

Comparison of fully three-dimensional optical, normally conducting, and superconducting interconnections

Haldun M. Ozaktas and M. Fatih Erden

Several approaches to three-dimensional integration of conventional electronic circuits have been pursued recently. To determine whether the advantages of optical interconnections are negated by these advances, we compare the limitations of fully three-dimensional systems interconnected with optical, normally conducting, repeatered normally conducting, and superconducting interconnections by showing how system-level parameters such as signal delay, bandwidth, and number of computing elements are related. In particular, we show that the duty ratio of pulses transmitted on terminated transmission lines is an important optimization parameter that can be used to trade off signal delay and bandwidth so as to optimize applicable measures of performance or cost, such as minimum message delay in parallel computation. © 1999 Optical Society of America

OCIS codes: 200.0200, 200.4650.

1. Introduction

Several approaches to three-dimensional integration of conventional electronic circuits have been pursued recently (see Ref. 1 and references cited therein). Some have claimed that these developments negate the advantages of optical interconnections. In this paper we compare the utilities and the limitations of fully three-dimensional circuit layouts based on optical, normally conducting, repeatered normally conducting, and superconducting interconnections. We show that, even if fully three-dimensional plain or repeatered normally conducting interconnections are possible, they are still inferior to optical interconnections. Fully three-dimensional superconducting interconnections, however, are comparable with optical interconnections.

Present-day VLSI technology is a very mature technology in the sense that, for a given lithographic patterning accuracy, the interconnections on a chip are packed almost as closely as fundamental limitations would allow for. Little room for improvement by modification of parameters such as the aspect ra-

tios of the lines and so forth is left, since these have already been fairly optimized. Further reduction of linewidths, even if achievable, are unlikely to offer significant returns, because, at such linewidths, it is not the linewidth but heat removal and other considerations that limit the size and the performance of the chip.

What we refer to as a fully three-dimensional circuit layout is best imagined as the three-dimensional version of a VLSI chip. Thus a fully three-dimensional VLSI "chop," as we might refer to it, is assumed to be the three-dimensional analog of a VLSI chip in the sense that active circuits and interconnections can occupy and be routed through three-dimensional space with the same kind of freedom and efficiency with which they are routed through (several layers of) two-dimensional space in VLSI chips.

A VLSI chop, as we have described it, may not be feasible in the near future, because there seems to be no effective way of manufacturing it. However, for highly interconnected circuits, the following result is of relevance²⁻⁵: Provided the interconnections are permitted to be routed through three-dimensional space, there is no disadvantage in restricting the active devices to a plane. That is, the overall volume and signal delay and the thus clock rate of a highly interconnected three-dimensional system in which the active devices are restricted to a plane will not be inferior to a system in which the active devices can be situated freely throughout the volume. This is because the volume and the delays in a highly interconnected system are determined primarily by the space

H. M. Ozaktas (haldun@ee.bilkent.edu.tr) is with the Department of Electrical Engineering, TR-06533 Bilkent, Ankara, Turkey. M. F. Erden is with Massana Ltd., 5 Westland Square, Dublin 2, Ireland.

Received 7 July 1998; revised manuscript received 30 August 1999.

0003-6935/99/357264-12\$15.00/0

© 1999 Optical Society of America

occupied by the interconnections so that additional restrictions on the active devices are of little or no significance. This result, which is based on purely geometrical considerations, is valid for all types of interconnection, including optical interconnections.

The above result is encouraging in that it seems to indicate that three-dimensional chips may be manufactured by use of the planar process to manufacture two-dimensional chips with a large number of wiring layers. However, it is important to realize that a many-layered chip thus manufactured will merit being called a fully three-dimensional chip only if the number of layers is comparable with the number of line tracks across the horizontal extent of the chip. Otherwise, it is merely a two-dimensional chip with a large number of layers and will exhibit the general properties and limitations of two-dimensional systems. Moreover, even if the number of layers can be as great as the number of line tracks across the chip, it is possible that the quality (capacitance, resistance, and thus speed, density, and coupling) of vertically running lines may be inferior to horizontally running lines.

An alternative approach that has been pursued is to stack a large number of two-dimensional chips (or wafers) on top of one another and connect them with a grid of regularly spaced interconnections from one chip to another.⁶ Although such structures have been referred to as three-dimensional, they are not fully three-dimensional in the sense we have defined, because the number of chips stacked is small compared with the number of line tracks across the chips. However, if each chip incorporates a large number of wiring layers and if the number of chips is further increased, it is conceivable that in the near future the total number of layers in all chips may approach the number of line tracks on a single chip, and thus such systems merit being called fully three-dimensional chips. For this, however, it will also be necessary to ensure that the density, speed, and routability of vertical lines approach that of those confined to a single chip.

Even if unrealizable, the fully three-dimensional VLSI chip serves as an appropriate basis for comparison with optically interconnected systems, because it represents the best case for normally conducting interconnections. Since we show that optical and superconducting interconnections both offer performance that is superior to even this best case, the difficulty in manufacturing fully three-dimensional normally conducting chips only strengthens our conclusion.

The above discussion also applies to repeated normally conducting interconnections. However, since repeated lines incorporate active devices, they can be optimally utilized only if active devices are permitted on all layers. Thus the use of repeaters, which, as we will see, offers a significant performance advantage over plain normally conducting lines, is more demanding from a manufacturing viewpoint. Given this, it is not possible to make a simple argument in favor of one over the other. However,

this point will not be of concern to us, since we will show that optical and superconducting interconnections are superior to both.

Having said this, it also seems necessary to clarify the realizability of fully three-dimensional optical and superconducting circuits. Fully three-dimensional superconducting circuits are subject to considerations similar to those of plain normally conducting circuits, apart from the fact that the materials involved are different. Thus it seems realistic to expect that similar circuit structures can be manufactured with both technologies. However, superconducting interconnections offer performance that is clearly superior to that of both plain and repeated normally conducting interconnections.

Optical interconnections offer performance that is superior to plain and repeated normally conducting interconnections and that is comparable with superconducting interconnections, if we assume similar structural and manufacturability constraints on all technologies. However, we can almost certainly get closer to full three dimensionality with optical interconnections than with any of the other technologies; so optics emerges as the best option.

Although most people are readily willing to accept the fact that only optics will allow for fully three-dimensional circuits, there is a point in more careful justification of this as well. Our justification rests on three points. The first point is based on the previously stated observation that there is nothing to be lost by constraining the active devices to lie on a plane, provided that the interconnections are able to occupy three-dimensional space. This is important because there does not seem to be an efficient and practical optical architecture that can provide interconnections among a three-dimensional array of optical sources and detectors (although one approach is suggested in Ref. 7); most optical architectures provide interconnections between planes of sources and detectors. Second, it is possible to realize arbitrary patterns of interconnections with optical interconnections with the effective interconnection density assumed in the models we use below to show the superiority of optics.⁸⁻¹⁰ Thus the density of optical communication is determined primarily by the density of the optical sources, modulators, or detectors. This brings us to our third point. Currently, self-electro-optic effect devices flip-chip bonded to silicon chips^{11,12} can be spaced at least as close as 50- μm apart and perhaps even closer. Although this is quite large compared with the size of an optical wavelength, in light of what is achievable today it seems fair to assume that the density of vertical connections in three-dimensional electrical technologies will not greatly exceed (if they do exceed) the density of this and other optical technologies so that optical interconnections will always allow us to construct systems that are comparatively closer to full three dimensionality.

To support our claim that optics is superior, we are relying only on optical systems' being at least as three-dimensional as normally conducting systems

(since we show that optics is superior when both are fully three-dimensional). However, superconductors offer performance that is comparable with optics when both are fully three-dimensional. Thus, if some kind of volume lithography or other technique allows for the construction of fully three-dimensional superconducting circuits in the future, these may contend with optical circuits. In that case the choice will be based on a number of factors not accounted for in this study, such as voltage isolation, impedance matching in the presence of multiple taps, electromagnetic interference, etc. Finally, we add that, even solely on the basis of the considerations of this study, there are a number of circumstances in which fully three-dimensional superconducting systems may be preferable to optical ones.

