Toward an optimal foundation architecture for optoelectronic computing. Part I.   Regularly interconnected device planes

Haldun M. Ozaktas

By systematically examining the tree of possibilities for optoelectronic computing architectures and offering arguments that allow one to prune suboptimal branches of this tree, I come to the conclusion that electronic circuit planes interconnected optically according to regular connection patterns represent an alternative that is reasonably close to the best possible, as defined by physical limitations.   Thus I propose that this foundation architecture should provide a basis for future research and development in this area.   © 1997 Optical Society of America

Key words: Optical interconnections, optical computing.

1.   Introduction

A.   Background

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 if one resorts to architectures with local connections only, but for problems that intrinsically require a 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 is dictated by the speed of light.

Although the above 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 permit pulse propagation.[1,2,4]   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, an increasing use of memory, the aspiration 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 that have contributed to the increasing importance of interconnections.   For these and other reasons (e.g., the possibility of nonplanar interconnections, voltage isolation, very little or no frequency-dependent cross talk and distortion, no impedance-matching problems even with multiple taps, etc.) that have been extensively discussed, it has been suggested that optical interconnections be used for implementing the longer connections in computing systems, especially when an electrical line 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 Ref. 5, the analysis, design, and demonstration of devices, materials, and components for optical interconnections has become a major part of the subarea of optics called optical computing, or optics in computing. Because of the intrinsic overlap with respect to the devices, architectures, and even systems employed (such as permutation networks), some of this re-

search has also taken place under the subarea 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 the provision of 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 I believe will bring the greatest rewards.

Fortunately, the need for general conceptual analysis, simulation, comparison, and optimization at the systems level has also been well recognized and has resulted in considerable research. I refer the reader to a sampling of papers, special issues, and conference proceedings that partly represent or include the work in this direction and in which further references may be found: see, for example, Refs. 6–43.

B. Nature of the Models Employed

The abstract models used for analyzing, comparing, and predicting the properties of a certain class of systems must capture the essential nature of the technology used to implement these systems. Let us consider that, if we were dealing with simple electronic logic circuits assembled from discrete components on a breadboard, the relevant parameters would include component count, logic depth, etc. On the other hand, for advanced digital integrated circuits, the relevant parameters include chip area, the longest connection length, etc. As technology evolves or when it is altered radically, it is necessary to reevaluate the models employed and change or replace them as appropriate. Electronics technology has come a long way since the transistor was conceived. Many of the technical and nonfundamental barriers determining cost and performance have been overcome. At the stage of digital electronics as we know it today,[44] further improvements are bringing us closer to the ultimate cost and performance possible, as determined by fundamental physical limits. We can expect the models appropriate to the present state of the technology, such as those employed in Refs. 1, 2, 45, and 46, to serve us until we actually reach the fundamental limits. Thus, these models can also be used to determine the ultimate cost and performance that can be attained when these fundamental limits are reached. This exercise has been carried out for systems interconnected with normally conducting interconnections, repeated interconnections, superconducting interconnections, and optical interconnections.[1,2] The same exercise was extended to systems employing both normally conducting and optical interconnections.[47]

The models employed for optical interconnections in the studies just referred to were also chosen to reflect the final stage in the development of optical interconnection technology, when we will be working against fundamental physical limits. We are already there in some respects, but not yet so in others. In general we are close enough that the assumptions of our models are plausible extrapolations of present trends and developments. The major exception is in the area of packaging, where the level of development is not yet pushing against fundamental limits. However, there is no fundamental reason why we cannot expect technical ingenuity to eliminate the obstacles in this area as well.

The essential character of the models employed is that they correspond to the case in which a system has been packed as tightly as possible insofar as physical limitations will allow. This is mostly the case for a modern electronic integrated-circuit chip and is more or less the case for a modern high-performance electronic computing system. The models we employ for electronic systems would have been inappropriate in the age of discrete components and also in the age of device-limited, rather than wire-limited, integrated components. However, as inappropriate as they were for the systems built in the past, these models would have enabled researchers to predict the limits of integrated-circuit technology 30 or 40 years ago with only minimal guesswork. All that those researchers had to do was to examine the fundamental physical limits involved and assume that the technical problems would eventually be overcome. However, meaningful predictions did not arrive until the late 1970's.[48–50] Earlier researchers did seek the fundamental limits involved, but they seem to have failed to appreciate the growing dominance of interconnections. They predicted the limits of digital computing systems on the basis of the fundamental limits imposed by devices, treating interconnections as mere parasitics or ignoring them altogether, whereas the opposite would have been more appropriate.[51] In other words, they failed to identify correctly what were fundamental limitations and what were merely technical problems to be overcome. Their example illustrates the difficulties and pitfalls inherent in trying to see the future. Without discounting such difficulties, I feel that, having observed the development of integrated electronic systems and witnessing the trend in optoelectronic systems in a similar direction, we are in a position to claim with reasonable confidence that optoelectronic systems will also converge toward the densely integrated models we employ.

It is important to underline that the models we use are not arbitrary; they are defined by physical and geometrical limitations that remain after technical obstacles have been surmounted and thus represent the natural end toward which technology should converge. In the early stages of a technology, the initial aim is to show that things can work with reasonable efficiency and to demonstrate that there exists a path for future progress. As more and more of the technical problems encountered are solved, performance aims are set higher and higher, and the technology tends to converge toward the point at which cost and performance are limited by fundamental physical

limitations only. This suggests the following strategy for predicting the shape of things to come:

- Clearly separate technical obstacles from fundamental limitations.
- Assume the technical obstacles will be overcome.
- Determine the system for which performance and cost attain their optimal values, as constrained by fundamental limitations.

The above argument suggests a deterministic theory of technological progress with the state of the art evolving teleologically toward the point at which it offers the ultimate possible performance–cost curve. However, technological progress is a more complicated process than is suggested by this simple theory. In fact, we are not at all automatically ensured of reaching the optimal final point. The final state of the art that we arrive at will most likely be *path dependent*, representing a local rather than a global optimum point. Various economic, corporate, and historical factors may limit the attention span of the research and development community, diverting development into paths that may lead to a globally suboptimal terminal point. After significant effort has been invested in following such a path, it might not be possible to back out.

Rosenberg[52] has discussed at length the path dependence of the telecommunications industry. There are sufficient parallels to allow us to generalize his results to the information-processing industry at large. This indicates the importance of the research community's having a clear picture, indeed a common long-term vision, of where it should be going so that it can consciously avoid drifting into the wrong path. It is one of the major purposes of this study to contribute to the discussion and development of such a vision.

The above remarks are applicable to the progress of a given technology. In some cases, a totally different competing technology may eradicate the given one in the middle of its progress, before it has even reached its fundamental limits. It seems, however, that optoelectronic computing will reach maturity before other technologies, such as atomic-scale quantum technology, molecular–biological engineering, etc., become feasible.

Let us conclude this subsection with an observation that is particularly applicable to our efforts[53]: Engineers usually consider their work to be hard science for which everything is quantifiable and all statements can be expected to be precise. However, the problems encountered in trying to predict future developments in a technology or exploring alternative paths of development are more similar to the problems of sociology, economics, or similar sciences. The problems are very complex; it is possible to deal quantitatively with only a small fraction of the very large number of parameters, some of which are not known or cannot be controlled or even measured. These circumstances require different standards of rigor and different standards of what is a valid argument.

C. Overview of the Paper

In Refs. 47 and 54 hybrid systems employing both optical and electrical interconnections have been analyzed and optimized. In this and previous works, the computing system as a whole was imagined to be a single uniform integrated system. Whereas this approach is useful for predicting the overall performance limits and the role of optics, it is not helpful in a constructive way for the design of systems.

