Design Rules for Highly Parallel Free-Space Optical Interconnects

Andrew G. Kirk, *Member, IEEE*, David V. Plant, *Member, IEEE*, Michael H. Ayliffe, Marc Châteauneuf, and Frédéric Lacroix

Invited Paper

***Abstract*—Recently, a number of successful free-space chip-to-chip and board-to-board optical interconnects have been demonstrated. Here, we present some of the design rules that can be derived as a result of this work and also as a result of numerical and theoretical analyzes. We draw a number of conclusions. In the area of optoelectronic very large scale integration (VLSI) design, we suggest that differential electrical and optical transceiver designs provide the best performance. In the area of optical design, we present scaling and system partitioning laws for clustered optical relays and determine the interconnect distances at which microlens or macrolens systems are more suitable. We also show that the ease with which two modules can be aligned can be related to the optical invariant of the system and is, thus, a function of the size of the detector and the numerical aperture of the detector optics. Finally, we show that when multiple optical components must be aligned, very high individual component tolerances are required if the system as a whole is to have a high chance of success.**

***Index Terms*—Free-space optics, hybrid integrated circuit packaging, microoptics, optical interconnection, optical receivers, optical transmitters, optoelectronic devices, vertical cavity surface emitting lasers (VCSELs).**

I. Introduction

PARALLEL optical interconnects are capable of providing high bandwidth communication links both within and between high-performance electronic systems. The advantages of optical communications for long-distance interconnects are well known, and provide the motivation for modern optical fiber networks. Optics is now challenging copper at shorter and shorter distances. The benefits of optical interconnects include reduced signal distortion and attenuation, lower power requirements, lighter components, potentially lower costs, and much greater immunity to electromagnetic interference. A thorough review of these physical issues is provided in [1] and [2]. In the commercial arena, several manufacturers now supply optical fiber ribbon-based parallel optical data links (PODLs) of 8 to 12 channels, operating at data rates of up to 3.25 Gb/s per channel over distances of 100–1000 m (depending on bit rate). However, there are applications where many more parallel channels are required and in some cases the interconnect configuration is more complex than a simple point-to-point link. For example multiprocessor computers, telecommunications switches, and embedded systems all require highly parallel interconnections. In particular, the concept of the direct sourcing and termination of optical signals on silicon has been proposed as a method to relieve the off-chip communication bottleneck. Fig. 1 is a schematic representation of this concept.

In recent years, significant progress has been made in the enabling technologies for parallel optical interconnects for digital systems. This includes the hybrid integration of vertical cavity surface emitting lasers (VCSELs) with silicon integrated circuits [also referred to as optoelectronic very large scale integration (OE-VLSI)], assembly techniques for free-space optical interconnects (FSOIs), the development of fiber arrays and the integration of optical waveguides with printed circuit boards.

In this paper, we will codify a set of design rules and guidelines for free-space optical interconnects for chip-to-chip and board-to-board communication. These rules have emerged after implementing a series of optical interconnect demonstrator systems and subsystems, including a 512 channel optical ring interconnect [3], [4], a 512 channel bidirectional interconnect [5], [6], a 1080 element OE-VLSI chip which implements a variety of data transmission protocols [7], [8] and a highly misalignment tolerant interconnect that makes use of spatial redundancy [9]. We have also implemented a range of alignment and assembly schemes [10]–[12] and performed analyses of misalignment tolerance [13], [14] and scalability for free-space optical interconnects [15]–[17]. In this paper, we also consider the published results of other researchers in order to obtain a complete a picture as possible of the issues involved in FSOI design and

Manuscript received December 16, 2002; revised February 17, 2003. This work was supported by BAE Systems under a contract with DARPA/Army Research Laboratory via the VLSI-Photonics program, DAAL01-98-C-0074, the Canadian Institute for Telecommunications Research under the NCE Program of Canada, the Canadian Microelectronics Corporation, and NSERC and FCAR fellowships.

A. G. Kirk and D. V. Plant are with the Department of Electrical and Computer Engineering, McGill University, Montreal, QC H3A 2A7, Canada (e-mail: andrew.kirk@mcgill.ca).

M. H. Ayliffe was with the Department of Electrical and Computer Engineering, McGill University, Montreal, QC H3A 2A7, Canada. He is now with Ignis Optics, San Jose, CA, USA.

M. Châteauneuf was with the Department of Electrical and Computer Engineering, McGill University, Montreal, QC H3A 2A7, Canada. He is now with the Defense R&D Canada, Val Belair, QC G3J 1X5, Canada (e-mail: marc.chateauneuf@drdc-rddc.gc.ca).

F. Lacroix was with the Department of Electrical and Computer Engineering, McGill University, Montreal, QC H3A 2A7, Canada. He is now with Ciena Inc., San Jose, CA (e-mail: flacroix@ciena.com).

Digital Object Identifier 10.1109/JSTQE.2003.812482

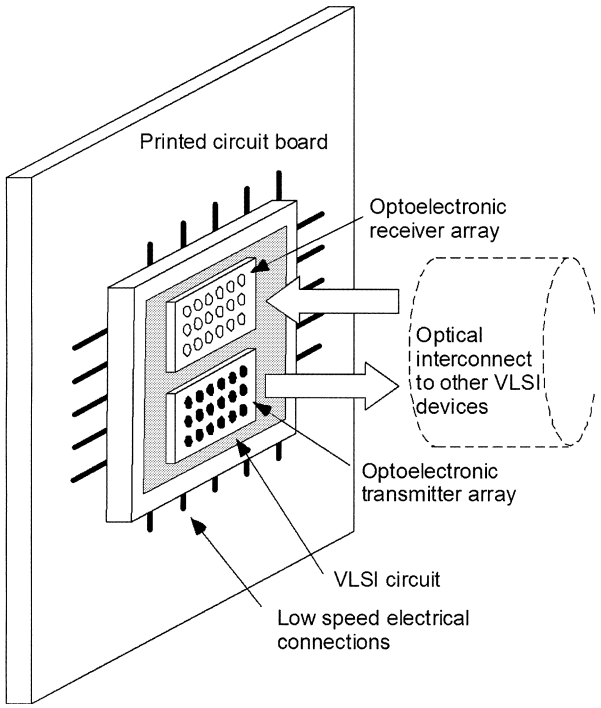

Fig. 1. Schematic representation of a board-to-board parallel optical link with direct termination of optical signals on-chip.

implementation. The rules that we introduce cover a range of areas, including optoelectronic device layout, electronic circuit design and operation, optical design and packaging.

In developing these rules we have made several assumptions:

1) optical sources in a free-space optical interconnect will be VCSELs, and may be either single mode or multimode;
2) heterogeneous integration using flip-chip bonding and substrate removal of optoelectronic devices with complementary metal-oxide-semiconductor (CMOS) based transceiver and processing logic;
3) the optoelectronic detectors are GaAs p-i-n diodes;
4) optical interconnect distance is less than 150 mm (i.e., characteristic of interchip and board-to-board spacings);
5) interconnects are all rigid free-space optical systems;
6) interconnect has a point-to-point transmissive topology;
7) interconnect should incorporate a realistic tolerance to misalignment;
8) interconnect is capable of assembly by passive techniques (i.e., should not be necessary to switch emitters on use active alignment to ensure that light arrives at the detectors).

These assumptions will be further clarified in the following section when we will introduce a generic free-space optical interconnect for chip-to-chip or board-to-board communication. Our research has been based on the use of GaAs optoelectronic devices and so the wavelength of operation is 850 nm. However, many of the guidelines can be adapted to other wavelengths.

II. FSOI Design Space

The first task is to identify the design space in which FSOI for chip-to-chip and board-to-board communication must operate. Since this technology is aimed at future high performance

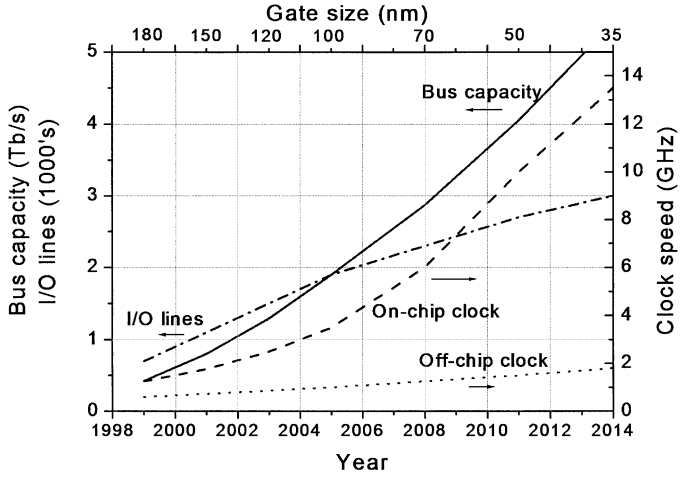

Fig. 2. Projected evolution of on-chip clock speed (dashed line, right axis), off-chip clock speed (dotted line, right axis), number of high speed I/O lines (dash-dot line, left axis) and total bus capacity (solid line, left axis) as a function of time and transistor size for high performance systems.

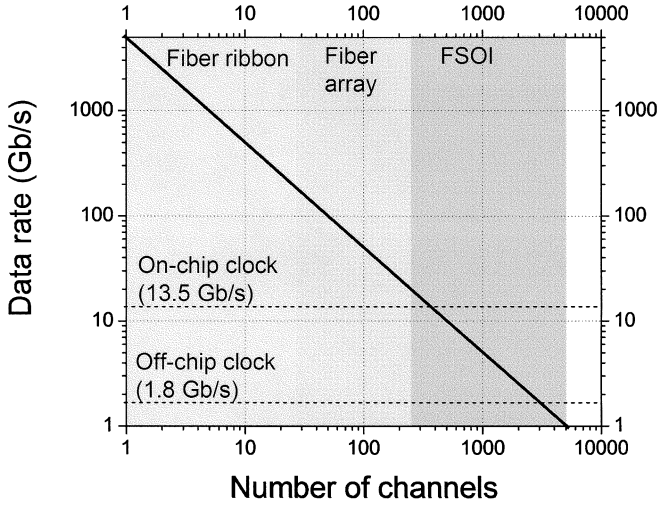

Fig. 3. Projected off-chip I/O requirements for 2014. The graph shows the data rate per channel necessary to achieve an aggregate data rate of 5 Tb/s.

electronic systems, we refer to the International Technology Road Map for semiconductors [18]. Fig. 2 shows the projected increase in VLSI transistor on-chip clock speed, off-chip clock speed, number of high-speed off-chip clock lines and total off-chip I/O capacity as a function of time for high-performance systems, taken from [18]. It can be seen that by 2014, the off-chip clock speed is projected to reach 1.8 GHz and the width of the off-chip bus is also projected to increase to 3000 high speed lines, with a total projected off-chip I/O capacity of 5 Tb/s. The on-chip clock speed is projected to reach 13.5 Gb/s. From this data, we can obtain a view of the design space for off-chip interconnects by the year 2014, under the assumption that the necessary off-chip bandwidth will be 5 Tb/s. This is shown in Fig. 3, where the data rate per channel necessary to achieve 5 Tb/s is plotted as a function of the number of parallel channels. We have shaded different regions corresponding to different possible optical interconnect formats. It seems apparent that one-dimensional (1-D) fiber ribbons will not be capable of delivering the necessary bandwidth as data rates of

more than 200 Gb/s per channel would be necessary when, for example, a 24-fiber ribbon is used. The step up in density is provided by fiber arrays. At present fiber arrays of 8×8 to 16×16 have been demonstrated [19], [20] but volume fabrication of these elements remains to be properly developed. Even a 256 channel fiber array would require per channel data rates of 19 Gb/s. Moreover, for the very short distances that we consider here, the limited bending radius and high fabrication cost of fiber arrays may render them impractical. From Fig. 3, it is apparent that if we assume that the optical interconnect channel data rate is not to exceed the projected on-chip data rate of 13.5 Gb/s (which would otherwise require the use serialization/deserialization circuits) then a 5 Tb/s aggregate data rate implies the presence of 370 optical lines. If we assume that the off-chip optical links run at the projected electrical off-chip clock speed of 1.8 Gb/s, then this implies approximately 2700 optical lines. We suggest that these two values represent the boundaries for optical solutions to the off-chip interconnect problem, and that they also represent the domain in which free-space optical interconnects represents a possible solution. Several experimental free-space interconnects have achieved the lower end of the parallelism range [4], [6], [21], [22] and we have recently reported an OE-VLSI application specific integrated circuit (ASIC) with 1080-optical I/O that approaches the middle of the required parallelism range [7], [8]. Other researchers have demonstrated matrix addressed VCSEL arrays with 4096 outputs [23].

Fiber image guides and fiber image conduits [24]–[26] represent another possible alternative high-density interconnection medium. However, fiber image guides suffer from transmission nonuniformity unless large channel spacings are used, and questions also exist as to their potential for low cost fabrication. Image conduits are closer to rivaling free-space optics in terms of performance. Since they are also a rigid technology, a comparison should be made in terms of cost and performance, but we will not consider them further here. It has been suggested that wavelength division multiplexing could be used as an alternative to spatial multiplexing, but we will also not consider that option [27].