2. General Considerations

Our comparison of different interconnection technologies will be based on the trade-off relations between the quantities $S = 1/\tau$ (inverse signal delay or latency), B (bit repetition rate or bandwidth), and N (number of devices or computing elements). These trade-off relations determine the largest values of S , B , and N that are simultaneously achievable; given any two of these parameters, we can determine the largest value attainable by the other. Another quantity of importance, denoted by H , is the bisection of the system. The bisection is the number of connections crossing an imaginary plane that divides the system into two approximately equal parts. The bisection-bandwidth product HB and the bisection-inverse-delay product HS are often more meaningful measures of performance or throughput than aggregate bandwidth NB or NS . It is often appropriate to model the dependence of the bisection on N by a power-law expression of the form $H = k\kappa N^p$, where κ is a constant, $2/3 \leq p \leq 1$ is a measure of the connectivity of the system (larger p means greater connectivity), and k is the average number of connections per element.¹³ [We can let k be absorbed in the definition of κ but prefer not to do so. As defined, κ is given in terms of p by $\kappa = (p - 2/3)^{-1}(p + 1/3)^{-1}$]. Yet another interesting quantity that is of importance in certain contexts is the first-to-last bit message delay $\tau_L = \tau + L/B$, where L is the bit length of a message (for definitions of all symbols used, please see Appendix A).

We assume that the system consists of N primitive elements or devices each of size d_d , arrayed in the form of a three-dimensional $N^{1/3} \times N^{1/3} \times N^{1/3}$ grid. The grid spacing will be denoted by d and the linear extent of the system by $\mathcal{L} = N^{1/3}d$. We further assume that there is an average of k connections per element and that the average length of each connection in grid units is \bar{r} . Each logical connection may be realized with $\chi \geq 1$ physical interconnections to increase the bandwidth B .

Heat-removal considerations imply that a system dissipating a total power \mathcal{P} must have a minimum cross-sectional area of at least \mathcal{P}/Q , where Q is the maximum amount of power that can be removed per

unit area.¹⁴ This implies that the linear extent of the system \mathcal{L} be at least $\mathcal{L} = (\mathcal{P}/Q)^{1/2}$.

Three basic considerations apply to all interconnection media. The total volume occupied by the elements is Nd_d^3 . The total volume occupied by the interconnections is $k\chi N\bar{r}dW^2$, since $k\chi N$ is the total number of physical connections, $\bar{r}d$ is their average length in physical units, and $W^2 = A$ is the cross-sectional area of an interconnection. Heat-removal considerations imply a volume of at least $(\mathcal{P}/Q)^{3/2} = (kNBE/Q)^{3/2}$, where E denotes the energy per transmitted bit, since there are kN connections, each dissipating power BE . Thus the linear extent of the system $\mathcal{L} = N^{1/3}d$ must satisfy

$$\mathcal{L}^3 = Nd^3 = \max[d_d^3 N, k\chi\bar{r}W^2 Nd, (kEB/Q)^{3/2} N^{3/2}]. \quad (1)$$

It is known that the bisection and the average connection length of a system are closely related. If the bisection can be modeled as $k\kappa N^p$, then the average connection length is $\bar{r} = \kappa N^{p-2}$.¹³ Inserting this into the above, we obtain

$$\mathcal{L} = N^{1/3}d = \max[d_d N^{1/3}, (k\chi\kappa)^{1/2} W N^{p/2}, (kEB/Q)^{1/2}]. \quad (2)$$

Further progress requires the introduction of physical interconnection models, which we do below.

3. Interconnection Models

The interconnection models used are summarized in Tables 1–4 below. Their justification and derivation have been given elsewhere² (also see Refs. 15 and 16 for further references). These models are simply relations between the *external* parameters of interconnections, which are

1. Interconnection length l .
2. Cross-sectional area A or transverse linear extent W , where $A = W^2$. These parameters define packing density and thus include any necessary line-to-line separations.
3. Signal delay τ , which is given by the greater of the propagation delay T_p and the minimum temporal pulse width T , which in turn is the greater of a line-imposed component T_l and a device-imposed component T_d :

$$\tau = \max(T_p, T), \quad T = \max(T_l, T_d).$$

4. Minimum pulse repetition interval T_r , which is usually equal to T , the minimum temporal pulse width along the interconnection

$$T_r = T.$$

5. The energy per transmitted bit, E .

The relationships in Tables 1–4 represent a full characterization of the physical properties of the interconnections, as far as this study is concerned. All other

Table 1. Optical Interconnection Model^a when $T_i \leq T_d \leq T_p$

W	τ	$T = T_r$	E
$f\lambda = \text{constant}$	l/c	$T_d = \text{constant}$	constant

^aThe delay $\tau = T_p$ is a function of length l only. The pulse width $T = T_d$ and the energy E are assumed independent of length l and width W .

internal parameters that are commonly used to characterize such lines, such as transverse aspect ratio, capacitance, resistance, etc., are assumed to be set (at least approximately) to their optimal values in the derivation of these models and thus do not appear in these models.

For optical interconnections (Table 1) the cross-sectional area 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 larger in practice. (It is important to note that for certain classes of optical interconnection architectures, which we collectively refer to as multifacet architectures, f is not a constant but increases with N . Such architectures are clearly

Table 2. Normally Conducting Interconnection Model^a when $T_d \leq T_i$

	τ	$T = T_r$	E	Termination
$W^2 \leq 16\rho\epsilon vl$	$16\rho\epsilon \frac{l^2}{W^2}$	$16\rho\epsilon \frac{l^2}{W^2}$	$2\epsilon V^2 l$	no
$W^2 \geq 16\rho\epsilon vl$	$\frac{l}{v}$	$16\rho\epsilon \frac{l^2}{W^2}$	$2\epsilon V^2 v T$	yes

^aThe delay $\tau = \max(T, T_p)$, pulse width $T = \max(T_i, T_d)$, and energy E are given as functions of length l and width W . $T_p = l/v$ is the propagation delay. The final column indicates whether the line is to be terminated in that region.

Table 3. Repeater Interconnection Model^a

	τ	$T = T_r$	E	Termination
$W \leq 4 \left(\frac{\rho R_0 C_0}{\mu} \right)^{1/2}$	$4(R_0 C_0 \rho \epsilon)^{1/2} \frac{l}{W}$	$R_0 C_0 = \text{constant}$	$2\epsilon V^2 l$	no
$W \geq 4 \left(\frac{\rho R_0 C_0}{\mu} \right)^{1/2}$	$\sqrt{\mu \epsilon} l$	$R_0 C_0 = \text{constant}$	$8\epsilon V^2 \left(\frac{\rho R_0 C_0}{\mu} \right)^{1/2} \frac{l}{W}$	yes

The delay τ , pulse width T , and energy E are given as functions of length l and width W . The final column indicates whether the stages of the line are to be terminated in that region.

Table 4. Superconducting Interconnection Model when $T_d \leq T_p$ ($T_d \leq T_i$ in the Underterminated Case)^a

	τ	$T = T_r$	E	Termination
$W \leq \frac{4V}{J_{sc} \sqrt{\mu/\epsilon}}$	$\frac{16\epsilon V \lambda_p}{J_{sc}} \frac{l}{W^2}$	$\frac{16\epsilon V \lambda_p}{J_{sc}} \frac{l}{W^2}$	$2\epsilon V^2 l$	no
$\frac{4V}{J_{sc} \sqrt{\mu/\epsilon}} \leq W \leq 4\lambda_p$	$\frac{4\lambda_p}{v} \frac{l}{W}$	$T_d = \text{constant}$	$2 \sqrt{\frac{\epsilon}{\mu}} V^2 \frac{W}{4\lambda_p} T_d$	yes
$W \geq 4\lambda_p$	$\frac{l}{v}$	$T_d = \text{constant}$	$2 \sqrt{\frac{\epsilon}{\mu}} V^2 T_d = \text{constant}$	yes

The delay $\tau = \max(T_p, T)$, pulse width T , and energy E are given as functions of length l and width W . The final column indicates whether the line is to be terminated or not in that region.

