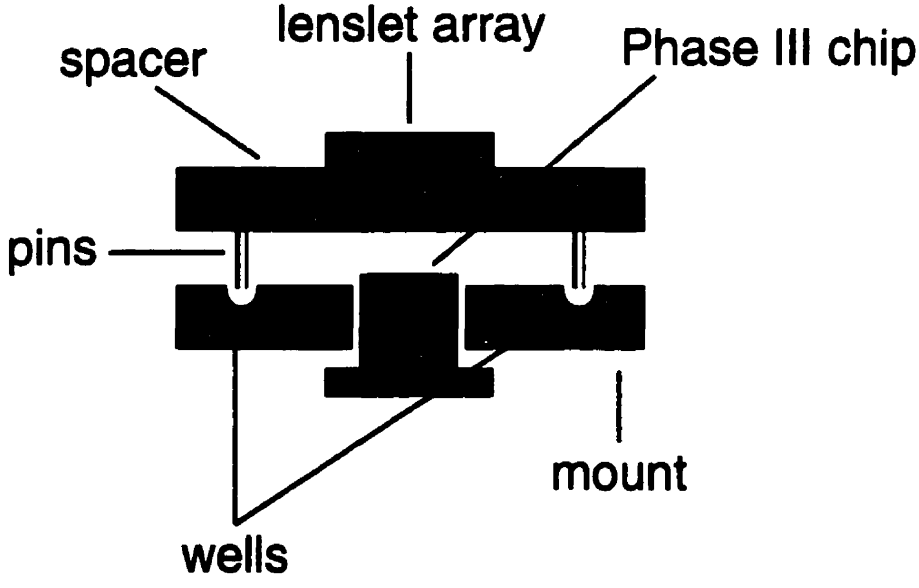

For this interfacing, it is clear that the chip to mount miasalignment must lie within several hard tolerance limits so that the pins do not locate outside of the wells. These limits forced the chip attach alignment to be within 100 μm in lateral alignment, and limited the maximum die tilt to 1°. The height of the chip attach was not critical nor limiting, as the pins permit a large vertical play.

This completed package module is aligned to the rest of the optical system using a kinematic mount, whose design and characterization will not be presented here.

4.7.3 Thermal management

The chip is attached to the thermally conductive mount using conductive epoxy adhesive (*Epoxy Technology* H20E). A thermoelectric (TE) cooler (*Marlow Industries* 1023T) is attached to the mount by sandwiching it between the mount and a custom-machined heat sink, using two screws to apply pressure to the TE cooler. Thermal grease is used to fill surface morphological gaps and pockets, making a solid thermal junctions between TE cooler and heat sink, and TE cooler and mount. Air gaps machined into the mount increase the resistance to thermal transfer between the chip area (pedestal) and the surrounding mount area, and also help to limit the allowed paths for thermal conduction. The complete thermal construction is illustrated in figure 4.14. A side view of the package module is shown in figure 4.15.

To complete the thermal management, an array of temperature sensors was incorporated onto the surface of the Phase III chip. This array's response allows the temperature of the chip to be controlled and also allows for temperature

Figure 4.14: Thermal management in package module

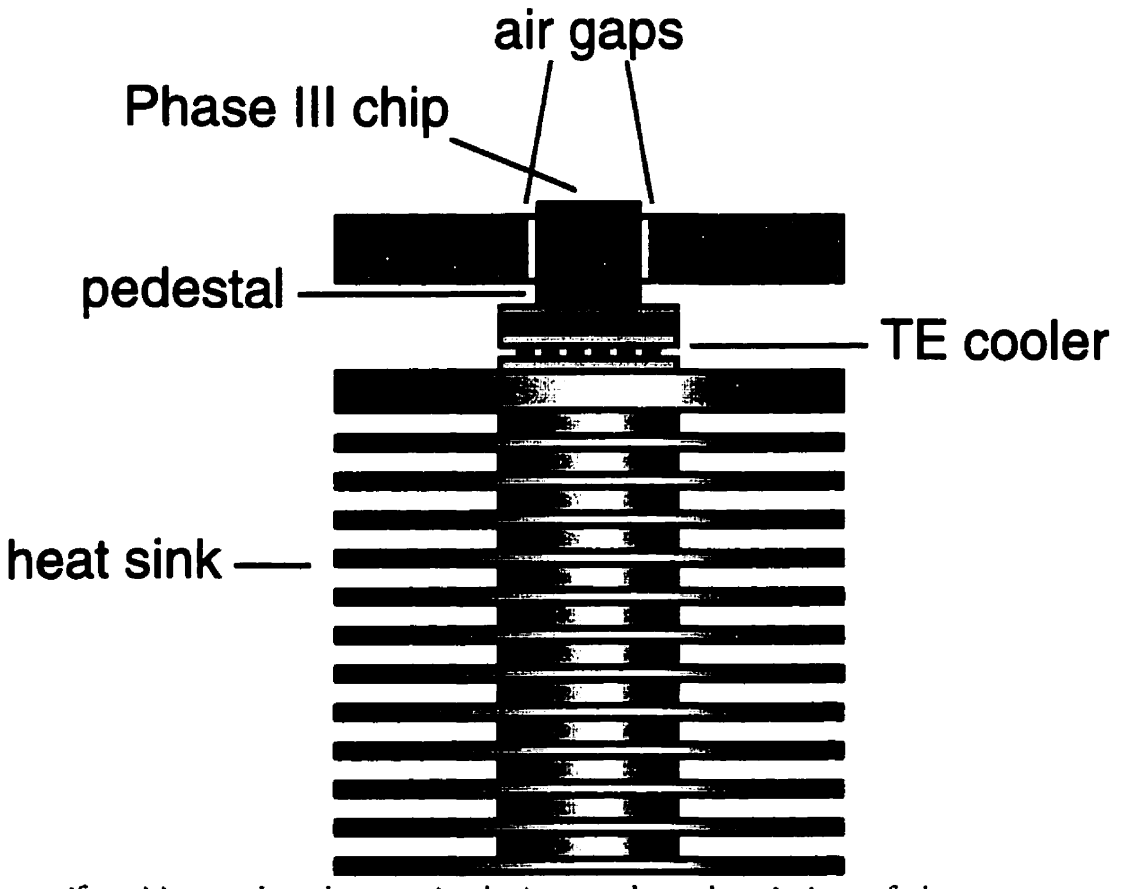

nonuniformities to be characterized. A complete description of the temperature sensor design and performance can be found in chapter 5.

4.8 Characterization of packaging

To test the packaging design, a custom board was designed and built, allowing sensor values to be sampled, as well as to allow high speed testing of a representative number of electrical I/O lines. This board is shown in figure 4.16. This board also permitted control of the chip temperature using a thermoelectric cooler integrated into the package module, as illustrated in figure 4.14.