An example of an experimentally realized point-to-point free-space optical interconnect is shown in Fig. 4 [6]. In this system, two OE-VLSI chips are bidirectionally interconnected over a distance of 86 mm. Each chip contains 256 transmitters (VCSELs) and 256 photodetectors and operates at a wavelength of 850 nm. The interconnection distance was selected such that by inserting two additional prisms into the beam path it could interconnect two boards in a bookshelf configuration [6]. The relay optics are based on clustered diffractive lens and the array density that was achieved was 28 channels/mm^2. A schematic diagram of this system is shown in Fig. 5 This will be used as a framework for the ensuing discussion of design rules as it represents the necessary components of a free-space optical interconnect. These are: a printed circuit board (PCB) [or multichip mode (MCM) substrate], an integrated circuit which contains control circuitry and the optical transceivers, the optoelectronic devices, collimation optics (necessary when VCSELs are used), relay optics and then various packaging levels, including electrical packaging (which concerns issues

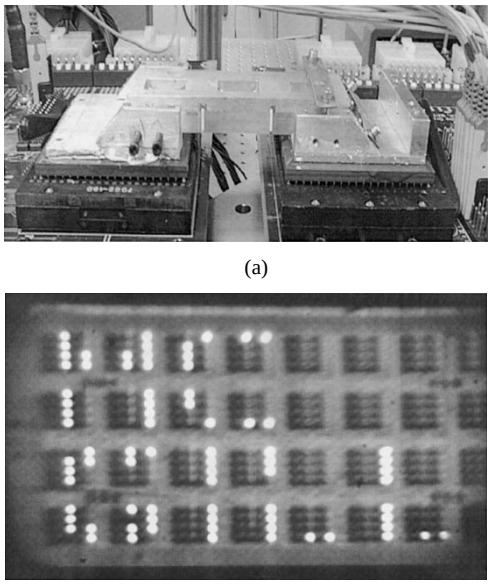

(a)

(b)

Fig. 4. (a) Experimentally realized 512 channel inter-chip interconnect system. (b) Transmitted beams.

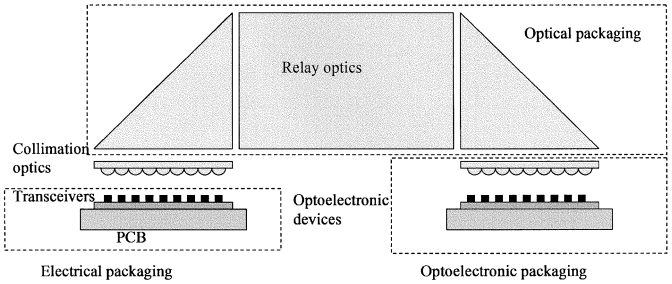

Fig. 5. Generic free-space optical interconnect.

such as signal integrity and power dissipation), optoelectronic packaging (which concerns the attachment of the optical components to the OE-VLSI chip), and the optical packaging (which concerns the assembly of the relay optics). In designing a system, some of the issues which must be addressed are: transceiver layout; the use of single-ended or differential circuitry; power consumption; the use of error coding; the incorporation of built-in self test; the optical interconnect topology; and the impact of optical misalignment. Finally, it is necessary to consider the way that these different aspects interact with each other. In the following sections, we will discuss these issues in detail.

III. OE-VLSI Chip Design

In the following sections, we describe OE-VLSI chips and associated packaging used in our system demonstrators. By example, we articulate appropriate design guidelines we have employed while constructing these devices. This section is organized as follows: in Section III-A we describe heterogeneous integration techniques we have used and in Section III-B we describe the two transceiver architectures and comment on the suitability for each. We also discuss in Section III-B the functionality we have embedded into our OE-VLSI circuits to

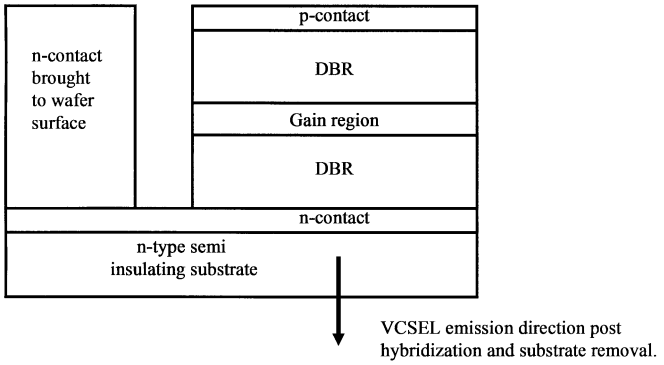

Fig. 6. Schematic of the VCSEL geometry indicating the emission direction post substrate removal and integration to the CMOS chip. Based on the partial transmittivity of the p-contact/DBR, the VCSELs could be probed prior to hybridization with the CMOS to verify functionality at the wafer level.

enhance the signal processing capability of this technology. Our overall approach in designing these OE-VLSI ASICs is consistent with the design philosophy that the optics are used for the interconnect and the electronics are used for transceiver functions and digital signal processing of data.

A. Heterogeneous Integration

We have used heterogeneous integration to achieve the high-density optical I/O required for achieving the FSOI interconnects realized to date. Key issues include using foundry CMOS for our ASICs. In doing so, it is important to assure compatibility of CMOS metals with the metals used in the flip-chip bump bonding. A related issue is the need to insure alignment of CMOS connection points with III–V device contact points. In the following sections, we summarize our flip-chip bump bonding based heterogeneous integration strategy. To achieve the OE-VLSI ASICs described throughout this paper, two-dimensional (2-D) arrays of VCSELs and photodetectors (PDs) were fabricated on separate substrates and subsequently integrated onto the silicon CMOS die. In order to support the compact high-density microoptical interconnects described above, the VCSELs and PDs were interleaved. We describe in this section: the design and target operating properties of the VCSELs and PDs, the OE device layout geometries, and heterogeneous integration techniques including flip-chip bonding and substrate removal of the interdigitated OE devices.

1) VCSEL and PD Design and Specifications: The VCSELs used in our designs were designed to operate at 850 nm with threshold currents of 1.0–4.5 mA and slope efficiencies of 0.25–0.35 mWs/mAs. The devices were also designed to be backside-emitting because of the desire to flip-chip bond them to CMOS driver circuits as described in the following. This necessitated removal of the GaAs substrate to minimize absorption of light. To achieve these objectives, VCSELs were fabricated with both the n-contact and the p-contact located on the top surface of the wafer to facilitate electrical contact to the CMOS circuits. Fig. 6 shows a schematic of the VCSEL geometry indicating emission direction after substrate removal and integration to the CMOS. The p-contact was formed above the top distributed Bragg mirror (DBR) and the n-contact was brought to the substrate surface through mesa isolation and ion implantation.

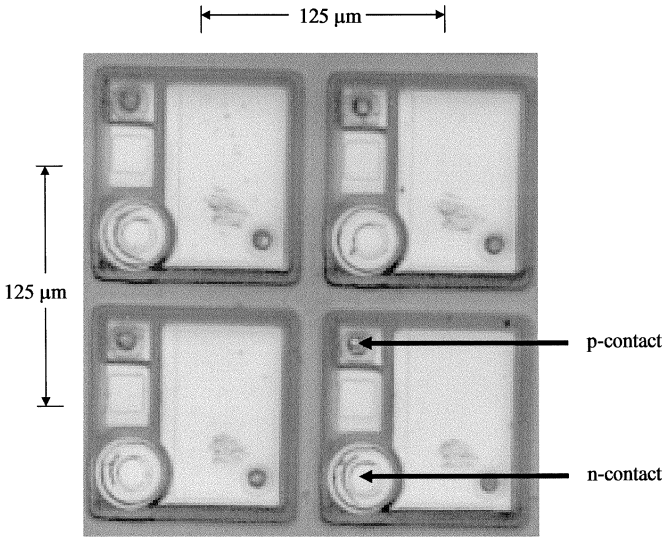

Fig. 7. Photomicrograph of four isolated VCSELs prior to flip-chip bonding. The n-contact and the p-contact are indicated. The scratch marks are a result of wafer probing prior to flip-chipping.

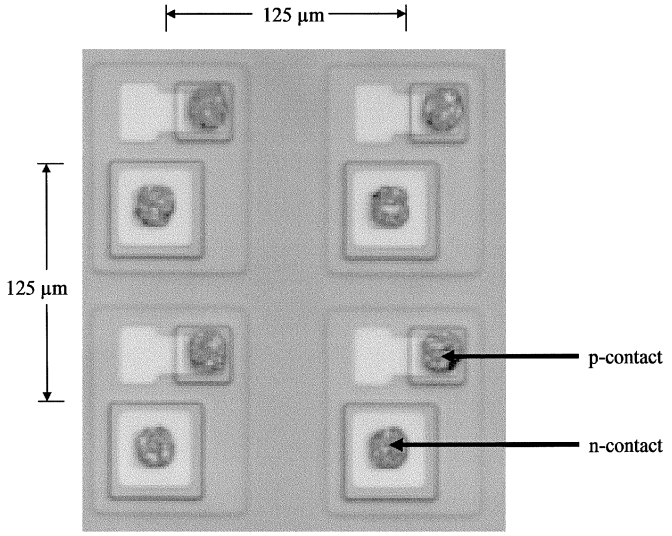

Fig. 8. Photomicrograph of four isolated PIN's prior to flip-chip bonding. The n-contact and the p-contact are indicated.

Fig. 7 is a photomicrograph of four isolated VCSELs prior to flip-chip bonding and substrate removal with the p-contacts and n-contacts indicated. The VCSELs in the photograph are on a 125- $\times$ 125-μm pitch. Once bonded to the CMOS as per the description given, the n-contact and DBR became the top (emitting) surface of the VCSEL.

The PDs were p-i-n structures designed to operate with a responsivity of 0.5 A/W. Fig. 8 is a photomicrograph of four isolated PDs on a 125- $\times$ 125-μm pitch prior to flip-chip bonding and substrate removal; the p-contacts and the n-contacts are indicated. The 2-D PD arrays were fabricated at the wafer level on a 125- $\times$ 125-μm pitch and were designed to be flip-chip bonded to the CMOS driver chip.

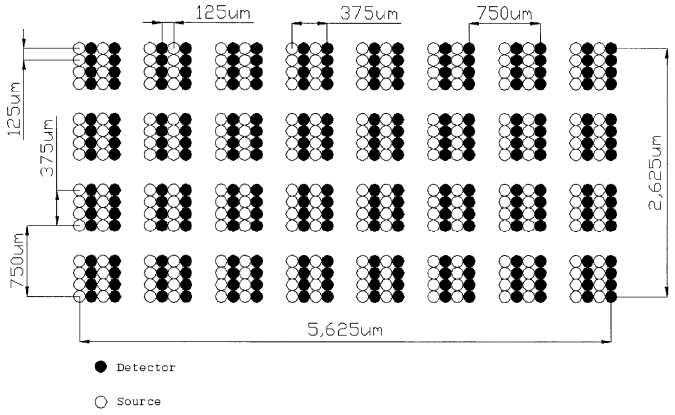

Fig. 9. Schematic of the VCSEL and PD placement requirements for an implemented free-space optical link.

2) OE Device Layout Geometry: In order to support a compact point-to-point optical interconnect system, as shown in Figs. 4 and 5, the VCSELs and PDs were interleaved and arranged in a clustered geometry. Fig. 9 is a schematic of the VCSEL and PD placement requirements. As is indicated, VCSELs and PDs were grouped together in clusters, and within each cluster, rows of VCSELs and PDs were interleaved. Specifically, a cluster consisted of eight VCSELs and eight PDs arranged in four rows. The pitch of the optoelectronic devices was 125 μm in both the horizontal and vertical directions; therefore, the VCSELs and PDs were on 125-μm horizontal by 250-μm vertical pitch. The complete 256-VCSEL and 256-PD array consisted of 32 clusters arranged in eight rows and four columns. The center-to-center spacing of clusters was 750 μm horizontally and 750 μm vertically. The CMOS driver ASIC was designed to accommodate this OE device pitch and placement.

3) Heterogeneous Integration and Substrate Removal: The VCSELs and PDs were integrated onto the CMOS driver using flip-chip bonding and substrate removal techniques. The VCSEL flip-chip contact area was 15 $\times$ 15 μm and the PD flip-chip contact area was 10 $\times$ 10 μm. The contact areas on the CMOS die for the VCSEL driver and PD receiver were identical to those on the optoelectronic devices.

