

Optoelectronic Interconnection Technology in the HOLMS System

Paul Lukowicz, *Member, IEEE*, Jürgen Jahns, *Member, IEEE*, R. Barbieri, P. Benabes, T. Bierhoff, A. Gauthier, M. Jarczynski, G. A. Russell, J. Schrage, W. Süllau, J. F. Snowdon, M. Wirz, and G. Tröster

Invited Paper

Abstract—The High-Speed Optoelectronic Memory Systems (HOLMS) project, sponsored by the European Union Information Society Technology program, aims to make the use of board level optical interconnection in information systems practical and economical by developing optoelectronic packaging technology compatible with standard electronic assembly processes. To demonstrate the potential of the technology, it develops a demonstrator system that addresses the most pressing problem of contemporary computer architecture: memory latency. This paper describes the key ideas and some preliminary results of the HOLMS projects focusing on the electronic interconnection technology, in particular optoelectronic packaging issues.

Index Terms—Integrated optics, optical computing, optical interconnections, optoelectronic packaging, waveguide interconnections.

I. INTRODUCTION

CONCENTRATED research efforts in recent years have lead to significant progress in optoelectronic components including the tight integration of optoelectronic input/outputs (I/Os) and silicon chips, novel microoptical devices, and advances in optomechanical integration [36]. A number of high-bandwidth interconnection demonstrators were also constructed. However, so far, optical interconnections have had nearly no impact on board or even rack level interconnections in real life systems. This failure can be mostly attributed to two factors: 1) inadequate and costly packaging technology and 2) inadequate system architectures.

Manuscript received February 5, 2003.

P. Lukowicz, R. Barbieri, M. Wirz and G. Tröster are with the ETH Zurich, Electronics Laboratory, Zurich CH 8805, Switzerland (e-mail: lukowicz@ife.ee.ethz.ch).

J. Jahns and M. Jarczynski are with the FernUniversitaet Hagen, Optische Nachrichtentechnik, Universitaetsstrasse 27/PRG, D-58084 Hagen, Germany.

P. Benabes and A. Gauthier are with the Supelec, Service des Mesures, 1192 Gif sur Yvette Cedex, France.

T. Bierhoff is with the University of Paderborn, Theoretische Elektrotechnik/C-LAB, D-33102 Paderborn, Germany.

G. A. Russell and J. F. Snowdon are with the Heriot-Watt University, Riccarton, Edinburgh EH14 4AS, UK.

J. Schrage is with the Siemens SBS C-LAB, D-33102 Paderborn, Germany.

W. Süllau is with the ILFA Feinleitertechnik GmbH, Hannover D-30559, Germany.

Digital Object Identifier 10.1109/JSTQE.2003.815492

a) *Packaging Technology*: Electronic assembly technology has become optimized for low-cost mass production. Lithography based on cheap materials for system-level printed circuit board (PCB) interconnection and chip packages suitable for attachment using pick and place machines allow low-cost automated production. Optical systems, as they are implemented today, are mostly incompatible with this assembly technology. This has created a barrier for the adoption of optical interconnection that neutralizes any potential benefits.

b) *System Architecture Issues*: The increasing gap between very large scale integration (VLSI) circuit speed and the I/O capabilities of system-level electronic interconnection technology has often been cited as a compelling reason for the transition to optical interconnections [35]. Obviously, extrapolating current trends in VLSI technology far enough shows that at a certain point, chip-to-chip bandwidth requirements that are going to exceed the physical limits of electronic links. On the other hand, the proponents of purely electronic technology argue that, for current system architecture, this point will not be reached for another 15 to 20 years [48]. At the same time, while implementing adequate interconnections is going to become increasingly harder with higher frequencies, in practice, evolutionary improvements in electronic signaling and detection techniques are more attractive than switching to a new technology. As a consequence, it is not clear when (if at all) mere replacement of electrical links with higher bandwidth optical pendant will become an attractive option for system designers. However, current computer systems have been optimized to circumvent the shortcoming of electronic interconnections. It is, thus, understandable that they cannot fully utilize the advantages of optical interconnections. As a consequence, system architectures specifically designed to take advantage of the properties of optical transmission are needed to make optical interconnection technology attractive for practical applications.

A. High-Speed Optoelectronic Memory Systems (HOLMS) Project

The HOLMS project sponsored by the European Union Information Society Technology program aims to overcome the

above obstacles to the practical use of system-level optical interconnections.¹ To this end, it develops an optoelectronic packaging technology that facilitates seamless integration of free space, waveguides, and fiber optical links with standard electronic packaging technology. The core of the technology is a novel optoelectronic multichip module (OE-MCM) concept.

In the architecture area, the key concept of the HOLMS approach can be summarized as *cutting through the interconnection hierarchy* [29] to reduce communication latency. The concept exploits the fact that in the relevant approximation the bandwidth and latency of an optical link is virtually independent of the length, the fanout/fanin, and the overall density of the interconnection. At the same time, it leverages new technological developments that allow close integration of large arrays of low-latency high-bandwidth optical I/Os with complex VLSI circuits. As a result, new interconnection architectures replacing high-latency multistage networks with direct low-latency links between all relevant components become possible. In addition, generous use of optical fanout allows the implementation of efficient cache coherence protocols.

To demonstrate the potential of our approach, we address what is generally regarded as the most pressing problem of contemporary computer architecture: memory latency. For this purpose, a memory architecture specifically tailored to utilize the advantages of optical interconnections has been developed. In the project, a fully functional demonstrator system based on this architecture is being built to show that optical interconnection technology is capable of providing system-level benefits in real life computers.

B. Paper Overview

This paper focuses on the optoelectronic interconnection packaging technology and its application in the demonstrator system. The description of the memory architecture is limited to an overview required for the understanding of the interconnection issues. After an introductory overview of the HOLMS concepts in Section II, we present the HOLMS demonstrator system in Section III. We, then, turn to a detailed discussion of the OE-MCM technology in Section IV, and the printed circuit board-based waveguide interconnections in Section V. Finally, we discuss the interface between OE-MCM devices and the PCB embedded waveguides that combines standard electronic attachment with an appropriate optomechanical coupling mechanism.

C. Related Work

Early implementation of optoelectronically connected computers include, among others, the Delft fully connected multiprocessor [52], the Cosine system [43], and the MEMSY machine [31]. At the same time, a number of architecture concepts such as a time division multiple access-based crossbar [37], a hierarchical hypercube based design [40], and different bus systems were studied (e.g., [27]).

More recent optoelectronic systems demonstrators include the opto field programmable gate array [3], [17] built by an EU

sponsored consortium, the optoelectronic backplane system realized by McGill lead Canadian consortium [6], [23], [49], the DARPA three-dimensional (3-D) optoelectronic processor system [11], the Heriot-Watt SPOEC high-speed interconnection system [1], and the RWC-1 massively parallel system [42]. In addition, different theoretical architecture concepts have been developed [13], [15], [26]. In this context, a considerable amount of attention has also been given to the issue of optical interconnections in shared memory systems [7], [13], [18], [21], [30], [45].

Most of the demonstrator projects mentioned above have developed their own packaging concepts. They include an OE-MCM technology similar to the one pursued in the HOLMS project [12], and different precision mechanical systems (e.g., [33], [38], [50]). In addition, a number of other groups have looked at novel optoelectronic packaging techniques (e.g., [14], [34]). Most notably, an image guide-based packaging concept has been recently proposed and demonstrated by several groups [5], [22], [51], that aims at similar applications as the HOLMS technology.

II. HOLMS OVERVIEW

A. HOLMS Rationale

As described in the introduction the main idea behind the HOLMS project is that optical interconnection technology can be used to cut through the traditional electrical interconnection hierarchy. This hierarchy is the main factor in the communication latency, which, in turn, is the main performance bottleneck in today's computer systems.

