Implications of interconnection theory for optical digital computing

Haldun M. Ozaktas and Joseph W. Goodman

Heat removal, rather than finite interconnect density, is the major mechanism that limits how densely we can pack three-dimensional computing systems of increasing numbers of elements. Thus highly interconnected approaches can be employed without a further increase in system size. The use of optical interconnections for implementing the longer connections of such systems is advantageous. In fact, if the optical communication energy is sufficiently low and large-bit repetition rates are employed, conductors are useful for only the shortest connections and can be dispensed with altogether with little disadvantage. This justifies consideration of an optical digital computer.

This paper is an initial attempt to understand whether and when an all-optical digital computer may prove useful. Several researchers have addressed this issue in the past, often with negative conclusions. We believe that an increasing understanding regarding the importance of communication in computing and the realization that the architectural-logical construction of a computing system can no longer be divorced from its physical construction justifies a reevaluation of previous arguments and a search for hitherto unexplored perspectives.

One cannot be overcautious in interpreting our discussion. Such studies can never be definitive and our arguments unavoidably rest on floating ground. We nevertheless present the following with the hope that it will provide a seed for further investigations.

1. Introduction

The possibility of an optical digital computer has attracted considerable attention. Despite the vast literature on devices, systems, architectures, and algorithms for optical digital computing,[1] there has been considerable distrust of its usefulness compared with other approaches. Switching-energy arguments on power-delay diagrams, as in Refs. 2 and 3, have resulted in the digital optical computer being viewed mostly as an esoteric curiosity. In the 1980's, interest in optical computing shifted toward optical interconnections.[4] This led to the notion of the hybrid optoelectronic computer. Nevertheless the assumptions underlying the negative arguments have not remained unchallenged[5] and optical computing research has been accelerating. Here we explore several issues relating to the potential usefulness of an optical digital computer and try to identify the conditions under which it would make sense to consider it as an alternative to existing approaches.

A digital computer can be viewed as a collection of nonlinear switching elements that are interconnected according to a certain graph. The nonlinear switching elements often rely on electronic interaction. Usually this is true even of so-called optical switches, since, in most optical switches, photons interact indirectly through electrons. On the other hand, communication along the elements is often established through photons, even in conventional computers.[6] The question is whether conductors are used to confine the wave fields. Thus we define an all-optical computer as one that does not employ conducting materials for the purpose of guiding signals among its switching elements. The all-optical computer, as defined, is a special case of hybrid optoelectronic computer that employs both conducting wires and optical communication for this purpose.

In Sections 2 and 3 we describe the connectivity and heat removal models used in this paper. After discussing how we characterize and evaluate our systems, in Section 4, we mention some crucial properties of various wiring technologies in Section 5. Then we discuss the physical mechanisms that limit the density at which we can pack the elements of our computing system in Section 6. In Sections 7–10 we

H. M. Ozaktas is with the Department of Electrical Engineering, Bilkent University, 06533 Bilkent, Ankara, Turkey; J. W. Goodman is with the Department of Electrical Engineering, Stanford University, Stanford, California 94305.

Received 3 June 1991.
0003-6935/92/265559-09$05.00/0.

argue that an all-optical computer need not be much worse than an optimal hybrid computer under certain conditions, so that an all-optical computer might be preferred for its relative simplicity compared with a hybrid computer. In Section 11 we discuss some objections to our argument. Sections 12 and 13 briefly discuss some issues pertaining to the construction and utilization of an all-optical computer. Section 14 summarizes our main argument.

2. Connectivity Model

Consider a simple model computing system consisting of N elements (graph nodes) laid out on an e dimensional ($e = 2$ or $e = 3$) regular cartesian grid with $N^{1/e}$ elements along an edge (Fig. 1). Let there be an average of k connections (graph edges) per element in our system. The interelement spacing is denoted by d.

We use the parameter $0 \leq p \leq 1$, which is known as the Rent exponent, to quantify the connectivity of our layouts. Refer to Fig. 2; Rent's rule states that kN'^p graph edges emanate from a group of N' elements with k edges each. The rule does not apply when N' is close to N, the total number of elements in the system. Statistical variations are expected. Rent's rule was originally established as an empirical relationship[7,8] and later shown to be a consequence of the logic design process.[9,10] Such a power law may also be justified based on a principle of self-similarity.[11,12] We now understand that Rent's rule is also related to the separator concept of a VLSI complexity theory,[13,14] which provides a formal basis for the layout of given graphs, and to the theory of fractals.[15,16] This relationship has been used widely as a wiring model for two decades.[17,18] In short, Rent's rule is a useful paradigm of connectivity that enables us to quantify to first order the communication requirements of computing systems.

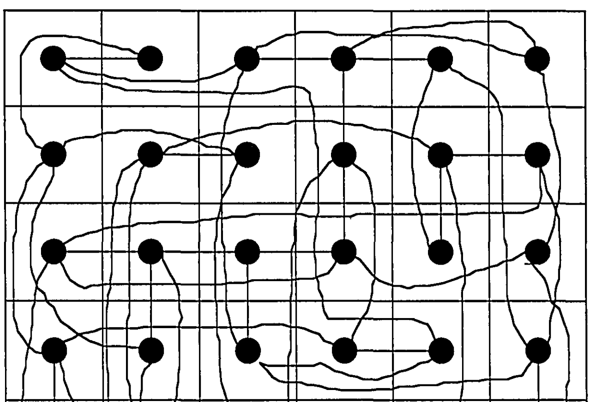

Fig. 1. Model computing system in $e = 2$ dimensions.