This is because even moderately complicated systems cannot be designed by specification of their logic function and then employment of a fully automated computer-aided design tool. Rather, the design of a computing machine takes place at several levels of abstraction ranging from materials and device engineering to system architecture to high-level software. This system of levels of abstraction enables the design problem to be broken down into manageable subproblems, much as in a procedural programming language. It is first necessary to show how certain elementary functional units (in the abstract sense) can be formed and then how these can form higher-level units and so on, until we arrive at some kind of high-level programming language that permits the problem description to be formulated. (For further discussion of these issues, see Refs. 55 and 56.)

Replacing the longer wires in existing digital electronic systems with optical interconnections is not necessarily the best way to realize an optoelectronic computer, even if it offers a certain degree of improvement. Examples of this approach might be the introduction of optical backplanes or chip-to-chip modules instead of their electrical counterparts, while leaving the architectural conception and logical structure of the machine intact. This approach is appealing in that we do not have to worry about the development of new architectural concepts. However, there is no reason why the existing concepts should be particularly congenial to optical technology. In fact, they have historically developed to benefit from the strengths and accommodate the weaknesses of electrical technology that are in some senses complementary to those of optics, so that this approach may not bring out the best of optical components. (VLSI architectures that try to minimize the length and number of chip-to-chip interconnections provide a good example.)

Thus, the existence of a feasible optical interconnection technology and the results of studies such as that reported in Ref. 47 are necessary but not sufficient. It is also necessary to come forward with an arguably efficient or optimal platform encompassing certain lower levels of abstraction on which higher-level design can take place. In this analysis our aim is to argue in favor of certain platforms encompassing the physical and architectural levels on which algorithm and circuit design can take place.

I first discuss what is meant by the term interconnection theory and give some examples of the types of problem it addresses. Then I discuss the various architectural choices for optical interconnections, first

considering two-dimensional systems and then moving on to three-dimensional free-space architectures. Among the various alternatives, I single out regularly connected multiple-device-plane architectures as a promising alternative and discuss its benefits.

This paper and its sequel serve as a review of several issues that have been discussed by my colleagues and me as well as by other researchers in previous publications, and it tries to unify them to construct an argument as to what the best architectural choices are. It is a point of convergence for several previous studies and also serves as a point of departure for recently completed or ongoing research. To make the exposition as accessible as possible, I have tried to simplify and streamline the discussion and to make it as transparent as possible, especially when more extensive discussions pertaining to the particular results and issues in question may be found in the references.

2. Interconnection Theory

A. Nature of Interconnection Theory

A set of concepts and methods of analysis need not have a name to be useful. However, a name can give cohesion and unity to these concepts and methods and make them more tangible and visible. For this reason it is useful to identify a set of mathematical and empirical models, observations, mathematical concepts and tools, and methods of analysis under the title of "interconnection theory." Interconnection theory is a *physical* theory of computation based on *interconnect*-dominated models.[1,2] It is a physical and architectural[57] theory as opposed to a logical or algorithmic theory in that it deals with the actual physical and material construction of computing systems, with the flow of information through real space as governed by geometrical and physical limitations, and with the problems of heat removal and power distribution. Indeed, interconnection theory may be called *physical* computer science. This does not mean that logical or algorithmic considerations are ignored; quite the contrary: It is found that these considerations are tightly coupled to physical considerations, necessitating an interdisciplinary treatment (as is discussed further below).

Interconnection theory is based on interconnect-dominated models rather than on device-dominated models on the basis of the understanding that computing systems of ever-increasing numbers of components are limited by the problems associated with transferring information within the system rather than with the intrinsic limitations of the devices themselves.[1] Interconnection theory does not treat interconnections as mere parasitics that degrade the expected performance of the devices; rather, it puts them at the center of its models.

Digital computers are made by the interconnection of nonlinear elements according to a certain graph. Interconnections are physical channels with width, length, energy consumption, delay, and bandwidth. Interconnection theory deals with the resulting system-level parameters such as size, power consumption, and speed and how these are affected by architectural and technological choices.

B. Architectural and Algorithmic Issues are Coupled—but Not Always

Before embarking on our main discussion, it is useful to give some examples that not only constitute building blocks of our main discussion but also illustrate the types of problems one can try to solve with interconnection theory.

There are several architectural and algorithmic decisions that must be made when contemplating a computing system. It is particularly difficult to arrive at the correct decision when these considerations are tightly coupled. In general we need a physical theory of computation with which we can formulate the various constraints and optimize jointly over the various architectural and algorithmic choices so as to optimize measures of performance and cost. VLSI complexity theory[58] combines these considerations to a limited degree, and some applications to optical systems can be found in previous studies.[59–62] A general discussion of what such a theory would look like is given in Ref. 2, but the theory itself does not really exist. Those dealing with the physical aspects of devices, those dealing with transmission lines, interconnections, and packaging, and those dealing with the architectural and logical aspects of computing systems often limit their attention to their own domains and remain necessarily naive about the concerns of those dealing with other domains. As a result, the solutions they find are optimal in a narrow sense, in that they may not be the optimum solutions that would be obtained from a theory that jointly considers all domains at once. No one can be found at fault for failing to address an inherently difficult problem, and indeed we are very far from the kind of theory we are alluding to, which can take into account the various factors all at the same time.

However, for certain issues a number of general assumptions may allow us to reach certain results that may be claimed to be optimal in a wider sense, although they are not obtained from a fully general theory. This is possible when a certain aspect of the problem can be isolated or separated such that consideration of other parts would have no effect on the result anyway. Let us consider two example problems for which some useful conclusions can be drawn.

C. Global Versus Local Interconnections

Our first example is the contention between global and local architectures,[3] which is summarized in itemized form as follows:

- Global architectures (example: a butterfly graph):
  - Algorithms with a *small* number of steps.
  - *Long* physical duration for each step.
- Local architectures (example: a mesh graph):
  - Algorithms with a *large* number of steps.
  - *Short* physical duration for each step.

If we employ global architectures, the available connections make it possible to employ an algorithm that computes the answer in a small number of time steps. However, the physical duration of a single time step is long because of the length of the interconnections. On the other hand, we may employ a locally connected architecture that will require a large number of time steps but in which the duration of the time steps will be short. To determine which alternative will result in the smallest overall time of computation, we must optimize jointly over all possible algorithms for both architectures. In fact, this is a highly simplified picture since there is actually a continuum of degrees of connectedness between the extremes of complete locality and complete globality, so the actual problem is even more difficult. However, by comparison of the limitations imposed by heat removal with those imposed by interconnection density, it is possible to make a general argument in favor of global architectures without getting into a discussion of algorithms.

Although the reader is referred to Ref. 3 for details, we can summarize the essential point of the argument as follows. The use of a globally connected architecture is advantageous since it minimizes the number of time steps, but it is disadvantageous because the long interconnections needed may take up too much space, forcing the elements constituting the system far apart and resulting in a large system size and long signal delays. However, it is possible to show that heat-removal considerations imply a growth rate for the system size that is proportional to $N^{1/2}$, where N is the number of elements in the system.[63] The heat-removal-imposed system size is almost always greater than the size needed to accommodate the interconnections in even the most globally connected systems, such as permutation networks. Since heat removal requires large interelement separation anyway, there is no additional penalty to pay for employing a global interconnection architecture.