1) *Interconnection Hierarchy and Latency*: Hierarchical multistage interconnect architectures are a fundamental concept of electronic system design. In general, electronic systems are hierarchically partitioned into subsystems consisting of physically adjacent components. Such subsystems can be the circuits on a single chips, chips on a single multichip module, or a local group of chips and multichip modules on a PCB. The interconnection architecture follows the subsystem structure favoring local links inside a subsystem over intersubsystem connections. In most subsystems, there are just a few designated components (interfaces) that are responsible for communicating with other subsystems. The majority of such links connect to subsystems on the same or, at most, one stage higher hierarchy level. Empirical studies have shown that, for most systems, the number of interconnections in any given subsystem exponentially decreases with increasing hierarchy level. Since increasing hierarchy levels correspond to increasing subsystem diameter, this also means that the link density exponentially decreases with distance. As a consequence, communication between far away components is mostly indirect, passing through the interfaces between many hierarchy stages. This leads to high latency and introduces complex communication protocols. Thus, for example, the message latency in a multicabinet parallel computer or a workstation cluster can be as high as a few microseconds, which is three orders of magnitude above the time of flight delay.²

¹The Project participants are: 1) Heriot-Watt University, Scotland; 2) ETH Zurich, Switzerland; 3) THALES Communications, France; 4) Siemens/C-Lab, Germany; 5) University of Hagen, Germany; 6) SUPELEC, France; 7) University of Paderborn, Germany; and 8) ILFA AG Germany.

²This assumes a signal speed of about 30 cm/ns and a cabinet-to-cabinet distance of one or two meters.

For many systems, the above hierarchical architecture is a natural and efficient implementation option, since it reflects their logical structure, which is dominated by local communication. However, in cases where global communication is important, the high latency of long range communication is a major performance bottleneck. As will be described later, shared memory multiprocessors (SMPs) are a good example of this type of system.

One of the main reasons for the interconnection hierarchy is the fact that the performance of electronic interconnections is strongly dependent on the physical properties of a link. To achieve high bandwidth, the links and the transceivers must be exactly impedance matched and care must be taken to route all signals in such a way that crosstalk is minimized. In addition, the interconnection complexity is limited by the planar nature of the interconnect, since the interconnect area is proportional to the second power of the interconnection bisection width.

2) *Cutting Through the Hierarchy Using Optics:* The principal advantage of optical interconnections lies in the fact that in the relevant approximation the bandwidth and latency of an optical link is virtually independent of the length the fanout/fanin and the overall density of the interconnection. In addition, with free space optics the interconnect area increases only linearly with the bisection width. Furthermore, new trends in optoelectronic integration are emerging that allow hundreds or even thousands of high-speed optical I/Os to be directly integrated with high-performance VLSI circuits. Thus, a single chip can have direct high-bandwidth low-latency links to a large number of devices both in its direct neighborhood and as far away as the next cabinet. This allows for new architecture concepts that cut through the traditional interconnection hierarchy providing direct time of flight limited links between all relevant components. In particular, the number of stages between the CPU and the memory chips in a multiprocessor system can be reduced leading to memory latency close to the chip access time. At the same time, appropriate use of broadcast channels can simplify cache coherence protocols.

B. Memory Architecture

1) *Memory Latency in Multiprocessor Systems:* One of the main problems of today's computer architecture is the widening gap between processor cycle time and memory latency (e.g., [44]). This problem is particularly grave in SMPs, where the CPUs of the parallel machine communicate through a common memory system. Depending on the number of CPUs, the memory access latency in such systems ranges from about 200 ns in dual CPU machines to more than 1000 ns in large scale servers.

Memory latency is often blamed on the gap between DRAM speed and CPU operating frequency. Although such a gap is undeniable, it explains only a small part of the above latency. Today DRAM memory modules with a latency of 22.5 ns are available (ESDRAM). Single transistor high-density SRAM (ESRAM) can provide a random access latency of 12.5 ns at similar density and only slightly higher cost. This is between 1% and 10% of the overall latency. The remaining 90% to 99% are due to interconnection delays [4].

HOLMS Memory Architecture




Fig. 1. Idea behind the HOLMS memory architecture. In this diagram, there are two processors, $NG = 6$ memory bank groups, each with $NB = 2$ memory banks. Each memory bank has $NC = 4$ RAM chips organized around a controller that is optically connected to the processor.

With physical system diameters between 50 cm and a few meters for large multirack systems, it is obvious that the distance limited time of flight of the signals does not significantly contribute to the above delay. Instead, the interconnection latency is due to a hierarchical, multistage network architecture. Such architecture introduces a cascade of driver and conversion delays as well as complex communication and memory access protocols.

2) *HOLMS Memory Architecture:* In the HOLMS project, the concept of cutting through the interconnection hierarchy is applied to the processor memory network in multiprocessor systems. As described in [28], high-latency multistage links are replaced by direct optical channels between optoelectronic I/Os closely integrated with the relevant VLSI components. To provide the lowest possible latency, the processor should have a direct optical connection to every RAM chip in the system. However, even using free-space optics implementing such a complex interconnection is not practical. As a compromise, the HOLMS architecture uses a single intermediate stage: a controller that is responsible for a small group of RAM chips. The processors and the controllers communicate with each other using optical I/Os closely integrated with electronic circuits. To further reduce the number of processor I/Os, broadcast, rather than point-to-point channels, are used with each transmitter on the processor. Thus, each transmitter is assigned to several memory controllers. In general, the HOLMS memory architecture is composed of NRAM banks, each consisting of NC RAM chips electronically connected to an optoelectronic bank controller. The memory banks are organized in NG groups with $NB = N/NG$ banks each. The memory controllers in a single memory bank group share a single fanout/fanin channel optical channel for each processor. An example of a system with $NG = 6$, $NC = 4$, $NB = 2$, and two CPUs is shown in Fig. 1.

C. Interconnection Technology and Packaging Concept

As described above, our approach to using optical interconnections in computer systems exploits the insensitivity of optical links to the length, fanout, and complexity to cut through the traditional interconnection hierarchy. The main practical problem facing this approach is the fact that, depending on the distance




Fig. 2. HOLMS optoelectronic packaging technology. An OE-MCM containing a PIFSO free-space optical system and electronic interconnection attached to the glass substrate is soldered on a PCB containing an optical waveguide layer.

and type of interconnection, different transmission technologies need to be used. In the HOLMS project, these are:

- 1) Parallel fiber arrays for transmission of data between different PCBs. Here, a commercial system based on the MT connector standard is used.
- 2) Waveguides integrated in a conventional electronic PCB for a robust, low cost point-to-point transmission across the system board. This technology has been developed by Siemens and University of Paderborn [8].
- 3) Planar integrated free-space systems (PIFSO) for fanin/fanout operations and the mapping of the optoelectronic I/Os to the fibers and waveguides. PIFSO systems, developed by University of Hagen [4], fold a free space optical system inside a thin glass substrate. They provide all the benefits of free-space optic without the alignment and stability problems usually associated with conventional free space systems. In addition, they can be used as an optomechanical interface to the MT connectors of the fiber arrays [32].

The key packaging issue faced by the HOLMS project is the development of an optomechanical interface that allows a single electronic component mounted on a printed circuit board to simultaneously access all the different transmission technologies. This needs to be accomplished in a way compatible with standard electronic assembly techniques. For this purpose, a novel optoelectronic multi-chip module (OE-MCM) packaging technology [39], [47] has been developed (see Fig. 2) As described in more detail in Section IV, it combines PIFSO systems with thin film on glass electrical MCM-D technology. Electronic and optoelectronic components are flip chip bonded to the thin film layers deposited on the glass substrate containing the free space optics. Appropriate optical windows in the thin film layers allow the optoelectronic components to be coupled into the optical system. The whole substrate can be packaged as a ball grid array (BGA) and attached to a PCB using conventional soldering techniques.