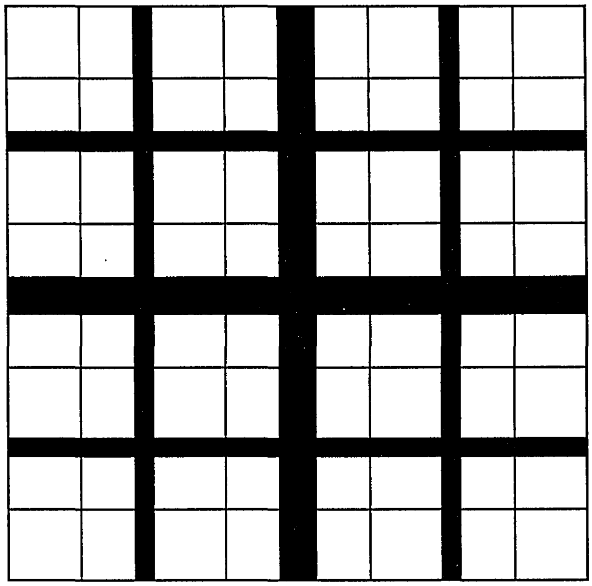

Fig. 2. Binary hierarchical partitioning of the elements. A group at the ith level has $N' = N/4^i$ elements.

An equivalent measure of connectivity is the fractal dimension of information flow n, which is related to the Rent exponent through the relation $p = (n - 1)/n$ (see Refs. 16 and 19).

Rent's rule is often broken in present digital systems, often because of the limitations of existing technology. It is not possible to accommodate the number of wires and input–output ports required by Rent's rule at higher levels of the interconnection hierarchy. As a result, serialism is employed and performance is sacrificed. We consider the implementation of layouts with p uniform throughout the hierarchy. We speak of systems with large values of p (or n) as highly interconnected. The larger p is, the greater is the fraction of longer connections in our system.

It is possible to show that Rent's rule implies an average connection length of $\bar{l} \simeq N^{1/e-1/n}d = N^{p-(e-1)/e}d$ when $n > e$ (i.e., when $p > (e - 1)/e$) (see Refs. 11, 20, and 21). Here l denotes the length of a particular interconnection and the overbar denotes averaging over all connections.

3. Heat Removal

Whatever the modality (conduction, convection, or radiation), heat transfer can take place only through a surface (energy conservation). Thus we assume that our heat removal capability can be characterized by a quantity Q, the maximum power that can be removed per unit cross-sectional area. This is most easily visualized by considering the flow of a cooling fluid through our system, as illustrated in Fig. 3. A fluid with heat capacity C_s and mass density ρ_m flowing at an effective mean velocity v_f may carry away at most $Q = v_f \rho_m C_s \Delta T$, where $\Delta T = T_{\max} - T_{\text{init}}$. $T_{\max}$ is the maximum permissible operating temperature of the devices and T_{init} is the initial temperature of the coolant.

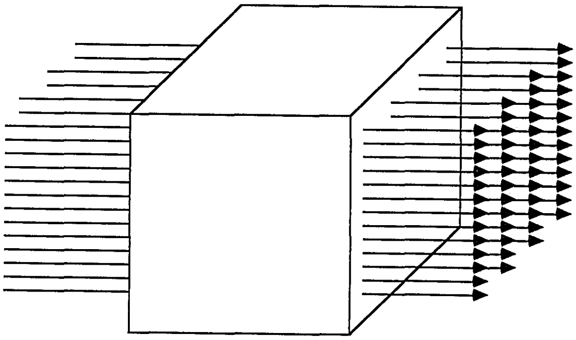

Fig. 3. Three-dimensional heat removal.

The crucial point is that the maximum power we can remove from a system is proportional to its cross-sectional area and not to its volume. Thus the edge length $\mathscr{L}$ of a cubic system dissipating a total power $\mathscr{P}$ must satisfy $Q\mathscr{L}^2 \geq \mathscr{P}$.

4. System Characterization and Evaluation

We characterize our processing systems systems with 3 parameters:

(1) The number of elements N;

(2) The bit repetition rate B along the connections. For simplicity we assume that the bit rate is the same for all connections;

(3) The inverse delay $S = 1/\tau$ across the linear extent of the system. τ will, in general, be the sum of a propagation (speed of light) component and a device component. The number of device delays that are incurred while an influence traverses the extent of the system is fixed by graph topology. For given topology, the best we can do is to minimize propagation delay, which is equivalent to minimizing the system linear extent.

How much we can increase all three of these quantities simultaneously depends on interconnect technology and is ultimately limited by the laws of physics. The particular application–algorithm at hand determines which we prefer to increase at the cost of the other(s).

Our first priority is to maximize performance. Quantitatively, we try to maximize S (i.e., minimize τ) for given N and B. (Our choice of N and B as the independent variables is arbitrary.) If additional degrees of freedom are left in our design, we try to minimize total power consumption.