D. Regular Versus Irregular Interconnections

In our second example, the multifacet architecture[18,64] that can provide an arbitrary pattern of connections between two device planes is compared with a nearly space-invariant interconnection architecture that can provide only a regular pattern of connections (see Figs. 1 and 2). Again, the main features of the trade-off involved may be summarized in itemized form:

- Device planes interconnected with the multi-facet architecture:
    - *Arbitrary* pattern of connections.
    - *Fewer* steps or iterations.
    - *Large* system size, connection length, and delay.
- Device planes interconnected with a regular (nearly space-invariant) architecture:
    - *Restricted* pattern of connections.
    - *More* steps or iterations.
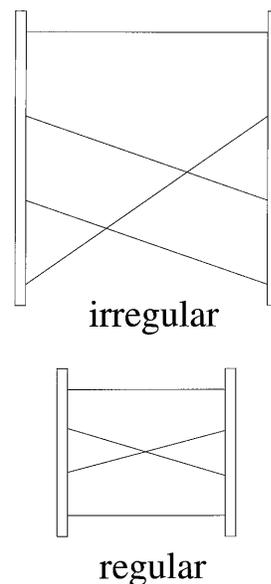   - *Small* system size, connection length, and delay.

Fig. 1.   Regularly and irregularly interconnected systems.

We imagine that the two device planes shown in both parts of Fig. 1 house a number of processors that are able to work together to solve a certain problem. These two device planes might represent the whole of a computer or only a section of it. Information might go back and forth between the two planes in an iterative fashion, or similar sections may be cascaded to form a pipeline. In fact, maybe there is only a single plane of devices instead of the two shown in Fig. 1, and the connection pattern is onto itself (as shown in Fig. 2). One of the device planes may consist of processors and the other of memories, or there might be some local memory in each processor. Such details are not relevant for our purpose here.

The architecture with the regular connection pattern may seem restrictive, but a system with such connections can solve the same problems as the other in an indirect manner through the action of shuffling the information back and forth several times. Despite the fact that this system will require a greater number of iterations or time steps to solve a given problem, it will also exhibit a smaller system size and
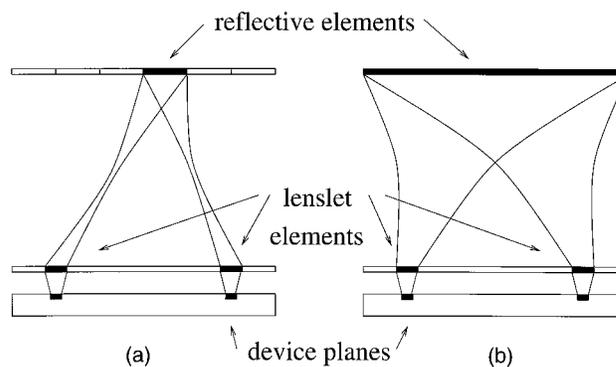

Fig. 2.   (a) Schematic depiction of a multifacet architecture.   (b) Schematic depiction of a single-facet space-invariant architecture.
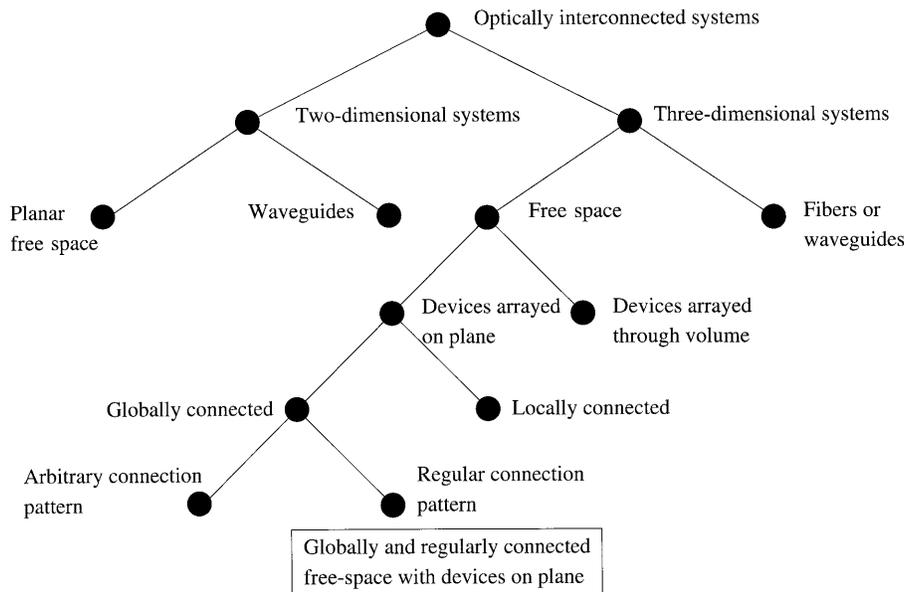
Fig. 3. Tree of alternative optical interconnection architectures.

shorter signal delay. Since the physical duration of each time step or iteration will be smaller, this system may perhaps exhibit a smaller overall time of computation. To say which system will be faster in general, it is necessary to carry out joint optimization over architectural and algorithmic choices.

However, once again it is possible to offer a general argument without embarking on such a joint optimization. We will return to this problem below and argue that the regularly connected system is better. The argument relies on the observation that, while the regularly connected system may incur a factor of $\log N$ slowdown in terms of the number of time steps needed to solve typical problems, the size of such a system and thus the propagation delays can be of the order of $N^{1/2}$, as opposed to N for the irregularly connected system. Since $N^{1/2} \log N < N$, the regularly connected system results in a smaller overall time of computation.

3. Architectural Choices for Optical Interconnections

After the somewhat extended introductory material, we can now embark on our main argument. We take a walk down the tree of alternative optical interconnection architectures. The labels of the options we examine are itemized below and also depicted in Fig. 3:

- 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 look first at two-dimensional systems and argue that they are of limited utility. Turning our attention to three-dimensional systems, it becomes evident that free-space systems offer the best promise. By further examining the alternatives, we decide that arraying the optical, electronic, or optoelectronic devices on a plane is preferable to arraying them throughout a volume. On comparing locally and globally connected systems, we decide that globally connected systems are preferable. We further argue that globally connected systems based on regular connection patterns constitute the best option. A system is considered superior to another if it can finish the same task in a shorter amount of time or finish a larger task in the same amount of time (cost may similarly be factored into the equation).

A. Two-Dimensional Systems

Three-dimensional systems are of course better than two-dimensional systems in terms of performance, but since they take up less space there is still a point to comparing two-dimensional optically interconnected systems with two-dimensional electronic systems.

Comparisons of the capabilities of two-dimensional optical and electrical interconnections do not significantly favor optics when we allow for active repeating stages in the electrical lines. In electrical systems repeaters can be used without significant penalty. (With submicrometer systems, the area the repeaters consume on the chip can be much less than that

consumed by the wires.[1,2]) Optical interconnections offer better performance only if effective interconnection widths can be brought down to the order of a few micrometers (which means of the order of an optical wavelength). Even then, they offer a noticeable advantage in very limited circumstances.[47]

To determine if it is indeed possible to bring the effective interconnection widths down to the order of a few micrometers for complex waveguide circuits, we have developed a computer-aided analysis and design tool that allows us to calculate the minimum waveguide spacings in complex circuits of arbitrary rectilinear topology, so as to maintain acceptable cross-talk levels.[65] We found that, as a result of the necessity of avoiding interwaveguide coupling, complex waveguide circuits force large effective widths. The results of this study indicate that effective widths cannot be brought down to a few micrometers for large circuits, so we conclude that optical waveguide circuits cannot compete with electrical integrated circuits. (Optics becomes even more disadvantageous when we consider the additional improvements possible by reducing the electrical resistance at low temperatures and the possibility of a greater number of interconnection layers in electrical systems, which may not be possible with optical waveguide circuits.)

