RZ 3658 (# 99678) 08/21/2006
Electrical Engineering 3 pages

Research Report

High bandwidth board-level parallel optical interconnects for server applications

B.J. Offrein, C. Berger, R. Beyeler, R. Dangel, L. Dellmann, F. Horst, T. Lamprecht, N. Meier

IBM Research GmbH
Zurich Research Laboratory
8803 Rüschlikon
Switzerland

J. Kash

IBM Research
T.J. Watson Research Center
Yorktown Heights, N.Y. 10598
USA

LIMITED DISTRIBUTION NOTICE

This report has been submitted for publication outside of IBM and will probably be copyrighted if accepted for publication. It has been issued as a Research Report for early dissemination of its contents. In view of the transfer of copyright to the outside publisher, its distribution outside of IBM prior to publication should be limited to peer communications and specific requests. After outside publication, requests should be filled only by reprints or legally obtained copies (e.g., payment of royalties). Some reports are available at http://domino.watson.ibm.com/library/Cyberdig.nsf/home.

Research
Almaden • Austin • Beijing • Delhi • Haifa • T.J. Watson • Tokyo • Zurich

High bandwidth board-level parallel optical interconnects for server applications

B.J. Offrein (1), C. Berger (1), R. Beyeler (1), R. Dangel (1), L. Dellmann (1),
F. Horst (1), T. Lamprecht (1), N. Meier(1), J. Kash (2)
1 :IBM Research GmbH, Säumerstrasse 4, 8803, Rüschlikon, Switzerland, ofb@zurich.ibm.com
2 : IBM Research, T.J. Watson Research Center, jeffkash@us.ibm.com

Abstract Optical waveguides embedded in a printed circuit board offer a substantial increase in the achievable bandwidth density compared with todays electrical interconnects. Specific challenges and solutions to integrate optical interconnects in a printed circuit board are addressed.

Introduction

The success story of optical communication is directly linked with the availability of high-quality optical fibers that enable the transport of tremendous amounts of data over large distances with low attenuation. The invention of the erbium-doped fiber amplifier boosted the application of fiber-based optical communication in long distance links as it enables a direct amplification of the optical data and supersedes the use of electronic repeaters. This makes a large-distance optical link highly cost-effective compared with an electronic link. The fundamental advantage of optics as compared to electrical communication is its higher bandwidth-distance product, which explains why optical interconnects were first applied in large-distance links. With increasing bandwidth requirements, optics also can penetrate into shorter-reach applications [1]. This trend is clearly visible in the optical revolution that started about 25 years ago. Today, we see that optical interconnects handle most of the large-bandwidth inter-system data-exchange but the intra-system communication is still performed in the electrical domain [2]. Future computing systems for high-performance applications such as number-crunching, biological modeling, or climate modeling typically consist of multiple racks, where each rack has a backplane or midplane, and multiple plug-in cards (or blades). The increasing clock speed and data throughput of the processors and switches that are used in these systems entail an increased bandwidth demand on the intra-system communication. In the intra-system link hierarchy, it is the card-to-card interconnect that could first evolve to optics. The link distance is between 30 cm and 1 m and is shown schematically in Figure 1. Besides the distance arguments for optics as discussed above, it is also the high signal density that drives the development of optical board technology.

The introduction of optical technology on or in the printed circuit board demands the development of various elements as can be derived from Figure 1.

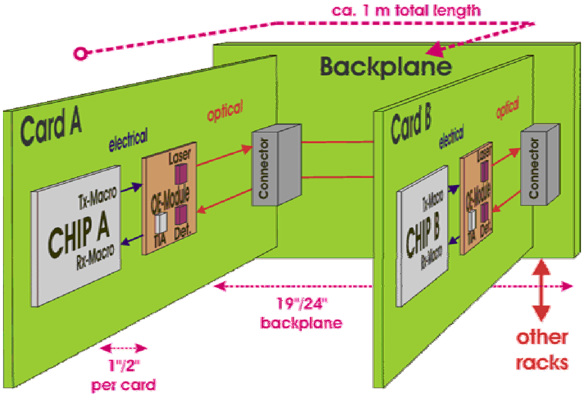

Fig. 1 Card-to-card optical interconnect.

These elements are:

1. Optical transport layers in the printed circuit board (PCB) to transport and distribute the optical signals.
2. Transceiver elements that perform the optoelectronic conversion.
3. Optical coupling solutions to transfer the optical power between the transceiver and the optical layers in the board.
4. Optical board-to-backplane connectors to couple the light between the daughter card and the backplane as well as board-to-fiber connectors to distribute the optical signals over a larger distance.