Another figure of merit of performance that jointly emphasizes parallelism, connectivity, and bandwidth is the bisection–bandwidth product, which is defined as the amount of information we can transfer in unit time across an imaginary surface, dividing our system in two. This figure of merit is particularly suited for certain communication-limited applications.[13,22] In terms of our parameters, the bisection H of our system is given approximately by $H \simeq kN^p$ so that the bisection–bandwidth product can be expressed as $HB \simeq kN^pB$.[23]

5. Wiring Technologies and Their Properties

The most common choice for computer interconnections involves the use of normal conductors. Normal conductors offer good confinement and submicrometer scaling. However, by nature they are lossy so that the energy per transmitted bit and cross section per independent channel must increase with increasing line length.[24–29] As a simple model, we take the energy per bit for a normally conducting line of length l as $E_n = \gamma l$, where $\gamma \simeq 100$ fJ/mm (see Ref. 26). We have assumed that an unterminated line is charged up to 1 V. The corresponding expression for terminated lines is more complicated; however, it is true, in general, that the energy per transmitted bit must increase with line length because of resistive loss. The increase in the cross-sectional area may be overcome by suffering a quadratic growth of delay with length or by the use of repeaters. We do not need these results however.

Optical interconnections rely on only dielectric inhomogenieties for confinement so that they suffer little loss compared with normal conductors. For this very reason, however, the minimum cross-sectional area per independent spatial channel must be greater than the wavelength λ squared.[29] Furthermore, again for the same reason, some coupling and radiation losses (spilling photons) are unavoidable. In conclusion, the energy and the cross section for optical interconnections are large but roughly independent of connection length for the length scales involved in a computing environment.

As a direct outcome of the above discussion, we conclude that normal conductors are better for shorter connections whereas optics is better for longer ones.

In this paper we do not consider the use of superconductors. At first, it may seem that superconductors offer the best of both worlds, as they can provide conductor confinement without loss. However, they suffer from critical current limitations and reduced velocity with scaling. We satisfy ourselves by noting that they offer an effective performance similar to that of optical communication for comparable communication energies.[23]

6. Lower Bounds on System Size

We already have noted that minimizing global propagation delays is equivalent to minimizing the system linear extent. Two major physical considerations lead to lower bounds on system size:

1. Wireability requirements.
2. Heat removal requirements.

For bounded degree connection graphs, wireability requirements dictate that the system linear extent $\mathscr{L}$ grow as $\mathscr{L} \propto N^q$ for e = two-dimensional systems and $\mathscr{L} \propto N^{q/2}$ for e = three-dimensional systems, where $q = \max[p, (e-1)/e]$ (see Ref. 30).

If we assume a constant power dissipation per element, heat removal requirements dictate that the system linear extent grow as $\mathscr{L} \propto N^{1/2}$, since the surface area of both two-dimensional and three-dimensional systems grow as the square of their linear extent.

Since $p \leq 1$, we conclude that with increasing N (unless $p = 1$), the linear extent of three-dimensional systems will become limited by heat removal considerations. In some situations, the power dissipation per element may increase with system size, strengthening our conclusion. On the other hand, if the power dissipation per element were to decrease with increasing N, this might invalidate our conclusion. We come back to this point in Section 9.

Given that there are upper limits to device speed, the best we can do to construct ever-powerful computing systems is to make them three-dimensional with an ever-increasing number of elements N. Thus, based on the argument of the preceding paragraph, we conclude that heat removal will be the major factor that determines how densely we can pack the elements of our system. This means that employing layouts with a larger fractal dimension (or Rent exponent) will not result in a greater system size and delay (still assuming constant power dissipation per element). This suggests that highly interconnected approaches, which offer greater functional flexibility, will be preferred in future large-scale systems.

Of course, there are applications that intrinsically require only limited or local communication and that would not benefit from the opportunity for direct global communication, even if it were available at no cost. However, there are other situations in which intermediate results of computation depend on information located at distant parts of the system. Such applications would benefit from direct global connections rather than from relying on the indirect transfer of information or influences over local connections by means of several elements.

7. Optimal Hybrid Partitioning

Minimization of system size and propagation delay is possible with a hybrid layout, which involves both normally conducting and optical interconnections.[31] Let us consider an e = three-dimensional system with $N = 10^6$, $k = 5$, and $p = 0.8$ ($n = 5$) operated at $B = 10$ Gbit/s. Let the optical communication energy be $E_o = 1$ pJ and the electrical energy be $\gamma = 100$ fJ/mm per unit length, so that the energy per bit for a normally conducting line of length l is $E_n = \gamma l$ (see Ref. 26). For this choice of parameters, optical interconnections require less energy per transmitted bit for lines longer than $l = 1$ cm. Finally, let us be capable of removing $Q = 10$ W/cm^2 of power per unit cross section of our system.

We now show that, under the above assumptions, employing an optimal hybrid combination of normal conductors and optics results in a system with a linear extent $\mathscr{L} \simeq 0.5$ m. The use of any other than the optimal mix of optics and normal conductors will result in a larger linear extent.

In the following simplified derivations we consider only the effects of heat removal. Wiring density, bandwidth, and rise time considerations are actually all coupled to energy considerations. However, since heat removal is the dominating consideration, detailed calculations give similar results.

We consider that a total of N elements is partitioned into N/N_1 groups of N_1 elements each. Connections internal to a group are established with conducting wires and external connections are established with optics.