The folded multifacet architecture,[61] which is a particular kind of optimized planar free-space architecture,[66,67] was devised as a way to achieve near-diffraction-limited effective widths by means of avoiding the intrinsic problems of integrated optical waveguide circuits. Although we have not done a detailed analysis of higher-order effects in such a system, it does seem that effective widths approaching a few micrometers can be achieved. Nevertheless, as we commented above, even in this case the use of two-dimensional optical circuits offers a noticeable advantage in very limited circumstances. Thus we may conclude that two-dimensional optically interconnected systems will not find widespread use in future high-performance computing systems.

B. 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 rather are 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 the system size in a unit that is constant in processing power, size, number of connections (pinouts), and power dissipation. This unit is what we refer to as an element. For clarity, we concentrate on one-to-one (pairwise) connections. (We

should note that some authors have suggested that architectures with one-to-many/many-to-one interconnections may be more advantageous. For instance, see Refs. 68 and 69.)

We now take a look at the factors that determine the smallest size of a three-dimensional computing system with N elements. Heat removal and interconnection density are the two major considerations that give lower bounds on the 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 $\propto N^{1/2}$.[63] The derivation is elementary. In Ref. 63 it is shown that the maximum total power $\mathscr{P}$ 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 $\mathscr{P} \propto N$ when we assume constant power dissipation per element, the linear dimension of the system must grow by at least $\propto \mathscr{P}^{1/2} \propto N^{1/2}$.

In some systems, the power dissipation per element may actually increase with N because the contribution of the interconnections to the 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. (Lower bounds to system size are also implied by power-distribution requirements. These need not be dealt with separately since they imply bounds similar to those for heat removal.[1]

We now turn our attention to the bounds on system size imposed by the space occupied by the elements themselves. The minimum size we impose by arraying the elements throughout a volume is $\propto N^{1/3}$, and the minimum size we impose 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 a 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 nevertheless requires a system size that is at least $\propto N^{1/2}$, 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 that also pays attention to the proportionality constants may be found in Refs. 1 and 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 that the communication paths are still free to use three-dimensional space.[1,2,70,71]

We therefore conclude that arraying the elements and devices on a plane is satisfactory. This is fortu-

nate, since arraying the elements throughout a volume would introduce considerable difficulties to 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,[72] but they are 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. (An exception might arise with a system in which the power dissipation is exceedingly small and the connectivity requirements are low. In this case, it is possible to do considerably better if the elements are arrayed throughout the volume.[72])

Throughout this work, we speak of *device* planes. With this term we refer to planar electronic circuits with optical input–output capability from their surface. (The term smart pixels[73] is also used to describe such device planes, but we find that term to have restrictive connotations and thus avoid using it.) For instance, flip-chip bonding of self-electro-optic-effect devices (SEED's) on silicon[74,75] or other smart-pixel technologies[76–84] would allow the construction of such device planes. A device plane may actually consist of several active device layers sandwiched together so as to constitute an effective single device plane. This would allow greater amounts of silicon circuitry per area if needed.

We now turn our attention to bounds on the 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. We have already discussed the general trade-offs involved in the contention between globally connected and locally connected systems in Subsection 2.C. Actually, there exists a continuum of degrees of connectedness between complete locality and complete globality. Some commonly used quantitative measures of connectedness are reviewed in Ref. 85. However, consideration of the extreme cases is 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 that 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.) do not require system sizes exceeding $\propto N^{1/2}$.[1,2] To understand this, consider an imaginary surface bisecting the system 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 a greater system size than the implementation of a locally connected system. Since there is no trade-off involved, a globally connected graph is 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 that do demand global information flow for their solution. For an introduction to the problem of calculating the amount of information flow needed for the solution of a given problem, we refer the reader to Ref. 58.)

Finally, combining our two arguments we conclude that we prefer globally connected systems with devices arrayed on a plane. (Or on 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 grants us considerable latitude in arraying the elements and providing the interconnections among them. Since heat removal requires that we space the elements considerable distances apart, we might as well utilize this space to array the devices conveniently and also to 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 in which 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 or 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 the proper choice of architecture.[62] However, suboptimal designs may in general result in larger system sizes. Thus we must discuss what types of architecture 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.

C. Free-Space Architectures for Globally Connected Systems

1. Arbitrary Connection Patterns with Multistage Architectures

Having argued in favor of globally connected systems with the elements arrayed on a plane (or on some number of planes), we now explore in more detail the various alternative architectures for providing interconnections between these elements. We 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 discuss so that both the optical sources and detectors lie on the same plane). For simplicity and precision, we 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 of $\propto N^{1/2}$ can be achieved quite straightforwardly by use of three-dimensional fibers or waveguides.[1,2,61,70,71,86] 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 that allow arbitrary patterns of connections to be implemented is the class of architectures that we might term multifacet architectures [Fig. 2(a)]. They 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.[9,61,86] On the other hand, it can be shown that Banyan-type (Fig. 4) multistage architectures can be employed to realize an arbitrary pattern of connections in the theoretically minimum size $\sim N^{1/2}$.[62]

To avoid confusion we must clarify the following point: Multistage architectures are often used as switching networks. Here, we are talking about a hardwired multistage architecture that is used to provide an arbitrary but fixed connection pattern. (Instead of dynamic exchange–bypass switches, we assume hardwired exchange–bypass components that determine the connection pattern.)

As a further comment, let us clarify why we have specified the Banyan among several other multistage networks, such as that based on the perfect shuffle.[87–92] Use of a perfect-shuffle-based network (Fig. 5) 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.[62] 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 are sometimes not specific about which particular regular connection network is used, remembering that the difference is a logarithmic factor in the length of the system (the
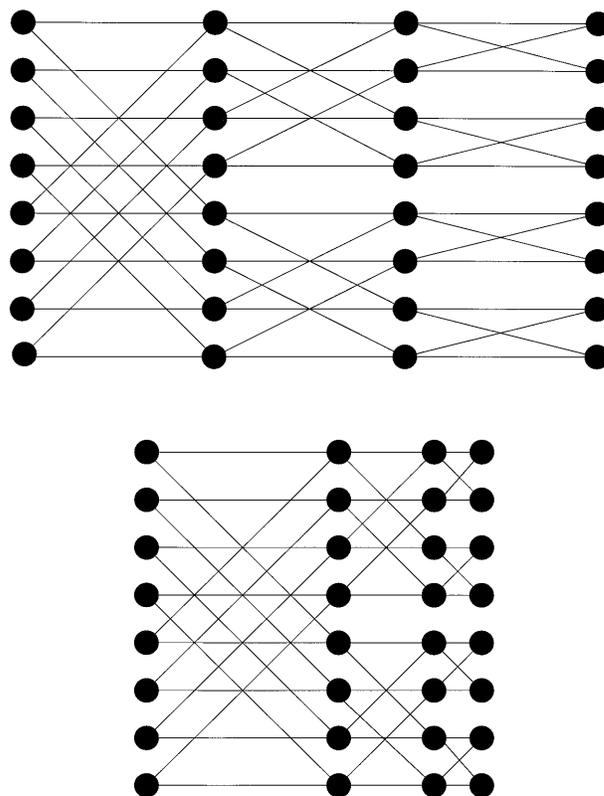

Fig. 4. Regular connection pattern of a one-dimensional Banyan (butterfly) multistage architecture. Top: conventional diagram. Bottom: diagram with angles of all connections drawn equal, which can fit into a box of approximate size $N \times N$. The two-dimensional Banyan is more difficult to draw but is similar in nature. Its optical realization can fit into a box of approximate size $N^{1/2} \times N^{1/2} \times N^{1/2}$.[62]

origin of which is evident on examination of Figs. 4 and 5).

