


AFRL-SN-WP-TR-1998-1055

**OPTICAL INTERCONNECT
TECHNOLOGY (OIT)
Multichip Module to Multichip Module**



Dr. Julian Bristow
Dr. Yue Liu (612-951-7843; fax: 612-951-7438; email: liu@htc.honeywell.com)
Dr. Klein Johnson

HONEYWELL TECHNOLOGY CENTER
3660 TECHNOLOGY DRIVE
MINNEAPOLIS, MINNESOTA 55416

JANUARY 1998

FINAL REPORT FOR PERIOD AUGUST 13, 1992 – JUNE 30, 1996

19990506 003

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION IS UNLIMITED

SENSORS DIRECTORATE
AIR FORCE RESEARCH LABORATORY
AIR FORCE MATERIEL COMMAND
WRIGHT-PATTERSON AIR FORCE BASE, OH 45433-7623


DTIC QUALITY INSPECTED 4


NOTICE

USING GOVERNMENT DRAWINGS, SPECIFICATIONS, OR OTHER DATA INCLUDED IN THIS DOCUMENT FOR ANY PURPOSE OTHER THAN GOVERNMENT PROCUREMENT DOES NOT IN ANY WAY OBLIGATE THE US GOVERNMENT. THE FACT THAT THE GOVERNMENT FORMULATED OR SUPPLIED THE DRAWINGS, SPECIFICATIONS, OR OTHER DATA DOES NOT LICENSE THE HOLDER OR ANY OTHER PERSON OR CORPORATION; OR CONVEY ANY RIGHTS OR PERMISSION TO MANUFACTURE, USE, OR SELL ANY PATENTED INVENTION THAT MAY RELATE TO THEM.

THIS REPORT IS RELEASABLE TO THE NATIONAL TECHNICAL INFORMATION SERVICE (NTIS). AT NTIS, IT WILL BE AVAILABLE TO THE GENERAL PUBLIC, INCLUDING FOREIGN NATIONS.

THIS TECHNICAL REPORT HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION.


GUY D. COUTURIER
Multi-Chip Integration Branch
Aerospace Components Division


ALAN J. TEWKSBURY, Chief
Multi-Chip Integration Branch
Aerospace Components Division


PAUL H. OSTDIEK, LT COL, USAF
Deputy, Aerospace Components Division
Sensors Directorate

Do not return copies of this report unless contractual obligations or notice on a specific document requires its return.

REPORT DOCUMENTATION PAGE*Form Approved*
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE January 1998	3. REPORT TYPE AND DATES COVERED FINAL AUGUST 13, 1992 - JUNE 30, 1996
4. TITLE AND SUBTITLE Optical Interconnect Technology (OIT) Multichip Module to Multichip Module		5. FUNDING NUMBERS C: F33615-92-C-1034 PE: 62712E/62301E/63739E PR: 2001 TA: 02 WU: AM	
6. AUTHOR(S) Dr. Julian Bristow, Dr. Yue Liu, Dr. Klein Johnson			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Honeywell Technology Center 3660 Technology Drive Minneapolis, Minnesota 55418		8. PERFORMING ORGANIZATION REPORT NUMBER C98057	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Sensors Directorate Air Force Research Laboratory Air Force Materiel Command Wright-Patterson AFB OH 45433-7623 POC: James G. Grote, AFRL/SNDM, 937-255-4557 x3372		10. SPONSORING/MONITORING AGENCY REPORT NUMBER AFRL-SN-WP-TR-1998-1055	
11. SUPPLEMENTAL NOTES			
12A. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.		12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) We have successfully demonstrated the first MCM-to-board-to-MCM optical interconnect based on VCSELs, polymer waveguide, and MCM-C packaging technologies. We demonstrated low-cost component fabrication and passive assembling techniques that are compatible with existing electronic manufacturing processes. We believe these approaches are critical to optical insertion into real system applications. We fabricated flexible polymer waveguide ribbons, board-integrated optical waveguides, and passively aligned flex-to-board waveguide connectors. We also demonstrated optical links based on conventional MCM-C packages with multiple data channels at data rates up to 1 Gbps per channel.			
14. SUBJECT TERMS optical interconnect;polymer waveguide;MCM		15. NUMBER OF PAGES 51	
		16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT SAR

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102

Table of Contents:

1	SUMMARY	1
2	INTRODUCTION TO OIT	2
2.1	INTERCABINET INTERCONNECTS	2
2.2	INTRACABINET INTERCONNECTS	4
3	PROGRAM GOALS	8
4	APPROACH TAKEN FOR OIT	9
5	INTERCONNECT MODELING AND DESIGN	10
5.1	WAVEGUIDE STRUCTURE	12
5.2	VCSEL-TO-WAVEGUIDE COUPLING	14
5.3	INTERCHANNEL CROSSTALK	14
5.4	RADIATION LOSSES DUE TO WAVEGUIDE CURVATURE	16
5.5	MODAL NOISE PHENOMENA AND THEIR IMPACT ON OIT	18
5.6	MODELING AND CHARACTERIZATION OF MODAL NOISE	19
6	POLYMER WAVEGUIDE DEVELOPMENT	22
6.1	POLYMER CHOICES FOR OPTICAL INTERCONNECTS	22
6.2	POLYMER IDENTIFICATION	22
6.3	POLYMER SELECTION	24
6.4	CHANNEL WAVEGUIDE PROPAGATION LOSS MEASUREMENT	25
6.5	FLEX WAVEGUIDE FABRICATION	26
6.6	BOARD LEVEL WAVEGUIDES	28
6.7	ENDFACE PREPARATION	30
6.8	BOARD-TO-FLEX CONNECTOR DESIGN	30
6.9	ENVIRONMENTAL TESTING	34
7	OPTICAL MULTICHIP MODULE DEVELOPMENT	36
7.1	MCM INTEGRATION	36
7.2	HDI WAVEGUIDE PROCESSING	37
7.3	HDI PACKAGING	38
7.4	HDI VCSEL PACKAGING	38
7.5	MECHANICAL PREPARATION OF END-FACES	41
8	OPTICAL INTERCONNECT DEMONSTRATION	43
8.1	TRANSMITTER (TX) AND RECEIVER (RX) MODULE:	43
8.2	OIT LINK BUDGET	44
8.3	OIT DEMONSTRATION SETUP	44
8.4	OIT DEMONSTRATION TESTING RESULTS	46
9	CONCLUSIONS	47

Table of Figures:

Figure 1 The complete interconnect modeled for OIT	11
Figure 2 The basic waveguide structure used for OIT	12
Figure 3 Selected modes of an Ultem/BCB waveguide system	13
Figure 4 Waveguide mask layout for evolution of waveguide bend losses	17
Figure 5 Measured bend losses for in-plane waveguide bends	18
Figure 6 Experimental arrangement used to investigate modal noise	20
Figure 7 A typical waveguide output showing predominantly low order excitation of multimode polymer waveguides	21
Figure 8 Experimentally derived optical characteristics for various polymer waveguide interconnect media	24
Figure 9 Process sequence for fabrication of the flexible polymer waveguide interconnect	27
Figure 10 A cross-section microphotograph (left), and a photograph of optical waveguides fabricated on a flexible Kapton carrier	27
Figure 11 Lamination of flexible waveguide layers into rigid printed wiring board materials	29
Figure 12 The connector assembly of a flexible waveguide layer with board-laminated and MCM level waveguides	31
Figure 13 A cross-section of the alignment features before (left) and after they are engaged	32
Figure 14 An assembled MCM-to-board connector	33
Figure 15 Measured loss of the flexible board-to-MCM connector	34
Figure 16 Schematics of MCM fabricated using GE's HDI process to implement both electrical (left) and optical (right) interconnects	38
Figure 17 An HDI packaged VCSEL array	39
Figure 18 Etching of polymer waveguides using laser ablation	39
Figure 19 Preparation of 45-degree angled facets	40
Figure 20 Mechanical preparation of angled end-faces	41

Figure 21 The test arrangement for evaluating the VCSEL-to-45-degree reflector coupling.	42
Figure 22 OIT Tx (upper) and Rx design schematics.	43
Figure 23. OIT Tx (left) and Rx (right) MCM and active devices.....	43
Figure 24 OIT link setup schematic and key link components and interfaces.	45
Figure 25 OIT demonstration link testing setup.	45
Figure 26 OIT link eye diagram at 1 Gbps.	46
Figure 27 OIT link demonstration: two link channel operating simultaneously at data rate of (a) 250 Mbps 16-bit fix pattern, and (b) 622 Mbps of PRBS.....	46

1 SUMMARY

The objectives of the Optical Interconnect Technology program were to demonstrate multichip module (MCM) to MCM optical interconnects to eliminate system I/O bottlenecks and hence to enable a tenfold increase in the throughput of backplane based systems. The program also set out to demonstrate cost effective packaging of optoelectronic components in standard multichip module technologies, and to demonstrate practical integration of polymer waveguides with standard printed circuit boards. A final goal was to demonstrate an optical interface from board laminated waveguides to multichip modules.

The approach adopted was to provide a passive multimode polymer waveguide interconnect which could be implemented using standard printed wiring board techniques, and which could also be incorporated in a standard multichip module fabrication scheme. In other words, the approach was to add optics to electronic systems by exploiting the existing infrastructure, rather than to replace all elements of an electronic system. In this respect, the resulting system would be an intermediate solution for electronic system integrators, retaining compatibility with existing system assembly practices and yet providing at least an order of magnitude improvement in system throughput. Longer term optical solutions might gain yet more benefit from optics by eliminating many of the currently accepted design philosophies, but such systems would be unlikely to find acceptance in the time frame addressed by this program.

Accomplishments in the program include the demonstration of the optoelectronic packaging in standard multichip modules using established processes. Integration of optical interconnects into standard printed wiring boards using established processes was also demonstrated. A new, easy to assemble optical interface between a multichip module and board-laminated waveguides was also demonstrated. This interface allows optically interconnected multichip modules to be assembled with the same ease and using the same assembly tools as are currently used to integrate electrically interconnected multichip modules into electronic systems. Compatibility of all components with military and commercial environments was demonstrated. In particular, the board laminated waveguides were evaluated for degradation after cycling according to a military standard board test. No degradation was noted as a result of exposure to these tests. Finally, a complete multichip module to multichip module interconnect using polymer waveguides was demonstrated. To our knowledge, this represents the first such demonstration reported.

2 INTRODUCTION TO OIT

Since the concept was first proposed, optical interconnects have offered the possibility of eliminating certain bottlenecks in electronic systems. Early uses of optical interconnects focused on the replacement of long (tens of kilometers) copper cables with lighter, more compact optical fiber media. The same considerations of installation cost, bandwidth, cable volume, and maintenance which drove the replacement of copper with fiber are now motivating the use of optical interconnects at progressively shorter distances in electronic systems.

Initial focus was on the provision of higher bandwidth than was attainable with state-of-the-art electrical interconnects, while more recent activity has extended the focus to encompass the provision of higher density interconnects following the realization that electrical interconnects may generally be adequate on the basis of bandwidth alone. At long distance, optics has been addressing near-term limitations in electrical interconnect bandwidth for cabinet-to-cabinet applications. As the interconnect distances become shorter, the limits of electrical interconnect bandwidth are not reached within a single cabinet, at least not for individual interconnects, however, optics offers advantages in the maximum attainable interconnect density. Electrical standards are emerging for gigabit per second interconnects over distances as great as 25 metres, while longer distances or greater bandwidths are possible if the characteristics of the copper are augmented by some form of signal processing. Individual electrical interconnects such as twisted pair are however bulky, and printed wiring boards do not offer bandwidth-length products as high as twisted pair.