First we find the minimum size $\mathscr{L}_1$ of an electrically connected cube of N_1 elements. The average connection length per element is given by $k\bar{l} \simeq kN_1^{p-2/3}d$ for an e = three-dimensional layout. Thus the power dissipation per element is $\gamma(kN_1^{p-2/3}d)B$. Also $d = \mathscr{L}_1/N_1^{1/3}$. Heat removal requires that the total power dissipation associated with the N_1 elements not exceed $Q\mathscr{L}_1^2$, where Q is the amount of power we can remove per unit cross section. Thus

$$Q\mathscr{L}_1^2 \geq N_1 k N_1^{p-2/3} \frac{\mathscr{L}_1}{N_1^{1/3}} \gamma B, \tag{1}$$

giving $\mathscr{L}_1 \geq kN_1^p \gamma B/Q$ and a total power dissipation of $\mathscr{P}_1 = (kN_1^p\gamma B)^2/Q$ per electrically connected cube of N_1 elements.

There are kN_1^p edges and hence the same number of optical connections per each cube of N_1 elements (Rent's rule). Thus the global heat removal condition requires that the system linear extent $\mathscr{L}$ satisfy

$$Q\mathscr{L}^2 \geq \frac{N}{N_1}(kN_1^p E_o B + \mathscr{P}_1). \tag{2}$$

The dependence of $\mathscr{L}$ on N_1 is shown in Fig. 4.

By substituting for $\mathscr{P}_1$ and minimizing over N_1 we find

$$\mathscr{L} \geq N^{1/2}\left(\frac{kB}{Q}\right)^{1/2p} E_o^{(p-1/2)/p}\gamma^{(1-p)/p}y^{1/2}, \tag{3}$$

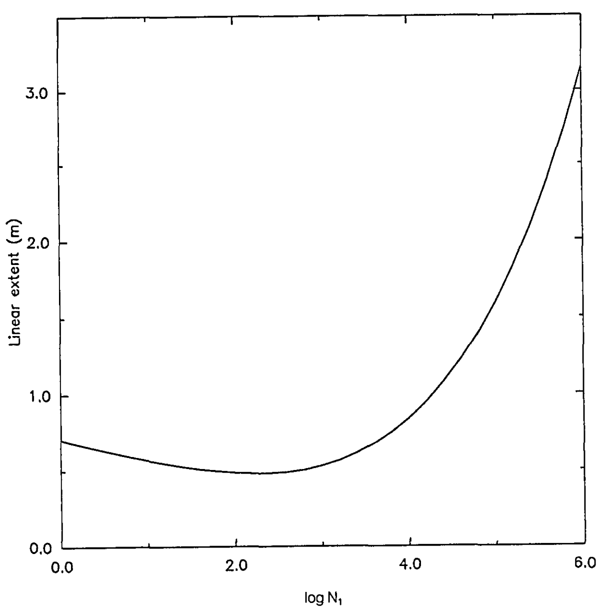

Fig. 4. Dependence of $\mathscr{L}$ on N_1 [relation (2)].

where $y = x^{(p-1)/p} + x^{(2p-1)/p}$ and $x = (1 - p)/(2p - 1)$. Notice that whereas a lower value of p would result in smaller $\mathscr{L}$, it would not change the dependence of $\mathscr{L}$ on N, which is still $\propto N^{1/2}$. Notice also that $\mathscr{L} > (N/N_1)^{1/3}\mathscr{L}_1$, i.e., we have relaxed the requirement of a uniform array of elements (Fig. 1). Permitting the clustering of the electrically connected elements increases the ratio of electrical interconnections, enabling an overall power savings. With our chosen parameters we find the optimal value $N_1 \simeq 190$, which results in a system linear extent $\mathscr{L} \simeq 0.5$ m.

We would now like to determine how worse off we are if we make all connections optically. In this case, the total power dissipation is kNE_oB, which leads to a system linear extent of $\mathscr{L} = (kNE_oB/Q)^{1/2} = 0.7$ m and a power dissipation that is twice as much as the optimal hybrid system. It is also possible to show that making all connections electrical would result in a system linear extent in excess of 3 m.

For the chosen value of p, it is readily shown that the ratio of the linear extent of the optimal hybrid system to that of the all-optical system is $(kB\gamma^2/QE_o)^{0.125}$ (1.32). Thus the disparity between the two systems is quite insensitive to the parameters involved.

If the disparity between 0.5 m and 0.7 m is not considered to be significant, we might as well make all connections optically. This might simplify the design and construction of our system. It should not be surprising that the all-optically connected system is almost as good as the optimal hybrid system in terms of size, delay, and power. After all, in our example the beneficial use of conductors for the shortest connections is an edge effect and can be neglected.

Until now we have not specified the function of the elements. Given that they have only a small number of pinouts, let us for the moment presume that they are simply switching devices or gates.

What does our all-optically connected system look like at this point? It is an array of N electrical switches or gates with sources/modulators at their outputs and detectors at their inputs.

8. Argument for Optical Digital Computing

The all-optically connected system described at the end of Section 7 already qualified as an all-optical digital computer according to our definition since it does not utilize any conductive wiring for communication among its elements. However, we can do even better by replacing the discrete detector–electrical switch–source/modulator combinations by their integrated versions. We speak of an integrated version of such a combination as an optical switch. A self-electro-optic effect device[32,33] is an example of such a device.

Such an integrated replacement can only reduce the overall energy consumption. Notice that there is no distinction between the optical communication energy and the optical switching energy, as the optical communication energy was that required for the optical modulator and detector, which we have now merged together to be the optical switch.

The derivation given in Section 7 reveals that the disparity between the size of the optimal hybrid system and the all-optical system decreases with increasing B and p. It is also necessary for N to be large for the validity of our analysis. Thus, in general, when N, B, and p are large, an all-optical system is almost as good as the optimal hybrid system.

