Frontiers of Computing Systems Research

Volume 2

Essays on Emerging Technologies, Architectures, and Theories

Contents

1 Limits of Information Processing Systems
R.W. Keyes 1
1.1 Introduction . . . 1
1.2 Physical Nature of Systems . . . 1
1.3 Devices . . . 9
1.4 Information Storage . . . 27
1.5 Communication . . . 29
1.6 Heat Removal . . . 49
1.7 Conclusion . . . 54
References . . . 56

2 The Limitations of Interconnections in Providing Communication between an Array of Points
H.M. Ozaktas and J.W. Goodman 61
2.1 Introduction . . . 63
2.2 Some Definitions . . . 65
2.3 System Model . . . 68
2.4 Optical Interconnections . . . 73
2.5 Normally Conducting Interconnections . . . 83
2.6 Repeatered Normally Conducting Interconnections . . . 100
2.7 Superconducting Interconnections . . . 107
2.8 Discussion . . . 114
2.9 Towards Unifying Physical and Algorithmic Approaches . . . 115
2.10 Summary and Conclusions . . . 118
A Appendix . . . 120
References . . . 124

3 Optics in Multiple-Instruction, Multiple-Data Stream Computers
E.E.E. Frietman 131
3.1 Introduction . . . 131
3.2 Interconnection Networks for Massive Parallel Computers . 139
3.3 An Electro-Optic Communication System . . . 141
3.4 A Reconfigurable Optical Interconnection Protocol . . . 149
3.5 Fabrication Processes Suitable to Realize an OELE . . . 157
3.6 A New Class of Opto-Electronic Logic Elements . . . 172
3.7 Future Research in the 'DPP9X' Computer Project . . . 186

A Appendix 189
References 190

4 Clouds, Computers and Complexity **197**
P. Christie
4.1 Interconnection Complexity 197
4.2 Clouds and Computers 198
4.3 Fractals and Entropy 213
4.4 Discussion 226
4.5 Conclusion 235
References 235

5 Experimental Studies of Analog Neural Networks **239**
R.C. Frye and E.A. Reitman
5.1 Introduction 239
5.2 Adaptive Networks and Analog Electronic Hardware 243
5.3 Optically Controlled Adaptive Network 244
5.4 Network Performance 252
5.5 Signal Prediction 256
5.6 Simulation Results 261
5.7 Scaling Considerations 266
5.8 An Application of an Analog Neural Network 270
5.9 Future Prospects for Analog Neural Networks 286
References 287

6 Complex Behavior in Networks with Distributed Routing **291**
A. Rucinski, P. Drexel, and B. Dziurla
6.1 Introduction 291
6.2 Physical Model of Distributed Networks 297
6.3 Dynamic Behavior 303
6.4 Phase Transition in Mesh Networks 306
6.5 Conclusions 310
References 313

7 New Approaches in System-Level Diagnosis **317**
A.K. Somani and O. Peleg
7.1 Fault Diagnosis in Multi-Processor Systems 318
7.2 System-Level Diagnosis 323
7.3 Classification of Diagnosable Systems 330
7.4 Diagnosability Algorithms 349
7.5 Diagnosis Algorithms 353
7.6 Summary and Conclusions 363
References 364

Index **369**

CHAPTER 2

THE LIMITATIONS OF INTERCONNECTIONS IN PROVIDING COMMUNICATION BETWEEN AN ARRAY OF POINTS

Haldun M. Ozaktas
Joseph W. Goodman
Stanford University
Information Systems Laboratory, Durand Building
Department of Electrical Engineering
Stanford, California 94305 USA

Abstract

We present a comparative analysis of optical, normally conducting, repeatered and superconducting interconnection performance in a very large scale digital computing environment. We derive tradeoff relations between delay, bandwidth and system size for each technology based on communication (wiring) volume and heat removal considerations and discuss their numerical and asymptotic properties. We show that the bisection-bandwidth and bisection-inverse delay products—which are appropriate measures of performance for communication limited applications—are bounded from above for normally conducting layouts, whereas they may be arbitrarily increased for repeatered, optical and superconducting layouts. The latter two are shown to suffer slower growth rate of signal delay with increasing system size in 3 dimensions than repeatered interconnections and thus offer the best performance. Based on the considerations of this paper, the comparison between optical and superconducting interconnections for same dimensional layouts reduces to a comparison of their respective communication energies.

Frontiers of Computing Systems Research, Volume 2
Edited by S.K. Tewksbury, Plenum Press, New York, 1991

List of Symbols

a	system radius in grid units
$(area)$	cross sectional area associated with each physical line
B	bit repetition rate along each edge of connection graph
c	vacuum velocity of light
C	capacitance per unit length
d	linear extent of a unit cell
d_d	linear extent of an element
d_{trans}	linear extent of a transducer
e	Euclidean dimension of layout space
E	energy associated with each transmitted bit of information
$f(\cdot)$	functional form of connection flux distribution
$g(\cdot)$	functional form of line length distribution
h	height of dielectric
H	bisection
J_c	volume critical current density
J_{sc}	surface critical current density
k	number of graph edges per element
K	number of wiring tracks per cell
ℓ	length of a line in real units
L	inductance per unit length
m	order of moment of line length distribution
M	number of wiring layers
n	fractal dimension of layout
N	number of elements
p	interconnectivity (Rent exponent) of layout
Q	maximum amount of power we can remove per cross section
r	length of a line in grid units
$\bar{r}$	average connection length in grid units
$< r^m >$	mth moment of line length distribution
R	resistance per unit length
R_d	drive impedance
R_0C_0	intrinsic delay of repeating devices
S	inverse of worst case signal delay
S_{ave}	inverse of average signal delay
t	height of conductor
T	minimum temporal pulse width associated with each transmitted bit of information

T_d	device imposed component of T
T_ℓ	line imposed component of T
T_p	propagation delay along a line
T_r	minimum pulse repetition interval along a line
V	nominal logic voltage level
w	width of conductor
$(width)$	transverse linear extent associated with each physical line
Z_0	characteristic impedance
α	attenuation constant
χ	number of parallel physical lines used to establish each graph edge
δ	classical skin depth
ϵ	permittivity of dielectric
κ	coefficient for average connection length
λ	optical wavelength
λ	superconducting penetration depth
μ	permeability of dielectric
ρ	resistivity of conductor
τ	worst case signal delay
τ_{ave}	average signal delay
υ	velocity of propagation
ω	fundamental frequency component
ξ	optimal number of repeater stages
$\zeta_m, \zeta'_m, \zeta''_m$	coefficients for the moments of line length distribution

2.1 Introduction

Interconnections are more and more becoming the factor limiting the performance of large scale digital computing systems [1] [2] [3] [4] [5] [6] [7] [8] [9]. Optical and superconducting interconnections are major candidates for alleviating this trend. Many authors have made comparative studies of optical and normally conducting interconnections [10] [11] [12] [13] [14] and superconducting and normally conducting interconnections [15] [16] [17]. In this work we combine physical models of interconnection media with an interconnect dominated, device independent system model which has its roots in VLSI complexity theory [18] [19] and empirical wiring models [20]. In this manner we hope to understand how the basic physical properties

of interconnection media affect system performance (inverse delay, bandwidth, bisection) and cost (energy, space) parameters. Our viewpoint in this work is more that of a physicist, rather than of an engineer. The main feature of our approach is its ability to simultaneously handle several of the interrelated physical and system aspects of the problem which are often isolated and treated separately.

Optical and superconducting technologies have not yet achieved full maturity, but are rapidly advancing. Some room for improvement for normal conductors also exists. In this study we look into the future and attempt to analyze and compare the basic limitations imposed by these technologies after all technical and practical difficulties associated with their implementation have been overcome. Determining the ultimate performance afforded by each technology is a formidable task, especially if one insists on numerical accuracy. Since our aim is to develop a general qualitative understanding, rather than to suggest practical design schemes, we will be crude in our handling of numerical factors. Nevertheless, we believe that our interconnection models represent, within a factor of the order of unity, the best achievable by each technology. Of course, we are to a certain extent conditioned by the current trends in these technologies and the way present computing systems are built, so we cannot exclude the possibility of breakthroughs or ingenuity not foreseen by us.

Some of our approximations are made with the interest of maintaining generality (i.e. we are reluctant to introduce system specific parameters) and others with the purpose of maintaining analytic simplicity and transparency. As an example, consider the equation $y = ay^{1/2} + b$ where all quantities are positive. The exact solution is $y = (a + \sqrt{a^2 + 4b}\,)^2/4$. An approximate solution may be written as $y \simeq a^2 + b$ by inspection and differs from the exact solution by at most a factor of 4/3. Moreover, when either a^2 or b is large compared to the other, the approximate solution will be nearly exact. Likewise, we will often use $\max(x, y)$ and $x + y$ interchangeably, where x and y are positive quantities. The form $y = \sqrt{x^2 + 1}$ will be approximated by $y = x$ for $x \geq 1$ and $y = 1$ for $0 \leq x \leq 1$. Needless to say, care must be exercised in employing such approximations, as raising such expressions to high powers or using them in the argument of an exponential function can lead to drastic errors.

An alternative approach would be to ignore constants and bounded variations altogether, as is common practice in VLSI complexity theory. We have preferred not to obscure the physical nature of the problem and, of course, order of magnitude information is better than none. Those inter-

ested in greater accuracy should be able to improve our results by incorporating the parameters or additional factors characterizing their particular application.

We start with some definitions, followed by a description of our system model. Sections 2.4, 2.5, 2.6 and 2.7 treat optical, normally conducting, repeatered and superconducting interconnections respectively. Descriptions of the physical models used are followed by analyses of the effects of information density, scaling and heat removal, leading to our major results. Section 2.8 provides some additional considerations and reservations. Section 2.9 introduces an abstract framework for relating the computational requirements of a problem to the physical limitations imposed by interconnect media. The final section summarizes important conclusions.

2.2 Some Definitions

For the purpose of this paper, a processing system is a collection of N given similar *primitive elements* connected to each other according to a prespecified *graph* [21]. The primitive elements may be simple switching devices or relatively complex subsystems. k will denote the number of connections (graph edges) per element[1], so that 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. We will assume that the number of input-output ports of each element does not vary greatly from element to element, so that each element has $\sim k$ ports. We will treat k as a given constant, although many of our results are easily extended to the case where k is a function of N. d_d will denote the linear extent of the elements (also referred to as *devices*). Of course, the elements should be at least large enough to accomodate their input-output ports (transducers in the case of optical interconnections).

One way to increase the processing power of the system is to increase the number of elements N. This may enable the system to handle larger problem sizes in a given amount of time, or given problems in a shorter amount of time (because of the increase in parallelism), or other intermediate combinations.

Another way to increase the processing power of a machine is to increase the rate at which information percolates among the elements of the system. The solution of a problem will in general require a certain number of time

[1] For simplicity we are considering pairwise connections only, the extension to fan-out and fan-in is discussed in appendix A.1.

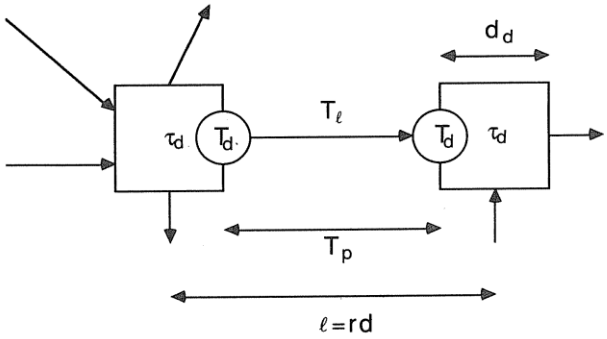

Figure 2.1: Factors determining the speed of the system. Two elements (graph nodes) sharing an (inter)connection (graph edge) of length ℓ are shown. τ_d is the time which elapses between arrival of new input to an element and updating of the output values. T is the minimum temporal pulse width. T_p is the propagation delay along the interconnection.

steps. The physical duration of a time step (measured in seconds) is set by one of several mechanisms, as illustrated in Figure 2.1. Information transfer takes place along the edges of the connection graph in the form of binary pulses of minimum temporal width $T = \max(T_d, T_\ell)$, as set by the greater of device (drive) or line imposed minimum pulse widths, T_d or T_ℓ, respectively. We assume that a pulse must be completely received for the value of a transmitted bit of information to be properly registered. T_p denotes the propagation delay along the interconnection. The meaning of these and subsequent quantities will become clearer when we specify them for specific technologies in later sections. τ_d denotes the time which elapses between arrival of new input at the elements and the updating of their output values accordingly. The largest of these quantities will determine the rate at which computational processes involving the cooperation of elements situated at a distance ℓ from each other will proceed. Let us denote this rate as $S = 1/\tau$ where τ, the signal delay, is defined by

$$\tau = \max(\tau_d, T, T_p) \tag{2.1}$$

and is a non-decreasing function of ℓ. In a synchronous system, the physical duration of a time step is determined by the worst case delay among all connections [3]. When we are speaking of an isolated connection, the quantities T, T_ℓ, T_p, τ and S defined in this paragraph, and the quantities T_r, χ and B which will be defined in subsequent paragraphs will refer to the properties of that particular connection. When we are talking about a system, these quantities will refer to the worst case over all connections. T_d, τ_d and T_{rd} (to be defined) will be assumed to be constants.

In certain cases, the worst case S may be a pessimistic measure. In general, each element will want to communicate with a certain set of other elements at different distances. Let τ_{ave} be defined as the average of τ over all connections. $S_{ave} = 1/\tau_{ave}$ is the inverse of the average delay over all edges of the connection graph and can be thought to be a measure of the speed (in nodes traversed per second) at which information flows through *paths* [21] of the connection graph. Whether S or S_{ave} is the relevant quantity will depend on how we operate our system. In this work we will limit our attention to S so as not to further lengthen our treatment. All of the analysis presented may be easily modified for S_{ave}. Most major qualitative conclusions will remain unchanged.

Another measure of speed is the rate B (in bits/sec) at which information is piped through the edges of the connection graph. We assume that this rate is kept constant for all lines and thus is determined by the worst case value of B over all lines. Let T_r denote the minimum pulse repetition interval, i.e. bits may be emitted into *each physical line* at a rate of one every T_r seconds. In most cases, T_r will approximately equal T, the minimum pulse width. If we desire to increase B beyond $1/T_r$, we may employ $\chi > 1$ parallel physical lines to establish each edge of the connection graph. Let $1/T_{rd}$ denote the maximum rate at which the elements can emit information into each edge of the connection graph when $\chi \to \infty$. Thus B may never exceed $1/\max(T_r/\chi, T_{rd})$.

The use of $\chi > 1$ physical channels per graph edge will require an increase in the number of physical input-output ports by a factor of χ. This may in turn dictate an increase in element size d_d. If the cross section of each port is not greater than the cross section of each physical channel, this increase in d_d will always be overshadowed by the increase in necessary communication (wiring) space, and thus need not be explicitly kept track of.[2] We will mostly assume this to be the case. In practice, however, input-output ports may be much larger than the cross section of the physical channels so that we must explicitly set the element size to be large enough to accomodate $\sim \chi k$ ports.

Notice that the effects of τ_d and T_{rd} are to simply hard limit S and B to $1/\tau_d$ and $1/T_{rd}$ respectively. In this work we are interested in the limits imposed by the interconnections, rather than the elements. Thus without further mention, we will assume T_{rd} to be negligibly small and that τ_d is no greater than T_d.

[2] The use of wavelength division multiplexing constitutes an exception and must be treated separately.

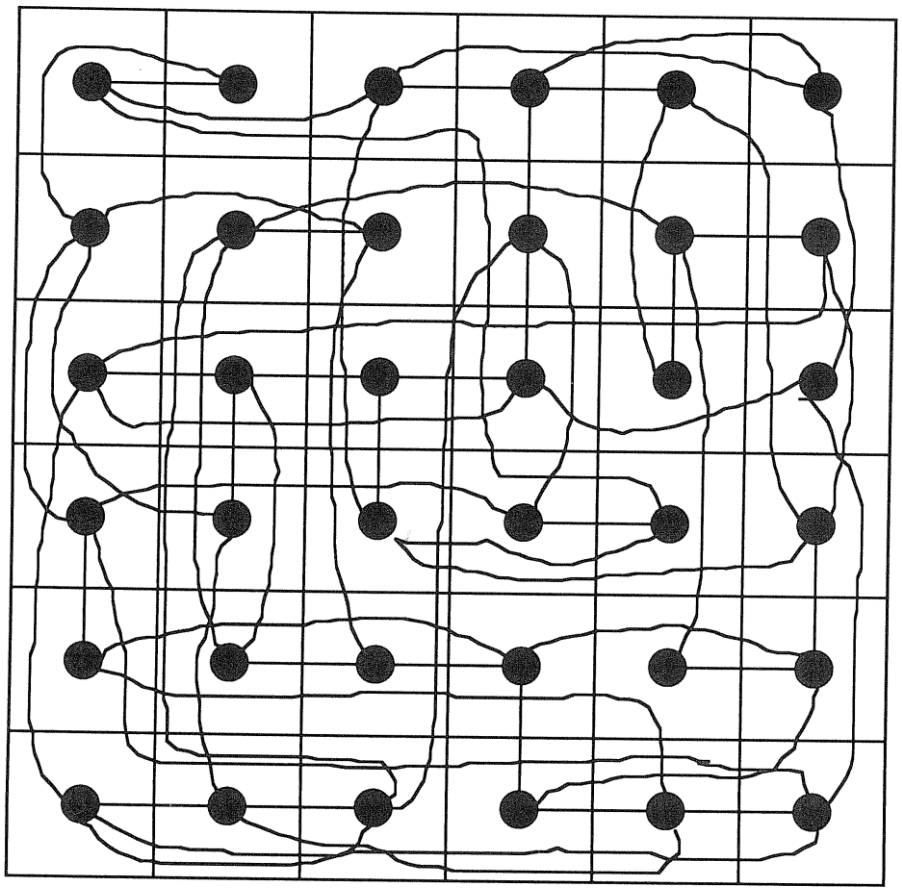

Figure 2.2: Embedding the connection graph. For convenience we lay out the elements in a cartesian array of cells. The size of a cell is to be determined according to the size of the elements, the space that must be provided for the wires and heat removal requirements.

Although it would certainly be desirable, it is not possible to arbitrarily increase S, B and N simultaneously due to physical limitations. In this paper we quantify this by deriving bounds of the form $\Phi(S, B, N) \leq C_\Phi$ for different interconnect technologies.

2.3 System Model

Let the $N \gg 1$ elements comprising our system be laid out on an e dimensional regular cartesian grid of as yet unspecified lattice constant d with $N^{1/e}$ elements along each dimension (Figure 2.2). In this work we do not attempt an interpretation of fractional values of e so that $e = 2$ or $e = 3$. Figure 2.3 shows a hierarchical partitioning of such a layout. We will quantify the communication requirements of connection graphs to first order by the Rent exponent $0 \leq p \leq 1$, which we also term as

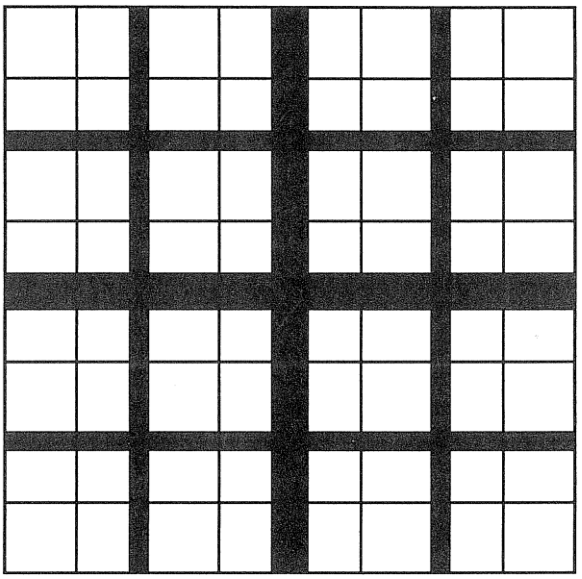

Figure 2.3: Binary hierarchical partitioning of the array of cells. After [22].

the *interconnectivity*. Referring to Figure 2.3, Rent's rule states that the number of graph edges P emanating from any subgroup containing $N' < N$ elements is given by $P = kN'^p$. Statistical variations from this formula are to be expected. The rule breaks down when N' is close to N, the total number of elements in the system. Rent's rule was originally established as an empirical relationship [23] [24] and later shown to be a consequence of the logic design process [25] [26]. Such a power law may also be justified based on a principal of *self similarity* [22] [27]. We now understand that Rent's rule is also related to the *separator* concept of VLSI complexity theory [18] [19], which provides a formal basis for the layout of given graphs, and to the theory of fractals [28]. This relationship has been used widely [20] for two decades as a wiring model.

Donath [30] and Feuer [31] showed that Rent's rule led to a *connection flux distribution* which is essentially of the form

$$f(r) = kr^{e(p-1)}(1 - r^e/a^e) \qquad 1 \leq r \leq a \tag{2.2}$$

where r denotes distances in units of grid spacing so that physical distances are given by $\ell = rd$. $a \gg 1$ denotes the system radius. $f(r_0)$ gives the *expected* number of connections originating from a certain element and emanating from a spherical surface of radius r_0 centered at that element (Figure 2.4). Note that $f(1) = k$ and $f(a) = 0$. The smaller p is, the quicker $f(r)$ decreases. Thus systems with small p are those involving local

communication, whereas those with large p involve global communication. The factor $(1 - r^e/a^e)$ has been introduced to account for the finite extent of the system and may be ignored either when r is not close to a or when p is not close to 1. The *line length distribution* is defined as

$$g(r) = -\frac{df(r)}{dr} \qquad 1 \leq r \leq a. \tag{2.3}$$

$g(r_0)\Delta r$ gives the *expected* number of connections originating from an element and terminating in the interval $[r_0, r_0 + \Delta r]$. $k^{-1}g(r)$ is a probability distribution defined over $[1, a]$. When p is small, it is more likely for a connection to be made to close by elements, rather than distant elements. When $p = 1$, we have $g(r) \propto r^{e-1}$, so that it is equally likely for connections to be made to elements at any distance (notice that there are $\propto r^{e-1}$ elements at distance r). This is consistent with the usual interpretation of $p = 1$.

Although $f(r)$ and $g(r)$ have been defined somewhat artificially for an element located at the center of a circular (or spherical) layout of radius a, ignoring edge effects and the distinction between cartesian and euclidean distances, we will assume $Ng(r)$ to be a good approximation to the distribution of line lengths in our system (Figure 2.2). That is, $Ng(r_0)\Delta r$ gives the expected number of connections in our system with lengths lying in the interval $[r_0, r_0 + \Delta r]$. Of course, we have $\int Ng(r)\, dr = kN$. Likewise, we will take a to be $\simeq N^{1/e}$ without concerning ourselves with precise geometrical factors.

The *fractal dimension* of the layout is given by the relationship $n = 1/(1 - p)$ [32] [33]. Although a justification of the use of the term 'dimension' is beyond the scope of this paper, we will use n interchangeably with p as a measure of the communication requirements of a system. The interested reader is referred to the work of Mandelbrot for a discussion of the relationship between inverse power law distributions and fractal forms [34] [35] [36].

We stress that although the relationships between layout theory, Rent's rule, fractal geometry and inverse power law distributions are compelling, they are not necessary for the purpose of this paper if the reader is willing to accept the inverse power law distribution of line lengths with parameter p as a starting point. Some authors have simply assumed similar distributions [37] without any underlying theory. The use of such a line length distribution may also be justified empirically [2] [38]. For an alternate approach more in the tradition of VLSI complexity theory the reader is referred to [12].

