Elements of a hybrid interconnection theory

Haldun M. Ozaktas and Joseph W. Goodman

We present a textbooklike treatment of hybrid systems employing both optical and electrical interconnections. We investigate how these two different interconnection media can be used in conjunction to realize a system not possible with any alone. More specifically, we determine the optimal mix of optical and normally conducting interconnections maximizing a given figure-of-merit function. We find that optical interconnections have relatively little to offer if the optical paths are constrained to lie on a plane (such as in an integrated optics system). However, if optical paths are permitted to leave the plane, they may enable considerable increase in performance. In any event the prize in terms of performance is accompanied by a penalty in terms of system power and/or size.

Key words: Optical interconnections, optical computing, optoelectronic computing.

1. Introduction

Computing systems are becoming increasingly limited by the signal delay, space consumption, and power dissipation associated with the complex network of resistive wiring that interconnects their switching elements.[1-4] It has been suggested that the use of optical or of superconducting interconnection media might alleviate this trend and enable the construction of computing systems that are superior to what can be constructed by use of normally conducting interconnections alone.[5-13]

In previous research we compared the system size, the signal delay, and the power dissipation of systems employing only one interconnection medium at a time.[14] We found that normally conducting interconnections were preferable for smaller numbers of elements, whereas optical and superconducting interconnections were preferable for larger numbers of elements. This suggests that we can do better by joint use of normal conductors (for the shorter connections) and optics or superconductors (for the longer connections). Indeed, the concept of use of optical interconnections for higher levels of the interconnection hierarchy has received more widespread attention than all optically connected systems. The question is, beyond what point should we start employing optics?

The way this problem was first addressed in the literature was by derivation of a breakeven distance beyond which the use of optical communication was preferable to the use of normal conductors. For instance, Feldman et al.[15] and Miller[9] claimed that optical communication is energetically favorable for connections of length $\ell > 1$ mm or so. There also have been attempts to compare the information density of optical and normally conducting interconnections in a similar manner.[7]

Whereas this approach to comparing various interconnect media can be instructive and useful, it is nevertheless unsatisfactory in many ways. First, it enables comparison of only one quantity at a time without attention being paid to the others. Whether information density or energy is of greater importance depends not only on whether the system is heat-removal, wireability, or device limited, but also on the relative emphasis we give to various measures of performance (signal delay, bit-repetition rate, etc.) and cost (system size and power dissipation). Since the length scale of the system is related to the properties of the interconnections through wireability and heat-removal requirements in a complicated manner, we do not know initially the physical length ℓ of a line of length r in (dimensionless) grid units. The comparison of isolated lines of given length has little meaning when these lines are embedded in a system.

Even the comparison of an all optically connected system with an all electrically connected system (as in Ref. 14) does not tell us which connections to implement optically in a hybrid system. Other research

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

falling into this category is that of Feldman *et al.*, who compare a three-dimensional optical system to a two-dimensional electrical system,[16] that of Stirk and Psaltis, who compare three-dimensional optical and electrical permutation network implementations based on yield considerations,[17] and that of Kiamilev *et al.*[18] The problem of how to use both media in conjunction has received less attention; Krishnamoorthy *et al.* discuss how a perfect-shuffle network should be partitioned into very-large-scale-integrated (VLSI) chips, which are then interconnected optically.[19]

For the reasons discussed above we take a more general approach to this problem. We start with a layout of N elements (gates or switches), which we partition into N/N_1 groups of N_1 elements each (Fig. 1). All connections internal to a group are made electrically, whereas connections between elements in different groups are made optically. Notice that $N_1 = N$ corresponds to an all electrically connected system, whereas $N_1 = 1$ corresponds to an all optically connected system. For given total number of elements N and bit-repetition-rate B we calculate the optimal value of N_1 maximizing our optimization function Γ, which in general can be a function of signal delay, total system size, and power consumption. We are interested mostly in high-performance systems for which system size and power dissipation are only secondary considerations, the primary consideration being minimization of signal delay. We also consider optimization functions putting a greater emphasis on cost of size and power.

In Section 2 we describe the models and the important variables used in this study. Section 3 outlines the general procedure, and Section 4 provides some remarks on the numerical examples. The major results of this paper are derived in Section 5 (two-dimensional systems), Section 6 (the effect of the cost of system size being taken into account), Section 7 (the effect of repeaters being used), and Section 8 (three-dimensional systems); our predominantly analytical treatment is illustrated by numerical examples. Section 9 provides discussion and conclusions.

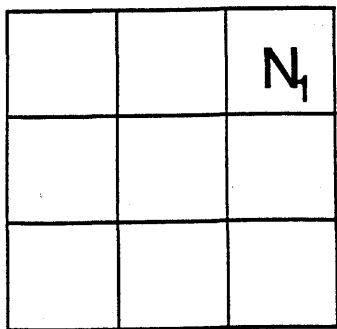Results of this study were first presented in Ref. 20 and subsequently in Ref. 21. This paper is a simplified and condensed version of the more elaborate analysis presented in Ref. 21. Unlike that study, in which the aim was a more general formulation, here we employ the simplest treatment that still leads to the same conclusions, so that our exposition is more transparent and instructive. The reader is referred to that study for extensions and generalizations.

Our analysis involves several approximations. Some are made to maintain analytic simplicity and transparency, others are made to maintain generality. In order to be more exact, we would have to introduce several new parameters and make many arbitrary assumptions. Dependent on the details of the design, additional factors or terms would have to be included in our expressions. We avoided doing so when the error incurred does not exceed a factor of the order of unity, since we believe that there is no point in trying to specify and to keep track of such factors in this type of (general) analysis.

2. Model Description

A. Wireability Limitations and Connectivity Model

The spacing between the elements (switches and gates) of a computing system must be large enough to permit sufficient space for the interconnections to pass between them. Systems employing longer connections require greater interelement spacing than systems employing shorter connections. We speak of systems employing greater fractions of longer connections as being highly connected. Here we present a model that enables us to quantify the connectivity of computer circuits to first order, and we can predict the interelement spacing necessary to ensure that there is enough space for implementing the desired pattern of connections. This is not the only such model possible, nor one that is suited for every possible situation. However, there are strong rational and empirical foundations for adopting it, in addition to its being well suited for the type of analysis we intend to pursue.[22,14,21]

For the purpose of this study a processing system is a collection of N given similar elements connected to each other according to a prespecified graph. A list of symbols is provided as Table 1. k denotes the number of connections (graph edges) per element (for simplicity we are considering pairwise connections only; the extension to fan-in and fan-out is not considered); thus there is a total of kN connections. Within a factor of 2, we may also interpret kN as the total number of input–output ports. d_d denotes the linear extent of the elements. Let the $N \gg 1$ elements comprising our system be laid out on an e (equal to two or three) -dimensional Cartesian grid of as yet unspecified lattice constant d with $N^{1/e}$ elements along each dimension (Fig. 2). Figure 3 depicts a hierarchical partitioning of this array of cells.

During the course of our analyses it is necessary to specify the following quantities in order to obtain explicit results:

(1) The average connection length $\bar{r}$ of the layout (in grid units).

Fig. 1. Partitioning a system of N elements into N/N_1 groups of N_1 elements each ($N/N_1 = 9$).

Table 1. List of Symbols

|	
A	Cross-sectional area associated with each physical line.
$\mathscr{A}$	Area (base area) of two (three) -dimensional system.
B	Bit-repetition rate along each edge of connection graph.
d	Linear extent of a unit cell, identical to interelement spacing.
d_d	Linear extent of an element.
d_m	Center-to-center spacing of modules.
d_{tr}	Linear extent of an optical transducer.
d_1	Linear extent of a module of N_1 elements.
e	Euclidean dimension of layout space.
E	Energy associated with each transmitted bit of information.
f	Identical to W/λ for optical systems.
k	Number of graph edges (connections) per element.
ℓ	Length of a line in real units.
$\mathscr{L}$	Linear extent of the system.
M	Number of interconnection layers.
N	Number of elements.
N_1	Number of elements in each module.
p	Rent exponent of layout.
$\mathscr{P}$	Total power dissipation.
Q	Maximum amount of power we can remove per area.
r	Length of a line in grid units.
S	Inverse of worst-case signal delay.
T_d	Response time of devices.
T_r	Minimum pulse-repetition interval along a line.
V	Nominal voltage level.
W	Transverse linear extent associated with each physical line.
W_{min}	Minimum manufacturable linewidth.
α	Defined in Eq. (9).
β	Defined in Eq. (12).
γ	Defined in Eq. (10).
Γ	Optimization function.
κ	Coefficient for average connection length.
λ	Optical wavelength.
τ	Worst-case signal delay.
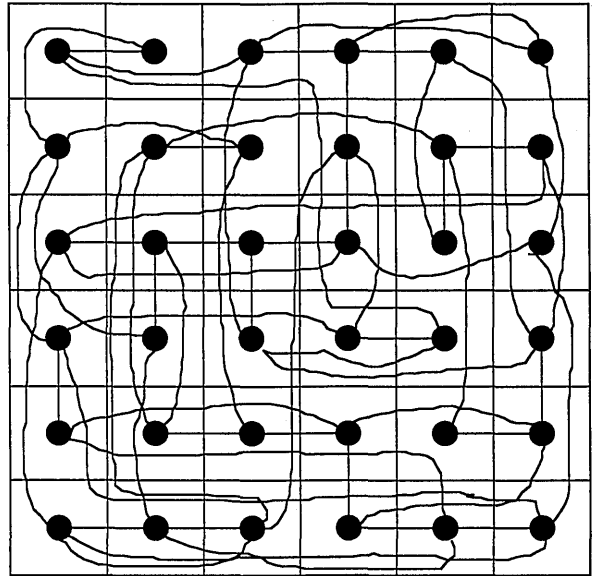
Fig. 2. Layout of the connection graph.

gives the expected number of connections in our system with lengths lying in the interval $[r_0, r_0 + \Delta r]$ and approximately satisfies $\int_{r=1}^{r_{max}} Ng(r)dr = kN$. {Thus $k^{-1}g(r)$ may be interpreted as a probability distribution defined over $[1, r_{max}]$.} $r_{max} \gg 1$ denotes the longest connection length (in grid units), assumed to be of the order of the linear extent of the system. We take $r_{max} = N^{1/e}$ without concerning ourselves with precise geometrical factors. (As we have already discussed, there seems to be little gained by our trying to specify and carry around factors such as $\sqrt{2}$.) The parameter p, known as the Rent exponent, is our measure of connectivity. Systems with large Rent exponents have a large fraction of longer connections.

(2) The longest connection length r_{max} (in grid units).

(3) The number of connections $P(N')$ emanating from a group of N' elements in the partitioning of Fig. 3.

All of the above quantities can be specified by postulation of the distribution of line lengths for our system. (The terms line, connection, and graph edge are used interchangeably.) Obviously we cannot hope to account for all possible connection patterns. Rather we seek a simple analytic distribution function with a few variable parameters, which we hope is representative of the wireability requirements of typical circuits.

Provided the parameter $0 \leq p \leq 1$ (to be discussed shortly) is not too close to unity (say $p \leq 0.9$), we assume the line-length distribution $g(r)$ to be of the form[22-24] (for an expression valid for all values of $0 \leq p \leq 1$, see Ref. 21)

$$g(r) = ke(1 - p)r^{e(p-1)-1}, \qquad 1 \leq r \leq r_{max}, \quad (1)$$

where r denotes distances in terms of grid spacing so that physical distances are given by $\ell = rd$. $Ng(r_0)\Delta r$
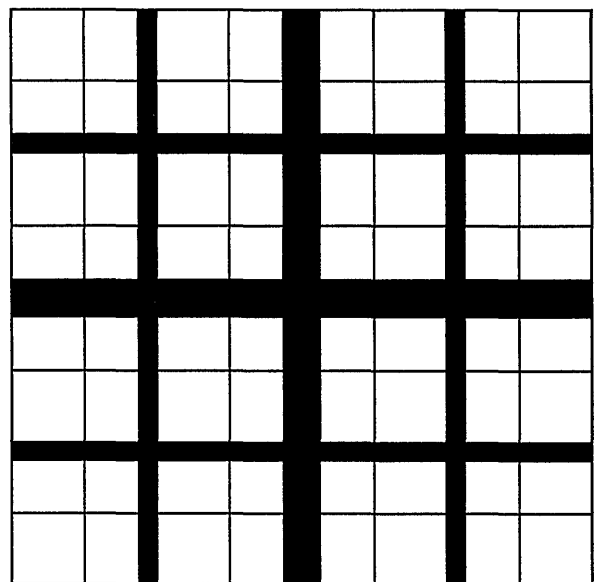

Fig. 3. Binary hierarchical partitioning of the array of cells. A group at the ith level has $N' = N/4^i$ cells.

