

Towards an Optimal Foundation Architecture for Optoelectronic Computing

Haldun M. Ozaktas
Bilkent University
Department of Electrical Engineering
TR-06533 Bilkent, Ankara, Turkey

Abstract

By systematically examining the tree of possibilities for optoelectronic computing architectures and offering arguments that allow us to prune suboptimal branches of this tree, we come to the conclusion that electronic circuit planes interconnected optically according to regular connection patterns, represents an alternative which is reasonably close to the best possible as defined by physical limitations. Thus we propose that this foundation architecture should provide a basis for future research and development in this area.

1. Introduction

The integration of larger numbers of primitive computing elements (switches, transistors, gates, processors, etc.) to produce computers of greater processing power requires the use of interconnections with greater length/width ratios [1, 2]. (This can be avoided by resorting to architectures with local connections only, but for problems which intrinsically require global flow of information this merely amounts to breaking down the necessary long distance communication paths into a large number of short hops, which is not necessarily optimal [3].) As the length of an interconnection is increased, the time it takes for a signal to propagate to the other end also increases, at least as much as dictated by the speed of light. While this limitation holds for all types of interconnections, normally conducting electrical interconnections have much more severe limitations. The signal delay is a quadratic function of the length/width ratio beyond a certain length/width ratio, since the line becomes too lossy to allow pulse propagation [1, 2]. The energy per transmitted bit also increases with line length, even when repeaters are used. It can also be shown that for systems employing normally conducting interconnections, there exists an upper bound beyond which it is not possible to further increase the *bisection-bandwidth product*, which is a measure of the rate of internal information transfer in a system [1, 2].

On the other hand, increasing use of memory, the am-

bitation of processing large amounts of information such as with images and video, the attraction of parallel computing, and purely geometrical and physical considerations are factors which have contributed to the increasing importance of interconnections. For these and other reasons, it has been suggested to use optical interconnections for implementing the longer connections in computing systems, especially when an electrical line to be used instead would have a high length/width ratio.

After the potential of optical interconnections for overcoming the communications bottleneck in digital electronic computing systems was brought to widespread attention by publications such as [4], the analysis, design, and demonstration of devices, materials, components for optical interconnections has become a major part of the sub-area of optics called "optical computing" or "optics in computing." Because of the intrinsic overlap as regards the devices, architectures, and even systems employed (such as permutation networks), some of this research has also taken part under the sub-area known as "photonics in switching."

The most widespread approach has been to replace the longer electrical interconnections with optical ones without otherwise modifying the logical architecture. Examples are optical backplanes, fixed free-space interconnections between circuit boards, etc. In this spirit, optoelectronic technologies can be used to help wire up electronic circuits designed in the conventional way, by providing a large number of pinouts and high performance long distance connections. Although this approach definitely has a certain promise, it is not the one that we believe will bring the greatest rewards.

Fortunately, the need for general conceptual analysis, simulation, comparison, and optimization at a systems level has also been well recognized and resulted in considerable research. We refer the reader to a small sampling of special issues and conference proceedings which partly represent or include the work in this direction and in which further references may be found: [5, 6, 7, 8].

2. Architectural choices for optical interconnections

We will take a walk down the tree of alternative optical interconnection architectures. The labels of the options we will examine are itemized below:

- Two-dimensional systems
 - Waveguides
 - Planar free-space
- Three-dimensional systems
 - Fibers or waveguides
 - Free-space

Free-space

- Devices arrayed through volume
- Devices arrayed on plane

Free-space with devices arrayed on plane

- Locally connected
- Globally connected

Globally connected free-space with devices on plane

- Arbitrary connection pattern
- Regular connection pattern

We will omit discussion of two-dimensional systems in this paper, simply noting that it can be argued that they are of limited utility. Turning our attention to three-dimensional systems, it will be evident that free-space systems offer the best promise. Further examining the alternatives, we will decide that arraying the optical, electronic, or optoelectronic devices on a plane is preferable to arraying them throughout a volume. Upon comparing locally and globally connected systems, we will decide that globally connected systems are preferable. We will further argue that globally connected systems based on regular connection patterns constitute the best option.