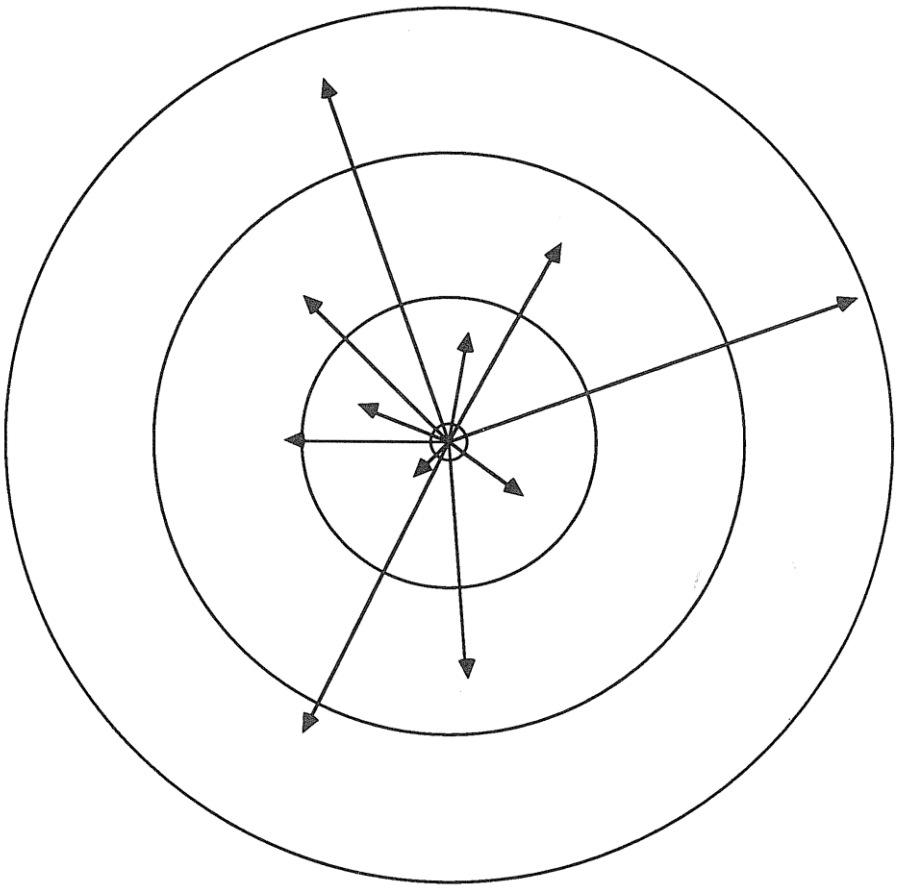

Figure 2.4: Connections made by an element to other elements.

A useful approximation for the mth moment $< r^m >= k^{-1} \int r^m g(r)\, dr$ may be derived as

$$\begin{aligned} e < mn \quad & < r^m >= \zeta_m N^{m/e-1/n} \\ e = mn \quad & < r^m >= \zeta'_m \ln N \\ e > mn \quad & < r^m >= \zeta''_m \end{aligned} \tag{2.4}$$

where the coefficients are functions of m, $n = 1/(1-p)$ and e. The first and second moments were already derived by Donath [22] [30]. A discussion of the approximations leading to these equations and expressions for the coefficients are given in appendix A.2. Because it is most often used, we will use the special symbols $\bar{r} =< r >$ and $\kappa = \zeta_1$ for the first moment.

Notice that r is actually a discrete quantity. It is possible to find graphs for which our continous approximation leads to erroneous results. For instance, according to (2.3), a simple planar mesh for which $p = 1/2$ laid out

in $e = 2$ dimensions has $g(r) \sim r^{-2}$. Though a quickly decreasing function, this is a very crude representation of the actual line length distribution, which is concentrated at $r = 1$, and would result in overestimates of higher order moments. Graphs exhibiting a high degree of regularity for which exact values of the moments may be calculated by combinatoric methods are best handled *per se*. The reader will notice that most of our results may be cast in a form that depends only on the first two moments of $g(r)$, without requiring a full specification of $g(r)$. The inverse power law distribution we are using is an attempt to describe the irregular nature of typical digital circuits we are likely to encounter in practice, as confirmed by earlier authors [22] [30] [31].

A significant quantity is what has been historically termed the number of *wiring tracks per cell*, which we denote by K. If $\chi \geq 1$ physical lines are being used to establish each edge of the connection graph, the cross section (or width) of each cell d^{e-1} must be wide enough to allow the passage of χK physical lines, in addition to accomodating the element itself [39]. A moment's reflection reveals that K is given by $K = k\bar{r}$ [40] [41], since $k\bar{r}$ is the total connection length per cell in grid units. Letting $(width)$ denote the linear extent of a single physical line, including its share of line to line spacings, the above condition may be expressed as $d^{e-1} \geq \chi K(width)^{e-1}$. Combining this with the condition $d \geq d_d$, we will write $d \geq \max(d_d, (k\chi\bar{r})^{1/(e-1)}(width))$. Notice that the error we incur in pretending that the interconnections and elements may co-occupy the same physical space is less than a factor of 2.

We will not consider statistical variations from cell to cell. We will ignore the fact that there will be a greater demand for wiring space towards the center, and also assume 100% utilization of the available wiring space. Of course, in practice, a less than unity efficiency factor will be involved. Typically, an approximately equal number of tracks will be running in each of the e orthogonal dimensions. We will not be concerned with this distinction and associated numerical factors. M will denote the number of wiring layers for 2 dimensional layouts.

The *bisection* H is defined as the number of graph edges crossing an imaginary surface dividing the system in two roughly equal parts[3] and is given by $H = N^{(e-1)/e}K$, since there are $N^{(e-1)/e}$ cells adjacent to this surface. When $n > e$, $\bar{r} = \kappa N^{1/e-1/n} = \kappa N^{p-(e-1)/e}$ so that $H = k\kappa N^p$.

[3] This surface is also referred to as the "bisection".

This must be multiplied by χ to obtain the number of physical lines crossing the surface in question.

The HS (bisection-inverse delay) and HB (bisection-bandwidth) products are appropriate figure of merit functions for communication limited applications. The communication complexity of many problems may be stated in terms of the amount of information that must pass through the bisection of the system [18] [42] [43], so that these products are direct measures of system performance.

More general figure of merit functions may be defined by incorporating the cost of energy or space. We will not go into the analysis of such cost based figure of merit functions in this paper.

In this work we will concentrate on highly interconnected systems, characterized by large values of p (or equivalently n). The method of analysis is easily extended to other cases.

2.4 Optical Interconnections

2.4.1 Physical Model Description

Conceptually, the simplest structure one might use to transmit optical signals is a single mode waveguide. If a sufficiently high numerical aperture is utilized (through use of a sufficiently high refractive index difference between the core and cladding), guide widths of the order of a wavelength are possible [44]. Use of a sufficiently high refractive index difference ensures that the evanescent fields in the cladding will decay within a short distance. In general, the crosstalk between adjacent guides is proportional to the product of the coupling constant and the length of the guides. This would mean that for increasing systems sizes, one would have to increase the separation between the guides in order to maintain an acceptable crosstalk level. However, the coupling constant is an exponentially decaying function of the guide separation [45]. This means that the required guide seperation is a slowly varying function of system size. For this reason, we will take the necessary guide separation to be constant and also of the order of a wavelength[4]. Decreasing the separation will increase crosstalk excessively with little gain in density. Increasing the separation somewhat beyond a wavelength may be desirable, but not by a factor much greater than unity.

[4] It is also possible to envision a design methodology for which the crosstalk does not increase with system size. This might be established through the use of design rules which exploit the periodic nature of coupling with distance and set the lengths of parallel runs accordingly.

In 2 dimensions the relevant quantity is the $(width)$ allocated to each line, whereas in 3 dimensions it is the cross sectional $(area)$ allocated to each line. For optical lines we will write

$$(width) = 2\lambda \tag{2.5}$$

$$(area) = (width)^2 = 4\lambda^2 \tag{2.6}$$

independent of all other parameters. Although it would be quite difficult to do any better than this in practice, theoretically there is still a little more room for improvement [46]. In this work we treat optical communication links as if they were solid wires of cross section $(2\lambda)^2$ (i.e. as is the case when waveguides are used—whether we allow them to intersect or not does not make a significant difference). The results thus obtained represent the limitations of all forms of optical communication (guided wave or free space) within a factor of the order of unity. This generalization is possible by virtue of a result that states that the minimum volume required for providing optical communication among an arbitrary array of points is $\sim \lambda^2 \ell_{total}$, where ℓ_{total} is the total interconnection length [46]. Because of the arbitrariness of the factor 2, we will never mix it with other constants so that the reader may modify our end results conveniently. The essential feature of our model is that the cross section need not be increased with increasing line length [46]. $\lambda \sim 1\,\mu$m will be used in numerical examples.

The energy E per transmitted bit will also be assumed to be constant and independent of line length. The length dependent attenuation component of loss can in principal be made very small in comparison to losses due to coupling and device inefficiency factors. We assume the use of light modulators as output transducers so that no threshold term is involved. Expressions for the energy required per transmitted bit were given previously by many authors [10] [11] [14]. Feldman et al. [11] also discussed the effects of fan-out on the energy.

For optical interconnections, T_d is simply the minimum pulse width the modulating and detecting devices can handle. We will be content with a smooth 'hump', rather than a square pulse with sharp edges, so that the highest frequency content need not be much greater than the inverse pulse width. This is consistent with our earlier requirement that a pulse be completely received before its value is registered. Thus, the minimum pulse width will often be approximately equal to (or twice) the slower of the rise times of the modulators or detectors. Electron-hole diffusion or transit time limitations may also contribute to T_d. T_ℓ will most probably be set by material dispersion, since we are assuming single mode guides. Even

if we launch an impulse, a hump of width T_ℓ will arrive at the detector. For free space systems, T_ℓ will be set by the (spatial) dispersive properties of the imaging elements. Pulses shorter than T_ℓ will not be allowed so that the imaging system performs its intended function. For the length scales involved in a computing environment, the effects of dispersion can be made negligible. Thus we will take $T = \max(T_d, T_\ell) = T_d$. Since the refractive indices of most materials are close enough to unity, we will take the propagation velocity as the vacuum speed of light c, so that $T_p = \ell/c$. If a single wavelength source is available, T_r will often be approximately equal to T, unless there are additional restrictions requiring an elapse of time between consecutive pulses, as might be the case with certain types of optical switches. If multiple wavelength sources are available, the use of wavelength division multiplexing might enable the effective value of T_r to be much less than T.

2.4.2 Relations between S, B and N

We postpone the inclusion of heat removal requirements until later. Thus the interelement spacing d is primarily set by the size of the elements and the number of 'wiring' tracks that must pass through each cell [39] [8] [41] [3]. When we speak of a 'volume', it will be understood that we mean an actual volume when $e = 3$ but an area when $e = 2$. A similar convention will apply for the use of the term 'cross-section'. To find the smallest possible value of d, we equate the total volume occupied by the interconnections and primitive elements to the total system volume:

$$\max(Nk\chi\bar{\ell}(width)^{e-1}, Nd_d^e) \simeq Nk\chi\bar{\ell}(width)^{e-1} + Nd_d^e = Nd^e. \quad (2.7)$$

where χ is the number of physical lines used to establish each edge of the connection graph and $\bar{\ell} = \bar{r}d$ is the average connection length in physical units. Ignoring d_d we find

$$d = (k\chi\bar{r})^{1/(e-1)}(width). \quad (2.8)$$

Of course, d may never actually be less than d_d. Note that the same result can be obtained directly by equating $K\chi(width)^{e-1}$ to the cell cross section d^{e-1}, as discussed earlier. When $n > e$, $\bar{r} \propto N^{1/e-1/n}$ so that we find $d \propto N^{(n-e)/ne(e-1)}$. The increase of the volume d^e per element with increasing N has been termed *space dilation* [9]. Space dilation occurs when $n > e$. Equation (2.8) is plotted in Figure 2.5. Given k and d_d, we can use these plots to predict beyond what value of N the system volume will be

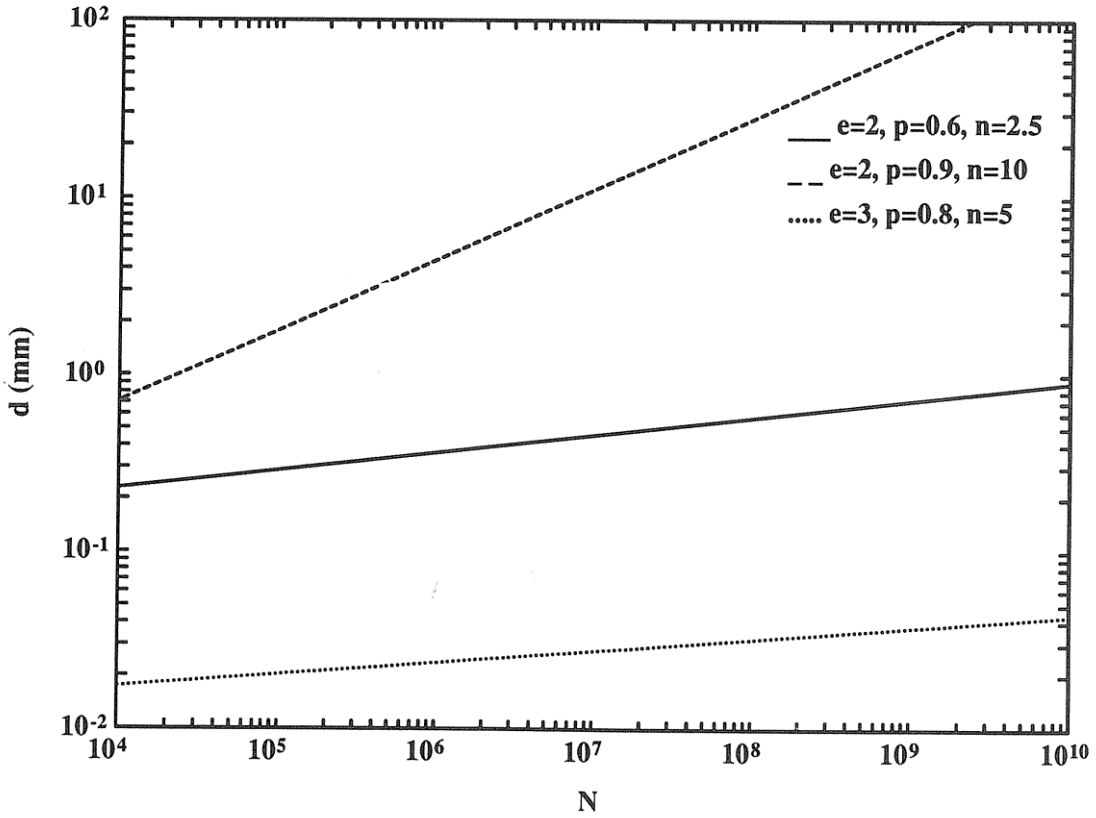

Figure 2.5: Cell size for communication limited systems. We take $k = 10$, $\chi = 1$ and $(width) = 2\lambda$.

determined by communication requirements, rather than by element size. The linear extent of the system may be obtained as

$$N^{1/e}d = N^{1/e}(k\chi\bar{r})^{1/(e-1)}(width). \tag{2.9}$$

Of course, the system linear extent may actually never be less than $N^{1/e}d_d$. In all numerical plots we will vary N from 10^4 to 10^{10} (for comparison, the human brain has about 10^{11} neurons [47]). One should keep in mind however that the larger values of N in this range may lead to unrealistic system sizes for 2 dimensional layouts.

When $T = T_d$ is small, $S = 1/\tau = 1/T_p = c/\ell_{max}$ satisfies

$$Sr_{max}\bar{r}^{1/(e-1)} = \left(\frac{c}{2\lambda}\right)(k\chi)^{-1/(e-1)} \tag{2.10}$$

where $\ell_{max} = r_{max}d$ denotes the length of the longest connection. When $n > e$, using $\bar{r} = \kappa N^{1/e-1/n}$ and $r_{max} \simeq a \simeq N^{1/e}$ we obtain

$$SN^{(n-1)/[n(e-1)]} = \left(\frac{c}{2\lambda}\right)(k\chi\kappa)^{-1/(e-1)}. \tag{2.11}$$

When $BT_r \leq 1$, we simply set $\chi = 1$. When $BT_r > 1$, we must choose[5] $\chi = BT_r$ since a single physical channel is incapable of transmitting information at a rate of B. Then

$$SB^{1/(e-1)} r_{max} \bar{r}^{1/(e-1)} = \left(\frac{c}{2\lambda}\right) (kT_r)^{-1/(e-1)} \qquad BT_r \geq 1 \qquad (2.12)$$

which becomes for $n > e$,

$$SB^{1/(e-1)} N^{(n-1)/[n(e-1)]} = \left(\frac{c}{2\lambda}\right) (kT_r\kappa)^{-1/(e-1)} \qquad BT_r \geq 1. \quad (2.13)$$

Of course, S may never actually exceed $1/T$ or $c/N^{1/e}d_d$. Thus in general there are three regions in the relationship between S and N. The leftmost (small N) region is the device speed limited region ($S = 1/T$), the middle region is the element size-speed of light limited region ($S = c/N^{1/e}d_d$) and the rightmost (large N) region is the communication volume-speed of light limited region (equation 2.11 with $\chi = 1$ or equation 2.13). If p is large, element size is small and/or devices slow, the central region may dissapear.

We remind the reader that the elements must be at least large enough to accomodate $\sim k\chi$ transducers. Also, if an m-fold reduction in T_r were made possible by the use of wavelength division multiplexing of m distinct wavelength sources, this number must be further multiplied by m. For given N, there is a limit to the usefulness of wavelength division multiplexing, since after a certain value of m, system size and delays are set by input-output limitations, rather than the necessary communication volume.

If feasible, the use of multiple layers can contribute to 2 dimensional system performance. The width of a cell d must now satisfy

$$d \geq \max(K\chi/M, 1)(width)$$

where M denotes the number of layers [39] and d must at least be wide enough to admit the passage of one physical channel. We will be justified in writing $d \geq K\chi/M$ since the requirement $d \geq d_d$ will always be stronger than $d \geq (width)$. It is important to note that there is a maximum useful value of M. Increasing the number of 'wiring' layers, without increasing the number of layers for input-output ports, leads to element size limited systems. Assuming this maximum useful value is not exceeded, the right hand sides of the above equations are improved by a factor of M. If the

[5] Strictly speaking, χ, being an integer quantity, is given by $\chi = \lceil BT_r \rceil$, which we approximate as $\max(1, BT_r)$.

number of layers is large, the effects of vertical runs must be taken into account. This is considered in detail in appendix A.3.

To illustrate the usefulness of our formulation, we consider a simple example derived from concurrent computer architecture. It is often the case that one desires to minimize the first-to-last bit communication latency τ_L of L bit messages. Thus, we desire to minimize

$$\tau_L = \tau + \frac{L}{B} = \frac{1}{S} + \frac{L}{B}. \tag{2.14}$$

Let us assume $N = 10^6$, $k = 10$, $e = 2$, $n = 3$, $T = T_r = 1\,\text{nsec}$ and $L = 20$. Using (2.13) we find that the optimum value of B is $\simeq 4\,\text{Gbit/sec}$ so that we choose $\chi = 4$. S and τ_L may be calculated as $\simeq 150 \times 10^6\,\text{sec}^{-1}$ and $\simeq 12\,\text{nsec}$ respectively.

2.4.3 Heat Removal

We assume that the energy E associated with each transmitted bit is dissipated and must be removed from the system. We also assume that the dissipation associated with the elements are negligible. If not, we simply need substitute $E \rightarrow E + E_d/k$ where E_d denotes the energy dissipation associated with an element.

The 2 and 3 dimensional cases are treated separately.

A. Two Dimensions

Heat removal considerations will also set a lower limit to the cell size d, and hence system size and delay. The total power dissipation is given by $kNEB$. Let Q denote the amount of power we can remove per unit area. Thus we must maintain $QNd^2 \geq kNEB$. Starting from this relation, the heat removal limited version of (2.13) may be derived as

$$SB^{1/2}N^{1/2} = \frac{c}{(E/Q)^{1/2}}k^{-1/2}. \tag{2.15}$$

When both communication volume and heat removal considerations are taken into account we have (for $p > 1/2$)

$$\frac{1}{S} = \frac{N^{1/2}d}{c} = \frac{N^{1/2}}{c}\max\left(k\chi\kappa N^{p-1/2}(2\lambda), \left(\frac{kEB}{Q}\right)^{1/2}\right). \tag{2.16}$$

where $\chi = \max(1, BT_r)$. Of course, as always, S can never be greater than $c/N^{1/2}d_d$ or $1/T$. Equation 2.16 is plotted in Figure 2.6 with EB/Q as

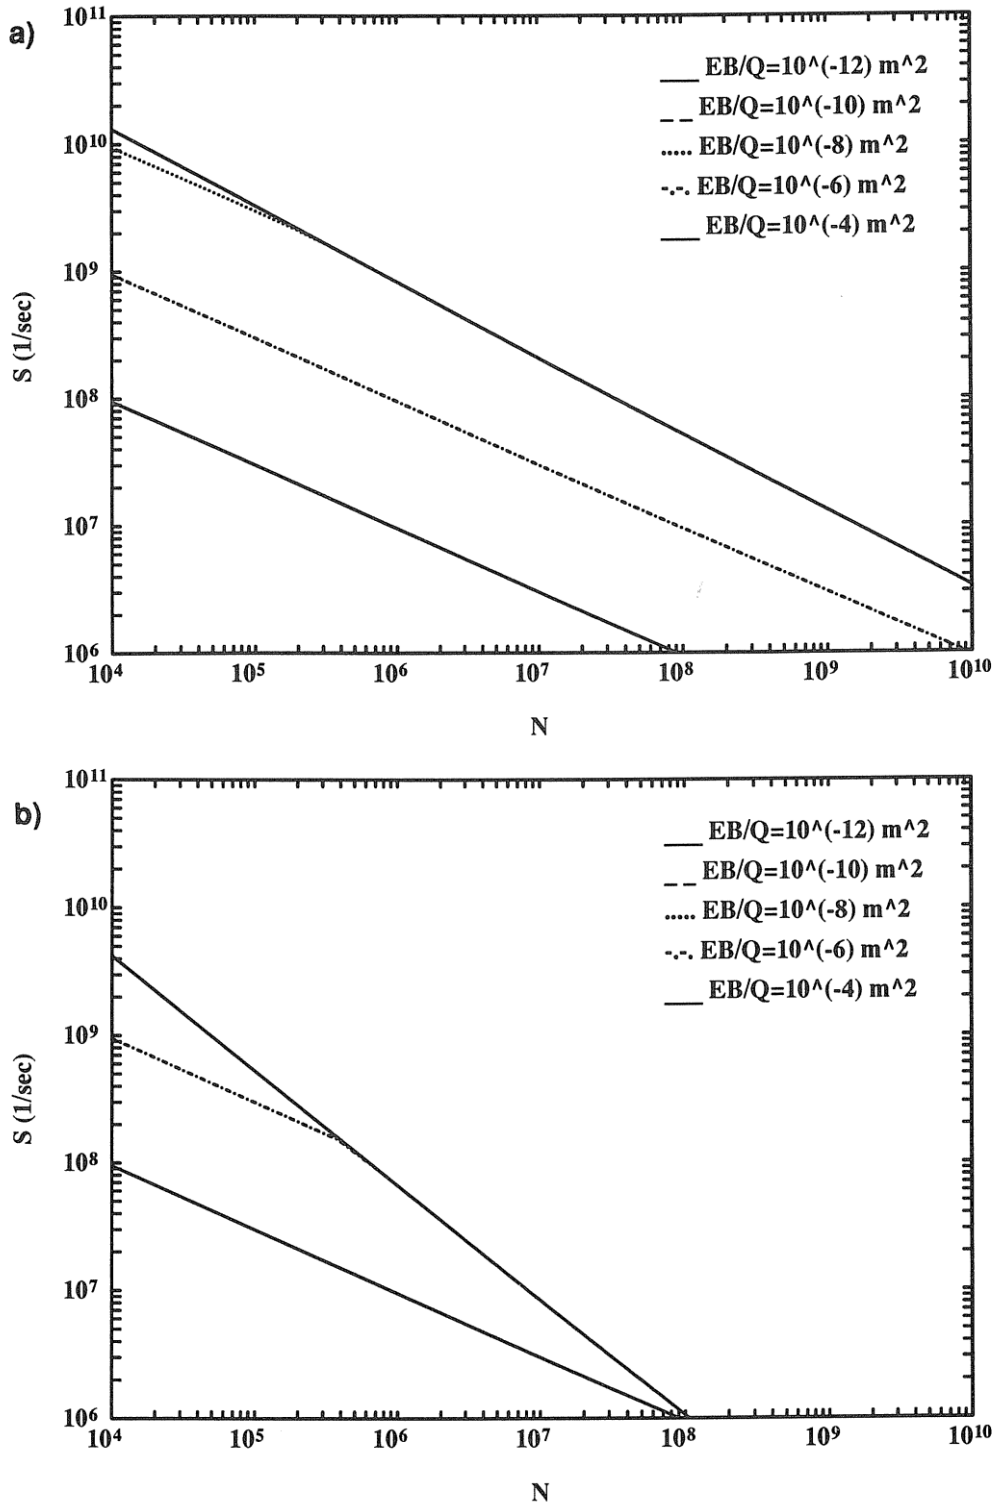

Figure 2.6: The effect of heat removal requirements in 2 dimensions. We take $k = 10$. d_d, T_d and T_r are assumed small enough to have no effect. The range of variation of EB/Q has been chosen based on the typical ranges of variation of the individual parameters. For smaller values of EB/Q, the system is communication limited so that the curves corresponding to these values coincide. (a) $e = 2$, $p = 0.6$, (b) $e = 2$, $p = 0.9$.