III. HOLMS DEMONSTRATOR SYSTEM

A. Overall Architecture

The HOLMS demonstrator is a four CPU shared memory heterogeneous multiprocessor with a low-latency optoelectronic

memory interconnection based on the transmission technologies and architecture concept sketched in the previous section. To demonstrate the capabilities of the system, a JPEG 2000 decoder software for large medical and satellite imaging applications will be implemented. For large images, it requires a lot of memory with low-latency random access, which means that it optimally utilizes the advantages of the planned system.

As shown in Fig. 3, the HOLMS demonstrator consists of four CPUs and 16 memory banks distributed over two PCBs. Each CPU is located on its own OE MCMs (CPU OE-MCM) together with an optoelectronic memory management unit. On the memory side, two memory bank controllers constituting a memory bank group are located on each OE MCM (MEM OE-MCM). Each CPU OE-MCM is connected to each MEM OE-MCM using a dedicated, point to point, duplex optical channel. Each channel consists of eight data, three control, and one clock line, each driven at 600 MHz. Depending on whether the OE-MCMs are on the same PCB or not, the channel is implemented using waveguides or fiber arrays. In any case, this link is shared by both memory bank controllers of the OE-MCM.

B. CPU OE-MCM

As shown in Fig. 4, CPU OE-MCM consists of the CPU, a mixed-signal controller chip, the optical I/Os, and a PIFSO optical system. The PIFSO optical system of the OE-MCM is used to image the optical I/Os of the controller onto the fiber and waveguide interface. This requires a 1:1 transmission through a light pipe system combined with appropriate pitch adjustment. For optimal execution speed of the JPEG 2000 applications, one CPU OE-MCM will contain a high end Analog Devices Tiger-Shark DSP, while the remaining three will be assigned custom CPUs implemented in XILINX XCV2V3000 FPGAs. Independent of the processor type, all CPU OE-MCMs have the same controller chip. It consists of two parts: a memory management unit (MMU) and the transmission circuits.

The MMU is identical for all four CPUs. It handles the interface to the optical bus and performs actions related to cache coherence, synchronization primitives, and conflict resolution in cooperation with the memory controller chips.

The main components of the transmission circuits are serialization/deserialization logic and analog VCSEL drivers and

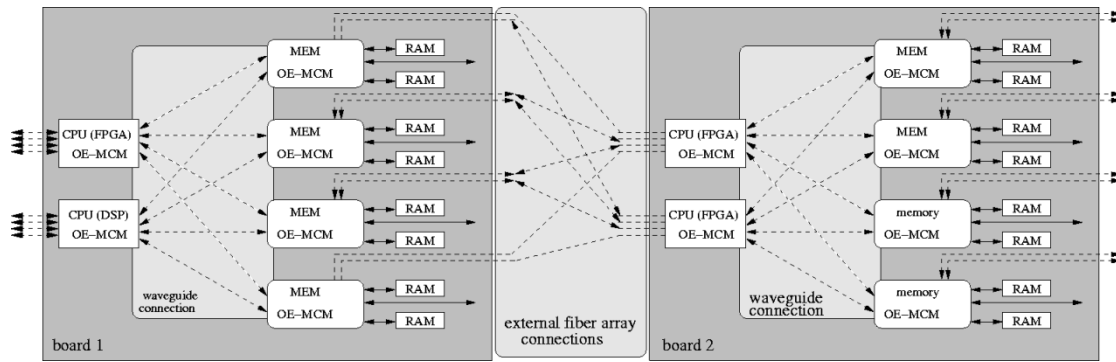


Fig. 3. Architecture of the HOLMS demonstrator system.

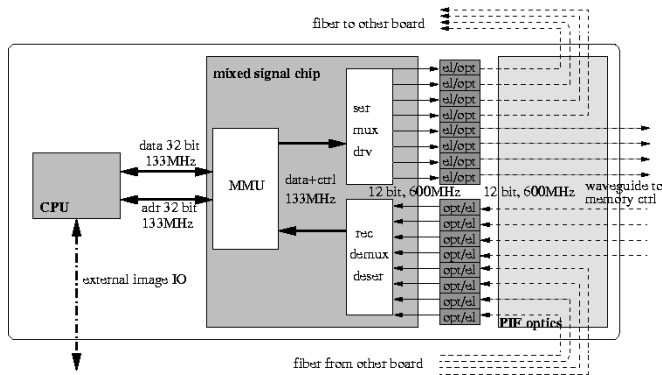


Fig. 4. Architecture of the CPU OE-MCM of the HOLMS demonstrator.

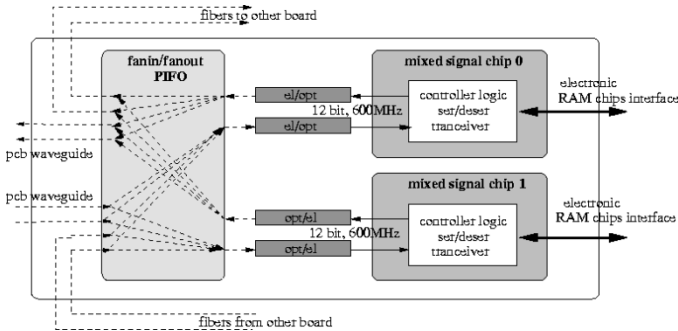


Fig. 5. Architecture of the memory OE-MCM of the HOLMS demonstrator.

photo receivers amplifiers. Since the CPU address and data bus are 32 bits wide at a frequency of 133 MHz and optical bus data is 8 bits wide, the optical bus frequency needs to be 600 MHz (Fig. 5). Due to the envisaged architecture, 48 drivers and 48 receivers will be integrated on the CPU side interface chip.

C. MEM OE-MCM

In the demonstrator system, there is a single MEM OE-MCM for each memory bank group. As shown in Fig. 5 it consists of two mixed signal optoelectronic memory bank controller chips: one for every memory bank of the group, PIFSO. In addition to pitch adjustment between the optical I/Os and the fibers and waveguides, the PIFSO performs the fanin/fanout operations required by the HOLMS's memory architecture.

Like the CPU side controllers, the mixed-signal chips consist of two parts: a purely digital memory controller unit and mixed-signal transmission circuits. The memory controller is

the logic responsible for address decoding, memory access, and conflict resolution. For memory access, there are two electrical channels to RAM chips on the printed circuit board, one for each controller. The system will be based on Enhanced Memory Systems' BSRAM SS2625-7.5, 2Mx36 chips with a random access latency of 12.5 ns and a capacity of nine MB per chip, of which only eight MB will be used for data. To achieve 512 MB total memory capacity with 16 memory banks in eight memory bank groups, four chips with a total of 32 MB have to be connected to each channel.

The transmission circuits take care of the frequency conversion and the analog interface (amplifiers and receivers) to the 24 optical channels required by the system architecture.

D. Transmission Logic and Optoelectronic I/Os

The design of an efficient transmission logic, drivers, and receivers is one of the major issues in the HOLMS's demonstrator realization. A review of published paper on drivers and photo receivers amplifiers shows that performances of CMOS technology with 3.3 V of power supply are too tight to obtain the desired performances. Taking into account that most of the commercial circuits are fabricated with a silicon-germanium technology, a SiGe BiCMOS 0.35 μm from Europractice technology has been chosen. The digital part of the chips will be designed using standard CMOS cells which are compatible with the BiCMOS technology.

1) *Optoelectronic Transmitters:* VCSEL's ULM 1x12 arrays at 250- μm pitch with a wavelength of 850 nm have been selected. The devices have a threshold current of 1.8 mA and work with an operating current of 5 mA. Due to variations between different VCSEL arrays, threshold and operating current need to be adjusted on a per array basis. As the drivers and the logic driving them are integrated on the same chip, these interface problems are drastically reduced. VCSEL are external components but since they will be, as well as the mixed chips, flip chip bonded onto a OE-MCM parasitic elements are minimized. Simulation with typical parameters showed that clock signal with frequency up to 1 GHz can be transmitted.