Nonetheless, interconnect distances of up to a metre are feasible using printed wiring board technologies and data rates to several hundred megahertz, albeit with limited density. Miller has defined a fundamental aspect ratio which indicates a maximum usable bandwidth-length product for various interconnect media. For individual interconnects, these aspect ratios are barely reached at the backplane level, and therefore not reached on individual boards implemented with the same PWB technology. At the backplane level, where interconnect densities are more critical than bandwidth, optics does however facilitate denser packaging of electronic circuits within an enclosure. Thus, optics is not justified on the basis of bandwidth alone at distances of a few meters. As interconnection distances are progressively decreased to the board level or individual package level, it is the density of optical interconnects which is of more impact than the bandwidth. Interconnect dimensions found in the backplane are also encountered in individual packages, thus on-board electrical interconnects also provide adequate bandwidth at this level.

Only in cases in which optics adds some new functionality is it likely that radically new implementations will be accepted. In other cases, the cost of addition of optics must be minimized by maintaining compatibility with the existing electronic system, as it will be traded against the incremental performance enhancements it offers. This applies both to interconnects between different cabinets of an electronic system and to interconnects within those cabinets.

2.1 INTERCABINET INTERCONNECTS

Some of the first demonstrations of optical interconnects involved the replacement of some form of copper cable with fiber. The problem to be solved was one of providing a bandwidth unattainable using contemporary electrical technology. Additional benefits propounded by the optical community included reductions in cable volume and weight, with the latter being critical for applications such as avionics local area networks, and being convenient, if nothing else, for many other applications. The limits of electrical interconnect bandwidths are reached earlier for longer interconnection distances, and the characteristics of interconnect media are better described by a (bandwidth x length) product.

In providing a specified aggregate bandwidth over a given length, the designer may choose to implement a large number of parallel interconnects at low speed, or serialize to a smaller number of higher speed interconnects. Once the bandwidth requirements exceed the capabilities of a given medium, the designer is forced into adopting parallel solutions. Optical interconnects offer extremely high bandwidths, and are exploited heavily by the long-haul telecommunications industry. At a single wavelength, bandwidths of 10 GHz and beyond are used, while wavelength division multiplexing techniques multiply these figures by the number of available wavelengths. A simple approach for electronic systems might be to serialize as many signals as possible into a single high-speed channel. Unfortunately, the issue of cost prevents adoption of long-haul telecommunication techniques in most electronic equipment. In long-haul telecommunications, the cost of installing a fiber may be thousands of dollars per kilometer, compared to which a \$500 transmitter is negligible. To look at this another way, the number of transmitters and receivers per unit interconnect length mandated by the architecture is low. In an electronic system consisting of a number of processors interconnecting to each other and to memory and peripherals, one transmitter and receiver is needed per interconnect, which may only be a few meters long. The maximum data rate in a serial link will therefore be dictated by the lowest cost for the complete interconnect solution.

Fiber does not save cost, therefore in comparing fiber to copper, the lowest cost solution will be that which minimizes the cost of electronics for all data rates below the maximum permitted by a serial electrical interconnect. The highest serial data rate attainable from a single serialization circuit is the same whether the interconnect medium be copper or fiber. Furthermore, optical interconnects will always require the addition of optoelectronics to an electronic system, thus putting optics at a disadvantage. Intercabinet interconnects will therefore be implemented in fiber only when the weight or volume reduction is important, or for interconnect distances exceeding those available from electrical media, or if the cost becomes comparable to copper.

Two areas of focus are likely to reduce the cost differential for using fiber. The first is that of parallel optical links in which the packaging and connectorization costs are amortized across a number of fibers, while the second addresses lower aggregate data rates by reducing the cost of discrete fiber interconnects by using plastic optical fiber. A major advantage of plastic fiber lies in its large diameter. Bandwidth length products of 2-300 MHz. km are possible with plastic fiber of 600 μ m core diameter if graded-index designs are used. In glass fiber interconnects, the high cost of connectorization represents a significant fraction of the cost of an interconnect. This cost is due to the need to maintain microns of precision in the mating components.

Relaxed alignment tolerances on the larger plastic fiber allow the implementation of less expensive connectors. Large diameter glass fibers would allow the same reduction in alignment precision and hence reduce connector cost, but would themselves be more expensive to produce. Plastic fibers have long been available with large core diameters, however, only recently have the losses been sufficiently low. PMMA fibers with losses of 180 dB/km represent the current state of the art. While the losses may appear high, many applications require interconnect distances of only 100 m or less. Losses of 18 dB or less may be acceptable for VCSEL-based links.

Exploitation of the low losses of PMMA-based fiber requires a source with a wavelength of 650 nm wavelength. LEDs offer insufficient power for most interconnect distances, and furthermore have too broad a spectral output for the narrow low-loss window in PMMA. Recent developments in red VCSELs offer the hope for low cost sources and hence low cost, short-distance plastic fiber interconnects. In summary, many targeted insertions of optical interconnects have focused on interconnection distances of tens of metres and on data rates of 1 Gbps or lower. These are within the capabilities of electrical interconnects, thus only when costs become comparable will system integrators begin to deploy optics on a broad basis.

2.2 INTRACABINET INTERCONNECTS

After initial hopes for an all-optical computer faded, researchers began to focus instead on optical interconnect technologies for ultimate use within cabinets. The areas addressed include board-to-board interconnects, module-to-module interconnects, and chip-to-chip interconnects. With the realization that most logic will be implemented in electronics rather than optics comes acceptance of the fact that any optical interconnects added to an electronic system must interface to ICs at some point. Optical interconnects are therefore likely to increase the cost of a system if simply used as replacement for copper, and will therefore only gain acceptance if they offer some advantage in system design to the system integrator.

Most current research activities focus on evolutionary insertion of optical interconnects wherein copper is replaced by electrical interconnects at progressively shorter distances. To illustrate the tradeoffs made by a system designer, consider first the case of board-to board interconnects within a single cabinet. If electrical interconnects can provide gigahertz of bandwidth at tens of meters then clearly they can provide the same bandwidth at the 1 meter distances characteristic of many Intracabinet applications. As interconnection distances decrease however, the number of interconnects increases such that discrete wiring solutions become prohibitively expensive. The electronics industry solves this problem by using printed wiring boards rather than hand-wiring each interconnect in, for example, a personal computer. Considerations of interconnect density also make printed wiring boards more attractive for most systems. Electronic systems running at clock speeds of tens of megahertz to 100 MHz frequently use linear buses in the backplane. These are easy to implement in the physical, medium dependent layer, and give rise to an interconnect density which is independent of the number of nodes on the bus. Multidrop buses however are limited in their ability to provide high speed interconnects in simple implementations. As interconnect clock speeds increase, reflections from the impedance mismatches caused by the various drops require that greater fractions of a bit period be used for settling before a decision

can be made. Based on these limitations, the optics research community targeted electronic multidrop buses for replacement with optics as projected data rates increase to several hundred megahertz and beyond. In the electrical community however, the inability to implement the required bandwidth started a movement towards different architectures. Switched interconnects in which point-to-point interconnects from each element in the system interface to an $N \times N$ switch allow reconfiguration of the interconnects as needed for program execution. With the elimination of the impedance discontinuities of a linear bus comes the possibility of higher speed operation. Indeed, electrical backplanes may operate with interconnect clock speeds of 1 Gbps or more in point-to-point architectures. Switched architectures however impose greater requirements on the interconnect density. The number of channels required into and from a centralized switch scales with the number of processors. Both the wiring density in the backplane and the number of pins required at each board-to-backplane interface can limit both the logical and physical size of the system that can be implemented using a switched architecture.

Consider a system in which N processors are to be fully connected, and in which communication channels are 32 bits wide. If a centralized switch is to be used and is mounted on one board in a multi-slot cabinet, then the number of physically independent channels which must be provided is $32 \times N$ where N is the number of processors in the system. For a 32 processor system, 1024 inputs and 1024 outputs would be required for a total of over 2,000 I/O. This number of I/O is also the maximum number of channels which must be implemented in the backplane. As more processors are added, more layers are required in the backplane. Once the number of pins required in a connector exceeds the capability of electrical technology, the connector becomes the bottleneck in extending the system size, and growth can only occur if lesser degrees of connectivity are provided. The maximum number of layers in the backplane is governed by the cost of implementation, whereas the connector I/O density is governed by more practical limitations.

Electrical technology appears well placed to provide gigabit interconnects or lower data rates at distances up to tens of centimetres or metres. Aggressive printed wiring board technologies offer line width of 25 μm with comparable spacing, while the mainstream of the industry focuses on 100 μm lines and 100 μm spaces (4 mil / 4 mil). If a sufficient number of interconnect layers are integrated, this density offers adequate capability for most electronic systems. The difficulty most often encountered is in maintaining this density across the various interfaces in the system. These interfaces include those from the package to the printed wiring board, and from a line replaceable module or removable card and the backplane. The problem encountered is one of poor impedance control at the interface and of crosstalk in this ill-controlled environment. In designs in which pins mate with corresponding sockets, insertion force of 100g per pin quickly adds up to unmanageable figures for connectors incorporating hundreds of pins. Addition of more pins would require such force that removal and insertion of the board cause physical damage. Aside from the matter of insertion force, it is the crosstalk in typical connectors which limits the maximum effective pin-count. As the pins are located closer together, crosstalk increases unless pins are periodically grounded. As the data rate increases, so must more pins be assigned to ground, and so the board-edge I/O reaches an asymptotic limit. For a SEM-E connector, this figure is somewhere of the order of 1 Gbps per pin pair for a total of 700 pins, or an aggregate data rate of 350 Gbps, or one third of a terabit per second. In a chassis incorporating 16 such cards, a total of 2.8 Terabits of information per second could be transmitted per second (accounting for a path

consisting of a source and a destination board). Boards with greater width at the board to backplane interface would allow proportionately greater aggregate data rate transfers, however simply using larger boards does not solve the problem as the I/O requirements for a uniformly populated board grow with the square of the board edge dimension, whereas the I/O capability grows only linearly with this dimension.

Systems which have already encountered this bottleneck are mostly high performance embedded processors in which the maximum processing power possible is located in a given volume. Since the major problem to be solved is one of density at the board-to backplane interface, if optics did nothing but provide a higher density at this interface, electronic systems could continue to exploit advances in processor technology. The addition of optical interconnects however requires a complete set of components from electrical input to electrical output to exploit this high density, and the advantages in performance of a system incorporating optics must be weighed against the additional cost of insertion of all the required interconnect components in any application in which the maximum size is not bounded. Optical backplanes are under serious investigation for spaceborne processing where single cabinet solutions are being developed, while fiber-based backplanes are under development for telecommunications switches. Military and commercial avionics developers are beginning to examine the technology for core avionics where reconfigurable avionics would allow a smaller number of processors, or where the functionality of several cabinets can be reduced to one as the I/O bottleneck is eliminated.

To gain near term acceptance, the minimum required of optics is that it offer a higher density than is attainable with electrical interconnects and that it affect as little of the remainder of the system as possible. Several solutions have been proposed and development of prototype systems is underway in several laboratories. The solutions include fiber, other guided wave media such as polymer waveguides, and free-space interconnects. Consider first fiber. While the core diameter of 50 or 62.5um offers compact size, the outer cladding of 125um fixes a minimum pitch of 8/mm, and the addition of the plastic buffers required to prevent long-term degradation of the fibers might increase the diameter of each line to 250um. A de facto standard for parallel fiber based interconnects is 250um pitch. In a single layer of fiber, this represents a maximum interconnect density of 4/mm, compared to a usable density in a multi-row (perhaps 5 rows) electrical connector. If multiple layers of fiber were used, the density would increase in proportion. In fact, not all the board edge can be associated solely with fiber, as most fiber based board-to-backplane connectors employ some form of insert to locate the fibers within a standard connector shell. This insert occupies space that is then unavailable for interconnects. For example, an MT connector of 5mm width supports 12 fibers in a single layer. This in a single layer gives a usable density which is comparable to the electrical case. By stacking several such inserts in a single connector, densities of 2-4x that of the electrical case can be obtained.

