A modular, guided-wave approach to plane-to-plane optical interconnects for multistage interconnection networks

○Alvaro Cassinelli(1), Makoto Naruse(2)(3), Masatoshi Ishikawa(1) and Fumito Kubota(2)
1: Dept. Information Physics and Computing, University of Tokyo.
2: Communications Research Laboratory, Nukui-Kita, Koganei, 184-8795 Tokyo.
3: PRESTO21, Japan Science and Technology Cooperation.
E-mail: *alvaro@k2.t.u-tokyo.ac.jp*

Guided-wave based elementary interconnection modules are proposed as reusable and scalable building blocks for regularly interconnected multistage networks. As an illustration of the approach, decomposition of shuffle-exchange permutations into elementary interconnection modules is shown, and a fiber-based prototype is presented.

1. Introduction

Bandwidth, cross-talk and volume consumption considerations, all advocate for optical interconnects within massively interconnected system as a promising substitute to conventional electronic circuitry [1]. Besides, regularly interconnected multistage architectures have been long proposed as an interesting alternative to the full crossbar and the shared bus (in terms of performance, cost and scalability) for use as telecom switching fabrics and as permutation networks for multiprocessors [2].

As a consequence of the above considerations, multistage architectures using optical plane-to-plane regular interconnections may well represent a theoretically optimal architecture for use within a massively interconnected system [1,3,4]. Furthermore, since optoelectronic integration is becoming a mature technology, it is likely that such paradigm will be adopted in near-future high-performance commercial hardware.

While innumerable demonstrators have been built to illustrate the advantage of free-space optics (including planar optics) over electronics for dense two-dimensional point-to-point interconnections [5], a relatively small number of papers has been produced so far on what may seems at first sight a more challenging issue: the use of point-to-point, plane-to-plane wave-guide interconnection arrays. Still, the issue should not be put aside too quickly: wave-guides can easily achieve a higher transmission efficiency while almost canceling cross-talk; and, contrary to the common belief, guided-wave optics is *more* volume efficient than free-space optics in both space-invariant and space-variant interconnect applications [6].

Summing up, **multistage architectures using regular plane-to-plane wave-guide based interconnections may represent (today) the most advantageous *hardware implementation* of a theoretically-optimal architecture for massively interconnected systems** (possibly including optoelectronic interconnected for 3D integrated circuits [7]).

The mayor drawbacks of the guided-wave approach appears to be (1) the packaging difficulties faced when implementing an arbitrary and even regular interconnection pattern (indeed, papers on the subject present demonstrators without point-to-point permutation capacity - a function implemented when needed by additional mirrors, holograms, etc [8,9,10]), along with the fact that (2) a specific interconnection module should be completely redesigned when scaling the number of channels, thus making the approach far from being cost-effective, and (3) the fact that good efficiency can only be achieved if *double* care is put on alignment, since both inputs (to) and outputs (from) the wave-guide interconnection array must precisely match the receiver *and* transmitter array geometries, in contrast with their holographic-based interconnection counterpart. In this paper we will concentrate on ways to overcome the first two obstacles. Alignment issues are discussed elsewhere [11].

2. Elemental modules

We will call **permutation module** a 3D structure consisting on two rectangular arrays representing inputs and outputs, interconnected trough point-to-point wave-guides and **folding** (or packaging) its implementation process. Complications arise from the fact that although many multistage networks use simple and regular interconnections, when folded, these materialize as non-regular, non-scalable and non-reusable modules. Because it should be possible to **cascade modules without the need of additional optics**, we may want to decompose a permutation as a product of easierto-fold permutations. We may also want these **elemental modules** to be of use when addressing the folding problem for a different permutation (**permutation reusability**), as well as being handy when scaling up the interconnection (**scaling reusability**). An interesting decomposition consists in separating the row and column permutation parts from the non-decomposable (or 'diagonal') part (Fig.1). This is appealing because the row and column permutations can be straightforwardly folded by simple stacking layers of laminated fibers or printed light-wave circuits for instance, and because row or column permutations are essentially plane topologies (only replicated along the third dimension), and can in principle achieve good scaling-reusability, by recursively building larger modules out of smaller ones [6].

