3D Optoelectronic Computer Architectures for the Conjugate Gradient and Multigrid Benchmark Algorithms

George A. Betzos and Pericles A. Mitkas

Optoelectronic Computing Systems Center and the Department of Electrical Engineering, Colorado State University, Fort Collins, CO 80523-1373, USA

Abstract. Optics can be used to build faster and more complex interconnection networks for parallel processors. Such systems will be composed of arrays of electronic processing elements interconnected with three-dimensional free space optics. We call this a 3D optoelectronic computer. We show that for the conjugate gradient benchmark an architecture with an adequately fast reconfigurable interconnection network can outperform all parallel supercomputers, but its performance is not as impressive when a fixed network is used. In the case of the multigrid benchmark, the 3D optoelectronic architecture can perform orders of magnitude better than the best parallel supercomputers.

1 Introduction

Free-space optical interconnects have significant advantages over electronic, due to lower power requirements, use of three-dimensional (3D) space and lack of cross-talk ([6], [4]). We expect the first systems that will make use of optical interconnects to be realized as arrays of electronic PEs with free-space optics interconnecting PEs of neighboring arrays [3]. We call this a 3D optoelectronic computer (OEC) [2]. To analyze the potential of such a system we use the NAS benchmark algorithms [1] because we can compare its performance to that of currently available parallel supercomputers [5]. This paper completes our study (see [2]) with two of the most complex benchmark algorithms, the Conjugate Gradient (CG) and the Multigrid (MG).

2 Architectures and Results

The CG computation (primarily matrix-vector multiplications) can be performed using a fixed interconnection network, but the sparse nature of the matrix is ignored. However, it can also be performed using a reconfigurable interconnection network. A 3D OEC consisting of an input/output array, a main PE array and a broadcasting/accumulator array interconnected with fixed networks, can be used to implement this algorithm. The network between the input/output array and the main PE array is straight and bidirectional. There is a unidirectional broadcasting network from the broadcasting PE to the main PE array, while

each half of the PEs across each column of the main array can transmit data to the corresponding upper or lower accumulator PE unit. Fig. 1 plots the total time. In this case a 3D OEC cannot surpass the performance of the best parallel supercomputer, mainly because it ignores the fact that the matrix is sparse. We expect that if the parallel supercomputers had to perform normal matrix-vector multiplications, they would be orders of magnitude slower that the 3D OEC.

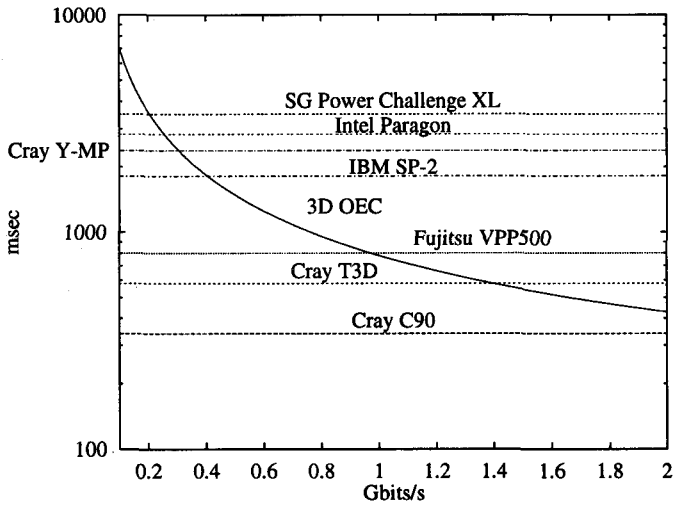

Fig. 1. Performance comparison of the 3D OEC and the best parallel electronic supercomputers on the CG benchmark

A 3D OEC consisting of an input/output array and two PE arrays interconnected with a bidirectional reconfigurable network can also implement this algorithm. In this case only the non-zero elements of the sparse matrix are processed. The total time it takes to execute this benchmark is computed in terms of t_r, the time it takes to reconfigure the network. In Fig. 2 we notice that for values of t_r less than approximately 1 μsec the 3D OEC with the reconfigurable network performs better than all the parallel supercomputers. For values of t_r greater than 10 μsec, the reconfiguration of the network becomes the bottleneck of the architecture. Since in this case the network reconfiguration is frequent, the advantages of using such a network, can be better seen when the reconfiguration occurs less frequently.