2.1. Three-dimensional systems

We denote the number of elements (switches, processors) in a computing system by N . We assume that the graphs specifying the connections between these elements are of bounded degree; that is, the number of connections (“pinouts”) emanating from each element does not increase

with N . We also assume constant or approximately constant power dissipation per element, and that the elements are of constant size.

These assumptions are not restrictive, but are rather needed to ensure consistency. If we are to compare systems of different sizes and discuss how certain quantities change as system size increases, we must measure system size in a unit which is constant in processing power, size, number of connections (“pinouts”), and power dissipation. This unit is what we refer to as an element. For concreteness, we will concentrate on one-to-one (pairwise) connections.

We will now take a look at the factors which determine the smallest size of a three-dimensional computing system with N elements. Heat removal and interconnection density are the two major considerations which give lower bounds on system size. The need to minimize size is important not only for its own sake, but also because of the need to minimize propagation delays which are becoming increasingly important.

The minimum system size imposed by heat removal requirements is $\propto N^{1/2}$ [9]. The derivation is elementary. In [9] it is shown that the maximum total power \mathcal{P} that can be dissipated by a system is proportional to the cross-sectional area of the system, since there is a bound to the amount of power that can be removed per unit cross-sectional area. Since $\mathcal{P} \propto N$ when we assume constant power dissipation per element, the linear dimension of the system must grow at least $\propto \mathcal{P}^{1/2} \propto N^{1/2}$. In some systems, the power dissipation per element may actually increase with N , because the contribution of the interconnections to total power dissipation increases with system size. In that case, the minimum system size will grow even stronger than $\propto N^{1/2}$. However, the lower bound $N^{1/2}$ will be sufficient for the purposes of our arguments.

We now turn our attention to the bounds on system size imposed by the space occupied by the elements themselves. The minimum size imposed by arraying the elements throughout a volume is $\propto N^{1/3}$ and the minimum size imposed by arraying them on a plane is $\propto N^{1/2}$, as dictated by simple geometry. (N elements of given size cannot be packed in a box of size smaller than $\propto N^{1/3}$, nor arrayed on a plane over an area of size smaller than $\propto N^{1/2}$.) Naturally, arraying the elements on a plane implies a greater system size than arraying them throughout a volume. However, since heat removal requires a system size which is at least $\propto N^{1/2}$ anyway, arraying the elements on a plane does not result in a larger system size than arraying them throughout a volume. A more careful discussion of the relative importance of these factors also paying attention to the proportionality constants may be found in [1, 2].

In other words, when heat removal considerations are the dominating factor, the minimum system size does not depend on the configuration of the elements. (A similar

conclusion can be reached when interconnection density considerations are the dominating factor. It can be shown in this case that confining the elements to a surface has little effect on system size provided the communication paths are still free to use three-dimensional space [1, 2].)

We therefore conclude that arraying the elements/devices on a plane is satisfactory. This is fortunate, since arraying the elements throughout a volume would introduce considerable difficulties in fabrication and packaging. Also, most practical optical interconnection schemes provide connections between points lying on planar surfaces. Schemes for interconnecting a three-dimensional array of elements would almost certainly be much more difficult to realize. Such schemes have indeed been devised [10], but more in the nature of an existence proof than a practical proposal. Since it is much more convenient to work with planar arrays of devices, it is useful to know that they are good enough.

Throughout this paper, we will be speaking of *device planes*. With this term we refer to planar electronic circuits with optical input/output capability from their surface. (The term smart pixels [11] is also used to describe such device planes, but we find that term to have restrictive connotations and will thus avoid using it.) For instance, flip-chip bonding of self electro-optic effect devices (SEEDs) on silicon [12] or hybrid optoelectronic GaAs technology would allow the construction of such device planes. A device plane may actually consist of several active device layers sandwiched together so as to effectively constitute a single device plane. This would allow greater amounts of silicon circuitry per area if needed.