Figure 4.15: Package module, side view

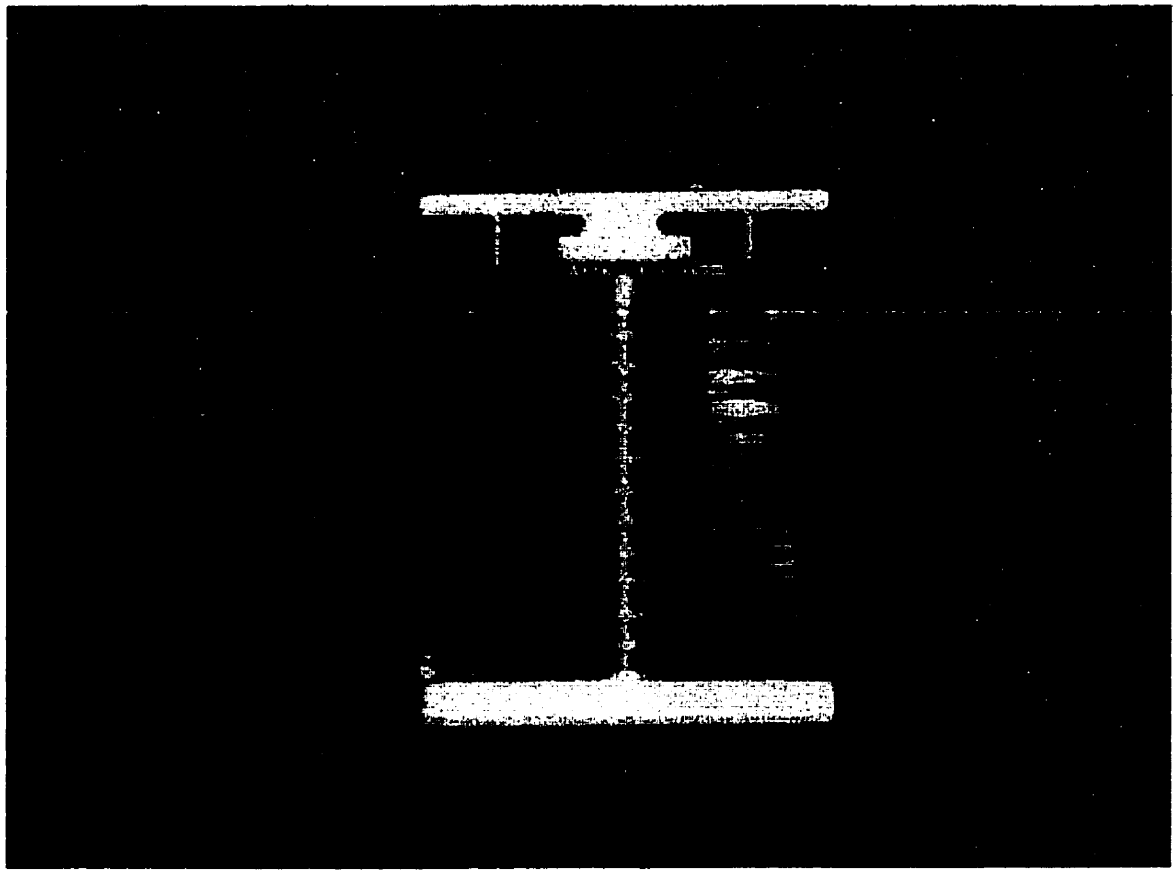

To test the characteristics of the packaging system, one complete package module was assembled. This chip module incorporated an electrical-only version of the Phase III chip, which was attached and wirebonded into the mount. The measured lateral misalignment of the chip in this configuration was less 50 μm, when placed under a microscrope and aligned to the air gaps in the mount.

The flexible printed circuit was prelaminated to the aluminum mount using a thermal set adhesive in a lamination press. Thermal grease was applied to fill interstitial gaps in the thermal junctions between heat sink, thermoelectric cooler

Figure 4.16: Packaging test and sensor sampling board

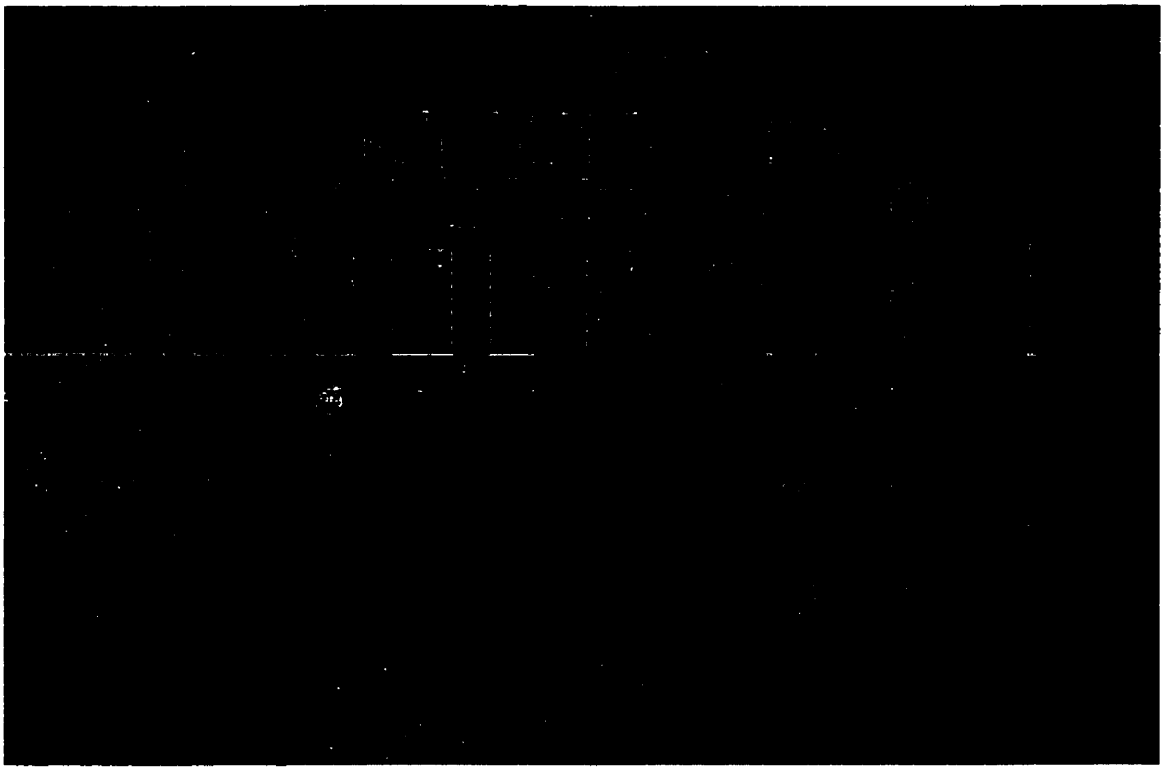

and mount. The connector and termination components were mounted using surface mount soldering techniques onto the flexible printed circuit. The calibration thermistor was attached and wirebonded to the flexible printed circuit using standard chip attach techniques. A front view of this completed package module is shown in figure 4.17 and the entire module is shown in figure 4.18.

4.8.4 Thermal characterization

The temperature sensors on the die surface are functional, and provide accurate active feedback to a temperature controller. This temperature controller (an IDT 5910), given high precision temperature sensor input, can stabilize the chip temperature to 0.01°C [4]. The limit of the resolution of the temperature sensors is 0.02°C [5] given the current computer sampling and calibration technology, but this

figure is more than required to ensure that the temperature lies within a biasing window as shown in section 2.4.3. The thermoelectric cooler is supplied with current through the flexible printed circuit. The leads for the thermoelectric cooler are soldered into the flexible printed circuit. Wide traces and multiple connector positions take the large current required for thermoelectric cooling to the connected board(s).