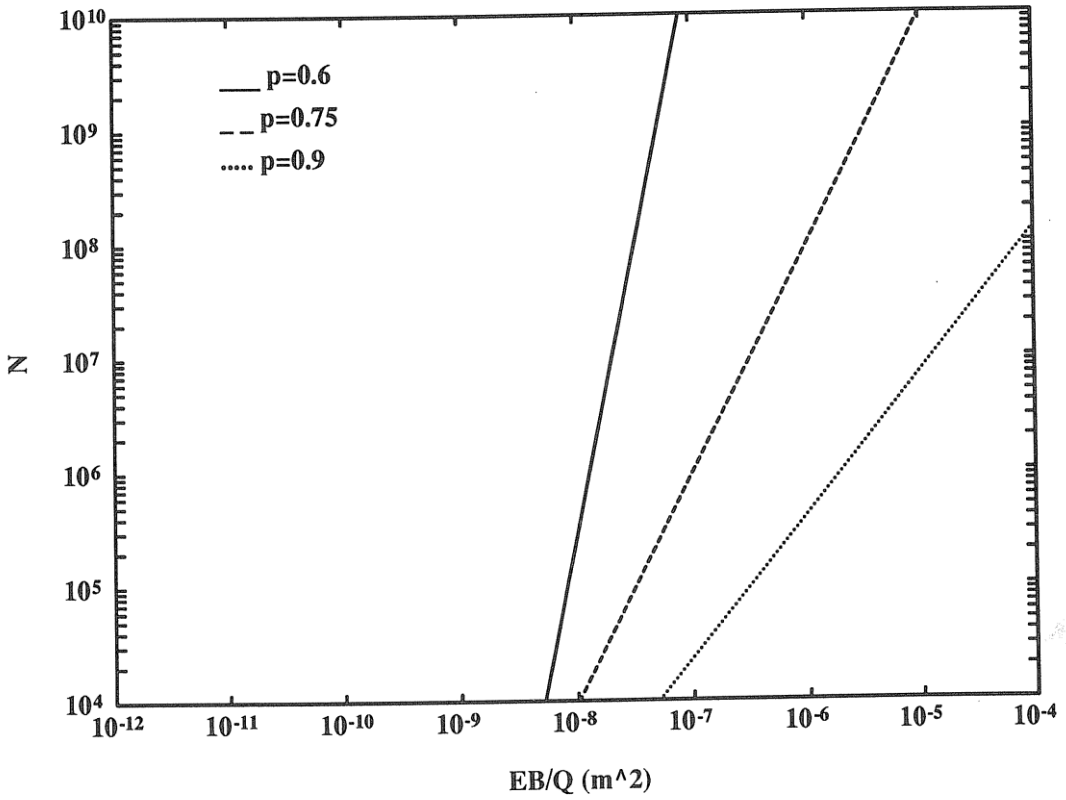

Figure 2.7: Critical value of N beyond which heat removal is not a limiting factor. We take $k = 10$ and assume T_r is small so that $\chi = 1$. EB/Q is varied over the same range as in the previous figure.

a parameter. We have assumed d_d and T to be negligible so as to make transparent the effects of heat removal. Notice that if B is kept constant, for large enough N the system is always communication volume limited, rather then heat removal limited. The critical system size beyond which heat removal is no longer a limiting factor is plotted in Figure 2.7 for various values of p. In some cases, the device speed and/or element size-speed of light limited regions may extend into the communication volume-speed of light limited region so that the heat removal limited region completely dissapears.

Many seemingly 3 dimensional optical architectures are actually as limited as the 2 dimensional case we have just considered [46] [48]. In particular, certain multi-facet holographic architectures [10] [49] [12] can be very inhibitive. In fact, when $p < 1$ but the longest interconnection is still of the order of the linear extent of the array of elements, these architectures are even worse than the fully 2 dimensional case we have considered. In this case, the linear extent of the system grows as $\propto N$ [12] [46] so that the

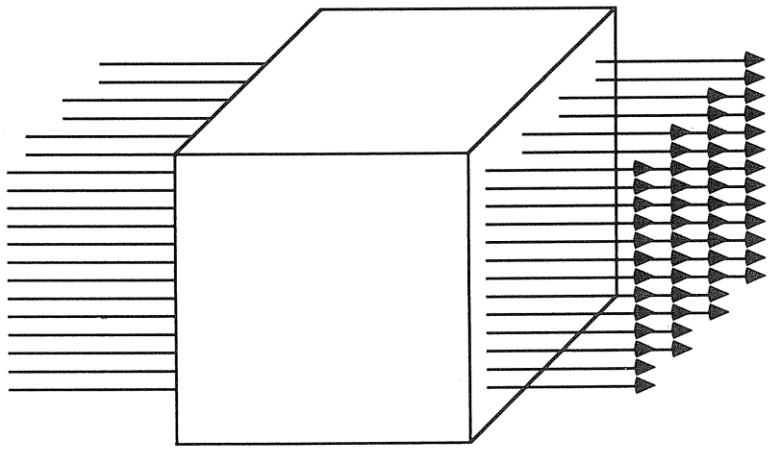

Figure 2.8: 3 dimensional heat removal.

asymptotic growth rate of signal delay is $\tau \propto N$ as opposed to $\propto N^p$ which we have found for the fully 2 dimensional case (equation 2.11). Because of their flexibility in providing an arbitrary pattern of interconnections, multi-facet architectures form the basis of many suggested optical computing schemes. Our results are also valid for systems where the primitive elements are optical switches, provided we interpret E as the switching energy. Thus we conclude that for large N, heat removal is not the limiting factor for such systems as well.

B. Three Dimensions

Whatever the modality (conduction, convection or radiation), heat transfer can only take place *through* a surface (flux conservation). Thus, just as in the 2 dimensional case, Q will be specified as the power which we can remove per unit cross sectional area. This is most easily visualized by considering the flow of a cooling fluid through our system, as illustrated in Figure 2.8. A fluid with heat capacity C_s, mass density ρ_m flowing at an effective mean velocity v_f may carry away *at most* $Q = v_f \rho_m C_s \Delta T$ where $\Delta T = T_{max} - T_{init}$ [50] [51]. T_{max} is the maximum permissible operating temperature of the devices and T_{init} is the initial temperature of the coolant.

The fluid flowing through a cross section d^2 must carry away the power dissipation associated with a stack of $N^{1/3}$ elements. Thus the heat removal condition in this case is

$$Qd^2 \geq kEBN^{1/3}. \tag{2.17}$$

We can also arrive at this by requiring $QN^{2/3}d^2$ to exceed the total power dissipation $kNEB$. The above equation results in a larger value of d than

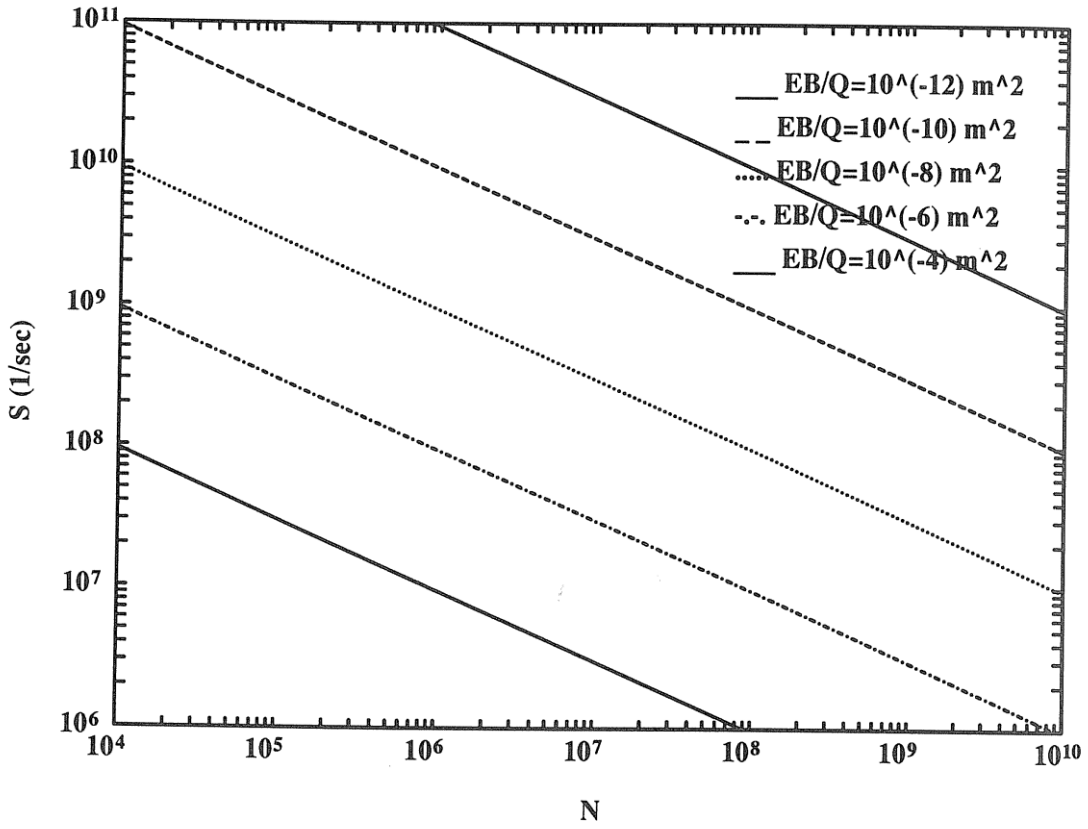

Figure 2.9: The effect of heat removal requirements in 3 dimensions. We take $k = 10$, $e = 3$ and $p = 0.8$. d_d, T_d and T_r are assumed to be small enough to have no effect. Unless p is very close to unity, heat removal is the major limiting factor for fully 3 dimensional layouts.

for the 2 dimensional case, but the same system linear extent $N^{1/3}d = (kEB/Q)^{1/2}N^{1/2}$.

Thus, using the above constraint and (2.8) we can show

$$\frac{1}{S} = \frac{N^{1/3}}{c} \max\left((k\chi\kappa N^{p-2/3})^{\frac{1}{2}}(2\lambda), \left(\frac{kEBN^{1/3}}{Q}\right)^{1/2}\right) \qquad (2.18)$$

which is plotted in Figure 2.9. Again we assume device related limitations to be negligible. If B is kept constant as N is increased, the heat removal term eventually dominates the communication volume term, unless $p = 1$, when they grow together. Thus, for large N, highly interconnected systems do not suffer greater delay than locally interconnected ones.

2.4.4 Asymptotic Properties

We make several observations regarding (2.13). For given N, B can be arbitrarily increased[6] by incurring a reduction in S. This would be desirable when massive data transfer is required but the time elapse in which this transfer takes place is not critical. Thus the $HB = k\kappa N^p B$ product can be arbitrarily increased by suffering a decrease in S. In fact, using (2.13), the tradeoff between HB and S may be written in the transparent form

$$S[HB]^{1/(e-1)} = \left(\frac{c}{2\lambda}\right) T_r^{-1/(e-1)}. \tag{2.19}$$

Unlike B, S *cannot* be arbitrarily increased by reducing B, since once B drops below $1/T_r$, (2.11) with $\chi = 1$ is applicable. The growth rate of the delay with increasing N is then given by $\tau \propto N^{p/(e-1)}$. Thus, for given B, the dependence of HS on N is given by

$$HS \propto N^{p(e-2)/(e-1)}. \tag{2.20}$$

We observe that the HS product may be arbitrarily increased by increasing N provided $e > 2$. If the largest possible value of $e = 3$ can be attained, we have $HS \propto N^{p/2}$. Despite the faster growth rate of delay, systems with large p have a faster increase of HS with N.

The above results must be modified if heat removal is accounted for. B and HB can again be arbitrarily increased at the expense of S, this time according to (2.16) or (2.18), for 2 and 3 dimensions respectively.

In both 2 and 3 dimensions heat removal considerations result in a growth rate of the delay $\propto N^{1/2}$. Thus

$$HS \propto N^{p-1/2}. \tag{2.21}$$

The resultant growth rate of HS is thus the slower of those given by (2.20) and the above.

2.5 Normally Conducting Interconnections

2.5.1 Physical Model Description

Our analysis will be based on the distributed parameters R, L and C; the resistance, inductance and capacitance of the line per unit length (Figure 2.10). As is mostly appropriate [52], the shunt dielectric conductance

[6] Of course, as far as the interconnection network is concerned, remember that we are assuming T_{rd} to be negligibly small.

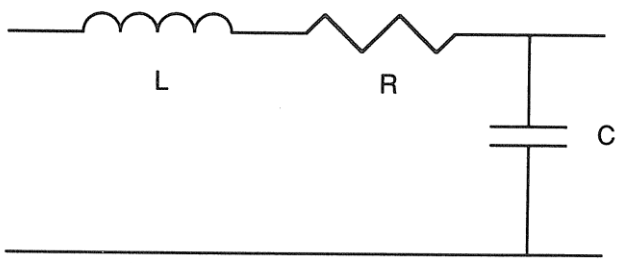

Figure 2.10: Section of distributed RLC interconnection model.

is ignored. Figure 2.11 depicts the physical cross section of our model. Other geometries are also possible and would change our results by only geometrical factors. It is well known that once a line starts becoming taller than it is wide, the line to line separation must be increased greatly to maintain acceptable crosstalk levels, whereas the capacitance and characteristic impedance are improved at most logarithmically [14]. For this reason, we will require that our lines satisfy $t \leq h$ and $h \leq w/2$. With these constraints, we will assume—based on a similar argument regarding crosstalk as in the optical case—that the minimum packing dimension is $(width) = 2w$ in 2 dimensions and $(area) = (width)(height) = 2w \times 2(h+t)$ in 3 dimensions, independent of length. Whereas the numerical factors involved are again somewhat arbitrary, they seem to be representative of the geometry to which technology is converging [17]. The resistance, inductance and capacitance per unit length; and propagation velocity and characteristic impedance of this line are approximately given by $R = \rho/w \min(t, \delta)$, $C = \epsilon w/h$, $L = \mu h/w$, $v = 1/\sqrt{LC} = 1/\sqrt{\mu\epsilon}$ and $Z_0 = \sqrt{L/C} = \sqrt{\mu/\epsilon}\ h/w$ respectively. ρ is the resistivity of the conductor and ϵ the permittivity of the dielectric. We will use room temperature aluminum resistivity and a relative permittivity $\epsilon_r = 4$ in our numerical examples. The permeability μ will be taken equal to that of free space. $\delta = \sqrt{2\rho/\omega\mu}$ denotes the classical skin depth at frequency ω. Unless otherwise stated, the voltage level will be taken as $V = 1$ Volt in numerical examples.

Based on this model, we will show that the line imposed minimum temporal pulse width for a normally conducting interconnection is given approximately by

$$T_\ell = (16\rho\epsilon)\frac{\ell^2}{(width)^2} = (16\rho\epsilon)\frac{\ell^2}{(area)}. \tag{2.22}$$

We will not mix the constant 16 with other constants so as to enable easy modification of end results. As in the optical case, the device imposed minimum temporal pulse width T_d is set by the intrinsic limitations of the

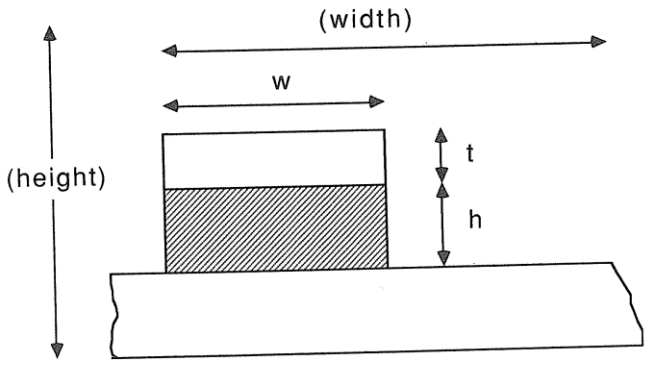

Figure 2.11: Physical cross section of interconnection model. t and h are the heights of the conductor and dielectric respectively. We constrain $t \leq h$ and $h \leq w/2$. These constraints not only ensure reasonable confinement of the fields, but also justify approximate use of a parallel plate model for calculating capacitances and inductances. $(width) = 2w$ is the two dimensional minimum packing dimension. $(area) = (width)(height) = 4w(h+t) \simeq 4w\max(h,t) = 4wh$ is the three dimensional minimum packing area.

transmitting and receiving devices and is assumed to be a given constant. Thus $T = \max(T_d, T_\ell)$ is the minimum temporal width associated with each bit[7]. For conducting interconnections, the minimum pulse repetition interval is simply given by $T_r = T = \max(T_d, T_\ell)$ so that we will drop the subscript r.

Equation 2.22 is valid for both RC lines which are left unterminated and charged up and for terminated transmission lines. For unterminated lines, $T = \max(T_d, T_\ell)$ is simply the RC rise time of the receiving end voltage and satisfies $T > T_p = \ell/v$ where v is the propagation velocity [53]. As in the optical case, we are not requiring sharp square pulses and are content with smooth 'humps'. Pulse transmission is not possible along high-loss lines (i.e. lines for which $R\ell > Z_0$), such lines must be charged up. In general, it is energetically wasteful to terminate a line if $T > T_p$, since in this case the energy per transmitted pulse $E = V^2T/Z_0$ would exceed that possible with an unterminated line $E = V^2C\ell$. However, when $T < T_p$ (which is possible only for lines with sufficiently low loss), it is beneficial to terminate the line so as to pipeline pulses through the line with less energy without worrying about reflections. In this case of terminated transmission, T corresponds to the minimum temporal width of a pulse travelling along the line. Thus we are agreeing to leave a line unterminated when $T > T_p$ and to terminate

[7] In practice, T_ℓ and T_d may be coupled, as in MOS VLSI technology. However, it is mostly possible to break the total pulse width into the maximum (or sum) of a line independent constant T_d and a device independent function of line parameters T_ℓ, enabling us to maintain a device independent model. This is further discussed in appendix A.4.

it when $T < T_p$. We assume perfect termination is possible. The signal delay for any normally conducting line can be written as $\max(T_d, T_\ell, T_p)$. If $T < T_p$ and the line is terminated, T_p/T pulses may be simultaneously in transit along the line.

We derive (2.22) first for unterminated lines. It is known that the skin effect need not be considered in this mode of operation [53]. The rise time of the line is given by $\simeq (R_d + R\ell)C\ell$ where R_d is the drive impedance [1]. Assume for the moment that the line is not drive limited, then the rise time and energy per bit are given by

$$T_\ell = RC\ell^2 = \frac{\rho}{wt}\frac{\epsilon w}{h}\ell^2 = \frac{\rho\epsilon\ell^2}{ht} = 4\rho\epsilon\frac{w}{h}\frac{w}{t}\frac{\ell^2}{(width)^2} \tag{2.23}$$

$$E = V^2C\ell = \epsilon V^2\frac{w\ell}{h} \tag{2.24}$$

where $(width) = 2w$ has been used. We are again assuming the energy dissipated by the devices to be negligible. It is evident from these equations that one should choose h/w and t/w as large as possible. Just as it is not beneficial to make lines tall and skinny, neither is it to make them flat and wide. Thus with $h = w/2$ and $t = h$ we obtain (2.22) and an expression for the energy: $E = 2\epsilon V^2\ell$. This discussion is consistent with and confirmed by the somewhat different approach of Masaki [54].

The performance of present day VLSI lines may be much worse than predicted by the above, because such lines are often drive limited [5] [55], i.e. $R_d > R\ell$. The above corresponds to what may be achieved with arbitrarily strong drivers. A more detailed discussion is given in appendix A.4.

Now we turn our attention to terminated transmission lines. We ignore the effects of dispersion, anomalous skin effect and assume the quasi-TEM approximation to be valid. We will show that the fundamental frequency satisfies $\omega L > R$ so that we may ignore the correction terms [56] $1/(1 + R^2/8\omega^2L^2)$ and $(1 - jR/2\omega L)$ associated with the propagation velocity and characteristic impedance respectively.

In this case, the minimum pulse width satisifes $T < T_p = \ell/v$. Since we are not insisting on sharp square pulses, but are satisfied with rounded 'humps', the highest frequency content need not be much greater than the inverse pulse width. Of course, since $T < T_p$, a frequency of at least $\omega \sim 2/T_p = 2v/\ell$ exists. Since the attenuation coefficient α of a transmission line is given by $R/2Z_0$ [57], we require approximately $R\ell \leq Z_0$ so that

attenuation is kept at an acceptable level[8]. Using these relations, we may immediately show $\omega L > R$ which we have promised above. Furthermore, one can show that the skin depth $\delta = \sqrt{2\rho/\omega\mu}$ satisfies $\delta^2 R \leq \rho h/w$. The resistance per unit length R is given by $R = \rho/w \min(t, \delta)$ so that $\delta^2 \leq h \min(t, \delta)$. Since $t \leq h$, this leads to $\delta \leq h$.

Since $(height)$ is already determined within a factor of 2 by h, we will agree never to set $t < \delta$ and unnecessarily increase the resistance. We can always do this without violating the constraint $t \leq h$ since we have just shown that $\delta \leq h$. Thus, we have $R = \rho/w\delta$.

Now the attenuation condition $R\ell \leq Z_0$ may be used to set a lower bound on $h\delta$ as

$$h\delta \geq \rho\sqrt{\frac{\epsilon}{\mu}}\,\ell. \tag{2.25}$$

A lower bound on the skin depth leads to an upper bound on the largest frequency component and hence to a lower bound on the minimum pulse width T_ℓ. Thus using $\omega \sim 2/T_\ell$ and the definition of the skin depth we obtain the minimum line imposed pulse width as

$$T_\ell = \rho\epsilon\frac{\ell^2}{h^2} = 4\rho\epsilon\left(\frac{w}{h}\right)^2\frac{\ell^2}{(width)^2} = (16\rho\epsilon)\frac{\ell^2}{(width)^2} = (16\rho\epsilon)\frac{\ell^2}{(area)} \tag{2.26}$$

where we have taken $h = w/2$ to keep T_ℓ as small as possible. Of course, no matter how small T_ℓ is, we cannot shape pulses shorter than T_d so that $T = \max(T_d, T_\ell)$. The energy is given by $E = V^2T/Z_0 = (w/h)\sqrt{\epsilon/\mu}\,V^2T$ which also indicates that we should choose h/w large. Thus we obtain $E = 2\sqrt{\epsilon/\mu}\,V^2T$ as the energy per transmitted bit in this case.

When T_ℓ exceeds T_d, we may express the condition $T = T_\ell \leq T_p$ as $16\rho\sqrt{\epsilon/\mu}\,\ell = 16\rho\epsilon v\ell \leq (width)^2 = (area)$. If this condition is not satisfied, the line is high-loss and pulse transmission is not possible. Our model equations are summarized in Table 2.1 and Figure 2.12 for the case $T_d \leq T_\ell$. Of course, neither T nor τ may actually be less than T_d. We also note that it is suboptimal to work with $T_d > T_\ell$. If for any given $(width)$ and ℓ we have $T_\ell < T_d$, we can reduce the $(width)$ of the line until $T_\ell = T_d$, ending up with a wire that occupies less space with the same pulse width and delay.