2) *Optoelectronic Receivers:* An identical 1x12 at 250 μm pitch configuration will be used for the detectors to match the VCSEL array. Operation will again be at 850 nm with an expected sensitivity as high as 0.7 AW^{-1} . The photodiodes will be constructed on a GaAs/AlGaAs strained-layer structure. They will be p-i-n devices and have two contacts each. It is suggested

that the bottom layer in the p-i-n structure is used as ground. However, this should not be connected to the substrate to avoid any potential reduction of bandwidth. The photodiodes will be circular with an optical window of 180 μm in diameter and active area of 220 μm in diameter. This gives a detector area of $3.8 \times 10^{-8} \text{ m}^2$. Since the intrinsic region will be 1.5 μm thick, a photodiode has a predicted capacitance of 2.69 pF. This should allow operation at up to 1.18 GHz. Series device resistance will be approximately 50 ohms, thus, the load should be equivalent. It is expected that a reverse bias of up to 5 V may be applied without any problem. To ensure an acceptable bit error rate (BER) of 10^{-9} in digital transmission, the current generated by the incident optical power on the photodiode must be 144 times the intrinsic device rms noise current. Since material purity is very high, noise from the device is expected to be low. Therefore, given a worst case scenario of 10 W of incident optical power, achieving such a BER should not be a particular problem.

A review of a published paper shows that all architecture of photo receivers are based on transimpedance amplifiers. The design of such amplifiers is highly dependant of the characteristics (capacitance and sensitivity) of the optical receiver. As the capacitance of the p-i-n diode to be used is rather high, the design of the amplifier is a tradeoff between expected performances and power dissipation due to bias current.

IV. OPTOELECTRONIC MULTICHIP MODULES

The HOLMS optoelectronic packaging concept calls for free-space systems to be used in three areas: 1) for high-density interconnection at the chip-and MCM level; 2) fanout/fanin operations; and 3) as an interface between long-distance low-density waveguide and fiber connections and the high-density device arrays on the optoelectronic chips. For practical implementations, suitable integration and packaging techniques for free-space optics are required. The specific approach used here is the concept of PIFSO as described in many publications (e.g. [20]). The HOLMS project builds on this concept to develop an optoelectromechanical integration technology that allows the combination of free-space optical systems with high-performance electronic substrates, and the required interfaces to fiber arrays and waveguides. This paragraph describes this technology that we refer to as OE-MCMs.

A. PIFSO Technology

The main feature of PIFSO is the monolithic integration of the optical elements on a single optical substrate by using a lithographic fabrication process. The essence of this concept is the folding of optical paths within a transparent substrate, where the passive optical components that make up the system (beam splitters, fanout gratings, lenses, mirrors, and so on) are fabricated on the outer surfaces of the substrate. The optics are kept small, in the sub-mm scale, and these microoptical components are compatible with established lithographic production. Lithography gives precise alignment of the devices on the surface to micrometer precision and the common substrate removes the concern about alignment retention. The further advantages of the technology include the ease of packaging of the flat, quasi-2-D, sub-

strates. In short, PIFSO combines the capabilities of free-space optics with the packaging simplicity of guided wave optics or standard electronic surface mount components.

Different demonstration experiments have been reported, one that is relevant for the HOLMS concept was demonstrated, for example, in [10]. The technological challenge is optimizing the transmission efficiency.

B. OE-MCMs

The PIFSO substrate can served as an open platform for the integration of optoelectronic devices [19] and used as a basis for simple electronic systems. Signal lines, pads for flip-chip bonding of chips (electronic and/or optoelectronic), and soldering pads for attachment to printed circuit boards can be fabricated with the same lithographic technique used for the optical components. In initial experiments, partners ETHZ and Hagen have tested and characterized some simple systems [39].


The HOLMS project goes beyond such simple systems toward the full integration of complex, high-density electronics with PIFSO substrates in OE-MCM. The concept is an extension of thin film multichip module technology (MCM-D). In MCM-D systems, a few micrometers-thick layers of metal conductor and dielectric (e.g., polyimide or BCB) are deposited on a few hundred μm thicker substrate. The technology allows a feature size of down to 10 μm and component placement accuracies of a few micrometers. A variety of substrates, some of them transparent (e.g., glass), can be used.

To combine thin film MCM technology with planar optical systems, the optoelectronic components are flip-chip bonded on the MCM module with the optical windows pointing toward the substrate (Fig. 2). A via hole through all the metal and dielectric layers is placed underneath the optical windows to allow light signals to reach the substrate. The substrate must be transparent to the wavelength used. The MCM system with the optoelectronic components is glued to a PIFSO system, forming a complete OE-MCM.


For mounting an OE-MCM on a printed circuit board, soldering balls are placed on appropriate pads of the MCM-D substrate. The OE-MCM component can then be attached to a printed circuit board much like a conventional BGA packaged chip. Currently, the suitability of different attachment techniques, including reflow soldering, laser soldering, and conductive glue, is being evaluated. The main technical challenge stems from thermal expansion issues.

In the OE-MCM concept, planar optics is just another interconnection layer of a conventional high-performance system. Thus, an OE-MCM module can be handled and packaged much like a conventional electronic component. At the same time, it provides complex electronic systems with a flexible interface to both free-space optics for inter- or intrachip communication and system optical data transmission through fiber arrays or PCB integrated waveguide (as described in Section V).

1) *Experimental Devices:* Currently, different experimental electronic substrates are being designed, produced, and evaluated. Fig. 6 shows a test substrate consisting of a three in and one out, 12-channel interconnection module combining one LD array, three PD arrays, the corresponding drivers, and receivers with all the necessary passive electronic components. It is im-



(a)



(b)

Fig. 6. (a) Electronic MCM-D test substrate produced for a simple OE-MCM component. (b) A closeup of the optical windows placed next to the VCSEL and PD bonding pads and alignment structures in the central transparent window.

plemented in three-layer thin film MCM-D technology based on $10\ \mu\text{m}$ -thick BCB dielectric, $2\ \mu\text{m}$, and thick copper lines with a minimal feature size of $10\ \mu\text{m}$. The glass substrate is quartz, which is normally used in PIFSO systems. For the alignment with the PIFSO, the module has alignment marks placed in a central transparent window. This allows the marks to be seen from both sides of the OE-MCM allowing different flexible assembly sequences.

C. Optical System of the Demonstrator OE-MCMs

Although a number of stringent placement constraints have to be considered, the electronic layout of the demonstrator OE-MCM is fairly straightforward and will not be discussed in this paper. Instead, we concentrate on the optical design issues.

1) *General Considerations:* To support the HOLMS architecture, two kinds of PIFSO systems with different special fea-

tures have to be developed: one for the interconnections at each CPU-MCM (to be called CPU-PIFSO in the following) and one for each memory-PIFSO (MEM-MCM). Because of geometric extent and electrical alignment, the CPU-PIFSO has to be divided into two parts, one for connecting the fibers and the other one for connecting the waveguides to the optoelectronic chips. Consequently, a total number of eight PIFSO systems are necessary to implement the interconnections on one PCB.

The CPU-PIFSO systems require direct interconnections of sources and targets without any fanin or fanout. The technological and design challenge consists of providing a large number of channels ($2 \times 4 \times 12 = 96$) arranged in a single line for realizing the physical interface to the PCB-integrated waveguides (see Fig. 4). A similar task is required for the fiber connections.

The memory-PIFSO system connects two sources with four targets by coupling from one MEM-MCM into a fiber and an integrated waveguide. That means $24 (= 2 \times 12)$ channels in two rows for each memory bank are subject to a fanout of four and one pair of channels of each row has to be coupled into one fiber or waveguide (equivalent to a fanin of two). For coupling into the detectors of the MEM-MCM, the reverse way has to be used.