When p is small, it is more likely for a connection to be made to close-by rather than to distant elements.

Since $e(p-1) - 1 < 0$, we observe that Eq. (1) is simply an inverse power law. Given an empirical line-length distribution, one may attempt to fit Eq. (1) by suitably choosing p. If this is not possible, some other functional form must be employed. We discuss our preference for this type of distribution in Ref. 22.

The average connection length $\bar{r} = k^{-1} \int_{r=1}^{r_{max}} r g(r) dr$ can be calculated as

$$\bar{r} \simeq \kappa(p, e) N^{p-(e-1)/e}, \qquad (2)$$

$$\kappa(p, e) = \frac{e(1-p)}{1 - e(1-p)}, \qquad (3)$$

for $p > (e-1)/e$ (which is the only case to be discussed here; locally connected systems with smaller values of p usually do not suffer from interconnect limitations and thus are not interesting from the viewpoint of this paper). Equation (3) is the definition of the coefficient κ appearing in relation (2).

In order to ensure that there is enough space between the elements for the connections, we need to set the interelement spacing d large enough to permit the passage of at least approximately $k\bar{r}$ connections through each elementary cell.[25,26] For instance, for an $e = 2$ -dimensional system, if the width of each connection is W, then the minimum value of d would be $k\bar{r}W$. The linear extent $\mathscr{L}$ of the system would then be at least $\mathscr{L} = N^{1/2} k\bar{r}W = k\kappa N^p W$, where in the last step we used relation (2). Of course, the linear extent must also be at least $\mathscr{L} = N^{1/2} d_d$. Thus the system linear extent would be given by the greater of the right-hand sides of these equations. If the technology permits M connection layers, this can be written as

$$\mathscr{L} = \max(N^{1/2} d_d, k\kappa N^p W/M). \qquad (4)$$

For larger values of N the second term dominates since we are restricting ourselves to the case $p < 1/2$. The same analysis can be repeated for an $e = 3$ -dimensional system, resulting in

$$\mathscr{L} = \max[N^{1/3} d_d, (k\kappa N^p)^{1/2} W]. \qquad (5)$$

Our line-length distribution is consistent with the following expression for the number of connections $P(N')$ emanating from a group of N' elements[23,24]:

$$P(N') = kN'^p. \qquad (6)$$

Equation (6), historically known as Rent's rule, is valid provided N' is not very close to N (say, $N' \leq N/4$). This equation implies that the number of connections emanating from the boundaries of a group increases sublinearly with the number of elements N' in the group. Larger values of p correspond to a stronger increase with N'.

The reader is referred to Refs. 14, 21, and 22 for discussion, elaboration, and justification of this model, as well as an extensive list of references. The essential assumption we make is that we can speak meaningfully about a distribution of line lengths for our system. Readers who prefer some other form for the function $g(r)$ may easily modify the results of this paper by using appropriate expressions for $\bar{r}$, r_{max}, and $P(N')$ corresponding to their particular choice of $g(r)$. The qualitative results of this paper (for instance, that a particular quantity increases with increasing connectivity) remain applicable regardless of how one chooses to model or measure connectivity.

B. Heat-Removal Model

The interelement spacing of a computing system must also be large enough that we can successfully remove the dissipated power. Packing the elements too densely may result in unacceptable temperature rises and destruction of the system. For systems in which the power-dissipating elements are confined to a planar surface we assume that there is an upper limit to the amount of power we can remove per unit area, denoted by Q (W/m²). Thus power dissipation $\mathscr{P}$ and linear extent $\mathscr{L}$ of a square system must satisfy $Q\mathscr{L}^2 \geq \mathscr{P}$.

For systems in which the power-dissipating elements are distributed throughout a volume, one can again quantify our heat-removal capability by the quantity Q, this time interpreted as the amount of power we can remove per unit cross-sectional area of the volume.[21,27] We do not require this result, however.

C. System Characterization

In this study we assume that the performance of our systems can be characterized in terms of those parameters: the number of elements, N, the bit-repetition rate, B, along each edge of the connection graph, and the inverse signal delay, $S = 1/\tau$. Although it would certainly be desirable, it is not possible to arbitrarily increase S, B, and N simultaneously because of physical limitations. For simplicity the rate B (bits/s) at which information is piped through the connections is assumed to be the same for all connections. Likewise, signal delay τ is taken as the worst case (maximum) over all connections, as is appropriate for synchronous systems.

D. Interconnection Models

We characterize physical interconnection media in terms of the following parameters: (i) Interconnection length ℓ. (ii) Cross-sectional area A or transverse linear extent W (where $A = W^2$). These parameters define packing density for three- and two-dimensional systems, respectively, and thus include any necessary line-to-line separations. (iii) Signal delay τ. (iv) Minimum pulse-repetition interval T_r, i.e., the minimum time interval between consecutive bits on the line. (v) The energy per transmitted bit, E. (Subscripts are used whenever necessary to clarify whether one of the above parameters is associated

with an optical or a normally conducting line, e.g., E_o or E_n.)

T_d denotes the response time of the devices (gates and optical transducers). Of course in general the value of T_d may be different for each of these, but for simplicity we assume that the rate at which the gates may be switched and the optical sources modulated are more or less the same. In almost all cases, minimum bit-repetition interval T_r cannot be less than T_d. This limits the rate B at which we can pipe information through each connection. Thus we assume that B is always specified so as to satisfy $B \leq 1/T_d$. (B can be increased beyond $1/T_d$ by use of several physical channels in parallel to establish each connection or by employment of wavelength-division multiplexing. Both complications, discussed in Ref. 14, are avoided in this paper.)

In the following sections we present models for optical and normally conducting interconnections in their simplest possible form. Detailed derivations and elaborations may be found in Refs. 14 and 21.

1. Optical Interconnections

The cross-sectional area per independent spatial channel is taken to be proportional to the wavelength squared: $A = W^2 = (f\lambda)^2$, where the constant f can be as small as ~ 1 for a diffraction-limited system but may be much larger in practice. In the context of a two-dimensional integrated optics system this simply means that waveguides can be packed at a transverse density of one every $W = f\lambda$. We have argued elsewhere that this is a suitable model for three-dimensional systems as well,[28,29,21] at least from a fundamental perspective. Of course, if one is confronted with a particular optical architecture, as in Refs. 30 and 16, one can determine the space required for communication directly without resorting to parameter W. The analyses presented in this paper can be easily adapted to such architectures.

The signal delay is taken to be the greater of the speed-of-light delay and the device response time: $\tau = \max(\ell/c, T_d)$. Since the effects of dispersion and attenuation can be made small for the length scales in consideration, minimum pulse-repetition interval $T_r \simeq T_d$ and energy per transmitted bit E_o are assumed to be constants. The value of E_o is determined largely by the properties of the light source/modulator and the detector.[9,15]

As a simple example, let us derive the signal delay for a wireability-limited (that is, the interelement spacing is set by wireability limitations) two-dimensional integrated optics system with very small elements (d_d negligible), very fast devices (T_d negligible), and only one connection layer. Using $\ell_{max} = \mathscr{L}$ and Eq. (4), we obtain[14]

$$\tau = \ell_{max}/c = \mathscr{L}/c = k\kappa N^p f\lambda/c, \qquad (7)$$

and $S = 1/\tau$. If the system is heat-removal limited, $Q\mathscr{L}^2$ must exceed the total power dissipation $\mathscr{P} = kNE_oB$ (since there are kN connections dissipating

E_oB each); thus the signal delay is given by

$$\tau = \ell_{max}/c = \mathscr{L}/c = (kNE_oB/Q)^{1/2}/c. \qquad (8)$$

2. Normally Conducting Interconnections

Shorter normally conducting interconnections can be left unterminated, whereas longer ones must be terminated. However, detailed analysis permitting termination of conducting lines shows that it is optimal to start employing optical interconnections at lengths for which it is not yet necessary to terminate lines.[21] Thus it makes little if any difference in our results if we restrict our attention to unterminated lines only. This permits considerable simplification.

The relationship between the rise-time, the length, and the cross-sectional parameters of an unterminated RC line is given by Refs. 14 and 21

$$\tau_{RC} = \alpha \frac{\ell^2}{W^2}, \qquad (9)$$

where α is a constant proportional to the resistivity of the conductor and the permittivity of the dielectric isolating the conductor from the ground plane. It is assumed that the optimal ratios between all transverse dimensions of the line are maintained for different values of W. (A line with width comparable to its height is close to optimal. Increasing the width increases capacitance and reduces packing density. Increasing the height increases fringe effects, again forcing a reduction in packing density, without improving the capacitance considerably.[9,14]) Although we refer the reader to the above references for a derivation, it is easy to convince oneself that this equation makes sense. τ_{RC} is proportional to $RC\ell^2$, where R and C are the resistance and the capacitance of the line per unit length. C is proportional to the ratio of the width of the line to the height of the dielectric; thus it is not affected by scaling of W. On the other hand, R is inversely proportional to both the width of the line and the height of the conductor; thus $R \propto 1/W^2$. Of course, no matter how small the rise time, the signal delay cannot be less than device delay T_d, so $\tau = \max(\tau_{RC}, T_d)$. The energy per transmitted bit is given by

$$E_n = \gamma\ell, \qquad (10)$$

where γ is a constant proportional to the permittivity of the dielectric and the square of the nominal voltage, V. The energy is proportional to $C\ell V^2$, and since C is not affected by scaling W, the above equation is justified.

For simplicity we assume that all lines in a given system are of the same width W. (This assumption is relaxed in Ref. 14.)

Again, as an example, let us derive the signal delay for a two-dimensional system with very small elements and very fast devices. Using $\ell_{max} = \mathscr{L}$ and Eq. (4), we find[14]

$$\tau = \alpha\ell_{max}^2/W^2 = \alpha(k\kappa N^p/M)^2. \qquad (11)$$

For an unterminated RC line the signal delay is also the minimum bit-repetition interval, which determines the maximum bit-repetition rate. Thus $B \leq S = 1/\tau$.

The above result is scale invariant[14]; i.e., it does not depend on W. If technology enables us to manufacture very fine lines, heat-removal requirements will determine how much we can scale down the system and thus its minimum linear extent, but this will have no effect on the delay (provided the scale of the system does not have to be so large that the lines become propagation limited[14]). To find the heat-removal-limited linear extent of the system, first note that the dissipation associated with each cell is $\gamma k \bar{\ell} B = \gamma k \bar{r} d B$ [obtained by Eq. (10) and because there are k connections per cell with average length $\bar{\ell} = \bar{r} d$, switched at a rate B]. Using relation (2) and requiring that the power dissipation per cell not exceed $Q d^2$, we can show that the system linear extent $\mathscr{L} = N^{1/2} d = \gamma k \kappa N^p B / Q$.

3. Repeatered Normally Conducting Interconnections

The inhibitive square-law behavior of normally conducting lines may be alleviated with the use of repeater structures. Bakoglu and Meindl derived the optimal configurations of such structures.[31] The delay along such a line is given by $\tau_{rep} \propto (R_0 C_0 R C \ell^2)^{1/2}$, where $R_0 C_0$ denotes the intrinsic delay of the repeaters. Following scaling arguments similar to those with ordinary normal conductors, we may write this in the form

$$\tau_{rep} = \beta \frac{\ell}{W}, \qquad (12)$$

where β is proportional to $(R_0 C_0)^{1/2}$ as well as the square roots of the conductor resistivity and the dielectric permittivity. The effect of repeaters being used on the energy per transmitted bit can be ignored so that Eq. (10) is still valid.

