

A PHYSICAL APPROACH TO COMMUNICATION LIMITS IN COMPUTATION

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF ELECTRICAL ENGINEERING
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

By
Haldun M. Ozaktas
June 1991

© Copyright 1991 by Haldun M. Ozaktas
All Rights Reserved

The authors full name is Haldun Memduh Özaktas. However, to facilitate indexing, please use the above form while referring to this thesis.

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and in quality, as a dissertation for the degree of Doctor of Philosophy.

Joseph W. Goodman
(Principal Advisor)

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and in quality, as a dissertation for the degree of Doctor of Philosophy.

Abbas El Gamal

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and in quality, as a dissertation for the degree of Doctor of Philosophy.

William E. Spicer

Approved for the University Committee on Graduate Studies:

Dean of Graduate Studies

Abstract

We develop a physical approach to modeling and analyzing communication limits in computation. We stress wireability and heat removal requirements as the major limiting mechanisms and explore the interplay between these considerations and the physical properties of optical, normally conducting and superconducting interconnection media, which we characterize in terms of the relations between their length, cross sectional area, delay, bandwidth and energy.

We discuss not only the limitations of each interconnection medium in providing communication among an array of computing elements, but also how they may be used in conjunction to realize hybrid systems exhibiting properties unachievable with any alone. In particular, the optimal combination of optical and normally conducting interconnections maximizing given figure of merit functions is derived. These considerations are extended to discuss the usefulness of optical digital computing and the relative merits of globally and locally connected systems. Other implications for computer architecture are also discussed.

A secondary theme is the study of the capabilities and limitations of optical interconnection architectures. A basic result regarding the limitations of optical wave fields in providing communication among an array of elements is presented and used to evaluate several specific optical interconnection architectures. We propose two architectures which may approach the best possible system size of any 2 and 3 dimensional architecture respectively. We also discuss the role of optical frequencies in balancing information density and heat removal imposed bounds in a dissipative computing environment.

Acknowledgements

I am most grateful to my research advisor Prof. Joseph Goodman for his constant input, support, trust and kindness. Thanks to him, I was able to pursue my doctoral research in an area matched to my background in a manner matched to my abilities.

I would like to extend my thanks to Prof. Abbas El Gamal for acting as my associate advisor. He has contributed to the development of several of the ideas in this thesis through many discussions. Also, I thank Prof. William Spicer, for acting as my academic advisor and reading this thesis, Prof. David Bloom, for being on my thesis defense committee, and Prof. Lambertus Hesselink, for supervising my earlier research at Stanford.

Throughout the course of this research, I have benefited from discussions with several individuals. Prof. Malcolm Beasley and Brian Langley provided me with feedback and suggestions on the superconducting interconnection model discussed in chapter 5. Hakan Oksuzoglu of the Department of Aeronautics and Astronautics helped me develop the heat removal model presented in chapter 3 through several indispensable discussions and also helped check the derivation. The comments and suggestions of Prof. Fabian Pease on the same chapter were invaluable. He also taught me how to shorten an overlength paper. Dr. Yaakov Amitai was a collaborator in devising the folded multi-facet architecture of chapter 7 and its 3 dimensional version and also helped me improve the arguments of chapter 6. The framework described in the latter part of chapter 4 was inspired by discussions with Prof. Abbas El Gamal and Prof. Jeffrey Ullman. I benefited greatly from discussions with Dr. David Miller of Bell Laboratories in maturing the arguments of chapter 12. Prof. Adolf Lohmann of the University of Erlangen-Nürnberg was the first to teach me the concepts of *isomorphic*,

global and *local* methods in a seminar he gave at Stanford several years ago. I would also like to thank Dr. Robert Keyes of the T.J. Watson Research Center for reading parts of my work and offering valuable suggestions and Prof. Stuart Tewksbury of West Virginia University for his interest and support for my work.

Of lasting benefit were courses I took from Profs. Joel Beinin, Ronald Bracewell, Jean-Pierre Dupuy, John Etchemendy, Joseph Goodman, Stephen Harris, Lambertus Hesselink, Janet Johnston, Joel Leivick and Essam Marouf.

I would also like to thank my past and present office and group mates Barry, Bob, Charlotte, Dorothy, Elfar, Hatsumi, Irenea, Jane, Jeff, Jing, Joice, Karin, Larry, Mieko, Mike, Pablo, Rick, Sylvia, Syuji, Takashi, Tim and Yaakov for their companionship and helpfulness throughout my stay in the Information Systems Laboratory.

For fear of leaving someone out, I will not attempt to name the people whose support, help and friendship have kept me alive throughout the past four years. I could not have survived without them. I thank them all for sharing their time with me and bearing with me through my difficult and uncertain times.

For better or for worse, my parents are those mainly responsible for my career choice. Both law school graduates convinced of the meaninglessness of their career, they systematically brainwashed me from early on to be a (physical) scientist or engineer. It is strange to look back and remember the days when I used to read about space exploration, computers, lasers and holography in a colorfully illustrated large format book titled ‘Young person’s science encyclopedia’ or something like that they put in front of me. I would like to thank them and also my brother for their lifelong friendship.

June 1991, Stanford, California

Haldun Özaktas

This thesis was typeset by the author using T_EX[89] with L^AT_EX[101].