An alternative approach to providing an arbitrary pattern of connections is the optical transpose interconnection system discussed in Refs. 35 and 93. The optical transpose interconnection system is a scalable
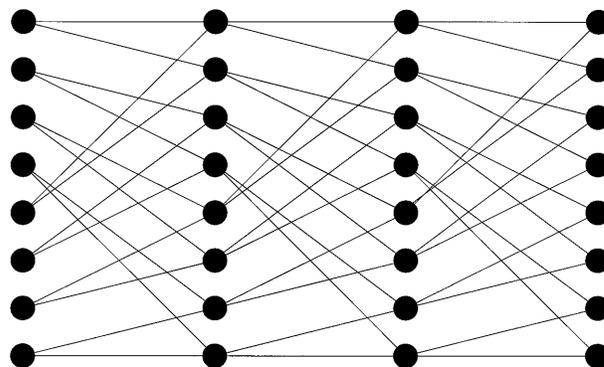

Fig. 5. Regular connection pattern of a one-dimensional perfect-shuffle multistage architecture that can fit into a box of approximate size $N \times N \log N$. The two-dimensional version is more difficult to draw but similar in nature. Its optical realization can fit into a box of approximate size $N^{1/2} \times N^{1/2} \times N^{1/2} \log N$.

dev plane   dev plane

(a)

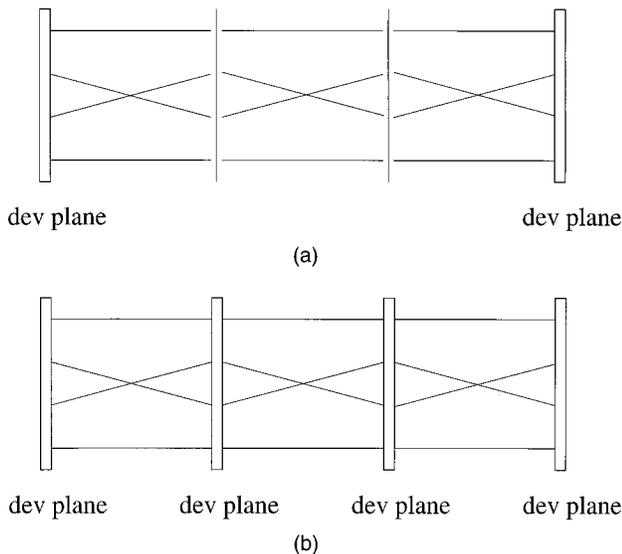dev plane   dev plane   dev plane   dev plane

(b)

Fig. 6. Replacement of passive intermediate planes with active device planes (dev): (a) Schematic diagram of a system in which the end planes house all the active devices. (b) Schematic diagram of a system in which active devices are distributed over all planes.

optical system that provides global connectivity when used with the appropriate electronics. The overall system volume grows by $\propto N^{3/2}$.

2. Introduction of Active Intermediate Planes

Let us consolidate our findings 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 multistage 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 is most likely to be the one to imply the largest proportionality factor and thus to determine the performance and size of the system.

We first consider a system whose length is $N^{1/2} \log N$ (for instance, on the basis of the two-dimensional version of the perfect shuffle shown in Fig. 5). From now on it will be simpler to refer to the schematic depiction shown in Fig. 6(a) rather than to the more detailed connection pattern shown in Fig. 5 or its equivalent for other multistage networks.

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 with 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 interme-

diate planes, especially if these planes must contain regeneration circuits regardless. In other words, if active circuits are needed anyway in the intermediate planes, we might as well make more efficient use of the silicon there. Furthermore, for instance, to construct a random-access parallel computer, we would be interested not merely in an arbitrary fixed connection pattern but in one that 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 in which the intermediate planes are expected to house active devices anyway, the argument in favor of full utilization of the intermediate planes becomes even stronger. Why should we only sparsely utilize the intermediate planes, while the end planes are strained to the limit? It is clearly beneficial to make the computational power uniform throughout all existing planes rather than to concentrate it at the ends and underutilize the intermediate planes. Thus we make the transition from Fig. 6(a) to Fig. 6(b).

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 multiple-device-plane computer with regular connections between its device planes. Such a system is 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 completely ignore 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 emphasize that we are measuring cost in terms of volume and area, not by the 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 were the major determinants of cost, and compare this with integrated circuits, for which essentially only the area counts; wires and devices do not have different costs in this uniform medium.

Introducing $\log N$ times as many circuits to 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. Heat-removal issues are discussed in greater detail in Part II of this paper (see pp. 5697–5705, in this issue).[94]

Our chain of arguments already shows that a system consisting of $\log N$ regularly connected device planes is better than a system based on a multifacet architecture. Nevertheless, a direct comparison would be instructive. It is almost always the case that a system with only regular connections between its planes—with modifications not affecting its essen-

tial 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 that could provide an arbitrary pattern of interconnections by employing a multifaceted interconnection architecture, but the total delay involved would be $\propto N$, since the size and delays of a multifacet system grow $\propto N$. Since $N^{1/2} \log N < N$, the regularly connected system is preferable.

Our argument relies essentially 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 on only 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 able to simulate a permutation network. However, the number 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 designed integrally for the regularly connected system so as to be able to guide 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.

3. Multistage Cascaded Versus Single-Stage Iterative Systems

One of the intrinsic capabilities of multistage systems is the potential for pipelining. New sets of data may be introduced at the left (input) end of the system before the first data set arrives at the right (output) end. (Of course, the whole section shown in Fig. 6(b) may be folded onto itself in an iterative or cyclical fashion or cascaded with similar sections to form a larger pipeline. That is, the object of our argument may be only a building block of some larger system. This would not alter the essence of our argument.)

We now discuss the possibility of collapsing a given multistage system into a system consisting of only one or two stages by assuming that the system is not pipelined. That is, we consider the case in which only one set of data is in transit through the several stages at any given moment. (If the consecutive data sets in a pipelined M-stage system do not interact with each other and travel independently through the pipeline, the same task can be achieved in the same time by employment of M identical nonpipelined systems working in parallel.)

First, let us consider the case in which the circuits and devices in all planes are identical, apart from certain dynamic parameters that can be set in real time. It is then evident that the multistage system can be collapsed into either a two-stage system shuffling the data between its two stages or a single-stage system iterating the data on itself. A simple example is a dynamic $\log N$-stage permutation network.

In the event that the circuits and devices on each stage are not identical, it is still possible to merge them into one or two stages; however this time it is possible that a moderate price would be paid in terms of the total time of computation. Although there are NM elements or devices in an M-stage system with N elements per stage, at most N elements are active at any given time, so that the total power dissipation is at most $\propto N$. Furthermore, the same N optical interconnections can be used for each iteration. Thus both heat-removal and optical interconnection density considerations still imply a system of size $\propto N^{1/2}$, so that the multistage system can be collapsed without a loss of performance.

In most cases the area occupied by the circuits and devices would not imply a system size larger than that dictated by heat removal or optical interconnection density since submicrometer-scaled multilayer circuits do not take up much space. However, if the number of stages M is an increasing function of N, the area occupied by the circuits could ultimately become the determinant of system size. Since the circuits in each stage are not identical, we now have to accommodate NM elements or devices in a single plane, implying a system size $\propto (NM)^{1/2}$. Thus M iterations will take a time of the order of $N^{1/2}M^{3/2}$, which is longer than the time $N^{1/2}M$ that we had for the multistage system. If $M = \log N$, the slowdown is by a factor of $(\log N)^{1/2}$ and may not be considered a very large price to pay if this system is otherwise convenient. Furthermore, if the system is designed as an iterative system in the first place, the actual system size might be much less then $\propto (NM)^{1/2}$ because of potential resource sharing made possible by the proximity of circuits that otherwise would have been situated on different planes.