We now turn our attention to bounds on system size imposed by interconnection density considerations. Since interconnections take up space, the minimum size of a system depends on the degree of connectedness of the graph specifying the connections between its elements. There exists a continuum of degrees of connectedness between complete locality and complete globality. Some commonly used quantitative measures of connectedness are reviewed in [13]. However, consideration of the extreme cases will be sufficient for the purposes of the present argument. In a locally connected system the space occupied by the interconnections can be neglected, and the minimum size is that needed to accommodate the elements. Thus the minimum system size of a locally connected system is $\propto N^{1/3}$. On the other hand, globally connected systems have longer interconnections which take up more space so that their elements must be spaced further apart, resulting in larger system sizes. However, even the most globally connected graphs (such as the butterfly etc.) will not require system sizes exceeding $\propto N^{1/2}$ [1, 2]. To understand this, consider an imaginary surface bisecting the system in two such that $N/2$ elements fall on both sides. Even if all connections were made between elements on opposite sides of this surface, the number

of connections that must pass through this imaginary surface would be $\propto N$. Thus the size of this surface must be $\propto N^{1/2}$. (Remember that we are assuming the number of connections per element to be bounded.)

Therefore, given the heat removal imposed minimum system size of $\propto N^{1/2}$, we conclude that the implementation of a globally connected system does not result in greater system size than the implementation of a locally connected system. Since there is no tradeoff involved, a globally connected graph is to be preferred because of its greater versatility.

(Certain operations do not demand much connectivity among the elements of the system designed to perform them. In such less demanding cases, it might not make much difference whether we use a locally or globally connected system. We are considering the more interesting set of operations or problems which do demand global information flow for their solution.)

Finally, combining our two arguments, we conclude that we will prefer globally connected systems with devices arrayed on a plane. (Or any constant number of planes, if that is more convenient. Let us remember that we have shown that it is not disadvantageous to array the elements on a plane, presuming that it is more convenient to do so. We can still choose to array the elements on any number of planes or even throughout a volume if that turns out to be more convenient.) The bottom line is that the rather stringent and uncircumventable requirement imposed by heat removal, allows us considerable latitude in arraying the elements and providing interconnections among them. Since heat removal requires that we space the elements considerable distances apart, we might as well utilize this space to conveniently array the devices and also provide global interconnections. This is a consequence of the fact that in three-dimensional systems, heat removal considerations tend to dominate interconnection density considerations. This is in contrast to two-dimensional systems, where interconnection density considerations dominate and a similar general argument in favor of globally connected systems is not possible. In that case, the determination of the optimal degree of connectedness cannot be decoupled from the information flow requirements of the specific problem/application as in the three-dimensional case, so that general statements cannot be made and each case must be treated individually.

As a final comment, we note that the minimum system size $\propto N^{1/2}$ for globally connected systems is the theoretical minimum, the best that can be achieved. This minimum can indeed be achieved with proper choice of architecture [14]. However, unthoughtful designs may in general result in larger system sizes. Thus we must discuss what types of architectures allow the minimum possible to be achieved, since our argument in favor of globally connected systems would fail if we could not achieve the minimum $\propto N^{1/2}$ system size.

2.2. Free-space architectures for globally connected systems

Arbitrary connection patterns with multi-stage architectures Having argued in favor of globally connected systems with the elements arrayed on a plane (or some number of planes), we now explore more concretely the various alternative architectures for providing interconnections between these elements. We will find it convenient to imagine two planes facing each other, between which a prespecified pattern of connections are to be implemented (although it is easy enough to fold the architectures we will be discussing so that both optical sources and detectors lie on the same plane). For simplicity and concreteness, we will assume that an arbitrary pattern of one-to-one connections (a permutation) between the N sources on the plane lying to the left and the N detectors on the plane lying to the right has been specified.