We already argued that future three-dimensional computing systems of an ever-increasing number of elements will tend to employ large values of p. If, in addition, we assume that large values of B are employed, we see that the conditions stated at the end of the previous paragraph coincide with the trend in constructing increasingly more powerful computing systems.

Large values of B and p are consistent also with other desirable properties of optics: ultrafast switching and interactionless cross through, respectively.

9. Importance of Bit Repetition Rate B

With increasing system size and delays, some algorithms may have a tendency to be bottlenecked by communication latencies, so that working with a high bit repetition rate B may not be of any utility. Thus there may be a tendency to operate at slower rates with increasing N. This would invalidate our argument in many ways:

1. First, as is evident from our analysis, when B is not high the discrepancy between the all-optical implementation and the optimal hybrid implementation increases.
2. If B decreases with N, the power dissipation per element decreases, which weakens our argument that three-dimensional systems become heat removal limited with increasing N.
3. When the value of B is less than the large intrinsic bandwidth of the optical communication channels, it is possible to employ various strategies to exploit this bandwidth in order to reduce system size. In such cases, the use of normal conductors for the shorter connections may be useful.[34]

Thus, for applications for which large values of B are not useful with increasing system size and propagation delays, an all-optical computer would probably not be useful for foreseeable values of N.

10. Comparison of Systems with Different Connectivity

In Section 6 we argued that since e = three-dimensional systems are heat removal limited, the value of p has no effect on the resulting system linear extent. However, since our assumption regarding constant power dissipation per element does not hold for hybrid designs, this is no longer precisely true, as is evident from relation (3). So as to provide a basis for comparison, let us also calculate the linear extent of a system with parameters that are identical to that considered above, except p. Let all connections in this system be to the nearest or second-nearest neighbors only, so that the average connection length per element is $\simeq kd$. (This layout has $n = e = 3$ and $p = 1 - 1/n = 2/3$.) Since all connections are short, we consider making all of them electrical. The heat removal condition may be written as

$$Q\mathscr{L}^2 \geq N\gamma(kd)B. \tag{4}$$

Remembering that $d = \mathscr{L}/N^{1/3}$, we can show that the linear extent of this system is $\mathscr{L} \simeq 0.5$ m. (It can also be shown that nonuniformizing the elements as in the hybrid case does not offer any advantage.)

Let us also consider the optimal hybrid implementation of this system. The derivation is similar to the case in which $p = 0.8$, with $p = 2/3$ instead. We find that the linear extent can be reduced to $\mathscr{L} = 0.27$ m with $N_1 \simeq 2800$.

The bisection–bandwidth product of this system is six times less than that of our original example. Since all connections are to nearest neighbors, the transfer of influences across the extent of this system takes $N^{1/3}$ device delays, which, with 100 ps switches, amounts to 10 ns. This is 1 order of magnitude greater than the speed of light delay across the linear extent $(0.27 \text{ m})/(3 \times 10^8 \text{ m/s}) \simeq 1$ ns. Our original example with $p = 0.8$ may, for instance, be a five-dimensional mesh. $N^{1/5}$ device delays are suffered in traversing the linear extent of such a system, adding up to 1.6 ns, which is comparable to the speed of light delay $(0.7 \text{ m})/(3 \times 10^8 \text{ m/s}) \simeq 2.3$ ns. The total delay is 3.9 ns. Our all-optical five-dimensional system is superior to the optimal hybrid three-dimensional system in terms of both bisection–bandwidth product and global delay, despite its greater linear extent. This strongly suggests that problems or applications exist for which the total time of computation would be less on the all-optical five-dimensional mesh. Of course, an optimal hybrid five-dimensional mesh is even better, but not by a significant amount (total global delay = 3.3 nsec).

Tables 1 and 2 summarize these results, which will be discussed further in Section 13.

11. Some Credible Objections

In this section we address some of many possible objections to our arguments. First, we must not forget that our discussion is based on certain basic physical considerations only. We cannot claim that other issues will not completely swamp them. Other optical considerations such as noise, cross talk, or aberrations may be shown to deem such large systems impossible. And, of course, a myriad of engineering issues must be faced in the construction of a real computer.

Our discussions assume dissipative computing. There is general consensus that dissipationless computing does not contradict the laws of physics.[35–38] This would radically alter the discussions of this paper.

The global delay across the linear extent of the system need not be a good indicator of performance. For instance, for some tasks that are divisible into a large number of relatively independent subtasks, the speed of local communication may be more significant than that of global communication. The worst-case delay may not bottleneck the operation.

A major objection to our argument rests on the following question: what if the elements are not just gates but more complicated circuitry, still of the same size and number of pinouts? In this case, even if making all connections optical is shown to be no worse than the optimal hybrid combination, we cannot replace these detector–element–source/modulator combinations with simple optical switches. In fact, this may be considered as proof that a hybrid system can always have more computational resources than an all-optical system. The size of an optical switch cannot be less than ~ 1 μm, whereas in a micrometer's space, several electrical switches can be squeezed if deep submicrometer technology is employed. However, Rent's rule will be broken; the switches will have limited communication with the outer world. Information flow at the original fractal dimension is not possible through the boundary of the elements and some form of serialism must be employed.