undesirable. See Refs. 7–10 for further discussion.) The signal delay is taken to be the greater of the speed-of-light delay and the device rise time: $\tau = \max(l/c, T_d)$. Since the effects of dispersion and attenuation can be made small for the length scales in consideration, $T = T_r$ and E are assumed to be constants. This model is valid for free-space interconnections as well as guided-wave interconnections, although the value of f can be much smaller with free-space interconnections.⁸

In Table 2 we see the relationships tying the length, cross-sectional area, delay, and energy for normally conducting lines for the case $T_d \leq T_i$. The symbols ρ , ϵ , μ , and $v = (\epsilon\mu)^{-1/2}$ denote the resistivity of the conductor, permittivity and permeability of the dielectric, and propagation velocity in the dielectric, respectively. V denotes the nominal voltage level. What is unique to our model is that it intrinsically accounts for the proper scaling effects that are due to the skin effect and deals with unterminated RC lines and transmission lines in a unified manner. One conclusion that may be derived on careful inspection of our model is that the use of many narrow lines is not more beneficial than a single wide line in terms of increasing information density. (Use of a single wide line amounts to setting $\chi = 1$.) Yet another conclusion is that it is beneficial to photographically scale down interconnection-density-limited layouts until we are in the unterminated region. After this, further reduction in scale does not further improve system signal delay. (This conclusion is related to a well-known argument stating that the rise time of a RC line remains constant when all of its dimensions are downscaled.)

The use of active repeater devices along the line changes the relationship for delay versus linewidth

as shown in Table 3. R_0C_0 denotes the intrinsic delay of the repeating devices. The optimal number of repeaters (which may be zero) is used at optimal spacings for each line.¹⁷

Table 4 is for superconducting lines. Our models take into account the proper scaling effects associated with the superconducting penetration depth λ_p and the critical current density J_{sc} . A conclusion that can be derived on careful examination of the table is that for wire-limited layouts it is optimal to scale down the system until we are in the intermediate region (the second line of the table).

We will assume the following in our numerical examples: $f\lambda = 10 \mu\text{m}$, $c = 3 \times 10^8 \text{ m/s}$, $T_d = 100 \text{ ps}$, $E = 1 \text{ pJ}$, $V = 1 \text{ V}$, $\rho = 0.0274 \Omega \mu\text{m}$ (aluminum), $\epsilon = 3.9 \times 8.85 \times 10^{-3} \text{ fF}/\mu\text{m}$ (silicon dioxide), $\mu = 4\pi \times 10^{-7} \text{ H/m}$ (nonmagnetic materials), $v = 1.52 \times 10^8 \text{ m/s}$, $R_0C_0 = 100 \text{ ps}$, $\lambda_p = 0.2 \mu\text{m}$, and $J_{sc} = 50 \text{ mA}/\mu\text{m}$.

4. Analysis

A. Optical Interconnections

In the case of optical interconnections, W and E are taken as constants. Assuming that T_d does not dominate the propagation delay $S = 1/\tau = 1/(l/c) = 1/(\mathcal{L}/c) = c/\mathcal{L}$, the trade-off relation between S , B , and N is obtained from Eq. (2) as

$$\frac{1}{S} = \frac{1}{c} \max[d_d N^{1/3}, (k\chi\kappa)^{1/2}(f\lambda)N^{p/2}, (kEB/Q)^{1/2}N^{1/2}],$$

$$\chi = \max(1, BT_d). \quad (3)$$

Clearly, we choose χ to be equal to $B/(1/T_d)$ when this ratio is greater than 1 to support that value of B , but of course we cannot have less than one channel even if this ratio is less than 1. This relationship defines a surface in the three-dimensional parameter space defined by S , B , and N . By examining this relationship, we can tell the price we have to pay in terms of a decrease in one or two of these parameters to increase the remaining one(s). Note that there are several regions that might be referred to as the element-size-limited region (when the first term dominates), the interconnection-density-limited region (when the second term dominates), and the heat-removal-limited region (when the third term dominates).

B. Normally Conducting Interconnections

Unlike with optical interconnections where W is constant, with normally conducting interconnections we are free in choosing W , provided it exceeds a certain minimum manufacturable value W_{\min} . If d_d is small and heat removal is not an issue, we would prefer to set W to this minimum so as to make d and the overall system as small as possible. However, the minimum manufacturable value of W is not the only determinant of how small the system can be made. If element size or heat removal require that we set $\mathcal{L}^3 = Nd^3 > kNrdW_{\min}^2$, we will agree to increase W

until $Nd^3 = kNrdW^2$. (If d and hence the lengths of the lines are already set by factors other than interconnection density, we increase W so as to fill up available space. In this way we reduce the resistance of the lines as much as possible. Keeping the lines narrow while we have extra space around is clearly suboptimal.) We also assume that device delays T_d are small so that $T = T_l$.

Thus, using $d^2 = k\bar{r}W^2$, $l = \mathcal{L} = N^{1/3}d$, and the interconnection model, we obtain

$$T = (16\rho\epsilon)(k\bar{r})N^{2/3}, \quad (4)$$

and the maximum value of B satisfies $B = 1/T$ or

$$BN^p = (16\rho\epsilon)^{-1}(k\kappa)^{-1}, \quad (5)$$

an expression in which S does not appear. By definition, $S = 1/\tau$ may never exceed $1/\max(T_l, T_d, T_p) = 1/\max(T_l, T_p) = 1/\max(T, T_p)$. Thus the above relation for B may be used to find $S = 1/\max(1/B, T_p) = \min(B, 1/T_p)$. From Table 2 we see that the condition for $T_p < T$ is $W^2 < 16\rho\epsilon vl$. As we scale down the system photographically, all linear dimensions are decreased in proportion. Thus, below a certain critical W , this condition is satisfied so that propagation effects need not be considered, and we have ensured that S is not worse than B .

However, a number of factors may be an impediment to downscaling. First, our lithographic accuracy may not enable us to pattern lines sufficiently fine. This is not an issue with the availability of submicrometer scaling. Second, the system cannot be made smaller than dictated by the size of the elements $N^{1/3}d_d$. This is also not a limitation in highly connected systems with a large number of elements. Finally, heat-removal limitations may prevent us from scaling the system down sufficiently.

For typical parameter values two-dimensional layouts may be downscaled to the extent that propagation effects need not be considered on the longest line (i.e., $T_p = l/v < T$) so that S is simply equal to B .² To understand this, note that T does not depend on the scale of the system as measured by the value of the grid spacing d . However, the propagation delay along the longest line $T_p = N^{1/3}d/v$ depends linearly on d . Thus, as we downscale the system by reducing d , eventually T_p will fall below T , and T_p will not be a limiting factor in determining τ and S . This may not be possible for three-dimensional layouts, especially for room-temperature voltages. This is because we cannot downscale the system below a certain value of d and still be able to remove the dissipated heat. Thus the propagation delay T_p may remain larger than T , and thus the value of S may be quite less than the value of B . This is the essential qualitative difference between two-dimensional systems and three-dimensional systems.

Now we give a complete analysis of the effects of heat removal, enabling us to determine the scale of the system as set by heat-removal considerations. We will assume that pulses of identical temporal width are launched into all lines regardless of their

length. Thus the minimum value of this pulse width is set by the longest connection. (In principle, there is nothing that stops us from launching shorter pulses into the shorter lines, resulting in some energy savings. The following analysis may be modified for this case if such an approach is deemed practical.) The minimum pulse width for the longest line is given by

$$T = (16\rho\epsilon) \frac{l^2}{W^2}, \quad (6)$$

with $l = \mathcal{L}$. The above expression for T is independent of the scale of the system. This is because both l and W will change by the same factor when the system is scaled. Pulses of this duration are emitted into lines of all lengths. According to our interconnection model, the shorter lines, for which $T \geq T_p$, will be left unterminated, whereas the longer lines, for which $T < T_p$, will be terminated. (On terminated lines, several pulses of length T might simultaneously be in transit along the line.) Let us denote the break-even length for termination (in grid units) as r_x . Thus lines for which $r_x d/v \leq T$ will be left unterminated and those for which $r_x d/v > T$ will be terminated. If the length $N^{1/3}d$ of the longest line in our system satisfies $N^{1/3}d \leq Tv$, all lines will be unterminated.