Referring to (2.22), we ask whether it is beneficial to use a bundle of narrow lines or a single wide line in order to achieve the greatest information throughput. First consider a 3 dimensional layout. Increasing $(width)$ by two (i.e. $(area)$ by four) decreases T_ℓ by four, corresponding to a potential

[8]This corresponds to degradation of the signal level by $e^{-\alpha\ell} = e^{-0.5} = 0.6$.

Table 2.1: Normally conducting interconnection model when $T_d \leq T_\ell$.

	$(width)^2 \leq 16\rho\epsilon v\ell$	$(width)^2 \geq 16\rho\epsilon v\ell$
delay, $\tau = \max(T, T_p)$	$16\rho\epsilon\frac{\ell^2}{(width)^2}$	$\frac{\ell}{v}$
pulse width, $T = \max(T_\ell, T_d)$	$16\rho\epsilon\frac{\ell^2}{(width)^2}$	$16\rho\epsilon\frac{\ell^2}{(width)^2}$
energy, E	$2\epsilon V^2\ell$	$2\epsilon V^2 vT$
termination	no	yes

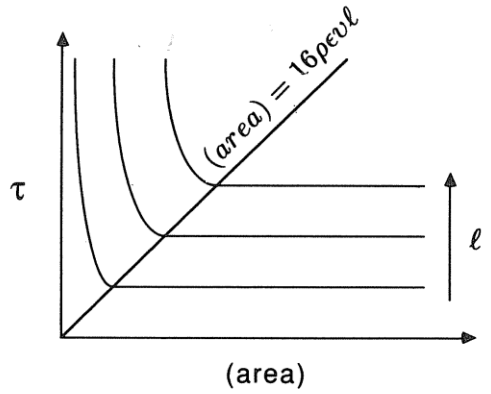

Figure 2.12: Normally conducting interconnection model when $T_d \leq T_\ell$.

increase in bit repetition rate by four. However, we are now able to pack only a fourth as many lines in the same cross sectional area. Thus in 3 dimensions, the same amount of information can be transmitted through given area in given time. Of course, we should never attempt to reduce T_ℓ below T_d, since then the increase in $(width)$ cannot be compensated by an increase in bit repetition rate. In 2 dimensions, T_ℓ is again reduced by four, but the linear packing density is reduced by only a factor of two, so that throughput is increased! Thus, as long as T_ℓ dominates T_d, we will agree to use a single wide line ($\chi = 1$) rather than many narrow ones to establish each edge of the connection graph.

2.5.2 Relations between S, B and N

In order to accomodate $K = k\bar{r}$ wiring tracks, the linear extent of a cell d must satisfy $d^{e-1} \geq K(width)^{e-1}$ [39]. We must also satisfy $d \geq d_d$ and the heat removal requirement, which will be discussed later. We are free in choosing $(width)$ provided it exceeds a certain minimum manufacturable value. If d_d is small and heat removal is not an issue, we would prefer to set

$(width)$ to this minimum so as to make d and the overall system as small as possible. In this case, $d^{e-1} = K(min.\ manuf.\ width)^{e-1}$. However, if element size or heat removal require that we set

$$d^{e-1} > K(min.\ manuf.\ width)^{e-1},$$

we will agree to increase $(width)$ until $d^{e-1} = K(width)^{e-1}$. If d and hence the lengths of the lines are already set by factors other than wiring density, we increase $(width)$ so as to fill up available space. In this way, we reduce the resistance of the lines as much as possible. Despite the fact that increasing $(width)$ no longer decreases the pulse width or delay once the line becomes device or propagation limited, we will never be at a disadvantage by choosing $(width)$ in this manner. As noted before, the error associated with assuming that the elements and wires may cooccupy the same physical space is bounded by a factor of 2 (i.e. we are using $\max(x, y) \simeq x + y$).

Thus using $d^{e-1} = K(width)^{e-1}$ and (2.22) with $\ell_{max} = r_{max}d$ we immediately obtain

$$T_\ell = (16\rho\epsilon)r_{max}^2(k\bar{r})^{2/(e-1)}. \tag{2.27}$$

When $T_\ell \geq T_d$ so that $T = T_\ell$, the maximum value of B satisfies $B = 1/T_\ell$ or

$$Br_{max}^2\bar{r}^{2/(e-1)} = (16\rho\epsilon)^{-1}k^{-2/(e-1)} \qquad BT_d \leq 1 \tag{2.28}$$

which becomes, for $e < n$, using $r_{max} \simeq a \simeq N^{1/e}$ and $\bar{r} = \kappa N^{1/e-1/n}$,

$$BN^{2(n-1)/[n(e-1)]} = (16\rho\epsilon)^{-1}(k\kappa)^{-2/(e-1)} \qquad BT_d \leq 1. \tag{2.29}$$

If the above equation predicts $B \leq 1/T_d$, then we are justified in assuming that $T_\ell \geq T_d$. Otherwise, we must use $\chi = BT_d > 1$ physical lines per graph edge in order to improve B beyond $1/T_d$. Since each physical line is bottlenecked by T_d, there is no use making T_ℓ any smaller than T_d. Thus with $T_d = T_\ell = (16\rho\epsilon)\ell_{max}^2/(width)^2$, $d^{e-1} = \chi K(width)^{e-1}$ and $\chi = BT_d$ we obtain

$$B^{2/(e-1)}r_{max}^2\bar{r}^{2/(e-1)} = (16\rho\epsilon)^{-1}k^{-2/(e-1)}T_d^{(e-3)/(e-1)} \qquad BT_d \geq 1 \tag{2.30}$$

which becomes for $e < n$,

$$B^{2/(e-1)}N^{2(n-1)/[n(e-1)]} = (16\rho\epsilon)^{-1}(k\kappa)^{-2/(e-1)}T_d^{(e-3)/(e-1)} \qquad BT_d \geq 1. \tag{2.31}$$

The above assumes the use of a constant (*width*) for lines of all lengths. Actually, since we are only interested in the delay and pulse width along the longest (worst case) connection, we may make shorter lines narrower with the objective of reducing cell size. Thus, for the case $BT_d \leq 1$, let us set

$$\frac{1}{B} = T_\ell = (16\rho\epsilon)\frac{r^2 d^2}{(width(r))^2} = constant \tag{2.32}$$

for all lines. That is, the width of each line is chosen in proportion to its length so that all lines have the same T_ℓ. Then, since a wire of length rd occupies volume (or area) $rd(width(r))^{e-1}$, the minimum value of d must satisfy

$$d^e = \int rd(width(r))^{e-1} g(r)\, dr. \tag{2.33}$$

Solving for $(width(r))$ from (2.32) and performing the integration we

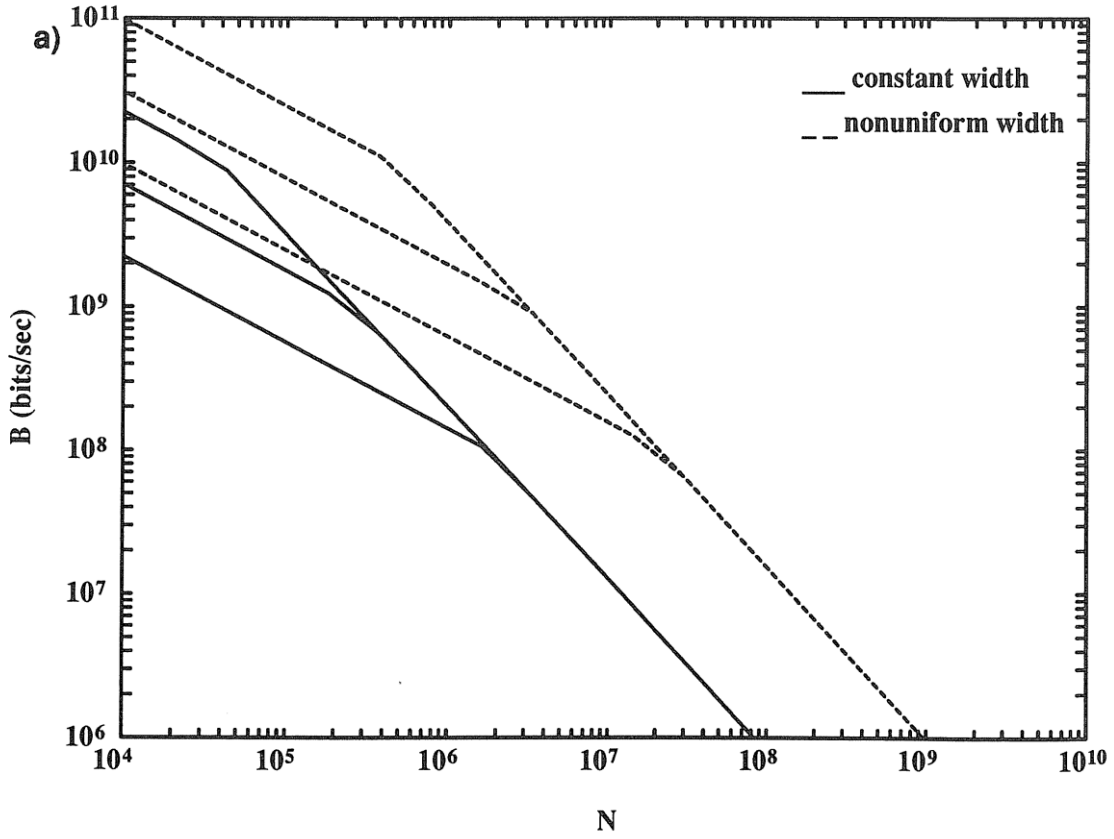

Figure 2.13: B versus N for normally conducting interconnections. The branches correspond to different values of T_d (0.1, 1 and 10 nsec). We take $k = 10$. $M = 10$ layers has been assumed for the 2 dimensional cases. (a) $e = 2$, $p = 0.6$. {Continued}

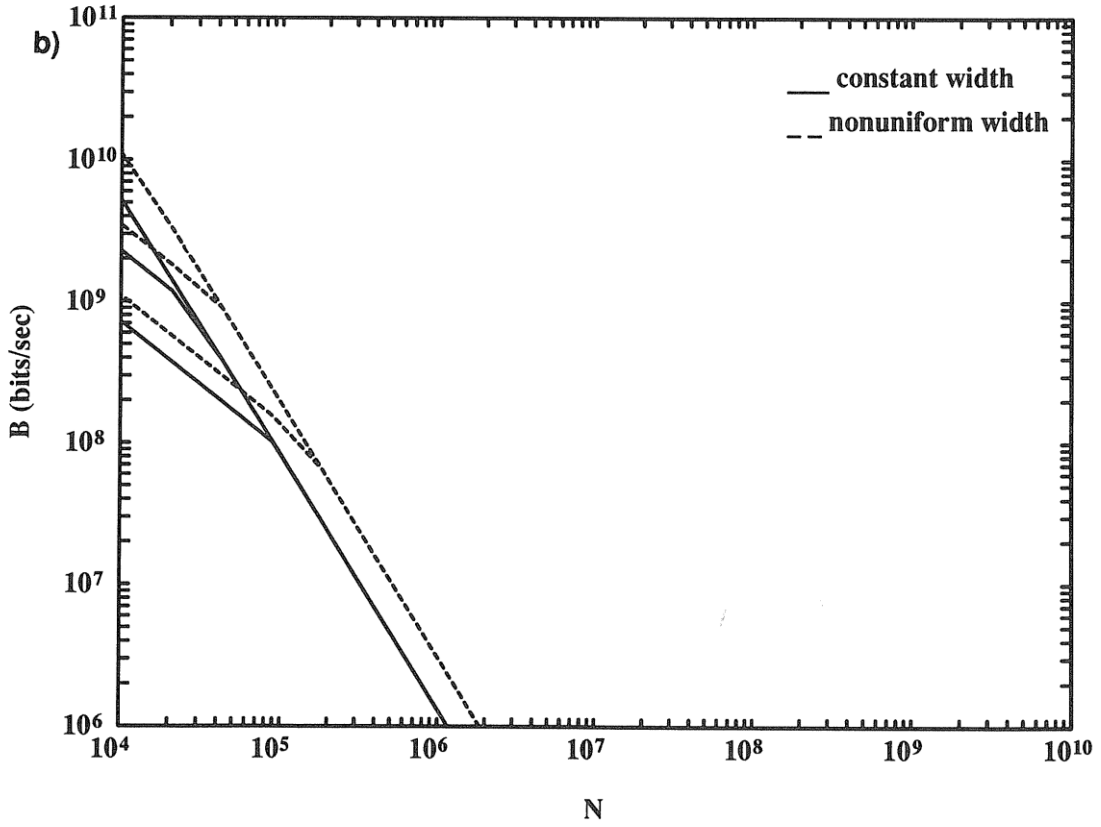

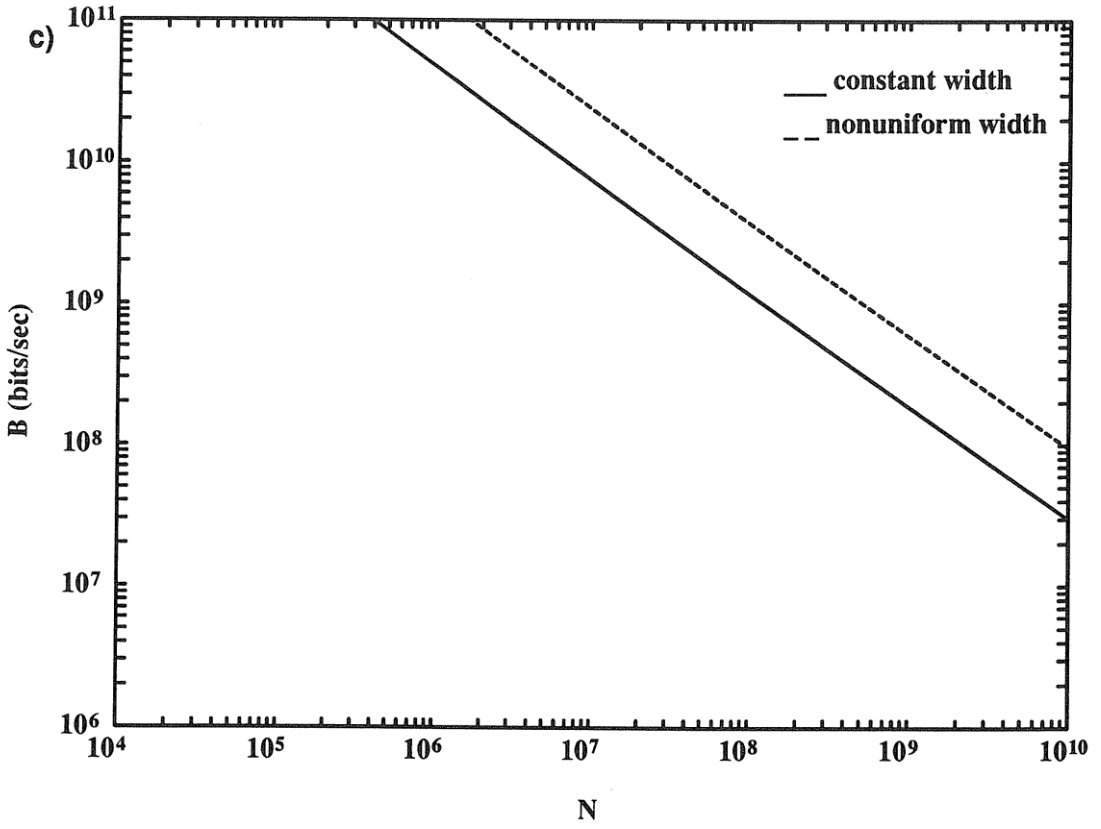

Figure 2.13: (Continued). (b) $e = 2$, $p = 0.9$, (c) $e = 3$, $p = 0.8$.

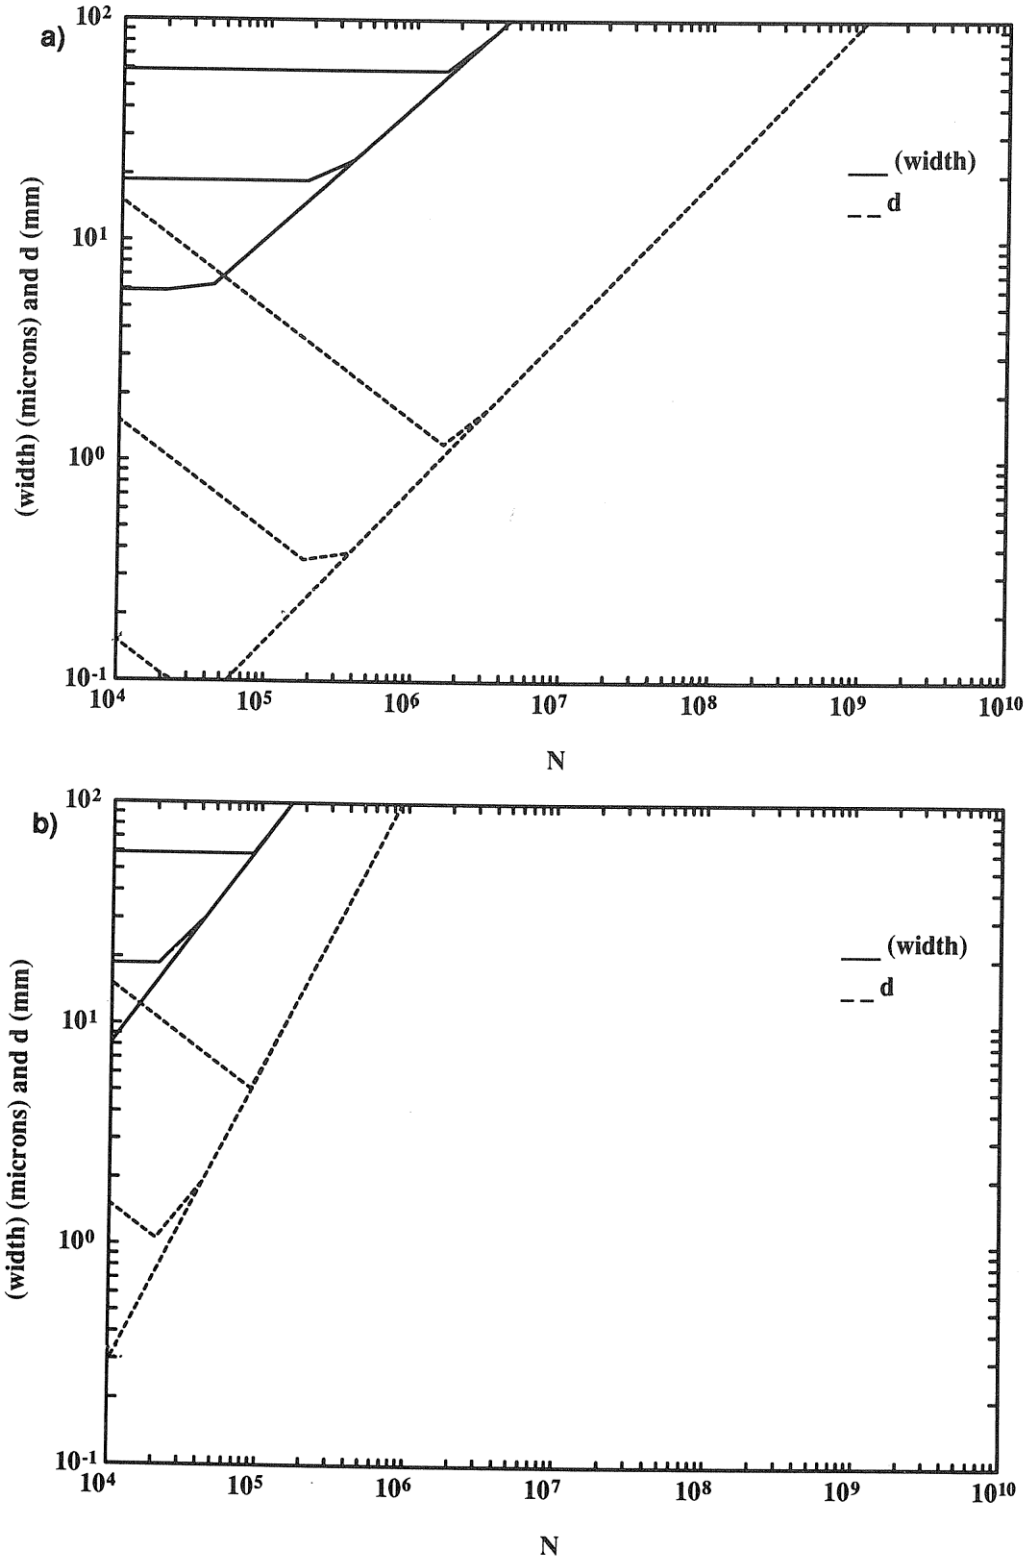

Figure 2.14: Critical width and cell size below which inductive effects need not be considered. Each part of this figure corresponds to those of the previous figure. (a) $e = 2$, $p = 0.6$, (b) be $= 2$, $p = 0.9$.{Continued}

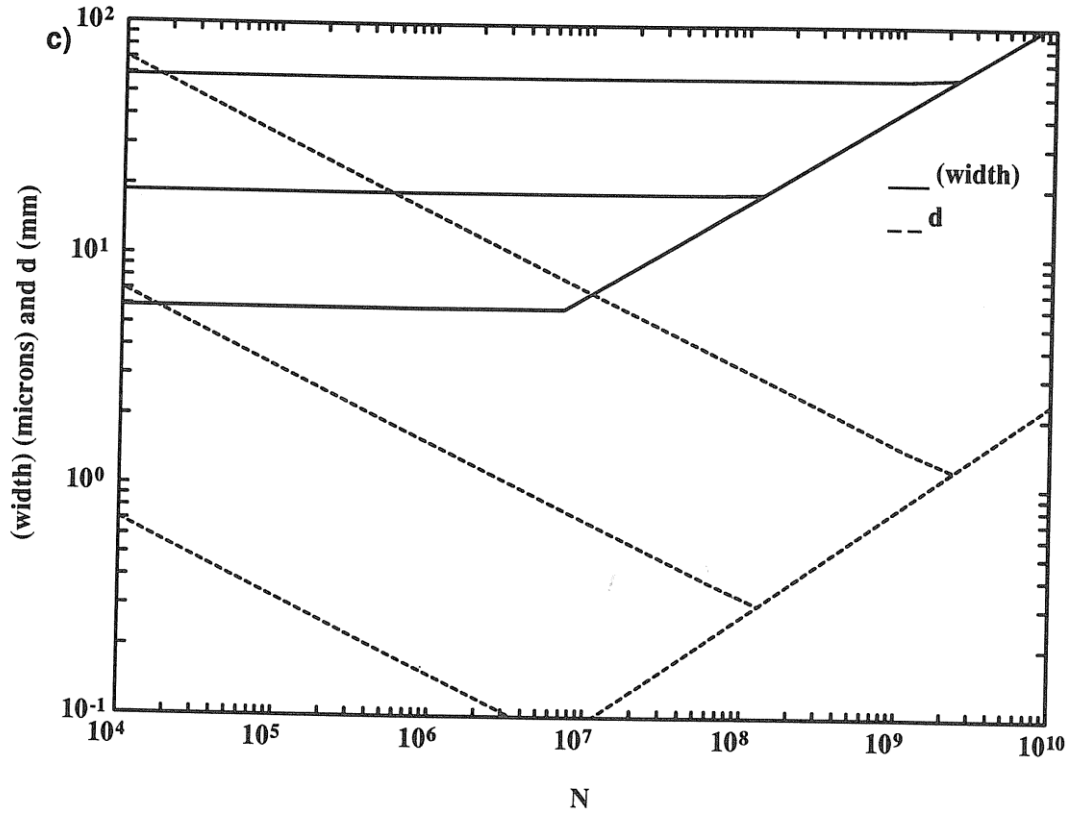

Figure 2.14: (Continued). (c) $e = 3$, $p = 0.8$.

find

$$\frac{1}{B} = T_\ell = (16\rho\epsilon)(k < r^e >)^{2/(e-1)} \qquad BT_d \leq 1 \qquad (2.34)$$

leading to, for $e < n$,

$$BN^{2(n-1)/[n(e-1)]} = (16\rho\epsilon)^{-1}(k\zeta_e)^{-2/(e-1)} \qquad BT_d \leq 1 \qquad (2.35)$$

which represents an improvement over (2.29) by only a constant factor! The asymptotic dependence of B on N remains unchanged. Similarly, when $BT_d \geq 1$, one can show that (2.31) is improved by the same factor. It is important to note that this calculation overestimates the improvement possible since we may not be able to manufacture the shortest lines as narrow as dictated by (2.32).