Thus this objection boils down to the question of what usefulness an element of internal sophistication but limited pinouts may have. Of course, an element with a sophisticated internal structure would always be desirable over a simple switch with the same number of pinouts, if it is available at no cost, since it could easily simulate the simple switch. However, the question is whether such elements will have sufficient utility to make their usage worthwhile. A simple example of an element with limited pinouts but with an arbitrarily large internal structure is a shift register, which can serve as a memory. How-

Table 1. System Linear Extent for n = Three-Dimensional and n = Five-Dimensional Meshes laid out in e = Three Dimensions[a]

n	All-Optical $\mathscr{L}$(m)	Optimal Hybrid $\mathscr{L}$ (m)	All-Electrical $\mathscr{L}$(m)
3	0.7	0.27	0.5
5	0.7	0.5	3.2

[a]$N = 10^6$, $k = 5$, $B = 10$ Gbit/s, $E_o = 1$ pJ, $\gamma = 100$ fJ/mm, and $Q = 10$ W/cm^2.

Table 2. Bisection–Bandwidth Product and Global Delay for n = Three-Dimensional and n = Five-Dimensional Meshes Laid Out in e = Three Dimensions[a]

n	All-Optical HB (Tbit/s)	Optimal Hybrid HB (Tbit/s)	All-Optical τ (ns)	Optimal Hybrid τ (ns)
3	500	500	12	11
5	3200	3200	3.9	3.3

[a]$N = 10^6$, $k = 5$, $B = 10$ Gbit/s, $E_o = 1$ pJ, $\gamma = 100$ fJ/mm, $Q = 10$ W/cm^2, and device delay = 100 ps.

ever, this memory will have a large access time. We might argue that a shift register is always preferable to a single gate, since a slow memory is better than none. On the other hand, this memory might have such limited usefulness in certain situations because of its slowness that we might not bother having it.

Another example is a microprocessor that has far less pinouts than Rent's rule implies.[39] Whereas such so-called functionally complete units are consistent with existing system approaches, it is not clear that this is the most advantageous way of constructing supercomputers.[40] It may be that a massive and homogeneous collection of switches connected with a uniform Rent exponent is more beneficial. Alternatively, it may be that optical nodes with some internal sophistication and function integrated into them do have some usefulness. Whether such an integrated structure is still considered to be an optical switch, or rather, an electronic circuit with optical ports, is a matter of definition. Indeed, Lentine *et al.*[41] have noted the fuzziness in the definition of an all-optical computer: "Because of the limited functionality achievable in 'all optical' logic gates, a growing interest is seen in 'optical' processing elements made using optoelectronic devices with greater functionality." Whether or not this trend is merely a legacy of existing design approaches, it seems that the first optical computers and switching systems will employ optical nodes with relatively sophisticated functionality.

We cannot arrive at any general conclusions. Whether limited pinout elements are useful will be determined mainly by the particular application. Nevertheless general observation does seem to indicate that the percolation of information is crucial in the working of both natural and artificial systems.[40]

12. What Might the Optical Computer Look Like?

The example of Section 7 illustrates the order of numerical values for which it is meaningful to consider an all-optical computer. For greater performance we might increase N as much as possible. For instance, a system with 10^8 elements would be 7 m in size.

The vision of the digital optical computer emerging from our considerations is as follows. It is large ($N \sim 10^7$–10^8 and $\mathscr{L} > 1$ m), highly interconnected, operated at large repetition rates (multi Gbit/s) and, because of its size, exhibits large speed-of-light-limited delays between its distant elements (~ 10 ns). Its bisection–bandwidth product might be of the order of 10^{16} bit/s. It is suitable for situations in which a large repetition rate is useful, despite large propagation delays. In other words, it will be appropriate for applications in which one prefers large values of B at the expense of N and S.

It is also interesting to note that the number of bits of information that such a system can remember at any given instant can be greater than the number of switches N since several bits are simultaneously in transit along the longer connections, which serve as memory.

How can such a system be actually constructed? It is possible to show for our original example that there is enough spacing between the elements to wire up the system by using discrete fibers of ~ 1 mm diameter. This approach is intimidating from a constructional viewpoint. Since the system is heat removal limited and we are interested in the worst-case (and not average) delay, it makes little difference if, instead of situating the elements on a three-dimensional grid, we lay them out on a plane, and use the third dimension for the purpose of communication. We simply lay the 10^6 elements ~ 0.7 mm apart in the form of a $10^3 \times 10^3$ array.

It is preferable to use free-space optics rather than discrete fibers. However, conventional imaging systems permit free-space interconnections at high density for a regular pattern of connections only. Since the system is heat removal limited anyway, multifacet holographic approaches[24,42] may be used to provide an arbitrary pattern of connections for smaller N. However, the resulting system size with these approaches is crudely $\mathscr{L} \simeq kN\lambda$ (see Ref. 29) so that, for $N > 10^5$, they would not be desirable.

We do not know if a system with a regular pattern of connections of the same computational power as that of an irregularly connected one is always possible. Huang seems to argue in favor of it[43] whereas the maximum entropy approach of Keyes would seem to indicate otherwise.[18,40]

If the three-dimensional multifacet architecture of Ref. 44 can be built, it would solve once and for all the problem of being restricted to a regular pattern of connections.

13. What is it Good For?

What do we do with an $N > 10^3 \times 10^3$ array of globally connected switches? It does not seem that straightforward mapping of conventional digital logic as it exists in today's electronic computers would make the most of such a system. We do not know what kind of functional implementation (existing or to be discovered) would be best, although one possibility is suggested below.