The problem is that initially we do not know d , which depends on the total power dissipated, which in turn depends not only on d but also on what fraction of the lines are unterminated. Let us assume initially that $N^{1/3}d \leq Tv$ so that all lines are unterminated. Then the average energy per bit is given by $2\epsilon V^2 \bar{r}d$, the power dissipation by $2\epsilon V^2 \bar{r}dB$, and the total power dissipation by $kN2\epsilon V^2 \bar{r}dB$. This must be less than the cross-sectional area of the system $N^{2/3}d^2$ times Q . Thus

$$d \geq \frac{2\epsilon V^2 k \kappa N^{p-1/3} B}{Q}. \quad (7)$$

Now, if indeed $N^{1/3}d \leq Tv$, justifying our assumption, we are done and d is given by relation (7). If not, this means that some of the longer lines will be terminated, for which the energy per bit is given by $2\epsilon V^2 vT$. Then the total power consumption and heat-removal condition may be expressed in terms of a piecewise integral

$$QN^{3/2}d^2 \geq \left[\int_1^{r_x} 2\epsilon V^2 r dg(r) dr + \int_{r_x}^{N^{1/3}} 2\epsilon V^2 v T g(r) dr \right] B, \quad (8)$$

where $g(r)$ is the line-length distribution,^{2,3} given approximately by $g(r) \approx (-d/dr)[kr^{3(p-1)}(1 - r^3/N)]$ for the analytical form of the bisection we have assumed. The first integral represents the power dissipated on the unterminated lines, and the second represents

the power dissipated on the terminated lines. Evaluating the above, we find

$$Qd^2 \geq N^{1/3}k[2\epsilon V^2 \kappa r_x^{3p-2}d + 2\epsilon V^2 v T r_x^{3(p-1)}z]B, \quad (9)$$

where $z \leq 1$ is a factor whose exact form will not be important. Now, using $r_x = vT/d$, it is possible to solve for d as

$$d^{3p-1} \geq \frac{2\epsilon V^2 (vT)^{3p-2} k \kappa N^{1/3} B}{Q}, \quad (10)$$

where, since $z \leq 1$, we replaced $\kappa + z \approx \kappa$ with little error. Note that this expression forms continuity with relation (7) at $N^{1/3}d = vT$. Thus, combining the two expressions, we may write the minimum interelement spacing d as set by heat removal in the form

$$d = \min \left\{ \frac{2\epsilon V^2 k \kappa N^{p-1/3} B}{Q}, \left[\frac{2\epsilon V^2 (vT)^{3p-2} k \kappa N^{1/3} B}{Q} \right]^{\frac{1}{3p-1}} \right\}. \quad (11)$$

The virtue of this equation is that it combines compactly all possible cases. Finally, the signal delay is given by $1/S = \max(N^{1/3}d/v, N^{1/3}d_d/v, T)$.

C. Repeated Interconnections

First, assume that element size and heat removal need not be considered. Using $\tau = 4(R_0 C_0 \rho\epsilon)^{1/2} l/W$, $l = N^{1/3}d$, and $d^2 = k\chi \bar{r}W^2$, we obtain

$$SN^{p/2} = [4(R_0 C_0 \rho\epsilon)^{1/2}]^{-1} (k\chi\kappa)^{-1/2}, \quad \chi = \max(1, BR_0 C_0), \quad (12)$$

where we assumed that the system has been down-scaled sufficiently so that $W \leq 4(\rho R_0 C_0/\mu)^{1/2}$. This relation is similar to the corresponding relation for optical communication in form [the second term of Eq. (3)].

Equation (12) is scale independent when $W \leq 4(\rho R_0 C_0/\mu)^{1/2}$, as we have assumed to be the case ($W \leq 5 \mu\text{m}$ for the assumed parameter values). When element size is accounted for, S is given by the minimum predicted by Eq. (12) and $1/S = \sqrt{\mu\epsilon} N^{1/3} d_d$.

When heat-removal considerations are taken into account, it may be the case that it is not possible to downscale the system so that $W \leq 4(\rho R_0 C_0/\mu)^{1/2}$. In this case it is possible to show that the minimum value of d is given by

$$d = \min \left\{ \frac{2\epsilon V^2 k \kappa N^{p-1/3} B}{Q}, \left[\frac{8\epsilon V^2 (\rho R_0 C_0/\mu)^{1/2}}{Q} \right]^{1/2} \times \chi^{1/4} (k\kappa)^{3/4} N^{3p/4-1/3} B^{1/2} \right\}. \quad (13)$$

Thus the resulting delay is the greater of $\sqrt{\mu\epsilon} N^{1/3} d$ with d given above, $\sqrt{\mu\epsilon} N^{1/3} d_d$, and that given by Eq. (12).

D. Superconducting Interconnections

As with normal conductors, we agree to choose W so that the condition $d^2 \geq k\chi\bar{r}W^2$ is always satisfied with equality. Mostly we will be at an advantage (because of the inverse dependence of τ on W for given l), and never at a disadvantage, by doing so.

When we refer to Table 4, some reflection reveals that if d_d is small and heat removal need not be considered, it is optimal to work in the intermediate region, assuming we can manufacture $W \leq 4\lambda_p$. (To see this, note that, as we scale the system photographically, l varies in linear proportion to W .) Assuming T_d is small, an analysis similar to that for repeaters results in

$$SN^{p/2} = \left(\frac{v}{4\lambda_p}\right)(k\chi\kappa)^{-1/2},$$

$$\chi = \max(1, BT_d). \quad (14)$$

This relation is independent of the specific choice of W , provided that it lies between $4V/(J_{sc}\sqrt{\mu/\epsilon})$ and $4\lambda_p$.

Heat removal has no effect on performance unless it requires that d be large enough that $W \geq 4\lambda_p$. The analysis and the results are similar to the optical case, provided that we replace the energy $E \leftarrow 2V^2\sqrt{\epsilon/\mu}T_d$:

$$\frac{1}{S} = \frac{1}{v} \max[d_d N^{1/3}, (k\chi\kappa)^{1/2}(4\lambda_p)N^{p/2},$$

$$(k2V^2\sqrt{\epsilon/\mu}T_d B/Q)^{1/2}N^{1/2}],$$

$$\chi = \max(1, BT_d). \quad (15)$$

If low voltage values are used, this energy can be much less than ever achievable with optical interconnections.

5. Comparisons

A. Qualitative Comparisons

For optical and superconducting interconnections there is no upper limit to B for any value of N . One can simultaneously choose N and B arbitrarily large. Thus one can also increase the bisection–bandwidth product $HB = k\kappa N^p B$ arbitrarily, at the expense of greater signal delay (smaller S). The value of S corresponding to a given value of B and N may be found from Eq. (3) or Eq. (15). A particularly simple trade-off between bisection–bandwidth product HB and S is obtained when interconnection density is the dominating consideration (the second term in Eq. (3) or Eq. (15)):

$$S(HB)^{1/2} = \text{constant}. \quad (16)$$

The bisection–inverse-delay product HS can also be determined from Eq. (3) or Eq. (15). Assuming d_d is small, we can show

$$HS \propto \frac{N^{p-1/2}}{B^{1/2}}, \quad (17)$$

from which we see that this measure of performance can also be increased without bound.

For normally conducting interconnections, B is related to N through Eq. (5) so that, for a given value of N , it is not possible to increase B beyond that dictated by this equation. This is in contrast with the optical and the superconducting cases where B could be arbitrarily increased. Any attempt at increasing B by using wider lines or $\chi > 1$ parallel channels is thwarted by the increase in line lengths, since $T \propto l^2/W^2$. We further see that the bisection–bandwidth product is a constant and cannot be increased beyond a certain value:

$$HB = \text{constant}. \quad (18)$$

Thus even when we assume very fast devices (T_d negligible), negligible element size d_d , arbitrarily small manufacturable linewidths W_{\min} , a fully three-dimensional layout, and ignore the effects of heat removal, we still find that there is a fundamental upper limit to the bisection–bandwidth product for normally conducting interconnections.