4. Banyan and Active Intermediate Planes

We have based our argument for introducing active intermediate planes on the schemes shown in Figs. 5 or 6. We said, however, that the Banyan allowed us to achieve an arbitrary pattern of connections in a system smaller than the one shown in these figures by a factor of $\log N$ (Fig. 4). Thus we must reconsider the same line of argument for a Banyan-based multistage system that can be fitted into a box of size $\sim N^{1/2} \times N^{1/2} \times N^{1/2}$.

Although we omit the details, it is possible to argue that either introducing active devices to the intermediate planes or attempting to collapse the system into a single stage, or—as is discussed in the sequel to this

paper[94]—attempting to lay out all the intermediate stages side by side on a single plane results in a loss of the intrinsic $\log N$ advantage of Banyan-based networks in comparison with perfect-shuffle-based networks.

Ultimately, the Banyan-based multistage system allows us to realize an arbitrary (although fixed) interconnection pattern in a box of size $\sim N^{1/2} \times N^{1/2} \times N^{1/2}$. Nearly the same size box is needed to implement each of the regular (e.g., perfect-shuffle) connection patterns appearing in the multistage networks of Figs. 5 or 6. Thus, we can automatically improve on the system depicted in Fig. 6(b) by replacing the regular connection patterns between the device planes with arbitrary fixed connection patterns implemented as Banyans. This would increase the system size by only a factor of the order of unity.[62] However, it is not clear to what extent this added flexibility would translate into an improvement at the user level. Perhaps the same tasks could be realized equally efficiently if the circuits and algorithms were designed appropriately in the first place. (In Section 4 we also argue that platforms based on fixed regular connection patterns not requiring customization might be more beneficial for the development and takeoff of optoelectronic systems technology.)

In addition to the fact that the benefits may be limited (although we do not know), the use of fixed Banyan networks has a number of drawbacks that could discourage us from preferring them to regular connections. First of all, most probably the signal will have to cross $\log N$ passive surfaces, which will result in an increase in attenuation with increasing system size for larger systems. We do not think this will be a major problem. However, it is quite possible that the construction of the Banyan might pose far greater complications and constraints in comparison to a regular pattern of connections. In particular, there might be some difficulties encountered when implementing the hardwired exchange–bypass switches without inflating system size.

From now on we assume the use of regular connection patterns between the device planes. But we do not exclude the possibility of replacing them with fixed Banyan units whenever the advantages of doing so outweigh the disadvantages.

5. 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 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 that implied by heat-removal considerations 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 multistage network and then noted that, if we have a multistage system, 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 multistage 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 on which to build 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 be after at most $\sim \log N$ stages. This number of stages is needed anyway for realizing arbitrary permutations with a multistage network (which takes up less space and results in less signal delay than does a single-stage multifacet arbitrary permutation architecture).

4. Discussion and Conclusion

When 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 have arrived at what seems the best approach. The option we advocate balances the various physical constraints while exploiting the strength of optics as much as possible. It is flexible enough to form the basis of several generic platforms,[94] 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 advocate balanced the major physical requirements nicely. The problem was to determine how much was lost when we restricted ourselves to regular connections. We have argued that we do not lose much. For instance, we argued that any parallel computer algorithm that runs on a reconfigurable permutation network can be distributed through multiple regularly connected stages and that it will be better than realizing the permutation network directly.

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

Such a system would be most successful if one were

to contemplate its higher-level organization and algorithms from the outset, such that it would rely on only a regular pattern of interconnections. Without the benefit of such integral design, simple-minded emulation of algorithms designed to work on architectures that are able to provide arbitrary patterns of connections may be inefficient. Several application platforms based on regularly interconnected device planes are discussed in the sequel to this paper.[94]

It is worth highlighting that customization of such a system involves customization of the electronic circuits in the device planes and whatever software is involved. Unlike the multifacet or fixed multistage architectures 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. This should also considerably simplify computer-aided design tools for optoelectronic systems, such as those described in Refs. 95 and 96.

One final advantage of regularly interconnected device planes is that architectures belonging to this class have already been studied extensively for use in switching systems[97] as well as for other applications.[98] Not only is knowledge of the mathematical aspects well developed,[97] but also optical implementations in the form of switching networks have been demonstrated.[99,100] What we have argued is that such systems are reasonably close to the best possible as defined by physical limitations.

Needless to say, it would be pretentious to claim that the arguments presented in this paper are definitive. And in any event, there are always situations and instances when alternative approaches are feasible or preferable; our arguments aim to capture the mainstream trend. However, with reference to the observation we made at the end of Subsection 1.B, 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 with 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.

References and Notes

1. H. M. Ozaktas, "A physical approach to communication limits in computation," Ph.D. dissertation (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*, S. K. Tewksbury, ed. (Plenum, New York, 1991), Vol. 2, pp. 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. D. A. B. Miller and H. M. Ozaktas, "Limit to the bit rate capacity of electrical interconnects from the aspect ratio of the system architecture," J. Parallel Distrib. Comput. (to be published).
5. J. W. Goodman, F. J. Leonberger, S.-Y. Kung, and R. Athale, "Optical interconnections for VLSI systems," Proc. IEEE **72,** 850–866 (1984).
6. A. W. Lohmann, "What classical optics can do for the digital optical computer," Appl. Opt., **25,** 1543–1549 (1986).
7. J. Jahns and M. J. Murdocca, "Crossover networks and their optical implementation," Appl. Opt. **27,** 3155–3160 (1988).
8. M. R. Feldman, S. C. Esener, C. C. Guest, and S. H. Lee, "Comparison between optical and electrical interconnects based on power and speed considerations," Appl. Opt. **27,** 1742–1751 (1988).
9. 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).
10. D. A. B. Miller, "Optics for low-energy communication inside digital processors: quantum detectors, sources and modulators as efficient impedance converters," Opt. Lett. **14,** 146–148 (1989).
11. N. Streibl, K.-H. Brenner, A. Huang, J. Jahns, J. Jewell, A. W. Lohmann, D. A. B. Miller, M. Murdocca, M. E. Prise, and T. Sizer, "Digital optics," Proc. IEEE **77,** 1954–1969 (1989).
12. A. W. Lohmann and A. S. Marathay, "Globality and speed of optical parallel processors," Appl. Opt. **28,** 3838–3842 (1989).
13. A. W. Lohmann, "Image formation of dilute arrays for optical information processing," Opt. Commun. **86,** 365–370 (1991).
14. A. Louri, "Three-dimensional optical architecture and data-parallel algorithms for massively parallel computing," IEEE Micro., 65–82 (April 24–27, 1991).
15. F. E. Kiamilev, P. Marchand, A. V. Krishnamoorthy, S. C. Esener, and S. H. Lee, "Performance comparison between optoelectronic and VLSI multistage interconnection networks," J. Lightwave Technol. **9,** 1674–1692 (1991).
16. A. Louri, "Optical content-addressable parallel processor: architecture, algorithms, and design concepts," Appl. Opt. **31,** 3241–3258 (1992).
17. A. V. Krishnamoorthy, P. J. Marchand, F. E. Kiamilev, and

S. C. Esener, "Grain-size considerations for optoelectronic multistage interconnection networks," Appl. Opt. **31,** 5480–5507 (1992).

18. G. E. Lohman and K.-H. Brenner, "Space-variance in optical computing," Optik (Stuttgart) **89,** 123–134 (1992).

19. H. M. Ozaktas and J. W. Goodman, "The optimal electromagnetic carrier frequency balancing structural and metrical information densities with respect to heat removal requirements," Opt. Commun. **94,** 13–18 (1992).