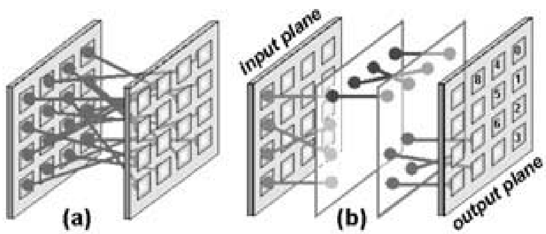

Figure 1. Folding the shuffle $\sigma_{(4)}$: (a) brute folding, (b) decomposed folding $\sigma_{(4)} = \sigma_{col}(2)\ \sigma_{row}(2)\ d(4)$.

In this presentation, we will focus on the particular class of 'regular' interconnections whose action actually corresponds to a permutation of the channel address bits, and/or a 'swap' of some of its bit values. Inter-stage permutations used in blocking Delta networks (in particular the shuffle-exchange networks) all fall into this category, as well as the permutations used in rearrangeable networks like Clos and Benes [12]. We will show then that for all of them, the permutation that exchanges the least significant bits of the column and the row address of a particular channel is a suitable diagonal permutation, both easy to implement and extremely scaling-reusable.

3. Module prototype

To experimentally validate the approach, a 16 input/output regular permutation was folded into a fiber-based 4x4 module [13]. Using similar modules, a butterfly network topology is being built to be used in a real-time pipelined processing system (Fig.2). Static and dynamic alignment of modules is being explored as well [14], and results will be presented soon in another communication.

4. Conclusion and further research

Multistage architectures using these guided-wave modules may find a very large range of applications. Thanks to the flexibility of the approach (scalability and modifiability), it will be easy to build several hybrid-optoelectronic demonstrators based on programmable SIMD arrays developed in our lab for real-time vision applications [15,16]. In the short term, we would like to demonstrate specialized real-time parallel data processing systems as well as packet switched multistage fabrics for networking applications. Design and fabrication of the necessary modules is underway.

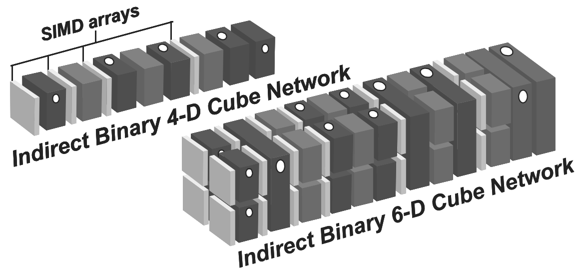

Figure 2. Module scaling-reusability at work on the 'butterfly' network (16x16 and 64x64)

A different, interesting issue worth to consider is the use of modules containing *several different, independent addressable* permutations (sharing the same folding properties). This can lead to a concurrent multistage interconnection paradigm not based on size-reduced (elemental) crossbars (as in most multistage networks) but on cascaded *permutation*-reduced crossbars instead. Promising architectures based on this paradigm as well as routing algorithms are currently being explored in our lab.

We wish that this presentation can contribute to popularize the idea of guided-wave based interconnection arrays in the optoelectronic hardware designer's mind.

References

[1] H.M.Ozaktas: Appl. Opt. **36** (1997) 5683.
[2] J.Duato et al.: *Int. Net.*, IEEE Comp. S. Press (1997).
[3] N.McArdle et al. : IEEE JSTQE **5** (1999) 250.
[4] H.M.Ozaktas: Appl. Opt. **36** (1997) 5697.
[5] *MEL-ARI OPTO REPORT*, H.Neefs Ed. (2000).
[6] Y.Li, J.Popelek: Appl. Opt. **39** (2000) 1815.
[7] D.Ludwig et al.: SPIE Chen&Guilfoyle Eds.(1996) 366.
[8] A.Goulet et al.: Proc. OC2001 (2001) 122.
[9] J.D.Bakos et al.: Proc. OC2002 (2002).
[10] A.Louri et al.: Appl. Opt. **34**, (1995) 6714.
[11] M.Naruse et al. : Appl.Opt. **39** (2000) 293.
[12] A.Varma et al.: *Int. Net.,* IEEE Comp.S.Press (1994).
[13] M.Naruse et al.: 2002 IEEE LEOS, *to appear*.
[14] M.Naruse et al., IEEE Ph.Tech. Lett. **13** (2001) 1257.
[15] H.Kawai et al.: Elect. Lett. **38** (2000).
[16] A.Cassinelli et al.: Proc. OC2002 (2002).