2) *Coupling Issues:* For the physical implementation, it is important to note that different types of light sources are used for the different systems: in the direction OE-chip to fiber/waveguide, VCSELs are the sources and fibers or waveguides play the role of targets, whereas in the reverse direction, fibers and waveguides are the sources and the targets are photodetectors. Coupling light from monomode sources into a PIFSO system and outcoupling it toward waveguides or fibers can be achieved by techniques as described in earlier publications [32], [46]. Waveguides and fibers are multimode with relatively large crosssections; they act as extended light sources, which emit incoherent radiation. Therefore, the coupling from a waveguide into a PIFSO system is more challenging, because the waveguide or the light guiding beak (LGB) have a relative large crosssection (approximately $100\ \mu\text{m}$) and numerical aperture (NA 0.3). For this task, an approach called pupil division using several lenses within the pupil is envisioned to steer the light emerging from the waveguide (see the right side of Fig. 7).

3) *Optical Design:* The optical realization of the CPU-PIFSO system may be achieved by a microchannel or hybrid imaging system [46]. As described before the PIFSO interface at the memory side is much more interesting and is discussed here in more detail. A memory-PIFSO system satisfying all the demands of the final application, but bridging a smaller distance, has been designed, simulated, and manufactured. It is based on a microchannel system. The fanout of four is achieved by a diffractive beam splitter, which is combined with an imaging lens in a single diffractive optical element (DOE). The lens has the task to collimate the divergent light of the source (VCSEL). For each channel, one more reflective imaging DOE (decreasing the tilt angle), a fanin DOE and two mirrors are used. Fig. 8 shows the paths of the light signals resulting from a ray-tracing simulation. The realization of the described design for six channels is shown in the center of Fig. 7. This picture includes the memory-PIFSO system and




Fig. 7. CENTER: The surface of the substrate with the elements for the memory-PIFSO system on the left and nine DOEs for pupil division on the right. LEFT TOP: three output DOEs consisting of two lens elements in each column. These two lens elements realize the fanin. LEFT BOTTOM: six input DOEs with four times fanout and focusing properties. The top row is used for memory bank 0 and the bottom one for memory bank 1. RIGHT: Blowup of one lens used for pupil division.




Fig. 8. 3-D ray-tracing simulation of the memory-PIFSO. An incoming collimated beam hits the splitting DOE and is divided into four parts. For simple description the propagation of only one beam (connection toward a waveguide) is calculated. The off-axis lenses are represented by a white-framed circle. The top picture shows an incoming beam from the left memory bank source, the one on the bottom represents a beam from the right memory bank source.

some lenses to test pupil division. In the center two rows with six elements correspond to the splitting DOE. Symmetrical to both sides, the reflective tilt changing DOE and the fanin DOE are arranged.

A detailed overview of the interconnections between and the arrangement of memory-MCM, PCB-integrated waveguides, and PIFSO is given by Fig. 9. Mounting these elements is another technological challenge that has to be solved in this project.




Fig. 9. Schematic view of the arrangement of PCB, MEM-MCM and memory-PIFSO system showing four channels. The ray trace for incoupling the fiber is shown by the pink and dark red lines toward the left MT-connector. Toward the right, waveguides the red and orange colored lines represent the light propagating through the PIFSO. The colored lines ending at the mirrors point out that the input element in the center is a four times beam splitter.




Fig. 10. Concept of PCB integrated optical waveguides fabricated with 45° micromirrors for coupling.

V. PCB INTEGRATED OPTICAL WAVEGUIDES

Optical interconnection technology based on board-integrated optical channel waveguides enables on-board data rates of several gigabytes per second [9], [41] over distances of up to one meter. As depicted in Fig. 10, these waveguides are embedded within an optical layer of a conventional multilayer PCB and, thus, represent an extension of today’s conventional PCB technology. A benefit of this technology is that it allows the same degree of freedom in routing optical interconnects as routing conventional electrical interconnects.

Within the HOLMS system, this technology is applied to realize high-speed optical data channels at board level between CPU and memory chips mounted on MCMs. Within the optical layer of the PCB, the optical signal is guided by multimode channel waveguides. The coupling from the waveguides to the top of the PCB and vice versa is carried out by so called LGB, explained in Section V-C, and 45° micromirrors that are integrated into the channel waveguides. In addition to technology development of board integrated optical waveguides, a simulation software is developed in the HOLMS project in order to predict the signal integrity of such optical interconnects.

A. Interconnection Topology

Within the HOLMS system, two CPU-OE-MCMs and four MEM OE-MCMs have to be connected by high-speed optical

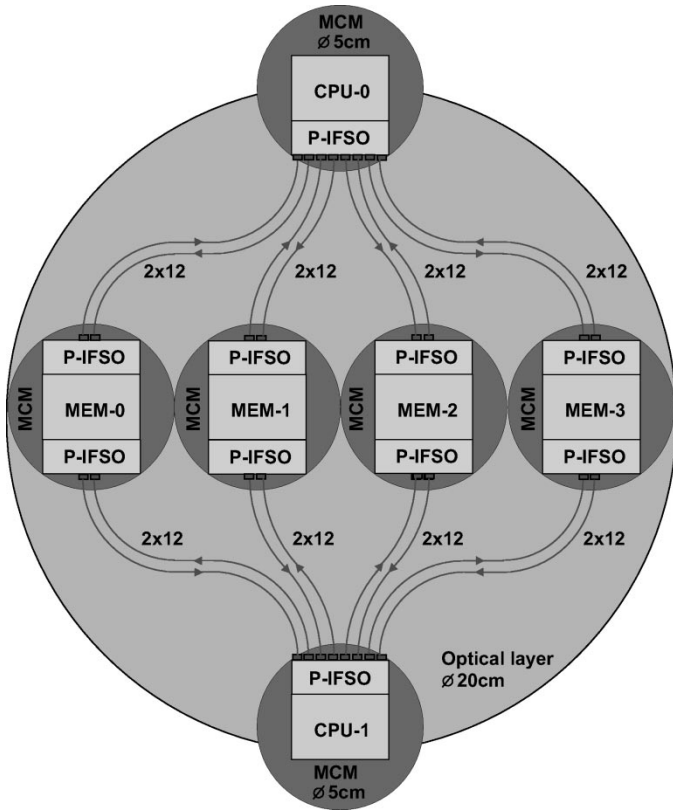


Fig. 11. Placement of MCMs and routing of PCB integrated waveguides.

data channels. Each of the high-speed optical data channels connecting one CPU and one memory MCM has to provide a total data rate of a minimum of 7.2 Gb/s per direction. This data rate has to be achieved over a maximum distance of 15 cm with a maximum loss of 0.2 dB/cm at an optical wavelength of 850 nm. This limited distance ensures a maximum optical signal skew of less than 1 ns, while one CPU chip may communicate with more than one memory chip at once. Due to the current manufacturing process, the maximum diameter of the optical layer is limited to 20 cm, which defines the routing space for the eight optical channels. Optical crossing of the channels is to be avoided in order to reduce losses. Since the system architecture requires bidirectional data flow between the CPUs and the memory banks, two physical channels (one for each direction) need to be used for each logical channel.