Heterogeneous integration was accomplished by employing relatively conventional photolithographic processes to deposit and lift off contact metals and the wafers followed by a precision assembly process using a flip-chip bonding tool. In the photolithographic step, a photoresist polymer is first spun out on the wafer and printed with the contact metal pattern and then developed. Indium is then evaporated onto the wafer and the photoresist is lifted off, leaving metal on the contact pads. This process was used for the VCSEL and PD wafers and the CMOS dies. The individual OE device dies were separated by mechanical dicing into arrays containing the necessary 256 elements and then integrated onto a CMOS die using the precision alignment hybridization tool. The VCSEL die was first attached to the CMOS chip followed by dry etching to remove the substrate; integration of the PD die was accomplished next followed by substrate removal. The bonding of the indium metal contacts on the CMOS chip and on the OE devices was accomplished through a combination of force and controlled temperature. The process resulted in electrical isolation of individual OE devices and allowed the interleaving of the VCSEL and PD devices onto a single CMOS die. Although individual dies were used to assemble this generation of OE-VLSI chips, migration of the process to the wafer level is relatively straightforward.

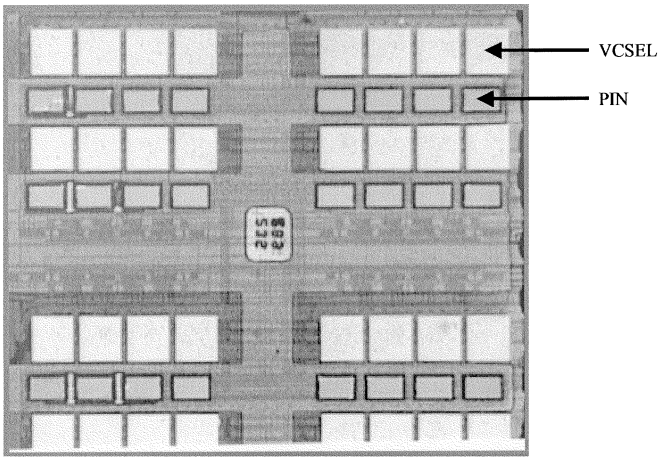

Fig. 10. Four clusters after heterogeneous integration and substrate removal. The VCSEL device was 100 $\times$ 100 μm and had a 10 μm in diameter active region; the PIN was 100 $\times$ 50 μm and had a 50 $\times$ 50 μm active region.

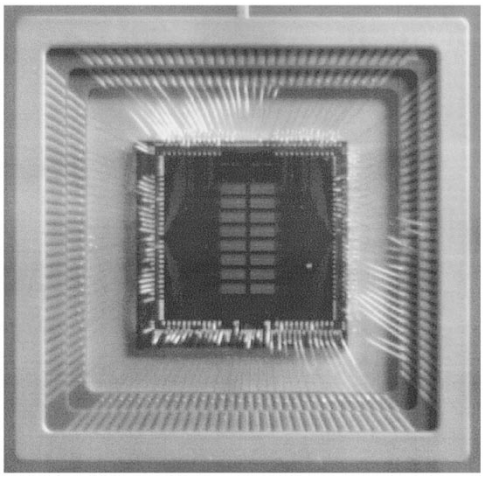

Fig. 11. Complete OE-VLSI chip. The rectangular section located in the middle of the die is the VCSEL and PD array. The 1 $\times$ 1 cm, 0.35 μm OE-VLSI ASIC was packaged in a 256-pin PGA.

Fig. 10 shows a photomicrograph of four clusters after heterogeneous integration and substrate removal, each cluster consisting of eight VCSELs and eight PDs hybridized to the underlying CMOS chip. Fig. 11 is a photograph of the complete OE-VLSI chip after VCSEL and PD integratation. Fig. 12(a) shows a group of four clusters with 32 VCSELs biased below the threshold, and Fig. 12(b) shows the entire VCSEL array biased above the threshold. Using continuous wave measurements, the VCSEL yield after heterogeneous integration was $>$98%. In the following two sections, we describe the transceiver circuits and the CMOS chip architecture which were implemented in this interconnect system.

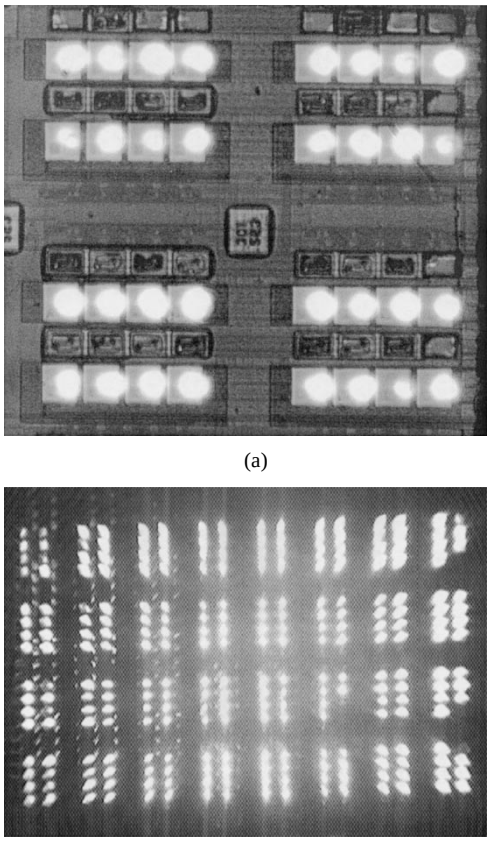

(a)

(b)

Fig. 12. (a) Four clusters with 32 VCSEL biased below threshold. (b) Entire VCSEL array biased above threshold. The astigmatism is caused by the optical system used to image the 3×6 mm array simultaneously.

We have found the process to be very effective in realizing the OE-VLSI ASICS used in our systems to date. As suggested, the process is relatively straightforward and can be extended to wafer level integration of CMOS devices with wafers of optoelectronic devices. The process lends itself to high levels of integration thus permitting large optical I/O counts. In a subsequent chip, we successfully integrated 1080 VCSELs and PDs on a 15×7.5 mm CMOS substrate [7], [8]. In this case, the optoelectronic devices were not interleaved but instead were integrated as blocks for transmitter and receivers.

B. Transceiver Architectures

There are numerous possibilities with respect to transceiver design for OE-VLSI technology. We have employed both single ended and differential designs. In the following section, we describe the two topologies including the merits of each.

In [5], the main objective of the transceiver circuit design was to provide enough flexibility to allow for the successful simultaneous operation of large numbers of transmitters and receivers. Although simulation results described in the following indicate high-speed operation was achievable, high data rate operation was not a principal design objective. It was expected that the VCSEL, PD, and CMOS characteristics would vary over a large device array; thus, the transceiver designs had to allow for statistical variations in device parameters and had to avoid dependence on parameters specific to the silicon and the OE process. The transceiver circuits were designed to keep their inherent switching noise generation at a practical minimum, as well as to be immune to the expected presence of the substantial amounts of aggregate switching noise generated from a large array of mixed analog and digital circuits.

Given these design objectives, the design of the laser driver was based on current-steering. Specifically, a VCSEL was dc-biased with a current IBIAS to a point above the VCSEL threshold current. Modulation current was provided by the current source IMOD and was steered through either the VCSEL or through an electrical dummy diode load, which was implemented as a diode-connected PMOS transistor. Current steering was achieved with switching transistors and complementary rail-to-rail digital CMOS inputs Vin and $Vinb$. The polarity of the laser driver circuit was noninverting; thus, when the inputs were logically low (Vin low and $Vinb$ high), the VCSEL was biased at IBIAS + IMOD and, therefore, produced a logically high-output power. When the inputs were reversed (Vin high and $Vinb$ low), the VCSEL was biased with only IBIAS and produced a logically low output power.

The current-steering nature of the laser driver allowed the total current drawn from the power supply to remain nominally constant at IBIAS + IMOD whether the VCSEL was in a high- or low-output power state. Power supply current transients could not be completely eliminated due to the mismatch in electrical parameters of the dummy load D1 and the VCSEL, but the approach allowed current transients (dI/dt noise) to be kept to a small fraction of the IMOD. The range of currents settable for IMOD and IBIAS was approximately 6 and 12 mA, respectively. The nominal voltage supply was 4.8 V. The power dissipation per laser driver circuit depended on the magnitudes of IBIAS and IMOD, and was estimated to be 86.4 mW in the worst case. An individual transmitter circuit was successfully simulated under worst-case (i.e., largest magnitude) conditions for IBIAS and IMOD at data rates in excess of 1 Gb/s.

The receiver circuit used in [5] were optically and electrically single-ended and was based on a common source transimpedance amplifier (TIA) front end. An offset-control stage was included to compensate for both the dc-coupled nature of working with CMOS amplifier stages and the decoupled nature of the optical input. This allowed properties of the receiver such as sensitivity (preamplifier feedback resistance) and the accommodation of various average optical power levels (offset control) to be dealt with independently, providing greater operational flexibility. The final stage of the receiver consisted of a Schmitt trigger that served as a final gain stage for decision-making and provided some hysteresis in its transfer function to help reduce the effects of power-supply switching noise in an array environment. Via the heterogeneous process described previously, each driver circuit was integrated with a VCSEL and each receiver was integrated with a PD; this resulted in a 2-D array of 256 transmitters and receivers.

In addition to the previous chip, we have designed, fabricated, and tested an OE-VLSI ASIC that employed a fully optically and electrically differential architecture. The details of this chip are presented in [7] and [8], and will only be summarized here. Transmitters and receivers were implemented as fully differential circuits, both optically and electrically. Using optically differential signaling allowed the receiver to determine a decision

threshold based on the optical input signals on a per-receiver basis. A fully differential electrical architecture allowed for reduced switching noise generation on the power supplies and enhanced rejection of common-mode noise. The transmitters and receivers were designed for operation at a data rate of 250 Mb/s. Additional circuit elements were added to the transmitter and receiver circuits to allow for circuit testability prior to optoelectronic device (OED) integration. These circuit elements were placed in parallel with the normal locations of OEDs, and could be specifically enabled by test inputs. To further enhance testability and operability, configurable parameters such as the magnitude of the modulation and bias current magnitudes and the magnitude of the feedback resistance the receivers could be set using digital inputs.

Through the design, fabrication, and testing of these two large-scale OE-VLSI ASICs and numerous test chips in which different receiver and transmitter topologies were employed, we draw upon analytical, experimental, and simulation-based work to arrive at a number of core conclusions. First, there needs to be some method of controlling the key set point parameters of the laser drivers and receivers across a subset of the entire array. This leads to higher operational yield. Secondly, the use of an optically and electrically differential architecture for the receiver and transmitter designs is optimal for OE-VLSI applications. Through our work, we have found that an optically and electrically differential architectures facilitate or optimizes the implementation of several critical aspects of OE-VLSI ASIC design, including:

1) design for testability (DFT) concepts and implementation for receiver and transmitter circuits;
2) the receiver and transmitter circuit generation of, and immunity to, switching noise on the voltage supply and ground rails and through the substrate;
3) the improvement of operational yield (percentage of functional circuits in a group of circuits that meet bit-error and data rate targets) in common bias and control receiver groups;
4) the reduction of intrachannel receiver skew in parallel digital synchronous OE-VLSI applications and the reduction of individual receiver latency.

These results are detailed in [7], [8], and [28].

IV. Optical Interconnect Design Issues

As described in the previous section, improvements in heterogeneous integration processes has resulted in the availability of large arrays of VCSELs and detectors on CMOS. The design of optics that can make efficient use of these devices is therefore critical. In the following sections, we attempt to address a number of important questions regarding the design of optical interconnects. First, it is important to determine the most suitable interconnect topology, which will transmit the optical signals over the required distance while also maximizing the interconnection density and also tolerance to misalignment. In Section IV-A, we briefly discuss the main classes of optical interconnect topologies that have been introduced, and describe their benefits and drawbacks. In Section IV-B, we will show that it is possible to obtain a quasi-analytical model for the density that is achievable for the clustered interconnect topology, which is one of the most popular classes of optical interconnects. With this model it is possible to directly obtain the relationship between interconnection distance and the spatial density that can be achieved. In Section IV-C, we will again compare interconnect topologies, and will introduce a model for misalignment tolerance. We will use this to draw conclusions as to which topologies provide the best misalignment tolerance and what the general characteristics of a misalignment tolerant system are. We will introduce the concept of the alignment product of an optical system and will show that this remains constant at all stages of the system. Finally, in Section IV-D, we will investigate the way in which misalignments of many individual components impact the performance of the system as a whole.