Fiber solutions offered little advantage over electrical technology until point-to-point architectures gained acceptance, since the bulk of fan-in and fan-out circuits prevented their use in electronic systems with tight physical constraints. Even for point-to-point architectures, the minimum bend radius of typical fibers results in large areas on each board being taken up with the routing to and from the connector housing. Furthermore, the assembly time required for the backplane scales with the total number of interconnect paths. These problems can be solved in two ways. The poor

scalability of the fabrication cost with interconnect count can be addressed by embedding the optical fiber with a flexible carrier to form an interconnect layer which can then be attached. Shahid has developed such a technology based on telecommunications fiber. The minimum bend radius however can only be addressed by fabricating fibers with smaller cladding diameter since the minimum bend radius is dictated by the maximum tolerable stress in the surface of the fiber. High humidity contributes further to the degradation, so that deployment in military environments might require relaxing of the bend radius. It is however these very same military processors which are likely to represent the first insertions of optical interconnects.

Embedded fiber approaches offer an assembly cost which is independent of the number of interconnects, however the cost of fabrication of that layer depends on the number of interconnects being implemented. The minimum feature size is defined by the fiber drawing process. Polymer optical backplanes represent an advance towards the techniques employed to reduce cost in the electronics industry where printed wiring boards have long displaced discrete wiring schemes. Several polymer waveguide interconnect media have been developed, and several attempts have been made to integrate the polymer waveguide media into complete optical backplanes. Polymer waveguides offer interconnect densities limited only by photolithography processes. As with electrical interconnect media however, the usable density in an electronic system depends on the density at the various interfaces between the different levels of packaging in a system.

In eliminating the bottlenecks in electronic systems, it is important that additional cost not be introduced, or at least that any additional cost be commensurate with the benefit provided. The MCM-to-board connector is integrated into the system using a passive alignment scheme. During fabrication of the waveguide layers, alignment features are defined and by virtue of the manufacturing process cannot help but be aligned to the waveguides. The features take the form of troughs immediately adjacent to the waveguide arrays. To assemble two components in the system with waveguide arrays defined, a rigid key is used. This key includes features which mate with the troughs defined on the MCM-level interconnect later and also with those on the printed wiring board level interconnect layer. By pressing the key piece into the troughs, the flexible element is brought into alignment with respect to the fixed waveguides. This clearly eliminates the need for active alignment, and furthermore offers an assembly time which is independent of the number of waveguides being connected or the dimensions of the waveguide.

3 PROGRAM GOALS

The OIT program had several specific goals relating to the development of optical interconnects for use in electronic systems. The first goal was to eliminate the I/O bottleneck in electronic systems by demonstrating MCM-to-MCM interconnects using polymer waveguides. The principal problem to be solved in eliminating these bottlenecks is to provide an interconnect density across the various interfaces in a system which matches the interconnect density of the media associated with individual components. For example, the board-level waveguide interconnect density must be maintained across the board-to-backplane interface to prevent the I/O density at that point becoming an impediment to system performance. In addition to the board-to-backplane I/O density, the density must be maintained across the MCM-to-board interface.

Polymer optical waveguide based interconnects allow the interconnect density to be preserved from MCM through the module-to-board interfaces, and on board.

4 APPROACH TAKEN FOR OIT

The overall approach taken for OIT was to add optics to electronic systems by exploiting the existing infrastructure, rather than to replace all elements of an electronic system. Specifically for the interconnect medium, the approach adopted was to provide a passive multimode polymer waveguide interconnect which could be implemented using standard printed wiring board techniques, and which could also be incorporated in a standard multichip module fabrication scheme. In this respect, the resulting system would be an intermediate solution for electronic system integrators, retaining compatibility with existing system assembly practices and yet providing at least an order of magnitude improvement in system throughput. Longer term optical solutions might gain yet more benefit from optics by eliminating many of the currently accepted design philosophies, but such systems would be unlikely to find acceptance in the timeframe addressed by this program.

5 INTERCONNECT MODELING AND DESIGN

For this program, we devised a high level model with which to characterize a complete optical interconnect from electrical input to electrical output. Parameters calculated include link risetimes, signal to noise characteristics, optical power budget, and total power dissipation for the optical interconnect. For the structures developed under OIT, we have used data from specific device models, enabling component specifications to be derived, and the effect on system performance of a particular set of component characteristics to be evaluated. For components not developed under this program, such as the receiver and the laser and its associated driver, existing data is used.

The purpose of modeling the optical components was to gain an understanding of the modal nature of energy propagation in highly multi-mode waveguides, and to use this information to predict device behavior. In order to optimize the performance of the link as a whole, we must be able to individually characterize each constituent component. For the MCM-to-MCM interconnect, this required the development of several distinct modules.

In modeling a complete interconnect, we must be able to predict the coupling between a VCSEL and waveguide via a 45 degree turning mirror, and determine the associated alignment tolerances. A similar analysis is required at the detector. The flexible optical jumper selected to connect between the MCM and the board will require bends in the waveguides which represent a mechanism for mode selective loss, so we must be able to evaluate the potential impact on modal noise.

The coupling between the flex jumper and the board or MCM level waveguides also may be mode selective, so this component and its interfaces must be examined for its potential to introduce modal noise. As these interfaces represent the assembly points of the MCM-to-MCM link, the necessary tolerances for obtaining passive alignment must be designed into the interconnect a priori with the aid of an appropriate model.

Lastly, we must be able to account for ancillary effects such as crosstalk between adjacent waveguide channels and scattering losses due to substrate roughness. These models and those listed above have all been developed under this program.

Table 1 shows the specific approaches adopted for each of the elements in the optical interconnect. Waveguide models using semi-vectorial finite difference analysis, matrix methods, and beam propagation methods have been established at Honeywell, in addition to routines for evaluating coupling efficiencies and modal noise effects. Our suite of routines can accommodate rectangular geometry of waveguides cross section, or irregular cross-sections. We are thus able to design and simulate waveguides made by reactive ion etching, laser ablation, and direct exposure.

Table 1. The various approaches used in modeling the elements of the system

<i>Design parameters</i>	<i>Modeling Method</i>
Waveguide Topology	Effective Index Method (EIM) 2D Matrix Method (MM) Semi-Vectorial Finite-Difference Method (SVFD)
VCSEL/Detector-to-waveguide	

coupling efficiency Modal Noise	CODE V SVFD/Overlap integral
Waveguide-to-waveguide coupling efficiency Modal Noise	Beam Propagation Method (BPM) SVFD MM/EIM SCFD/Overlap Integral BPM
Waveguide Bends Radiation Loss Modal Noise	Overlap Propagation Method (OPM) BPM OPM/MM/EIM/BPM
Crosstalk	BPM Approximate Analytical Model
Scattering Loss	Approximate Analytical Model

For the multimode waveguides chosen for this program, each element of the model determines the total number of modes propagating, and calculates the effect of that particular element on each mode. Specific parameters calculated include loss and propagation delay, together with the optical field distribution at the input, and if likely to be different, at the output. Having determined the field distributions of the modes and their relative phases, coupling efficiency evaluation routines are used to determine the power transfer between the propagating modes of different sections, thus enabling determination of the performance of an entire link from a knowledge of the performance of individual elements.

Components/Interfaces Contributing to Modal Noise




Figure 1 The complete interconnect modeled for OIT.

Figure 1 illustrates the complete interconnect structure modeled, including the elements for which existing data is used. For each subsystem, the major characteristics evaluated are shown.

5.1 WAVEGUIDE STRUCTURE

The basic waveguide structure consists of a rectangular Ultem core encapsulated in the BCB cladding polymer. The initial objective was to fabricate guides with core dimensions of 25 x 25 microns and 50 x 50 microns.




Figure 2 The basic waveguide structure used for OIT

Modeling of the electric field mode profiles and propagation constants was accomplished using the semi-vectorial finite-difference technique. In the method, the waveguide cross section is discretized into many points on a non-uniform grid. Over that grid, the Helmholtz wave equation

$$\nabla^2 \mathbf{E} - k^2 \mathbf{E} = 0 \quad (1)$$

is approximated in scalar form using the matrix defined by the second order differencing of the partial derivatives. The imposition of boundary conditions at the interfaces between media of differing indices yields solutions which are valid for quasi-TE and quasi-TM polarization states. The result is a simple eigenvalue problem

$$\nabla^2 \mathbf{E} + k_0^2 \mathbf{E} = \beta^2 \mathbf{E} \quad (2)$$

whose eigenvectors $\mathbf{E}$ and eigenvalues β are the desired field profiles and propagation constants of the guided modes. These are solved with the aid of UNIX based implementations of the popular and commercially available NAG and IMSL numerical routine libraries.

The main problem encountered in using the SVFD method in the modeling of multi-mode guides is the limitations placed on the computation by the exceedingly large amount of fast storage required. The amount of memory necessary for one calculation is proportional to $(NM)^2$, where N and M are the number of node points along the two grid axes. In a typical calculation, one would like to place at least 10 node points per wavelength in each linear dimension. For very large, i.e., 50 x 50 micron structures, this is simply not possible without the aid of a supercomputer. In some

cases it was therefore necessary to perform the calculations on a smaller domain and extrapolate the results out to the full sized structure.

As an example, several intermediate order modes for a 25 x 25 micron Ultem/BCB waveguide are shown in Fig. 3(a-f). These plots represent the transverse distribution of electric field intensity. Because of the 2 degrees of symmetry in the problem, the solution to (2) consists of degenerate pairs of eigenvalues. Therefore, the eigenvectors produced by the solver routines are linear combinations of the true eigenvectors. This was corrected with the addition of a slight perturbation to the guide geometry to break the degeneracy.




Figure 3 Selected modes of an Ultem/BCB waveguide system

In a twofold symmetric problem, there is also degeneracy between the TE and TM polarization states of the guided modes. When the symmetry is broken, the polarization split into distinct modes with slightly different propagation constants. Strictly speaking, structures such as that shown in Fig. 3 are not capable of supporting purely TE or TM polarization, but rather hybrid modes comprised of a combination of the two. The SVFD method solves for the individual modes by assuming they are principally polarized along a given axis, hence the designations quasi-TE and quasi-TM.

For the case of the buried channel Ultem/BCB waveguide, even though the quasi-TE and quasi-TM modes exhibit slightly different propagation constants, the corresponding E-field profiles differ by very little. Since it is the field profiles and not the propagation constants we are interested in, we chose to neglect polarization effects and use the quasi-TE results for all calculations.


5.2 VCSEL-TO-WAVEGUIDE COUPLING

To develop a model which would enable the effects of imperfect mirror fabrication and alignment to be predicted, we undertook the modeling of coupling between a VCSEL and waveguide with a 45-degree end facet. The objective of this modeling was to estimate the VCSEL to waveguide coupling efficiency at various separation distances. The basic approach was to use using ray-tracing under Code-V, a widely used and commercially available tool. The optical medium was defined to be similar to that shown in Figure 2. A total of 10,000 rays were used, these originating from the VCSEL in a distribution that simulates the observed VCSEL beam shape (both width and angle). The total coupling loss is calculated by counting only those rays still guided inside the waveguide core after propagating through an extended section of waveguide.

The results of this simulation are presented in Table 2. For a separation distance $Z=150$ microns, the data shows the best achievable coupling efficiency to be -3.64dB. This corresponds to the case in which Kapton-backed waveguide samples have been terminated such that the coupled light must first pass through the substrate before encountering the mirror. This level of loss is excessive, and does not account for any imperfections in the mirror surface arising from the sawing process.

Table 2. Simulated coupling efficiency versus separation between the VCSEL and the waveguide lower surface

Z (μm)	Loss (dB)
50	-1.01
100	-2.53
150	-3.64



The data shows that for $Z < 50$ microns, better than -1 dB coupling efficiency can be expected. This situation can be achieved by terminating the flex arrays "upside down", so that the upper BCB cladding can now be placed in direct contact with the VCSEL. This result is also applicable to HDI processed modules with up to two layers of interconnects.