In order to avoid optical channel crossing and obtaining almost equal signal propagation delays with a small signal skew (≤ 1 ns), the symmetric routing layout depicted in Fig. 11 was developed in the HOLMS project. Each of the unidirectional channels shows an s-bend formation with a minimum radius of 2.5 cm and consists of 12 parallel highly multimode optical step-index waveguides. It allows a total data rate of 12×600 MB/s for every CPU memory interconnections which meets the required data rate of 7.2 GB/s. In order to be compatible to commercially available vertical cavity surface emitting laser- and photo diode arrays, the waveguides of one channel have a pitch of 250 μm . Their crosssectional sizes of 75×75 μm is comparable to those of microstrip lines. The large crosssections of the optical interconnects are chosen to be compatible to the conventional

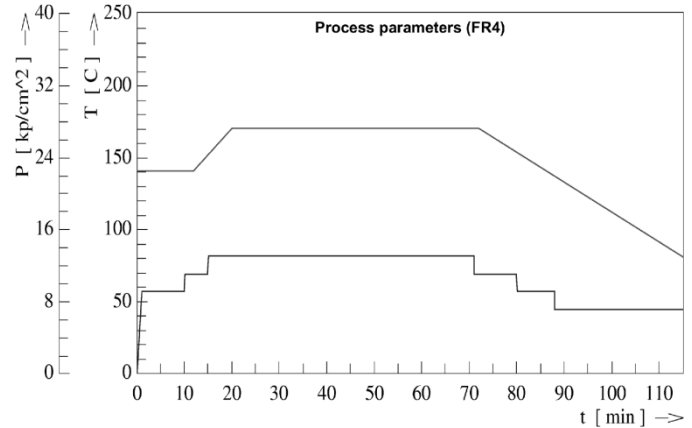


Fig. 12. Temperature (upper graph) and pressure (lower graph) profile during the lamination process of the optical layer.

PCB technology, allowing reduced tolerance requirements and simplifying the light coupling into the waveguides and vice versa. Further simplification of the coupling results from the large crosssectional sizes of the waveguides and the large numerical aperture of $\text{NA} = 0.3$.

B. Technology

The manufacturing of board integrated optical polymer waveguides is subdivided into two main processing steps. Within the first step, an optical polymer layer, containing the embedded channel waveguides, is fabricated, and within the second step, this optical layer is laminated into the multilayer PCB. Several processing technologies are known to manufacture the step-index optical channel waveguides, respectively, hot embossing [24], [25], photo lithographical techniques [2], and direct writing techniques [16]. Each of these techniques can produce a polymer optical layer with integrated multimode optical waveguides. One common property of these manufacturing processes is that the dielectric interfaces of the fabricated waveguides show surface roughness up to one tenth of the optical wavelength that causes unavoidable losses. Within the subsequent extended PCB lamination process, this optical layer is integrated into a multilayer PCB. Therefore, a recess is molded into a FR-4 layer of the PCB and the optical layer is inserted. Afterwards, the FR-4 layer is laminated into the PCB using a similar temperature and pressure profile as depicted in Fig. 12.

C. OE-MCM Waveguide Coupling: LGB

The coupling the board-integrated waveguides and the PIFSO integrated in the OE-MCMs is one of the key packaging problems of the HOLMS system. The vertical distance that has to be bridged between the PIFSO and the waveguides, is in range of 1.5 to 2 mm. Compared with the horizontal pitch between elements of the PIFSO and the waveguide of 250 μm , this vertical distance is very large. In addition, to avoid excessive losses stringent constraints on the horizontal and angular alignment, accuracy must be satisfied. At the same time, the optomechanical coupling mechanism must be




Fig. 13. Assembly concept of the PCB, the LGB, and the MCM.

compatible with standard electronic attachment technology (soldering, conductive gluing, etc.).

1) *LGB Approach*: The coupling from the PIFSO to the waveguides can be realized by free-space optics as the PIFSO provides focusing elements. To this end, the focusing elements have to be aligned in a precise manner that the angle deviation between the focused beam and the vertical axis is smaller than 1° for a vertical distance of 2 mm.

The coupling from the waveguides to the PIFSO cannot be achieved by free-space optics, as there is no possibility to integrate focusing elements within the waveguide. Without these elements, the large numerical aperture of the waveguides combined with the long vertical coupling distance makes it impossible to couple the optical power efficiently and causes significant optical cross talks and power losses. In order to increase the coupling efficiency, the LGB approach, as depicted in Fig. 13, was developed within the HOLMS project. This approach represents an imaging system based on vertical aligned multimode step-index fibers embedded within the LGB. The vertical coupling distance can, thus, be reduced to the height of the cladding material (≤ 50 m) between the waveguides and the fiber as pointed out by Fig. 13. The optical power flow within the waveguides is turned round by 45° micromirrors, transmits vertically through the cladding material and is coupled into the fiber. The multimode fiber guides the power flow to the top of the PCB and then the power flow is coupled into the PIFSO.

2) *Technology*: The LGB is built up by a synthetic block using conventional FR-4 PCB material. As each optical data channel is built up by twelve parallel channel waveguides, the LGB contains twelve fibers as depicted in Fig. 14. Therefore, it shows twelve drilling holes with a pitch of $250 \mu\text{m}$ and a diameter of $150 \mu\text{m}$. The height of the block can be matched to the vertical distance, that has to be bridged between the PIFSO and the waveguides. Within each drilling hole, one step-index multimode fiber with a core diameter of $100 \mu\text{m}$ and an overall diameter of $140 \mu\text{m}$ is glued in. Afterwards, the top and bottom face surface is grinded and polished.




Fig. 14. Concept of the LGB.




Fig. 15. LGB manufactured by ILFA using FR-4 material.

One of the manufactured LGBs is shown in Fig. 15. The fibers have a larger core diameter than the waveguides in order to reduce horizontal tolerance requirements and to improve the coupling efficiency to the fiber. The horizontal alignment of the fibers and the waveguides is achieved using MT-pins at the LGB and drilling holes within the optical layer. This technique permits a passive adjustment of the LGB to the waveguides.

VI. CONCLUSION AND FUTURE WORK

The HOLMS project aims to make the use of board-level optical interconnection in information systems practical and economical by developing optoelectronic packaging technology compatible with standard electronic assembly processes. The technology enables seamless integration of PIFSO, waveguides, and fiber arrays into high-performance electronic systems. To this end, a novel OE-MCM concept has been developed that adds PIFSO systems as just another layer interconnection layer to high-performance MCM-D electronic packages. In addition, a coupling technique that allows optical signals to cross from the OE-MCM package to PCB integrated waveguides has been introduced.

Currently, different aspects of the above technology are being optimized, fabricated, and tested by the various groups involved with the project. Key concerns being address include reducing losses, alignment tolerance, thermal control, and the optimization of the overall system assembly sequence.

Once completed, the development will provide information and computing systems with high chip I/O, high bandwidth, and

low-latency connections. This is crucial to enable novel computer system architectures focused on using the fundamental advantages of optical interconnection to cut through the traditional electronic interconnection hierarchy. To demonstrate the potential of our approach, we address the most pressing problem of contemporary computer architecture: memory latency. To this end, a fully functional multiprocessor signal processing system, in which the optical interconnection reduces the memory latency to the range of raw chip latency and signal propagation time, has been designed as the overall project demonstrator.