We have seen that, because heat removal may not allow for such systems to be scaled down sufficiently, the inverse delay S may be even smaller than B . Thus the bisection–inverse-delay product HS may be even more inferior than the bisection–bandwidth product HB . The use of normal conductors is inhibitive for applications for which these products are suitable figures of merit.

The behavior of repeated systems is similar to that of optical and superconducting systems, when heat removal is not considered. However, the situation is worse when heat removal is considered. We refer to Eq. (13), which accounts for the effects of heat removal. For large N and B , the second term in this equation will be applicable so that $d \propto \bar{r}^{3/4}B^{3/4}N^{1/6}$, which is larger than $d \propto \bar{r}^{1/2}B^{1/2}$ dictated by interconnection density considerations. Thus the resulting growth rate of signal delay becomes $\tau \propto \bar{r}^{3/4}B^{3/4}N^{1/2}$, which is worse than $\tau \propto B^{1/2}N^{1/2}$, which we found in the optical case. For given B the growth rate of the bisection–inverse-delay product is then found to be

$$HS \propto N^{p/4}, \quad (19)$$

which is inferior to the optical $HS \propto N^{p-1/2}$ (since $p > 2/3$). If we do not terminate each stage of the repeaters individually and charge up the segments, as would most likely be the case in practice, then the first term in Eq. (13) becomes applicable so that $\tau \propto \bar{r}BN^{2/3}$. In this case we find that, for given B , the bisection–inverse-delay product is bounded from above and cannot be increased with increasing N , an inhibiting situation.

B. Quantitative Comparisons

There are many ways to present quantitative comparisons of S – B – N surfaces for different interconnection media. First, we will plot S as a function of N

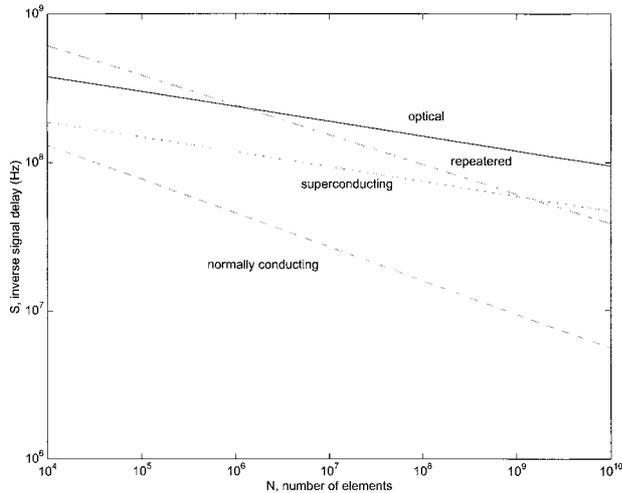


Fig. 1. Comparison of optical (solid curve), normally conducting (dashed curve), repeatered (dotted-dashed curve), and superconducting interconnections (dotted curve). We take $k = 5$, $p = 0.8$, $Q = 10 \text{ W/cm}^2$ and assume d_d , T_d , and T_r to be small enough to have no effect.

for the four media, with B set to the largest value allowed for by normally conducting interconnections for each value of N [as given by Eq. (5)]; For the three other media, B can be specified independently. However, since this is not the case for normal conductors, choosing B in this manner allows for a fair comparison of the values of S that correspond to the same values of N and B .] S , B , and N are related by Eq. (3) for optics and by Eq. (15) for superconductors. For repeaters the relation between S , B , and N is calculated as described by the comment following Eq. (13). For normal conductors S is determined as a function of N as described by the comment following Eq. (11). We see in Fig. 1 that the decrease of S with increasing N occurs more slowly with optical and superconducting interconnections so that they are superior for larger values of N . For lower values of N , repeatered interconnections may offer superior or comparable values of S .

On comparison of the coefficients of Eq. (12) and the second term of Eq. (3), we see that repeaters are comparable with optics in terms of interconnection density considerations, but the picture changes when heat-removal considerations are accounted for.

We also note that Eq. (15) is identical in form to Eq. (3) derived for optical interconnections. The numerical factors are also comparable, as is also evident from Fig. 1. We stress that, despite the similarity of the final relations, the physics involved is quite different. The scale of the optical system is fixed, whereas the scale of the superconducting system may be reduced, resulting in much smaller system size (this does not result in greater performance though, since the reduction in line lengths is precisely canceled by the inverse dependence of the delay on W .)

An alternative comparison is provided in Fig. 2, whose four panels each correspond to a different (constant) value of B , constituting sections of the S - B - N

space. The comments made for the previous figure hold in this case as well. Note, however, that the curve for normal conductors terminates at a certain value of N , reflecting the fact that information cannot be transmitted at the indicated value of B beyond that value of N , regardless of the value of S .

We have also examined the effect of reducing the nominal voltage value V from 1 V down to 0.1 and 0.01 V, which may be possible at lower temperatures (the improvements saturate near the lower level, and further improvements are not obtained for even lower voltages). Since the energy dissipation for the three conducting technologies will also be consequently reduced, this is expected to alter the comparison in favor of these technologies. From Fig. 3 we see that this is indeed the case, especially for higher values of B for which the dissipated power is higher. All three conducting technologies are able to offer much higher values of S ; however, the curve for normal conductors still terminates at a certain value of N : The fact that information cannot be transmitted at the indicated value of B beyond that value of N is not based on heat-removal considerations and so is not alleviated by a reduction of the voltage level. Reduction of the voltage level makes superconductors look especially attractive. It is important to note that in this figure we have assumed the reduction in voltage to have no effect on optical interconnections. However, it is likely that a reduction of temperature and voltage will also reduce the energy per transmitted bit for optics. If the reduction in energy is, for instance, 10, the curves for optics should simply be moved upward by $\sqrt{10}$ in Fig. 3.

Finally, we discuss the effect of varying p . Choosing larger values of p favors optics and superconductors, whereas choosing lower values of p has the opposite effect. The values of p for problems requiring global information flows, such as sorting, permutation and interconnection networks, discrete Fourier transforms, and global filtering, are usually close to unity. In other words, the bisections are proportional to N . The value of $p = 2/3$ represents complete locality in three dimensions. The value $p = 0.8$ used in our numerical examples thus represents an intermediate value.

C. Optimal Operating Point on the S - B - N Surface

In Fig. 1 we set B to its largest possible value for normal conductors, as determined by Eq. (5). Of course, it is not necessary to transmit bits along the lines at this maximum possible rate; one can transmit bits at lower rates as well (as is the case in Fig. 2). Referring back to Eq. (11), we see that if we choose B to be smaller, d will be smaller, resulting in larger values of S . Thus we see that there is a trade-off between S and B . If we set B to its largest possible value of $B = 1/T$, this will result in a particular value of S , as shown above. However, by choosing B to be smaller (i.e., by operating at a smaller duty ratio), we can reduce power dissipation, pack the elements more densely (i.e., reduce d), and thus decrease propagation delays, resulting in a larger value of S . (There is no purpose in reducing B beyond a certain extent, however, since once the scale of the system is reduced to