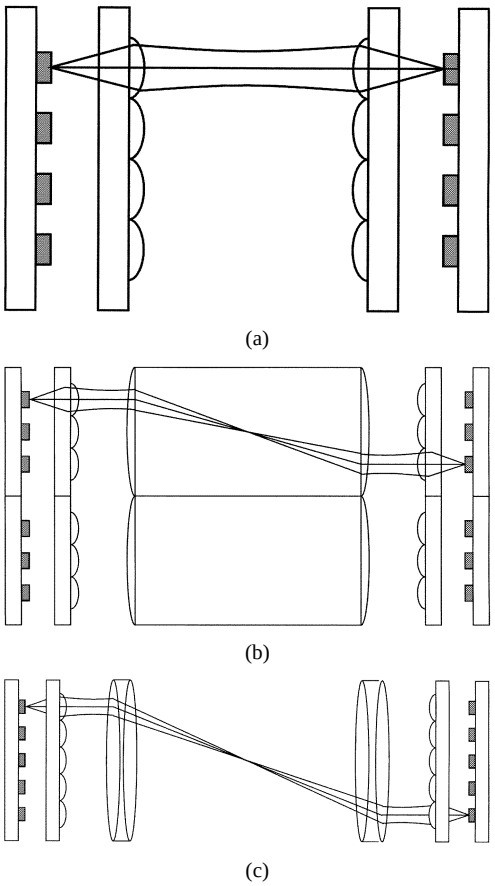

Fig. 13. Optical interconnect topologies. (a) Microchannel relay. (b) Clustered interconnect. (c) Macrolens.

A. Free-Space Interconnect Topologies

Several previous studies have dealt with the choice of optical interconnect topology [29]–[31]. The three most widely employed optical interconnect technologies are the microchannel relay [32] the clustered interconnect [29], [33] and a conventional bulk or macrolens solution [21]. These are shown Fig. 13. All of these schemes have different advantages and disadvantages. The microchannel relay [Fig. 13(a)], in which each optical channel is relayed by a single series of microlenses is limited by diffraction. As the interconnect distance (or optical throw) increases, it is, therefore, necessary to increase

the microlens aperture, which limits the device density, or add additional relay stages, which makes alignment and assembly more difficult. However, the design scales well with increase device array field size as additional channels can be added without changing the optical design. Microchannel relays typically operate in the maximum lens to waist configuration, in order to maximize the interconnect distance for a given lens focal length. Fig. 13(c) shows the macrolens design, which is also simple. All of the channels are share a common optical aperture. This design scales well with increasing optical throw distances, and can make use of multielement lens designs in order to improve performance. However, the design does not scale well as field size increases as this directly impacts the aberration performance of the lens. Furthermore, the incorporation of additional lens elements will increase the cost, although molded aspheric lenses could represent a cheaper alternative. The clustered interconnect (or minilens) design, shown in Fig. 13(b), is intermediate to these two designs. The channels are arranged into clusters, each of which is relayed through a single optical aperture. When VCSELs are used, a microlens array is also employed to collimate the VCSEL beams and to collect the relayed beams onto the detector array. Clustered systems are usually telecentric. At its extremes, a clustered interconnect could be said to subsume both the microlens relay (in which there is one channel per cluster) and the macro-lens approach (in which there is one cluster per array). Therefore, it is necessary to determine how the partitioning of the optical field impacts the performance of the system. In the following section, we will introduce an optimal technique to partition the optical field into clusters. However, we will show in Section IV-C that even more improvements can be made by introducing additional elements.

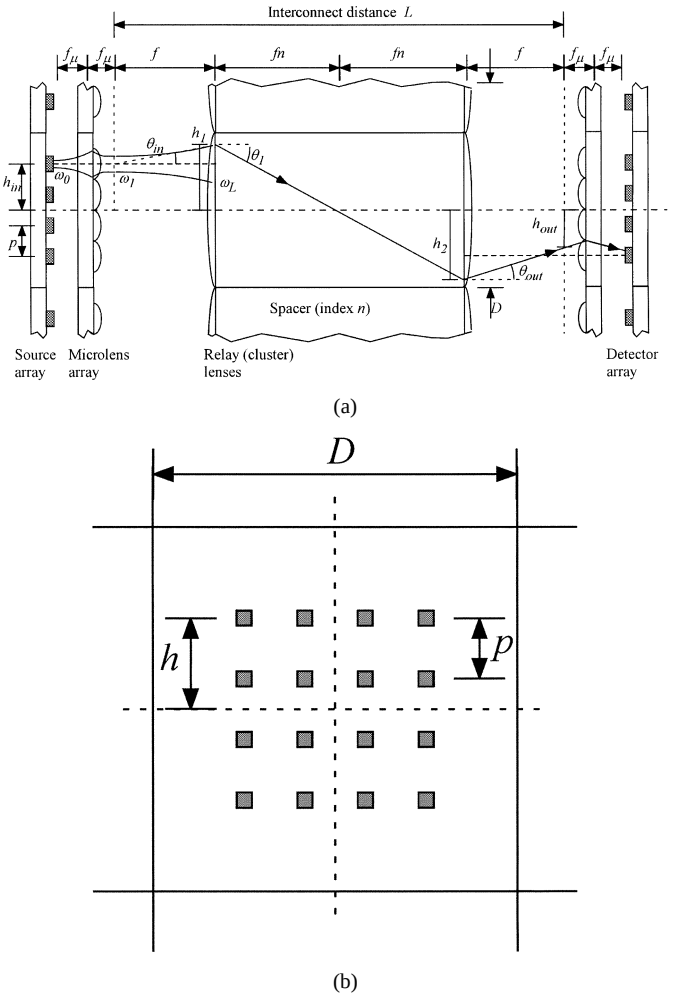

Fig. 14. Optical interconnect parameters. (a) Array layout. (b) Optical system.

B. Analytical Model for Clustered Optical Interconnects

There are three main factors which influence the performance and scalability of the clustered free-space optical interconnect. These are diffraction (which places an upper limit on lens spacing), geometric aberrations (which places an upper limit on the field size), and the speed of the relay lens (which will also limit field size for a given focal length) [15], [16]. Fig. 14 represents the important parameters of a clustered interconnect. The optical sources are arranged, in a square $N \times N$ array of pitch p. The distance from the center of the array to the outermost line of sources is h (i.e., this is the distance to the edge of the optical field of the cluster lens, measured along a vertex). The light from the sources is collimated by a microlens which has an aperture equal to the pitch of the channels and a focal length f_μ which for a given pitch is determined by the divergence of the VCSELs and the degree of clipping that can be tolerated). The clustered relay lenses are confocal with the microlenses and are assumed to have square apertures (side length D) and focal length f. They are attached at either end of a glass block, of index n, and which has a total length of $2fn$ (in order to maintain telecentricity). The optical system is symmetrical in that the microlenses at the detector end are identically located relative to the detectors. This ensures that the system is bidirectional. The optical throw L (or interconnect length) is defined as the distance between the focal planes of the clustered relay lenses [i.e., $L = 2f(n+1)$]. It should be noted that the interconnect distance can be increased by introducing additional glass elements between the cluster lenses and the microlenses (this is the case that for the system shown in Fig. 5). We can now determine the maximum field size h for a given cluster lens size D, optical throw L and a set of source properties. From this, we can calculate the number of channels that is supported by the interconnect and hence the channel density (defined as N^2/D^2 since the relay lenses are assumed to have a 100% fill factor).

We have determined the maximum off-axis source position h_{in} as limited by the speed of the relay lens and the aberrations of the relay system. We treat the system by modeling the light as Gaussian beams at the input and output side and via ray-tracing in the relay block. It can be shown that for a ray that enters the relay block with angle θ_{in}, originating at height h_{in}, the third-order expansion of the ray intercept error ε at the output is given by

$$\varepsilon = \frac{3h_{\text{in}}^2 \sin\theta_{\text{in}}}{f\,n^2} \tag{1}$$

Here the angle θ_{in} is the angle of the ray that corresponds to the 99% power asymptote of the incident multimode Gaussian beam after collimation by the microlens array. This assumes that

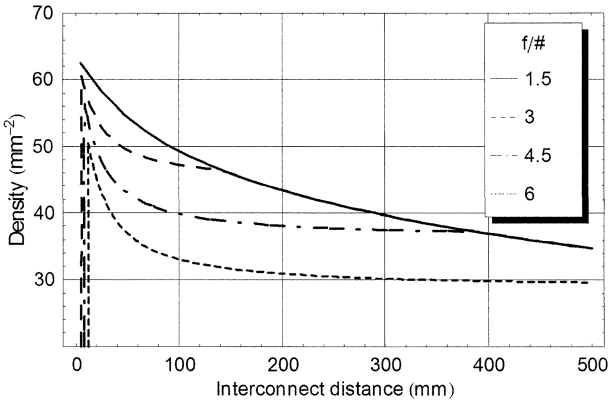

Fig. 15. Channel density as a function of interconnect distance.

the diffractive lenses are thin elements with a quadratic phase profile [16]. From this, we can obtain an approximate expression for the increased clipping loss γ that will be undergone at the output microlens [15]. For a given degree of tolerable clipping loss, we can obtain an allowable aberration $\varepsilon_{\max}$. This is given by

$$\varepsilon_{\max} = p\left(\frac{3}{4}\sqrt{\frac{2}{-\log(\gamma)}} - \frac{1}{2}\right) \tag{2}$$

from which we can obtain an expression for the maximum possible source height, as limited by aberrations. This is given by

$$h_{ab} = n\sqrt{\frac{f\,\varepsilon_{\max}}{6\sin\theta_{\text{in}}}}. \tag{3}$$

Note that an additional factor of $\sqrt{2}$ has been incorporated since h measures distances along the axis rather than along the diagonal (which will have the most aberrated beams).

The second limit to the source height is the aperture of the relay lens. If we assume that the $f/\#$ of the lens is defined as the ratio between the focal length and the diagonal aperture size, and that a 1% clipping condition again obtains, the maximum source height in this case is given by

$$h_{\text{ap}} = f\left(\frac{1}{2\sqrt{2}\,f/\#} - \tan(\theta_{\text{in}})\right). \tag{4}$$

The maximum source height $h_{\max}$ will, therefore, be determined by the minimum of (3) and (4). Therefore, the maximum number of channels (in one dimension) is given by

$$N = 2\,\frac{h_{\max}}{p} + 1. \tag{5}$$

Finally, we need to determine the area of the relay lens, in order to calculate the channel density. The relay lens aperture is given by

$$D = 2\left[h_{\max} + f\tan(\theta_{\text{in}})\right]. \tag{6}$$

We can now calculate the channel density ρ from

$$\rho = \frac{N^2}{D^2}. \tag{7}$$

Fig. 15 shows the calculated density as a function of interconnection distance for an interconnect which uses 850-nm multimode VCSELs with $M^2 = 2$ and $\omega_0 = 3\ \mu$m. The array pitch is 125 μm and the loss tolerance γ is 5%. It can be seen that at very short distances and for fast lenses a density of almost 64 channels/mm^2 can be achieved (which is the limit for devices on a 125 μm pitch). As the distance increases however, the achievable density decreases. Slow lenses offer a lower density than fast lenses at short distances, but eventually the performance is dominated by aberrations (as occurs at a distance of 120 mm for $f/3$ lenses) and so the relay lens speed is no longer an advantage.

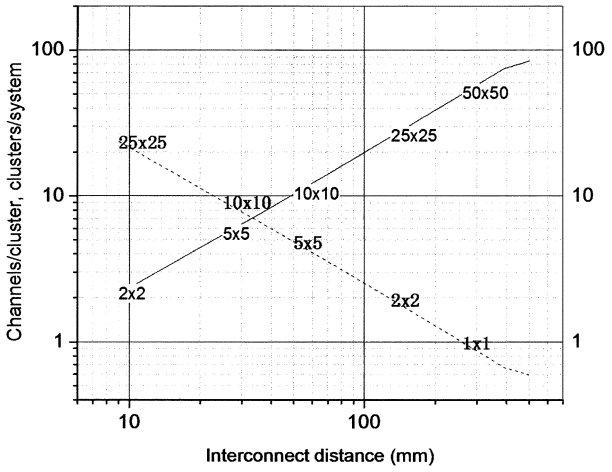

(a)

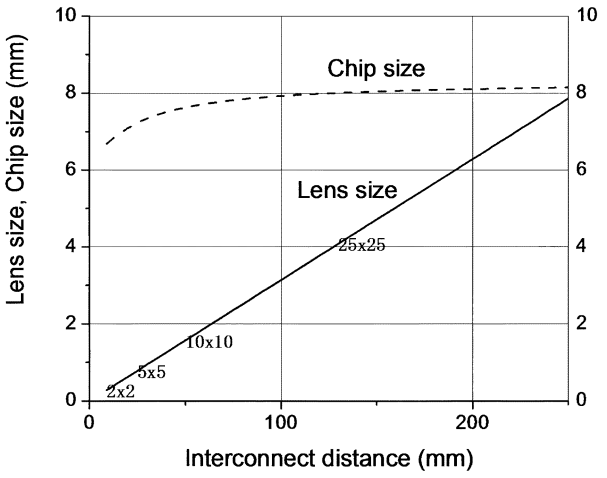

(b)

Fig. 16. System design example for a 50 × 50 $f/4.5$ system on a 125-μm pitch. (a) Channels/cluster (solid line) and clusters/system (dashed line). (b) Lens size (solid line) and chip size (dashed line).

We can also use this model to calculate the impact of VCSEL divergence on performance. It has been found that for devices on a 125-μm pitch at distances above 100 mm a single mode VCSEL (with $\omega_0 = 3\ \mu$m, $M^2 = 1$, $1/e^2$ half-beam divergence $= 5.2°$) allows a 50% increase in channel density when compared to a multimode VCSEL which has twice the divergence (i.e., $\omega_0 = 3\ \mu$m, $M^2 = 2$ and $1/e^2$ half-beam divergence $= 10.4°$).