In principle, a system size $\propto N^{1/2}$ can be achieved quite straightforwardly by the use of three-dimensional fibers or waveguides [1, 2, 15]. However, this alternative is not attractive because even if it were considered feasible from an engineering viewpoint, the constant of proportionality would be too large. The most common and conceptually simple class of architectures which allow arbitrary patterns of connections to be implemented, is the class of architectures which we might term *multi-facet architectures* (figure 1a), which all rely on aperture division to realize arbitrary space-variant connection patterns. It is well known that the system size imposed by this class of architectures is proportional to N , which is significantly larger than the theoretical minimum [15]. On the other hand, it can be shown that Banyan type (figure 2) multi-stage architectures can be employed to realize an arbitrary pattern of connections in the theoretical minimum size $\sim N^{1/2}$ [14].




Figure 1. a. Schematic depiction of a multi-facet architecture. b. Schematic depiction of a single-facet space-invariant architecture.

To avoid confusion, we must clarify the following point. Multi-stage architectures are often used as switching networks. Here, we are talking about a hardwired multi-stage




Figure 2. Regular connection pattern of a one-dimensional Banyan (butterfly) multi-stage architecture. Top: conventional diagram. Bottom: diagram with angles of all connections drawn equal, which can be fitted in a box of size $\sim N \times N$. The two-dimensional Banyan is more difficult to draw but similar in nature. Its optical realization can be fitted into a box of size $\sim N^{1/2} \times N^{1/2} \times N^{1/2}$ [14].

that is used to provide an arbitrary but fixed connection pattern. (Instead of dynamic exchange-bypass switches, we assume hardwired exchange-bypasses which determine the connection pattern.)

As a further comment, let us clarify why we have specified the Banyan among several other multi-stage networks, such as that based on the perfect shuffle. Use of a perfect shuffle based network (figure 3) results in a system whose size is larger than the theoretical minimum by a factor of $\log N$, whereas use of a Banyan based network allows us to achieve the theoretical minimum within a constant [14]. In most cases this might not be considered a significant difference, and other considerations might result in the choice of a perfect shuffle based or other network, rather than a Banyan. We will sometimes not be specific about which particular regular connection network is used, remembering that the difference is a logarithmic factor in the length of the system (the origin of which is evident upon examination of figures 2 and 3).

Introducing active intermediate planes Let us consolidate before we continue our argument. So far, we have argued in favor of a plane of electronic circuits, perhaps "smart pixels," interconnected to another plane of electronic circuits according to an arbitrary connection pattern, provided by the multi-stage network. Heat removal considerations, the volume occupied by the interconnections, and the area occupied by the devices, all imply a system linear extent $\propto N^{1/2}$. Of these three considerations, heat removal will most likely be the one to imply the largest proportionality factor and will thus determine the performance and size of




Figure 3. Regular connection pattern of a one-dimensional perfect shuffle multi-stage architecture which can be fitted in a box of size $\sim N \times N \log N$. The two-dimensional version is more difficult to draw but similar in nature. Its optical realization can be fitted into a box of size $\sim N^{1/2} \times N^{1/2} \times N^{1/2} \log N$.

the system.

We will consider a system whose length is $N^{1/2} \log N$ (for instance based on the two-dimensional version of the perfect shuffle shown in figure 3). From now on it will be simpler to refer to the schematic depiction shown in the top part of figure 4 rather than the more detailed connection pattern shown in figure 3 or its equivalent for other multi-stage networks.




Figure 4. Replacing passive intermediate planes with active device planes.