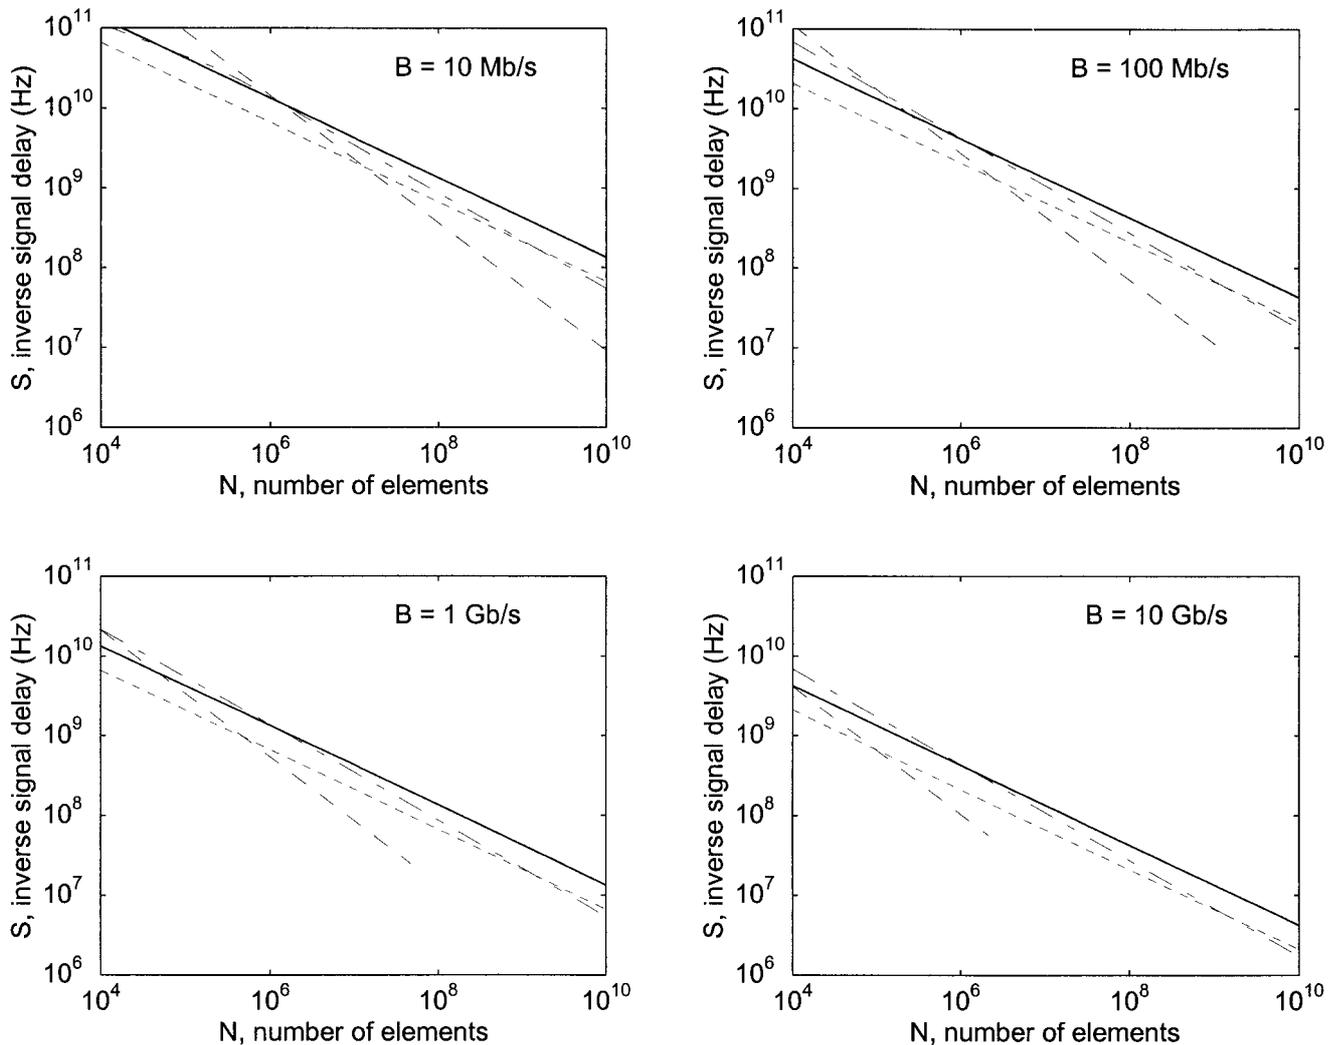


Fig. 2. Comparison of optical (solid curve), normally conducting (long-dashed curve), repeatered (long-short-dashed curve), and superconducting interconnections (short-dashed curve). We take $k = 5$, $p = 0.8$, $Q = 10 \text{ W/cm}^2$ and assume d_d , T_d , and T_r to be small enough to have no effect. (a) $B = 10 \text{ Mbit/s}$. (b) $B = 100 \text{ Mbit/s}$. (c) $B = 1 \text{ Gbit/s}$. (d) $B = 10 \text{ Gbit/s}$.

the extent that all lines become unterminated, further reduction in scale will not improve S . There is an upper limit to both quantities S and B regardless of the other; however, they can be traded off for each other over a certain range.) Given any optimization function involving B and S , we can find the optimum duty ratio and the associated values of S and B .

As a first example, let us assume that we would like to maximize $S = B$ (that is, maximize S and B under the constraint that they are equal). This optimization function might be appropriate for a synchronous system whose clock rate is set by the signal delay along the longest connection. The duty ratio will be denoted by $x \leq 1$ so that the bit repetition rate may be expressed as $B = x/T$. In Fig. 4(a) we plot the optimum duty ratio x that maximizes $S = B$. We observe that rather small duty ratios are optimal for a wide range of N .

A similar trade-off between S and B exists for the other interconnection media as well. Figure 4(b) shows the resulting comparison of the four media

when the duty ratio is chosen such that $S = B$ is maximized. Despite the fact that the asymptotic superiority of optical and superconducting interconnections remains and similar break-even values are observed, the performance offered by repeatered interconnections is much more comparable with that offered by optics and superconductors in this example. The major strength of optics and superconductors is their ability to provide large values of N and B simultaneously with minimal sacrifice in terms of S . Thus their superiority is less pronounced when S is emphasized as strongly as or more strongly than B .

We now consider a second example. In certain parallel computing contexts, it is the case that one desires to minimize the first-to-last bit communication latency τ_L of L bit messages, given by