As mentioned above, a five-dimensional mesh with radix $N^{1/5}$ (see Ref. 45) is an example of a graph with $p = 0.8$ and $k = 5$. This graph may represent the connection pattern of a five-dimensional nearest-neighbor (in five dimensions) connected cellular automaton. In Section 10, we also considered a three-dimensional mesh, which may represent the connection pattern of a three-dimensional nearest-neighbor connected cellular automaton. The resulting linear extent, the bisection–bandwidth product, and the global delay for these systems were presented in Tables 1 and 2. On the basis of a comparison of their bisection–bandwidth products and global delays, we conjecture that concrete problems (such as sorting, etc.) exist for which the time it takes to solve these problems on our five-dimensional examples is less than on our three-dimensional examples, despite the fact that our five-dimensional examples have greater linear extent and connection lengths. If one can find an example of a problem that can be solved in a shorter time on our optimal hybrid five-dimensional mesh of $\mathscr{L}$ = 0.5 m than on the optimal hybrid three-dimensional mesh of $\mathscr{L}$ = 0.27 m, this would serve as proof that indeed, for some applications, highly interconnected approaches are preferable for large e = three-dimensional systems. It would be even more interesting to provide an example problem that can be solved in a shorter time on the all-optical five-dimensional mesh with $\mathscr{L}$ = 0.7 m than on the optimal hybrid three-dimensional mesh with $\mathscr{L}$ = 0.27 m.

Multidimensional cellular automatons offer an alternative to computing based on functionally complete entities (assuming that one finds a better way of performing useful computation other than simulating conventional logic functions in the automatons). The latter approaches allow us to view functionally complete entities at a lower level as black boxes, which greatly facilitates design and construction. Without the benefit of such structuring, the conception and design of large-scale arbitrarily connected systems would be intimidatingly complex. One way to avoid this is to resort to systems that exhibit some form of regularity. This brings us back to the question of whether requiring such regularity takes back any of the potential advantages.

Another alternative to functionally complete approaches is that of neural computing, which is beyond the scope of this discussion.

14. Conclusion

To build increasingly powerful processing systems, we must increase the number of elements N. If the bit repetition rate B along the connections of our system is large, heat removal considerations tend to dominate wireability considerations for three-dimensional systems, which suggests (but does not prove) that highly interconnected approaches may be preferred for increasing parallelism and functional flexibility. For such systems the fraction of connections with lengths greater than the break-even energy between normally conducting and optical interconnections will be large, so that we might as well make all connections optically. It is meaningful to consider such a system because the use of conductive wiring to establish the shorter interconnections will not result in a considerable improvement in system size and global delays. (Nor will it necessarily make them worse, and it is a subjective issue whether it is considered more simple to keep it all optical or to keep the number of optical connections minimum by using conducting wires as much as possible.)

Thus, if large values of B are useful despite large propagation delays, it might be meaningful to consider the construction of an all-optical digital computer.

Needless to say, our arguments cannot be exhaustive or definitive. Our discussions must be looked at as being more of a thought experiment than a conclusive argument. One important implicit assumption that we made is that the Rent exponent is uniform over all hierarchical levels of the system. A more conclusive discussion of this issue will require a deeper understanding of how the solution of a problem relates to the percolation of information at various levels of the system.

Ultimately, more solid answers to the questions raised in this paper will emerge when the theory of algorithms is merged with a physically realistic theory of computer construction to create a physical theory of computation. Then it will be possible to compare and optimize jointly over many possible constructional and algorithmic approaches.

The authors gratefully acknowledge the benefit of discussions with David A.B. Miller of Bell Laboratories. This work was supported in part by the U.S. Air Force Office of Scientific Research under grant AFOSR-88-0024.

References

1. W. Cathey, K. Wagner, and W. J. Miceli, "Digital computing with optics," Proc. IEEE **77**, 1558–1572 (1989).
2. R. W. Keyes, "Power dissipation in information processing," Science **168**, 796–801 (1970).
3. P. W. Smith, "On the physical limits of digital optical switching and logic elements," Bell Syst. Tech. J. **61**, 1975–1993 (1982).
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. A. Huang, "Optical digital computing at Bell Labs," presented at the 1990 International Topical Meeting on Optical Computing, Kobe, 1990.
6. D. A. B. Miller, "Quantum wells for optical logic and interconnection," in *Microelectronic Interconnects and Packaging*, G. Arjavalingam and J. Pazaris, eds., Proc. Soc. Photo-Opt. Instrum. Eng. **1389**, 496–502 (1990).
7. B. S. Landman and R. L. Russo, "On a pin versus block relationship for partitions of logic graphs," IEEE Trans. Comput. **20**, 1469–1479 (1971).
8. R. L. Russo, "On the tradeoff between logic performance and circuit-to-pin ratio for LSI," IEEE Trans. Comput. **21**, 147–153 (1972).
9. W. E. Donath, "Stochastic model of the computer logic design process," Tech. Rep. RC 3136 (IBM Thomas J. Watson Research Center, Yorktown Heights, N.Y., 1970).
10. W. E. Donath, "Equivalence of memory to 'random logic,'" IBM J. Res. Dev. **18**, 401–407 (1974).
11. W. E. Donath, "Placement and average interconnection lengths of computer logic," IEEE Trans. Circuits Syst. **26**, 272–277 (1979).
12. L. Pietronero, "Fractals in physics: introductory concepts," in *Order and Chaos in Nonlinear Physical Systems*, S. Lundqvist, N. H. March, and M. P. Tosi, eds. (Plenum, New York, 1988).
13. J. D. Ullman, *Computational Aspects of VLSI* (Computer Science, Rockville, Md., 1984).
14. C. E. Leiserson, *Area-Efficient VLSI Computation* (MIT, Cambridge, Mass., 1983).
15. B. B. Mandelbrot, *Fractals: Form, Chance and Dimension* (Freeman, San Francisco, Calif., 1977).
16. P. Christie and S. B. Styer, "Fractal description of computer interconnection distributions," in *Microelectronic Interconnects and Packaging: System and Process Integration*, S. K.