20. D. A. B. Miller, "Computing with light," in *1995 Yearbook of Science and the Future* (Encyclopedia Britannica, Chicago, Ill., 1994), pp. 134–147.

21. T. J. Cloonan, "Comparative study of optical and electronic interconnection technologies for large asynchronous transfer mode packet switching applications," Appl. Opt. **33,** 1512–1523 (1994).

22. J. Giglmayr, "Locality and decomposition of regular optical interconnection patterns," Appl. Opt. **33,** 6157–6167 (1994).

23. L. J. Camp, R. Sharma, M. R. Feldman, "Guided-wave and free-space optical interconnects for parallel-processing systems: a comparison," Appl. Opt. **33,** 6168–6180 (1994).

24. T. J. Drabik, "Optoelectronic integrated systems based on free-space interconnects with an arbitrary degree of space variance," Proc. IEEE **82,** 1595–1622 (1994).

25. H. F. Jordan, V. P. Heuring, and R. Feuerstein, "Optoelectronic time-of-flight design and the demonstration of an all-optical, stored program, digital computer," Proc. IEEE **82,** 1678–1689 (1994).

26. M. P. Y. Desmulliez, B. S. Wherrett, J. F. Snowdon, and J. A. B. Dines, "Optical, algorithmic and electronic considerations on the desirable smartness of optical processing pixels," Inst. Phys. Conf. Ser. **139,** 489–492 (1995).

27. K. Wagner, "Fault-tolerant design in digital optical computing," Inst. Phys. Conf. Ser. **139,** 7–10 (1995).

28. S. Araki, M. Kajita, K. Kasahara, K. Kubota, K. Kurihara, I. Redmond, E. Schenfeld, and T. Suzaki, "Massive optical interconnections (MOI): interconnections for massively parallel processing systems," in *Optical Computing*, Vol. 10 of 1995 OSA Technical Digest Series (Optical Society of America, Washington, D.C., 1995), pp. 8–10.

29. V. N. Morozov, J. A. Neff, H. Temkin, A. S. Fedor, "Analysis of a three-dimensional computer optical scheme based on bidirectional free-space optical interconnects," Opt. Eng. **34,** 523–534 (1995).

30. J. Fan, B. Catanzaro, V. H. Ozguz, C. K. Cheng, and S. H. Lee, "Design considerations and algorithms for partitioning optoelectronic multichip modules," Appl. Opt. **34,** 3116–3127 (1995).

31. L. J. Irakliotis, S. A. Feld, F. R. Beyette, Jr., P. A. Mitkas, and C. W. Wilmsen, "Optoelectronic parallel processing with surface-emitting lasers and free-space interconnects," J. Lightwave Technol. **13,** 1074–1084 (1995).

32. A. V. Krishnamoorthy and D. A. B. Miller, "Scaling optoelectronic–VLSI circuits into the 21st century: a technology roadmap," IEEE J. Select. Top. Quantum Electron. **2,** 55–76 (1996).

33. A. Louri and S. Furlonge, "Feasibility study of a scalable optical interconnection network for massively parallel processing systems," Appl. Opt. **35,** 1296–1308 (1996).

34. A. Louri, S. Furlonge, and C. Neocleous, "Experimental demonstration of the optical multimesh hypercube: scalable interconnection network for multiprocessors and multicomputers," Appl. Opt. **35,** 6909–6919 (1995).

35. P. J. Marchand, A. V. Krishnamoorthy, G. I. Yayla, S. C. Esener, and U. Efron, "Optically augmented 3-D computer: system technology and architecture," J. Parallel Distrib. Comput. (to be published).

36. A. V. Krishnamoorthy and D. A. B. Miller, "Firehose architectures for free-space optically interconnected VLSI circuits," J. Parallel Distrib. Comput. (to be published).

37. Special issue on Optical Computing Systems, Proc. IEEE, **82,**(11) (1994).

38. Special issue on Optical computing, Appl. Opt. **33,** (8) (1994).

39. Special issue on Optical computing, Appl. Opt. **35,** (8) (1996).

40. Special issue on Massively Parallel Processing Using Optical Interconnections, J. Parallel Distrib. Comput. (1997).

41. E. Schenfeld, ed., *Proceedings of the Workshop on Massively Parallel Processing with Optical Interconnections*, (IEEE Computer Society, Los Alamitos, Calif., 1994).

42. E. Schenfeld, ed., *Proceedings of the Second International Conference on Massively Parallel Processing Using Optical Interconnections*, (IEEE Computer Society, Los Alamitos, Calif., 1995).

43. A. Gottlieb, Y. Li, and E. Schenfeld, eds., *Proceedings of the Third International Conference on Massively Parallel Processing Using Optical Interconnections*, (IEEE Computer Society, Los Alamitos, Calif., 1996).

44. Digital electronics as it is known today may be defined as that which employs a collection of some kind of nonlinear switches or logic elements interconnected to each other according to some kind of circuit with the aid of electrically conducting paths.

45. H. B. Bakoglu, *Circuits, Interconnections, and Packaging for VLSI* (Addison-Wesley, Reading, Mass., 1990).

46. C. W. Stirk, "Cost models of components for free-space optically-interconnected systems," in *Photonics for Computers, Neural Networks, and Memories*, W. J. Miceli, J. A. Neffs, and S. T. Kowel, eds., Proc. SPIE **1773,** 231–241 (1993).

47. H. M. Ozaktas and J. W. Goodman, "Elements of a hybrid interconnection theory," Appl. Opt. **33,** 2968–2987 (1994).

48. R. W. Keyes, "The wire-limited logic chip," J. Solid State Circuits **17,** 1232–1233 (1982).

49. R. W. Keyes, *The Physics of VLSI Systems* (Addison-Wesley, Reading, Mass., 1987).

50. P. M. Solomon, "A comparison of semiconductor devices for high-speed logic," Proc. IEEE **70,** 489–509 (1982).

51. R. W. Keyes, "Physical limits in digital electronics," Proc. IEEE **63,** 740–767 (1975).

52. N. Rosenberg, *Exploring The Black Box: Technology, Economics, and History* (Cambridge Univ. Press, Cambridge, U.K., 1994).

53. This argument is based on a paragraph from a source we are now unable to identify.

54. H. M. Ozaktas and J. W. Goodman, "Implications of interconnection theory for optical digital computing," Appl. Opt. **31,** 5559–5567 (1992).

55. H. M. Ozaktas, "Levels of abstraction in computing systems and optical interconnection technology," in *Optical Interconnections and Parallel Processing: The Interface*, P. Berthomé and A. Ferreira, eds. (Kluwer Academic, Dordrecht, The Netherlands, 1997).

56. A. S. Tanenbaum, *Structured Computer Organization*, 3rd ed. (Prentice-Hall, Englewood Cliffs, N.J., 1990).

57. The use of the term architecture in this study is more similar to its common use, which refers to the design of physical structures, etc., rather than to its use in computer science, which may often refer to functional organization.

58. J. D. Ullman, *Computational Aspects of VLSI* (Computer Science, Rockville, Md., 1984).

59. H. M. Ozaktas and J. W. Goodman, "Organization of information flow in computation for efficient utilization of high information flux communication media," Opt. Commun. **89,** 178–182 (1992).

60. H. M. Ozaktas and J. W. Goodman, "Lower bound for the communication volume required for an optically interconnected array of points," J. Opt. Soc. Am. A **7,** 2100–2106 (1990).

61. H. M. Ozaktas, Y. Amitai, and J. W. Goodman, "Comparison of system size for some optical interconnection architectures and the folded multifacet architecture," Opt. Commun. **82,** 225–228 (1991).

62. H. M. Ozaktas and D. Mendlovic, "Multistage optical interconnection architectures with least possible growth of system size," Opt. Lett. **18,** 296–298 (1993).