Following the derivation of earlier results, one can show that the signal delay for a two-dimensional repeatered system with very small elements and very fast devices is given by[14]

$$\tau = \beta \ell_{max}/W = \beta (k \kappa N^p / M). \qquad (13)$$

The effect of heat-removal requirements is similar to that in the unrepeatered case.

4. Superconducting Interconnections

Superconducting interconnections are not treated in this paper. Let it suffice to say that superconductors lead to similar results as optical interconnections, both analytically and numerically.[14] However, there is one exception: The energy per transmitted bit can be made much smaller than with optical interconnections, especially at lower temperatures; thus they may be advantageous in heat-removal-limited systems. On the other hand they suffer from termination

problems, and the prospects for three-dimensional circuits are no better than for normal conductors.

3. Outline of the Analysis

Now we actually outline the steps of our analysis (Fig. 4). The linear extent d_1 of an electrically connected group of N_1 elements must be large enough to do the following:

(1) Accommodate N_1 elements.
(2) Accommodate the electrical wires connecting them.
(3) Accommodate $k N_1^p$ optical transducers [Eq. (6)].
(4) Satisfy heat-removal requirements.

Then d_m, the intergroup spacing of the electrically connected groups (also referred to as modules) must be large enough to do the following:

(1) Accommodate d_1 (i.e., $d_m \geq d_1$).
(2) Accommodate the optical channels connecting the groups.
(3) Satisfy any additional heat-removal requirements.

Note that the N elements (switches and gates) are no longer uniformly laid out as in Fig. 2 but are clustered into modules. This enables considerable energy savings since the electrical wires can be made much shorter.

On the basis of these considerations we can write expressions for the signal delay (which is taken as the worst case over all connections), the system size, and the power dissipation as functions of N, B, and N_1. Then we can pick the value of N_1 maximizing our figure-of-merit function.

We carried out this analysis for a variety of layout constraints, combinations of media, and physical parameters. It is not possible (and perhaps not useful) for us to reproduce all of our results. Rather, we try to present representative examples chosen for their illustrative qualities and qualitatively discuss
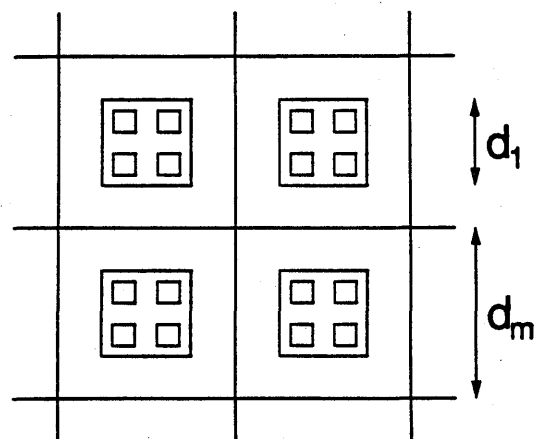

Fig. 4. Analysis of optimal hybrid layouts ($N_1 = 4$).

several general conclusions deduced from the study of a large number of cases.

Although not attempted in this paper, systems with three or more hierarchical levels can also be analyzed.[21] This would permit systems involving normal conductors, superconductors, and optics all together.[21] It would also enable the different parameters of electrical interconnections at different levels (such as on-chip and multichip substrates) to be taken into account.

4. Regarding Numerical Examples

In our numerical examples we try to look into the future and to select reasonably optimistic parameters for each interconnection media. We also consider the effects of degradation of the optical parameters from what seems to be their best possible values, as it is not yet a mature technology.

We assume 10-GHz devices, i.e., $T_d = 0.1$ ns. We permit a maximum of $M = 10$ conducting connection layers but only one optical connection layer for two-dimensional layouts. We assume a nominal voltage level of $V = 1$ V and room-temperature aluminum conductivity. The parameters α, β, and γ are taken as 1.5×10^{-17} s, 3.9×10^{-14} s, and 6.9×10^{-11} J/m, respectively.[21] These values roughly correspond to the best possible at room temperature, with 10-GHz repeaters. Since our analysis is approximate, these numbers can also be rounded to the nearest order of magnitude; however, we use them as they are so that consistency with Ref. 21 is maintained. We permit a maximum power dissipation per unit area of $Q = 10$ W/cm^2. The minimum manufacturable value of W for conducting interconnections is taken as $W_{\min} = 0.2$ μm. Element size d_d is assumed to be ten times this value. (This implies that our elements are of the simplest type, such as logic gates.) We assume there to be $k = 5$ connections per element. We take $\lambda = 1$ μm as the optical wavelength. The optical transducers are assumed to be $d_{tr} = 5$ μm in diameter. We assume the best possible optical communication energy to be $E_o = 1$ pJ, but we also consider degradation of this value by a factor of 100. Likewise, we consider near-diffraction-limited operation ($f = 2$), but we also consider degradation of this by a factor of as much as 100 ($f = 200$). We consider two different values of the Rent exponent, $p = 0.6$ and $p = 0.8$, to observe the effects of connectivity on the results. The values chosen above for $W_{\min}$, d_d, and d_{tr}, which seem realizable in the near future, are already small enough that in most cases, totally ignoring the effects of these parameters will have little or no effect on our results.

Our primary objective is minimization of signal delay. As a secondary objective we try to minimize total power dissipation. In other words we try to minimize signal delay as much as possible, and only then, to the extent possible without making any sacrifice from the minimum signal delay, do we try to minimize power dissipation. This approach emphasizes performance, and cost is minimized only if this

can be achieved without performance being sacrificed. Analytically this can be realized by maximization of

$$\Gamma = \frac{S}{\mathscr{P}^\epsilon}, \qquad \epsilon \to 0. \qquad (14)$$

Other figures of merit emphasizing cost are also discussed.

Computationally, we chose $\epsilon = 10^{-10}$. We observed that choosing values as large as 10^{-5} or as small as 10^{-12} makes no difference. Values larger than 10^{-5} start changing the results in favor of systems exhibiting somewhat larger signal delay but less power dissipation. Values lower than 10^{-12} start causing numerical problems.

5. Two-Dimensional Systems

In this section we consider fully two-dimensional systems, i.e., systems that are confined to a two-dimensional surface, including optical paths. One implementation of such a system may involve a two-dimensional array of VLSI chips with optical transducers located on a topmost layer or on dedicated islands. The optical imaging system may be a glass waveguide overlay or the folded multifacet architecture described in Ref. 32.

A. Analysis

We refer to Fig. 4. We assume that there are k connections per element. Some of these connections are made to elements in the same module and are implemented with normal conductors. Other connections are made to elements in other modules. One can establish such connections optically by tying optical transducers at the to-be-connected terminals and by guiding the light emanated from the source terminal to the target terminal with some type of optical imaging system. It matters little whether the transducers are on a separate layer or side by side with the electrical wiring [since $\max(x, y) \simeq x + y$ within a factor of 2]. The intergroup spacing d_m may have to be larger than d_1 because of the space necessary to accommodate the optical channels. Again, it makes little difference whether we assume that the optical channels are on a separate layer or that they compete for the same space with the modules. As discussed earlier, there is little purpose in specifying such details, as they ultimately change the results by factors such as 2, $\sqrt{2}$, etc. If one is confronted with a particular system for which such details are specified, such factors can be readily introduced in order to obtain more accurate results.

Remember that our purpose is to determine, for every N and $B \le 1/T_d$, the value of N_1 maximizing our figure-of-merit function. We can immediately set an upper bound on N_1 since we know that the maximum bit-repetition rate, B, is a decreasing function of the number of electrically connected elements [Eq. (11) and the following remarks]. We can solve for N_1 from Eq. (11) (rewritten for a module with N_1, rather than N elements, and since $B \le 1/\tau$ for

unterminated conducting lines) as

$$N_1 \leq (M/k\kappa)^{1/p}(1/\alpha B)^{1/2p}. \qquad (15)$$

The right-hand side of this inequality is the largest value of N_1 compatible with given B. We denote this value of N_1 as $N_{1_{max}}$. For given B, N_1 must be chosen to lie between 1 and $N_{1_{max}}$.

Ignoring heat-removal requirements for the moment, we see that the linear extent of each module d_1 must satisfy

$$d_1 \geq \max[N_1^{1/2}d_d, k\kappa N_1^p W_{min}/M, (kN_1^p)^{1/2}d_{tr}, kN_1^p(f\lambda)]. \qquad (16)$$

The first term is trivial. d_1 must be at least large enough that the module can accommodate $N^{1/2} \times N^{1/2}$ elements of linear extent d_d each. The second term is simply the smallest module linear extent that still provides enough room for routing the electrical connections at the minimum manufacturable linewidth W_{min}. These two terms correspond to Eq. (4). The third term accounts for the fact that the module must be at least large enough to accommodate kN_1^p transducers of linear extent d_{tr}, since this many connections are made to elements in other modules [according to Rent's rule, Eq. (6)]. The final term accounts for the fact that the linear extent of the module must be large enough to permit the passage of kN_1^p optical channels through it; otherwise, the channels emanating from the transducers would not be able to get out of the boundaries of the module. (Such a problem does not arise if out-of-plane optical communication is permitted.) The values of d_d and W_{min}/M are often small enough that the first two terms can be ignored. Technological improvements will further decrease d_d and W_{min} and increase M so that it is unlikely for these terms to determine d_1. Also, since d_{tr} need not be much larger than $f\lambda$ (i.e., the transducers need not be much larger than the transverse extent allocated for each channel), the third term will often be shadowed by the fourth. Thus in most cases we will be left with the single term $d_1 \geq kN_1^p(f\lambda)$.

Heat-removal considerations will require that

$$Qd_1^2 \geq \text{(electrical dissipation plus optical dissipation)}. \qquad (17)$$

First, let us calculate the electrical dissipation. We can find the total connection length by multiplying the total number of connections kN_1 by the average connection length $\bar{\ell} = \bar{r}d$ (where, as before, d is the interelement spacing of the elements). Then, with Eq. (10), the total energy involved in one cycle is $\gamma k N_1 \bar{\ell}$ and the power dissipation is $\gamma k N_1 \bar{\ell} B$. Using relation (2) and also $d_1 = N_1^{1/2}d$, we find the total electrical power dissipation to be $\gamma k\kappa N_1^p B d_1$. The optical power dissipation per module is simply the number of optical connections per module times the

power dissipation $E_o B$ per connection, giving $kN_1^p E_o B$. Using these results, simplifying inequality (17) with $x + y \simeq \max(x, y)$, and solving for d_1, we obtain

$$d_1 \geq \max\left[\left(\frac{kN_1^p E_o B}{Q}\right)^{1/2}, \frac{\gamma k\kappa N_1^p B}{Q}\right]. \qquad (18)$$

Combining this with $d_1 \geq kN_1^p(f\lambda)$, we obtain the minimum value of d_1 as

$$d_1 = \max\left[kN_1^p(f\lambda), \left(\frac{kN_1^p E_o B}{Q}\right)^{1/2}, \frac{\gamma k\kappa N_1^p B}{Q}\right]. \qquad (19)$$

The minimum value of the intermodule separation, d_m, is then given by

$$d_m = \max[d_1, k\kappa N_1^p(N/N_1)^{p-1/2}(f\lambda)], \qquad (20)$$

where we assume only one optical connection layer. Apart from being large enough to accommodate the module, the intermodule spacing must be large enough to permit the passage of at least $kN_1^p\kappa(N/N_1)^{p-1/2}$ optical channels. (Remember that the spacing between elementary cells had to be large enough to permit the passage of at least $k\bar{r} = k\kappa N_1^{p-1/2}$ connections. Now replace $k \rightarrow kN_1^p$ and $N_1 \rightarrow N/N_1$, since now instead of N_1 elements with k connections each, we have N/N_1 modules with kN_1^p connections each.)