The use of multiple layers in 2 dimensions contributes greatly to performance. The number of layers M may exceed the order of $\sim$ 10 [58]. Of course, our previous comments regarding the existence of a maximum useful value of M apply to this case as well. Increasing M beyond this value will no longer contribute to performance. Assuming this value is not exceeded, the right hand sides of the above equations are improved by a factor of M^2.

Equations 2.29 and 2.31 define the relation between the maximum possible value of B and N over the whole range of N. This relation has been plotted in Figure 2.13 along with the corresponding relation derived using (2.32). The improvement possible using nonuniform linewidths is greatest when p is small and least when p is large. When p is small, there exists a larger fraction of shorter lines so that greater reduction in cell size is possible.

The above relations are *scale invariant* in the sense that they do not depend on the actual choice of $(width)$, provided $(width)$ is chosen large enough to fill available wiring space, as discussed at the beginning of this subsection. This result, based on interconnect scaling, is in contrast to those based on device scaling, which predict ever increasing performance as the scale is reduced [1].

Until now, we refrained from mentioning S. By definition, S may never exceed $1/\max(T_\ell, T_d, T_p)$. The above relations for B may be used to find $S = 1/\max(1/B, T_d, T_p)$. Remember that the condition for $T_p < T_\ell$ was $(width)^2 < 16\rho\epsilon v\ell$. As we scale down the system photographically, all linear dimensions are decreased in proportion. Thus, below a certain critical $(width)$, this condition is satisfied so that propagation effects (i.e. inductive effects) need not be considered. Indeed, downscaling is recognized as a useful tactic to ensure that S is not worse than $1/\max(T_\ell, T_d) = 1/\max(1/B, T_d)$ [54]. But is this critical value of $(width)$ manufacturable? Using the above condition with $\ell_{max} = r_{max}d$ and $d = (k\chi\bar{r})^{1/(e-1)}(width)$, we can show that for the longest line not to be propagation limited, the scale must be reduced down to

$$(width) \leq 16\rho\epsilon v r_{max}(k\chi\bar{r})^{1/(e-1)} \tag{2.36}$$

which is plotted in Figure 2.14 for $e < n$ along with the corresponding cell size d. These $(width)$ values are certainly manufacturable. Somewhat different considerations apply if we employ the nonuniform width distribution; nevertheless, the qualitative behavior remains unchanged, inductive effects need not be considered for large N.

Two other factors may be an impediment to downscaling. One is the size d_d of the elements. Since d may not be less than d_d, the critical value of d_d below which inductive effects need not be considered may be directly determined from Figure 2.14, from which we see that this does not become a problem for large N. The other is heat removal requirements, which will be considered below.

2.5.3 Heat Removal

A. Two Dimensions

Heat removal has no effect on the relations between B and N derived in the preceding subsection, which are scale invariant. If the system can be downscaled sufficiently so that the longest line is not propagation limited (i.e. $T_p \leq T$), it does not have any effect on S either. In 2 dimensions, this is often possible. Since the energy dissipated per bit along a line of length ℓ is $2\epsilon V^2 \ell$, the heat removal condition becomes

$$Qd^2 \geq k(2\epsilon V^2 \bar{\ell})B = 2\epsilon V^2 k\bar{r}dB \tag{2.37}$$

$$d \geq \frac{2\epsilon V^2}{Q} k\bar{r}B \tag{2.38}$$

where $\bar{\ell} = \bar{r}d$. Let us assume that $T_d \to 0$ and assume the maximum possible bit repetition rate is employed, as given by (2.29). Thus, we may calculate the minimum value of d as set by heat removal. For $e < n$ we find

$$d = \frac{2\epsilon V^2}{Q(16\rho\epsilon)}(k\kappa)^{-1} M^2 N^{-p-1/2} \tag{2.39}$$

which quickly drops below the critical cell size presented in Figure 2.14 for voltages $V \sim 1$ Volt and the modest $Q = 1\,\mathrm{W/cm^2}$. Thus, with increasing N, heat removal is not a limiting factor in 2 dimensions.

Figure 2.15 provides a comparison of the S versus N curves for optical and normally conducting interconnections. We assume d_d and T_d to be small so as to push the element limited regimes as far as possible to the left. Based on the discussion of the preceding paragraph, we assume V^2/Q is small enough to enable the scale to be reduced to the extent that $T_p < T$ on the longest line. Thus $S = B$ for the normally conducting case. The curves for the optical case are in terms of the parameter EB/Q, as in Figure 2.6.

These curves do not provide a fair comparison, as they assume that B is kept constant with increasing N for optical interconnections, whereas it must be involuntarily decreased for normal conductors. Thus in Figure 2.16 we set B to the largest possible value allowed by normal conductors with nonuniform widths. Remember that our analysis of the nonuniform width case was optimistic in the sense that it was strictly applicable only in the limit of arbitrarily small manufacturable linewidths.

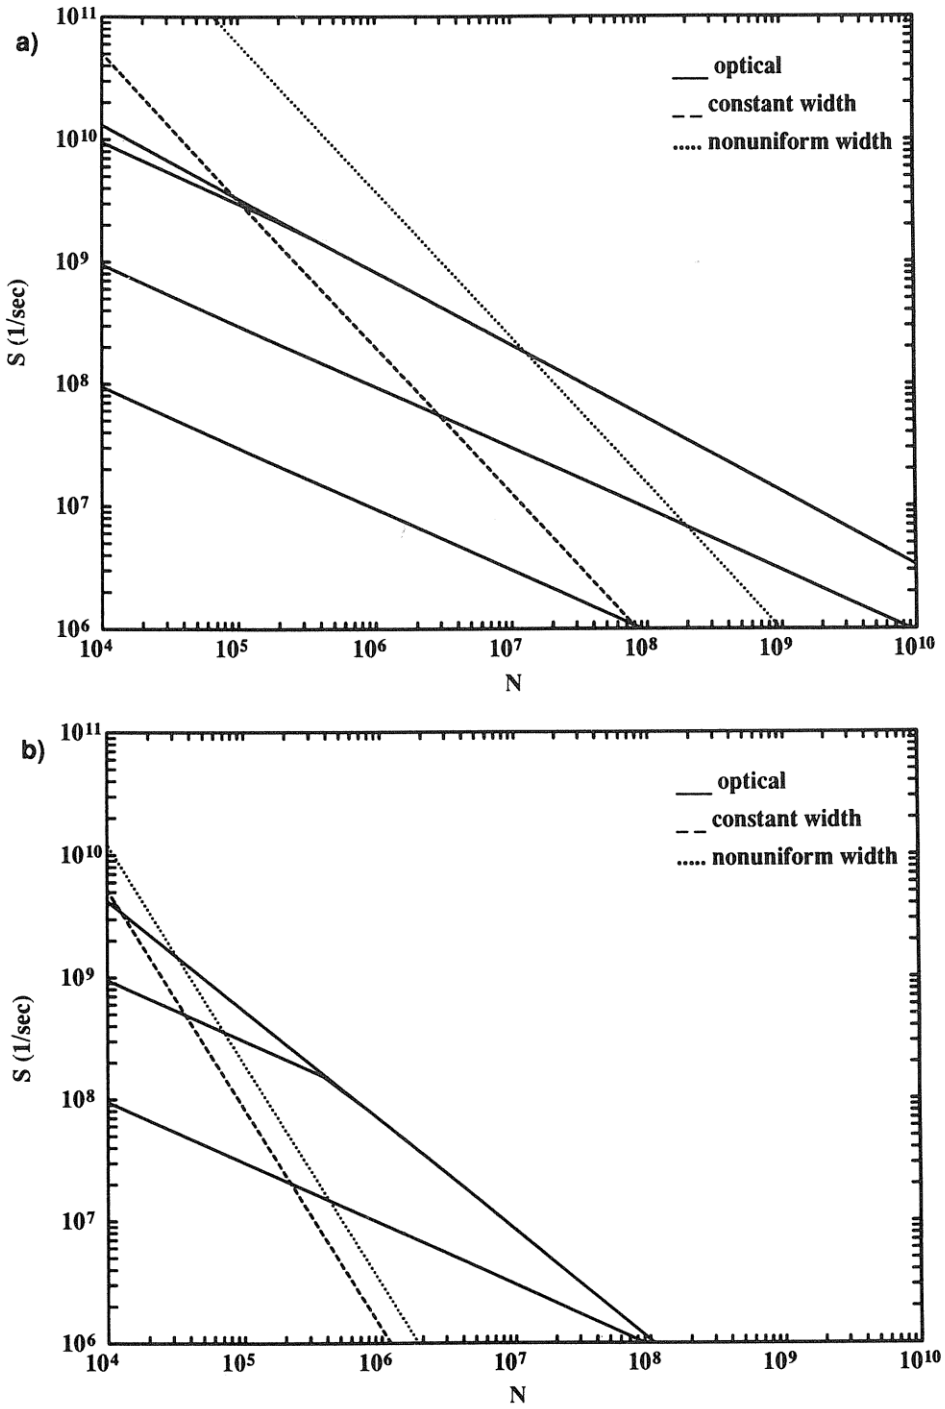

Figure 2.15: Comparison of optical and normally conducting interconnections. We take $k = 10$ and assume d_d, T_d and T_r to be small so as to isolate their effects. Heat removal is not considered for normal conductors. We take $M = 10$ for normal conductors and $M = 1$ for the optical layout. Both the constant width and nonuniform width cases are shown for the normally conducting case. For the optical case, EB/Q fans from $10^{-12}\,\mathrm{m}^2$ to $10^{-4}\,\mathrm{m}^2$ in increments of $10^{-2}\,\mathrm{m}^2$, as in Figure 2.6. (a) $e = 2$, $p = 0.6$. (b) $e = 2$, $p = 0.9$.

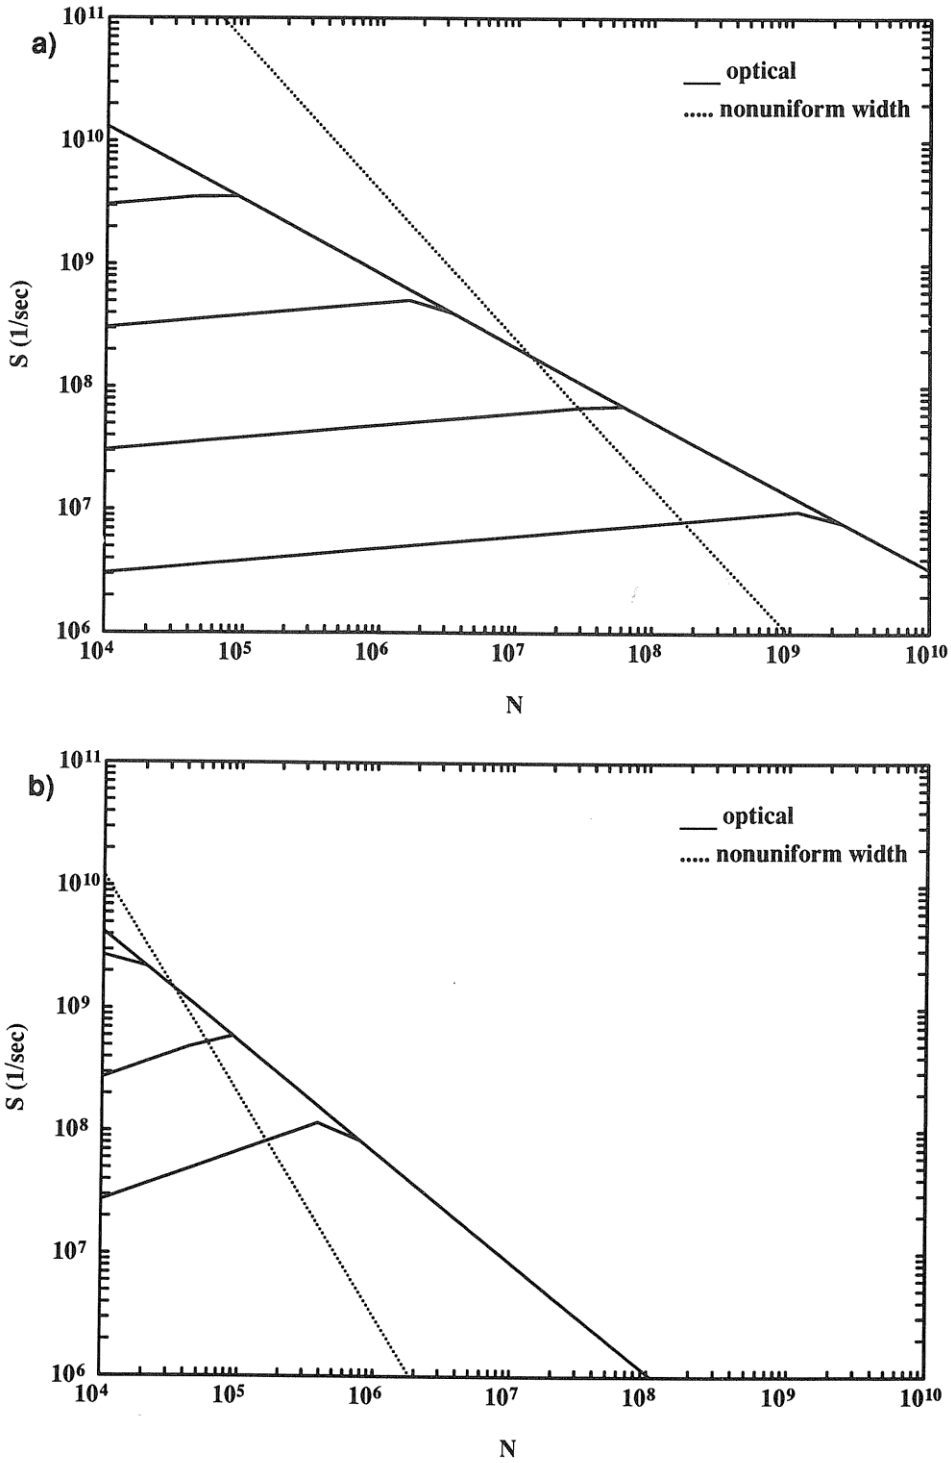

Figure 2.16: Comparison of optical and normally conducting interconnections for the same values of B. We assume the nonuniform width distribution for normal conductors. Similar parameter values are taken as in the previous figure. E/Q fans from 10^{-21} m^2 sec to 10^{-13} m^2 sec in increments of 10^{-2} m^2 sec. (a) $e = 2$, $p = 0.6$, (b) $e = 2$, $p = 0.9$.

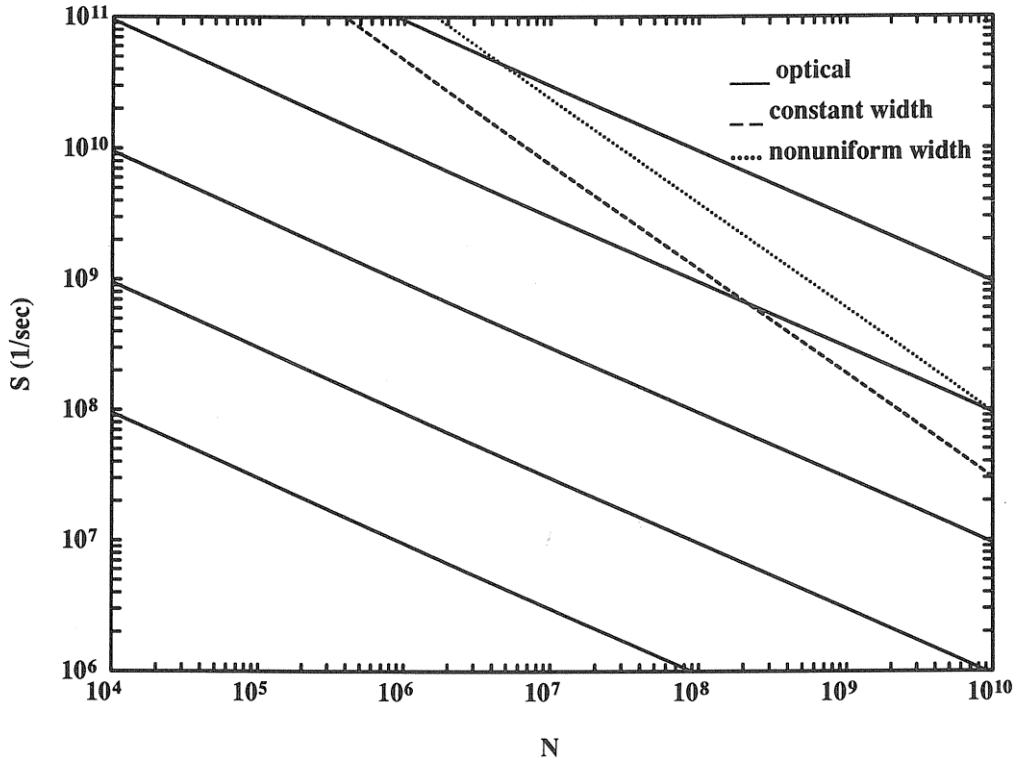

Figure 2.17: Comparison of optical and normally conducting interconnections in 3 dimensions. We take $k = 10$, $e = 3$, and $p = 0.8$ and assume d_d, T_d and T_r to be small. Heat removal is not considered for normal conductors. For the optical case, EB/Q fans from $10^{-12}\,\mathrm{m}^2$ to $10^{-4}\,\mathrm{m}^2$ in increments of $10^{-2}\,\mathrm{m}^2$, as in Figure 2.9.

In general we observe that there is a critical system size beyond which optical communication offers superior performance over normal conductors. Normal conductors are beneficial for small systems since the linewidths can be reduced much below than ever possible with optical lines. However, with increasing system size and line lengths, we must either i.) keep linewidths constant and suffer quadratic increase of delay; or ii.) increase linewidths so as to keep attenuation at an acceptable level and maintain linear growth rate of delay with length, once again resulting in quadratic growth rate of delay with system size (since the growth of line lengths are compounded by the increase in linewidths). Optical communication has the advantage of enabling us to keep the effective communication cross section constant with increasing system size [46].

For large p, certain multi-facet holographic optical architectures [12] [48], as we have discussed in an earlier section, exhibit similar behavior to

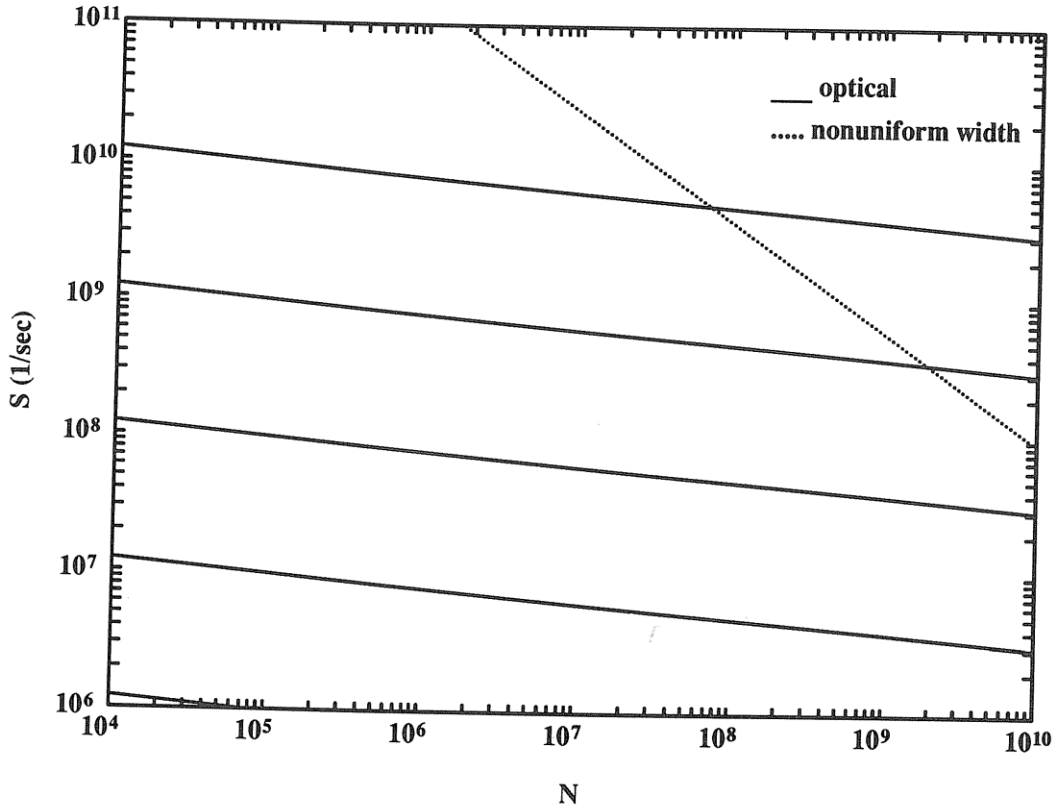

Figure 2.18: Comparison of optical and normally conducting interconnections for the same values of B in 3 dimensions. Similar assumptions as in the previous figure are made. E/Q fans from 10^{-21} m^2 sec to 10^{-13} m^2 sec in increments of 10^{-2} m^2 sec.

the fully 2 dimensional case we have considered. Thus we see that despite their inhibitive nature, these architectures are superior to 2 dimensional normally conducting layouts for large N.

The heat removal limited normally conducting case is of historical importance and has been subject to many previous studies. For instance, see [59] [60] [61] [62] [29].

B. Three Dimensions

We saw that for typical parameter values, 2 dimensional layouts may be downscaled to the extent that inductive effects need not be considered on the longest line (i.e. $T_p < T$) so that S is given by $1/\max(1/B, T_d)$. This may not be possible for 3 dimensional layouts. A complete analysis of this situation is somewhat involved and thus will be treated at length elsewhere. For room temperature voltages and for the range of N in consideration, it may be the case that the cell size need be greater than the values given in part b.) of Figure 2.14. Thus, the value of S may be quite worse than

$\min(B, 1/T_d)$. Of course, for ever increasing values of N, heat removal will eventually cease to be an issue since B must necessarily drop, and, wiring requirements become more stringent.

By reducing the temperature, we may reduce V to an extent that the system may be downscaled sufficiently. Then, a similar comparison as for the 2 dimensional case is possible and is presented in Figures 2.17 and 2.18.

The advantage of optical communication is greater when p is large and/or the dimension e is low.

2.5.4 Asymptotic Properties

In this subsection we assume arbitrarily fast devices ($T_d \to 0$), negligible element size d_d and arbitrarily small manufacturable linewidths. As N increases heat removal ceases to be a problem. Thus, the scale of the system may be reduced sufficiently so that inductive effects need not be considered and $S = B$ as given by (2.29).

Equation 2.29 immediately leads to an important conclusion: For given N, there is an upper limit to B. This is in contrast to the optical case where B could be arbitrarily increased by suffering a reduction in S. Any attempt at increasing B by using wider lines or $\chi > 1$ parallel channels is thwarted by the increase in line lengths, since $T_\ell \propto \ell^2/(width)^2$.

Let us also investigate the dependence of the $HS = HB$ product on N. We find

$$HS = HB \propto N^{p(e-3)/(e-1)} \tag{2.40}$$

which cannot be improved by increasing N since $e \leq 3$. In fact, since fully 3 dimensional circuits are very difficult to realize, these quantities will often decrease with increasing N. Thus, once N is large enough to bring us into the interconnect dominated regime for which our analysis is applicable, further improvement in these products is not possible. The use of normal conductors is inhibitive for applications for which these products are a suitable figure of merit.