Tewksbury, ed., Proc. Soc. Photo-Opt. Instrum. Eng. **1390**, 359–367 (1990).

17. R. R. Tummala and E. J. Rymaszewski, eds., *Microelectronics Packaging Handbook* (Van Nostrand Reinhold, New York, 1989).
18. R. W. Keyes, *The Physics of VLSI Systems* (Addison-Wesley, Reading, Mass., 1987).
19. P. Christie, J. E. Cotter, and A. M. Barrett, "Design and simulation of optically interconnected computer systems," in *Interconnection of High Speed and High Frequency Devices and Systems,* A. P. DeFonzo, ed., Proc. Soc. Photo-Opt. Instrum. Eng. **947**, 19–24 (1989).
20. M. Feuer, "Connectivity of random logic," IEEE Trans. Comput. **31**, 29–33 (1982).
21. W. E. Donath, "Wire length distribution for placements of computer logic," IBM J. Res. Dev. **25**, 152–155 (1981).
22. A. Orlitsky and A. E. Gamal, "Communication complexity," in *Complexity in Information Theory,* Y. S. A.-Mostafa, ed. (Springer-Verlag, New York, 1988).
23. 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. 61–124.
24. R. K. Kostuk, J. W. Goodman, and L. Hesselink, "Optical imaging applied to microelectronic chip-to-chip interconnections," Appl. Opt. **24**, 2851–2858 (1985).
25. 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).
26. 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).
27. R. C. Frye, "Analysis of the trade-offs between conventional and superconducting interconnections," IEEE Circuits Devices Mag. **5**(3), 27–32 (1989).
28. H. Kroger, C. Hilbert, U. Ghoshal, D. Gibson, and L. Smith, "Applications of superconductivity to packaging," IEEE Circuits Devices Mag. **5**(3), 16–21 (1989).
29. 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).
30. 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).
31. H. M. Ozaktas and J. W. Goodman, "Optimal partitioning of very large-scale optoelectronic computing systems," in *Annual Meeting 1990,* Vol. 15 of OSA 1990 Technical Digest Series (Optical Society of America, Washington, D.C., 1990), p. 87.
32. D. A. B. Miller, "Quantum-well self-electro-optic effect devices," Opt. Quantum Electron. **22**, 61–98 (1990).
33. D. A. B. Miller, S. E. Henry, A. C. Gossard, and J. H. English, "Integrated quantum well self-electrooptic effect devices: 2×2 array of optically bistable devices," Appl. Phys. Lett. **49**, 821–823 (1986).
34. H. M. Ozaktas and J. W. Goodman, "Multiplexed hybrid interconnection architectures," in *Optical Computing,* Vol. 6 of OSA Technical Digest Series (Optical Society of America, Washington, D.C., 1990), pp. 45–48.
35. C. H. Bennett and R. Landauer, "The fundamental physical limits of computation," Sci. Am. **253**(1), 48–56 (1985).
36. R. Landauer, "Uncertainty principle and minimal energy dissipation in the computer," Int. J. Theor. Phys. **21**, 283–297 (1982).
37. E. Fredkin and T. Toffoli, "Conservative logic," Int. J. Theor. Phys. **21**, 219–253 (1982).
38. H. J. Caulfield and J. Shamir, "Wave particle duality considerations in optical computing," Appl. Opt. **28**, 2184–2186 (1989).
39. D. K. Ferry, "Interconnection lengths and VLSI," IEEE Circuits Devices Mag. **1**(4), 39–42 (1985).
40. R. W. Keyes, "Communication in computation," Int. J. Theor. Phys. **21**, 263–273 (1982).
41. A. L. Lentine, L. M. F. Chirovsky, M. W. Focht, J. M. Freund, G. D. Guth, R. E. Leibenguth, G. J. Przybylek, L. E. Smith, L. A. D'Asaro, and D. A. B. Miller, "Integrated self electro-optic effect device photonic switching nodes," in *Photonic Switching,* H. S. Hinton and J. W. Goodman, eds., Vol. 8 of OSA Proceedings Series (Optical Society of America, Washington, D. C., 1991), pp. 170–173.
42. M. Feldman, C. C. Guest, T. J. Drabik, and S. C. Esener, "Comparison between electrical and free space optical interconnects for fine grain processor arrays based on interconnect density capabilities," Appl. Opt. **28**, 3820–3829 (1989).
43. A. Huang, "Optical digital computers: devices and architecture," in *Concurrent Computations,* S. K. Tewksbury, ed. (Plenum, New York, 1988), Chap. 2.
44. 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).
45. W. J. Dally, *A VLSI Architecture for Concurrent Data Structures* (Kluwer, Norwell, Mass., 1987).