5.3 INTERCHANNEL CROSSTALK

For parallel interconnects, the polymer optical waveguides must consist of at least two parallel optical channels. As the interaction length between the waveguides is necessarily long, there exists the possibility of significant crosstalk occurring. Two mechanisms can give rise to energy exchange between two waveguides in close proximity to one another.

First, if the waveguides are identical, the phase-matched condition of the propagation constants allows energy to be coupled from one guide to the other by interaction of the evanescent tails of the modes, similar to the operation of a directional coupler. Evanescent coupling between two parallel guides is minimized by assuring that they are placed on a pitch sufficient to assure negligible overlap of the guided mode tails. For the Ultem/BCB buried channel waveguides, this was done using an effective index approximation for determination of the guided modes of 2-dimensional waveguides based upon Ghatak's matrix method [Theodor Tamir, Ed. "Guided wave optoelectronics", 2nd Ed., pp. 35-50, 1990, Spring Valley.] for 1-dimensional dielectric stacks. It was determined that a pitch of 50 microns was sufficient to assure negligible crosstalk arising from evanescent mode coupling.

The other form of crosstalk arises from radiation mode coupling induced by the presence of surface irregularities on the interface between the core and cladding materials. Some degree of surface scattering is unavoidable due to the etching of the waveguide cores. In our process, the RIE etching of Ultem results in a residual RMS surface roughness. The same scattering mechanism which leads to loss via radiation mode coupling in one waveguide can lead to radiation-to-guided mode coupling in an adjacent guide, thus giving rise to crosstalk.

A simplified model to predict crosstalk levels arising from surface scattering in parallel slab waveguides was devised based upon the assumptions that all observed losses were due to such scattering, and that all scattered light was subsequently coupled into the adjacent waveguide. The resulting model consists of the two coupled-wave equations

$$dP_1/dz = -\alpha P_1 + \alpha/2 P_2 \quad (3)$$

$$dP_2/dz = -\alpha P_2 + \alpha/2 P_1 \quad (4)$$

where P_1 and P_2 are the optical power propagating in the two slabs and α is the scattering loss coefficient. The factor of $1/2$ arises from the fact that only that energy scattered in the direction of the other slab can be coupled. Since the slab waveguides are infinite in lateral extent, the crosstalk levels predicted by this model are independent of spacing.

Equations (3) and (4) were solved simultaneously subject to the boundary conditions $P_1(0)=1$ and $P_2(0)=0$, and the quantity $10\log(P_2(z)/P_1(z))$ can be calculated as a function of α in dB/cm. The slab waveguide model is the limiting case of two infinitely wide channels, and so represents the "worst case" crosstalk level which can be expected. Our lowest channel waveguide loss obtained using the Ultem/BCB system has been 0.24 dB/cm. Some of this loss is likely to be due to absorption, but even if all of this loss were due to scattering, the model predicts a crosstalk level of less than -60 dB for a 30 cm interconnect.

We have performed measurements of crosstalk levels in parallel waveguide channels to provide experimental verification of the theoretical crosstalk predictions. The measurements were performed by launching signal light into one waveguide, and measuring the optical power emitted from the guides on either side using an aperture to block the primary channel.

We evaluated a number of samples in which etching of the Ultem waveguide core was incomplete in areas, and measured substantial crosstalk as expected. This illustrates the importance of a process with minimal undercut during etching. Since complete etching of the waveguide layer between cores is required, some over-etching is necessary in order to ensure that the thickest regions of the waveguide layer are fully etched. A process with appreciable undercut would result in regions which were initially thinner also becoming narrower.

Measurements performed on a fully etched 10cm sample containing 25 μm guides on 25 μm spacings revealed optical crosstalk levels of at most -31 dB. We believe that the actual levels were significantly lower than this, but the figure represents the dynamic range of our experiment. At least part of the measured crosstalk was due to the limited resolution of the output coupling optics and scattering after the waveguide output facet. For a digital system, a figure of -20 dB electrical crosstalk is typically acceptable. We have thus confirmed experimentally that crosstalk in the waveguides will not limit system performance. In practice, crosstalk in receiver arrays is likely to be the dominant crosstalk mechanism, and is likely to far exceed any effects due to the waveguides.

5.4 RADIATION LOSSES DUE TO WAVEGUIDE CURVATURE

Radiation losses from waveguides can be caused by in-plane bends required for signal routings or from bends caused by flexing the carrier or substrate upon which the guides have been fabricated. It is important to characterize this source of loss, because in large part it determines the necessary area or volume required for a particular network topology.

The origin of radiation loss from waveguide bends can best be visualized by considering the case of a non-uniform plane wave with a gaussian distribution propagating along a curve, with phase fronts parallel to the radial direction. It is clear that the outward tail of the wave must travel faster than that closer to the bend origin if the phase front is to be preserved. For sufficiently small bend radii, the speed of light in the media may be exceeded. Since this is clearly not possible, that portion of the energy contained in the tail becomes unguided and is radiated away.

No practical models exist for the prediction of radiation losses due to large (i.e. 90 degree) bends in highly multi-mode waveguides. Beam propagation codes are good when the radii of curvature is small, but fail to give accurate results when the paraxial assumption implicit in these methods is violated.

Gu, et al., has presented a method suitable for few mode guides based upon the solution of the scalar Helmholtz in cylindrical coordinates ("Novel Method for Analysis of Curved Optical Rib Waveguides", Electronic Letters, Vol 25, No 4, 16th February 1989, p278). Using this technique, the extinction coefficient of each propagating mode can be determined. However, the extension of this approach to

large core guides was not feasible due to the unavailability of sufficient computational resources. Even if the model were available, it would be of limited utility because the evaluation of overall bend loss requires an accurate knowledge of the propagating mode spectra.

Without benefit of a model, we were forced to rely upon experimental measurements and heuristically derived formulae for the prediction of waveguide curvature loss. To carry out these measurements, we fabricated Ultem/BCB 15 x 50 micron core waveguides using a specially designed mask (Fig. 4) containing a series of 90 degree bends ranging from 0.5 to 10 mm in radius. Each waveguide on the mask consists of four identical bends connected by five straight segments.




Figure 4 Waveguide mask layout for evolution of waveguide bend losses

Throughput measurements on this sample were used to determine the excess curvature loss per 90 degree in-plane bend as a function of radius. The data collected from the analysis is presented in Fig. 5. Despite a fairly high degree of scatter in the data attributable to variations in input coupling efficiency, it can be seen that the minimum bend radius for 1dB of curvature loss is approximately 3mm, while 3dB loss is obtainable for radii as small as 1mm. These numbers are consistent with high density board-level routing.

Experiments were also performed to determine the excess loss induced by flexing the waveguide substrate. Samples were fixtured with their ends held secure while wrapped partially around various sized mandrels. It was observed that bend radii as small as 1 cm introduced a negligible degree of excess loss. This is significant because a 1cm bend radius is consistent with the

dimensional requirements of our interconnect scheme. Smaller bend radii tended to manifest higher loss due to microcracking of the BCB cladding.




Figure 5 Measured bend losses for in-plane waveguide bends

5.5 MODAL NOISE PHENOMENA AND THEIR IMPACT ON OIT

An important concern in the design of multi-mode optical links is the potential degradation of signal-to-noise ratio due to modal noise. This source of noise is observed as fluctuations in throughput due to the changing modal spectrum of energy propagating in a waveguide containing an interface which does not uniformly transmit all modes. Unlike other forms, modal noise is directly proportional to received power, so signal-to-noise ratios cannot be improved with an increase in power budget. If unaccounted for, modal noise can be so dominant that acceptable error rates are impossible to achieve.

There are three factors that must be present simultaneously in order for modal noise to be an issue. First, the link must have some mechanism for mode-selective loss, where modes are attenuated to varying degrees depending upon the structure of their transverse field profiles. This can be an offset connector, fabrication defect, a bending structure, longitudinal non-uniformity, or almost any form of discontinuity transverse to the cross-sectional plane of the waveguide.

Second, the optical source must possess a coherence length longer than the distance to the first mode selective element. If this is the case, dephasing of the modes incident on the interface will not occur, resulting in a spatially non-uniform interference pattern. In the case of an offset connector, the origin of the modal noise can easily be visualized. If under a particular excitation condition the interference pattern has a significant portion of the total energy outside the aperture of the coupled guide; transmission will be low. A slight change in the modal spectra of the incident guide may then yield an interference pattern with a higher proportion of its energy within the aperture, leading to greater transmission with no change in incident power.

Lastly, there must be a temporal dependence of the mode distribution in order for mode selective transmission to manifest itself as "noise". This can be due to environmental changes, vibration, and bending, but these are generally at low frequencies which can easily be filtered out. More significant are the fluctuations due to the characteristics of the source itself. When directly modulated, semiconductor lasers exhibit a shift in wavelength known as "chirp". Also, there may be present fluctuations in the far field pattern of the emitted radiation due to mode-hopping or filamentation. These mechanisms give rise to changes in the modal excitation at the data rate, a much more serious problem.

The effects of modal noise are minimized by : 1) Using a highly multi-mode waveguide geometry and exciting as many modes as possible. Note that it is the number of modes actually present, not the maximum number supported, which dictates the magnitude of the modal noise. 2) Minimizing the coherence length of the source. Incoherent LED's are far superior to semiconductor lasers in this respect. This solution is not of interest when the particular source to be used is specified and not subject to change. 3) Minimizing the loss of all components in the system and in particular minimizing the mode dependence of any loss mechanisms, whether distributed or localized.

The optical sources specified for use in this program are vertical cavity surface emitting lasers with coherence lengths roughly on the order of the overall link dimension. Nevertheless, it will be shown that the effects of modal noise can be made inconsequential, even in the presence of a high degree of mode selective loss, through the use of large core waveguides supporting a very high volume of guided modes. This requires that the guide cross-section and/or the difference in refractive indices between the core and cladding materials be large, consistent with our selection of Ultem and BCB as our baseline polymers. A large relative index difference has the added advantage of reducing radiation losses from small bend radii. This is important so that individual components can be made compact, reducing the overall volume required by the complete interconnect.

5.6 MODELING AND CHARACTERIZATION OF MODAL NOISE

In an experiment designed to verify the low levels of modal noise predicted by the simulations, a mock connector was set up as shown in Fig. 6 in order to measure the obtainable BER as a function of connector loss due to misalignment. In the experiment, two identical 15 x 50 micron core Ultem/BCB waveguides were butt-coupled together with the aid of xyz translation stages. Light from an 830 nm semiconductor laser diode (not a VCSEL) driven directly by the data output port of the Anritsu error-rate tester was coupled into and out of the guides using microscope objectives, and detected with a

silicon APD. The output of the bit error rate tester included a variable dc offset, allowing biasing of the laser at threshold to minimize pulse width distortion and rise/fall time. The signal from the detector was amplified using an HP broadband amplifier, then supplied to the signal port of the BERT. The experiment was limited to a data rate of 500 Mbps by the amplifier.




Figure 6 Experimental arrangement used to investigate modal noise

Due to the fact that modal noise is independent of power level, an examination of bit-error-rate as a function of received optical power in the presence of modal noise should reveal a BER floor, representing the minimum achievable error rate regardless of the signal level available at the receiver. By translating the two halves of the “connector” in the setup shown above, we were able to introduce varying degrees of mode selective loss. Also, by deliberately misaligning the waveguides with respect to the focussed input spot of the laser we were able to simulate the effects of different launch conditions. While it is difficult to excite a different number of modes, it is relatively easy to selectively excite groups of high order modes or low order modes. What we observed was that even for very high levels of connector loss we were unable to detect any such floor within the limits of our experiment. This provides evidence that significant levels of modal noise were not present.