The intermediate planes may be passive in a small system. In larger systems, signal attenuation through the several stages might require regeneration of the signals as they go through several of the hardwired exchange-bypass modules. In any event, the intermediate planes have little function compared to the busy and bustling device planes, where all the processing elements reside. This unbalanced distribution of circuits and activity is clearly suboptimal, as we can obtain additional flexibility and function, without incurring any penalty in terms of system size, by adding circuits to the intermediate planes, especially if these planes must anyway contain regeneration circuits. In other words, if ac-

tive circuits are needed anyway in the intermediate planes, we might as well make more efficient use of the silicon there. Furthermore, for instance, in order to construct a random access parallel computer, we would be interested not merely in an arbitrary fixed connection pattern, but one which is dynamically programmable (a reconfigurable permutation network). In that case, the $\log N$ stage network we use would employ dynamically programmable exchange-bypass elements in the intermediate planes. In this case where the intermediate planes are expected to house active devices anyway, the argument in favor of fully utilizing the intermediate planes becomes even stronger. Why would we only sparsely utilize the intermediate planes, while the end planes are strained to the limit? It is clearly beneficial to uniformize the computational power throughout all existing planes rather than concentrating it at the ends and underutilizing the intermediate planes. Thus we make the transition from the top part of figure 4 to the bottom part of the same figure.

The system thus obtained occupies the same amount of space and is clearly equal to or greater in power than the previous system, since if nothing else, it can simulate the passive interconnection network. What we obtain as a result is a multi-device-plane computer with regular connections between its device planes. Such a system is of the same size as a system with only two device planes connected according to a fixed arbitrary pattern, and is much more versatile.

It is possible to bring forth the objection that we are totally ignoring the cost of furnishing the additional device planes. But the fact that we are adding more devices and circuitry does not mean that we will increase overall cost per performance. More fundamentally, we should underline that we are measuring cost in terms of volume and area, not number of devices or what is in the volume. This is the measure of cost that we expect to be relevant in future systems. To convince ourselves of this, we might think of the days of discrete electronic circuits, when component count and type was the major determinant of cost, and compare this with integrated circuits, where essentially only area counts; wires and devices do not have different costs in this uniform medium.

Introducing $\log N$ times as many circuits into the system means that the total power dissipation will also be increased by this factor, if all devices are active at the same time. However, any of the side faces of the system of area $N^{1/2} \times N^{1/2} \log N$ is sufficient to remove this power.

Our chain of arguments already shows that a system consisting of $\log N$ regularly connected device planes is better than a system based on the multi-facet architecture. Nevertheless a direct comparison will be instructive. It is almost always the case that a system with only regular connections between its planes—with modifications not affecting its essential properties—can simulate a system with an arbitrary

pattern of connections between its planes, with $\log N$ stages or iterations (as elaborated in the next paragraph). Thus, since the size and delay for a single stage or iteration is $\propto N^{1/2}$, the total delay involved is $N^{1/2} \log N$. The same could be realized in a single step on a system which can provide an arbitrary pattern of interconnections by employing a multi-facet interconnection architecture, but the total delay involved would be $\propto N$, since the size and delays of a multi-facet system grow $\propto N$. Since $N^{1/2} \log N < N$, the regularly connected system is preferable.

Our argument essentially relies on the fact that we can simulate a system whose elements are connected by an arbitrary connection pattern, with a system connected by a regular connection pattern, in $\log N$ stages or iterations. The “proof” is relatively easy. It is known that an arbitrary permutation network can be realized in $\log N$ stages or iterations, relying only on regular connections between the stages or iterations. Thus, the least the regularly connected system can do is to simulate the arbitrarily connected system in $\log N$ stages or iterations. If the existing circuits or processors are not already capable of such functions, exchange-bypass switches may have to be introduced to make a given regularly connected system capable of simulating a permutation network. However, the amount of circuits per plane needed for these switches is proportional to N , which can be absorbed into the area occupied by the N elements or processors.

We emphasize that the introduction of exchange-bypass switches is only a fiction employed in our “proof.” In practice, the circuits and algorithms would be integrally designed for the regularly connected system so as to be capable of guiding the information in the necessary manner through the regular pattern of connections; there would be no reason to first design the circuits and algorithms for an arbitrarily connected system, and then simulate the arbitrarily connected system on a regularly connected system.