This model can be used to design an interconnect system that will maximize density for a given interconnect distance. As an example, Fig. 16 shows the variation in the number of channels per cluster and clusters per system for an interconnect that requires 50 × 50 channels, with multimode ($M^2 = 2$) VCSELs

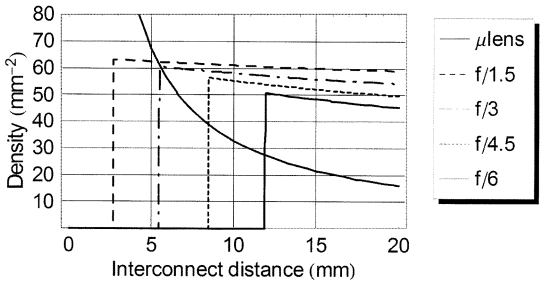

Fig. 17. Comparison of interconnect density for microlens relay (μlens) and clustered relay, for a range of $f/\#$. This assumes the clustered relay has sources with a pitch of 125 μm and that the VCSEL M parameter is two for all systems.

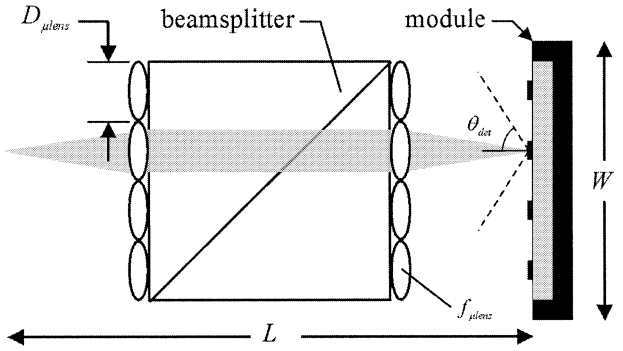

Fig. 18. Free-space interconnect example.

on a 125-μm pitch with $f/4.5$ relay lenses, as a function of interconnect distance. It can be seen that for a short distance (10 mm) there will be a small number of channels per cluster (2×2) and, hence, many clusters per system (25×25). As the interconnect distance increases it eventually becomes most efficient to use just one relay lens (i.e., a macrolens) with all 50×50 channels traversing it. This holds for distances above 250 mm. The relay lens size increases almost linearly with distance and the chip size (determined as the product of the relay lens size and the number of clusters per system) remains almost constant at 8 mm. Obviously, there exist intermediate distances where it would not be possible to have an integer number of clusters per system. This indicates that at these distances, it is not possible to obtain the maximum density. However, by selecting a different source array pitch it may be possible to improve the density which can be obtained.

We can also make a comparison between the clustered relay system and a microchannel interconnect. As we have seen in Fig. 16, as the interconnection distance decreases the number of channels per cluster also decreases until there is eventually only one channel per cluster. However, the clustered interconnect remains telecentric, and so does not provide the same interconnection distance (as a function of relay lens focal length) as can be achieved by a microlens relay in the maximum lens-to-waist configuration [32]. This effectively means that at very short interconnection distances the microchannel interconnect should offer better densities. This is shown in Fig. 17 for a multimode ($M^2 = 2$) system for which the clustered system has a device pitch of 125 μm (the pitch of the microchannel relay is a function of interconnection distance [34]). As can be seen, for distances below 5–15 mm (as a function of relay lens $f/\#$) the microchannel relay offers a better density. However, for interchip distances the clustered interconnect offers much better density.

C. Optical Invariant and Misalignment Tolerance

The clustered interconnects described previously represent one possible solution to the design of optical interconnects. However, it is interesting to ask whether these provide the best possible tolerance to misalignment. Although many researchers have calculated the misalignment tolerance of individual systems there has been little systematic investigation into which classes of design provide the greatest tolerances and of general principles for designs which have a high tolerance to misalignment. Research in this direction has recently been initiated by Neilson [35]; in this work, the author calculated the coupling efficiency between two misaligned FSOI components by considering the overlap integral between their respective optical modes. Neilson's approach focuses on the fundamental properties of Gaussian beams and his conclusions are independent of the lens configuration used. Here, we will addresses the issue of alignment from the opposite perspective: by investigating how certain lens configurations can lead to FSOI systems that are inherently more tolerant to misalignment.

The first step toward such an analysis [13], [36] consists of the definition of a suitable misalignment metric. Typically this is a loss measurement, due to beam clipping or misalignment of a beam on a detector (for example it may be decided that a 5% loss is tolerable). Having done this, we can then define a figure of merit (FOM) for alignment that should measure the ease of aligning an optical interconnect. Components within FSOIs can be misaligned within all six degrees of freedom. However, in most practical cases only two degrees are of interest. These are lateral misalignment in x (represented as Δx) and tilt misalignment in x (represented as $\Delta\theta_x$). The others can be neglected for the following reasons: 1) most optical interconnects are symmetrical so that $\Delta x = \Delta y$ and $\Delta\theta_x = \Delta\theta_y$; 2) rotational misalignment ($\Delta\theta_z$) has the same impact locally as a lateral misalignment; and 3) longitudinal misalignment tolerance (Δz) is typically much greater than lateral tolerance. Therefore, the alignment FOM is defined as the product of lateral and tilt misalignment tolerance, i.e.,

$$\text{Alignment FOM} = \Delta x \cdot \Delta\theta \tag{8}$$

where the subscript on $\Delta\theta$ has been dropped for simplicity. In general, the larger the figure of merit, or alignment product, the easier the system is to align.

In the following section, we calculate the alignment product for five different optical relay systems which are designed to perform an identical task, which is to relay light through a cube beam-splitter and focus onto a detector (see Fig. 18). The parameters of the optical system are given in Table I. The clipping ratio k_l is a measure of the degree to which beams are clipped at the lens apertures and is defined as the ratio

$$k_l = \frac{D_{\mu\text{lens}}}{2\omega_{\mu lens}} \tag{9}$$

where $D_{\mu\text{lens}}$ is the effective lens aperture and $\omega_{\mu\text{lens}}$ is the $1/e^2$ Gaussian beam radius [32]. Belland and Crenn [37] have shown

TABLE I
DEFINITION OF PARAMETERS

Parameter	Definition
λ	Operating wavelength
N	Linear dimension of the array
W	Linear dimension of the chip
L	Interconnect length
m	Cluster size (where appropriate)
d	Linear dimension of the detector
ω_d	Beam waist at detector plane
k_l	Minimum clipping ratio at a lens

that when the clipping ratio is greater than 2.12 the clipping losses are less than 0.1% and diffraction effects do not significantly modify the beam propagation characteristics.

For each system, it is assumed that the system is fabricated in a pair of modules, one of which contains some of the relay optics and the other of which contains a detector array (and possibly some optics also). Within each module the components are assumed to be perfectly aligned, and so the alignment product must be calculated for module to module alignment. The five designs are shown within Table II and are:

1) microlens relay (i.e., the lenses are perfectly aligned to each other but not to the detector array);
2) lens relay in which the focusing microlenses are integrated with the detector module (and thus assumed to be perfectly aligned to it), but where the first set of microlenses are not aligned to the second;
3) minilens clustered interconnect;
4) microlens telescope (in which the final microlens is a short focal length collecting lens);
5) microlens relay with a field lens (where the final collecting lens is at the focal plane of the second relay lens).

In each diagram, the solid bars connected to the detectors show which components are assumed to be part of the detector module.

Table II also shows closed forms for the lateral and tilt tolerances for these five designs. The maximum array size in each case is limited by diffraction beam clipping and is based on a loss tolerance of 50%. Details of the reasoning which underlies the calculations is contained in references [13], [36]. It can be seen that even configurations which are optically identical (such as design 1 and design 2) can have a very different tolerance distribution, as a function of where the system is broken. Thus, design 2 has much better lateral tolerance than design 1 because as the second microlens remains aligned with the detector as the first lens shifts laterally, so that the beam remains focused on the detector. This is traded for decreased tilt tolerance, however. The table also shows the results of a calculation example based on a 10-mm square chip with an interconnection distance of 25 mm. The individual lateral and tilt tolerances are given, in addition to the misalignment product. It can be seen that the minilens clustered system (design 3) and the field microlens system (design number 5) provide the highest misalignment product. Design 4 (the microlens telescope) provides an identical misalignment product to Design 1 but the tolerances are more equally distributed between Δx and $\Delta\theta$. Although the clustered system (Design 3) provides almost as much total tolerance as the field lens system (design 5) the tilt tolerance is extremely tight (only 14 arcmin). This low tilt tolerance has previously been recognized as a disadvantage of clustered systems [6], [3]. Design 5, therefore, provides the best partitioning of lateral and tilt tolerances. This advantage is maintained when the interconnect distance is increased to 50 mm. In this case, Designs 1 and 2 would not perform at all and the clustered system (Design 3) has a misalignment product of only 315-μm arcmin (due to a very low tilt tolerance of 3 arcmin). Design 4 has a misalignment product of 2750-μm arcmin and Design 5 has a misalignment product of 4683-μm arcmin (divided between 90 μm of lateral tolerance and 52 arcmin of tilt tolerance) [13], [36]. The excellent performance of the field lens system is to be expected. Several authors have previously proposed the use of field lenses to increase misalignment tolerances in free-space optical interconnects. More fundamentally, the field lens system is acting as a Gaussian relay in which the detector is imaged in the aperture of the second relay microlens. Neilson [38], [35] has shown that Gaussian relays (for which $f_{\mu\text{lens}} = \pi\omega_0^2/\lambda$) provide an optimum tradeoff between tilt and lateral misalignment tolerance. This is achieved by having slow beams on both sides of the interface at which misalignment occurs, which is the case in Design 5.

This line of inquiry can be further extended to consider the nature of the misalignment product of a module. Consider Fig. 19, which is a module that contains a single detector and a lens. The lens produces an image of the detector in space (which is essentially the entrance pupil of the system). It is assumed that the magnification of the system is such that the size of the detector image is smaller than the lens aperture. The alignment tolerance of the module is given by the product of Δx and $\Delta\theta$. The image of the detector is drawn using one oblique ray and one axial ray and has a total size of d'. The angle the oblique ray makes with the optical axis is θ' and θ in the image plane and object plane respectively. These angles correspond to the entrance and exit numerical apertures (NA) of the module. Within the laws of geometrical optics, it can be shown that the following relationship is true at all points within an optical system [39]:

$$\frac{d}{2}\theta = \frac{d'}{2}\theta'. \tag{10}$$

This is referred to as the optical invariant. We can now relate this to the misalignment product. If we assume a 50% loss tolerance, then we can argue that the lateral tolerance Δx is equal to $d'/2$ and that the angular tolerance $\Delta\theta$ is equal to the NA of the module θ'. Therefore, from this and (10), we can write

$$\Delta x \Delta\theta \propto \frac{d}{2}\theta. \tag{11}$$

Therefore, the alignment product is directly related to the optical invariant, which is a product of the detector size and the NA of receiving optics. This is an invariant of the system, regardless of the complexity of a module. We can, therefore, conclude that we can maximize the alignment product firstly by increasing d, which means using large detectors (up to the limit imposed by capacitance and bit rate) and secondly by maximizing the NA of the receiving optics. This provides a quick and easy way to determine the alignment tolerance of an entire interconnect

TABLE II
SUMMARY OF RESULTS: EXPRESSIONS FOR LATERAL TOLERANCE, TILT TOLERANCE, MAXIMUM ARRAY SIZE, AND OPTICAL INVARIANT