The system signal delay is finally given as the maximum over all connections:

$$\tau = \max[(N/N_1)^{1/2}d_m/c, \alpha(k\kappa/M)^2 N_1^{2p}, T_d). \qquad (21)$$

The first term is the speed-of-light delay along the longest optical connection. The second term is the delay along the longest normally conducting interconnection, and the last term accounts for device delay. Employing previous equations, we obtain

$$\tau = \max[k\kappa N^p(f\lambda/c), (N/N_1)^{1/2}(kN_1^p E_o B/c^2 Q)^{1/2},$$
$$(N/N_1)^{1/2}(\gamma k\kappa/N_1^p B/cQ), \alpha(k\kappa/M)^2 N_1^{2p}, T_d] \qquad (22)$$

for the resulting signal delay for a hybrid system. The signal delay for an all-optical system under the same approximations is given by the maximum of the right-hand sides of Eqs. (7) and (8) and T_d. For an all-electrical system it is given by the maximum of the right-hand side of Eq. (11) [Eq. (13) with repeaters and T_d.] The first term, which is independent of N_1, is also the delay of an all optically connected system of N elements [Eq. (7)]. Thus we conclude that the use of a hybrid layout cannot reduce the system size and delay below that of a wireability-limited all-optical system.

Since $N_1 \leq N_{1_{max}}$, the first term in the above equation eventually dominates the others with increasing N (remember that our expressions are valid

for $p > 1/2$). When this is the case, the choice of N_1 has no effect on the delay. Thus we choose that value resulting in minimum power dissipation. The total power dissipation is expressed as

$$\mathscr{P} = (N/N_1)(kN_1^p E_o + \gamma k\kappa N_1^p d_1)B$$

$$\simeq (N/N_1)\max(kN_1^p E_o, \gamma k\kappa N_1^p d_1)B. \quad (23)$$

When d_1 is given by the first term of Eq. (19), the value of N_1 minimizing total power dissipation is given by

$$N_1^p = \frac{E_o}{\gamma(f\lambda)k\kappa}. \quad (24)$$

For larger values of B either of the two latter terms in Eq. (19) may dominate. Interestingly, whichever of these two terms dominates, it is possible to show that the value of N_1 minimizing relation (23) is given by the same expression:

$$N_1^p = \frac{kE_oQ}{(\gamma k\kappa)^2 B}. \quad (25)$$

[Incidentally, we also note that this value of N_1 minimizes the combined second and third terms of Eq. (22).] We see from this equation that the optimal value of N_1 increases with the optical communication energy and our heat-removal ability and decreases with the bit-repetition rate and the Rent exponent. It obviously increases with increasing E_o, as optical communication becomes energetically more expensive. It also increases with increasing Q: if we are able to remove larger amounts of power per unit area, this means that the scale of the electrically connected groups can be reduced, reducing the energy cost of electrical interconnections. On the other hand, increasing B reduces N_1, since it results in an increase in power dissipation and d_1, making electrical interconnections more expensive. N_1 also decreases with increasing p. Systems with larger p have a larger fraction of longer connections, so it is beneficial to make fewer electrically.

Of course, since we are giving full precedence to minimizing signal delay, keeping the latter terms of Eq. (22) at less than its first term has priority over minimizing power. For instance, so that the fourth term of Eq. (22) does not exceed the first, we must maintain

$$N_1 \le \left(\frac{f\lambda/c}{\alpha}\frac{M^2}{k\kappa}\right)^{1/2p} N^{1/2}. \quad (26)$$

Thus the optimal value of N_1 is determined by three considerations. For larger N it is that value which minimizes total power consumption, as given by either of Eqs. (24) or (25), with the restriction that it can never exceed $N_{1_{\max}}$. These two considerations determine the optimal value of N_1 for large N. The optimal value of N_1 should also not exceed the

envelope given by inequality (26). This last restriction relaxes with increasing N. Notice that the effect of variation of p on all three of these considerations is the same: larger Rent exponents favor the use of more optics.

Naturally, it is possible for there to be combinations of parameters and variables for which other terms that we neglect dominate. But in most cases the above expressions agree with the calculations of Ref. 21, in which all terms were maintained. Thus they suffice for our purposes.

B. Numerical Examples

Throughout this paper we keep all but three of the physical parameters the same for all examples, as discussed in Section 4. Thus in the following we only specify the values of p, f, and E_o for each example.

In our first example we consider a system with Rent exponent $p = 0.6$, optical communication density only $f = 2$ times worse than diffraction limited, and an optical communication energy $E_o = 1$ pJ. Figure 5 shows the optimal value of N_1 as a function of N with B as a parameter. In all numerical plots we vary N from 10^4 to 10^{10}. One should keep in mind, however, that the larger values of N in this range may lead to unrealistic system sizes and/or power dissipations for some combinations of the parameters. For the two lower values of B it is optimal to make all connections electrical (i.e., $N_1 = N$) until $N \simeq 2 \times 10^5$, after which the optimal value of N_1 is independent of N and is that value that minimizes total power consumption. For these relatively low values of B the size of each electrically connected group d_1 is given by $d_1 = kN_1^p(f\lambda)$. Thus the total power dissipation is given by relation (23) with $d_1 = kN_1^p(f\lambda)$. The value of N_1 that minimizes power dissipation is given by Eq. (24) and is indeed consistent with that observed in Fig. 5 for the two lowest values of B. For the two larger values of B the latter terms of Eq. (19) dominate so that the optimal value of N_1 is given by Eq. (25), which indeed predicts the optimum values of N_1 for the two larger values of B in Fig. 5. [The optimal values of N_1 for smaller values of N are determined by the competition of the various terms in Eq. (22), until the first term dominates the others, after which N_1 is given by Eq. (25).]

Figure 6 illustrates the resulting dependence of S on N for $B = 100$ Mbits/s. The solid curve corresponds to the optimal choice of N_1. The dashed curve corresponds to all connections being made optical and coincides with the solid curve for larger values of N; it is given by $1/S = \tau = k\kappa N^p f\lambda/c$. The dotted curve, which overlaps with the solid curve initially, corresponds to all connections being made electrical and is given by $\alpha(k\kappa/M)^2 N^{2p}$. We cannot make all connections electrical once $N > N_{1_{\max}} = 7 \times 10^6$; thus the dotted curve terminates at this value of N. After a certain value of N, making all connections optical is as good as the optimal hybrid combina-
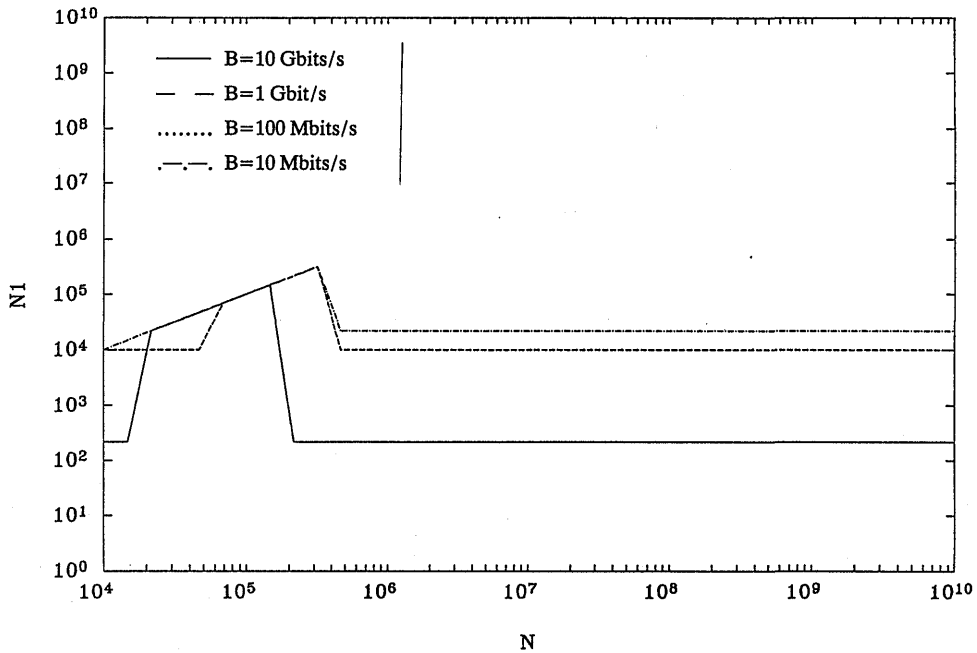
Fig. 5. N_1 versus N: $p = 0.6, f = 2, E_o = 1$ pJ. The plots for the two lower values of B coincide.

tion in terms of minimization of system size and delay, as discussed previously [following Eq. (22)].

However, making all connections optical results in power dissipation ~1 order of magnitude larger than the optimal hybrid combination for the largest value of B and ~2 orders of magnitude larger for the smallest value of B. The disparity is greater for smaller values of B because the optimal value of N_1 is larger when B is smaller. In other words the all-optical system $(N_1 = 1)$ is farther away from the optimum. These considerations have no effect on the resulting system size and signal delay, since two-dimensional systems tend to be wireability rather than heat-removal limited. Another consequence of this is that the resulting values of S for different values of B are identical.

The entrance of optical interconnections after $N \sim 2 \times 10^5$ elements is accompanied by a drastic increase in system size, as illustrated in Fig. 7. The linear extent of the all-electrical system is given by
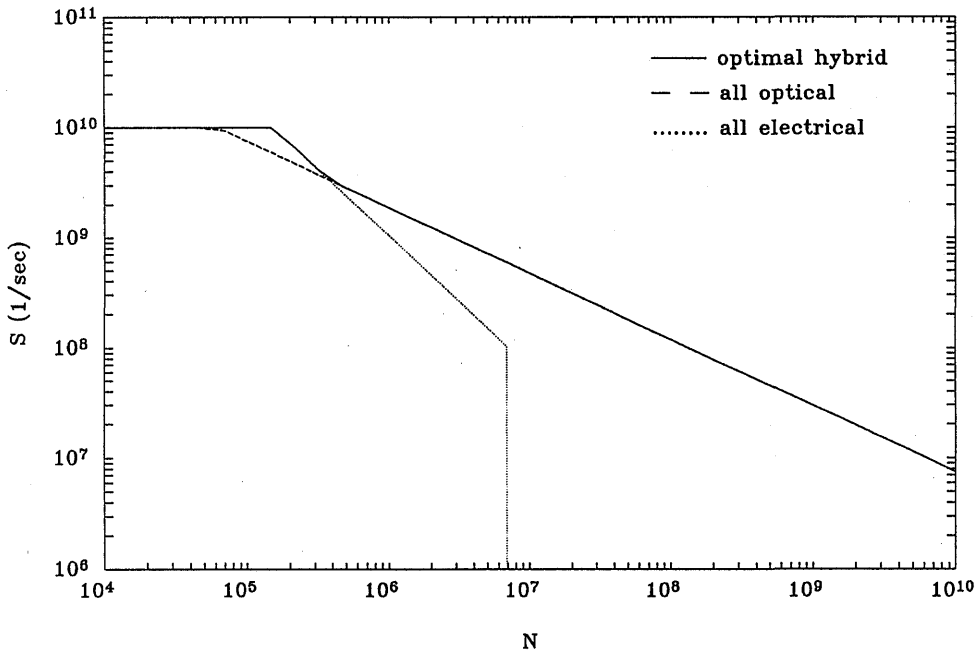