2.6 Repeatered Normally Conducting Interconnections

2.6.1 Physical Model Description

The inhibitive square law behavior of normal conductors may be alleviated with the use of repeater structures. In our treatment, we will consider a highly idealized situation. We will not address the issue of power distribution to the repeating devices and not be concerned with the discrete

structure of such lines, treating them as if they were a continous structure. We will assume repeatered lines are used for all connections, although the optimal number of stages for the shortest connections will be less than unity. We also assume that the system is large enough so that the longest line requires more than one stage. Bakoglu [63] derived the optimal configurations of such interconnections for the lumped case (i.e. when inductive effects need not be considered). Within numerical factors the optimal number of stages ξ and the resulting delay is given by

$$\xi \simeq \sqrt{\frac{RC\ell^2}{R_0C_0}} \tag{2.41}$$

$$T_p \simeq \sqrt{R_0C_0RC\ell^2}\,. \tag{2.42}$$

In this section T_p denotes the time it takes a single bit to ripple through the ξ stages, rather than an electromagnetic propagation delay as in earlier sections. R_0C_0 is the intrinsic delay of the repeaters. Following similar arguments as in the preceding section, our model equations may be derived as

$$\xi \simeq 4\sqrt{\frac{\rho\epsilon}{R_0C_0}}\,\frac{\ell}{(width)} \tag{2.43}$$

$$T_r = T \simeq T_\ell \simeq T_d \simeq R_0C_0 \tag{2.44}$$

$$T_p = \xi T \simeq 4\sqrt{R_0C_0\rho\epsilon}\,\frac{\ell}{(width)} \tag{2.45}$$

$$E \simeq 2\epsilon V^2\ell. \tag{2.46}$$

As before, the numerical factors are crude. We are assuming bits may be pipelined through each line at a rate of one every T seconds, the time it takes one bit to traverse a single stage. Depending on R_0C_0, the value of T may be low enough to challenge other high bandwidth approaches. Most importantly, it is independent of other line parameters. The optimum value of ξ is proportional to $\ell/(width)$; the number of stages increases linearly with distance. The length of each stage depends only on $(width)$ and may be found to be

$$\ell_{stage} = \frac{\ell}{\xi} = \frac{1}{4}\sqrt{\frac{R_0C_0}{\rho\epsilon}}\,(width) \tag{2.47}$$

which together with $(width)^2 \leq 16\rho\epsilon v\ell_{stage}$ dictates that approximately $(width) \leq 4\sqrt{\rho R_0C_0/\mu}$ is necessary for inductive effects not to be con-

Table 2.2: Repeatered interconnection model.

	$(width) \leq 4\sqrt{\frac{\rho R_0 C_0}{\mu}}$	$(width) \geq 4\sqrt{\frac{\rho R_0 C_0}{\mu}}$
delay, τ	$4\sqrt{R_0 C_0 \rho\epsilon}\ \frac{\ell}{(width)}$	$\sqrt{\mu\epsilon}\ \ell$
pulse width, T	$R_0 C_0$	$R_0 C_0$
energy, E	$2\epsilon V^2 \ell$	$8\epsilon V^2 \sqrt{\frac{\rho R_0 C_0}{\mu}}\ \frac{\ell}{(width)}$
termination	no	yes

sidered. This evaluates to about $(width) < 5\,\mu$m for $R_0 C_0 = 100\,$psec and our usual choice of physical parameters. If $(width)$ is set to $4\sqrt{\rho R_0 C_0/\mu}$, we obtain $T_p \simeq \sqrt{\mu\epsilon}\ \ell$. If $(width)$ is greater than this value, each stage will become propagation limited and the delay will still be given by this expression independent of $(width)$. When this is the case, we will agree to individually terminate each repeater stage. The energy for this repeatered transmission case may be calculated in a similar manner as for repeaterless transmission. We multiply the number of stages ξ by $V^2 T/Z_0$, the energy per stage. Thus, we find that the energy per transmitted bit is given by

$$E \simeq \min\left(2\epsilon V^2 \ell, 8\epsilon V^2 \sqrt{\frac{\rho R_0 C_0}{\mu}}\ \frac{\ell}{(width)}\right). \tag{2.48}$$

The aspect ratio of a single stage is given by

$$\ell_{stage}/(width) = (1/4)\sqrt{R_0 C_0/\rho\epsilon}$$

which evaluates to ~ 2500 for $R_0 C_0 = 100\,$psec. If $(width) \sim 1\,\mu$m, all lines up to a few millimeters will be single stage and longer ones multistage.

We finally inquire whether it is safe to ignore the space occupied by the repeaters in comparison to the wires. For concreteness, let us consider CMOS VLSI repeaters. The optimum transistor strength (and hence area) was derived to be $s = \sqrt{R_0 C/C_0 R} = \sqrt{R_0/C_0}\sqrt{\epsilon/\rho}\sqrt{t/h}\ w$ times that of a minimum sized transistor [63]. Remember that we took $t = h$ and $w = (width)/2$. With $R_0 = 20\,\text{K}\Omega$ and $C_0 = 5\,$fF consistent with $R_0 C_0 = 100\,$psec, and $(width) \sim 1\,\mu$m, we find $s \sim 35$ which we compare with $\ell_{stage}/(width)$ found above. Thus the space occupied by the drivers may be absorbed into that occupied by the wires with little error.

Our model is summarized in Table 2.2 and Figure 2.19.

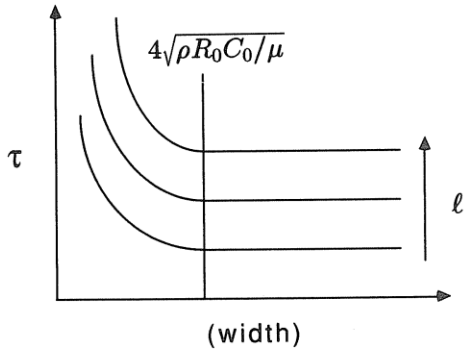

Figure 2.19: Repeatered interconnection model.

2.6.2 Relations between S, B and N

First assume that element size and heat removal need not be considered. As in the optical case, $T \simeq R_0 C_0$ is a constant independent of other line parameters. Thus if $BT \leq 1$ we set $\chi = 1$. Then, for $e < n$, using $d^{e-1} = k\kappa N^{(n-e)/ne}(width)^{e-1}$ and

$$T_p = 4\sqrt{R_0 C_0 \rho\epsilon}\,\frac{N^{1/e} d}{(width)} \tag{2.49}$$

we obtain

$$SN^{(n-1)/[n(e-1)]} = (4\sqrt{R_0 C_0 \rho\epsilon}\,)^{-1}(k\kappa)^{-1/(e-1)}. \tag{2.50}$$

where we assumed $T_p \geq T$ so that $S = 1/T_p$. Of course, S may actually never exceed $1/T$. This relation is similar to the corresponding relation for optical communication in form (equation 2.11 with $\chi = 1$), despite being numerically inferior. The relation between S, B and N when $BT > 1$ is also similar to that derived previously for the optical case (equation 2.13).

Upon comparison with the coefficient of (2.11), we see that repeaters are worse than optical communication by less than a factor of 10, assuming $R_0 C_0 = 100$ psec and an optical wavelength of about a micron. Thus, if fast devices are available, we may approach the performance offered by optical communication within an order of magnitude. Such a system may be more compact than the corresponding optical system, if deep submicron scaling is employed.

Notice that (2.50) is scale independent, assuming $(width)$ is small enough so that inductive effects need not be accounted for ($\leq 5\,\mu$m for $R_0 C_0 =$ 100 psec). When element size is accounted for, S is given by the minimum predicted by (2.50) and $1/S = \sqrt{\mu\epsilon}\, N^{1/e} d_d$.

2.6.3 Heat Removal

A. Two Dimensions

Since (2.50) is scale invariant, heat removal has no effect unless it requires that the scale be chosen large enough to lead to inductive effects. The power dissipation per cell is $k\bar{E}B$ where $\bar{E}$ is the average of E given by (2.48) over lines of all lengths. The power dissipation per cell must not exceed Qd^2. Just as in the repeaterless case, we are agreeing to fill up available wiring space by choosing $k\chi\bar{r}(width)/M = d$. Since $\ell/(width) = rd/(width) = rk\chi\bar{r}/M$, we find that we must maintain

$$Qd^2 \geq k\bar{r} \min\left(2\epsilon V^2 d, 8\epsilon V^2 \sqrt{\frac{\rho R_0 C_0}{\mu}} \frac{k\chi\bar{r}}{M}\right) B \tag{2.51}$$

so that d must be at least

$$d = \min\left(\frac{2\epsilon V^2 k\bar{r}B}{Q}, \left(\frac{8\epsilon V^2\sqrt{\rho R_0 C_0/\mu}}{Q}\right)^{\frac{1}{2}} \frac{\chi^{\frac{1}{2}} k\bar{r}B^{\frac{1}{2}}}{M^{\frac{1}{2}}}\right). \tag{2.52}$$

This is to be compared with the critical value $d = (k\chi\bar{r}/M)\, 4\sqrt{\rho R_0 C_0/\mu}$ below which inductive effects need not be considered. If the first term is less than the second, we can show that d is less than the mentioned critical value. Then, (2.50) is applicable. If the second term is less than the first, we find that d is greater than the mentioned critical value. Then

$$\frac{1}{S} = T_p = \sqrt{\mu\epsilon}\, N^{\frac{1}{2}} d \tag{2.53}$$

where d is given by the preceding equation (which is dominated by its second term in this case). In general, S is given by the smaller predicted by (2.50) and (2.53). The resulting dependence of S on N is presented in Figure 2.20.

Let us take a closer look at the dependence of the heat removal limited value of d on the various parameters. First of all note that $d \propto \bar{r}$, a direct consequence of the fact that the energy always increases with line length. Assuming $B \geq 1/T$ so that $\chi = BT$, we may reexpress d as follows

$$d = k\bar{r}B \min\left(\frac{2\epsilon V^2}{Q}, \left(\frac{8\epsilon V^2\sqrt{\rho R_0 C_0/\mu}\, R_0 C_0}{QM}\right)^{\frac{1}{2}}\right) \tag{2.54}$$

which we may simply write as $d = k\bar{r}B(cons.)$. For $Q/V^2 = 10\,\mathrm{W/cm^2 volt^2}$

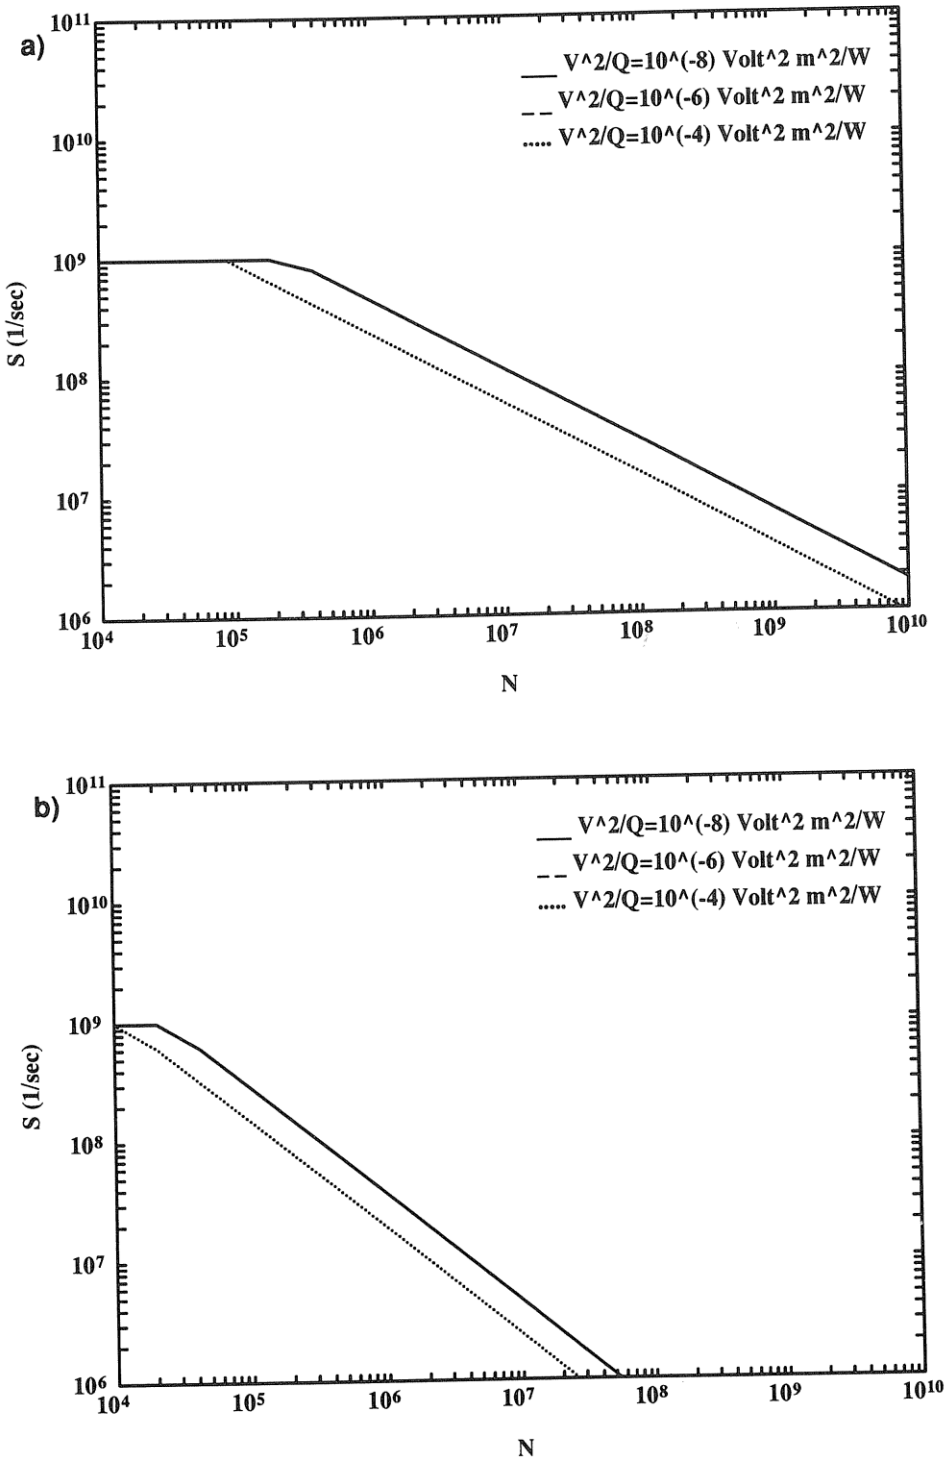

Figure 2.20: S versus N for $B = 1\,\text{Gbit/sec}$ for repeaters in 2 dimensions. We take $k = 10$, $M = 10$ and $T = R_0 C_0 = 1\,\text{nsec}$ so that $\chi = 1$. d_d is assumed to be small enough to have no effect. The curves corresponding to the two smaller values of V^2/Q overlap. (a) $e = 2$, $p = 0.6$, (b) $e = 2$, $p = 0.9$,

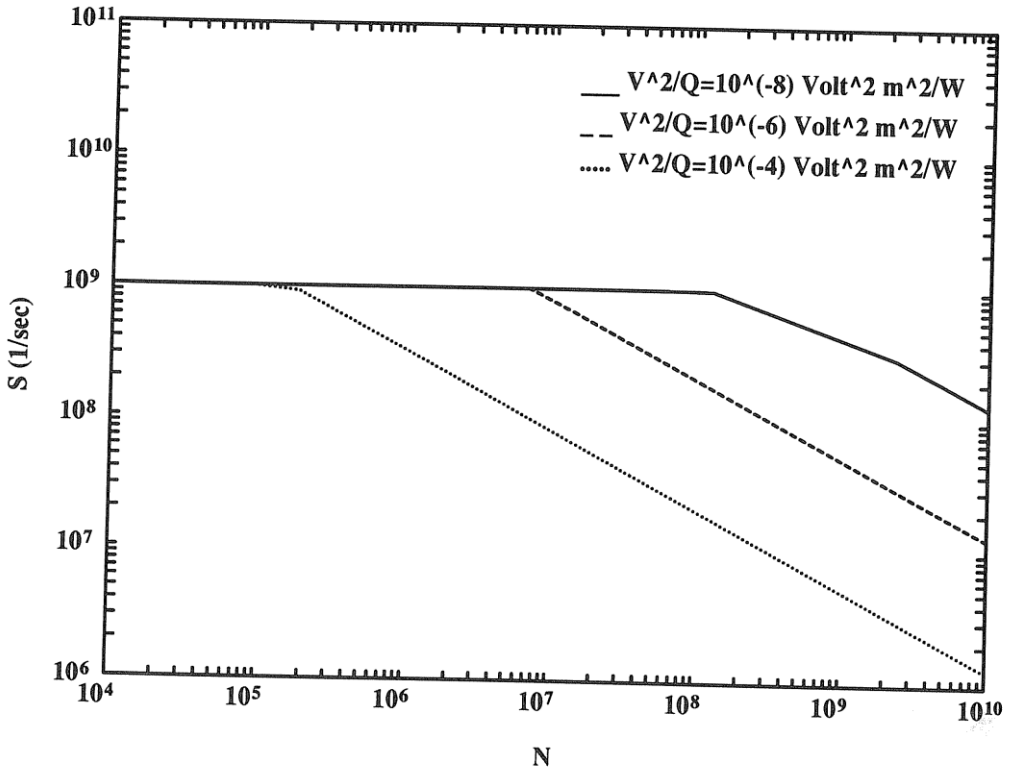

Figure 2.21: S versus N for $B = 1$ Gbit/sec for repeaters in 3 dimensions. We take $k = 10$, $e = 3$, and $p = 0.8$ and $T = R_0C_0 = 1$ nsec so that $\chi = 1$. d_d is assumed to be small enough to have no effect.

and our usual parameters, (*cons.*) may be approximately expressed as $\simeq \min(0.7, 3.5\,(R_0C_0)^{3/4}/M^{1/2})$(fm sec) where R_0C_0 is in nsec. Thus, if R_0C_0 is not small or M is not large we will most likely be operating in the lumped regime so that (2.50) is applicable. The growth rate of d as imposed by heat removal is $\propto \bar{r}B$ whereas wiring requirements dictate $d \propto \chi\bar{r}$ which is also $\propto \bar{r}B$ when $BT \geq 1$. Thus which mechanism will dominate depends on numerical factors, in contrast to the optical case where heat removal requirements were always overshadowed by communication requirements with increasing N.

B. Three Dimensions

We repeat the analysis of the preceding section. Now the minimum cell size is found to be

$$d = \min\left(\frac{2\epsilon V^2 k\bar{r}N^{\frac{1}{3}}B}{Q}, \left(\frac{8\epsilon V^2\sqrt{\rho R_0C_0/\mu}}{Q}\right)^{\frac{1}{2}}\chi^{\frac{1}{4}}(k\bar{r})^{\frac{3}{4}}N^{\frac{1}{6}}B^{\frac{1}{2}}\right). \quad (2.55)$$

Thus the resulting delay is the greater of $T_p = \sqrt{\mu\epsilon}\, N^{1/3} d$ and that given by (2.50) and is illustrated in Figure 2.21.

2.6.4 Asymptotic Properties

The asymptotic behavior of repeatered lines are similar to those presented for optical interconnections, when heat removal is not considered. It remains the same for 2 dimensional layouts even when heat removal *is* considered. However, the situation is worse when heat removal is considered in 3 dimensions.

For the 3 dimensional case we refer to (2.55). For large N and B, the second term in this equation will be applicable so that $d \propto \bar{r}^{3/4} B^{3/4} N^{1/6}$ as opposed to $d \propto \bar{r}^{1/2} B^{1/2}$ dictated by wiring requirements. Thus the resulting growth rate of signal delay becomes $\tau \propto \bar{r}^{3/4} B^{3/4} N^{1/2}$ as opposed to $\propto B^{1/2} N^{1/2}$ possible with optical communication. For given B, the growth rate of the bisection-inverse delay product is then found to be (for $e < n$)

$$HS \propto N^{p/4} \tag{2.56}$$

which is inferior to the optical $HS \propto N^{p-1/2}$ in its range of applicability ($e < n$ or equivalently $p > 2/3$).

If we do not terminate each stage of the repeaters individually and charge up the segments, then the first term in (2.55) becomes applicable so that $\tau \propto \bar{r} B N^{2/3}$. In this case we find that for given B, the bisection-inverse delay product can*not* be increased with increasing N, an inhibiting situation.

2.7 Superconducting Interconnections

2.7.1 Physical Model Description

The propagation delay and characteristic impedance of a superconducting transmission line are essentially given by [56]

$$T_p = \sqrt{\mu\epsilon} \left[\frac{\lambda}{h} \coth\left(\frac{t}{\lambda}\right) + 1 \right]^{\frac{1}{2}} \ell \tag{2.57}$$

$$Z_0 = \sqrt{\frac{\mu}{\epsilon}} \frac{h}{w} \left[\frac{\lambda}{h} \coth\left(\frac{t}{\lambda}\right) + 1 \right]^{\frac{1}{2}} \tag{2.58}$$

where all parameters are defined as previously, except λ which denotes the superconducting penetration depth throughout this section. We again refer to Figure 2.11 and invoke similar geometrical constraints. Attenuation and dispersion are small enough to be safely ignored for the length scales in consideration [17]. Thus, just as in the optical case, we assume that the minimum temporal pulse width $T = \max(T_d, T_\ell) = T_d$ is set by device limitations in transmission mode. During lumped operation, T_ℓ will correspond to the rise time of the output end voltage and may or may not be greater than T_d. As with normally conducting lines, the minimum pulse repetition interval is just $T_r = T$, so that we drop the subscript r. We again assume perfect termination in transmission mode is possible.

Throughout our analysis, we will use as an example the high critical temperature superconductor Ba-Y-Cu-O with $T_c = 92.5\,\mathrm{K}$, absolute penetration depth $\lambda_0 = 1400\,\text{Å}$, normal resistivity $\rho_n = 200\,\mu\,\Omega\,\mathrm{cm}$ operated at 77 K [17]. The value of the penetration depth at $T = 77\,\mathrm{K}$ may be calculated as $\lambda = \lambda_0/\sqrt{1-(77/92.5)^4} = 1942 \simeq 2000\,\text{Å}$. We will assume these materials to have standard superconducting behavior below critical current, critical field, critical temperature and energy gap frequency.

In order to maintain desirable superconducting behavior, both the flux entry field and the critical current density should not be exceeded. If a surface barrier to flux entry is not present and breakdown at edges can be neglected, the flux entry field is just H_{c1}, the lower critical field of the superconductor. If the thickness of the conductor is larger than the penetration depth, current only flows through a sheet of thickness λ. The maximum surface current density that can be allowed before vortices enter the superconductor is just $J_{sc} = H_{c1}$. A value of $J_{sc} = 8\,\mathrm{mA}/\mu\mathrm{m}$ was estimated [15] based on earlier experimental results (The author extrapolated a critical field value of 500 Oersted measured at a lower temperature to 100 Oersted at 77 K). Kwon et al. [17] estimated $50\,\mathrm{mA}/\mu\mathrm{m}$ for low temperatures based on the same data. In general, a few hundred Oersteds seems to be a value which one might reasonably expect to achieve. When the penetration depth exceeds the conductor thickness it is preferable to speak of a volume critical current density J_c. Based on intuitive grounds, we would expect J_c to satisfy $J_c\lambda \sim J_{sc}$. Indeed, with $\lambda = 2000\,\text{Å}$, the above mentioned values for J_{sc} are consistent with often cited values for the volume critical current density ($J_c = 10^6$-$10^7\,\mathrm{A/cm^2}$). However, for films this thin, edge effects become increasingly important so that one must be careful in interpreting the physical origin of J_c and the implications of our simple model.

The energy per transmitted pulse is given by $E = V^2 T/Z_0$ with $T = T_d$. Thus (2.57) and (2.58)lead to

$$T_p E = \epsilon V^2 T \frac{w}{h} \ell. \tag{2.59}$$

This product does not depend on $(width) = 2w$ or the cross sectional $(area) = (width)(height) \simeq 4wh$, but only on the ratio (w/h). For any given cross section, it is optimal to set w/h to its smallest value of 2, leading to $T_p E = 2\epsilon V^2 T \ell$. A knowledge of T_p directly leads to a knowledge of E and vice versa.