	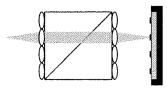 Design 1 Module with no optics	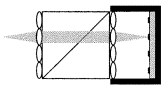 Design 2 Microchannel design	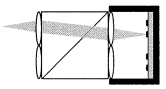 Design 3 Clustering configuration	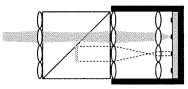 Design 4 Microchannel telescope	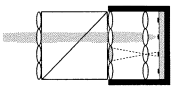 Design 5 Field lens configuration
Lateral tolerance (Δx)	$\frac{d}{2}$	$\frac{W}{2N}\left(1-\frac{N}{N^*}\right)$	$\frac{W}{2N}\left(m-\frac{N}{N^*}\right)-(m-1)d$	*MINIMUM:* $\frac{Md}{2}, \frac{W}{2N}-k_I\sqrt{\frac{\lambda L}{2\pi}}$	*MINIMUM:* $\frac{Md}{2}, \frac{W}{2N}-k_I\sqrt{\frac{\lambda L}{2\pi}}$
Tilt tolerance ($\Delta\theta$)	$\sqrt{\frac{2d}{W}}$	$\frac{2d}{L}$	$\frac{2d}{L}$	*MINIMUM:* $\frac{M2d}{L}, \frac{2W/NL-k_I\sqrt{4\lambda/\pi L}}{1+M}$	$\frac{2W}{NL}-k_I\sqrt{\frac{4\lambda}{\pi L}}$
Max. array size (N_{max})	$\frac{2\pi\omega_d W}{\lambda L k_I}=N^*$	$\frac{2\pi\omega_d W}{\lambda L k_I}=N^*$	$\frac{mN^*}{[2d(m-1)N^*]/W+1}$	$\frac{W}{k_I\sqrt{2\lambda L/\pi}}$	$\frac{W}{k_I\sqrt{2\lambda L/\pi}}$
Optical invariant (I)	$\frac{d}{2}\theta_{det}$	$\frac{dW}{NL}=I^*$	$I^*\left(m-\frac{2d(m-1)N}{W}\right)$	$M\times I^*$	$M\times I^*$
Example W=10mm, λ=850nm L=25mm, d=50μm k_I=2.12, N=16, m=4	Δx=25 μm, $\Delta\theta$=344 arcmin, $\Delta x\Delta\theta$=8594 μm.arcmin	Δx=97 μm, $\Delta\theta$=14 arcmin, $\Delta x\Delta\theta$=1339 μm.arcmin	Δx=882 μm, $\Delta\theta$=14 arcmin, $\Delta x\Delta\theta$=12129 μm.arcmin	Δx=125 μm, $\Delta\theta$=69 arcmin, $\Delta x\Delta\theta$=8594 μm.arcmin	Δx=100 μm, $\Delta\theta$=124 arcmin, $\Delta x\Delta\theta$=12394 μm.arcmin

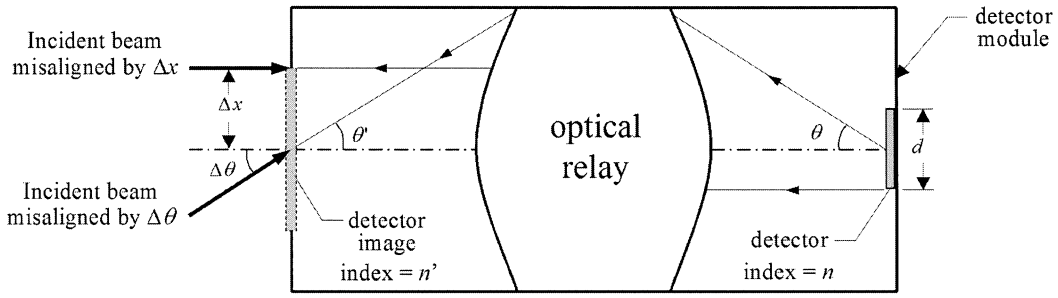

Fig. 19. Relationship between the optical invariant and the misalignment product.

system. One simply needs to determine the value of the optical invariant at the detector plane; if the invariant is small then the design is inherently difficult to align. With this in mind, we can now reassess the five designs in Table II. Designs 1 and 2 both have small optical invariants (and, hence, alignment products) due to the low NA of the microlenses (which are required to have long focal lengths for relaying and yet have small apertures). The clustered configuration (Design 3) increases the NA of the lenses by a factor of m and so the alignment product increases by the same degree. Designs 4 and 5 both have a large alignment product due to the fast microlenses that are placed in front of the detectors.

D. Stochastic Analysis of Misalignment Tolerance

The optical invariant approach described above represents a powerful way of understanding the misalignment tolerance of an optical interconnect. However, it does not necessarily tell us everything we need to know about the ease or difficulty of assembly a particular system when many individual components are randomly misaligned. The interconnect described in reference [6] contains two consecutive 4-f relays and requires the alignment of eight independent optical device planes. It is particularly important to determine the way in which misalignments of individual components stack-up to determine the performance of the system as a whole. In particular, we need to be able to answer the following question: How precisely must each individual component be aligned in order to obtain a high probability that the system as a whole will work? The starting point is a sensitivity analysis in which each component is misaligned in each degree of freedom in a ray-tracing or Gaussian beam simulation and the impact on throughput is calculated [40], [41]. The degree of misalignment that can be tolerated for a given throughput reduction is described as

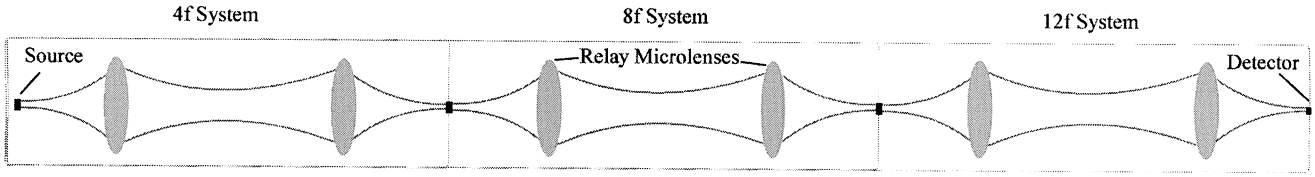

Fig. 20. Multielement optical relay.

TABLE III
COMPONENT TOLERANCES AS A FUNCTION OF SENSITIVITY

	Slow system				Fast system			
Parameters	10%	1%	0.1%	0.01%	10%	1%	0.1%	0.01%
Source (lateral)	±27μm	±20μm	±15μm	±10μm	±5.7μm	±4.25μm	±3.2μm	±2.25μm
Source (angular)	±1.7°	±1.15°	±0.75°	±0.375°	±7.2°	±4.8°	±3°	±1.8°
Relay lens (lateral)	±27μm	±20μm	±15μm	±10μm	±5.7μm	±4.25μm	±3.2μm	±2.25μm
Detector (lateral)	±27μm	±20μm	±15μm	±10μm	±5.7μm	±4.25μm	±3.2μm	±2.25μm

the sensitivity of the component. The selection of the tolerance metric is a critical part of the sensitivity analysis. For example, a lens in an optical system may be found to have a lateral misalignment tolerance of 5 μm at the 10% power loss level, indicating that a 5-μm misalignment will result in a 10% throughput reduction. Systems assembled to a relaxed metric might be easier and cheaper to assemble but might also provide lower performance in the assembled system due to accumulated losses. A more severe metric will mean that the tolerances are tighter, providing a lower loss but also increasing the fabrication cost. Once the individual tolerances have been calculated it is then necessary to determine the way in which individual component misalignments interact and stack up. There are a variety of ways by which this can be done. The most direct approach is to obtain an expression for throughput as a function of the positions of all of the individual components. For systems with more than a few components this rapidly becomes an intractable calculation. The simplest alternative is a worse-case misalignment tolerance calculation in which all components are assumed to be maximally misaligned. Another technique that is commonly used is a root-sum-of-square (RSS) analysis [42], [43] in which each component is assumed to misaligned such that it causes a given reduction in throughput and the total throughput reduction is then calculated as the root sum of squares of the individual values. RSS analysis is fast, but it is known that this method is not valid as error functions for optical systems are not additive and tolerances are not linearly independent. Finally, it is possible to obtain an accurate determination of misalignment tolerance by performing a Monte Carlo analysis [44], [14]. Many simulated systems are generated in which all components are assumed to be misaligned with some suitable probability distribution and the final throughput in each case is calculated. In this way, a probability distribution of throughputs can be obtained [45].

In order to investigate the points stated previously, two parallel free-space optical interconnects possessing identical properties but different f-numbers and with different numbers of relay blocks (i.e., with one, two, and three blocks corresponding to 4-, 8-, and 12-f systems) were selected for simulation purposes (see Fig. 20) [14]. The first system had an f-number of $f/16$ (note that this refers to the f/number of the Gaussian beams traveling in the interconnect) and was based on a telecentric relay system used in a FSOI [46]. The relay lenses were 8.5-mm focal length minilenses with a square aperture of 800 μm. The source was placed on-axis at the focal point of the lens and was assumed to emit a Gaussian beam possessing a waist radius of 13.1 μm. The detector was assumed to be square and 70 μm on the side. The relay lenses had an f-number of $f/10.6$ (which is necessarily faster than the f-number of the each individual beam). This design is referred to as the "slow" system. The second system had $f/3.38$ lenses and was a scaled version of the first, in that the focal length, minilens apertures and detector size are scaled to provide the same clipping ratio (the ratio of the beam diameter over the minilens or detector aperture) for both systems. In the second system, the focal length was 3.0 mm and the minilenses were 1341.5 μm on each side. The source was assumed to emit a Gaussian beam possessing a 2.75 μm waist radius (corresponding to the radius of the beam emitted from a singlemode fiber operating at 850 nm). The detector was assumed to be square and 14.65 μm on the side. This system is referred to as the "fast" system. Aberrations have no significant effect on the performance of either system.

A sensitivity analysis followed by a Monte-Carlo analysis was performed to assign positioning tolerances to the components of the optical systems and calculate the probable performance when assembling a system using components specified to those tolerances. Four tolerance sets were calculated with the help of four loss metrics: 10%, 1%, 0.1%, and 0.01% power loss relative to the maximum. These are shown in Table III. This table shows for example that for the slow system the throughput drops by 10% when the source module is misaligned by ± 27 μm. The tolerance limit as calculated in the sensitivity analysis is used to set the limits of the distributions

that control the misalignments of the components induced by fabrication and positioning errors. The shape of these distributions is highly coupled to the manufacturing process used to fabricate the component an it was decided to use the same statistical distributions as are employed in a commercial optics simulation package (see [14] for more details).

It was found during the course of the sensitivity analysis that the tolerances for each parameter did not vary significantly as the system complexity was increased from 4 to 8 to $12f$. This is true for all the parameters except the source tilt which display a 0.15° decrease for each $4f$ addition. This is to be expected as the minilenses are oversized with respect to the Gaussian beam passing through them which means that most of the clipping and power loss occurs at the detector and not at the aperture of each component. During the Monte-Carlo analysis, 25 000–50 000 samples were calculated in each run. In order to compare the output of the different Monte-Carlo runs, normalized cumulative histograms were generated. The distributions calculated by the Monte-Carlo simulations for the slow and fast systems are shown in Fig. 21 in the form of normalized cumulative distribution plots for each of the four tolerance sets and for 4-f (a) and 12-f (b) systems. Note that the dashed curves represent data for the slow system while the solid curves represent data for the fast system and that they are paired for each tolerance. In the case of the 4-f system the dashed curves are always topmost. Since these are cumulative histograms, this shows that more samples possess small throughput values for the slow system. The fast system suffers from slightly less severe stack-up effects. For example, in Fig. 21(a), it can be seen that there is a probability of 0.5 of obtaining a throughput of 0.5 or better when constructing a slow 4-f system (solid curves) with components having tolerances specified using the 10% tolerance set. This means that about half of the systems constructed will possess a throughput above 50% while the rest will have a throughput below 50%. In contrast, when a system is assembled from components specified using the 0.01% tolerance metric, there is a much smaller chance (close to zero) of obtaining a throughput of less than 50%. The probable performance is slightly better for a fast system than for a slow system. For example, there is a 0.55 probability that systems built using a 10% tolerance set will possess a throughput of 50% or more while there is a 0.95 probability that a system specified to 0.01% tolerances will possess a throughput of 90% or more. The curves are noticeably shifted upwards as the system length increases, indicating that the effect of errors accumulate as the system length is increased. These graphs also show that as system complexity increases (4–12-f) the probability of success decreases, due to the increased number of parameters. Thus, for slow systems with a single 4-f relay assembled to the 1% tolerances, 50% the samples have a throughput of at least 90%. However, for a triple (12-f) relay, this throughput has fallen to 60% for 50% of the samples.

These results were also compared with RSS. It was found that RSS was dramatically over-optimistic, particularly for large tolerance metrics (for example, for a 4-f system with a 1% metric, the mean system throughput was predicted by RSS to be 96%, whereas, the Monte Carlo simulation predicted a mean system throughput of 69%.