We have also investigated the effects of modal noise on passive multimode polymer waveguides. We have modeled the effect of propagation through a series of waveguides fabricated using a mask defined by GE’s Adaptive Lithography process. These adaptive bends provided a total of six waveguide offsets grouped into three pairs. We have earlier published a discussion of the suitability of various models used to predict the effects of modal noise. In this case, our calculations predict a signal to noise ratio for a 30 um offset occurring in 300 um adaptive length of 124 for full modal excitation, or 2630 for excitation with a Gaussian input. Assuming the noise associated with each bend to be uncorrelated, we would expect the signal to noise ratio for six adaptive bends with these dimensions to be 58 or 1073 respectively. In neither case does the model predict that the signal to noise ratio from these components would prevent bit error rates of 10^{-13} or better being obtained. In reality, the input is approximately Gaussian, and will excite a small number of modes centered around a mode order which varies according to the excitation conditions, but two competing effects will tend to increase the number of modes excited with significant amplitude at any given point. Sidewall scattering will tend to convert

energy from low-order modes to high order modes, while the waveguide loss is mode dependent. Depending on the magnitude of the two effects, the waveguide may carry a large number of modes. Figure 7 shows the output from a test waveguide.

Bright lobes of low spatial periodicity indicate that the energy is confined mainly in low order modes, while a background of higher order modes indicates either a small amount of scattering, or a large amount of mode selective loss together with a large amount of scattering. The low modal noise and reasonable baseline loss of our waveguide structure tends to support the former explanation.




Figure 7 A typical waveguide output showing predominantly low order excitation of multimode polymer waveguides

In order to evaluate the potential degradation in bit-error-rate due to the HDI written adaptive waveguide bends necessary for chip-last packaging, an experiment was carried out using waveguides fabricated from a quartz/chrome mask which had been written using the HDI direct-write process. On the mask were a series of 25 and 50 micron wide waveguides, each offset by three identical pairs of adaptive bends. Each bend simulates an interconnect as it would actually be written on a multi-chip module. The total offset was held constant at 50 microns while the severity of the bends was increased by reducing the length of the adaptive region from 500 to 50 microns in 25 micron increments. Moreover, the nominally straight sections of the waveguides incorporated stitching errors, which might be expected to contribute additional modal noise.

The loss on all but the most severe bend (50 μm offset of a 25 μm wide guide in 50 μm adaptive length) allowed sufficient power to reach the detector that receiver noise would permit a bit error rate better than 10^{-12} in the absence of modal noise. Waveguides containing adaptive features defined by 50 micron offsets in 50 micron adaptive lengths exhibited a per bend loss of 1.6dB, for a total mode selective loss of 9.6dB for the link. This loss is similar to that we expect from a complete MCM-to-MCM interconnect. Even for this severe case it was not possible to measure a bit error rate in excess of our practical limit of 10^{-12} . (At 500 Mbps, a single error would on average be detected every 33 minutes, while a reasonable degree of confidence in the measured result requires perhaps 10 detected errors. Thus BER measurements at small bit error rates are time consuming).

6 POLYMER WAVEGUIDE DEVELOPMENT

The interconnect medium used for OIT is a passive polymer multimode waveguide. This offers compatibility with printed wiring board fabrication in particular and with the existing infrastructure for electronic system assembly in general.

6.1 POLYMER CHOICES FOR OPTICAL INTERCONNECTS

Optical interconnects for insertion into current MCM or backplane systems must exhibit the following characteristics. They must offer performance better than that of the existing electronic technology, or offer a leap forward in functionality. They must be cheaper, or at least show the potential of becoming so in the near term, and must be free of exotic processing methods which are beyond the capabilities of existing foundries. And, they must be compatible with established processes and materials if widespread acceptance is to be found.

Few would claim that optical interconnects do not hold the potential for vast performance improvements over their state-of-the-art electrical counterparts. Indeed, fundamental limitations in throughput and power dissipation are already being felt which may only be resolved through a transition to optical components. It is the goal of projects such as this to demonstrate strategies which meet the demands for performance, manufacturability, and compatibility. But, the question remains, "Can it be done cheaply enough to justify the insertion of optical interconnects in the near term?" One answer to this question lies in the use of easily processed, low cost polymer-based materials as the optical media. Polymers have been used in MCM processes for some time for use as inter-metal dielectric and passivation layers, so compatibility is not an issue. Many vendors are also currently developing low-loss materials in a range of refractive indices specifically for optical applications, including a few of which are photo-definable using standard lithographic techniques.

6.2 POLYMER IDENTIFICATION

A considerable amount of effort under this program was directed towards identifying appropriate materials for use in passive polymeric interconnects. Several key factors were important in selecting candidates for evaluation. First, any material chosen was to be commercially available or available in developmental form but shortly to be released as a commercial product. This was to insure a low cost technology attractive from the point of view of its potential for near-term insertion. Second, the product was to be compatible with established low loss waveguide fabrication techniques. Lastly, the materials had to be compatible with current board integration materials and processes.

The fundamental component in this program which facilitates the interconnection of multiple optoelectronic die via optical interconnects is the buried channel waveguide(Fig.2). In this type of structure, light is confined by total internal reflection to a high index core region bounded on all sides by a cladding of lower index. Use of a buried, as opposed to an air-bounded or ridge-loaded

structure, leads to a reduction in scattering losses associated with surface imperfections, allows for sufficient planarization to isolate optical layers from the topographical structure of under-lying or over-lying layers, and prevents the deleterious effects of environmental exposure in addition to providing a robust interconnect capable of withstanding handling in standard assembly schemes.

Two polymers are necessary for the fabrication of such a structure, one for the waveguide core and the other for the necessary cladding. For reasons to be outlined later, the indices of refraction of the two materials should differ by as wide a margin as possible. In addition to exhibiting acceptably low intrinsic optical absorption at the wavelength of interest (here 830 nm), they must show sufficient mechanical and thermal stability as well as resistance to environmental contamination.

An extensive list of polymer materials were evaluated with respect to the above criteria, including materials from GE, DOW, OCG, Hitachi, Allied Signal, Ciba Geigy, Dupont, PhotoNeece, UpJohn, and Amoco. Small quantities of developmental and commercialized products were readily made available by the above manufacturers for our evaluation. All of the materials examined are either currently used within the HDI process at GE, or are suitable for direct integration.

As a preliminary step, each material was characterized for indices of refraction and for optical loss due to intrinsic material absorption in order to determine potential suitability. Films of each material, 4 to 7 microns thick, were deposited by spinning on polished silicon wafers onto which had been grown a thick layer of thermal oxide to serve as a buffer layer to isolate the film from the highly absorbing substrate. The films were then thermally cured per manufacturers recommendations. Mode dependent loss measurements were obtained via the sliding prism coupling technique using an 830-nm semiconductor laser diode as the light source. This is a standard technique for evaluating the quality of planar film (slab) waveguides. Since attenuation in 2-dimensionally confined waveguides is typically greater than that of their 1-D slab counterparts (due principally to sidewall scattering associated with fabrication processes) the loss data for the films was taken to be lower bound for that obtainable in a 2-D structure.

The refractive indices of the films were taken to be the measured effective indices of the lowest order propagating TE modes. This is a valid assumption because the films in all cases were optically "thick", supporting a large number of guided modes. Under this condition, the fundamental mode is very strongly guided, and its effective index asymptotically approaches that of the film. This is the case for both TE and TM polarization.

The experimentally derived data for selected samples is shown graphically in Fig. 8. In the chart, both the index of refraction and the material absorption coefficient in dB/cm are given for each material. Only those polymers of potential significance are shown. The horizontal line represents the maximum loss permissible for initial systems applications of 0.3dB/cm. This would enable systems to be implemented in which the backplane is 12" long or so. In fact, the desire to reduce receiver power dissipation motivates the development of optical interconnect systems in which each receiver receives a similar signal level, since demands on automatic power control circuits are reduced. Ideally then in the future, the interconnect medium will have losses less than 0.1

dB/cm with 0.05 dB/cm or lower being even more attractive. Since we required a loss of no more than 0.3dB/cm for initial demonstration, this was taken as the goal for OIT.




Figure 8 Experimentally derived optical characteristics for various polymer waveguide interconnect media

6.3 POLYMER SELECTION

Based on our experimental evaluation, Ultem was selected as the baseline polymer for fabrication of the waveguide cores. Ultem 1010LC is a liquid version of a polyetherimide thermoplastic resin produced by GE which has long been used in the injection molding of high quality plastic parts and optical components. Ultem offers relatively low loss and a high index of refraction. The lowest propagation loss measured in this material was approximately 0.36dB/cm, comparable to our design goals. The only material with significantly lower absorption was the Allied Signal acrylate-based material, but at the time of selection, development of this product had not yet advanced to the point where a commercial release can be assured.

Ultem films as thick as 15 microns can be spin coated using standard equipment, and ramp curing is affected in a vacuum oven at temperatures in the range of 220 to 280 degrees C for periods up to 1 hour.

Cyclotene (benzocyclobutane or BCB) was chosen as the baseline cladding polymer due to its low exhibited loss of 0.28dB/cm and the fact that its refractive index of 1.56 provides the desired differential between that of Ultem. What is it principally used for? How is it cured? There are photoimable versions of BCB available, but these exhibit significantly higher loss, as do all of the other photosensitive polymers examined.

As part of the process development phase of the waveguide development task, extensive experimentation was undertaken with the goal of improving the optical qualities of Ultem and BCB. Deposition techniques and cure time/temperature cycles were optimized to yield the lowest possible losses, and the results obtained represent the state of the art for currently available materials. It goes without saying that further improvements will significantly broaden the impact of polymer waveguides on future product developments.

6.4 CHANNEL WAVEGUIDE PROPAGATION LOSS MEASUREMENT

An experimental test station was assembled on a breadboard table with which to evaluate the multimode polymer waveguides. Light from a collimated 40mW single transverse mode 830nm GaAs/AlGaAs semiconductor laser diode was focused on the endface of a waveguide sample using a low numerical aperture lens. The sample was mounted on an xyz positioning stage equipped with piezo-actuators to facilitate accurate and reproducible alignment. At the output, the transmitted light was collected by a high numerical aperture lens and focused down through a variable aperture onto a CCD or silicon PIN diode detector connected to a calibrated optical power meter.

Using this setup, sample throughput and nearfield measurements were obtained in straightforward manner. It was also possible to observe and quantify variations in modal behavior under changing source coupling conditions.

The propagation losses of the various waveguide samples were evaluated using the cutback technique. For each sample, measurements were made of the optical throughput for as many lengths as were available, while maintaining the launch conditions as constant as possible. When plotted on a semi-log scale, the slope of the resultant line directly gives the attenuation constant for the waveguide in dB/cm. The input coupling efficiency can be deduced from the intercept of the linear data fit with the abscissa. The degree to which the data departs from linear suggests the degree of confidence one should place in a particular measurement.

For the flex samples, typical measured attenuation constants for the 10, 25, and 50 micron wide waveguides were 2.78, 1.0, and 0.35 dB/cm respectively. The higher loss of the narrow channels is to be expected due to the increased influence of sidewall scattering. The 0.35dB/cm attenuation in the 50 micron wide guides was an excellent result for RIE etched polymer waveguides, and came very close to meeting our design objective of 0.3dB/cm. In addition, the high coupling efficiency observed testified to the quality of the waveguide endfaces prepared via the sawing process.

Throughput measurements taken of the board laminated waveguides showed no deterioration when compared to data taken before undergoing the lamination process. In some cases, a slight increase in throughput was actually observed, large enough not to be attributable to the minor reduction in sample length required for re-termination of the waveguide endfaces. The reasons for this improvement are undetermined, but could possibly be due to an improvement in the facet quality achieved in the sawing operation.

Being rigid, the laminated samples are less subject to chatter induced by the friction of the diamond blade. This would undoubtedly yield cuts with smoother surface figures resulting in lower optical scatter. In any case, the concern that conformal adherence to the underlying weave pattern of the board material would lead to excessively high loss proved to be without merit.

6.5 FLEX WAVEGUIDE FABRICATION

An attractive feature of this waveguide fabrication technology is that it relies upon readily available commercialized materials and industry standard processing techniques. Required operations involve spin-coating of resists and polymers, controlled baking, e-beam metallizations, photolithographic pattern transfers, and reactive-ion etching. This makes it possible to insert optics into a process flow with a minimal capital investment.