First assume that $h \geq t \geq \lambda$. Within factors close to unity,

$$T_p = \sqrt{\mu\epsilon}\, \ell = \frac{\ell}{v} \tag{2.60}$$

$$Z_0 = \frac{1}{2}\sqrt{\frac{\mu}{\epsilon}} \tag{2.61}$$

$$E = 2V^2 \sqrt{\frac{\epsilon}{\mu}}\, T. \tag{2.62}$$

Since $(width) = 2w$ and we take $h = w/2$, the condition $t \geq \lambda$ may be expressed as $(width) \geq 4\lambda(h/t)$. By choosing $t = h$, the region of validity of the above equations may be extended down to $(width) = 4\lambda$.

So that the critical current density is not exceeded we require $V/Z_0 \leq J_{sc} w$. Using the above expression for Z_0, this condition can be expressed as $4V/(J_{sc}\sqrt{\mu/\epsilon}\,) \leq (width)$. If the critical current is high enough and the operating voltage low enough this condition is less restrictive than $(width) \geq 4\lambda(h/t)$. Even for $t = h$, with presently achievable critical currents as cited above, voltage values somewhat less than 1 Volt are sufficient to ensure this. One expects much lower voltage values to be used at these low temperatures. Also, we might expect materials with even higher critical current densities to be produced. Hence we will assume $\lambda > V/(J_{sc}\sqrt{\mu/\epsilon})$ throughout our analysis. This means that we need not be concerned with the critical current density in this regime of operation.

Now, let us consider the case $t \leq h \leq \lambda$. Again within numerical factors close to unity,

$$T_p = \sqrt{\mu\epsilon}\, \lambda\, \frac{\ell}{h^{\frac{1}{2}} t^{\frac{1}{2}}} \tag{2.63}$$

$$Z_0 = \sqrt{\frac{\mu}{\epsilon}}\,\lambda\,\frac{h^{\frac{1}{2}}}{wt^{\frac{1}{2}}}. \tag{2.64}$$

The critical current condition can now be written as $V/Z_0 \leq J_c wt$ which translates into $V/(J_c\lambda\sqrt{\mu/\epsilon}\,) \leq h^{1/2}t^{1/2}$, making it desirable to choose h and t as large as possible. If this condition *is* violated, we can use a sufficiently large drive impedance $R_d > Z_0$ to limit the current and charge up the line

$$R_d = \frac{V}{J_c wt} > Z_0. \tag{2.65}$$

The lumped delay and energy are then expressed as $T_\ell = R_d C\ell$ [5] and $E = V^2 C\ell$,

$$T_\ell = \frac{\epsilon V}{J_c}\frac{\ell}{ht} = \frac{16\epsilon V}{J_c}\frac{\ell}{(width)^2}\frac{h}{t} \tag{2.66}$$

$$E = 2\epsilon V^2 \ell \tag{2.67}$$

where $(width) = 2w = 4h$ was used. We immediately observe from the above that $t = h$ is the optimal choice, leading to $T_\ell = (16\epsilon V/J_c)\ell/(width)^2$ in the region $(width) < 4V/(J_{sc}\sqrt{\mu/\epsilon})$. If the critical current condition is *not* violated, then using (2.63) and 2.64 the delay, characteristic impedance and energy are expressed as

$$T_p = \frac{4\lambda}{v}\frac{\ell}{(width)}\left(\frac{h}{t}\right)^{\frac{1}{2}} \tag{2.68}$$

$$Z_0 = \frac{1}{2}\sqrt{\frac{\mu}{\epsilon}}\,\frac{4\lambda}{(width)}\left(\frac{h}{t}\right)^{\frac{1}{2}} \tag{2.69}$$

$$E = \frac{V^2 T}{Z_0} \tag{2.70}$$

where $v = 1/\sqrt{\mu\epsilon}$ was used. We previously showed $T_p E = 2\epsilon V^2 T\ell$. For given ℓ, any pair of T_p and E compatible with this equation determines the value of $(width)(t/h)^{1/2}$. Thus to minimize $(width)$ we choose $t = h$, leading to

$$T_p = \frac{4\lambda}{v}\frac{\ell}{(width)} \tag{2.71}$$

$$Z_0 = \frac{1}{2}\sqrt{\frac{\mu}{\epsilon}}\,\frac{4\lambda}{(width)} \tag{2.72}$$

Table 2.3: Superconducting interconnection model when $T_d \leq T_p$ (or $T_d \leq T_\ell$ in the lumped case). Delay ($\tau = \max(T_p, T)$), pulse width (T), energy (E), and termination ("ter") are given for various ranges of (width).

	$(width) \leq \frac{4V}{J_{sc}\sqrt{\mu/\epsilon}}$	$\frac{4V}{J_{sc}\sqrt{\mu/\epsilon}} \leq (width) \leq 4\lambda$	$(width) \geq 4\lambda$
τ	$\frac{16\epsilon V\lambda}{J_{sc}} \frac{\ell}{(width)^2}$	$\frac{4\lambda}{v} \frac{\ell}{(width)}$	$\frac{\ell}{v}$
T	T_ℓ	T_d	T_d
E	$2\epsilon V^2 \ell$	$2\sqrt{\frac{\epsilon}{\mu}}\, V^2 \frac{(width)}{4\lambda} T_d$	$2\sqrt{\frac{\epsilon}{\mu}}\, V^2 T_d$
ter.	no	yes	yes

$$E = \frac{V^2 T}{Z_0} \tag{2.73}$$

valid in the region $4V/(J_{sc}\sqrt{\mu/\epsilon}) \leq (width) \leq 4\lambda$.

Finally, we consider the case $t \leq \lambda \leq h$. If $\lambda^2 \leq ht$, then a very similar analysis as for the case $h \geq t \geq \lambda$ applies. If $\lambda^2 \geq ht$, then a very similar analysis as for the case $t \leq h \leq \lambda$ applies. In both cases we again find that $t = h$ is the best choice (or at least as good as anything else) so that this case collapses.

Thus we agree to set $t = h = w/2$. In some cases, a smaller value of t may do just as well, but since this can improve (*height*) by at most a factor of 2, we will not be overlooking any significant room for improvement in this direction. As remarked before, due to the arbitrariness of our constraints, the actual optimum dimensions may be somewhat, but not greatly different from these. Table 2.3 and Figure 2.22 summarize our superconducting model when $T_d \leq T_p$ ($T_d \leq T_\ell$ for the lumped case), i.e. when T_d has no effect in determining the delay. Of course, the delay may never actually be less than T_d. Note that it is suboptimal to work with $T_d > T_p$ (or $T_d > T_\ell$ in the lumped case), since we can reduce $(width)$ until $T_d = T_p$ (or $T_d = T_\ell$), ending up with a line occupying less space with the same delay.

2.7.2 Relations between S, B and N

As was with normal conductors, we agree to choose $(width)$ so that the condition $d^{e-1} \geq k\chi\bar{r}(width)^{e-1}$ is always satisfied with equality. We will mostly be at an advantage (because of the inverse dependence of τ on $(width)$ for given ℓ), and never at a disadvantage by doing so (within the limits of our abstraction).

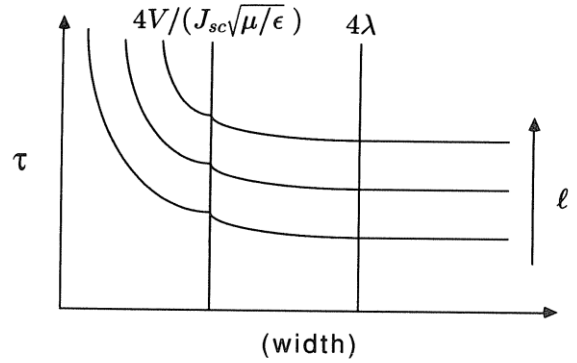

Figure 2.22: Superconducting interconnection model when $T_d \leq T_p$ (or $T_d \leq T_\ell$ in the lumped case).

We refer to Figure 2.22. If d_d is small and heat removal need not be considered, a moments reflection reveals that it is optimal to work in the intermediate region, assuming we can manufacture $(width) \leq 4\lambda$. To see this, simply notice that as we scale the system photographically, ℓ varies in linear proportion to $(width)$. The key quantity to be calculated is $\ell_{max}/(width)$. If $B \leq 1/T_r = 1/T = 1/T_d$ so that $\chi = 1$, then for $e < n$

$$\frac{\ell_{max}}{(width)} = r_{max}(k\kappa)^{1/(e-1)} N^{(n-e)/[ne(e-1)]}. \tag{2.74}$$

Since in this region $T_p = (4\lambda/v)\ell/(width)$, and assuming T_d is small so that $\tau = T_p$, we find, with $r_{max} \simeq N^{1/e}$

$$SN^{(n-1)/[n(e-1)]} = \left(\frac{v}{4\lambda}\right)(k\kappa)^{-1/(e-1)}. \tag{2.75}$$

This relation is independent of the specific choice of $(width)$, as long as it lies between $4V/(J_{sc}\sqrt{\mu/\epsilon})$ and 4λ. Can a nonuniform distribution of linewidths help? 4λ is already less than a micron. Unless we can manufacture $(width)$ less than $0.1\,\mu$m or so, there is not much room for variation, *even if* we assume $4V/(J_{sc}\sqrt{\mu/\epsilon})$ to be small. Thus we do not consider this case.

Notice that this relation is identical in form with (2.11) derived for optical interconnections. The relation for $BT_d \geq 1$ is likewise similar to (2.13). If $\epsilon_r \simeq 4$ for the superconducting interconnections and optical wavelengths ($\lambda \sim 0.5\,\mu$m) are utilized for the optical interconnections, the relations become numerically identical within a factor of 2.

We stress that despite the similarity of the above relation to the corresponding optical relation, the superconducting system may be smaller in size. The reduction in line lengths is precisely cancelled by the inverse

dependence of the delay on $(width)$, so that there is no performance advantage. The potential advantage in terms of cost of area is limited by how small we can manufacture $(width)$ and how much the voltage can be reduced and/or critical current increased.

2.7.3 Heat Removal

A. Two Dimensions

Heat removal has no effect on performance unless it requires that d be large enough that $(width) \geq 4\lambda$. Wiring requirements dictate $d \geq k\chi\bar{r}(width)/M$ whereas heat removal dictates $d \geq (kEB/Q)^{1/2}$. Thus with increasing N, heat removal is not a problem in 2 dimensions. For finite values of N, heat removal may require $(width)$ to be larger than 4λ. The analysis is very similar to the optical case; in fact, Figure 2.6 is approximately applicable to superconductors as well, provided we interpret $E = 2V^2\sqrt{\epsilon/\mu}\, T_d$. If low voltage values are used, this energy can be much less than ever achievable with optical interconnections.

When heat removal or element size is not a limiting factor, reducing $(width)$ also reduces ℓ_{max}, keeping $\ell_{max}/(width)$ and the system delay constant. This not only reduces system size, as mentioned earlier, but results in less total energy consumption (since $Z_0 \propto 1/(width)$), making it desirable to choose $(width)$ as small as possible. However, this increases power dissipation per unit area. Thus, if we use the expression for energy $E = 2\sqrt{\epsilon/\mu}\; V^2 T_d$ valid in the region $(width) \geq 4\lambda$ and find that heat removal requires that the scale be large enough that $(width) \geq 4\lambda$, we know that we are not excluding any room for improvement in the $(width) \leq 4\lambda$ region. When heat removal allows $(width) \leq 4\lambda$, S does not depend on heat removal parameters anyway and is given by (2.75).

B. Three Dimensions

This case is likewise similar to the corresponding optical case, with the above remarks in mind. We only need interpret Figure 2.9 with the appropriate superconducting energy E.

2.7.4 Asymptotic Properties

The asymptotic behavior is identical to that of optical interconnections whether heat removal is considered or not and is not repeated.

2.8 Discussion

In this work we have emphasized certain basic physical considerations which we believe are major factors limiting the computational capacity of large scale processing systems. Apart from a multitude of engineering issues [20], one must also be aware of other physical limitations which we have not considered. For instance, it may be that power distribution is the major limiting factor, especially if resistive wires are employed.

Of the many implementation related issues and limitations we have not considered, we briefly mention a few. Satisfactory termination of transmission lines may prove to be very difficult, putting optical systems at an advantage. We are still far from being able to construct fully 3 dimensional systems with conductors. In fact, we do not know of a fully 3 dimensional optical architecture that can provide an arbitrary pattern of interconnections [46]. The performance of stacked 2 dimensional or other 'quasi' 3 dimensional layouts would lie between the fully 2 dimensional and fully 3 dimensional cases we have considered. Presently we are far from being able to realize waveguide circuits approaching the diffraction limited $(width) \simeq \lambda$ we have assumed. The construction of efficient, reliable and small size transducers is another major difficulty with optical interconnections. A mature thin film technology for superconducting interconnections is still to be developed. On the other hand, the use of laminated conductors promises improvement (even if by only a constant factor) for normal conductors [64]. Another issue we did not directly account for is that of fan-out. Architectures involving large fan-outs tend to favor optical communication [11].

We have exclusively concentrated on highly interconnected systems characterized by large p (or equivalently n). Not all applications require such a system. For instance, to add a million pair of numbers, all we need is a million adders. However, it is generally recognized that the solution of many interesting problems, including those that are often loosely associated with human 'intelligence', require greater degrees of communication between the primitive elements.

In this work we have assumed p to be constant throughout the system hierarchy. More generally, p may be a function of hierarchical level. Although it is possible to extend our analysis to this more general case, here we have not attempted to do so as this greatly complicates the analysis without contributing any additional understanding.

Perhaps our most important reservation regards the underlying paradigm

of computation inherent in our models, which is essentially related to the way electronic digital computers have been traditionally built. For instance, we are assuming the energy associated with the transmission of each bit of information to be irreversibly dissipated. For some applications, this need not be the case at all [66]. Although our results are applicable when the primitive elements are optical switches, we must bare in mind that the way we process information with optical computing systems may be quite different than with a digital electronic computer.

2.9 Towards Unifying Physical and Algorithmic Approaches

We desire to solve problems of ever increasing size, despite the fact that the human life span is more or less constant. Let us say we are interested in solving a problem of certain size in a certain finite amount of time. We will construct our processing system by assembling together a set of primitive elements with given function. We will agree on a certain procedure in which the number of elements N may be increased by introducing new elements to the system in a useful way, and on how the computation is to be performed (i.e. the algorithm). In conjunction, we will agree on a family of connection graphs, one for each value of N, with interconnectivity p.

It is possible to an extent to trade off between the three quantities S (or S_{ave}), B and N in solving a given problem. For instance, it may be possible to solve a given problem in a given amount of time with a small yet fast system, or alternatively with a large yet slower system. In general, the set of all possible triplets (S, B, N) which will enable us to solve the given problem in the given amount of time will define a region in S-B-N space, satisfying the following property: if (S_0, B_0, N_0) is an element of this region, so is $(\iota_1 S_0, \iota_2 B_0, \iota_3 N_0)$, where $\iota_i \geq 1$. This region may be described as $\Psi(S, B, N) \geq C_\Psi$. We will speak of this region as the Ψ-region and the surface defining this region as the Ψ-surface. Simple examples of such considerations are the area-time bounds of VLSI complexity theory [18].

These *lower* bounds should be interpreted in conjunction with the *upper* bounds of the form $\Phi(S, B, N) \leq C_\Phi$ derived in this paper, for which the terms Φ-regions and Φ-surfaces will be used. If there exists a triplet compatible with both bounds, we will be able to solve the given problem in the given amount of time with the given interconnection media.

By comparing the regions $\Phi(S, B, N) \leq C_\Phi$ and $\Psi(S, B, N) \geq C_\Psi$ it is not only possible to decide whether a given interconnect medium is capable of handling given problems, but also to determine the appropriate choice

of S, B and N. Figure 2.23 illustrates the various possibilities, where, for simplicity in illustration we assume that B is not involved in the tradeoff. Figure 2.23a illustrates a situation where the Φ-surface completely lies below the Ψ-surface. This interconnect medium is not capable of performing the prespecified task in the given amount of time. It is necessary to relax the Ψ-surface by increasing the time allowed for computation until a point of intersection is reached (Figure 2.23b). Thus it is possible to determine the minimum time in which the task may be performed. There will be a certain value of N for which this minimum time value can be achieved. Figure 2.23c shows a situation where we have the flexibility of choosing S and N from a finite interval where the technology curve exceeds the requirement curve.

We now see how we may compare various interconnect media characterized by the functional forms $\Phi(S, B, N) \leq C_\Phi$ with reference to the computational requirements of a given problem-algorithm. We illustrate this in Figure 2.24 where we have again assumed that B is not involved. Figure 2.24a shows the upper bounds for two different interconnect media. The curve lying to the upper right is superior to the other regardless of problem requirements and operating point. Figure 2.24b illustrates a situation where one medium is superior to the other for N greater than a

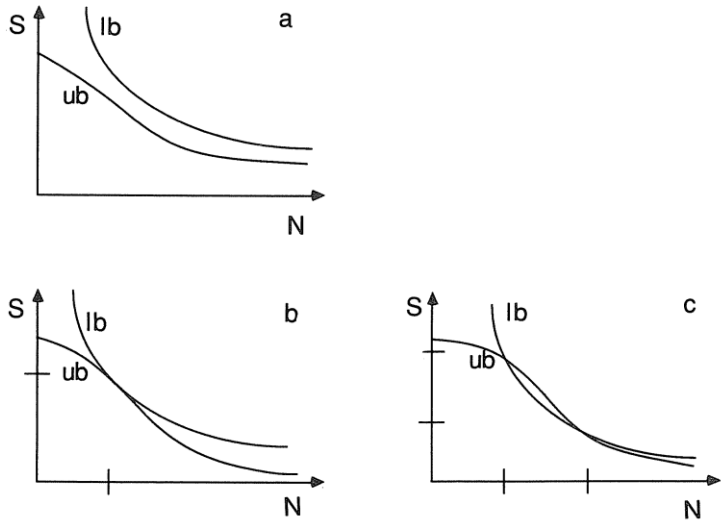

Figure 2.23: Ψ-surfaces versus Φ-surfaces. (a) illustrates a situation where there are no common points among the regions defined by the lower and upper bounds. In (b), the requirement on computation time has been relaxed so that a point of intersection is obtained. In (c) we are free to choose from a range of possible values of S and N.

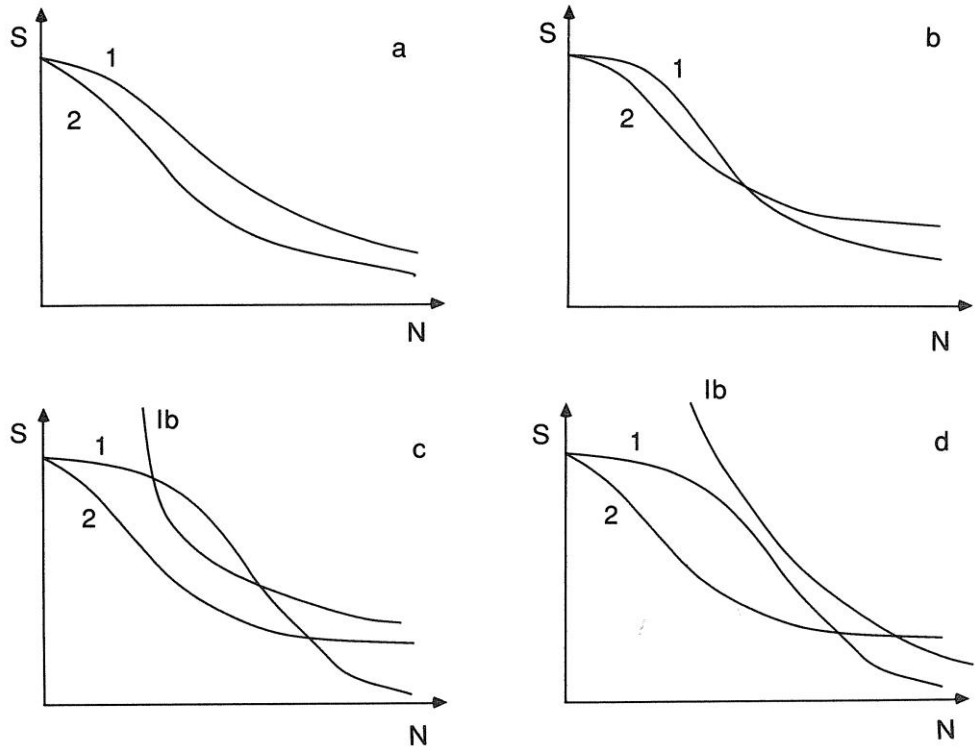

Figure 2.24: Comparison of interconnect media. (a) illustrates the Φ-surfaces associated with two different interconnect media, one of which is superior to the other. In (b), again the surfaces of two media are shown, one is superior for values of N larger than a critical value, the other for lower values. (c) shows a requirement curve which has an intersection with only the curve belonging to the first media, whereas (d) shows one which intersects only the second.

certain critical value. Neither, one, or both may be able to perform the stated task in the stated amount of time. Figure 2.24c and d illustrate two different requirement curves, one of which has an overlap with only the first technology curve, the other with the second.

The use of Ψ-surfaces is only one of many possible ways to characterize the computational requirements of a problem-algorithm, but one which we believe is especially suitable for interfacing the algorithmic and physical aspects of computation.

Hillis [68] has noted the disparity between traditional abstractions of computing systems and their physical implementation. For instance, much of the literature on parallel computation is based on models which may not be possible to directly implement [18]. He has argued in favor of a *physical* theory of computation. Such a theory should take into account basic properties of the universe, such as the impossibility of action at a

distance, and in the case of dissipative systems, limitations associated with heat removal.

The mentioned abstractions have been inherited from a time when all computing systems were device limited. Although this is no longer true, there is still a tendency to consider wires to be mere parasitics degrading the intended performance of the devices they interconnect [5]. With increasing system sizes, it is probably more appropriate to base our expectations of intended performance on interconnect oriented models and consider device limitations as parasitics degrading these expectations. This has been the philosophy of this paper.

(Parallel) algorithms must be developed in conjunction with their physical implementation. As an example, we consider the work of Feldman et al. [69]. They consider the problem of matrix vector multiplication, which may be solved in parallel on many graph topologies. They, however, introduce a new family of graphs on which this problem may be solved which also have an efficient optical implementation due to their space invariant properties.

In this paper we limited ourselves to an examination of the behavior of S (i.e. inverse signal delay) as a function of N for constant B and to an examination of the asymptotic properties of the bisection-inverse delay and bisection-bandwidth products. Based on the discussion of this section, the reader will realize that these are mere examples, which we have chosen for their simplicity and general interest. A complete treatment employing the formalism of this section is beyond the scope of this paper.

2.10 Summary and Conclusions

In this work we quantified the communication requirements of a processing system by its interconnectivity (Rent exponent) p, or equivalently the fractal dimension n. The distribution of line lengths in our system obeys an inverse power law distribution, with p as a parameter.

We derived physical models of interconnection media which we believe approximately represent the best achievable. We provided a unified description of terminated and unterminated cases for conducting lines which inherently account for major physical mechanisms such as the skin effect, superconducting penetration depth and critical current limitations. We showed that it is preferable to use a single wide normally conducting line rather than many narrow ones, as long as the pulse width is not device limited. The interconnection models presented in this work enable results

which are more physically realistic than those of VLSI complexity theory, which are based on physically naive models.

Combining our system and physical models, we derived relations of the form $\Phi(S, B, N) \leq C_\Phi$ for each interconnection technology. These relations bound the largest simultaneously possible values of S (inverse signal delay), B (bit repetition rate) and N (number of elements). We presented an abstract formalism enabling us to relate these bounds to the computational requirements of given applications.

In discussing the limitations of conducting interconnections, we allowed arbitrarily small scaling and arbitrarily fast devices. We saw that normal conductors, whether unterminated or terminated, did not allow B to be kept constant with increasing system size. Both B and S were found to sharply decrease with increasing N. Making longer lines wider leads to improvement by only a constant factor and does not change the asymptotic dependence on N. The bisection-inverse delay and bisection-bandwidth products were found to be bounded from above. This is in contrast with the other technologies with which it is possible to arbitrarily increase B and the bisection-bandwidth product for any given N, by suffering a decrease in S.