$$\tau_L = \tau + \frac{L}{B} = \frac{1}{S} + \frac{L}{B}. \quad (20)$$

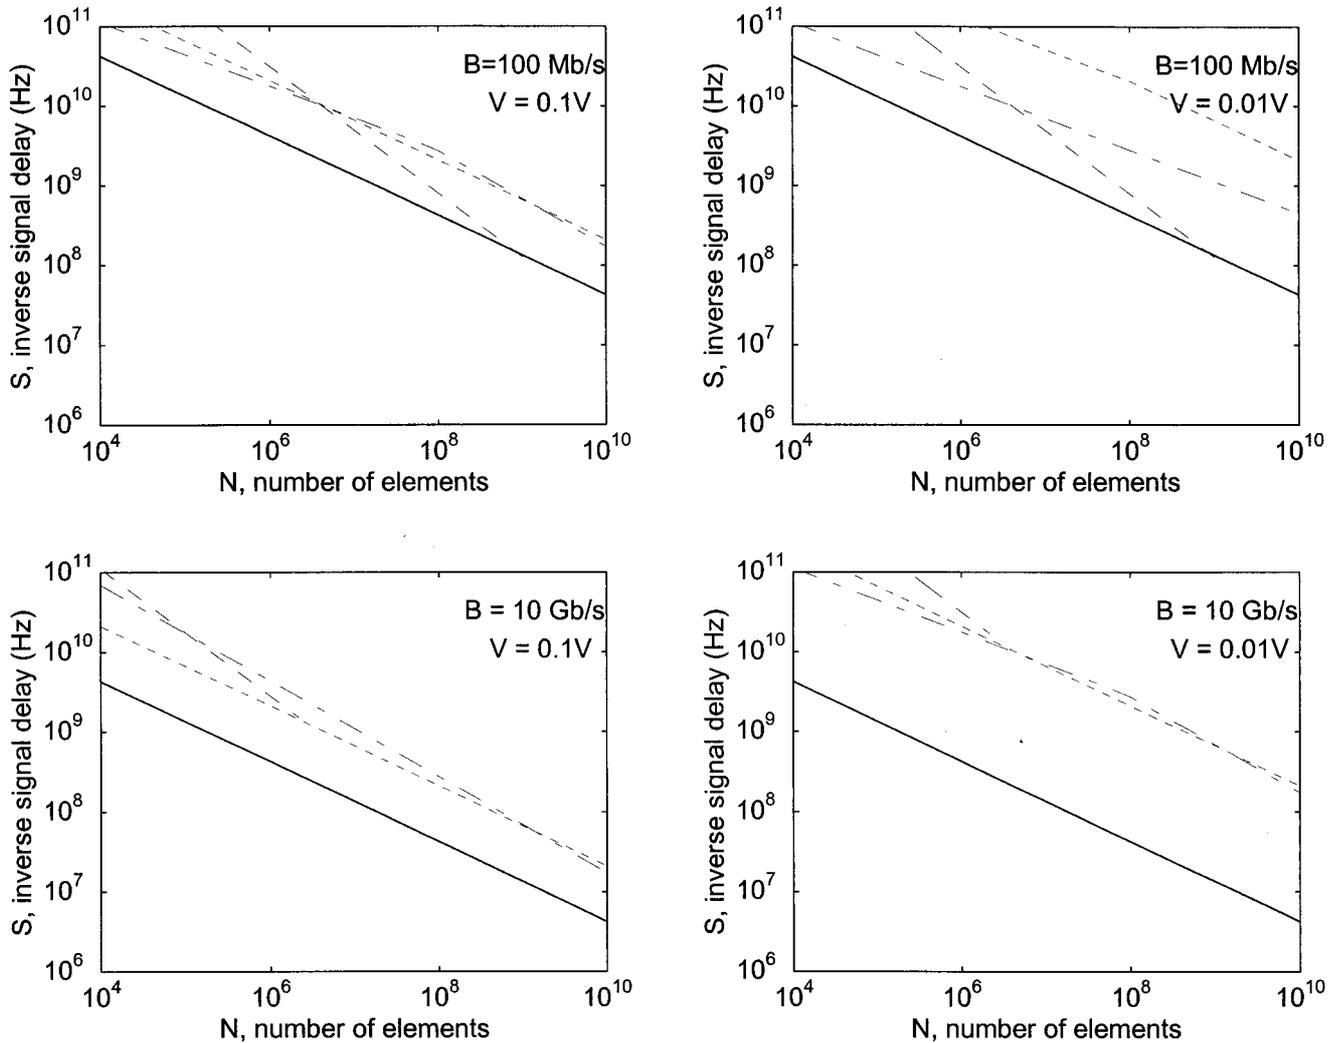


Fig. 3. Comparison of optical (solid curve), normally conducting (long-dashed curve), repeated (long-short-dashed curve), and superconducting interconnections (short-dashed curve). Similar assumptions are made as in the previous figures. (a) $B = 100$ Mbit/s, $V = 0.1$ V; (b) $B = 100$ Mbit/s, $V = 0.01$ V; (c) $B = 10$ Gbit/s, $V = 0.1$ V; (d) $B = 10$ Gbit/s, $V = 0.01$ V.

The trade-off relations we have derived between S and B allow us to find the optimum operating point resulting in the smallest value of τ_L and also to compare the resulting values of τ_L for the different media. The results are shown in Fig. 5. Although this example does not add significant new information with regard to the comparison of the technologies, it serves to illustrate the usefulness of our models and analysis in obtaining quantitative results.

An alternative discussion of related issues may be found in Refs. 18 and 19.

6. Conclusions

In this paper we have compared hypothetical fully three-dimensional normally conducting, repeated normally conducting, superconducting, and optical systems. We find that optical and superconducting interconnections are comparable with each other and superior to the others. Since optics seems to allow

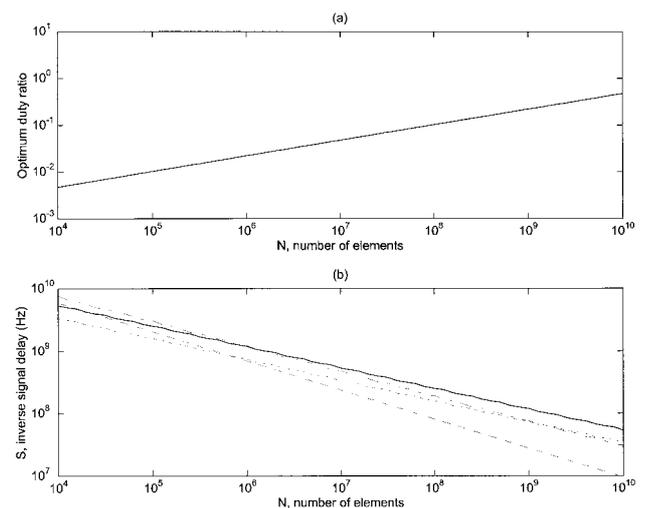


Fig. 4. Comparison of optical (solid curve), normally conducting (dashed curve), repeated (dotted-dashed curve), and superconducting interconnections (dotted curve) when $S = B$ is maximized. Similar assumptions are made as in the previous figures, but $T_d = 100$ ps. (a) Optimum duty ratio. (b) Resulting $S = B$ versus N .