63. 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. Electron. **73,** 1227–1232, (1992).

64. H. M. Ozaktas, K.-H. Brenner, and A. W. Lohmann, "Interpretation of the space–bandwidth product as the entropy of distinct connection patterns in multifacet optical interconnection architectures," J. Opt. Soc. Am. A **10,** 418–422 (1993).

65. G. Önal, A. Altıntaş, and H. M. Ozaktas, "Computer-aided analysis and simulation of complex passive integrated optical circuits of arbitrary rectilinear topology," Opt. Eng. **33,** 1596–1603 (1994).

66. J. Jahns and S. J. Walker, "Imaging with planar optical systems," Opt. Commun., **76,** 313–317 (1990).

67. J. Jahns, "Planar packaging of free-space optical interconnections," Proc. IEEE **82,** 1623–1631 (1994).

68. P. S. Guilfoyle and J. M. Hessenbruch, "Reconfigurable N^4 optical interconnects using the PROMAC architecture," Int. J. Opt. Mem. Neural Networks **3,** 99–109 (1994).

69. P. S. Guilfoyle, "Digital optical computing architectures for compute intensive applications," Inst. Phys. Conf. Ser. **139,** 285–288 (1995).

70. A. L. Rosenberg, "Three-dimensional VLSI: a case study," J. Assoc. Comput. Mach. **30,** 397–416 (1983).

71. F. T. Leighton and A. L. Rosenberg, "Three-dimensional circuit layouts," J. Comput. Sys. Sci. **15,** 793–813 (1986).

72. 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); erratum, **88,** 569 (1992).

73. Special issue on Smart Pixels, J. Quantum Electron. **29,**(2), 1993.

74. K. W. Goossen, J. E. Cunningham, and W. Y. Jan, "GaAs 850 modulators solder-bonded to silicon," IEEE Photon. Technol. Lett. **5,** 776–778 (1993).

75. 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," Photon. Technol. Lett. **7,** 360–362 (1995).

76. S. Esener, "Smart pixels: technology and applications to parallel computing," in *Spatial Light Modulator Technology*, U. Efron, ed., (Marcel Dekker, New York, 1994).

77. D. J. McKnight, M. A. Follett, and K. M. Johnson, "Liquid crystal over silicon spatial light modulators," Inst. Phy. Conf. Ser. **139,** 535–538 (1995).

78. M. P. Y. Desmulliez, J. F. Snowdon, A. J. Waddie, and B. S. Wherrett, "Critical issues in smart pixel design," in *Optical Computing*, Vol. 10 of 1995 OSA Technical Digest Series (Optical Society of America, Washington, D.C., 1995), pp. 96–98.

79. T. L. Worchesky, K. J. Ritter, R. Martin, and B. Lane, "Large arrays of spatial light modulators hybridized to silicon integrated circuits," Appl. Opt. **35,** 1180–1186 (1996).

80. M. K. Smit, "Compact components for semiconductor photonic switches," in *Proceedings of the 1996 International Topical Meeting on Photonics in Switching*, Sendai, Japan (April 1996), paper PWB1.

81. J. A. Neff, C. Chen, T. McLaren, C.-C. Mao, A. Fedor, W. Berseth, Y. C. Lee, and V. Morozov. VCSEL/CMOS smart pixel arrays for free-space optical interconnects. in *Proceedings of the Third International Conference on Massively Parallel Processing Using Optical Interconnections* (*MPPOI '96*), (IEEE Computer Society, Los Alamitos, Calif., 1996), pp. 282–289.

82. T. Kurokawa and T. Ikegami, "Optical interconnection technologies based on vertical-cavity surface-emitting lasers and smart pixels," in *Proceedings of the Third International Conference on Massively Parallel Processing Using Optical Interconnections* (*MPPOI '96*) (IEEE Computer Society, Los Alamitos, Calif., 1996), pp. 300–305.

83. A. Kirk, H. Thienpoint, V. Baukens, N. Debaes, A. Goulet, P. Heremans, M. Kuijk, G. Borghs, R. Vounckx, and I. Veretennicoff, "Demonstration of parallel optical data input for arrays of PnpN optical thyristors," in *Proceedings of the Third International Conference on Massively Parallel Processing Using Optical Interconnections* (*MPPOI '96*) (IEEE Computer Society, Los Alamitos, Calif., 1996), pp. 360–366.

84. A. Z. Shang and F. A. P. Tooley, "Design of smart pixel receivers and transmitters for free-space optical backplane." paper presented at the Optical Society of America Annual Meeting, 20–25 October 1996, Rochester, N.Y.

85. H. M. Ozaktas, "Paradigms of connectivity for computer circuits and networks," Opt. Eng. **31,** 1563–1567 (1992).

86. D. Mendlovic and H. M. Ozaktas, "Optical-coordinate transformation methods and optical-interconnection architectures," Appl. Opt. **32,** 5119–5124 (1993).

87. A. Lohmann, G. Stucke, and W. Stork, "Optical perfect shuffle," Appl. Opt. **25,** 1530–1531 (1986).

88. K.-H. Brenner and A. Huang, "Optical implementations of the perfect shuffle interconnection," Appl. Opt. **27,** 135–137 (1988).

89. C. Stirk, R. A. Athale, and M. W. Haney, "Folded perfect shuffle optical processor," Appl. Opt. **27,** 202–203 (1988).

90. Q. W. Song and F. T. S. Yu, "Generalized perfect shuffle using optical spatial filtering," Appl. Opt. **27,** 1222–1223 (1988).

91. M. W. Haney and J. J. Levy, "Optically efficient free-space folded perfect-shuffle network," Appl. Opt. **30,** 2833–2840 (1991).

92. M. W. Haney, "Pipelined optoelectronic free-space permutation network," Opt. Lett. **17,** 282–284 (1992).

93. G. C. Marsden, P. J. Marchand, P. Harvey, and S. C. Esener, "Optical transpose interconnection system architectures," Opt. Lett. **18,** 1083–1085 (1993).

94. H. M. Ozaktas, "Toward an optimal foundation architecture for optoelectronic computing. Part II. Physical construction and application platforms," Appl. Opt. **36,** 5697–5705 (1997).

95. J. Fan, B. Catanzaro, F. Kiamilev, S. Esener, and S. H. Lee, "The architecture of an integrated computer aided design system for optoelectronics," Opt. Eng. **33,** 1571–1580 (1994).

96. D. Fey, M. Degenkolb, G. Grimm, D. Herzog, and T. Körbs, "HADLOP—A design and simulator tool for digital optoelectronic systems," in *Proceedings of the 1996 International Topical Meeting on Optical Computing*, Sendai, Japan (April 1996).

97. H. S. Hinton, *Introduction to Photonic Switching Fabrics* (Plenum, New York, 1993).

98. K. S. Huang, C. B. Kuznia, B. K. Jenkins, and A. A. Sawchuk, "Parallel architectures for digital optical cellular image processing," Proc. IEEE **82,** 1711–1723 (1994).

99. F. B. McCormick, T. J. Cloonan, A. L. Lentine, J. M. Sasian, R. L. Morrison, M. G. Beckman, S. L. Walker, M. J. Wojcik, S. J. Hinterlong, R. J. Crisci. R. A. Novotny, and H. S. Hinton, "Five-stage free-space optical switching network with field-effect transistor self-electro-optic-effect-device smart-pixel arrays," Appl. Opt. **33,** 1601–1618 (1994).

100. H. S. Hinton, T. J. Cloonan, F. B. McCormick, A. L. Lentine, and F. A. P. Tooley, "Free-space digital optical systems," Proc. IEEE **82,** 1632–1649 (1994).