If repeater structures employing ultrafast devices are possible, the performance for 2 dimensional layouts may approach that possible with the other technologies we have considered within an order of magnitude. For large system sizes, they will still be more costly in terms of energy consumption. In 3 dimensions, repeaters are inferior to optical and superconducting technologies 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 very similar performance for same dimensional layouts and similar switching 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 line width, once conductor thickness drops below the penetration depth. Optical interconnections may enable a 3 dimensional layout and freedom from termination problems. On the other hand, superconductors may offer much lower energies, especially if the voltage level is reduced. In 2 dimensions, this leads to improved performance for only a limited range of N, since wiring density becomes more important than heat removal as N increases. In 3 dimensions however, this enables lower signal delay for given N and B.

In this work we compared the ability of given technologies in providing

communication between a given array of elements. Elsewhere, we have discussed how these technologies may be used in conjunction to achieve performance not possible with any alone [70] [71].

Acknowledgements

We gratefully acknowledge the benefit of discussions with Prof. A. El Gamal of the same laboratory, Prof. F. Pease and B. Langley of the Solid State Electronics Laboratory and Prof. M. Beasley of the Applied Physics Department. They are however not responsible for any possible errors.

This work was supported by the Air Force Office of Scientific Research under Grant No. AFOSR-88-0024.

A Appendix

A.1 Extension to Fan-Out and Fan-In

One can model the equivalent of a fan-out or fan-in situation using only pairwise links. This is depicted in Figure 2.25b. The layout minimizing connection length is shown in Figure 2.25a. Notice that the inefficiency in using pairwise links is bounded between 1 (when two target elements are located as in Figure 2.25d) and $1/F$ (when the target elements are as in Figure 2.25c), where F is the maximum fan-out or fan-in. Thus our analysis based on the total interconnection length $kN\bar{\ell}$ may be modified by the introduction of an appropriate average factor $1/F \leq \eta_F \leq 1$, if specific characteristics of such an architecture are specified. For a discussion of the effects of fan-out and fan-in on energy, the reader is referred to [72] [11].

A.2 Coefficients for the Moments of g(r)

The coefficients appearing in (2.4) are given by

$$\zeta_m = \frac{me}{(m - e/n)(e - e/n + m)} \tag{2.76}$$

$$\zeta'_m = \frac{1}{n}$$

$$\zeta''_m = \frac{1}{1 - mn/e}.$$

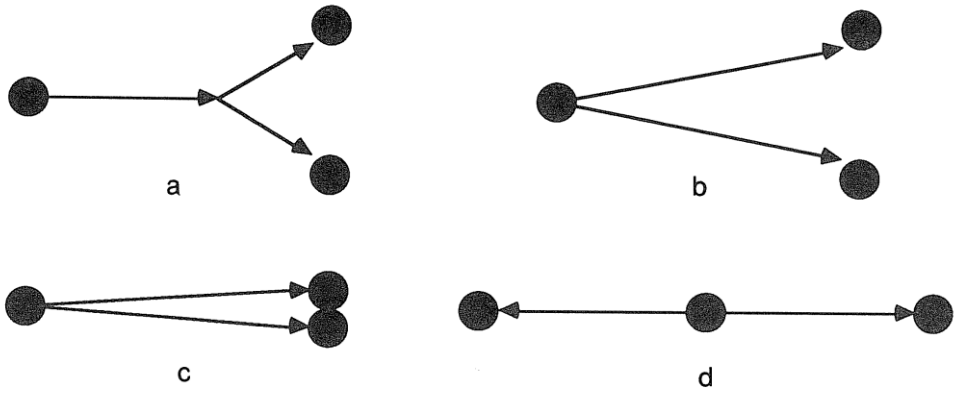

Figure 2.25: Extension to fan-out and fan-in. (a) shows optimal power splitting which minimizes the total connection length. (b) illustrates the same connections wired using pairwise connections only. (c) and (d) show the two extreme cases illustrating the bounds on the inefficiency incurred by using only pairwise interconnections.

The major approximation made in the derivation of (2.4) is to ignore 1 with respect to $N^{m/e-1/n}$ when $e < mn$ and vice versa when $e > mn$. Thus, if $N^{m/e-1/n}$ is at least ~ 2 when $e < mn$ (or at most $\sim 1/2$ when $e > mn$), our error is less than about a factor of 2. If $m/e - 1/n \sim 0$, it is more appropriate to use the logarithmic dependence. In this work only the first, second and third moments are used. Since we restrict ourselves to highly interconnected systems for which $n > e$, the condition $mn > e$ is always satisfied.

As an example, let us calculate the values of $\kappa = \zeta_1$ for the special case $n = \infty$ ($p = 1$). We find 2/3 and 3/4 for $e = 2$ and $e = 3$ respectively. Assuming a cartesian metric, the exact values of these coefficients for a square (or cubic) grid are 2/3 and 1.

A.3 Three Dimensional Optical Layouts Where the Elements Are Confined to a Plane

Here we consider the case where the elements are constrained to lie on a planar $N^{1/2} \times N^{1/2}$ grid, as in the fully 2 dimensional layout, but the 'wires' are allowed to leave the plane. Let the system be confined to a square prism of volume $N^{1/2}d \times N^{1/2}d \times M(2\lambda)$. That is, we are measuring the height of the system in units of (2λ) and denoting it by M. If a sandwich of planar waveguides is used, M may be interpreted as the number of layers. Unlike in the fully 2 dimensional case, where M was specified as a constant, here we will be free to choose M as large as we wish. The contributions of the vertical runs will be taken into consideration.

Due to finite element size d_d and heat removal considerations, d must satisfy the conditions $d \geq d_d$ and $d^2 \geq kEB/Q$.

The horizontal contribution to the total interconnection length is just $kN\chi\bar{r}d$ as before. For the vertical contribution, let us first consider the worst case, that all signals must travel up and down a total length of $2M(2\lambda)$. This would be the case (within factors like $\sqrt{2}$ etc.) if communication is established by a hologram or other reflective imaging system located $M(2\lambda)$ above the device plane. Thus by multiplying the total connection length by $(2\lambda)^2$ [46] we find the total volume needed for communication to be $kN\chi(\bar{r}d(2\lambda)^2 + 2M(2\lambda)^3)$. Requiring that the total available communication volume $Nd^2M(2\lambda)$ exceed this, we obtain, in addition to previous requirements,

$$d^2 \geq \frac{k\chi(2\lambda)}{M}(\bar{r}d + 2M(2\lambda)) \tag{2.77}$$

$$d \geq \frac{k\chi(2\lambda)\bar{r}}{M} + (2k\chi)^{\frac{1}{2}}(2\lambda). \tag{2.78}$$

We immediately notice that the second term may be ignored, if the transducers are restricted to a single planar layer on the surface of the elements, as would almost always be the case in practice. We cannot expect the transducers to be packed denser than one per $(2\lambda)^2$ so that $d^2 \geq d_d^2 \geq k\chi(2\lambda)^2$. Thus, the condition $d \geq d_d$ already covers this requirement (within a factor of $\sqrt{2}$). Now we remember that we had assumed the worst case for the contribution of the vertical runs. The vertical contribution, which we saw can be ignored even in the worst case, may actually be much less.

Thus including all considerations, the minimum value of d is given by

$$d = \max\left[\left(\frac{kEB}{Q}\right)^{\frac{1}{2}}, \frac{k\chi\bar{r}(2\lambda)}{M}, d_d\right]. \tag{2.79}$$

In passing, we notice that it is of no utility to choose M to be any greater than $k\chi\bar{r}(2\lambda)/\max[(kEB/Q)^{1/2}, d_d]$.

Assuming the propagation delay $T_p \geq T$, the signal delay $\tau = T_p$ may now be expressed as

$$\tau = \frac{1}{c}\left[N^{\frac{1}{2}}d + 2M(2\lambda)\right], \tag{2.80}$$

again assuming worst case contribution of the vertical runs. For the moment assuming that d is given by $d = k\chi\bar{r}(2\lambda)/M$, the optimum value of M

minimizing the delay is found as $M = (N^{1/2}k\chi\bar{r}/2)^{1/2}$. Of course, we never need set M to a value greater than that mentioned in our passing remark above. We then find, within a factor of two that the resulting delay is given by (for $p > 1/2$)

$$\frac{1}{S} = \tau = \frac{1}{c}\max\left[N^{\frac{1}{2}}\left(\frac{kEB}{Q}\right)^{\frac{1}{2}}, (k\chi\kappa)^{\frac{1}{2}}N^{\frac{p}{2}}(2\lambda), N^{\frac{1}{2}}d_d\right] \tag{2.81}$$

where $\chi = \max(1, BT_r)$ and $d_d^2 \geq k\chi d_{trans}^2$ where d_{trans} denotes the extent of a transducer. The reader will notice that apart from the last term, this equation is very similar to (2.18).

Unless $p = 1$, the second term falls behind with increasing N. Let us consider the case $p = 1$ and rewrite the above equation for $BT_r \geq 1$ as

$$\frac{1}{S} = \frac{1}{c}(kNB)^{\frac{1}{2}}\max\left[\left(\frac{E}{Q}\right)^{\frac{1}{2}}, T_r^{\frac{1}{2}}\lambda, T_r^{\frac{1}{2}}d_{trans}\right] \tag{2.82}$$

where we ignored all numerical factors and assumed that $d_d^2 \sim k\chi d_{trans}^2$, i.e. the element size is transducer limited. Of course, the second term is redundant since $d_{trans} \geq \lambda$. Thus, we conclude that, given that the elements are to be arrayed on a planar surface, circuits with $p = 1$ do not lead to greater delay than those with smaller p, since the system is transducer surface limited [46] anyway. Another conclusion is that, when $p = 1$, only a constant increase in delay (by at most a factor of $d_{trans}/2\lambda$) is incurred by constraining the elements to lie on a plane, instead of a 3 dimensional grid.

As noted earlier, our results are valid for a system employing optical switching if one interprets E as the switching energy.

The reader is referred to [12] for further discussion of situations where the elements are constrained to lie on a plane.

A.4 Signal Delay for VLSI Circuits

In this work we have attempted to derive models of conducting interconnections which represent the basic limitations of the wires and which are independent of device technology. For an alternate approach, the reader is referred to [3] [4]. Essential to our analysis is the assumption that T_d, the intrinsic delay of the switching devices can be specified as a given constant independent of line length. In practice T_d may be coupled to T_ℓ and/or may depend on whether the line is terminated or not.

As an example, let us consider the rise time $\tau = T$ of an unterminated VLSI line driving a very high impedance load

$$T \simeq R_d C_d + (R_d + R\ell)C\ell = R_d(C_d + C\ell) + T_\ell \tag{2.83}$$

where C_d is the drive capacitance and we have replaced $RC\ell^2 = T_\ell$. Now, if we can argue that the first term can be kept constant independent of line length, we may define $T_d = R_d(C_d + C\ell)$ and thus write the total delay in the form $T = T_d + T_\ell \simeq \max(T_d, T_\ell)$.

The term $R_d(C_d + C\ell)$ can be kept constant, as argued by Thompson. We simply agree to increase the driver size in proportion to ℓ, thus reducing $R_d \propto 1/\ell$ and increasing $C_d \propto \ell$. Since the area occupied by the wire also increases, the area of the driver can always be absorbed in the area of its wire [73].

Such arguments may not always be possible. However, remember that it is suboptimal to work with lines which satisfy $T_d > T_\ell$ and that in any case, we are mainly interested in small values of T_d. If T_d is *very* small ($<$ 10-100 psec or so), dispersion and attenuation models may have to be introduced for optical and superconducting interconnections.

References

[1] K. C. Saraswat and F. Mohammadi, *Effect of scaling of interconnections on the time delay of VLSI circuits,* IEEE Transactions on Electron Devices, vol. 29, pp. 645–650 (1982).

[2] D. S. Gardner, J. D. Meindl, and K. C. Saraswat, *Interconnection and electromigration scaling theory,* IEEE Transactions on Electron Devices, vol. 34, pp. 633–643 (1987).

[3] H. B. Bakoglu, *Circuit and System Performance Limits on ULSI: Interconnections and Packaging,* PhD thesis, Stanford University, Stanford, California (1986).

[4] H. B. Bakoglu, *Circuits, Interconnections and Packaging for VLSI,* Addison-Wesley, Reading, Massachusetts (1990).

[5] T. A. Schreyer, *The Effects of Interconnection Parasitics on VLSI Performance,* PhD thesis, Stanford University, Stanford, California (1989).

[6] J. W. Goodman, F. J. Leonberger, S-Y. Kung, and R. Athale, *Optical interconnections for VLSI systems,* Proc. IEEE, vol. 72, pp. 850–866 (1984).

[7] W. D. Hillis, *New computer architectures and their relationship to physics or why computer science is no good,* Int. J. Theoretical Physics, vol. 21, pp. 255–262 (1982).

[8] R. W. Keyes, *Communication in computation,* Int. J. Theoretical Physics, vol. 21, pp. 263–273 (1982).

[9] A. C. Hartmann and J. D. Ullman, *Model categories for theories of parallel systems,* in *Parallel Computing: Theory and Experience,* G. J. Lipovski and M. Malek (Eds), John Wiley and Sons (1986).

[10] R. K. Kostuk, J. W. Goodman, and L. Hesselink, *Optical imaging applied to microelectronic chip-to-chip interconnections,* Applied Optics, vol. 24, pp. 2851–2858 (1985).

[11] 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,* Applied Optics, vol. 27, pp. 1742–1751 (1988).

[12] M. R. Feldman, C. C. Guest, T. J. Drabik, and S. C. Esener, *Comparison between electrical and free space optical interconnects for fine grain processor arrays based on interconnect density capabilities,* Applied Optics, vol. 28, pp. 3820–3829 (1989).

[13] P.R. Haugen, S. Rychnovsky, A. Husain, and L. D. Hutcheson, *Optical interconnects for high speed computing,* Optical Engineering, vol. 25, p. 1076 (1986).

[14] D. A. B. Miller, *Optics for low-energy communication inside digital processors: Quantum detectors, sources and modulators as efficient impedance converters,* Optics Letters, vol. 14, pp. 146–148 (1989).

[15] R. C. Frye, *Analysis of the trade-offs between conventional and superconducting interconnections,* IEEE Circuits and Devices Magazine, pp. 27–32 (May 1989).

[16] H. Kroger, C. Hilbert, U. Ghoshal, D. Gibson, and L. Smith, *Applications of superconductivity to packaging,* IEEE Circuits and Devices Magazine, pp. 16–21 (May 1989).

[17] 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 Letters, vol. 8, pp. 582–585 (1987).

[18] J. D. Ullman, *Computational Aspects of VLSI,* Computer Science Press, Rockville, Maryland (1984).

[19] C. E. Leiserson, *Area-Efficient VLSI Computation,* The MIT Press, Cambridge, Massachusetts (1983).

[20] R. R. Tummala and E. J. Rymaszewski (Eds), *Microelectronics Packaging Handbook,* Van Nostrand Reinhold, New York, New York (1989).

[21] C. Berge, *The Theory of Graphs,* Wiley, New York (1962).

[22] W. E. Donath, *Placement and average interconnection lengths of computer logic,* IEEE Transactions on Circuits and Systems, vol. 26, pp. 272–277 (1979).

[23] B. S. Landman and R. L. Russo, *On a pin versus block relationship for partitions of logic graphs,* IEEE Transactions on Computers, vol. 20, pp. 1469–1479 (1971).

[24] R. L. Russo, *On the tradeoff between logic performance and circuit-to-pin ratio for LSI,* IEEE Transactions on Computers, vol. 21, pp. 147–153 (1972).

[25] W. E. Donath, *Stochastic model of the computer logic design process,* Technical Report RC 3136, IBM Thomas T.J.Watson Research Center, Yorktown Heights, New York (1970).

[26] W. E. Donath, *Equivalence of memory to 'random logic',* IBM Journal of Research and Development, vol. 18, pp. 401–407 (1974).

[27] L. Pietronero, *Fractals in physics: Introductory concepts,* in S. Lundqvist, N. H. March, and M. P. Tosi (Eds), *Order and Chaos in Nonlinear Physical Systems,* Plenum Press, New York (1988).

[28] B. B. Mandelbrot, *Fractals: Form, Chance and Dimension,* W.H. Freeman, San Francisco (1977).

[29] R. W. Keyes, *The Physics of VLSI Systems,* Addison-Wesley, Reading, Massachusetts (1987).

[30] W. E. Donath, *Wire length distribution for placements of computer logic,* IBM Journal of Research and Development, vol. 25, pp. 152–155 (1981).

[31] M. Feuer, *Connectivity of random logic,* IEEE Transactions on Computers, vol. 31, pp. 29–33 (1982).

[32] P. Christie and S. B. Styer, *Fractal description of computer interconnection distributions,* in S. K. Tewksbury and J. Carruthers (Eds), *Microelectronic Interconnects and Packaging: System and Process Integration*, volume 1390. SPIE (1990).

[33] P. Christie, J. E. Cotter, and A. M. Barrett, *Design and simulation of optically interconnected computer systems,* in A. P. DeFonzo (Ed), *Interconnection of High Speed and High Frequency Devices and Systems*, volume 947, pp. 19–24. SPIE (1989).

[34] B. B. Mandelbrot, *The Fractal Geometry of Nature,* W.H. Freeman, New York (1983).

[35] B. B. Mandelbrot, *Information theory and psycholinguistics: A theory of word frequencies,* in P. F. Lazarsfeld and N. W. Henry (Eds), *Readings in Mathematical Social Science,* MIT press, Cambridge, Massachusetts (1968).

[36] B. B. Mandelbrot, *The Pareto-Levy law and the distribution of income,* International Economic Review, vol. 1, pp. 79–106 (1960).

[37] A. C. Hartmann, *Computational metrics and fundamental limits for parallel architectures,* in S. K. Tewksbury (Ed), *Frontiers of Computing Systems Research, Volume 1,* Plenum Press, New York (1990).

[38] A. El Gamal, J. L. Kouloheris, D. How, and M. Morf, *BiNMOS: A basic cell for BiCMOS sea-of-gates,* in *Proceedings of the IEEE Custom Integrated Circuits Conference,* pp. 8.3.1–8.3.4 (1989).

[39] R. W. Keyes, *The wire-limited logic chip,* IEEE Journal of Solid State Circuits, vol. 17, pp. 1232–1233 (1982).

[40] W. R. Heller, W. F. Mikhail, and W. E. Donath, *Prediction of wiring space requirements for LSI,* Journal of Design Automation and Fault Tolerant Computing, vol. 2, pp. 117–144 (1978).

[41] A. El Gamal, *Two-dimensional stochastic model for interconnections in master slice integrated circuits,* IEEE Transactions on Circuits and Systems, vol. 28, pp. 127–134 (1981).

[42] A. Orlitsky and A. El Gamal, *Communication complexity,* in Y. S. Abu-Mostafa (Ed), *Complexity in Information Theory,* Springer-Verlag, New York (1988).

[43] R. Barakat and J. Reif, *Lower bound on the computational efficiency of optical computing systems,* Applied Optics, vol. 26, pp. 1015–1018 (1987).

[44] H. Kogelnik, *Theory of optical waveguides,* in T.Tamir (Ed), *Guided-Wave Optoelectronics,* Springer-Verlag, Berlin Heidelberg (1988).

[45] A.Yariv, *Introduction to Optical Electronics, Second Edition,* Holt, Rinehart and Winston, New York (1976).

[46] H. M. Ozaktas and J. W. Goodman, *Lower bound for the communication volume required for an optically interconnected array of points,* Journal of the Optical Society of America A, vol. 7, pp. 2100–2106 (1990).

[47] R. F. Thompson, *The Brain,* W.H. Freeman and Company, New York (1985).

[48] H. M. Ozaktas, Y. Amitai, and J. W. Goodman, *Comparison of system size for some optical interconnection architectures and the folded multi-facet architecture,* Optics Communications (1991 – accepted for publication).

[49] M R. Feldman and C. C. Guest, *Interconnect density capabilities of computer generated holograms for optical interconnection of very large scale integrated circuits,* Applied Optics, vol. 28, pp. 3134–3137 (1989).

[50] H. M. Ozaktas, H. Oksuzoglu, R. F. W. Pease, and J. W. Goodman, *The effect on scaling of heat removal requirements in 3 dimensional systems,* Submitted to IEEE Electron Device Letters, 1990.

[51] D. B. Tuckerman and R. F. W. Pease, *High performance heat sinking for VLSI,* IEEE Electron Device Letters, vol. 2, pp. 126–129 (1981).

[52] T. L. Michalka, *Models for Wafer Scale Integration Implementation,* PhD thesis, Stanford University, Stanford, California (1988).

[53] P. M. Solomon, *A comparison of semiconductor devices for high-speed logic,* Proc. IEEE, vol. 70, pp. 489–509 (1982).

[54] A. Masaki, *Electrical resistance as a limiting factor for high performance computer packaging,* IEEE Circuits and Devices Magazine, pp. 22–26 (May 1989).

[55] J. T. Watt and J. D. Plummer, *Effect of interconnection delay on liquid nitrogen temperature CMOS circuit performance,* Proceedings of the International Electron Devices Meeting, pp. 393–396 (1987).

[56] R E. Matick, *Transmission Lines for Digital and Communication Networks,* Mc-Graw Hill, New York (1969).

[57] S. Ramo, J. R. Whinnery, and T. Van Duzer, *Fields and Waves in Communication Electronics,* John Wiley and Sons Inc., New York, second edition (1984).

[58] A. J. Blodgett, Jr., *Microelectronic packaging,* Scientific American, vol. 249, pp. 86–96 (July 1983).

[59] R. W. Keyes, *Physical limits in digital electronics,* Proc. IEEE, vol. 63, pp. 740–767 (1975).

[60] R. W. Keyes, *Fundamental limits in digital information processing,* Proc. IEEE, vol. 69, pp. 267–278 (1981).

[61] R. W. Keyes, *The evolution of digital electronics towards VLSI,* IEEE Transactions on Electron Devices, vol. 26, pp. 271–279 (1979).

[62] R. W. Keyes, *A figure of merit for IC packaging,* IEEE Journal of Solid State Circuits, vol. 13, pp. 265–266 (1978).

[63] H. B. Bakoglu and J. D. Meindl, *Optimal interconnection circuits for VLSI,* IEEE Transactions on Electron Devices, vol. 32, pp. 903–909 (1985).

[64] O-K. Kwon, *Chip-to-Chip Interconnections for Very High-speed System-level Integration,* PhD thesis, Stanford University, Stanford, California (1988).

[65] H. H. Zappe, *Josephson quantum interference computer devices,* IEEE Transactions on Magnetics, vol. 13, pp. 41–47 (1977).

[66] H. J. Caulfield and J. Shamir, *Wave particle duality considerations in optical computing,* Applied Optics, vol. 28, pp. 2184–2186 (1989).

[67] J. Hopcroft, W. Paul, and L. Valiant, *On time versus space,* Journal of the Association for Computing Machinery, vol. 24, pp. 332–337 (1977).

[68] W. D. Hillis, *The Connection Machine,* The MIT press, Cambridge, Massachusetts (1985).

[69] M. R. Feldman and C. C. Guest, *Nested crossbar connection networks for optically interconnected processor arrays for vector-matrix multiplication,* Applied Optics, vol. 29, pp. 1068–1076 (1990).

[70] H. M. Ozaktas and J. W. Goodman, *Optimal partitioning of very large scale optoelectronic computing systems,* Optical Society of America 1990 Annual Meeting Technical Digest (1990).

[71] H. M. Ozaktas and J. W. Goodman, *Multiplexed hybrid interconnection architectures,* Proceedings of the 1991 OSA Topical Meeting on Optical Computing (1991).

[72] J. W. Goodman, *Fan-in and fan-out with optical interconnections,* Optica Acta, vol. 32, p. 1489 (1985).

[73] C. D. Thompson, *Area-time complexity for VLSI,* Proceedings of the 11th Annual ACM Symposium on the Theory of Computing, pp. 81–88 (1979).