Figure 4.17: Completed package module (front view)

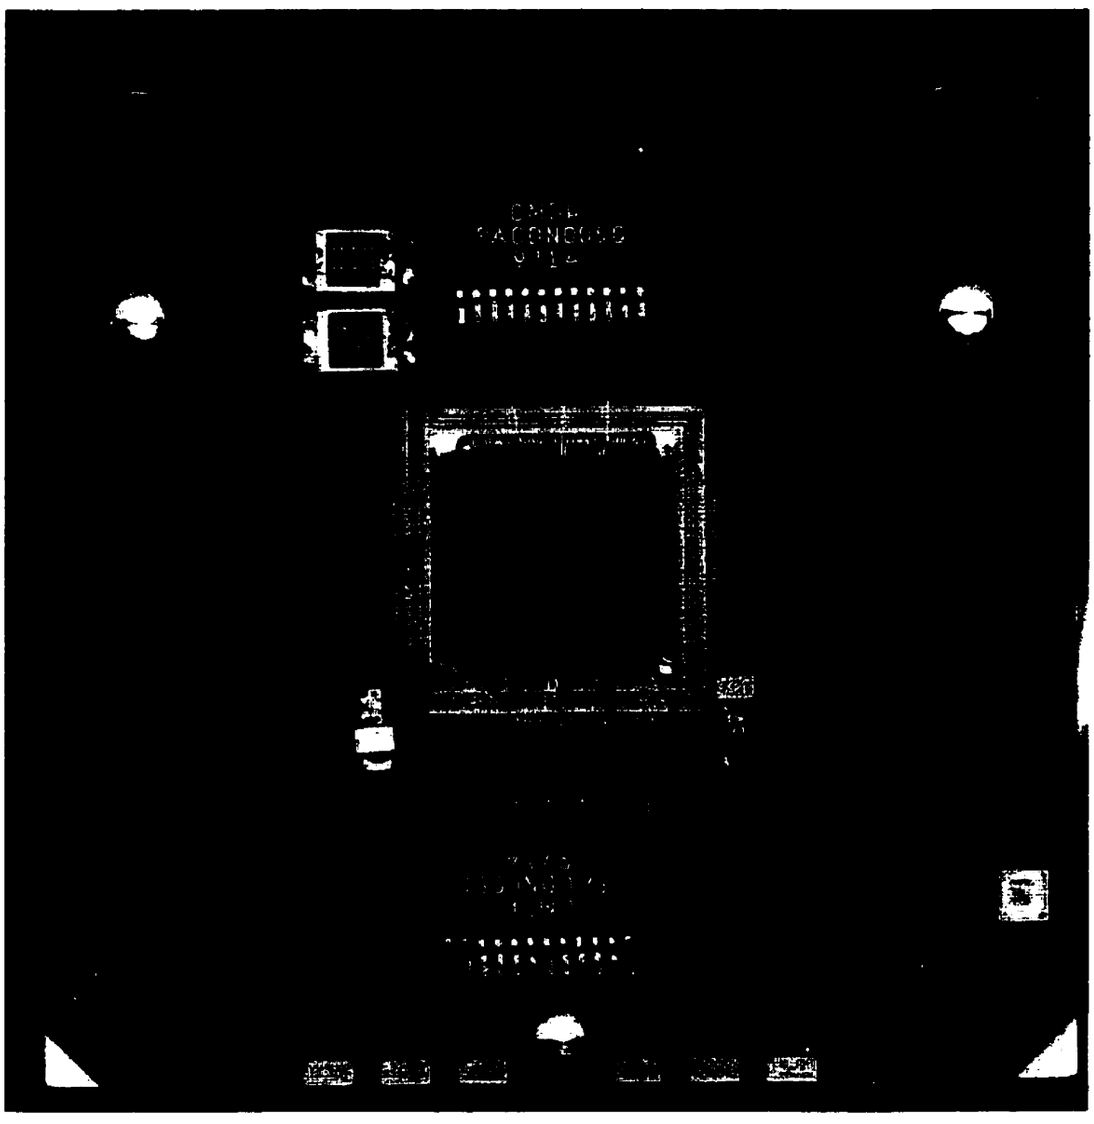

Figure 4.18: Package module

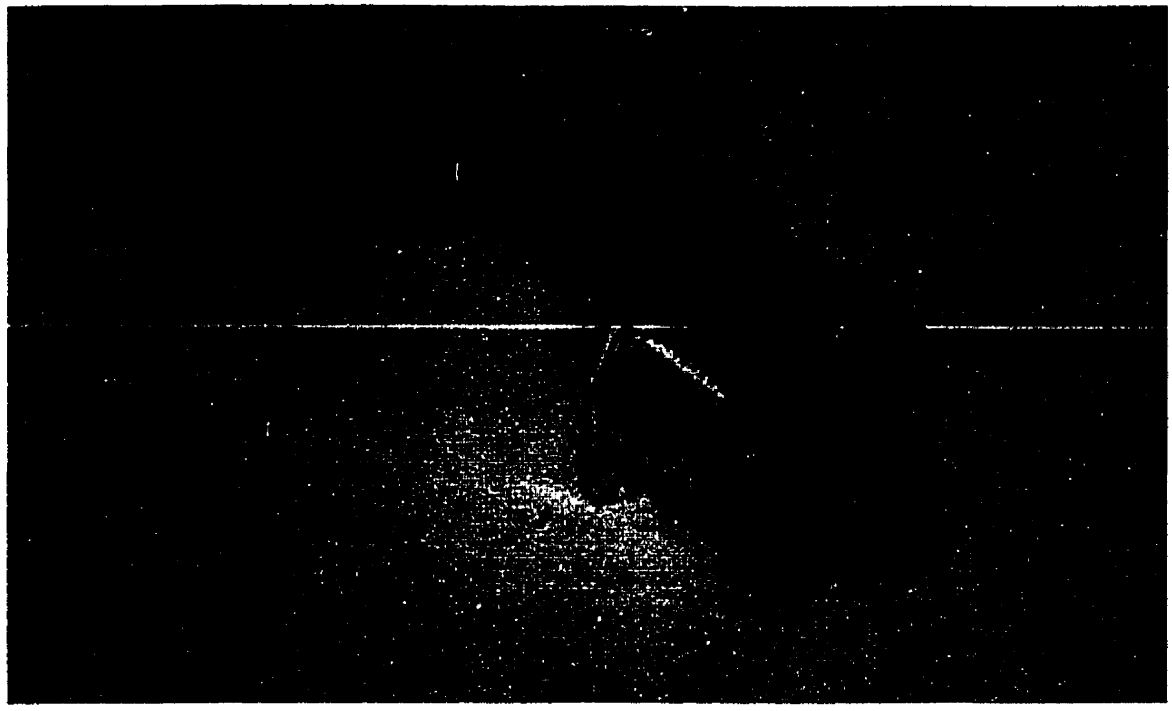

4.8.5 Electrical characterization

The test board allows high-speed connection via SMA-type connectors to a representative number of inputs on the packaging. Using a network analyser, the performance of these specific traces could be measured. A description of these

traces is shown in Table 4.3. Each of these traces was designed to have a

Table 4.3: Test traces for electrical parameters

Trace (name)	Termination	Description
ELECINFAR	Schottky active diode termination	An high-speed input trace that must travel a long distance between connector and chip pads.
ELECINNEAR	Schottky active diode termination	An high-speed input trace that must travel a short distance between connector and chip pads
ELECINNOTERM	None	An average, unterminated trace (used for control)
CLOCK	50 ¾ to 2.5 V	Very high-speed clock
RECCLK	50 ¾ to 2.5 V	Very high-speed clock

characteristic impedance of 50 ¾..

Individual trace performance is to be measured using the time domain reflectometer features of the network analyser, and an analysis of the characteristic impedance of the traces was obtained using the Smith chart features of this same instrument. These results are presented below.

The average characteristic impedance of these traces was 50 ¾. The peak deviation from 50 ¾. was observed in the clock traces, with a 20% error distributed over the entire length (from 41 ¾. to 60 ¾.) The resistance of the traces, measured from test board connector through to the end of the trace on the packaging, was less than 5 ¾ in all cases.

The delay of the traces, measured from the test board, was approximately 2 ns. With calibration of the network analyser velocity factor for the exact length of each trace, the relative permittivity of the stackup, combining both coverlayer and dielectric permittivities, was 4.7. This relationship is derived from:

$$v = \frac{c}{\sqrt{\varepsilon_R}} \tag{4.1}$$

where v is the velocity in the medium, c is the speed of light in a vacuum, and ε_R is the relative permittivity of the medium, such that:

$$\varepsilon = \varepsilon_0 \times \varepsilon_R \tag{4.2}$$

The permeability is assumed fixed at:

$$\mu = \mu_0 \tag{4.3}$$

Solving for a representative trace, given its exact length (available in the Mentor Graphics CAD environment):

$$\varepsilon_R = \left(\frac{c \times t}{l}\right)^2 \tag{4.4}$$

where l is the length of the trace, t is the transit time obtained by measurement with the TDR.