In this paper, we will focus on the first three aspects.

Cost is the main challenge in the development of a parallel optical interconnect technology on a PCB but advances are also required in technical aspects such as speed, density and reliability. In order to understand just what the cost challenges are for optics, we have extrapolated the expected costs of optical vs. electrical links for future 10 Gbps per channel links, as shown in Figure 2 [3]. These estimates are based on relatively high performance computing systems which need to support very high aggregate bandwidth (and therefore high density) interconnects.

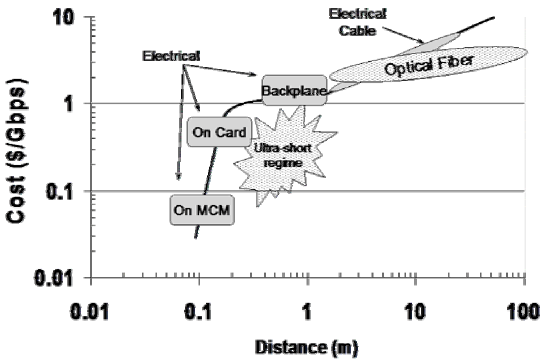

Fig. 2. Comparison of the cost (in $/Gbps) of electrical links versus optical links projected for the 10 Gbps per channel generation in high performance (high aggregate BW) computing systems.

These estimates were derived by extrapolating the known costs of electrical links out to 10 Gb/s, taking into account projected cost increases due to the required improvements in boards and cables. In some cases the electrical links may be actually composed of 2 x 5 Gb/s (e.g. single ended) links as a less expensive alternative, and it is not clear that electrical cabled links can be practically extended beyond 10m at 10 Gb/s channel rates. Similarly known costs of optical links were extrapolated, assuming, e.g. module and VCSEL yields at these speeds can reach costs and a state of maturity equivalent to today's 2.5 Gb/s parallel optics modules.

For cabled links, the slope of the optical interconnect cost curve is shallower than that for electrical links, due to the much lower cost of optical fiber over electrical cable. However, the cost of the optics transceivers is much higher than their electrical counterparts. We believe that optics will be a cost effective solution in the 10 Gb/s per channel generation for links in the few meter range. In the ultra-short link range of less than 1m (backplanes and cards) the story is much more challenging. Here the cost of the card level traces and associated packaging and connectors is very inexpensive, and electrical links drop into the $1/Gbps range for backplanes and into the 10's of cents per Gbps range for cards. This cost target represents a tremendous challenge for optics to meet.

A large product volume is a key condition for low-cost production. Therefore we develop a generic optical interconnect technology that can be applied to various interconnect classes, ranging from high-density and high-speed links for high-end servers, routers and gaming systems to lower-density and moderate-speed links as required for sensor applications, for example.

Optical printed circuit board technology

For the transport of the optical signals in the printed circuit board, we chose an integrated waveguide technology based on polymers. Multimode polymer waveguides with a size of 50 μm can be realized in a low-cost way and the alignment tolerances are much more relaxed than for single mode waveguides. Furthermore, the short distance links in intra-system communication are not limited by modal dispersion effects for the signaling speeds of 10 to 20 Gb/s that are envisioned.

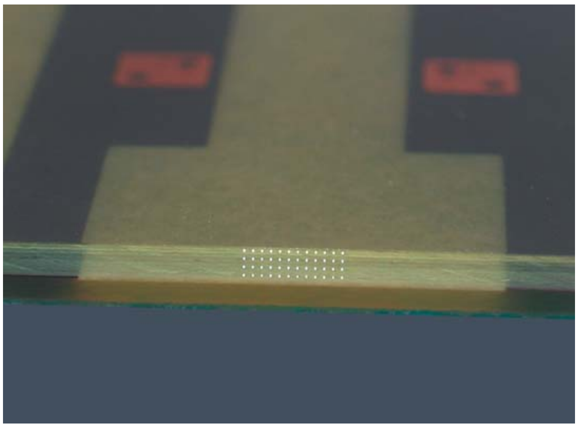

Fig. 3. Printed circuit board with a two dimensional array of polymer waveguides.

Fig. 3. shows a printed circuit board with four polymer waveguide layers. The waveguides have a rectangular shape and a size of 50 μm, the lateral and vertical pitch is 250 μm and the propagation losses are 0.05 dB/cm at 850 nm. The potential bandwidth density of this optical structure is enormous and could be increased even further by reducing the waveguide pitch. Depending on the application, the optical waveguides can be realized on top of the printed circuit board or embedded in the board stack.