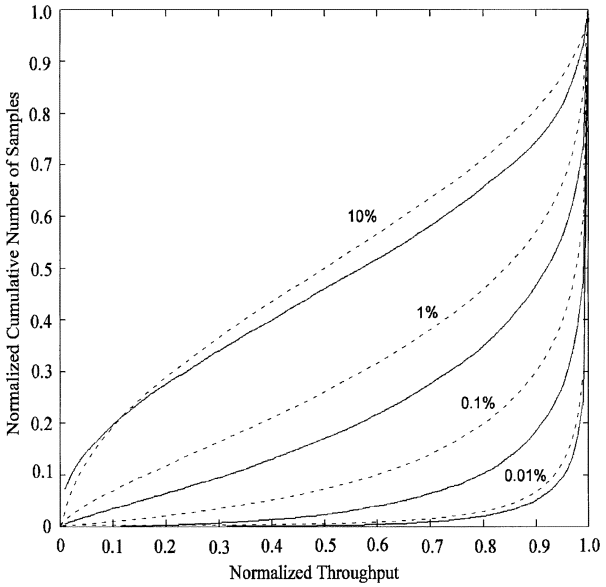

(a)

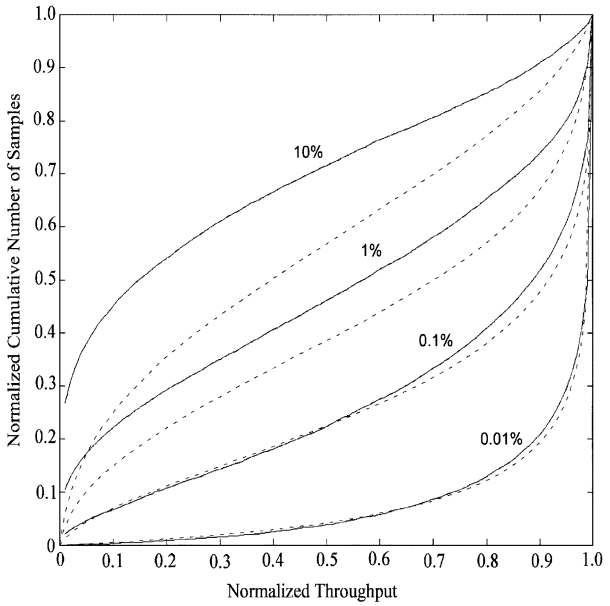

(b)

Fig. 21. Cumulative loss for slow (solid) and fast (dashed) (a) $4f$ and (b) $12f$ optical relays.

Several conclusions can be drawn. First, tolerance stackup effects are important and must be included as part of system design. Many systems are designed with individual component tolerance metrics of 10%. Fig. 21 shows that this will often lead to poor throughputs in real systems if passive assembly is used and that tighter tolerance (1%) must be applied. This study also shows that an RSS approach does not accurately predict loss as the losses are coupled in an optical relay system and that a Monte Carlo analysis must be used. Software tools for this purpose are being developed [47]. As could be expected, systems containing several consecutive relay blocks (and, thus, many components) have a lower probability of successful assembly than short systems (and will result in lower throughput). If a system does require many components this suggests that they should be assembled actively into a module which can then be

inserted passively. Finally, there is little difference in the tolerance stackup of slow and fast systems. Fast systems display a slightly less severe stackup but slow systems have looser individual tolerances.

This approach can be extended further. Once a Monte-Carlo analysis has been performed, it is then possible to perform a regression analysis in order to obtain a relationship between the system variables and the performance metric, as described in [48]. This can provide more insight into the sensitivity of the system in each degree of freedom and thus informs the system designer as to the areas in which the greatest care is required.

V. Discussion

In this paper, we have sought to introduce some of the design rules for free-space optical interconnects that we have discovered as a result of designing, building and testing interconnect. Although we have presented these rules within the context of individual system components (for example analog circuit design, optical design, etc.), it is obvious that there will be a high degree of interaction between them. Thus, when a clustered optical design is selected, in order to obtain good tolerance to misalignment with long optical throw and high density, the optoelectronic devices must also be laid out in a clustered fashion. The presence of gaps between clusters can then be used to advantage to route traces to the edge of the chip.

Another interesting area of interaction between the optical portions and the electronic portions of the system lies in the selection of the optical transmitters. At present, the most widely used VCSEL emitters for free-space interconnects are multimode devices. These provide high output powers and can deliver good modulation rates because they are driven far above threshold. However, they also have a much higher beam divergence than single mode VCSELs. As a result, the array density that can be achieved is not as large, as we have seen in Section IV-B. For a well-designed and assembled FSOI, transmission losses of only 3 dB should be achievable. As a result the $\sim$1 mW power that is emitted by a single mode VCSEL should be adequate to drive a typical receiver at data rates up to 10 Gb/s. Furthermore, because multimode VCSELs display spatial mode switching and random polarization switching as they are driven up and down, they may induce more noise and less predictable behavior. However, if the system is limited by the power available at the receiver, the higher power delivered by multimode VCSELs may result in an improved performance. In this situation, a full system simulation at the device level is required. Examples of suitable tools are provided by [47]. Such a simulation should also include thermal modeling, so that the impact of modifying the device pitch on the system as a whole can be properly understood.

In the area of optical system design, there are a number of areas that we have not had sufficient space to discuss. This includes the use of techniques such as eikonal analysis [49] which can be very powerful in developing designs that have low distortion and other aberrations. The telecentric systems that we have considered are inherently distortion-free, but more general designs will require more advanced analyses. Examples of designs of this sort include reflective designs [22] and planar optics [50], [51].

Although we have discussed techniques for the design of misalignment tolerant interconnects, and have also discussed the impact of misalignment stack-up, we have not discussed methods by which these systems can be assembled to the required tolerances. One of the biggest challenges in this area is to align device planes in all six degrees of freedom which are separated by several mm or tens of mm. A variety of techniques have been developed to solve these challenges, including the use of interferometric alignment techniques [52] and diffractive features that can provide alignment in all six degrees of freedom [11], [12]. In particular, we have recently described a technique which can provide precise diagnostic information on the degree of misalignment on each axis independently to within the tolerances required for FSOI, thus, opening the way to automated assembly. However, although these techniques are successful in producing small numbers of precisely aligned components, they are not efficient for volume production. However, they may be useful as a method of producing master components that can then be molded into a high-grade optical polymer, as discussed in [42] and [53].

There remain many other challenges that must be overcome before FSOIs can be adopted as a solution to the chip-to-chip communication bottleneck. In particular, issues such as thermal and vibrational stability have not been properly addressed as yet. However, for many of these issues it always possible to obtain an improved tolerance to misalignment by decreasing the density of the interconnect. Another solution is to use spatial redundancy to increase tolerance to misalignment. For example, the system described in [9] offers up to ± 1 mm and $\pm 1^\circ$ of tolerance. However, in both cases, the density of the interconnect is reduced and, thus, it may not provide the capacity advantages that make FSOI attractive.

VI. Conclusion

In this paper, we have attempted to codify a number of design rules for FSOIs at the chip-to-chip and board-to-board level. These have been derived partly from experience in constructing a number of different interconnect systems and partly from analytical and numerical analyses of the optics and optoelectronics which underlie these systems.

One of the most obvious conclusions is that heterogeneous integration of optoelectronic devices onto CMOS is a critical aspect of this approach. The high density 2-D optical I/O which are seeking to achieve necessitates the use of 2-D surface emitting device arrays which are only achievable through this approach. The close integration of optoelectronics with CMOS also opens a route to the incorporation of advanced signal processing functions at the interconnect layer, including forward error correction and clock acceleration. Yield can also be improved through the adoption of built-in self-test structures. Differential electrical and optical transceiver designs provide the best performance due to immunity to electrical and optical crosstalk through common mode rejection techniques. Although they entail the use of more space and resources the performance gains so obtained outweigh the costs. Finally, we have shown that the ability to control the bias currents and voltages of individual sections of large arrays greatly improves performance. By defining common control groups

(CCGs) whose dimensionality is governed by factors such as the number of channels and the optical I/O topology, we have found that high-operational yield can be achieved using this approach.

In the area of optical design, we have investigated scaling and partitioning laws for clustered interconnects. We have introduced an analytical model which includes diffraction and aberration effects and which allows designers to determine the optimum clustering configuration which will maximize device density for a given distance. We have also determined the distance at which microchannel interconnects are more suitable, and conclude that imaging (clustered) designs are superior for distances above about 15 mm. We have also investigated the problem of system partitioning modularization. The selection of interfaces that are used to define module boundaries has a critical impact on the misalignment tolerance of a system, and in particular in the way that it is partitioned between lateral and angular terms. Furthermore, we have shown that it is possible to relate the misalignment tolerance of the entire system to the optical invariant that is defined by the size of the detector and the numerical aperture of the collecting optics. Finally, in this area we have also investigated the way that system misalignment tolerances stack-up as more components have are added, using Monte-Carlo simulation. We have shown that in order to obtain a high probability that the system as a whole will function well, individual components must be aligned with sufficient accuracy that each individual component misalignment contributes to a reduction in throughput of 1% or less. The Monte-Carlo analysis also confirms that multistage relays significantly reduce the probability of success. Since the optimal cluster design analysis also shows that single stage relays can support the greatest interconnection density, this suggests that single stage (4-f) relays should be used wherever possible.

ACKNOWLEDGMENT

The authors would like to thank the many current and past members of the McGill Photonic Systems Group who have contributed to this research.

REFERENCES

[1] "Optical interconnections for digital systems," *Special Issue Proc. IEEE*, vol. 88, pp. 721–876, June 2000.

[2] J. H. Collet, W. Hlayhel, and D. Litaize, "Parallel optical interconnects may reduce the communication bottleneck in symmetric multiprocessor computers," *Appl. Opt.*, vol. 40, no. 20, pp. 3371–3378, 2001.

[3] F. Lacroix, E. Bernier, M. H. Ayliffe, F. A. P. Tooley, D. V. Plant, and A. G. Kirk, "Implementation of a compact, four-stage, scalable optical interconnect for photonic backplane applications," *Appl. Opt.*, vol. 41, no. 8, pp. 1541–1555, 2002.

[4] A. G. Kirk, D. V. Plant, T. H. Szymanski, Z. G. Vranesic, F. A. P. Tooley, D. R. Rolston, M. H. Ayliffe, F. K. Lacroix, B. Robertson, E. Bernier, and D. F. Brosseau, "Design and implementation of a modulator-based free-space optical backplane for multiprocessor applications," *Appl. Opt.*, vol. 42, pp. 2465–2481, 2003.

[5] D. V. Plant, M. B. Venditti, E. Laprise, J. Faucher, K. Razavi, M. Chateauneuf, A. G. Kirk, and J. D. Ahearn, "A 256 channel bi-directional optical interconnect using VCSELs and photodiodes on CMOS," *IEEE J. Lightwave Technol.*, vol. 19, pp. 1093–1103, Aug. 2001.

[6] M. Châteauneuf, A. G. Kirk, D. V. Plant, T. Yamamoto, J. D. Ahearn, and W. Luo, "512-channel vertical-cavity surface-emitting laser based free-space optical link," *Appl. Opt.*, vol. 41, pp. 5552–5561, 2002.

[7] M. B. Venditti, E. Laprise, J. Faucher, P.-O. Laprise, J. S. Ahearn, and D. V. Plant, "Design and verification of an OE-VLSI chip with 1080 VCSELs and PDs heterogeneously integrated with CMOS," in *Proc. IEEE/LEOS 2001 Annual Meeting*, 2001, PD-1.4.

[8] M. B. Venditti, E. Laprise, J. Faucher, P.-O. Laprise, J. E. Lugo, and D. V. Plant, "Design and test of an optoelectronic-VLSI chip with 540-element receiver/transmitter arrays using differential signaling," IEEE Journal of Select. Topics Quantum Electron., to be published.

[9] E. Bisaillon, D. F. Brosseau, T. Yamamoto, M. Mony, E. Bernier, D. Goodwill, D. V. Plant, and A. G. Kirk, "Free-space optical link with spatial redundancy for misalignment tolerance," *IEEE Photon. Technol. Lett.*, vol. 14, pp. 222–224, Feb. 2002.

[10] M. H. Ayliffe, D. R. Rolston, A. E. L. Chuah, E. Bernier, F. S. J. Michael, D. Kabal, A. G. Kirk, and D. V. Plant, "Design and testing of a kinematic package supporting a 32 × 32 array of GaAs MQW modulators flip-chip bonded to a CMOS chip," *IEEE J. Lightwave Technol.*, vol. 19, pp. 1543–1599, Oct. 2001.

[11] M. H. Ayliffe, M. Chateauneuf, D. R. Rolston, A. G. Kirk, and D. V. Plant, "Six-degrees-of-freedom alignment of two-dimensional array components using *in-situ* off-axis diffractive structures," *Appl. Opt.*, vol. 40, pp. 6515–6526, 2001.

[12] M. Chateauneuf, M. H. Ayliffe, and A. G. Kirk, "*In Situ* technique to measure the orthogonality of a plane wave to a substrate," *Opt. Lett.*, vol. 28, pp. 677–679, 2003.

[13] M. H. Ayliffe and D. V. Plant, "Design of misalignment-tolerant free-space optical interconnects," *Proc. SPIE*, vol. 4089, pp. 905–916, 2000.

[14] F. Lacroix and A. G. Kirk, "Tolerance stack-up in free-space optical interconnects," *Appl. Opt.*, vol. 40, no. 29, pp. 5240–5247, 2001.

[15] M. Châteauneuf and A. G. Kirk, "Determination of the optimum cluster parameters for a clustered free-space optical interconnect," in *Proc. LEOS Annu. Meeting*, Glasgow, U.K., Nov. 2002, pp. 899–900.

[16] A. G. Kirk and M. Châteauneuf, "Optical design choices for VCSEL-based free-space optical interconnects," in *Proc. SPIE 4942 (VCSELs and Optical Interconnects)*, Bruges, Belgium, 2002.