Kapton film was chosen as a fabrication platform for the polymer waveguides and to provide increased mechanical stability. Kapton is a low cost polyimide material marketed by DOW Chemical for a multitude of uses, including electrical flex wiring and MCM packaging. Kapton exhibits excellent mechanical and thermal characteristics, particularly in high temperature environments. Available thickness range from 2 to 5 mils in sheets or rolled strips.

Prior to any processing, 5 mil Kapton sheets were initially trimmed to a workable size approximately 5 inches square, consistent with the 4" contact masks used for the subsequent photolithography exposures. The sheets were first mounted to rigid glass carriers to facilitate handling. This also avoided the tendency of unmounted Kapton to curl later in processing from the stresses caused by the addition of polymer coatings.

The remainder of the process flow is as depicted in Figure 9. First, a 15 micron thick BCB lower cladding was deposited on the Kapton substrates by spinning, followed by a bake cycle to cure the film. Next, a nominally 20 micron thick ULTEM waveguide core layer was deposited on top of the BCB in the same fashion. After curing in a vacuum oven, a layer of SiO₂ was RF sputtered on top of the ULTEM to serve as a hard mask for the definition of the waveguide cores by dry etching.

Positive photoresist was spun onto the samples and exposed through the light field waveguide mask. This pattern was then transferred into the underlying SiO₂ by reactive ion etching. Using the SiO₂ as a mask, waveguide cores were then etched into the ULTEM using oxygen RIE, etching all the way down to the lower BCB cladding. The RIE etch process is highly anisotropic,

so the waveguides maintained a rectangular cross section. This etch process has been optimized to yield the low degree of sidewall roughness necessary for low loss waveguides.





Figure 9 Process sequence for fabrication of the flexible polymer waveguide interconnect

After etching the waveguide cores, a final overcoat of BCB was applied. This layer serves as the upper and lateral waveguide cladding, and is also effective in planarizing the surface. After undergoing a final bake cycle to fully cure the BCB cladding, a layer of titanium was e-beam evaporated over the entire surface to encapsulate the waveguides and protect them from environmental contamination.

After metallization, the flexible Kapton substrates were demounted from the rigid carriers. Figure 10 is a cross-section microphotograph (left), and a photograph of optical waveguides fabricated on a flexible Kapton carrier.



Board integrated waveguide cross-section



Flexible waveguide ribbon

Figure 10 A cross-section microphotograph (left), and a photograph of optical waveguides fabricated on a flexible Kapton carrier

6.6 BOARD LEVEL WAVEGUIDES

An important part of this program is the ability to integrate polymer waveguides with board level electronics for inter-MCM signal routing and card-edge backplane interfaces.

Results of our survey of commonly employed board materials indicate that the greatest payoff will result from demonstrating our polymeric interconnects on the glass:polyimide board system. Glass:epoxy finds widespread use, but is subject to degradation of mechanical properties at temperatures of 125C- a target for many military systems. Glass:polyimide offers comparable or better electrical performance while maintaining mechanical integrity in high temperature environments. Indeed, operating temperatures up to 250C are sustainable. Cost of this system is approximately double that of glass:epoxy. Limitations of the glass:polyimide system are primarily in the maximum data rate attainable. A maximum figure of 1 GHz over a 1 foot length appear feasible. Beyond these data rates and lengths, more exotic board materials are used, such as PTFE based systems. These are more expensive still than glass:polyimide. Given the high data rates attainable for polyimide waveguides (10s of GHz), the introduction of optical interconnects will eliminate the need for more exotic materials than glass:polyimide. Good candidates for demonstration of the board level waveguides then appear to be glass:epoxy and glass:polyimide.

Two approaches are possible for achieving the integration of electronic and optical signal routing. The first approach is to fabricate the optical interconnects directly on the circuit board. The other is to fabricate the optics as a separate entity which can then be laminated on at a given process step. We believe there are numerous advantages to adopting the latter approach.

In the on board processing, the circuit board material must be able to withstand the high temperature waveguide process, eliminating glass:epoxy as a potential substrate material. Also, the circuit boards must be assembled in a non-standard clean fabrication facility to preserve the integrity of the optical materials, leading to increased costs. Lastly, the inability to completely planarize partially fabricated circuit boards yields high additional loss associated with surface scattering.

Adopting the prefabricated optical layer approach means that the circuit board materials will not need to be exposed to the high temperature cycling necessary for curing the polymer films, making possible the use of all currently available substrates. As the optical layer is fabricated separately in a clean room environment, no additional facilities are required for board processing. Since the waveguide substrate is now initially smooth, additional scattering losses are minimized.

Selected samples of the Kapton-backed polymer flex waveguides were laminated to 60 mil thick glass:polyimide circuit board blanks in a high temperature/high pressure autoclave via a standard process at a third-party vendor (PCI, Inc., of Bloomington, MN.) with no previous experience in the processing of optical materials. At 375 degrees C and 150 psi, adhesion of the Kapton to the PCB blanks was excellent over the entire surface, with no visible damage or discoloration. The lamination processes and results are shown in, Figure 11.

In order to minimize expenditure on the board lamination procedure, we typically scheduled our runs to share an autoclave process with standard electrical circuit board runs. This had the added benefit of ensuring that fabricated samples were truly representative of industry-standard processes.




Figure 11 Lamination of flexible waveguide layers into rigid printed wiring board materials

We applied our interconnect design model to optimize the design of the waveguide system for the periodic undulation in the board surface. Two effects can give rise to a non-planar surface. The first is a rapid change in height with a "saw-tooth" like appearance on materials from which copper cladding layers have been removed. Such a situation is likely when polymer waveguides are directly fabricated on a circuit board on which electrical interconnects have already been defined.

A second effect is caused by the weave of the glass reinforcing fabric. With a periodicity of 0.5mm for typical boards, standard planarization techniques cannot be used to yield a flat surface. We have measured best case peak-to-peak undulations of 1.5 microns after planarizing standard glass:epoxy boards with BCB. The effect is also evident in samples which have undergone the post-process board lamination procedure. The high pressure experienced by the polymer over-layer in the lamination process causes it to replicate the underlying topography of the board material with exceptional fidelity.

We have modeled this effect using BEAMPROP, a commercially available beam propagation software package, and found that under certain circumstances board weave does have the potential to impart a high level of loss. The effect is most evident if a significant portion of the total energy is propagating in high order modes, which makes sense because they are the least confined and most subject to surface scattering effects. The simulation assumes a 1.5 micron peak-to-peak square wave perturbation to the substrate, with a periodicity of 0.3mm. This is consistent with the observed characteristics of board laminated waveguides as described above. The model predicts a minimal effect when only low order modes are excited, but an added 1.2dB/cm loss in the case of uniform mode excitation, clearly unacceptable given our design goals.

6.7 ENDFACE PREPARATION

In order to achieve efficient coupling between waveguides, high optical quality waveguide endfaces must be achieved. For preparing the facets of the fabricated samples, we desired a method that was capable of producing an optical quality surface finish, fast, and able to mass terminate many waveguides simultaneously. We wished to avoid mechanical polishing, as this is usually a difficult, involved, and dirty process not amenable to large scale manufacture. In initial experiments, 248nm excimer laser ablation proved incapable of producing the desired results. Somewhat surprisingly, we found that excellent results could be obtained using a common production grade dicing saw.

After mounting on an adhesive-backed carrier, waveguide substrates (either Kapton or Kapton/epoxy) were first cut on a dicing saw into 1cm wide strips parallel to the waveguide arrays. Each strip contained 15 to 32 waveguides. The endfaces were then prepared on the same apparatus, rotating the samples 90 degrees so as to cut perpendicularly across the waveguides. With proper blade selection, feed rate, and blade speed, the quality of the resulting waveguide faces was excellent, without any need for polishing. The waveguide cross-section photograph shown in Figure 10 is an array of waveguides simultaneously terminated by sawing.

The waveguides from any given sample were cut into several different lengths, ranging from 1 to 10 centimeters. This was done to allow subsequent characterization of the facet quality and waveguide attenuation constants.

6.8 BOARD-TO-FLEX CONNECTOR DESIGN

The complete MCM-to-board interconnect requires a means with which to interface polymer waveguides with on-chip VCSELs and detectors. There were two approaches undertaken in this respect under OIT. One approach was to fabricate waveguides directly on or as part of the MCM package in a chip-first approach, terminating the waveguides with 45 degree reflectors in situ using the laser ablation capabilities of HDI. The other approach was to fabricate the reflectors on flexible waveguide arrays which could then be aligned to the MCM and secured using ceramic pedestals. This method is amenable to chip-last or chip-first packaging

The design of the connector for interfacing the flexible waveguide ribbons with the board-level and MCM-level guides was based upon needs for multi-channel capability, passive assembly, low loss, reliability, low cost, and if needed, reconfigurability. Under OIT, Honeywell has developed an interconnect suitable for both interfaces capable of meeting all of these performance goals. The approach is based upon the use of a rigid "key" piece as shown in Figure 12 and 13 into which has been incorporated raised alignment features that mate with similar relief patterns etched into the MCM, board, or flex waveguide arrays.

In the lateral direction, the critical guide-to-guide alignment is accomplished with the aid of two sets of interlocking channels, one set on either side of the waveguide array. These are defined by reactive ion etching or laser ablation in the same step used for the definition of the waveguide cores. In this way, proper etch depth and absolute registration of the fiducials with the waveguides are guaranteed.

The "key" pieces are made of 15 mil quartz which is coated with Ultem and etched to form the mating ridges. The mask used for this is almost an exact negative of the array mask, the only differences being the incorporation of a slight shrink-factor to account for

linewidth reduction in the etching process and the absence of the waveguides themselves. Many pieces can be batch fabricated from a single 4-inch quartz substrate, cutting each to the desired proportions by diamond saw. In the future, these parts might be made of plastic via injection or compression molding, but the quartz version has served for this demonstration.

When assembled, the surface of the quartz key serves as a reference level for the vertical alignment of the arrays. Proper alignment only requires that the thicknesses of the polymer layers be consistent and reproducible, both across sample and from run-to-run. This is not terribly restrictive, because quality control in thin film deposition has advanced to the point where such things can be regulated to very strict tolerances.

Interfacing with board-level waveguides requires the presence of a recessed slot to accommodate the body of the flexible ribbon. Our initial efforts focused on using 248 nm laser ablation to process this receptacle. Using a local company (Spectronics), we characterized cuts in polyimide waveguide layers produced with excimer lasers under various process conditions. Spectronics is a small local company with capabilities representative of those available in a large number of such vendors across the country.

Using a Lamda Physik excimer laser with 850 mJ pulse energy at 60 Hz repetition rate, two experiments were performed on standard PCI laminated waveguides. In the first, an aperture of 0.04" was used in conjunction with a lens to write a line approximately 150 microns in width. In the second, the full power of the laser was used in conjunction with a metal contact mask in a lensless system. In the first case the board material with waveguide layers was translated under the laser at a speed of 0.05"/sec, while in the second the board plus mask were moved at 1"/sec. Results were similar for both processes, with sidewalls being smooth but not vertical, and with a small amount of carbonization on




Figure 12 The connector assembly of a flexible waveguide layer with board-laminated and MCM level waveguides

the surface. The optical quality of the waveguide endfaces appeared good, but the fact that they were not vertical suggested that guide-to-guide butt coupling losses in the connector would be high.

The excimer laser ablation studies demonstrated a clear potential for use in board-level processing of slots and waveguide terminations. However, further development was required to achieve results consistent with program goals. As we were unwilling to commit the resources necessary to refine the process, an alternate approach was chosen for the sake of expediency.

The method adopted was to machine the receptacle using standard metalworking machinery, and then use our dicing saw technique to establish an optical grade end-face on the waveguides. This approach, while slightly less elegant than laser ablation, proved to be more than adequate for demonstration purposes.