Low-cost coupling solutions

The alignment of optical elements to the polymer waveguides is an essential part of the board-level optical interconnect technology. We developed a procedure that enables passive alignment of any type of optical or opto-electronic element to the waveguides. The procedure relies on common printed circuit board technology processes and tools and hence enables a low-cost realization and assembly of optical boards. Fig. 4 shows the alignment of an MT adapter to a printed circuit board with embedded optical waveguides. The prototype adapter is made out of brass using a high-precision CNC tool and can be replaced by a low-cost molded part in a product application. The adapter contains alignment studs that snap into alignment slots in copper markers in the printed circuit board. The alignment accuracy of this concept was demonstrated to be well below 8 μm such that the alignment error related coupling loss to a fiber will be less than 0.5 dB. More details can be found in [4].

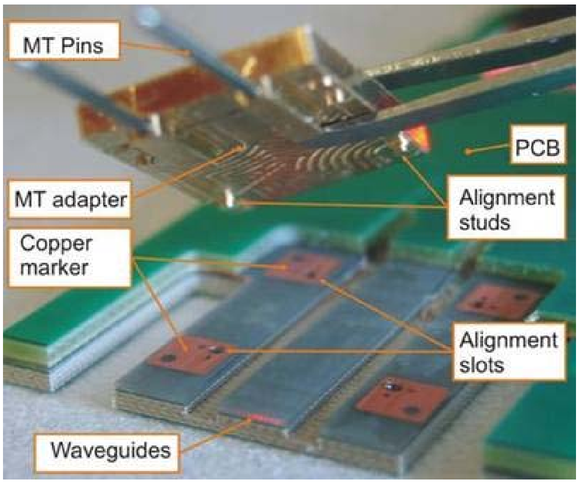

Fig. 4. Alignment of an MT adapter into slots in the printed circuit board with embedded waveguides.

Transceiver elements

The opto-electronic conversion is performed by a transceiver subassembly. Fig. 5 depicts the two basic transceiver configurations to couple the light into the polymer waveguides. The printed circuit board (green) contains one or more waveguide layers (light blue). The electrical signal (red) is converted to the optical domain (blue) with a vertical cavity surface emitting laser (purple). Depending on the configuration, there is a requirement for a 90 deg turn in the optical or in the electrical signal path.

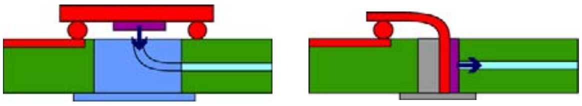

Fig. 5. Transceiver assemblies on a printed circuit board, based on a configuration with a 90 degree turn in the optical domain (left) or in the electrical domain.

At IBM, we pursue both packaging concepts. Advantages of the '90°-in-optics'-approach include better compatibility with standard packaging and assembly processes, the possibility for a clear separation of the OE-package and the passive PCB (including a hermetically sealed board surface with 'optical pads' than can be standardized and serviceable), and the option to realize large packages with multiple densely packed opto-electronic components in close vicinity of the rest of the electronics. This concept is used in the Darpa-Terabus project for the realization of an opto-electronic link with a density of 16 signal channels/mm [5].

Advantages of the 90°-in-electronics approach include the fact that an architecture with an electrical flex-circuit is already being used in commercial fiber-optic transceiver modules, which allows to build upon that experience. Since the lasers emit in the direction of the waveguides, there is no need for a mirror. As 850 nm VCSELs are usually top-emitters they cannot be directly placed on the waveguide facets, i.e. through-substrate waveguides or microlenses are still required. An extension to 2D-arrays is straightforward, at least as far as the optical coupling concept is concerned.

Experimental results on high density optical interconnect systems using polymer waveguides will be reported for both transceiver configurations.

Summary

The increasing bandwidth requirements of processors and switches drive the development of board-level optical interconnect technology. Compared to electrical signaling, optics offers a higher bandwidth-distance product and a higher bandwidth density. Especially the latter is of importance for intra-system communication. We develop a generally applicable board-level optical interconnect technology. Specifications, manufacturablity, maturity and overall cost are factors that determine the potential of optical interconnects for short distance applications.

References

1 D.A.B. Miller, Int. J. Optoelectronics, Volume 11, (1997), Page 155
2 A. Benner et al., IBM J. Res. & Dev. Volume 49 (2006), Page 755.
3 M. Taubenblatt, Photonics West (2006), Paper 6124-06
4 T. Lamprecht et al., 56 ECTC, San Diego (2006, Paper 371
5 L. Schares et al., IEEE J. Select. Topics Quantum Electron., Volume 12 (2006).