Fig. 6. S versus N: $B = 100$ Mbits/s, $p = 0.6, f = 2, E_o = 1$ pJ. The plot for the optimal hybrid case first coincides with that for the all-electrical case and then with that for the all-optical case.
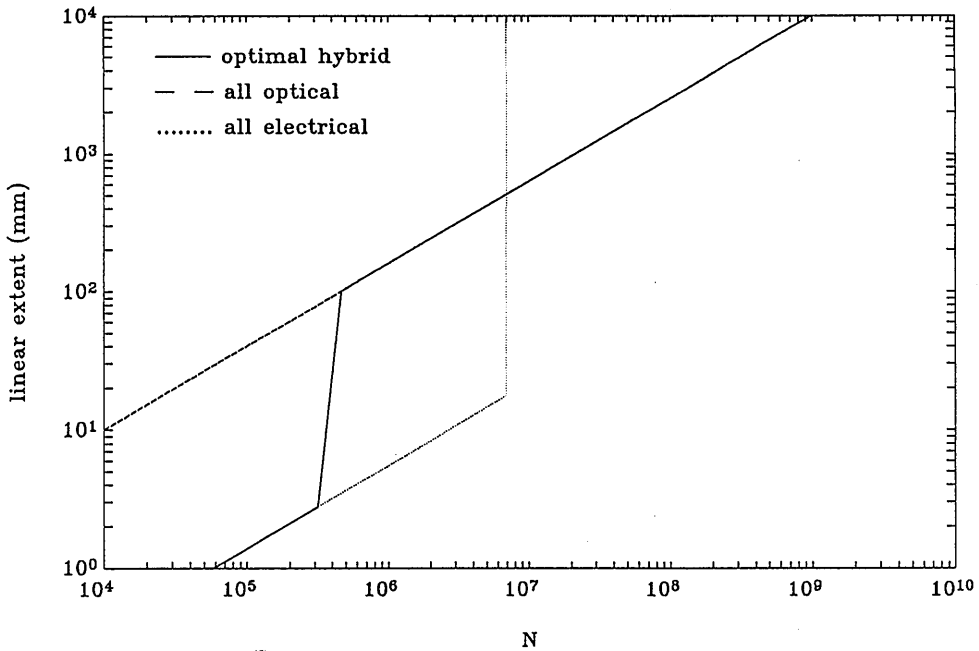
Fig. 7. $\mathscr{L}$ versus N: $B = 100$ Mbits/s, $p = 0.6$, $f = 2$, $E_o = 1$ pJ. The plots for the all-optical and the optimal hybrid cases coincide for larger N.

$\gamma k \kappa N^p B/Q$. Once we start using optical communication for the longer connections, they dominate the system area, leading to a linear extent given by $k \kappa N^p f \lambda$. The curve for S is continuous because the fast velocity of propagation of optical interconnections compensates for the increase in system size. We can avoid the jump in system size by keeping all connections electrical; however, in this case the value of S is less than that possible with a hybrid system. This is one example of a situation in which the use of

optics permits performance not possible with normal conductors alone but at a significant penalty in terms of system size.

Let us now explore the effects of degradation of the optical parameters. The optimal value of N_1 when $f = 50$ is plotted in Fig. 8. We observe that it is beneficial to stick to an all-electrical system $(N_1 = N)$ until $N > N_{1_{max}}$, after which the entrance of optics is unavoidable. (If f is too large, the increase in $\mathscr{L}$ accompanying the entrance of optics may be unreason-
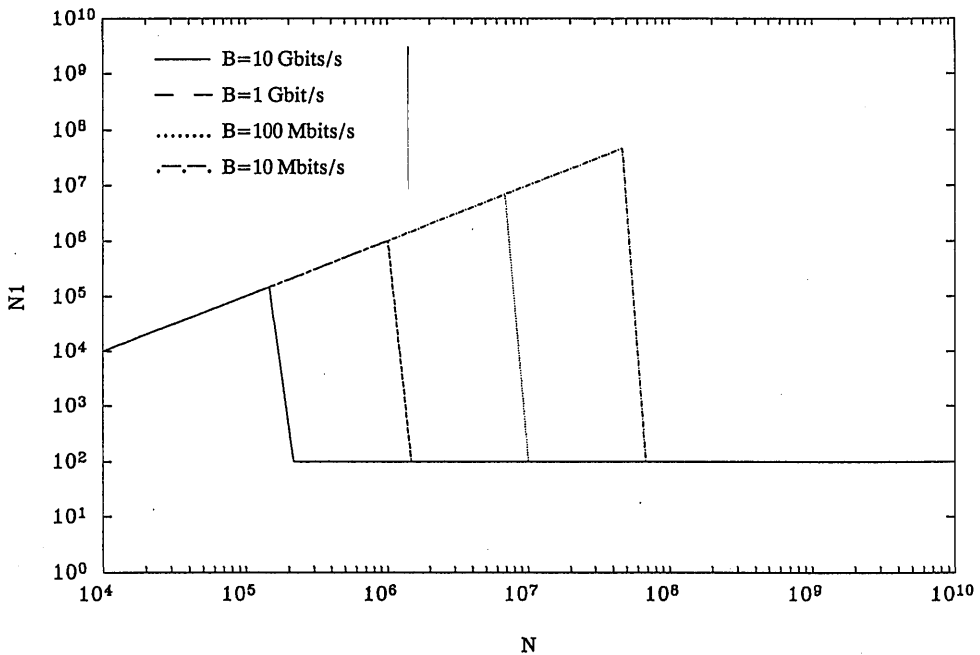

Fig. 8. N_1 versus N: $p = 0.6$, $f = 50$, $E_o = 1$ pJ. All plots coincide for the smallest and the largest values of N.

ably large; thus increasing N beyond $N_{1_{max}}$ while we maintain the given value of B may not be feasible.) Once this occurs, module size d_1 is limited by the term $kN_1^p(f\lambda)$ for all values of B. The resulting large value of d_1 makes electrical connections expensive, leading to a small optimal value of N_1, as given by Eq. (24). Notice that although increasing f leads to an overall degradation in performance and cost, it results in greater use of optics when $N > N_{1_{max}}$.

Figure 9 illustrates the resulting dependence of S on N. Notice that in this case, unlike in the previous example, a sudden drop in S is observed. This is because we are forced to use optical interconnections prematurely so as to maintain the given value of B, before the value of S for an all-electrical system falls below that for an all-optical system (as was the case in Fig. 6).

In other words, when $f > 10$ or so, we have a region in S–B–N space in which a small increase in N or B is accompanied by a large increase in system size and a large decrease in S. This behavior may have algorithmic implications.[14] Among several algorithms designed to solve a given problem it may be preferable to employ those requiring relatively smaller values of N and/or B, if possible, since even small increases in these parameters require a large sacrifice in terms of S. For instance, if we are trying to maximize a figure-of-merit function of the form $S^x B^y$, where $x, y > 0$ and x is not much smaller than y, it is likely that we will settle for an operating point not involving any optical interconnections.

In conclusion, if the use of optics is to be worthwhile for two-dimensional systems, it is of paramount importance to bring f as close as possible to unity.

The folded multifacet architecture described in Ref. 32 was devised to meet this requirement.

Now we discuss the effects of an increase in the optical communication energy hundredfold (Fig. 10). For large N the optimal value of N_1 is seen to hit $N_{1_{max}}$. For smaller N we observe the envelope $N_1 \propto N^{1/2}$ given by inequality (26). Despite the fact that increasing the optical energy results in a drastic shift in N_1, it has no effect on the resulting value of S, which is still as given in Fig. 6. This is because the system is still wireability limited. (An exception occurs for the largest values of B, for which the system may be heat-removal limited up to a certain value of N.)

Of course, now the total power dissipation is much larger, and the discrepancy between the optimal system and the all-optical system in this respect is even greater than before. The entrance of optics results in a large increase in total power dissipation (Fig. 11). The all-optical and the all-electrical system power dissipations are given by kNE_oB and $(\gamma k\kappa N^p B)^2/Q$, respectively. The total power dissipation for the optimal hybrid case for larger N is given by $NkN_1^{p-1}E_oB$ with $N_1 \sim 5 \times 10^6$.

Let us now consider that p is increased to 0.8 (Fig. 12). The downward shift in the optimal values of N_1 is easily explained by the changes in the values of $N_{1_{max}}$. Once again the envelope given by inequality (26) is observed.

6. Cost-Based Optimization

Until now we have concentrated on the optimization function given by Eq. (14), which gave full precedence to minimizing signal delay and only secondarily tried
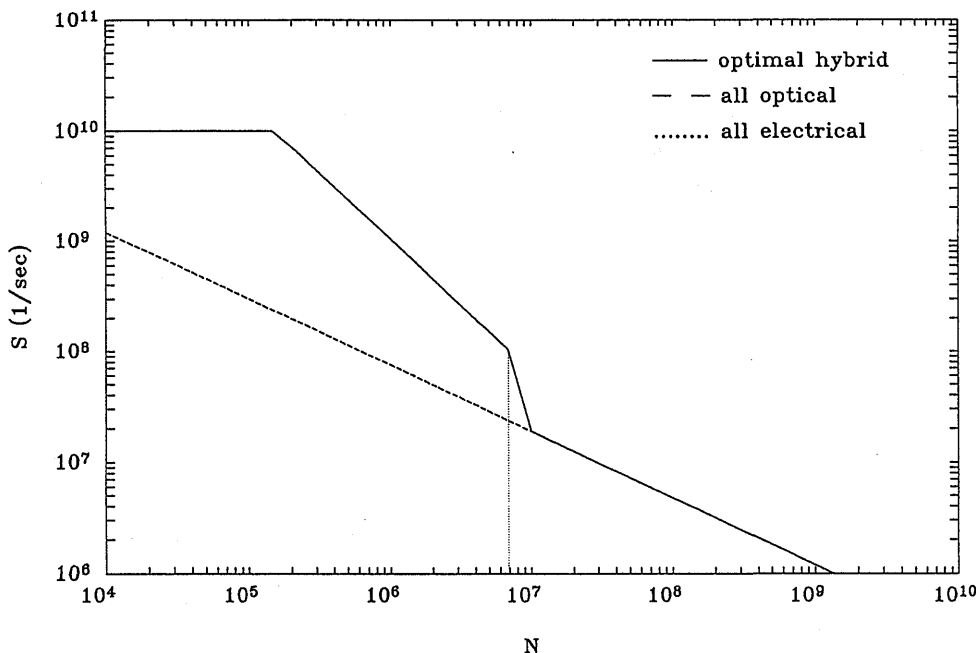

Fig. 9. S versus N: $B = 100$ Mbits/s, $p = 0.6$, $f = 50$, $E_o = 1$ pJ. The plot for the optimal hybrid case first coincides with that for the all-electrical case and then with that for the all-optical case.
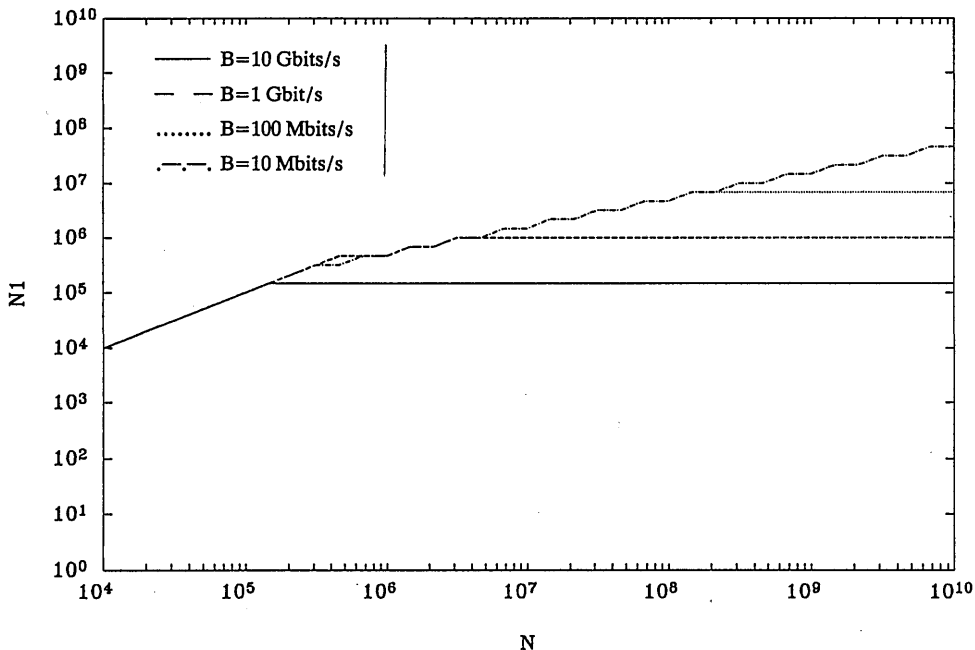
Fig. 10. N_1 versus N: $p = 0.6, f = 2, E_o = 100$ pJ.

to minimize power dissipation. The cost of system size was not accounted for at all. Now we consider another example optimization function, which accounts for the cost of system size.