REFERENCES

- [1] M. G. Forbes, S. J. Fancy, M. R. Taghizadeh, J. A. B. Dines, C. R. Stanley, G. Pennelli, A. R. Boyd, P. Horan, D. Byrne, J. Hegarty, S. Eitel, H.-P. Gauggel, K.-H. Gulden, A. Gauthier, P. Benabes, J.-L. Gutzwiller, A. C. Walker, M. P. Y. Desmulliez, and M. Goetz, "Design and construction of an optoelectronic crossbar switch containing a terabit per second free-space optical interconnect," *IEEE J. Select. Topics Quantum Electron.*, vol. 5, pp. 236–249, Mar.-Apr. 1999.
- [2] B. T. Lee and M. S. Kwon *et al.*, "Fabrication of polymeric large-core waveguides for optical interconnects using a rubber molding process," *IEEE Photon. Technol. Lett.*, vol. 12, pp. 62–64, Jan. 2000.
- [3] M. Brunfaut, W. Meeus, J. Dambre, R. Bockstaele, J. Van Campenhout, H. Melchior, J. Hall, A. Neyer, P. Heremans, J. Van Koetsem, R. King, H. Thienpont, L. Vanwassenhove, and R. Baets, "Demonstrating pof based optoelectronic interconnect in a multi-fpga prototype system," in *Proc. 27th Eur. Conf. Opt. Commun.*, Amsterdam, The Netherlands, Oct. 2001, pp. 298–299.
- [4] V. Cuppu and B. Jacobb, "Concurrency, latency or system overhead: which has the highest impact on uniprocessor dram-system performance," in *Proc. 28th Annu. Symp. Comput. Architect.*, Goeteborg, Sweden, June 2001.
- [5] P. Derr, R. Hofmann, B. Greiner, D. Chiarulli, S. Levitan, and M. Robinson, "Demonstration and characterization of multi-channel interconnections using imaging fiber bundles butt coupled to optoelectronic devices," *Appl. Opt.*, vol. 39, no. 5, p. 698, 2000.
- [6] M. H. Ayliffe, F. A. P. Tooley, D. V. Plant, A. G. Kirk, F. Lacroix, and E. Bernier, "Implementation of a compact, four-stage, scalable optical interconnect for photonic backplane applications," *Appl. Opt.*, vol. 41, no. 8, p. 1541, Mar. 2002.
- [7] A. Ghose, "OPTIMUL: a hybrid multiprocessor for distributed and shared memory multiprocessing with WDM optical fiber interconnections," in *Parallel Architectures and Languages Europe*. Munich, Germany: Springer-Verlag, 1994.
- [8] E. Griesse, "Optical interconnections on printed circuit boards," *Proc. SPIE Opt. Comput.*, vol. 4089, 2000.
- [9] —, "A high-performance hybrid electrical-optical interconnection technology for high-speed electronic systems," *IEEE Trans. Adv. Packag.*, vol. 24, pp. 375–383, Aug. 2001.
- [10] M. Gruber, S. Sinzinger, and J. Jahns, *Planar-Integrated Opt. Vector-Matrix Multiplier*, vol. 39, p. 5367, 2000.
- [11] E. Yuceturk, P. J. Marchand, S. C. Esener, V. H. Ozguz, Y. Liu, G. Li, and D. Huang, "Three-dimensional optoelectronic stacked processor by use of free-space optical interconnection and three-dimensional VLSI chip stacks," *Appl. Opt.*, vol. 41, no. 2, p. 348, 2002.
- [12] V. Baukens, H. Ottevaere, G. Verschaffelt, P. Tuteleers, P. Vynck, B. Volckaerts, A. Hermanne, I. Veretennicoff, H. Thienpont, and C. Debaes, "Plastic micro-optical interconnection modules for parallel free-space intra mcm data communication," *Proc. IEEE*, vol. 88, pp. 769–779, June 2000.
- [13] J.-H. Ha and T. M. Pinkston, "SPEED DMON: Cache coherence on an optical multichannel interconnect architecture," *J. Parallel Distributed Comput.*, vol. 41, no. 1, pp. 78–91, 1997.
- [14] H. J. Haumann, H. Kobolla, F. Suer, J. Schmidt, J. Schwider, W. Stork, N. Streibl, and R. Vikel, "Optoelectronic interconnection based on a light-guiding plate with holographic coupling elements," *Opt. Eng.*, vol. 30, 1990.
- [15] D. C. Hoffmeister, J. Chu, J. A. Perreault, and P. Dowd, "Lightning network and systems architecture," in *Parallel Computing Using Optical Interconnections*, K. Li, Y. Pan, and S. Q. Zheng, Eds. Norwell, MA: Kluwer, 1998, pp. 4–21.
- [16] H. Huber, J. Moisel, and J. Guttmann, "Polymer waveguides for 100 cm (40") optical backplanes," *Proc. SPIE*, pp. 4805–4815, 2002.
- [17] J. Depreitere, J. Van Campenhout, H. Van Marck, and J. Dambre, "Optoelectronic FPGAs," *IEEE J. Select. Topics Quantum Electron.*, vol. 5, pp. 306–315, Mar.-Apr. 1999.
- [18] J. Van Campenhout, C. Jesshope, M. Desmulliez, H. Thienpont, J. Goodman, J. H. Collet, D. Litaize, and A. Louri, "Architectural approach to the role of optics in monoprocessor and multiprocessor machines," *Appl. Opt.*, vol. 39, no. 5, 2000.
- [19] J. Jahns, "Planar packaging of free-space optical interconnections," *Proc. IEEE*, vol. 82, pp. 1623–1631, Nov. 1994.
- [20] J. Jahns and A. Huang, *Planar Integration Free-Space Opt. Components*, vol. 28, pp. 1602–1605, 1989.
- [21] K.-Y. J. Li and K. Jenkins, *Opt. Parallel-Access Shared-Memory Syst.: Anal. Exper. Demonstr.*, vol. 34, no. 2, p. 358, 1995.
- [22] Y. Li and J. Ai, "Polymer fiber-image-guide-based embedded optical circuit board," *Appl. Opt.*, vol. 38, no. 2, p. 325, 1999.
- [23] A. G. Kirk, D. V. Plant, T. H. Szymanski, Z. G. Vranesic, J. A. Trezza, F. A. P. Tooley, D. R. Rolston, M. H. Ayliffe, F. Lacroix, D. Kabal, B. Robertson, E. Bernier, D. F. Brosseau, F. S. J. Michael, and A. Chuah, "A modulator-based multistage free-space optical interconnection system," *Proc. Opt. Comput.*, pp. 449–459, 2000.
- [24] D. Krabe and W. Scheel, "Optical interconnects by hot embossing for module and PCB technology," in *Proc. 49th Electron. Components Technol. Conf.*, San Diego, CA, 1999, pp. 1164–1166.
- [25] S. Lehmachner and A. Neyer, "Integration of polymer optical waveguides into printed circuit boards," *Electron. Lett.*, vol. 36, no. 12, pp. 1052–1053, 2000.
- [26] A. Ahmed Louri and B. Brent Weech, "Scalable optical interconnection networks for large-scale parallel computers," in *Parallel Computing Using Optical Interconnections*, K. Li, Y. Pan, and S. Q. Zheng, Eds. Norwell, MA: Kluwer, 1998, pp. 48–74.
- [27] P. Lukowicz, "Design of an efficient shared memory architecture using hybrid opto-electronic VLSI circuits and space invariant optical busses," in *Proc. 3rd Int. Conf. Massively Parallel Process. Using Opt. Interconnect.* Maui, Hawaii, Oct. 1996, pp. 231–238.
- [28] —, "Giga-cache: A high performance opto-electronic memory architecture," in *Proc. Workshop Critical Technol. Comput., SPIE Annu. Meeting*, San Diego, CA, 2000.
- [29] —, "Reducing communication and memory latency in computer systems using opto-electronic interconnections," in *Proc. Progress Electromagnet. Res. Symp.*, Boston, 2000, p. 320.
- [30] P. Lukowicz, S. Sinzinger, K. Dunkel, and H.-D. Bauer, "Design of an opto-electronic VLSI/parallel fiber bus," *J. Eur. Opt. Soc. Part-A*, pp. 367–370, 1999.
- [31] M. Dal Cinand, W. Hohland, T. Dalibor, S. Eckert, A. Hessenauer, H. Grygier, U. Hildebrand, J. Hnig, F. Hofmann, C.-U. Linster, E. Pataricza Michel, A. Sieh, V. Thiel, and T. Turowski, "Architecture and realization of the modular expandable multiprocessor system memsy," in *Proc. 1st Int. Conf. Massively Parallel Comput. Syst.*, 1994, pp. 7–15.
- [32] J. Jahns, S. Sinzinger, M. Gruber, and E. ElJoudi, "Practical realization of massively parallel fiber-free-space optical interconnects," *Appl. Opt.*, vol. 40, pp. 2902–2908, 2001.
- [33] D. R. Rolston, A. G. Kirk, D. V. Plant, M. H. Ayliffe, and M. Teauneuf, "Six-degrees-of-freedom alignment of two-dimensional array components by use of off-axis linear fresnel zone plates," *Appl. Opt.*, vol. 40, no. 35, p. 6515, 2001.
- [34] P. Milojkovic, J. Ekman, P. Chandramani, R. Rozier, F. Kiamilev, Y. Liu, M. Hibbs-Brenner, M. W. Haney, and M. P. Christensen, "Multichip free-space global optical interconnection demonstration with integrated arrays of vertical-cavity surface-emitting lasers and photodetectors," *Appl. Opt.*, vol. 38, no. 29, p. 6190, 1999.
- [35] D. A. B. Miller, "Optical interconnects to silicon," *IEEE J. Select. Topics Quantum Electron.*, vol. 6, pp. 1312–1317, Nov.-Dec. 2000.
- [36] D. A. B. Miller, "Optical interconnects to silicon," *IEEE J. Select. Topics Quantum Electron.*, vol. 6, pp. 1312–1317, Nov.-Dec. 2000.
- [37] A. G. Nowatzky and P. R. Pruncal, "Are crossbars really dead? The case for optical multiprocessor interconnect systems," in *Proc. 22nd Int. Symp. Comput. Architect.*, Gold Coast, Italy, 1995.
- [38] V. Ozguz, J. Carson, M. Hibbs-Brenner, Y. Liu, F. Kiamilev, P. Marchand, S. Esener, and L. Coldren, "3-D-optoelectronic stacked processors: Technology development," in *Proc. Opt. Comput.*, Apr. 1999.
- [39] R. Barbieri, G. Tröster, S. Fancey-Matthias, G. Jurgen, J. W. F. Tichy, P. Lukowicz, and J. Grzyb, "Making optical interconnection packaging compatible with electronic assembly technology," *Opt. Memory Neural Networks*, vol. 4, 2002.