To test the transmission of the clocks into the termination, an HP80000 stimulus machine generated a 3.1 V swing signal about 2.5 V (from 0.95 V to 4.05 V). Active probes, (Picoprobes, FET probes from *GGB*) were used to sample the resulting waveform at the termination on the flexible printed circuit. The input waveform had a 200 ps rise and fall time, at an operating clock frequency of 320 MHz (square

wave at 640 MBit/s). This result is shown in figure 4.19. This waveform has a 53%

Figure 4.19: Transmitted clock signal at 320 MHz

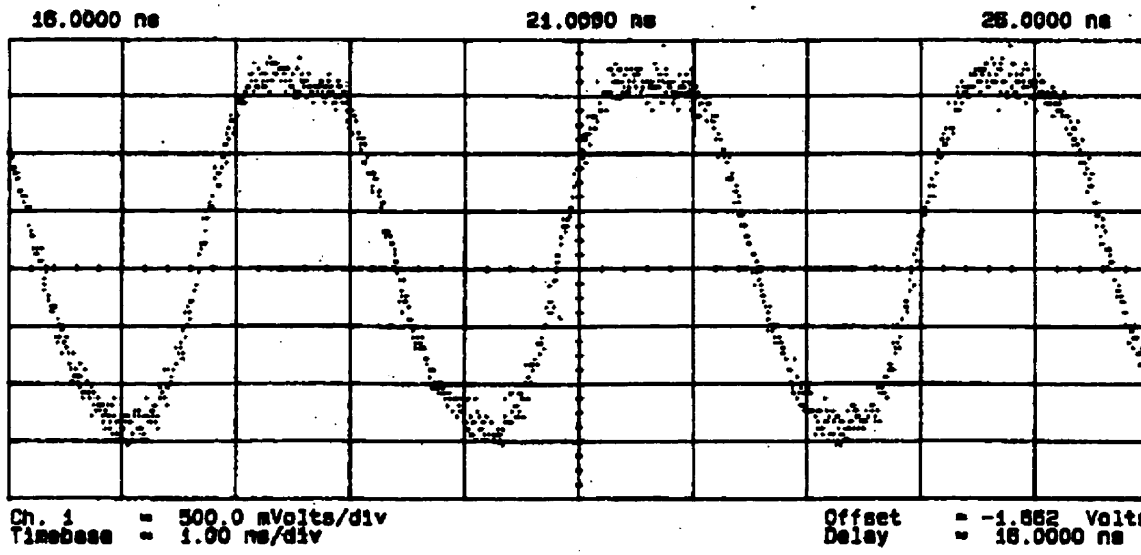

duty cycle compared to the input 50% duty cycle, and has a rise time of 1.3 ns and a fall time of 1.5 ns, making it nearly sinusoidal.

The TDR results for all the representative signals are presented in figures 4.20 through 4.24. In each case, the first peak represents the lumped 190 pin connector along with the extremely close SMA connector, and the second peak represents the termination. The vertical scale is 100 mUnits/div, where 1 Unit represents a reflection over incident factor of 1.0. These values indicate the absolute value of ρ, the reflection coefficient, where a value of 1 would indicate 100% reflection (for an unterminated transmission line), and 0 indicates no reflection. In each of the TDR graphs, the peak value at the line termination is marked.