2.3. Summary

In essence, we have argued that a certain degree of physical interconnectivity is optimal. Global interconnections are better, but regular ones are sufficient. This degree of interconnectivity is precisely that provided by regular interconnection patterns such as the perfect shuffle or the most significant stage of the Banyan. Anything less connected (more locally connected) does not save space, since heat removal considerations will force things apart anyway. On the other hand, architectures providing an arbitrary pattern of connections directly are not beneficial since they require more space than implied by heat removal considerations, but without offering any compensating advantage. (To argue this last point, we first showed that architectures providing an arbitrary pattern can be simulated by a hardwired

multi-stage network, and then noted that if we have a multi-stage, there is no point in underutilizing the intermediate stages while crowding the computational elements at the end planes. Clearly, it is better to put some processing power in the intermediate stages as well. Thus, we ended up with a multi-stage system with regular interconnections between its stages, and processing power distributed uniformly throughout all stages.)

In conclusion, we have decided that the best foundation architecture upon which to build on is that consisting of regularly interconnected device planes. Instead of trying to provide arbitrary patterns of connections with the hardware, we should provide global regular connections—an approach that balances almost every physical requirement harmoniously, and then we should design the circuits and algorithms so that the information flows as it should. The lack of arbitrary connections is not a loss, since in such systems the information can be propagated to where it should after at most $\sim \log N$ stages; and this many stages are needed anyway for realizing arbitrary permutations with a multi-stage network (which takes up less space and results in less signal delay than a single-stage multi-facet arbitrary permutation architecture.)

3. Regularly interconnected device planes

In advocating regularly interconnected device planes as a foundation architecture, what we are essentially saying is that instead of trying to provide an arbitrary pattern of connections in hardware, it makes more sense to provide the opportunity for global flow of information in a physically efficient way, and let the information be guided where it needs to by the algorithm, if necessary in several steps or iterations.

Such a system will be most successful if one contemplates its higher level organization and algorithms from the outset such that it relies on only a regular pattern of interconnections. Without the benefit of such integral design, simple-minded emulation of algorithms designed to work on architectures which are able to provide arbitrary patterns of connections may be inefficient.

It is worth underlining that customization of such a system involves customization of the electronic circuits in the device planes and whatever software is involved. Unlike the multi-facet or fixed multi-stage architecture whose optical components must be customized, the optical interconnection pattern for this architecture and thus the optical components are always the same no matter what purpose the system is designed for. Delegating the customization to the well established VLSI and software technologies should be beneficial from the optical design and manufacturing viewpoint, and should enable the production of robust and well-optimized optical interconnection modules. The fact that VLSI and

computer systems designers do not have to worry about the optics involved should greatly increase the interest in this technology and contribute to its rapid takeoff.

One final advantage of regularly interconnected device planes is that architectures belonging to this class have already been extensively studied for use in switching systems [16] as well as for other applications. Not only is knowledge of the mathematical aspects well developed [16], but optical implementations in the form of switching networks have been demonstrated [17]. What we have argued is that such systems are reasonably close to the best possible as defined by physical limitations.

4. Application platforms based on regularly interconnected device planes

The regularly interconnected device planes architecture advocated in this paper is a foundation architecture upon which to build platforms on which more specific applications can be implemented. The general features of the regularly interconnected device plane architecture were determined by purely physical considerations. The more specific features of an evolved platform must be determined by the class of applications one wishes to implement. An analogy with programming languages may be useful. The regularly interconnected device plane architecture is in some ways analogous to a low-level programming language, such as machine language. It is not a very efficient process to design a system for a particular application by starting directly from the bare regularly interconnected device plane concept, much as it is inefficient to write an application directly in machine language. The platforms we describe in the full version of this paper may be likened to high-level programming languages, with which it is much more efficient to implement a desired application.

Although all programming languages are equivalent in a general sense, they are not equally efficient for a given purpose. The choice of a suitable platform, much like the choice of a suitable programming language, depends on the type of application that is to be implemented. New classes of applications might require the development of new platforms. Basing such platforms on regularly interconnected device planes will assure us that it exhibits desirable or optimal characteristics from a physical point of view.