A. Analysis

We consider dividing Eq. (14) by the system area $\mathcal{L}^2$:

$$\Gamma = \frac{S}{\mathcal{L}^2 \mathcal{P}^\epsilon}. \qquad (27)$$

Of course, it is always possible to employ more complicated functions if desired, depending on the relative importance we attach to speed, system size, and power cost.

B. Numerical Examples

Figure 13 shows how the optimal values of N_1 are changed for the layout considered earlier (Fig. 5). As discussed before, the entrance of optics is accompanied by a drastic increase in system size. Thus with
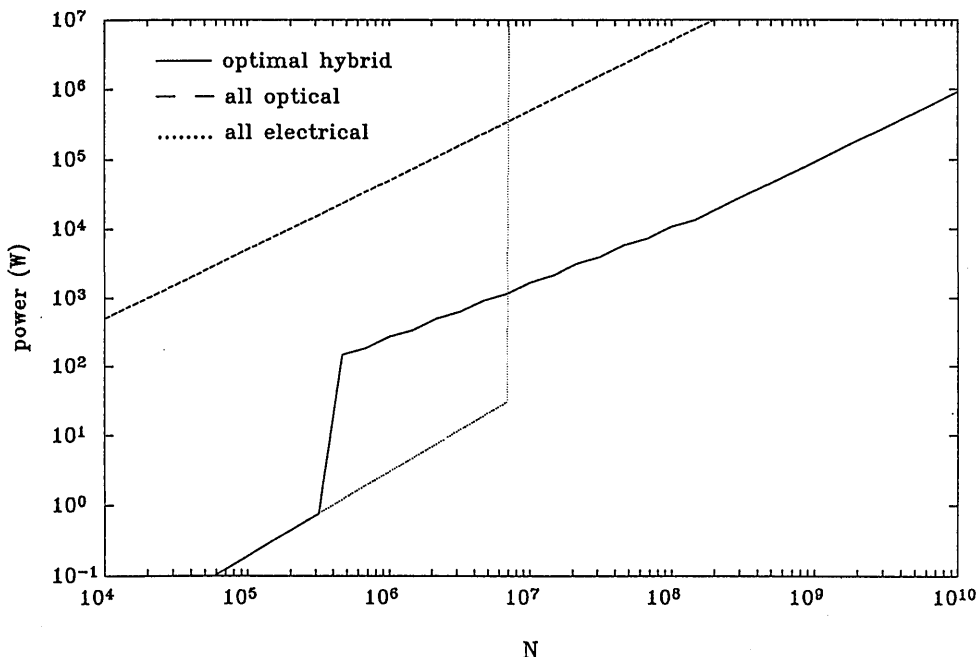

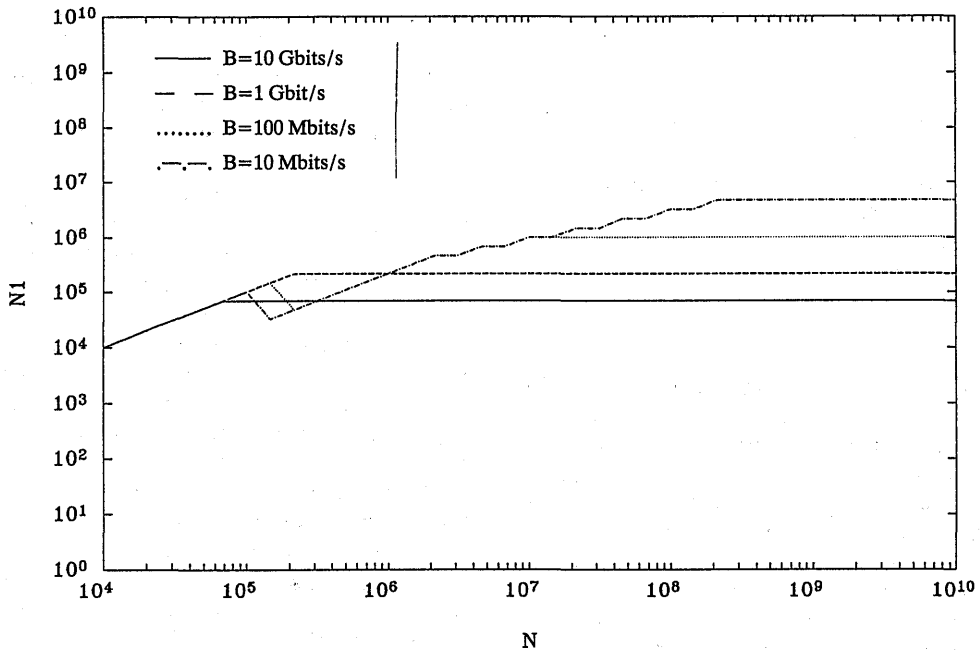Fig. 11. $\mathcal{P}$ versus N: $B = 100$ Mbits/s, $p = 0.6, f = 2, E_o = 100$ pJ.

Fig. 12. N_1 versus N: $p = 0.8, f = 2, E_o = 100$ pJ.

our new figure of merit it is beneficial to stick to an all-electrical system until larger values of N despite the fact that the resulting signal delay will be worse than that possible with a hybrid system. However, once N tends to exceed $N_{1_{max}}$, the entrance of optics is unavoidable. Once optical interconnections are entered, the steep increase in system size is observed. For $N > N_{1_{max}}$ the optimal values of N_1 are identical to those in Fig. 5.

One other possible figure-of-merit function that we do not deal with here is given by $\Gamma = S/\mathscr{P}$, which is

discussed in Ref. 21. The general conclusion is that accounting for the cost of system size and power dissipation results in all-electrical systems being preferred until very large numbers of elements.

7. Effect of Using Repeaters

A. Analysis
There is no maximum value of N_1 for given B when repeaters are used; thus there is no analog of inequality (15). The linear extent of each module must
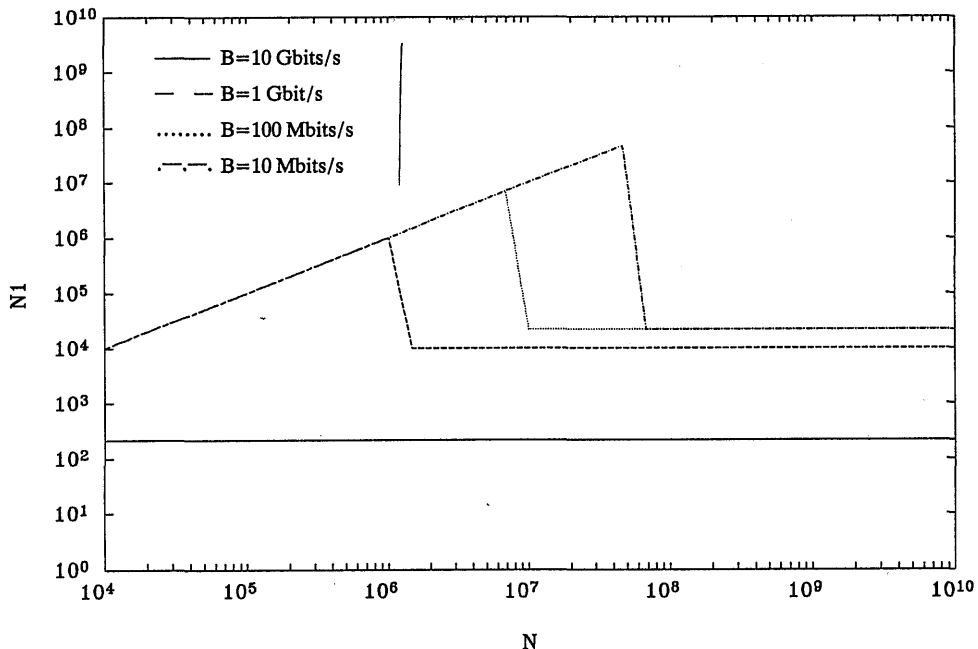

Fig. 13. N_1 versus N (cost-based optimization): $p = 0.6, f = 2, E_o = 1$ pJ.

again satisfy Eq. (19), and the intermodule separation d_m is calculated similarly. Finally the system signal delay is given by

$$\tau = \max[(N/N_1)^{1/2}d_m/c, \ \beta(k\kappa/M)N_1^p, \ T_d]. \quad (28)$$

The total power dissipation is again given by relation (23), and Eqs. (24) and (25) are still applicable. Again, on the basis of arguments similar to those leading to inequality (26), we must ensure that $\beta(k\kappa/M)N_1^p \leq k\kappa N^p(f\lambda/c)$.

B. Numerical Examples

Comparing Eqs. (7) and (13), we observe that the relations between $S = 1/\tau$ and N for repeatered and optical layouts are identical in form and differ only by the numerical factors $(f\lambda/c)$ versus (β/M). Optical interconnections have larger linewidths and faster propagation velocity. Repeatered interconnections can (and must) be scaled down to smaller linewidths but have a slower propagation velocity. Let us first consider a system with optical communication density $f = 50$ times worse than diffraction limited and an optical communication energy of $E_o = 1$ pJ. The same results are obtained whether $p = 0.6$ or $p = 0.8$. (For these parameters an all-repeatered layout results in signal delay more than 40 times less than that possible with an all-optical layout). We find that an all-electrical system $(N_1 = N)$ is best for the range of N and B in consideration (Fig. 14). What essentially happens is that with 10-GHz devices, repeaters permit fast propagation. Much smaller linewidths and ten connection layers result in much smaller system size, more than making up for the deficiency compared with the speed of light.

It is also interesting to note that no optical communication is used even when the system size exceeds a centimeter, which is the breakeven length for energy for this choice of parameters. This illustrates the inadequacy of breakeven-length approaches. A submicrometer repeatered line cannot be replaced with an $f \sim 50$ optical line even though the latter may have a smaller communication energy. (Of course, for large N the total power dissipation for the all-repeatered system will be greater than that possible with a hybrid system. Remember, however, that our original figure-of-merit function gives full priority to minimizing signal delay.)

Now, let us assume the value of f is reduced to 2, which we are assuming to be its best possible value. In this case the discrepancy between the use of all repeaters and all optics in terms of signal delay is reduced to less than a factor of 2.

(An exception may occur for the largest values of B for which the system tends to be heat-removal limited. Then, the heat-removal-limited all-repeatered system may result in signal delay greater than that given by the wireability-limited Eq. (13). However, the demonstration of this would require consideration of effects not taken into account in this paper and would greatly complicate our analysis; thus we satisfy ourselves by referring the reader to Ref. 21.)