Figure 4.20: ELECINFAR TDR

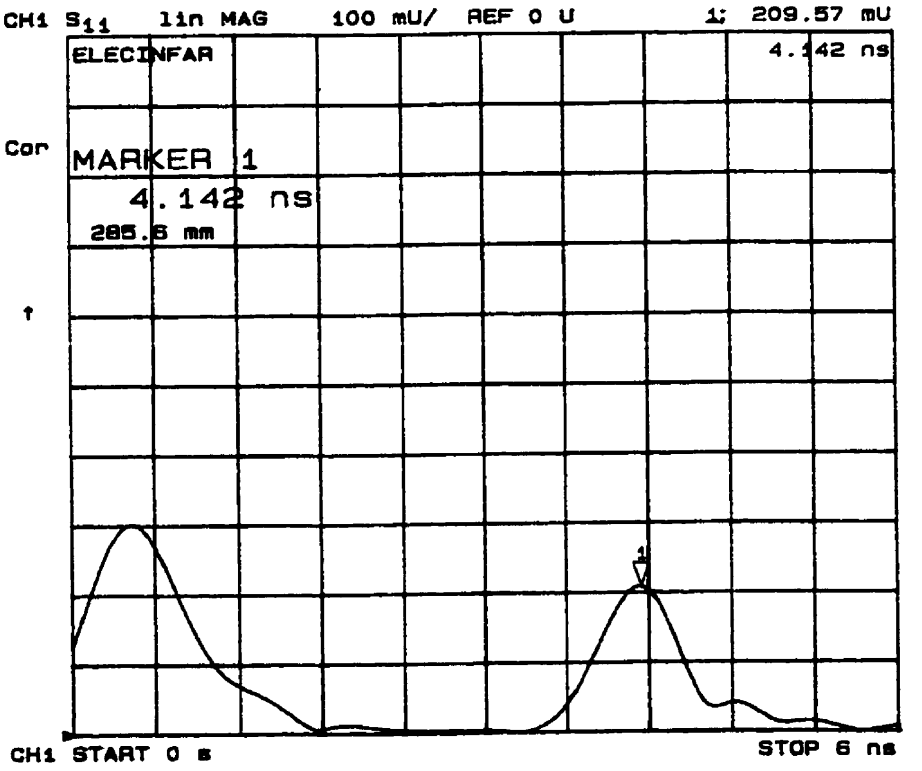

Figure 4.21: ELECINNEAR TDR

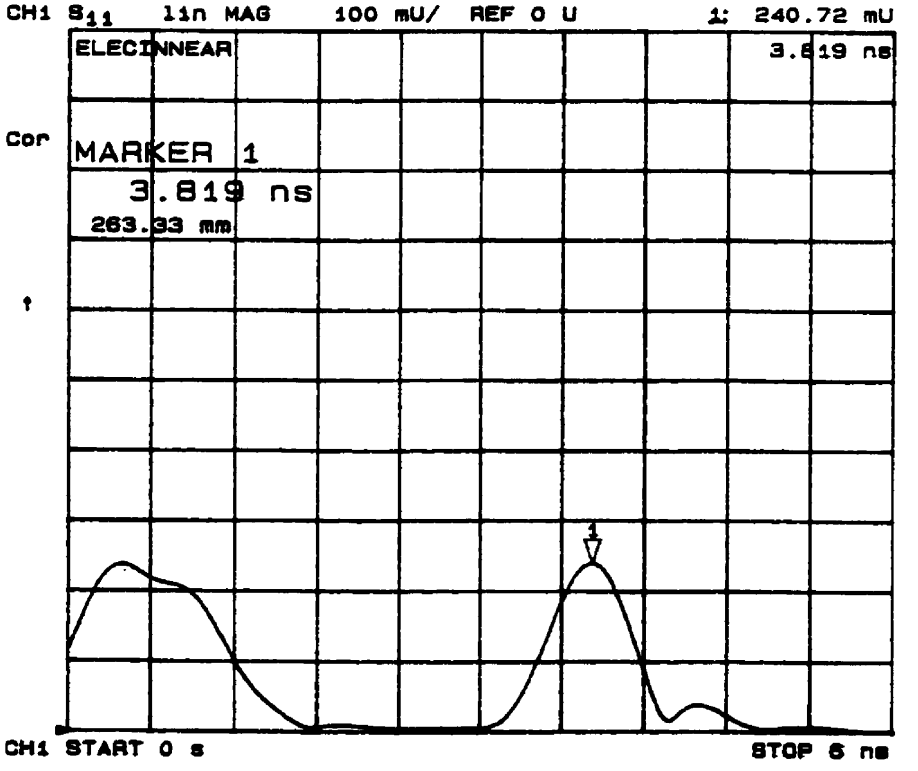

Figure 4.22: ELECINNOTERM TDR

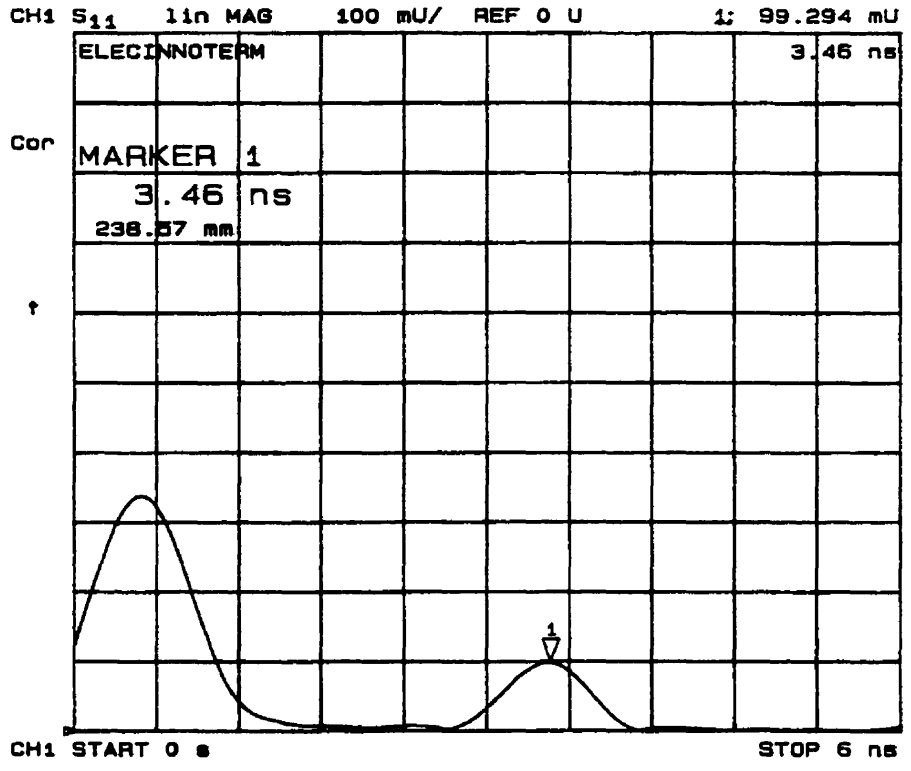

Figure 4.23: CLOCK TDR

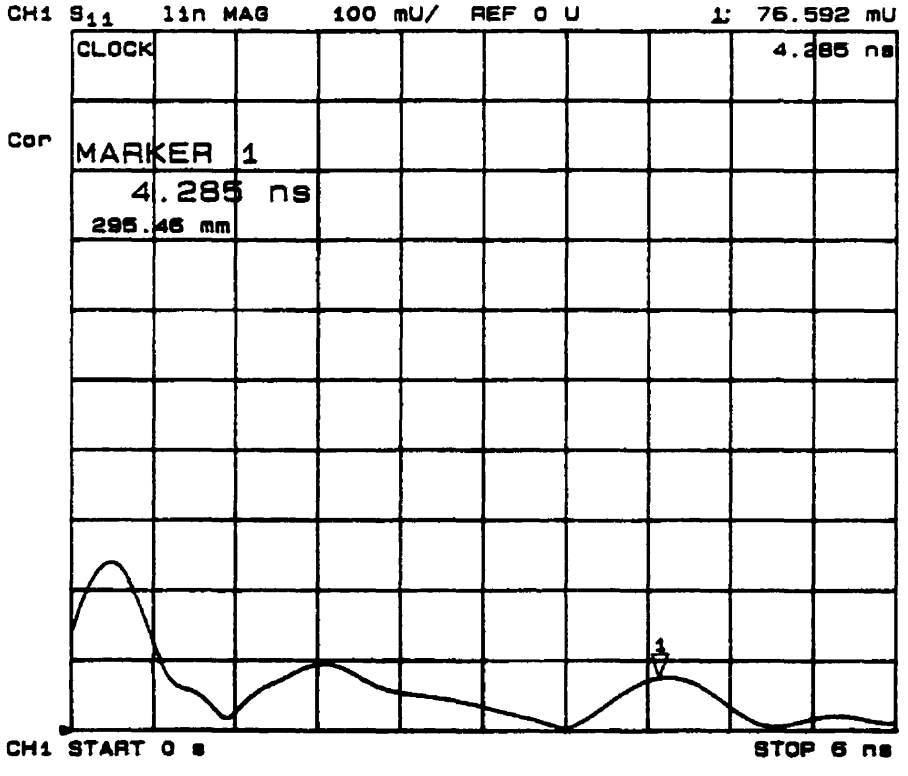

Figure 4.24: RECCLK TDR

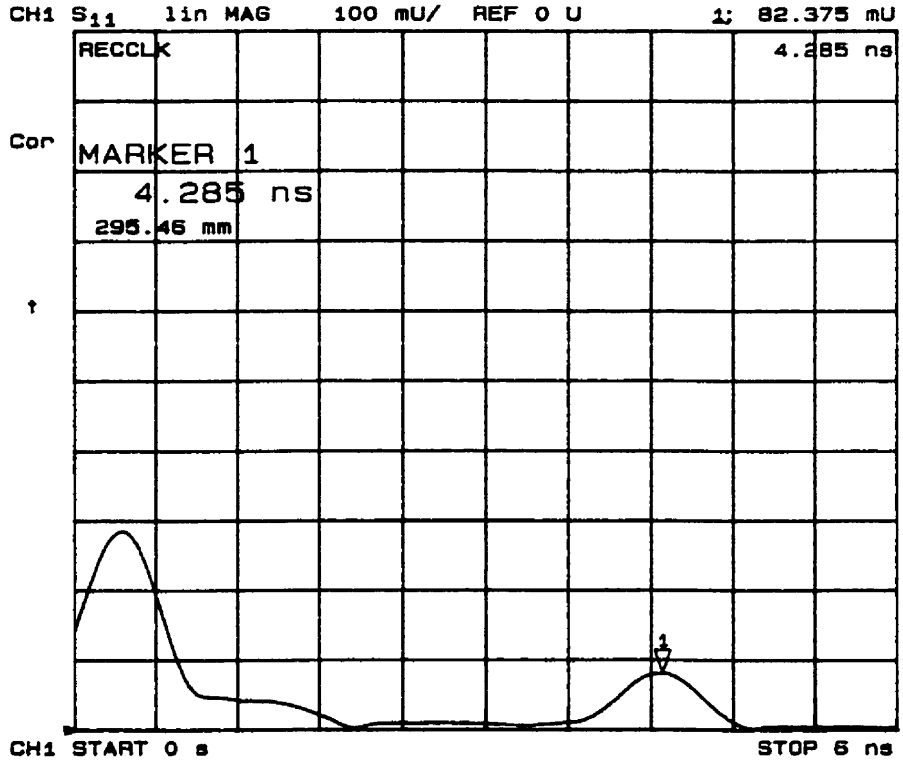

These results indicate good signal integrity over the flexible printed circuit and connectorization, with the results satisfying the system design constraints in speed and termination.

4.8.6 Thermal characterization

The thermoelectric cooler was activated, and was able to control the die surface to 0.01°C with a temperature nonuniformity of less than 0.1°C, when controlled through the computer and a thermoelectric temperature controller.

4.9 Conclusions

The *packaging module* for this demonstrator was designed, assembled and tested. It satisfies the electrical, thermal and optomechanical packaging requirements. The

system is not package bandwidth limited, and is able to effectively dissipate the expected chip power without sacrificing QCSE performance.