5. Conclusion

It is clear that large arrays of very fast and low-energy optical devices integrated with established electronic technology and interconnected with free-space optics has very large computational power in the raw sense, but realizing this potential may not be so easy. The difficulty stems from

the fact that a whole system of paradigms and levels of abstraction has been constructed around the capabilities and limitations of purely electronic systems, and the dominance of this system resists the introduction of a new technology with completely different capabilities and limitations. There does not seem to be much point in trying to build an optical microprocessor, and the user level improvements obtained by replacing the longer wires in conventional systems with optics may be limited. On the other hand, starting with an array of smart pixels, we are too many levels of abstraction away from being able to write a program that plays chess. Clearly, considerable research is needed to determine how optoelectronic computing systems should be contemplated and to develop platforms which could guide future efforts.

In contemplating the design of some system, it is common to choose an ad hoc starting point. Instead, we have carefully and systematically examined the tree of possibilities for optoelectronic computing systems, and by offering arguments that allow us to prune suboptimal branches of this tree, we arrived at what seems the best approach. The option we are advocating balances the various physical constraints while exploiting the strength of optics as much as possible. It is flexible enough for forming the basis of several generic platforms which should stimulate further development. Some of these platforms had already been studied, but mostly on an ad hoc basis.

It was quite clear that the architecture we advocated balanced the major physical requirements nicely, and thus was in some sense optimal. The problem was to determine how much was lost by restricting ourselves to regular connections. We have argued that we do not lose much. For instance, we argued that any parallel computer algorithm which runs on a reconfigurable permutation network, can be distributed through multiple regularly connected stages, and that will be better than realizing the permutation network directly.




Figure 5. The foundation architecture (multi-stage version).

The foundation architecture we have argued for is depicted in figure 5. It consists of an array of primitive elec-

tronic elements lying on a plane, with a certain number of optical inputs and outputs corresponding to each element. The power dissipated by these elements is removed by fluid convection with the paths of fluid flow perpendicular to the device plane. The thickness of the coolant paths are chosen according to the analysis in [9]. The optical interconnection system provides global regular interconnections (such as the perfect shuffle or the most significant stage of the Banyan) among the optical pixels, perhaps by using a nearly space-invariant optical system as described in [18, 19, 15] and shown in figure 1b. What we have purported to show in this paper is that provided we use this type of architecture properly by employing the proper algorithm and so forth, no other architecture can do much better. We believe that development of application platforms based on our foundation architecture will constitute promising and fruitful avenues for further research. These are discussed in the full version of this paper.

Needless to say, it would be pretentious to claim that the arguments presented in this paper are definitive. And in any event, there will always be situations and instances when alternative approaches are viable or preferable; our arguments aim at capturing the mainstream trend. However, we believe we have maintained a level of rigor commensurate with the complexity of the problem. Indeed, predicting the future of optoelectronic computing should be likened to problems such as predicting the future of some aspect of the world economy or the like. Although experience shows us that little success is achieved in such endeavors, they are nevertheless not considered futile exercises because of the useful thinking they stimulate, and we hope the same can be concluded for this work.

6. Acknowledgments

I would like to acknowledge the benefit of extended interaction with David A. B. Miller of Bell Laboratories, which has helped me develop or clarify several ideas and issues which appear in this paper. I would also like to extend my thanks to Ashok Krishnamoorthy and John Ford of Bell Laboratories, Fouad Kiamilev of The University of North Carolina at Charlotte, and Cevdet Aykanat of Bilkent University for useful discussions. Some of the research that constitutes the background for the present work were realized in collaboration with Joseph W. Goodman of Stanford University, Adolph W. Lohmann of the University of Erlangen-Nürnberg, Yaakov Amitai of The Weizmann Institute, and David Mendlovic of Tel-Aviv University.