[17] C. Alleyne and A. G. Kirk, "Transmission uniformity of diffractive parallel optical interconnect relays: A numerical analysis based on rigorous coupled wave theory," in *IEEE LEOS Annu. Meeting*, Glasgow, U.K., 2002, pp. 901–902.

[18] Semiconductor Industry Association, *International Technology Roadmap for Semiconductors: 1999 Edition*. Austin, TX: International SEMATECH, 1999.

[19] J. M. Sasian, R. A. Novotny, M. G. Beckman, S. L. Walker, M. J. Wojcik, and S. J. Hinterlong, "Fabrication of fiber arrays for optical computing and switching applications," *OSA Int. Conf. Opt. Comput. Techn. Dig.*, pp. 229–230, 1994.

[20] C. V. Cryan, "Two-dimensional multimode fiber array for optical interconnects," *Electron. Lett.*, vol. 34, pp. 586–587, 1998.

[21] C. Berger, J. Ekman, X. Wang, P. Marchand, H. Spaanenburg, F. Kiamilev, and S. Esener, "Parallel distributed free-space optoelectronic computer engine using flat 'plug-on-top' optics package," in *Proc. SPIE, Optics in Computing 2000*, vol. 4089, 2000, pp. 1037–1045.

[22] M. W. Haney, M. P. Christianson, F. Milojkovic, G. J. Fokken, M. Vickberg, B. K. Gilbert, J. Rieve, J. Erkman, P. Chandramani, and F. Kiamilev, "Description and evaluation of the *Fast-Net* smart pixel-based optical interconnection prototype," *Proc. IEEE*, vol. 88, pp. 819–828, June 2000.

[23] K. M. Geib, K. D. Choquette, D. K. Serkland, A. A. Allerman, and T. W. Hargett, "Fabrication and performance of two-dimensional matrix addressable arrays of integrated vertical-cavity lasers and resonant cavity photodetectors," *IEEE J. Select. Topics Quantum Electron.*, vol. 8, pp. 943–947, July 2002.

[24] T. Maj, A. G. Kirk, D. V. Plant, J. F. Ahadian, C. G. Fonstad, K. L. Lear, K. Tatah, M. S. Robinson, and J. A. Trezza, "Interconnection of a two-dimensional array of vertical-cavity surface-emitting lasers to a receiver array by means of a fiber image guide," *Appl. Opt.*, vol. 39, no. 5, pp. 683–689, 2000.

[25] D. M. Chiarulli, S. P. Levitan, P. Derr, R. Hofmann, B. Greiner, and M. Robinson, "Demonstration of a multichannel optical interconnection by use of imaging fiber bundles butt coupled to optoelectronic circuits," *Appl. Opt.*, vol. 39, no. 5, pp. 698–703, 2000.

[26] Y. Li, T. Wang, H. Kosaka, S. Kawai, and K. Kasahara, "Fiber-image-guide-based bit-parallel optical interconnects," *Appl. Opt.*, vol. 35, pp. 6920–6933, 1996.

[27] A. E. Willner, C. J. Chang-Hasnain, and J. E. Leight, "2-D WDM optical interconnections using multiple-wavelength VCSELs for simultaneous and reconfigurable communication among many planes," *IEEE Photon. Technol. Lett.*, vol. 5, pp. 838–841, July 1993.

[28] M. B. Venditti and D. V. Plant, "On the design of large receiver and transmitter arrays for OE-VLSI applications," IEEE/OSA J. Lightwave Technol., vol. 21, Dec., 2003, submitted for publication.

[29] A. W. Lohmann, "Image formation of dilute arrays for optical information processing," *Opt. Commun.*, vol. 86, pp. 365–370, 1991.
[30] F. A. P. Tooley, "Challenges in optically interconnecting electronics," *IEEE J. Select. Topics Quantum Electron.*, vol. 2, pp. 3–13, Jan./Feb. 1996.
[31] D. Plant and A. G. Kirk, "Optical interconnects at the chip and board level: Challenges and solutions," *Proc. IEEE*, vol. 88, pp. 806–817, June 2000.
[32] F. B. McCormick, F. A. P. Tooley, T. J. Cloonan, J. M. Sasian, and H. S. Hinton, "Optical interconnects using microlens arrays," *J. Opt. Quantum Electron.*, vol. 34, pp. 6471–6480, 1992.
[33] D. R. Rolston, B. Robertson, H. S. Hinton, and D. V. Plant, "Analysis of a microchannel interconnect based on the clustering of smart-pixel-device windows," *Appl. Opt.*, vol. 35, no. 8, pp. 1220–1233, 1996.
[34] H. Thienpont, C. Debaes, V. Baukens, H. Ottevaere, P. Vynck, P. Tuteleers, G. Verschaffelt, B. Volckaerts, A. Hermanne, and M. Hanney, "Plastic microoptical interconnection modules for parallel free-space inter-and intra-MCM data communication," *Proc. IEEE*, vol. 88, pp. 769–779, June 2000.
[35] D. T. Neilson, "Tolerance of optical interconnections to misalignment," *Appl. Opt.*, vol. 38, pp. 2282–2290, 1999.
[36] M. H. Ayliffe, "Alignment and packaging techniques for two-dimensional free-space optical interconnects," Ph.D. dissertation, McGill Univ., Montreal, QC, Canada, 2001.
[37] P. Belland and J. P. Crenn, "Changes in the characteristics of a Gaussian beam weakly diffracted by a circular aperture," *Appl. Opt.*, vol. 21, pp. 522–527, 1982.
[38] D. T. Neilson and E. Schenfeld, "Free-space optical relay for the interconnection of multimode fibers," *Appl. Opt.*, vol. 38, no. 14, pp. 2944–2952, 1999.
[39] W. J. Smith, *Modern Optical Engineering: The Design of Optical Systems*, 2nd ed. New York: McGraw-Hill, 1990.
[40] Y. S. Liu, G. C. Boisset, M. H. Ayliffe, R. Iyer, and D. V. Plant, "Design, implementation and characterization of a four stage hybrid optical system for a free-space optical backplane demonstrator," *Appl. Opt.*, vol. 37, no. 14, pp. 2895–2914, 1998.
[41] F. Lacroix, M. Chateauneuf, X. Xue, and A. G. Kirk, "Experimental and numerical analysis of misalignment tolerances in free-space optical interconnects," *Appl. Opt.*, vol. 39, no. 5, pp. 704–713, 2000.
[42] D. T. Neilson and E. Schenfeld, "Plastic modules for free-space optical interconnects," *Appl. Opt.*, vol. 37, no. 14, May 10, 1998.
[43] D. Zaleta, S. Patra, V. Orguz, J. Ma, and S. H. Lee, "Tolerancing of borad-level-free-space optical interconnects," *Appl. Opt.*, vol. 35, no. 8, pp. 1317–27, 1996.
[44] W. L. Hendrick, N. S. F. Ozkan, P. J. Marchand, and S. C. Esener, "A Monte-Carlo simulation for tolerancing microbeam free space optical interconnects," in *Proc. Optics in Computing*, vol. OC-99, Snowmass, CO, Apr. 1999.
[45] S. D. Nigam and J. U. Turner, "Review of statistical approaches to tolerance analysis," *Computer-Aided Design*, vol. 27, no. 1, pp. 6–15, 1995.
[46] B. Robertson, "Design of an optical interconnect for photonic backplane applications," *Appl. Opt.*, vol. 37, no. 14, pp. 2974–84, 1998.
[47] S. P. Levitan, T. P. Kurzweg, P. J. Marchand, M. A. Rempel, D. M. Chiarulli, J. A. Martinez, J. M. Bridgen, C. Fan, and F. B. McCormick, "Chatoyant: A computer-aided-design tool for free-space optoelectronic system," *Appl. Opt.*, vol. 37, no. 26, Sept. 10, 1998.
[48] N. S. F. Ozkan, W. L. Hendrick, P. J. Marchand, and S. C. Esener, "Misalignment tolerance analysis of free-space optical interconnects via statistical methods," *Appl. Opt.*, vol. 41, pp. 2686–2694, 2002.
[49] M. P. Christensen, P. Milojkovic, and M. W. Haney, "Analysis of a hybrid micro/macro-optical method for distortion removal in free-space optical interconnections," *JOSA A*, vol. 19, pp. 2473–2478, 2002.
[50] D. Fey, W. Erhard, M. Gruber, J. Jahns, H. Bartelt, G. Grimm, L. Hoppe, and S. Sinzinger, "Optical interconnects for neural and reconfigurable VLSI architectures," *Proc. IEEE*, vol. 88, pp. 838–848, June 2000.
[51] M. Gruber, J. Jahns, E. Mehdi, E. Joudi, and S. Sinzinger, "Practical realization of massively parallel fiber-free-space optical interconnects," *Appl. Opt.*, vol. 40, no. 17, pp. 2902–2908, 2001.
[52] B. Robertson, Y. Liu, G. C. Boisset, M. R. Taghizadeh, and D. V. Plant, "*In situ* interferometric alignment systems for the assembly of microchannel relay systems," *Appl. Opt.*, vol. 35, pp. 9253–9260, 1997.
[53] P. Tuteleers, A. Kirk, M. Chateauneuf, H. Ottevaere, V. Baukens, C. Debaes, M. Vervaeke, A. Hermanne, I. Veretennicoff, and H. Thienpont, "Investigation of the replication quality of plastic micro-optical interconnection components," in *Proc. 6th Ann. Symp. IEEE/LEOS Benelux Chapter*, Brussels, Belgium, 2001, pp. 73–77.

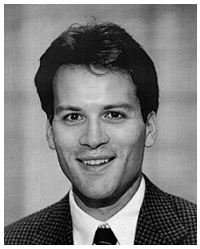

Andrew G. Kirk (M'97) received the B.Sc. degree in physics from the University of Bristol, Bristol, U.K., in 1989 and the Ph.D. degree in physics from the University of London (King's College), London, U.K., in 1992.

From 1992 to 1994, he held a Visiting Research Fellowship with the Department of Mathematical Engineering and Information Physics, the University of Tokyo, Tokyo, Japan, and in January 1994, he began work as a Postdoctoral Research Assistant in the Department of Applied Physics, Brussels Free University, Brussels, Belgium. In 1996, he joined the Department of Electrical and Computer Engineering, McGill University, Montreal, QC, Canada, where he currently holds the position of Associate Professor and William Dawson Scholar. His research interests include microoptics, diffractive optics, and microelectromechanical systems. He has investigated the application of microoptics to free-space optical interconnects in high-speed digital systems and he is also interested in the use of three-dimensional microoptical systems in telecommunications network components.

Dr. Kirk received the Principal's Award for Excellence in Teaching in 2000.

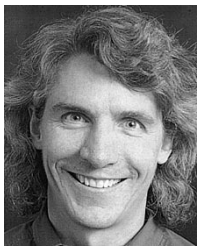

David V. Plant (S'86–M'89) received the Ph.D. degree in electrical engineering from Brown University, Providence, RI, in 1989.

From 1989 to 1993, he was a Research Engineer with the Department of Electrical and Computer Engineering, the University of Los Angeles at California (UCLA). He has been a Professor and Member of the Photonic Systems Group, the Department of Electrical and Computer Engineering, McGill University, Montreal, QC, Canada, since 1993. He was tenured and promoted to Associate Professor in 1997. During the 2000 to 2001 academic years, he took a leave of absence from McGill University to become the Director of Optical Integration at Accelight Networks, Pittsburgh, PA. He is the Scientific Director of the Canadian National Science and Engineering Research Council (NSERC) funded research network entitled Agile All-Photonic Networks.

Dr. Plant received the Outstanding Departmental Teaching Award and the Faculty of Engineering Teaching Award in 1996. In recognition of his accomplishments in research and teaching, he was named a James McGill Professor in 2001.

Michael H. Ayliffe received the B.Eng. and Ph.D. degrees in electrical engineering from McGill University, Montreal, QC, Canada, in 1996 and 2001, respectively. His Ph.D. research focused on packaging and alignment of two-dimensional arrays of optoelectronic devices.

He was with Picolight Inc., Boulder, CO, from 2001 to 2002 and then joined Ignis Optics, San Jose, CA, as a Senior Design Engineer where he currently works on the design of 10-Gb/s singlemode optical transceivers.

Marc Châteauneuf received the B.Eng. degree (with honors) in electrical engineering from McGill University, Montreal, QC, Canada, in 1998, where he is currently working toward the Ph.D. degree, which focuses on design and alignment methodologies for freespace optical interconnects.

Since September 2002, he has been with the Defense Research and Development Canada, Valcartier, QC.

Frédéric Lacroix received the baccalauréat degree in engineering physics from Laval University, Quebec, QC, Canada, in 1996, and the M.Eng. and Ph.D. in electrical engineering from McGill University, Montreal, QC, Canada, in 1998 and 2001, respectively. His Ph.D. research focused on modeling of free-space optical interconnects.

He joined ONI Systems Corporation (now Ciena Corporation), San Jose, CA, in 2001 as Senior Optical Design Engineer, and has worked on design of 10- and 2.5-Gb/s optical transmission systems.