4.10 References

[1] Brian Robertson, "Design of a Compact Alignment-Tolerant Optical Interconnect for Photonic Backplane Applications," Proceeding of the fourth international conference on Massively Parallel Processing Using Optical Interconnections, June 22-24, 1997, Montreal, Canada, IEEE Computer Society.

[2] David R. Rolston, B. Robertson, H.S. Hinton and D.V. Plant, "Analysis of a micro-channel inteconnect based on the clustering of smart-pixel device windows," *Applied Optics*, 1996, Vol. 35, pp 1220–1233

[3] Data sheet, Epo-Tek H20E, Electrically conductive, Silver Epoxy, Epoxy Technology Inc.

[4] LDT 5910B Instruction manual, Thermoelectric Temperature Controller, ILX Lightwave, June 20, 1997

[5] Lab-NB data acquisition board manual, National Instruments, 1995.

Chapter 5: Temperature sensors

5.1 Introduction

For several classes of optoelectronic devices, including QCSE MQW diodes, the performance of emitter, modulator or detector elements is dependent on accurate control of the device temperatures. New optoelectronic packaging techniques are addressing this particular aspect of the performance, but diagnostic tools are needed to characterize temperature sensitivity, uniformity, and to actually control the device temperatures.

In this section, an array of temperature sensors integrated into a two-dimensional BiCMOS-QCSE optoelectronic device array, the Phase III chip, is presented, as well as results indicating its performance.

5.2 Temperature sensor description

In section 2.4, the temperature dependence of QCSE modulators and P-I-N detectors was introduced and optimization and control applications were presented. From this discussion, it is clear that direct measurement of the modulator and detector temperatures would be a valuable diagnostic tool, especially in cases where interconnect applications are being considered. The following section describes the design and testing of a temperature sensor array directly integrated into the Phase III chip. Using a clustered optical window design [1], the BiCMOS-QCSE chip has an 8 x 8 array of clusters. A temperature sensor has been located under each of these QCSE clusters as shown in figure 5.1, allowing a direct

Figure 5.1: Temperature sensor placement

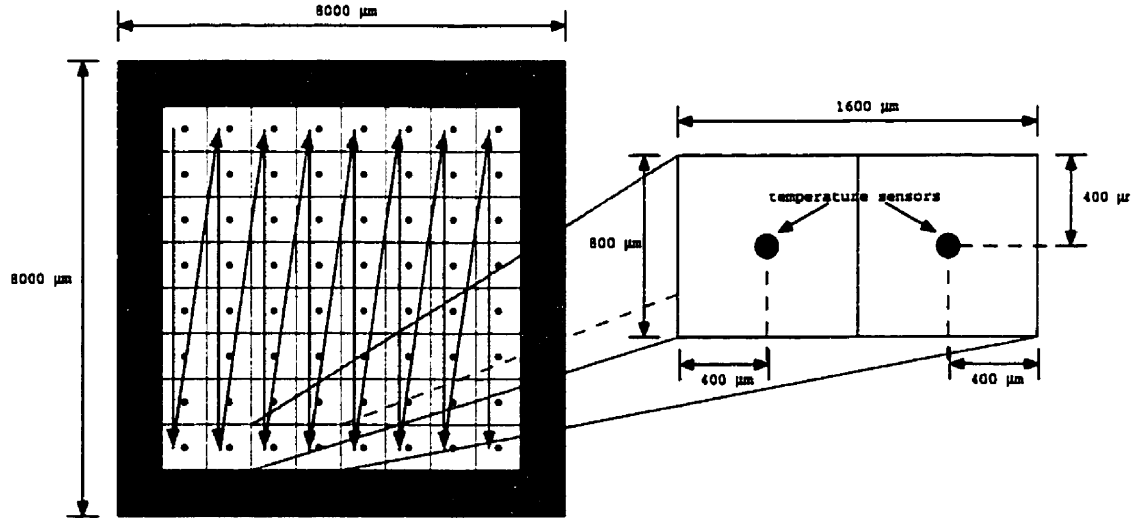

measurement of the surface temperature of the BiCMOS device at that point.

Due to bond pad limitations, the entire array is accessed via five bond pads and the chip ground, in a scanning architecture. Each temperature sensor is turned on in turn, and its output connected to a common bond pad, while all other temperature sensors are disabled. Scanning is accomplished in a columnwise fashion by a 6 bit counter and quad 3-to-8 bit demultiplexers, under the control of a *Clock* and *Reset* signal, which advance to the next temperature sensor and restart the the scan at the first temperature sensor, respectively.

5.2.1 Design

The temperature sensor circuitry is shown in figure 5.2. The (bootstrap) current

Figure 5.2: Sensor circuitry

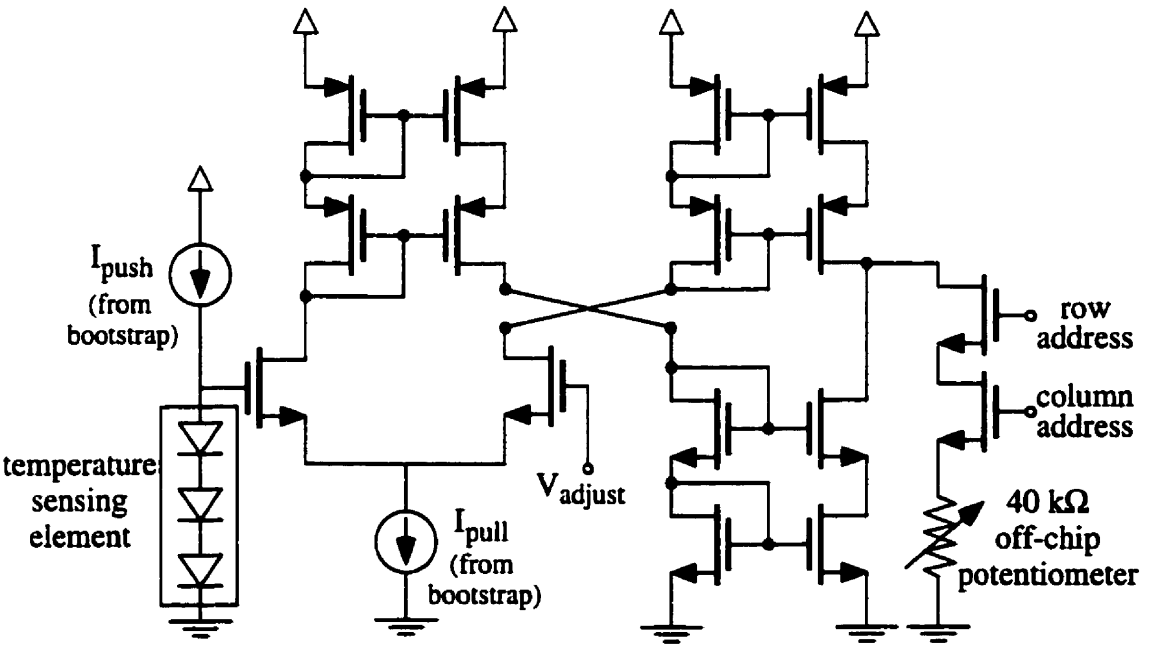

sources for the circuit are shown in figure 5.3. The sensor uses three diode-connected *npn* transistors to form the temperature sensing element. Using a differential amplifier, the difference between this voltage and an externally applied