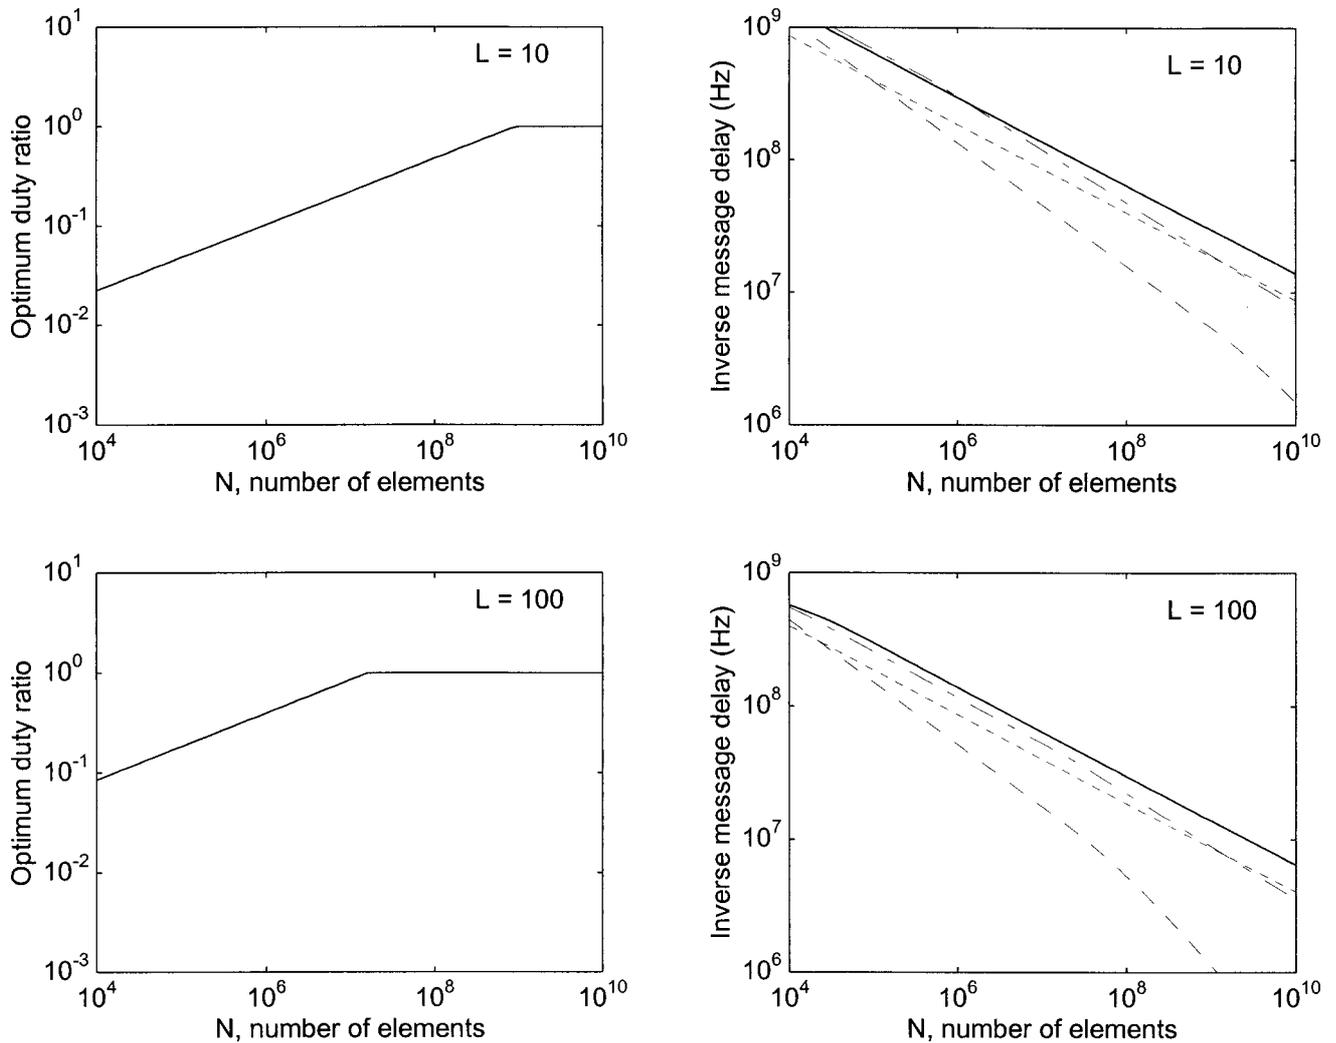


Fig. 5. Comparison of optical (solid curve), normally conducting (long-dashed curve), repeatered (long-short-dashed curve), and superconducting interconnections (short-dashed curve) when τ_L is minimized. Similar assumptions are made as in the previous figures, but $T_d = 100$ ps. (a) Optimum duty ratio for $L = 10$. (b) Resulting τ_L^{-1} versus N for $L = 10$. (c) Optimum duty ratio for $L = 100$. (d) Resulting τ_L^{-1} versus N for $L = 100$.

us to approach full three dimensionality more closely, it represents the most superior option.

It is important to understand that none of the interconnection media we have considered enables continual reduction of signal delay by downscaling. (Optical interconnections cannot be downscaled, and all kinds of conducting interconnections exhibit an inverse dependence of delay on linewidth, below a certain linewidth.)

In discussing the limitations of conducting interconnections, we allowed for arbitrarily small scaling and arbitrarily fast devices. We saw that normal conductors, whether terminated or not, did not allow for B to be kept constant with increasing system size (which is possible with the other media). Both B and S were found to decrease sharply with increasing N . The bisection-inverse-delay and bisection-bandwidth products were found to be bounded from above. This is in contrast to the other media with

which it is possible to arbitrarily increase B and the bisection-bandwidth product for any given N .

Repeaters are inferior to optics and superconductors, since they result in faster growth of signal delay and slower growth of the bisection-inverse-delay product with increasing N .

Optical and superconducting interconnections lead to similar performance for similar communication energies. Although superconducting layouts may be much smaller than optical layouts, they do not result in smaller delay because of the inverse dependence of delay on linewidth, once conductor thickness drops below the penetration depth. Optical interconnections may allow us to more closely approach full three dimensionality and offer freedom from termination problems. Superconductors may offer much lower energies, especially if the voltage level is reduced. This in turn might enable reduction of signal delay.

Appendix A: Definitions of Symbols Used

S	inverse signal delay = $1/\tau$,
B	bit repetition rate (bandwidth),
N	number of devices or elements,
H	bisection of the system = $k\kappa N^p$,
HB	bisection–bandwidth product,
HS	bisection–inverse-delay product,
k	average number of connections per element,
p	system connectivity measure (Rent exponent),
κ	$= (p - 2/3)^{-1} (p + 1/3)^{-1}$,
d_d	linear size of each device or element,
d	grid (lattice) spacing of the layout,
\mathcal{L}	linear extent of the system = $N^{1/3}d$,
r	length of an interconnection in grid units (dimensionless),
\bar{r}	average connection length in grid units = $\kappa N^{p-2/3}$,
χ	number of physical interconnections per logical connection,
\mathcal{P}	total power dissipated by the system,
Q	maximum amount of power that can be removed per unit area,
l	length of an interconnection in physical units = rd ,
W	transverse linear extent of an interconnection,
A	cross-sectional area of an interconnection = W^2 ,
τ	signal delay (latency) = $\max(T_p, T)$,
T_p	propagation delay along an interconnection,
T	minimum temporal pulse width = $\max(T_i, T_d)$,
T_l	line-imposed component of T ,
T_d	device-imposed component of T ,
T_r	minimum pulse repetition interval, usually = T ,
E	energy dissipated per transmitted bit,
W_{\min}	minimum manufacturable value of W ,
λ	optical wavelength,
c	speed of light in free space,
f	effective f -number of optical interconnection system,
ρ	resistivity of conductor,
ϵ	permittivity of dielectric,
μ	permeability of dielectric,
v	propagation velocity in dielectric,
V	nominal voltage level,
R_0C_0	intrinsic delay of a repeater,
λ_p	superconducting penetration depth,
J_{sc}	superconducting critical current density.

We acknowledge the benefit of discussions with Volkan H. Özgüz of Irvine Sensors Corporation, Costa Mesa, California, and Sadık C. Esener of the University of California, San Diego, California.

References

1. IEEE, eds., *Proceedings of the 45th Electronic Components and Technology Conference (ECTC)* (Institute of Electrical and Electronics Engineers, Piscataway, N.J., 1995).

2. H. M. Ozaktas and J. W. Goodman, "The limitations of interconnections in providing communication between an array of points," in *Frontiers of Computing Systems Research*, S. K. Tewksbury, ed. (Plenum, New York, 1991), Vol. 2, pp. 61–130.
3. H. M. Ozaktas, "A physical approach to communication limits in computation," Ph.D. dissertation (Stanford University, Stanford, Calif., 1991).
4. A. L. Rosenberg, "Three-dimensional VLSI: a case study," *J. Assoc. Comput. Mach.* **30**, 397–416 (1983).
5. F. T. Leighton and A. L. Rosenberg, "Three-dimensional circuit layouts," *J. Comput. Sys. Sci.* **15**, 793–813 (1986).
6. M. J. Little and J. Grinberg, "The 3-D computer: an integrated stack of WSI wafers," in *Wafer-Scale Integration* (Kluwer, New York, 1988), Chap. 8.
7. H. M. Ozaktas, Y. Amitai, and J. W. Goodman, "A three dimensional optical interconnection architecture with minimal growth rate of system size," *Opt. Commun.* **85**, 1–4 (1991); errata **88**, 569 (1992).
8. 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).
9. 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).
10. H. M. Ozaktas and D. Mendlovic, "Multistage optical interconnection architectures with least possible growth of system size," *Opt. Lett.* **18**, 296–298 (1993).
11. K. W. Goossen, J. E. Cunningham, and W. Y. Jan, "GaAs 850 modulators solder-bonded to silicon," *IEEE Photonics Technol. Lett.* **5**, 776–778 (1993).
12. K. W. Goossen, J. A. Walker, L. A. D'Asaro, S. P. Hui, B. Tseng, R. Leibenguth, D. Kossives, D. D. Bacon, D. Dahringer, L. M. F. Chirovsky, A. L. Lentine, and D. A. B. Miller, "GaAs MQW modulators integrated with silicon CMOS," *IEEE Photonics Technol. Lett.* **7**, 360–362 (1995).
13. H. M. Ozaktas, "Paradigms of connectivity for computer circuits and networks," *Opt. Eng.* **31**, 1563–1567 (1992).
14. H. M. Ozaktas, 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).
15. H. M. Ozaktas, "Toward an optimal foundation architecture for optoelectronic computing. Part I. Regularly interconnected device planes," *Appl. Opt.* **36**, 5682–5696 (1997).
16. H. M. Ozaktas, "Toward an optimal foundation architecture for optoelectronic computing, Part II. Physical construction and application platforms," *Appl. Opt.* **36**, 5697–5705 (1997).
17. H. B. Bakoglu, *Circuits, Interconnections, and Packaging for VLSI* (Addison-Wesley, Reading, Mass., 1990).
18. W. Nakayama, "On the accomodation of coolant flow paths in high density packaging," *IEEE Trans. Component Hybrids Manuf. Technol.* **13**, 1040–1049 (1990).
19. W. Nakayama, "Heat-transfer engineering in systems integration—outlook for closer coupling of thermal and electrical designs of computers," *IEEE Trans. Components Packag. Manuf. Technol. Part A* **18**, 818–826 (1995).