The MG benchmark, is the most complex NAS benchmark we are considering, due to its 3D nature, the many different processing stages and types of communication patterns [1]. The 3D OEC is made up of three stages. In the first stage the outside border of the array is obtained. The second PE array also performs the apply correction operation. The third PE array performs the restrict

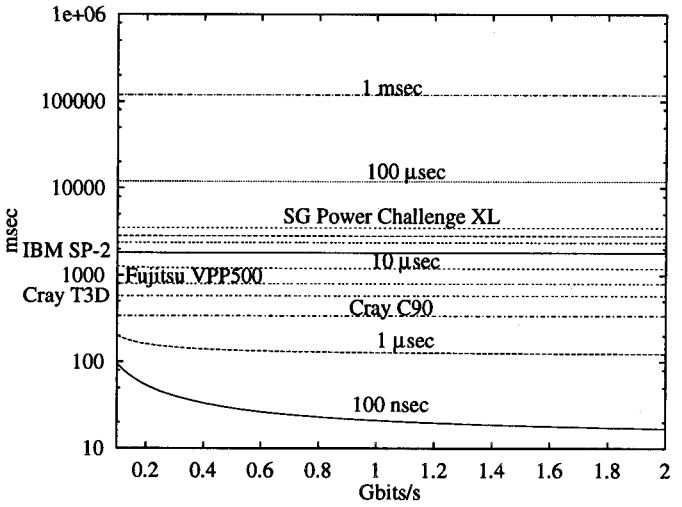

Fig. 2. Comparison of the performance of the reconfigurable 3D OEC against the best parallel supercomputers for certain values of t_r and for the CG benchmark

residual operation. The interconnection network between the second and third PE arrays basically reduces each received array of data to a quarter of its size. The fourth and last PE array performs the prolongate operation. The interconnection network between the third and fourth PE array expands (quadruples) the smaller received array of data. On the opposite direction there are straight-pass interconnection networks. We also assume that the first and the fourth PE arrays have straight pass bidirectional input/output interconnections.

Making similar assumptions as before [2], we can compute the total time it would take such a 3D OEC architecture to execute the benchmark. Fig. 3 shows the result and the comparison with the best parallel supercomputers. As the plot indicates, the 3D OEC performs one to two orders of magnitude better than the best parallel supercomputer despite the many additional data transfers required to reduce or enlarge the size of the 2D arrays.

The performance analysis we have presented, based on the CG and MG benchmarks, produced "mixed" results. On one hand the most complex benchmark, the MG, produced results consistent with the analysis in our previous paper [2]. On the other hand, the analysis of the CG benchmark showed that a 3D OEC, without the advantage of a fast reconfigurable interconnection network, may not outperform such powerful supercomputers.

With electronics becoming increasingly faster but reaching their limits, optical interconnects can provide a better solution to the problem of exchanging data between PEs. Special purpose systems can be build sooner and according to our analysis can provide considerable improvement in performance compared to

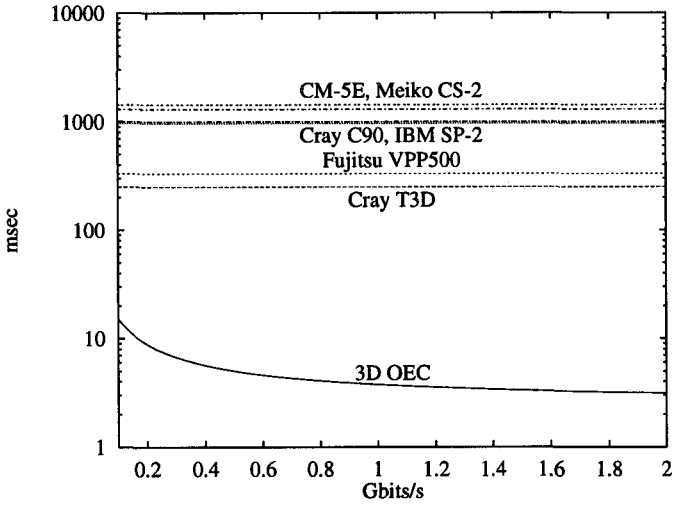

Fig. 3. Comparison of the performance of the 3D OEC and the best parallel supercomputers on the MG benchmark

electronic solutions. General purpose systems require the use of a reconfigurable network. As long as this is not feasible, electronics remain the only option.

Acknowledgements. The authors would like to acknowledge the partial support of National Science Foundation through contract EEC-9015128 and of Colorado Advanced Technology Institute through contract GEA-95-0002.

References

1. D. Bailey, *et al.*: The NAS Parallel Benchmarks. RNR Tech. Rep. RNR-94-007 (1994)
2. G. A. Betzos and P. A. Mitkas: Performance evaluation of 3D optoelectronic computer architectures bAsed on the FFT and sorting benchmarks. *Proc. MPPOI'95* (1995) 160–167
3. V. Morozov, J. Neff, H. Temkin, and A. Fedor: Analysis of a three-dimensional computer optical scheme based on bidirectional free-space optical interconnects. Opt. Eng., **34** (1995) 523–534
4. R. A. Nordin: Optical interconnects in electronic processing systems. Chapter 9 in *Photonics in Switching*, **1**, Academic Press, 1993
5. S. Saini and D. Bailey: NAS parallel benchmark results 12-95. NAS Tech. Rep. NAS-95-021 (1995)
6. E. Schenfeld: Massively parallel processing with optical interconnections: what can be, should be and must not be done by optics. Opt. Comp., **10** (1995) 16–18