[40] T. M. Pinkston, "Design considerations for optical interconnects in parallel computers," in *Proc. 1st Int. Workshop Massively Parallel Process. Using Opt. Interconnect.*, 1994, p. 151.

[41] C. Choi, R. Chen, and L. Lin, "Fully embedded board-level guided wave optoelectronic interconnects," *Proc. IEEE*, vol. 88, pp. 780–792, June 2000.

[42] H. Matsuoka, S. Nishimura, H. Inoue, and T. Yokota, "Optical interconnection subsystem used in the RWC-1 massively parallel computer," *IEEE J. Select. Topics Quantum Electron.*, vol. 5, pp. 360–367, Mar.–Apr. 1999.

[43] T. Sakano, T. Matsumoto, K. Noguchi, and T. Sawabe, *Design Perform. Multiprocess. Syst. Employing Board-to-Board Free Space Opt. Interconnect.*: *Cosine-1*, vol. 30, no. 17, pp. 2334–2342, 1991.

[44] A. Saulsbury, F. Pong, and A. Nowatzky, "Missing the memory wall: The case for processor/memory integration," in *Proc. 23rd Annu. Int. Symp. Comput. Architect.*, May 22–24, 1996, pp. 90–103.

[45] S. Shibayama, K. Hamaguchi, T. Fukui, Y. Sudo, T. Shimoyama, and S. Nakamura, "An optical bus computer cluster with a deferred cache coherence protocol," in *Proc. Int. Conf. Parallel Distributed Syst.*, June 3–6, 1996, pp. 175–182.

[46] S. Sinzinger and J. Jahns, *Integrated Microopt. Imag. Syst. High Interconnect. Capacity Fabricated Planar Opt.*, vol. 39, pp. 4729–4735, 1997.

[47] J. Grzyb, S. Fancey, J. Jürgen Jahns, and P. Paul Lukowicz, "Low latency opto-electronic processor memory interconnection demonstrator," in *Proc. SPIE Annu. Meeting*, San Diego, CA, July 2001.

[48] C. Svensson. (2001) Electrical Interconnects Revitalized. Linköping University. [Online] Available: <http://www.ek.isy.liu.se/christer/CS Papers.htm>

[49] T. H. Szymanski, "Parallel computing with 'intelligent optical networks'," in *Parallel Computing Using Optical Interconnections*, K. Li, Y. Pan, and S. Q. Zheng, Eds. Norwell, MA: Kluwer, 1998, pp. 26–44.

[50] D. Huang, O. Kibar, N. S. E. Ozkan, S. C. Esener, X. Zheng, and P. J. Marchand, "Optomechanical design and characterization of a printed-circuit-board-based free-space optical interconnect package," *Appl. Opt.*, vol. 38, no. 26, p. 5631, 1999.

[51] H. Kosaka, S. Kawai, K. Kasahara, Y. Li, and T. Wang, *Fiber-Image-Guide-Based Bit-Parallel Opt. Interconnects*, vol. 35, no. 35, p. 6920, 1996.

[52] E. E. E. Frietman and F. Zhao, "Optical multi-faceted free-space image distributor for massively parallel processing," in *Proc. MPOI* Las Vegas, NV, June 1998.



Paul Lukowicz (M'95) was born in Krakow, Poland, in 1967. He received the M.Sc. in computer science, the M.Sc. in physics, and the Ph.D. in computer science, from the University of Karlsruhe, Karlsruhe, Germany, in 1992, 1993, and 1999, respectively.

Since 1999, he has led the Computer Architecture and Wearable Computing Group, Electronics Laboratory, Swiss Federal Institute of Technology, Zurich, Switzerland. He has been working on the applications of optical interconnections in novel computer architectures since 1995, focusing on low-latency memory

system. He has been among the initiators and is currently the Scientific Coordinator of the HOLMS project. His research interests also include innovative mobile computer architectures and wearable sensor systems.

Dr. Lukowicz is a member of the Association for Computing Machinery and the Optical Society of America.



Jürgen Jahns (M'87) was born in Erlangen, Germany, in 1953. He received the diploma and doctorate in physics from the University of Erlangen-Nürnberg, Erlangen-Nuernberg, Germany, in 1978 and 1982, respectively.

From 1983 to 1986, he was with Siemens AG, Munich, Germany, where he worked on robotics, sensors, and optical fiber communications. From 1986 to 1994, he was a Member of the Technical Staff in the Optical Computing Research Department, AT&T Bell Laboratories, Holmdel, NJ. During that time, he worked on optical interconnections, diffractive optics, and microoptic packaging. Since 1994, he has been a Full Professor of Optical Information Technology at the University of Hagen, Hagen, Germany. He is the author of some 60 journal and more than 150 conference publications and holds 16 international patents. He has contributed several book chapters, co-edited a book "Optical Computing Hardware" (New York: Academic, 1994), co-authored a book "Microoptics" (New York: Wiley-VCH, 2nd ed., 2003) and authored a book "Photonik" (Munich, Germany: Oldenbourg, 2000, in German). Prof. Jahns is a member of the German Society for Applied Optics (DGaO), where he leads a Technical Group on Microoptics, as well as a Fellow of the Optical Society of America.

P. Barbieri, photo and biography not available at time of publication.

R. Benabes, photo and biography not available at time of publication.

Th. Bierhoff, photo and biography not available at time of publication.

A. Gauthier, photo and biography not available at time of publication.

M. Jarczyński, photo and biography not available at time of publication.

G. A. Russell, photo and biography not available at time of publication.

J. Scharge, photo and biography not available at time of publication.

W. Süllau, photo and biography not available at time of publication.

J. F. Snowdon, photo and biography not available at time of publication.

M. Wirz, photo and biography not available at time of publication.

G. Tröster, photo and biography not available at time of publication.