Figure 5.3: Bootstrap current source

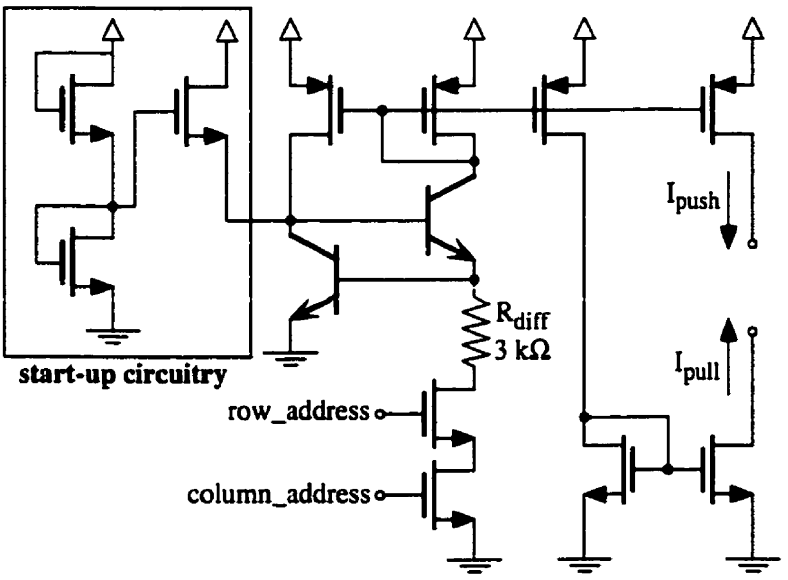

V_{adjust} is turned into a pair of currents which are multiplied by current mirrors (by a factor of 4). The difference between these multiplied currents appears at the temperature sensor output that is common to all the sensors. Both the row and column address lines of a given sensor must be active (high) for the bootstrap current source and the output to be active.

All the temperature sensor circuitry except for the column and row address lines is independent and localized to the temperature sensor, confining local temperature effects to the circuitry of one sensor. The sensor circuitry is covered on the topmost metal layer in the chip design, thereby protecting the sensor from light intrusion and therefore sensitivity. The complete temperature sensor layout is shown in figure 5.4. The overall dimensions of the temperature sensor, in 0.8 μm Nortel BiCMOS technology, are 186.4 x 42.3 μm, in total representing less than 3 % (1.7 mm^2) of the

Figure 5.4: Temperature sensor layout

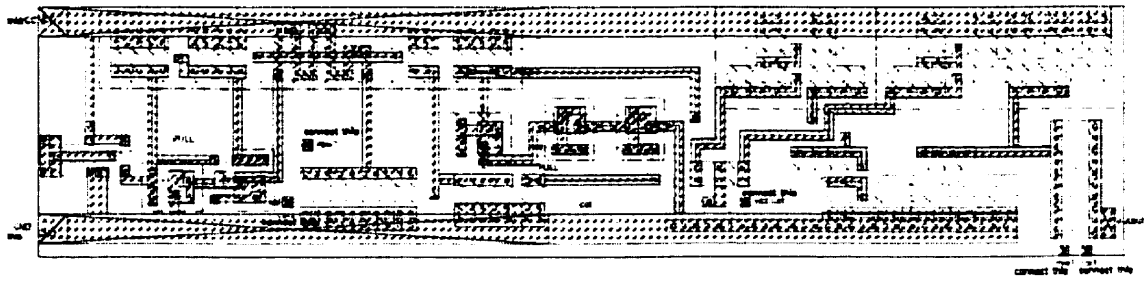

total chip area, including the required bond pads, scanning architecture and routing.

The design allowed for a 3 V swing on the output before saturation, when the output temperature sensor current is dropped over a 40 k¾ resistor. The externally applied V_{adjust} allows the temperature sensor range of sensitivity to be modified. The graphs of simulated and experimental results for values of V_{adjust} ranging from 2.2 to 2.8 volts in 0.1 V increments for the array are shown in figure 5.5. These graphs indicate that the sensor performance agrees very closely with the results predicted by the SPICE simulation of the circuit.

5.3 Experimental setup and temperature calibration

For the experimental measurements shown in figure 5.5, a prototype wiring board was built, with a precision voltage reference (*Linear* LT1021) supplying the +5 V power supply and unity-gain op-amps (*National Semiconductor* LM6134) buffering the inputs and outputs to the chip, which was mounted in a 68-pin PGA cavity, along with a calibration thermistor. A potentiometer was used for the 40.0 k¾

Figure 5.5: Simulated and experimental results for temperature sensors

Simulated (HSPICE)

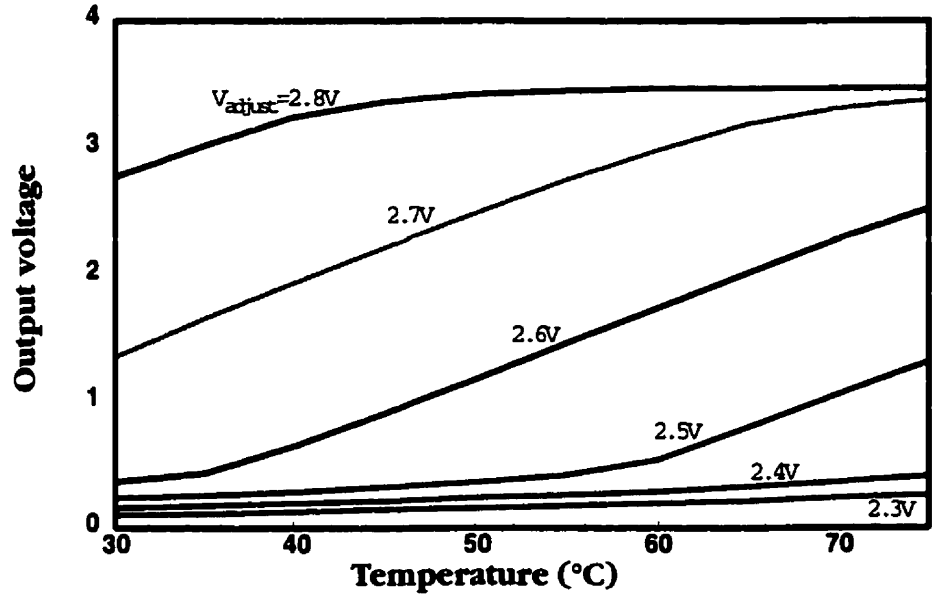

Experimental

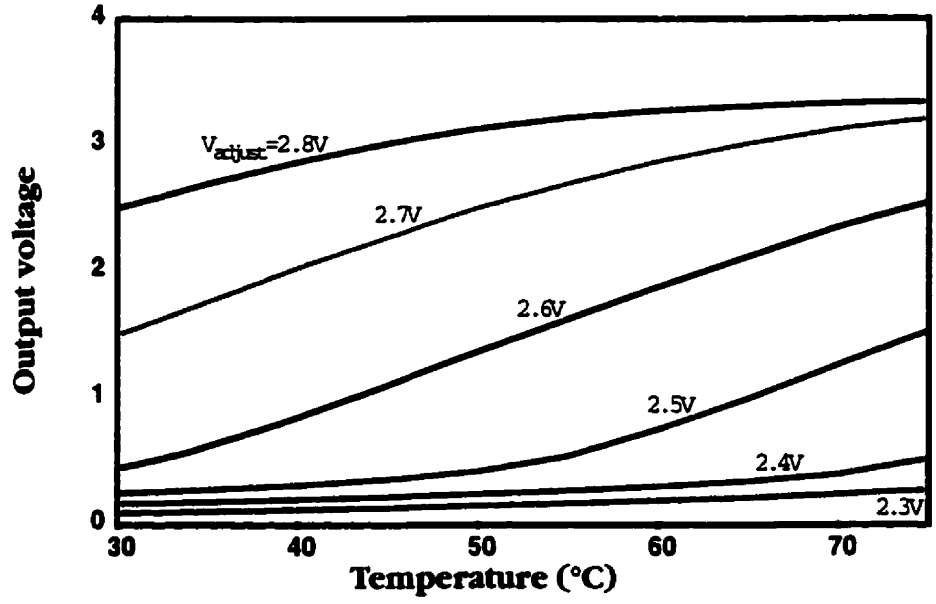

output, to change the current output into a voltage. The resulting output voltage from the array was sampled by a Macintosh computer equipped with a I/O board (*National Instruments* Lab-NB) and in-house custom software. The software

provided both the *Clock* and *Reset* signals and synchronized the sampling process. The software also provided a graphical user interface (GUI) for the array's temperature display, testing and calibration. A screen shot of this software in operation is shown in figure 5.6.