Connector assembly is accomplished in two steps. First, the flexible waveguide array is affixed to the quartz key with the aid of a low viscosity room temperature curing adhesive. The pieces are held together with manual pressure to engage the alignment features while a drop of adhesive placed at the edge of the key is allowed to spread to all mating surfaces through capillary action. This is done to avoid any misalignment due to the presence of a glue layer between the components. The depth of engagement of the alignment features is less than one mil, nevertheless the tactile sensitivity required for proper assembly is well within the capability of the average human on less than 12 cups of coffee. Pressure is maintained until the adhesive is fully cured.

Figure 13 is a cross-sectional view depicting the engagement of the alignment features before and after gluing. The uniform contact between the Ultem keys and the etched channels can clearly be seen. Also evident in the picture is a "dovetail" structure caused by the undercut of the keys during etching. The adhesive subsequently filled in these regions, locking the two pieces together and giving the joint its high observed strength.




Figure 13 A cross-section of the alignment features before (left) and after they are engaged.

Next, the protruding end of the key is positioned over the recessed slot in the board, roughly in the proper orientation. Applying slight pressure, the assembly is rocked back and fourth until again the positive engagement of the channels with the key is felt. The key is then slid longitudinally until positive contact between the two waveguide interfaces is achieved. The assembly is then secured with adhesive as before.

A 24 channel board-to-flex connector, passively aligned and secured in this fashion was evaluated in order to determine the uniformity of the coupling efficiency across the waveguide array. The assembled connector is as shown in Fig.14. The measurement was performed by endfire coupling 830nm light from a semiconductor laser into the rigid board laminated waveguides and collecting




Figure 14 An assembled MCM-to-board connector

the light transmitted at the far end of the flexible MCM coupler. The measurements were compared to those taken using a standard sample of the same net overall length in order to factor out the losses due to waveguide attenuation.

The data from this experiment is presented in Fig.15. Over the entire array, an average per waveguide coupler insertion loss of only 1.26 dB was observed with a standard deviation of 0.58 dB. The lowest recorded loss was a mere 0.6 dB, corresponding to an efficiency of over 87%. This is an outstanding result, given that it was achieved using only passive alignment and relatively crude assembly techniques. The adhesive used did provide some degree of index matching, thus improving the coupling efficiency over that obtained with an air gap between the components. A more careful choice of adhesive with better index matching properties should yield an even better result.

As can be seen, there is a systematic deterioration of the coupling efficiency towards the outermost edges of the array attributable to curling of the sample. This effect can be reduced or eliminated by improving the assembly procedure and with stricter process control of the alignment feature etching.

Prior and subsequent to testing, the above sample underwent a great deal of manual handling and bending, stressing the connector assembly in order to assess its rigidity. We are able to report that the device showed no sign of failure or degradation in either its optical or mechanical properties.




Figure 15 Measured loss of the flexible board-to-MCM connector

6.9 ENVIRONMENTAL TESTING

One concern was the potential for deterioration of the waveguides due to moisture absorption in high humidity ambient. According to manufacturers specifications, Ultem should absorb a maximum of 1% of water by weight, while BCB should absorb a maximum of 0.2% at 80% RH. Our metal-coated system should restrict or prevent the ingress of water into the waveguides. Nevertheless, it was necessary to verify the performance of the Ultem/BCB waveguide system with respect to the appropriate military-standard tests.

In accordance with MIL-spec.139-49/13, an experiment was performed to evaluate the performance of the BCB-ULTEM system in high humidity environments. In the experiment, standard process board-laminated waveguide samples were placed in a sealed environmental chamber and ramped up to 85% RH while holding the temperature constant at 24 degrees C. This condition was maintained for a period of 48 hours, after which the samples were removed for immediate testing. In order to demonstrate worst-case exposure, the tested samples were not processed with the encapsulating metal sealing layer. Examination of the data showed an average *reduction* in loss of 0.05dB/cm across all channels, not indicative of any adverse effects from the exposure and possibly attributable to experimental error. No physical deterioration of the samples was observed.

In a more severe test, similar samples were immersed in boiling water for extended periods of time and periodically removed in order to measure sample throughput. Initial tests in which waveguides were boiled for 13.5 hours showed ambiguous results, with an increase in loss after boiling being reversed by drying. We were however unable to demonstrate any increase in loss by boiling the

sample a second time for the same period. Even after an additional 36 hour exposure to the boiling water, no degradation was detectable in either the optical or mechanical characteristics.

Tests were also performed in order to access any potential degradation in sample quality due to temperature extremes. In accordance with MIL spec 139-49/13. waveguides identical to those used in the humidity exposure experiments were temperature cycled from room temperature to 105 degrees C, held at this temperature for 15 minutes, ramped down to -60 degrees C over a period of 15 minutes, held again for 15 minutes, and ramped to room temperature. Each sample underwent a total of 10 such cycles before testing. Physical examination and comparisons of throughput measurements to those taken before thermal cycling showed no evidence of deterioration in structural or optical quality.

7 OPTICAL MULTICHIP MODULE DEVELOPMENT

The OIT program focussed on the addition of optical interconnects into electronic systems employing multichip modules. Our approach centered on the application of existing fabrication techniques for multichip modules to optoelectronic packaging rather than the application of optical fabrication techniques to electronic packaging. By adopting this approach we believe we would reduce the risk, cost, and time to provide usable optically interconnected electronic systems.

7.1 MCM INTEGRATION

For short distance interconnects where moderate propagation loss and modal dispersion can be tolerated, the high costs associated with single-mode interconnect packaging can be avoided through the use of multi-mode devices. The resultant large dimension components are much more easily interfaced with each other, and coupled to external devices such as optoelectronic emitters and detectors. This is evidenced by the recent interest in fluorinated plastic fiber as a medium for short-haul data links which can be assembled and repaired in the field. Indeed, passive alignment was a key motivation in OIT, where all assembly can be performed by an untrained technician using currently available pick-and-place machinery, without the necessity for any sort of feedback to optimize necessary optical alignments.

Our design for multi-channel MCM-to-MCM interconnection is depicted in Figure 12. The idea is to use a flexible waveguide to interface an MCM to board level waveguides over which data can be exchanged with other modules integrated on the same carrier. This design is compatible with standard MCM packages in which arrays of leads pass through the module walls. We wished to make the optics look as much as possible like the electrical interconnects with which industry is familiar. This requires that registration features be incorporated for passive alignment of all optical interfaces, enabling the use of existing tools for insertion and removal of the MCMs in a typical assembly environment.

There were four primary tasks involved in the development of the MCM-to-MCM link. The first task was the realization of a polymer-based waveguide technology on which the entire interconnect strategy would be based. We have already described the selection process, but a key requirement in the design of the polymer waveguide system was the identification of materials with suitable optical properties, being simultaneously compatible with MCM fabrication processes. With the material system specified, the waveguide structure was defined, and the necessary fabrication expertise developed with which to implement it.

The next task was to devise a method of incorporating optical layers in the packaging of multichip modules, and interfacing with VCSEL and detector die on the modules. With an eye towards versatility, both MCM-C "chip last" and General Electric's HDI (High Density Interconnect) chip-first MCM packaging schemes were investigated.

The third task in the connector design was the development of a means with which to align and affix the flexible waveguide array to the board and MCM level waveguides. This was probably the highlight of the program, demonstrating truly passive assembly of multi-channel waveguide arrays with low insertion loss.

The fourth task was the development of the MCM modules themselves. If the functionality and performance benefits of optical interconnects was to be asserted, it was deemed important to demonstrate multi-gigabit data transmission over parallel channels at low bit-error-rates.

In the design of the MCM-level waveguides and MCM-to-board connector, as with the board level waveguides, the design of both the electrical and optical components relied extensively on computer-based modeling. As the availability of commercial software in the area of photonics design is severely limited, many of the necessary codes had to be developed under this program as described previously.

Discrete device packaging and system interconnectivity are prevalent concerns in the development of photonic components for the telecommunications industry. The need to use long distance transmission media, EDFA's, and devices based on coherence effects requires the use of single mode optics, where stable sub-micron component alignments are necessary for reliable operation. Currently, such tolerances demand that all critical assembly steps be performed manually using high precision equipment and active feedback for optimization. It is not surprising then that assembly costs typically surpass all others by a wide margin. Therefore, reducing packaging costs will be critical to the long range acceptance of many new and proposed technologies. By using multimode waveguides throughout, and by employing passive alignment and assembly techniques to the greatest extent possible, we aimed to develop an optical interconnect and packaging technology which would allow low cost and environmentally rugged implementation of optical interconnects.

7.2 HDI WAVEGUIDE PROCESSING

One of the two MCM approaches used in OIT was the chip-first High Density Interconnect system developed by GE. This packaging system is unusual in that the chips are integrated into the substrate by placing them in receptacles in the substrate to form a nominally planar surface. A layer of Kapton is then laminated onto the substrate using a thermoplastic adhesive (GE's Ultem), and metal deposited. The metal is then coated with resist, and a laser scanning system is used to define an interconnect pattern in the photoresist. The position of the written waveguides is governed by a computer controlled system which derives placement information from a visual inspection of the substrate. In this way the system is able to write an interconnect which accommodates the displacement of the die from their ideal position. Each substrate then is likely to be very slightly different when completed, although with the same basic interconnect pattern. To implement vias where needed, excimer laser ablation is used prior to metal deposition. Further layers of Kapton are added and the process repeated to form a multilayer interconnect. Finally, the completed substrate is trimmed and mounted in a package, and leads attached.




Figure 16 Schematics of MCM fabricated using GE's HDI process to implement both electrical (left) and optical (right) interconnects.

The attraction of the process compared to other technologies lies in the high coverage of the substrate with ICs, and in the planar geometry which allows the exploitation of IC-type fabrication processes in which economies of scale can be exploited. In this program we designed a process by which the HDI equipment could be used to write optical waveguides and electrical interconnects with acceptable performance. The goal was not to implement the lowest possible loss optical waveguide in an MCM but rather to exploit the available infrastructure to produce an adequate interconnect.

7.3 HDI PACKAGING

The HDI approach to MCM integration requires the deposition and definition of optical waveguides directly on the MCM substrate, in a manner similar to that used for the electrical interconnects. HDI is a chip first assembly where the optical interconnects must be adaptively written after the electronic die have been embedded in the MCM substrate. In the HDI process, laser etching is the standard method used in the definition of inter-layer via interconnects. Using thick electroplated copper masks defined by direct write laser lithography, an excimer laser is used to ablate material, thus replicating the mask pattern in the underlying polymer layers. The method is very versatile, using large beam blanket laser exposures for field etching or small beam sizes for highly localized material removal.

The HDI direct write laser process can be directly extended to the fabrication of the on-MCM waveguide components necessary for the interconnect demonstration. Instead of vias, the copper mask can be used to define the on-chip waveguides and adaptive bends necessary for interfacing to the optoelectronic die. This can be done rapidly using the wide field etching capabilities of HDI. Using a narrow beam and a dithered scan, the required 45 degree turning mirrors can also be fabricated.

7.4 HDI VCSEL PACKAGING

The initial step in the HDI process is to locate the various die which comprise the MCM in processed slots which have been recessed into a ceramic carrier. The entire assembly is then overcoated with a polymer encapsulation through which vias are defined to enable the electrical interconnections. This is a multi-layer process, allowing for several levels of metalizations, each separated by an additional

polymer coating. As semiconductor lasers are typically very sensitive to processing history and environmental conditions, it was necessary to demonstrate that the VCSELs proposed for use in the MCM interconnect demonstration are able to survive the HDI process.

Several VCSEL arrays were supplied to GE for electrical packaging using their standard HDI process. The arrays used were developed under DARPA's Optoelectronic Technology Consortium program (OETC), and consist of 32 VCSELs on a pitch of 140 microns complete with bonding pads. The VCSELs each have an active region of 15 μm diameter, typical full angle at half maximum (FAHM) of no greater than 20-degree at a typical driving current of 10 mA.