Apart from such cases, when $f \sim 2$, the signal delay of an all-optical system is close to that of an all-repeatered system [i.e., Eqs. (7) and (13) are numerically close. This does not mean, however, that we can use optics or repeaters interchangeably for the individual connections in the same system because of scale incompatibility.] The all-repeatered system has much smaller area. When $p = 0.6$, the power dissipation of the all-repeatered system lies quite below that of the all-optical system for the range of N in consideration (although it grows faster with increasing N).
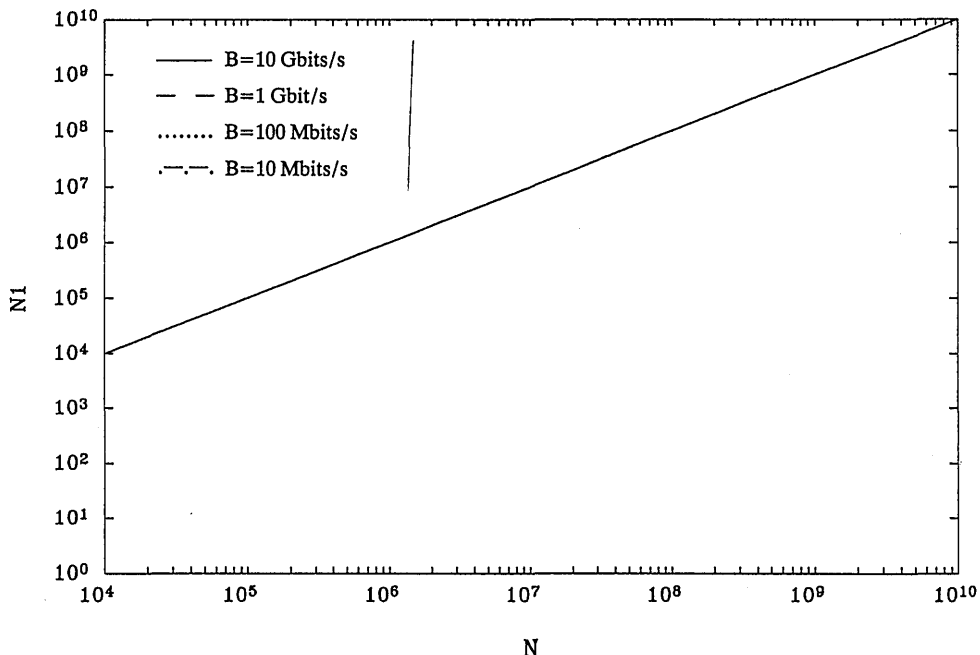

Fig. 14. N_1 versus N (repeaters): $p = 0.6$ or $p = 0.8$, $f = 50$, $E_o = 1$ pJ. The plots for all values of B coincide.

Thus unless B is very large so that the system is heat-removal limited (as discussed above), an all-repeatered system would be preferred. When $p = 0.8$, the power dissipation for the all-repeatered case grows faster so that hybridization according to Eqs. (24) and (25) may be beneficial, if minimizing power dissipation is considered to be more important than minimizing system size.

In conclusion, the only situations in which it is desirable to use optics is either when B is very large so that the system is heat-removal limited or when $f \sim 2, p$ is large and we give priority to minimizing the cost of power over the cost of system size. In practice it is relatively difficult to achieve $f \sim 2$ imaging. Furthermore, a larger number of conducting connection layers and faster devices may become possible in the near future (the latter of which will reduce the value of β). Thus it seems that as long as the optical communication paths are constrained to the plane, optics may have limited usefulness.

8. Three-Dimensional Systems

In this section we consider systems in which the modules are still confined to a plane, as in Fig. 4, but the optical communication paths are permitted to leave the plane. Such a system would most likely employ some type of free-space interconnection architecture.

Various other configurations are also possible. For instance, the modules may be arrayed throughout the volume of the system instead of being confined to the plane.[29] The N_1 electrically interconnected elements constituting each module may be arrayed in a three-dimensional lattice, if such a cubic chip can be manufactured. Such alternatives have been considered in Ref. 21. The general conclusion is that such choices have little or no effect on the optimal value of N_1 and the resulting performance and cost of the system.

Why is this so? Utilizing a volume for the establishment of the interconnections relaxes wireability requirements greatly, whereas it hardly has any effect on heat-removal requirements. (In Refs. 21 and 27 we show that the requirement $Q\mathscr{L}^2 \geq \mathscr{P}$ holds regardless of whether the system is two or three dimensional.) Thus in three-dimensional systems, heat removal becomes the dominant issue.[33,27,21] As a result, the particular choice of architecture, which essentially has to do with wireability requirements, has little or no effect.

A. Analysis

The analysis essentially follows the same lines as for two-dimensional systems; however, we concentrate exclusively on heat-removal requirements, on the basis of the above discussion.

Once again, N_1 must satisfy $N_1 \leq N_{1_{max}}$ [inequality (15)] and a condition similar to that given by inequality (26). However, we do not take these into consideration since the former condition can be avoided by use of repeaters, and the latter can be relaxed to an extent that it is not a limiting factor by use of either three-dimensional or two-dimensional modules with a large number of connection layers. We assume that these conditions are satisfied so as to be able to concentrate on the essential issues.

d_1 is given by the purely heat-removal-limited Eq. (18). Since we are ignoring wireability requirements, $d_m = d_1$, and since we are ignoring the electrical delay term (the latter condition in the
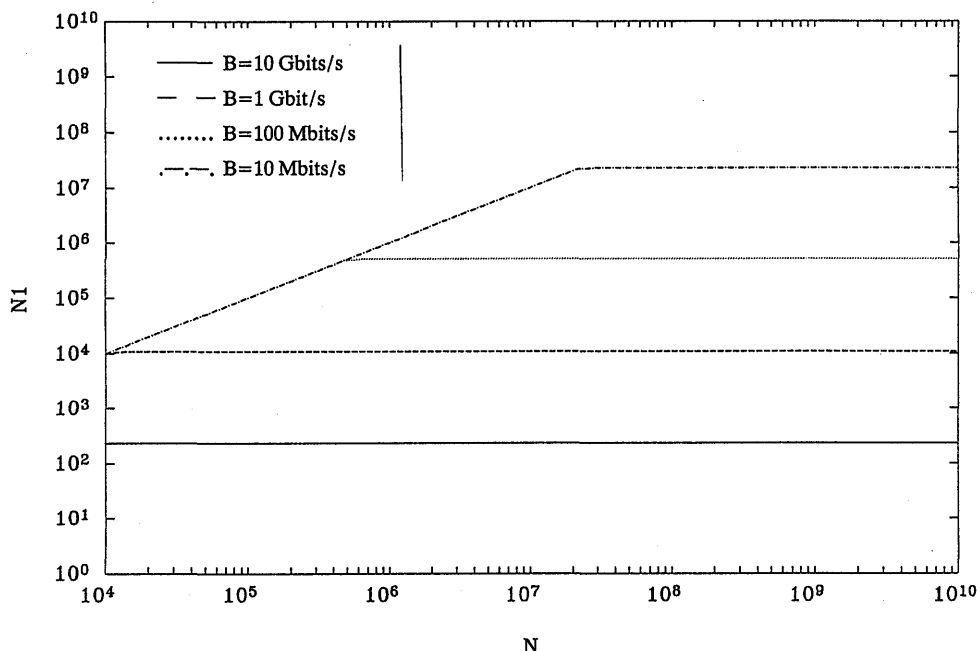

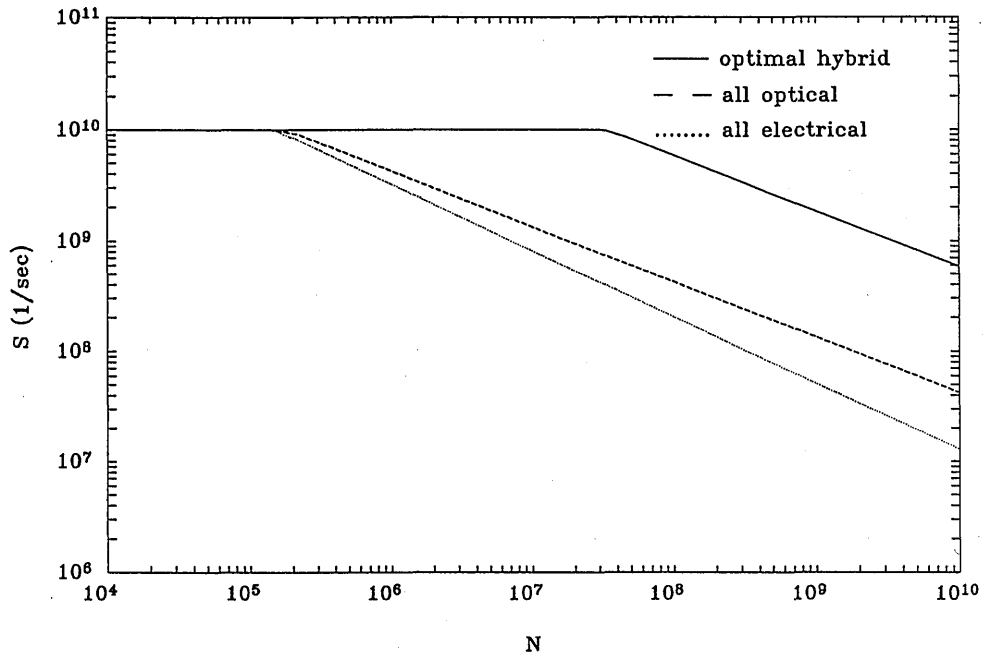Fig. 15. N_1 versus N (three dimensions): $p = 0.6, E_o = 1$ pJ.

Fig. 16. S versus N (three dimensions): $B = 100$ Mbits/s, $p = 0.6$, $E_o = 1$ pJ. Repeaters are assumed for the all-electrical case.

previous paragraph),

$$\tau = \max[(N/N_1)^{1/2}d_m/c, T_d]. \qquad (29)$$

Total power $\mathscr{P}$ is again given by relation (23). The value of N_1 minimizing $\mathscr{P}$, $\mathscr{L}$, and τ simultaneously is then easily shown to be given by Eq. (25). With this value of N_1 we find the system linear extent, total power, and delay to be given by

$$\mathscr{L} = N^{1/2}(kB/Q)^{1/2p}E_o^{1-1/2p}(\kappa\gamma)^{1/p-1}, \qquad (30)$$

$\mathscr{P} = Q\mathscr{L}^2$, and $1/S = \tau = \max(\mathscr{L}/c, T_d)$, respectively. Notice that $\mathscr{L} \propto N^{1/2}$ regardless of p.

B. Numerical Examples

First consider $p = 0.6$ and $E_o = 1$ pJ (Figs. 15 and 16). Unlike the two-dimensional layout discussed earlier in which the optimal value of S was equal to that for the all-optical case (Fig. 6), here the value of S for the optimal hybrid case is better than that for both all-electrical and all-optical alternatives. With Eqs.
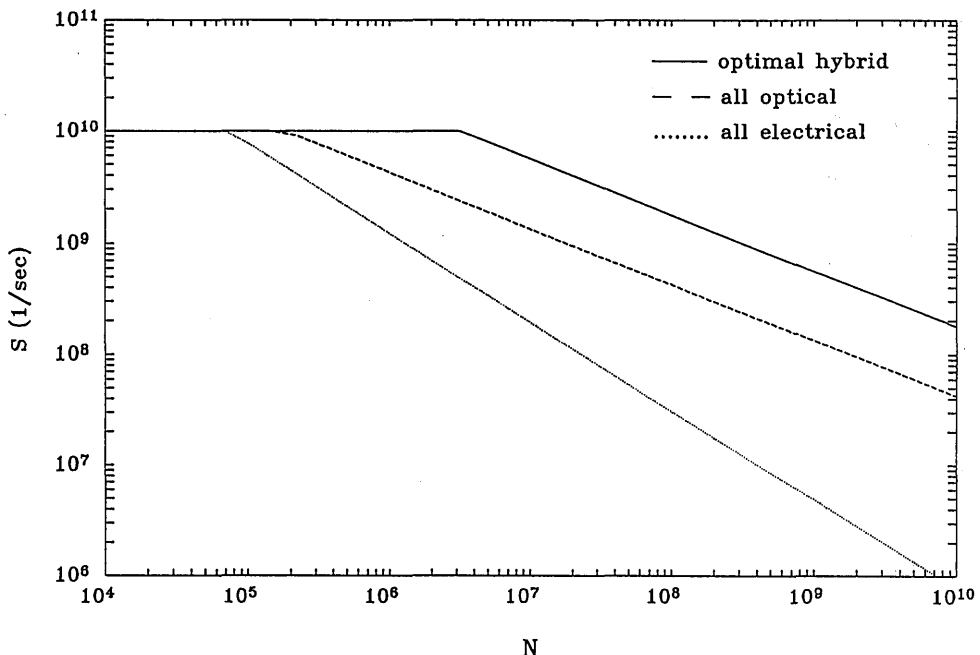

Fig. 17. S versus N (three dimensions): $B = 100$ Mbits/s, $p = 0.8$, $E_o = 1$ pJ. Repeaters are assumed for the all-electrical case.