Figure 5.6: Screen shot of sensor software

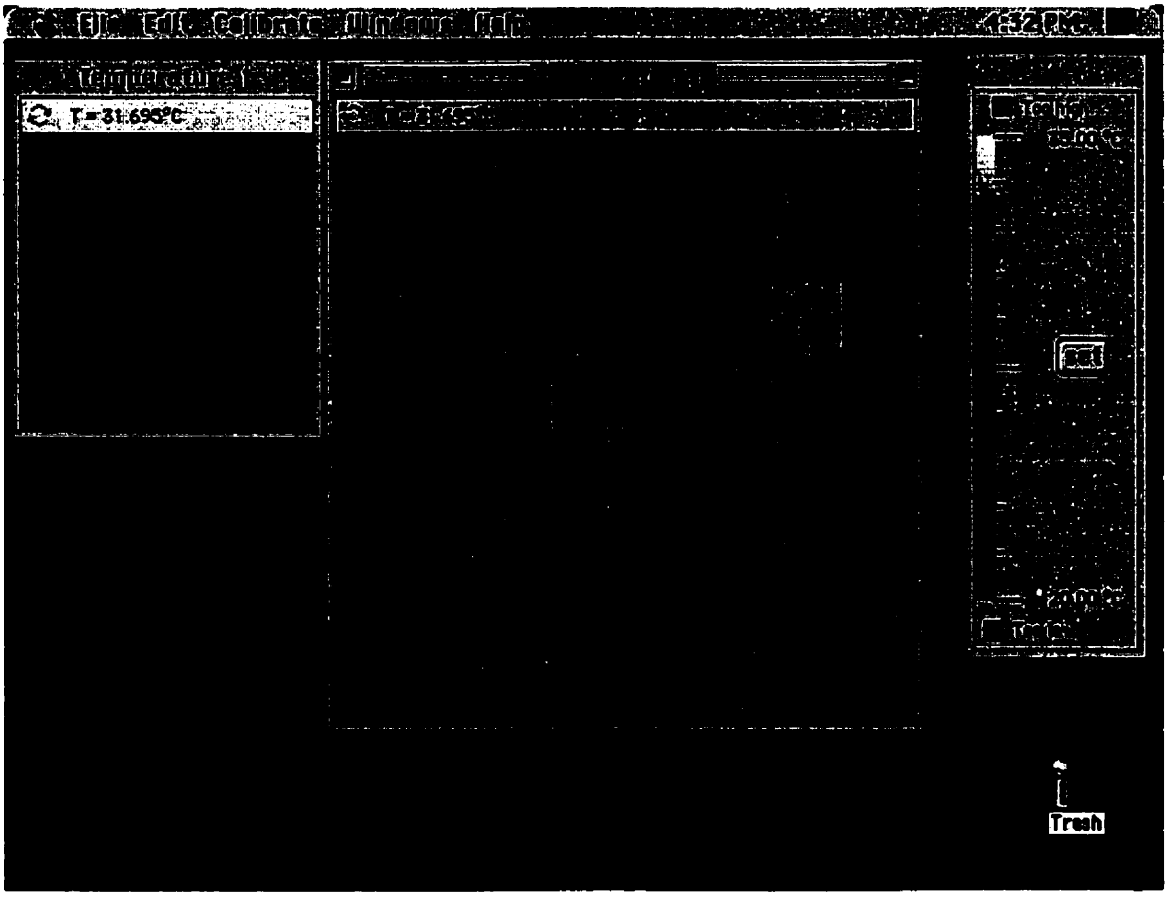

For calibration, each temperature sensor's voltage output was acquired at a given temperature, as measured by a precision chip thermistor (*Betatherm* 10KA3). The thermistor was located within the same cavity as the BiCMOS chip, for the purposes of calibration. The software subsequently converts each temperature sensor's output voltage into a temperature, using linear interpolation between calibration data

points. This calibration and measurement process allows for variations in the manufacturing process across the die, although no significant variations in the output voltages (<0.002 V) were observed. The use of a precison voltage reference to supply the chip and op-amp power generated <0.02 °C repeatability for a temperature measurement (for repeated power cycling).

For the system integration, the average value of the array is converted into a simulated LM335 temperature sensor output by the computer software. This voltage value is provided to a temperature controller where a thermoelectric cooler's output current can be controlled.

5.4 Conclusions

This application shows that it is possible to incorporate a temperature sensor array within the structure of an advanced optoelectronic device array, without interfering with the chip's primary function as a communication array for an interconnect. This diagnostic tool has been validated by experimental verification of a temperature sensor array integrated into the functioning optoelectronic device array.

5.5 References

[1] David R. Rolston, B. Robertson, H.S. Hinton and D.V. Plant, "Analysis of a micro-channel inteconnect based on the clustering of smart-pixel device windows," Applied Optics 35, pp 1220–1233 (1996).

Chapter 6: Conclusions and future work

6.1 Two dimensional optoelectronic device array packaging

In chapters 1 and 2, the background of and constraints for an optically interconnected backplane were presented, as well as the design methodology that can be used to satisfy these constraints. Specific past approaches are discussed, as well as promising future approaches that have not yet been undertaken in their entirety. A new optimization procedure is developped for a specific class of optoelectronic devices, QCSE modulators and detectors, that can be used to optimize the performance of an optical interconnect, when coupled with thermal control and stabilization techniques. The specific issues inherent to the electrical connectivities of an FSOI backplane are highlighted, where electrical connectivity and bandwidth form the dominant areas of concern.

In chapters 3 and 4, system demonstrator packaging was presented that specifically addresses these constraints and demonstrates the use of the approaches. In chapter 3, a 1996 packaging system for a optical backplane demonstrator is presented that integrates chip-on-board packaging techniques with optomechanical interfacing, for a 4 board interconnect. The packaging techniques used are expanded and farther restricted in chapter 4, where a larger optoelectronic device with more difficult optical and electronic constraints is integrated into a smaller packaging system. A recurring theme in these two chapters is the task of satisfying constraints. It is clear that these two packaging systems satisfied the system integration constraints specific to system development, in terms of space, thermal management, bandwidth, optical interfacing and connectivity. In no case was the packaging limiting to the speed or operation of the backplane, ensuring that a design bottleneck was not introduced in the packaging of the optoelectronic device. Making these guarantees in terms of system integration required the careful consideration of the design issues and an application, in each case, of a new or novel technology to the packaging. In both demonstrator systems, the packaging used borrowed from previous electronic and optoelectronic packaging techniques and expanded them or modified them to meet the specific constraints of the system.

6.2 Future and ongoing work

The optomechanical aspects of a packaging system are crucial to its success and feasibility when applied to an optical backplane demonstrator. Only token regard is paid here to this aspect of the packaging, and only where a discussion of the

optomechanics was vital to the explanation of the electronic or mechanical packaging constraints. The bulk of this work focussed on the electronic, mechanical and optical constraints, but future work by researchers within the Photonic Systems Group should demonstrate the optomechanical viability of the packaging presented here.

Current optoelectronic devices have not yet stretched the packaging techniques presented here to their limit, but much work is needed to ensure that this will not happen. Chips will require additional optical and electrical requirements to "fill" their available speed, and further packaging work is needed to make this happen. Without an eye on the future of optoelectronic devices, as well as an eye on current electronic packaging techniques, it is apparent that a future optoelectronic device could easily exceed the bandwidth requirements of the demonstrators presented here, as well as the ability of the packaging to satisfy these needs. An important part of ensuring the viability of optical backplanes is the ability to say, at every point in the design process, that the packaging system is not a limiting factor, a constraint unto itself, or a bottleneck.