Two samples were fabricated, one at a 200 C lamination temperature, and the other at 260 C. The HDI interconnect lines were designed to give a controlled 50-ohm impedance, although matching components were not added to this circuit to compensate for the mismatch at the laser. These were not deemed to be necessary because the laser impedance is actually quite close to 50 ohms. The top view of the HDI packaged OETC VCSEL array chip is shown in Figure 17.




Figure 17 An HDI packaged VCSEL array

The returned samples were characterized for electrical and optical performance, and the data compared to that taken before packaging. It is obvious that no change occurred in this device. We also found no change in threshold current or voltage, or optical output characteristics, demonstrating that optoelectronic devices can be packaged using the standard HDI process without suffering measurable degradation.

Process

- Sputter Cu on Ultem
- Pattern Cu using adaptive laser lithography (left)
- Laser ablated waveguides in Ultem LC-1010 (right)




Figure 18 Etching of polymer waveguides using laser ablation

We have demonstrated waveguides and reflectors fabricated by GE using HDI laser processing of Ultem layers deposited on Kapton and on oxidized silicon substrates. The etched structures include straight waveguides, curves of various radii, and adaptive bends, structures similar to those that would be required for MCM packaging. Figure 18 shows the results of etching a 50 micron wide by 15 micron thick waveguide in Ultem 1010LC using a copper mask. As can be seen, the sidewalls are vertical, and no rougher than would be expected from laser etching.

Preparation of the waveguide termination in the HDI process involves adaptively laser ablating the 45 degree mirrors over VCSEL and receiver die embedded in the MCM substrate. This can be done either by angling the incident laser beam with respect to the substrate, or by dithering a beam oriented normal to the material surface. Direct angled ablation requires complex fixturing and




Figure 19 Preparation of 45-degree angled facets

precise control over the substrate orientation, and so represents a significant modification to the standard HDI process. Laser dithering only requires control over beam scan speed and intensity, and so was determined to be the more viable approach.

The 45 degree reflectors in our standardized 15x50 micron BCB/Ultem/BCB channel waveguides prepared by dithered laser ablation at GE CRD were evaluated by measuring the output coupling of 830nm laser light which had been endfire coupled into the butt terminated end of the waveguides. Several samples were measured in order to factor out variations due to input coupling efficiency and sample variability. The efficiency of the reflectors was deduced through comparison to unterminated waveguides of the same length.

Results of the measurements indicated an excess coupling loss in the range of 15-19 dB. In order to understand the high excess loss, we inspected the laser-ablated facets with a microscope and an alpha-step profiler. The following defects contributed to the large excess loss: (a) The ablation depth was not deep enough, and the angled facet did not cut through the waveguide core. Under the optical microscope, traces of waveguide core were still visible under the bottom of the laser ablated region. Alpha-step profiles also showed that the deepest spot of the ablated region was no greater than 28 um. The waveguide sample had a core of 20 um, and an upper BCB cladding layer of 8-12 um. The measurements suggest there was up to 4 um of unablated core, and that an

additional 2-4 μm bottom region had less than the intended 45-degree slope. The ablation depth should be at least 35-38 μm from the upper surface to assure the waveguide core is totally cut through, and thus that reflection surface in the core region is straight (free of rounding). (b) The actual slopes of the facets are much less than 45-degrees. We measured the slopes of the ablated facets using an alpha-step profiler, and found they vary from 34-39 degrees. This was much less than the intended 45-degrees, and hence resulted in excess loss. This is due in part to difficulties in achieving uniform etching rates for both BCB and Ultem.

Similar to the turning mirrors, waveguide components fabricated using the HDI process also exhibited less than desirable performance. Straight and curved waveguides processed by direct laser ablation exhibited exceedingly high propagation loss. Visual observation of the samples indicated problems with the algorithms used for controlling the laser scan, resulting in grossly stepped approximations to angled waveguides and severe stitching errors.

The above results indicate the degree of complexity required to implement on-MCM optical components using the standard HDI process. We have, however, demonstrated all of the necessary components for an HDI processed MCM interconnect. While the fabrication can be optimized, we were unable to do so under OIT due to the fact that GE exhausted contract funds.

7.5 MECHANICAL PREPARATION OF END-FACES

The alternative method adopted for fabrication of the turning mirrors was to terminate the waveguide ribbons before attachment to the MCM using a sawing technique on a dicing machine, similar to the method used for bulk waveguide termination. This approach is suitable to either chip-first or chip-last packaging. The reflectors were processed offline on Kapton backed waveguide samples, cutting across the polymer layers at 45 degrees to the surface normal. Experimentation was able to establish the correct blade selection and cutting parameters necessary to yield cuts of very high optical quality. This method has the added advantages of being fast and easy to implement using existing equipment.

The test setup used for evaluating the quality of the sawed reflectors is shown in Figure 22. The flexible waveguide is a standard sample consisting of an ULTEM (core) and BCB(cladding) fabricated at Honeywell's facility on a 5 mil (125 μm) Kapton substrate. One end of the waveguide was prepared with a 45-degree facet, which was then coated with metal (Ti/Au) as the final reflector surface. The particular sample under test had a waveguide core 17 μm thick and 50 μm wide. The VCSEL chip used was the standard device described previously.




Figure 20 Mechanical preparation of angled end-faces

During the testing, the waveguide sample was placed directly on top of the VCSEL chip in order to minimize the distance between the source and the waveguide core, and hence maximize coupling efficiency. The estimated distance was about $Z=150\text{ }\mu\text{m}$. The best coupling loss measured was -4 dB, comparable to the modeled theoretical maximum coupling of -3.6dB obtainable with our configuration. We also measured throughput vs. lateral alignment (both x and y direction) of the waveguide with respect to the VCSEL. The results are shown in Figure 21. As can be seen, the tolerable transverse and longitudinal misalignments for -3dB coupling were measured to be $\pm 15\text{ }\mu\text{m}$ and $\pm 10\text{ }\mu\text{m}$ respectively. This level of tolerance can easily be achieved using our passive alignment scheme.




Figure 21 The test arrangement for evaluating the VCSEL-to-45-degree reflector coupling.

The reflectors produced by angled sawing were far superior to those produced by laser ablation. In agreement with the model, the loss has subsequently been reduced even further by terminating the waveguide samples upside down, so that the polymer layers are closest to the acute angle of the reflector. In this configuration, the light from the VCSEL does not pass through the Kapton substrate material, and the spot size at the waveguide aperture is reduced correspondingly. Coupling mirrors with excess losses as low as 1.0 dB/mirror have been fabricated in this fashion.

The inherent simplicity in fabricating pre-terminated ribbons of polymer waveguides, and the good optical characteristics observed suggest that the attachment of these prefabricated elements on an individual basis to features defined on each optoelectronic component array would be a low cost solution for all but the most complex MCMs.


8 OPTICAL INTERCONNECT DEMONSTRATION

8.1 TRANSMITTER (TX) AND RECEIVER (RX) MODULE:

As an alternative to GE's chip-first MCM package, we used more conventional chip-last thick film MCM-C package for the Tx and Rx modules (fabricated at Honeywell's Central Technology Organization (CTO)). These modules are designed to use Honeywell's 850 nm wavelength VCSELs, Si p-i-n photodetectors from UDT, and a chip set of laser driver, transimpedance amplifier, and variable gain amplifier from HP.

The schematics of the Tx and Rx channel are shown in the upper and lower part of Figure 22, respectively. Photograph of finished Tx and Rx MCMs are shown in Figure 23. Microphotographs of all the active chips in each MCM are also shown as inserts in Figure 23.

Tx Design Schematic:



Rx Design Schematic:





Figure 22 OIT Tx (upper) and Rx design schematics.

OIT Tx module:



OIT Rx module:




Figure 23. OIT Tx (left) and Rx (right) MCM and active devices.

8.2 OIT LINK BUDGET

Each Tx MCM and Rx MCM consist of three data channels, designed at 1 Gbps per channel. The link loss budget is shown in Table 3.

Table 3. OIT link loss budget

Link element	Remark	Unit	Worst	Typical	Best
Tx Output power		mW	1	1.5	2
VCSEL	Output power @ 25C	dBm	0.00	1.76	3.01
VCSEL/WG	45-deg reflector interface	dB	-3	-2	-1
Flex WG	WG propagation loss 3cm	dB	-1.8	-1.2	-0.75
WG/WG	Flex-board WG connector	dB	-1.5	-1	-0.5
Board WG	Board WG losses, 4cm	dB	-2.4	-1.6	-1
WG/WG	Board-flex WG connector	dB	-1.5	-1	-0.5
Flex WG	WG propagation loss 3cm	dB	-1.8	-1.2	-0.75
VCSEL/WG	45-deg reflector interface	dB	-3	-2	-1
	Power at Rx	dB	-15.00	-8.24	-2.49
Rx Sensitivity		dBm	-18	-18	-18
Margin		dB	3.00	9.76	15.51
	Modal noise penalty		-1	0.5	0
	RIN Penalty		-1	0	0
	Finite Extinction Penalty		-2	-1	0
	Total penalties		-4	-0.5	0
	Revised margin		-1.00	9.26	15.51

8.3 OIT DEMONSTRATION SETUP

The OIT final demonstration link schematics are shown in Figure 24. They consist a pair of Tx and Rx MCMs inserted in corresponding modules evaluation boards. The Tx and Rx modules are connected by a polymer optical waveguide assembly. The waveguide assembly is made of three segments, two flexible waveguide ribbons and one rigid board integrated waveguide piece. The three segments are jointed by two flex-to-board waveguide connectors. Arrays of waveguide channels are connected, simultaneously and passively, at the two connectors interfaces using the same connector scheme as described in Section 6.8. The other two ends of the flex waveguide ribbons are finished with 45 degree end facets using a simple polishing process. The 45 degree facets are then aligned actively to the VCSELs in the Tx module and pin detectors in the Rx module and permanently attached.




Figure 24 OIT link setup schematic and key link components and interfaces.




Figure 25 OIT demonstration link testing setup.

The total length of the waveguide link is about 11 cm long, and the total loss was measured to be less than 10 dB. We estimated the waveguide propagation loss is about 0.3 dB/cm or total loss of 3.3 dB. This leaves about 6.5 dB of loss to be attributed to the four coupling interfaces – two 45 degree facets-to-device interfaces and two flex-to-board connector interfaces. This implies that we have, on average, approximately 1.6 dB/interface of coupling loss.

8.4 OIT DEMONSTRATION TESTING RESULTS.

Figure 26 shows the demonstration setup. Each Tx channel needs a single power supply, dc bias and modulation amplitude controls, is driven with a pair of differential signals inputs. Each Rx channel requires a single power supply, a dc gain control, and a single ended signal output.

We were able demonstrate the link at data rate up to 1 Gbps. Figure 26 shows a eye diagram of the link running at 1 Gbps. We were able to continuously run the link at this data rate for several hours without error count, or bit-error-rate (BER) of less than 10^{-12} .






Figure 26 OIT link eye diagram at 1 Gbps.

We were able to run more than one channel simultaneously at various speed. Figure 27 shows two (a) two channels running at 250 Mbps in a 16-bit fixed pattern, and (b) two channel running at 622 Mbps at pseudo-random bit stream (PRBS).



(a) Ch.1 + Ch.2, 250 Mbps



(b) Ch.1 + Ch.2, 622 Mbps

Figure 27 OIT link demonstration: two link channel operating simultaneously at data rate of (a) 250 Mbps 16-bit fix pattern, and (b) 622 Mbps of PRBS.

9 CONCLUSIONS

We have successfully demonstrated the first MCM-to-board-to-MCM optical interconnect based on VCSELs, polymer waveguide, and MCM-C packaging technologies. We demonstrated low-cost component fabrication, and passive assembling techniques that are compatible with existing electronic manufacturing processes. We believe these approaches are critical to optical insertion into real system applications. We fabricated flexible polymer waveguide ribbons, board integrated optical waveguide, and passively aligned flex-to-board waveguide connector. We also demonstrated optical links based on conventional MCM-C package with multiple data channels at data rate up to 1 Gbps per channel.