(8) and (30) it is easily shown that an order of magnitude improvement in $S = 1/\tau$ and a hundred-fold reduction in power is achieved over the all-optical case, for $B = 100$ Mbits/s. This example provides a good illustration of how both media can be used in conjunction to obtain performance and cost much better than those possible with any alone.

We also note that the discrepancy in S between the all-optical and the optimal hybrid implementations is less for larger values of B. To see this, form the (greater than unity) ratio of the right-hand side of Eq. (8) to that of Eq. (30), which is $\propto B^{(1-1/p)/2}$, a decreasing function of B. This is also consistent with the fact that for larger B, the optimal value of N_1 is smaller; hence the optimal system is closer to the all-optical system ($N_1 = 1$).

Now let us look at the effects of an increase in p to 0.8 (Fig. 17). Remember that the value of S for an all-optical layout is independent of p [Eq. (8)]. For the optimal hybrid layout the value of S is somewhat worse when $p = 0.8$ than when $p = 0.6$. However, the effect of an increase of p on S is relatively small compared with the two-dimensional case.

Now let us consider an increase in E_o by a factor of 100. The value of $S = 1/\tau$ for the all-optical case [Eq. (8)] degrades by an order of magnitude. The value of S for the optimal hybrid case does not undergo a full order of magnitude degradation, however, since the system adapts itself by employing fewer optical interconnections (larger N_1). To see this, note that the power of E_o in Eq. (30) is $1 - 1/2p \le 1/2$.

Although we have concentrated on systems satisfying Rent's rule, our general method can be applied to the problem of location of the optimal value of N_1 for graphs with specific topologies, such as various types of permutation networks.

9. Discussion and Conclusions

In this section we present some of the more general conclusions derived from the study of a large number of examples, including those presented above. Up to a certain value of N (which usually includes the device-limited region for which $\tau = T_d$), it is preferable to make all connections electrical. In Ref. 14 it is shown that beyond a critical value of N, even an all-optical layout is preferable to an all normally conducting layout. Thus it is beneficial to start employing optics for the longer connections on or before this critical value.

Broadly speaking, we have found that the use of optics may enable performance (to increase performance may be interpreted to mean to be able to access previously unaccessible points in S-B-N space in a desired direction[14]) unachievable otherwise, however, possibly at a significant cost of space and/or power, depending on the information density and energy per bit achievable. The use of optical interconnections may enable smaller signal delay and larger bandwidth in comparison to an all-electrical system. However, they will usually not enable a reduction in our measures of cost. The system size and the power dissipation of the higher-performance hybrid system will often exceed that of the all-electrical system.

This behavior can be the result of a number of reasons. First of all, remember that optical lines cannot be scaled down to submicrometer dimensions; thus they lead to large system sizes. However, they may still enable greater performance than an all-electrical system, since the signal delay along a longer optical line can be less than the rise time of a shorter normally conducting line. A similar argument applies to heat-removal-limited systems. A hybrid system may dissipate more power than an all-electrical system, leading to larger system size, but may still have smaller delay.

Another reason why we must use optics despite the large cost of system size and power dissipation is because we may want to increase N and B simultaneously beyond what is possible with normal conductors. In such cases we might be willing to employ optical communication regardless of how large f, E_o, and the resulting cost of space and power is. We can avoid falling into this situation by using repeaters; however, there are still points in S–B–N space that are not accessible without the use of optics, no matter at what cost.

Figure-of-merit functions emphasizing the cost of system size and power dissipation tend to favor all-electrical systems until very large numbers of elements are used. If full priority is given to minimizing cost of system size and/or power dissipation, almost always an all-electrical system will be preferred. (The major exception arises when we want to increase N beyond $N_{1_{max}}$.) Even if equal priority is given to minimizing signal delay and measures of cost, as in our example (Section 6), all normally conducting systems are usually preferred until $N > N_{1_{max}}$. When repeaters are used, all-electrical layouts are preferred until even higher values of N, until eventually a hybrid layout is preferred because the signal delay becomes much worse than that possible with the hybrid system, outweighing the emphasis put on measures of cost.

Two-dimensional systems tend to be wireability limited. Consequently the resulting performance is found to depend strongly on the connectivity, as measured by p. How much we can approach diffraction-limited information densities has a significant effect on whether optics is worth using at all, especially if repeatered connections are employed. When hybridization is desirable, the optimal value of N_1 for large N is usually that which minimizes total power dissipation. As long as it is not very large, the value of E_o has little or no effect on the resulting performance (since the system is wireability limited), although it strongly affects the optimal value of N_1.

If $f \sim 2$ for two-dimensional systems, then a value of $E_o \sim 1$ pJ leads to optimal values of N_1 of the order of roughly 10^3–10^4. Increasing E_o hundredfold results in $N_1 \sim 10^5$–10^6. If repeaters are used, optical interconnections are useful only in a limited set of

circumstances. If $f > 10$ or so, there hardly seems to be any room for the beneficial use of optics, if repeaters are available. If not, we may resort to optics so as to achieve large values of N and B simultaneously. However, if f is too large, the resulting system size may quickly reach unrealistic proportions (> 10 m); thus such a system would not be feasible anyway.

Three-dimensional systems tend to be heat-removal limited; thus the value of f is relatively unimportant. A hundredfold degradation from what seems the best possible ($f \sim 2$) has little effect on the results. On the other hand, the value of E_o not only determines how much optics should be used but also has considerable effect on the resulting performance. If $E_o \sim 1$ pJ, the optimal value of N_1 is of the order of roughly 10^4–10^5, whereas a hundredfold increase in E_o leads to $N_1 \sim 10^7$.

Systems with smaller values of p are hurt less by an increase in E_o, since such systems have a smaller fraction of longer connections and are relatively less dependent on use of optics. On the other hand, for systems with large p an increase in E_o by 100 may lead to nearly an order of magnitude degradation in signal delay, nearly as much as the all-optical case. [To see this, note the dependence $\propto E_o^{1-1/2p}$ in Eq. (30).]

The value of p has less effect on resulting signal delay for three-dimensional systems. The power dissipation and the system size of a heat-removal-limited all-optical system are independent of p. For a hybrid system, smaller values of p enable some energy savings and consequent reduction in system size, since a greater fraction of connections will fall below the breakeven length for energy and can be replaced with normal conductors.

Hybridization is most important when the system is heat-removal limited. For two-dimensional systems, which are often wireability limited, hybridization has little effect on resulting signal delay. An all-optical system is just as good in this respect (although it may have much larger power dissipation).

When E_o is small, even the longest normally conducting lines in an optimal hybrid layout are usually short enough to be left unterminated (of the order of a centimeter or less). Terminated lines are not observed; RC lines are followed directly by optical lines. It is fortunate that the breakeven occurs at a point at which terminated lines are not yet necessary, since satisfactory termination can be a significant problem. For large values of E_o (or f in two-dimensional layouts), however, the module size may be forced to be large enough that the longer normally conducting lines have to be terminated. (The conclusions of this paragraph rely on the calculations of Ref. 21, in which terminated lines were taken into consideration.)

Based on the considerations of this paper, it seems fair to conclude that if the optical communication energy can be reduced to the order of 1 pJ, and if three-dimensional architectures with $f \sim 100$ or less can be realized, the use of optics has the potential to contribute significantly to the performance of large-scale systems. As a final remark, we note that there is considerable latitude in the choice of N_1, leaving room for other technological considerations. In other words, we can deviate quite a bit from the optimal value of N_1 while deviating only a little from the optimal performance and cost.

References

1. K. C. Saraswat and F. Mohammadi, "Effect of scaling of interconnections on the time delay of VLSI circuits," IEEE Trans. Electron Devices **ED-29**, 645–650 (1982).
2. H. B. Bakoglu, *Circuits, Interconnections and Packaging for VLSI* (Addison-Wesley, Reading, Mass. 1990).
3. W. D. Hillis, "New computer architectures and their relationship to physics or why computer science is no good," Int. J. Theor. Phys. **21**, 255–262 (1982).
4. R. W. Keyes, "Communication in computation," Int. J. Theor. Phys. **21**, 263–273 (1982).
5. J. W. Goodman, F. J. Leonbergere, S.-Y. Kung, and R. Athale, "Optical interconnections for VLSI systems," Proc. IEEE **72**, 850–866 (1984).
6. R. K. Kostuk, J. W. Goodman, and L. Hesselink, "Optical imaging applied to microelectronic chip-to-chip interconnections," Appl. Opt. **24**, 2851–2858 (1985).
7. P. R. Haugen, S. Rychnovsky, A. Husain, and L. D. Hutcheson, "Optical interconnects for high-speed computing," Opt. Eng. **25**, 1076 (1986).
8. L. D. Hutcheson and P. Haugen, "Optical interconnects replace hardwire," IEEE Spectrum (March 1987), pp. 30–35.
9. 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).
10. W. H. Wu, L. A. Bergman, A. R. Johnston, C. C. Guest, S. C. Esener, P. K. L. Yu, M. R. Feldman, and S. H. Lee, "Implementation of optical interconnections for VLSI," IEEE Trans. Electron Devices **ED-34**, 706–714 (1987).
11. H. Kroger, C. Hilbert, U. Ghoshal, D. Gibson, and L. Smith, "Applications of superconductivity to packaging," IEEE Circuits Devices Mag. (May 1989), pp. 16–21.
12. O. K. Kwon, B. W. Langley, R. F. W. Pease, and M. R. Beasley, "Superconductors as very-high-speed system-level interconnects," IEEE Electron Device Lett. **8**, 582–585 (1987).
13. R. C. Frye, "Analysis of the trade-offs between conventional and superconducting interconnections," IEEE Circuits Devices Mag. (May 1989), pp. 27–32.
14. 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.
15. 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).
16. M. R. Feldman, C. C. Guest, T. J. Drabik, and S. C. Esener, "Comparison between electrical and free space optical intercon-

nects for fine grain processor arrays based on interconnect density capabilities," Appl. Opt. **28**, 3820–3829 (1989).

17. C. W. Stirk and D. Psaltis, "Comparison of optical and electronic three-dimensional circuits," in *Microelectronic Interconnects and Packaging*, G. Arjavalingam and J. Pazaris, eds., Proc. Soc. Photo-Opt. Instrum. Eng. **1389**, 580–593 (1990).

18. F. E. Kiamilev, P. Marchand, A. V. Krishnamoorthy, S. C. Esener, and S. H. Lee, "Performance comparison bewteen optoelectronic and VLSI multistage interconnection networks," J. Lightwave Technol. **9**, 1674–1692 (1991).

19. 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).

20. H. M. Ozaktas and J. W. Goodman, "Optimal partitioning of very-large-scale optoelectronic computing systems," in *Annual Meeting*, Vol. 15 of 1990 Technical Digest Series (Optical Society of America, Washington, D.C., 1990), p. 87.

21. H. M. Ozaktas, "A physical approach to communication limits in computation," Ph.D. dissertation (Stanford University, Stanford, Calif., 1991).

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

23. W. E. Donath, "Wire length distribution for placements of computer logic," IBM J. Res. Dev. **25**, 152–155 (1981).

24. M. Feuer, "Connectivity of random logic," IEEE Trans. Comput. **C-31**, 29–33 (1982).

25. A. El Gamal, "Two-dimensional stochastic model for interconnections in master slice integrated circuits," IEEE Trans. Circuits Syst. **CS-28**, 127–134 (1981).

26. W. R. Heller, W. F. Mikhail, and W. E. Donath, "Prediction of wiring space requirements for LSI," J. Des. Autom. Fault Tolerant Comput. **2**, 117–144 (1978).

27. 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).

28. 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).

29. 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).

30. M. R. Feldman and C. C. Guest, "Interconnect density capabilities of computer-generated holograms for optical interconnection of very-large-scale-integrated circuits," Appl. Opt. **28**, 3134–3137 (1989).

31. H. B. Bakoglu, and J. D. Meindl, "Optimal interconnection circuits for VLSI," IEEE Trans. Electron. Devices **32**, 903–909 (1985).

32. 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).

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