References

- [1] H. M. Ozaktas. *A Physical Approach to Communication Limits in Computation*, Ph.D. thesis, Stanford University, Stanford, California, 1991.
- [2] H. M. Ozaktas and J. W. Goodman. The limitations of interconnections in providing communication between an array of points. In *Frontiers of Computing Systems Research, Volume 2*, S. K. Tewksbury, editor, Plenum Press, New York, 1991, pages 61–130.
- [3] H. M. Ozaktas and J. W. Goodman. Comparison of local and global computation and its implications for the role of optical interconnections in future nanoelectronic systems. *Opt. Commun.*, 100:247–258, 1993.
- [4] J. W. Goodman, F. J. Leonberger, S.-Y. Kung, and R. Athale. Optical interconnections for VLSI systems. *Proc. IEEE*, 72:850–866, 1984.
- [5] Special issue on optical computing systems. H. F. Jordan and M. J. Murdocca, editors. *Proc. IEEE*, Vol. 82, No. 11, November 1994.
- [6] Special issue on optical computing. H. S. Hinton, B. Soffer, F. A. P. Tooley, and K.-i. Yukimatsu, editors. *Appl. Opt.*, Vol. 33, No. 8, 10 March 1994.
- [7] Special issue on optical computing. *Appl. Opt.*, to appear in Vol. 35, 1996.
- [8] *Massively Parallel Processing Using Optical Interconnections*, Proceedings of the Second International Conference, San Antonio, Texas, October 1995, E. Schenfeld, editor, IEEE CS Press, Los Alamitos, California, 1995.
- [9] H. M. Ozaktas, H. Oksuzoglu, R. F. W. Pease, and J. W. Goodman. Effect on scaling of heat removal requirements in three-dimensional systems. *Int. J. Electronics*, 73:1227–1232, 1992.
- [10] H. M. Ozaktas, Y. Amitai, and J. W. Goodman. A three dimensional optical interconnection architecture with minimal growth rate of system size. *Opt. Commun.*, 85:1–4, 1991, Errata in 88:569, 1992.
- [11] Special issue on smart pixels. S. R. Forrest and H. S. Hinton, editors. *IEEE J. Quantum Electronics*, Vol. 29, No. 2, March 1993.
- [12] K. W. Goossen, J.A. Walker, L. A. D’Asaro, S. P. Hui, B. Tseng, R. Leibenguth, D. Kossives, D. D. Bacon, D. Dahringer, L. M. F. Chirovsky, A. L. Lentine, and D. A. B. Miller. GaAs MQW modulators integrated with silicon CMOS. *IEEE Phot. Tech. Lett.*, 7:360–362, 1995.
- [13] H. M. Ozaktas. Paradigms of connectivity for computer circuits and networks. *Opt. Eng.*, 31:1563–1567, 1992.
- [14] H. M. Ozaktas and D. Mendlovic. Multi-stage optical interconnection architectures with least possible growth of system size. *Opt. Lett.*, 18:296–298, 1993.
- [15] H. M. Ozaktas, Y. Amitai, and J. W. Goodman. Comparison of system size for some optical interconnection architectures and the folded multi-facet architecture. *Opt. Commun.*, 82:225–228, 1991.
- [16] H. S. Hinton. *Introduction to Photonic Switching Fabrics*. Plenum, New York, 1993.
- [17] H. S. Hinton, T. J. Cloonan, F. B. McCormick, A. L. Lentine, F. A. P. Tooley. Free-space digital optical systems. *Proc. IEEE*, 82:1632–1649, 1994.
- [18] M. R. Feldman, C. C. Guest, T. J. Drabik, and S. C. Esener. Comparison between optical and electrical interconnects for fine grain processor arrays based on interconnect density capabilities. *Appl. Opt.*, 28:3820–3829, 1989.
- [19] G. E. Lohman and K.-H. Brenner. Space-variance in optical computing. *Optik*, 89:123–134, 1992.
- [20] H. M. Ozaktas and J. W. Goodman. Elements of a hybrid interconnection theory. *Appl. Opt.*, 33:2968–2987, 1994.

We acknowledge the support of NATO under the Science for Stability Program.