Contents

I	An Interconnect Dominated Model of Computation	1
1	Introduction	2
1.1	Background	2
1.1.1	The increasing importance of interconnections	2
1.1.2	Different interconnection media	5
1.2	Purpose	5
1.3	Approach	6
1.4	Overview	9
1.4.1	Grid model	9
1.4.2	Tube model	10
1.4.3	Graph distinctions	11
1.4.4	Problem categories	12
1.4.5	Outline	12
2	Wireability Limitations and Connectivity Model	14
2.1	Description of the model	14
2.2	Justification of the model	18
2.3	Graph layout	22
2.4	Conclusion	23
3	Heat Removal Model	24
3.1	2 dimensional systems	24
3.2	3 dimensional systems	25

3.3	Derivation of the model	25
3.4	Numerical example	29
3.5	How further can Q be increased?	31
3.6	The effect on scaling of heat removal requirements	32
3.7	An alternate approach to 3 dimensional heat removal	33
3.8	Analogies between heat removal and wireability requirements	33
3.9	Power distribution	34
3.10	Conclusion	35
4	System Characterization	37
4.1	Some definitions	37
4.2	Towards a unification of physical and algorithmic approaches	41
4.3	Conclusion	44
5	Tube Models of Interconnections	46
5.1	Survey of interconnection alternatives	46
5.2	Presentation of the models	47
5.3	Justification and derivation of the models	53
5.3.1	Optical interconnections	53
5.3.2	Normally conducting interconnections	56
5.3.3	Repeatedly normally conducting interconnections	61
5.3.4	Superconducting interconnections	63
5.4	Extensions	67
II	Optical Systems	68
6	Lower Bound for the Communication Volume Required for an Optically Interconnected Array of Points	69
6.1	Introduction	70
6.2	Analysis	73
6.3	A derivation based on tubes of unit degree of freedom	75
6.4	Discussion	77

6.5	Conclusions	81
7	Optimal Optical Interconnection Architectures	83
7.1	Introduction	83
7.2	System size considerations	84
7.3	The folded multi-facet architecture	85
7.4	The 3 dimensional multi-facet architecture	89
7.5	Conclusion	93
8	The Optimal Electromagnetic Carrier Frequency Balancing Structural and Metrical Information Densities with Respect to Heat Removal Requirements	94
8.1	Introduction	94
8.2	Preliminaries	96
8.3	Analysis	97
8.4	Discussion	99
8.5	Relation between S , B and N	101
8.6	Conclusion	101
9	A Fundamental Consideration for Dissipative Computing	104
9.1	A fundamental consideration for dissipative computing	104
III	Physical Limits to Communication in Computation	107
10	Basic Analysis	108
10.1	Optical interconnections	108
10.1.1	Relations between S , B and N	108
10.1.2	Heat removal	111
10.1.3	Asymptotic properties	115
10.2	Normally conducting interconnections	116
10.2.1	Relations between S , B and N	116
10.2.2	Heat removal	121

10.2.3	Asymptotic properties	129
10.3	Repeated interconnections	130
10.3.1	Relations between S , B and N	130
10.3.2	Heat removal	131
10.3.3	Asymptotic properties	133
10.4	Superconducting interconnections	134
10.4.1	Relations between S , B and N	134
10.4.2	Heat removal	135
10.4.3	Asymptotic properties	136
10.5	Discussion	136
10.6	Summary and conclusions	137
11	Optimal Hybrid Implementations	139
11.1	Introduction	139
11.2	The breakeven distance approach	139
11.3	Outline of the analysis	141
11.4	Notation and nomenclature	144
11.5	Regarding numerical examples	144
11.6	Some analytical considerations regarding hierarchical analysis	146
11.7	Layouts confined to a plane: normal conductors and optics	147
11.7.1	Description	147
11.7.2	Analysis	148
11.7.3	Numerical examples	152
11.8	Layouts confined to a plane: repeaters and optics	158
11.8.1	Description	158
11.8.2	Analysis	158
11.8.3	Numerical examples	159
11.9	Out of plane optical communication: normal conductors and optics	161
11.9.1	Description	161
11.9.2	Analysis	162
11.9.3	Numerical examples	163

11.10	Out of plane optical communication: repeaters and optics	167
11.10.1	Description	167
11.10.2	Analysis	167
11.10.3	Numerical examples	168
11.11	Other systems	170
11.12	Cost based optimization	173
11.13	Discussion and conclusions	175
11.14	Related work	179
12	On the Usefulness of Optical Digital Computing	180
12.1	Introduction	181
12.2	System characterization and evaluation	182
12.3	Lower bounds to system size	183
12.4	Optimal hybrid partitioning	184
12.5	An argument for optical digital computing	187
12.6	The importance of bit repetition rate B	187
12.7	Comparison of systems with different connectivity	188
12.8	Some credible objections	190
12.9	What might the optical computer look like?	192
12.10	What is it good for?	194
12.11	Beyond dissipative computing	195
12.12	Conclusion	195
13	Indirect Implementations	197
13.1	Introduction	197
13.2	The multiplexed grid architecture	200
13.3	The multiplexed global interconnection architecture	202
13.4	The multiplexed fat tree architecture	203
13.5	Conclusion	206
13.6	Related work	206

14 Comparison of Local and Global Methods for the Simulation of Physical Phenomena	208
14.1 Locality, globality, physics and computation	209
14.2 Implications to future nanoelectronic computing systems	211
14.3 Solution of a problem of high information content	212
14.4 Quantum diffusion as a prototype physical problem	214
14.5 Solution by isomorphic simulation on a locally connected array	216
14.6 Solution by systolic grid convolution on a locally connected array . .	219
14.7 Solution by Fourier techniques on a globally connected array	219
14.8 Comparison	222
14.9 Discussion and extensions	223
14.10 Conclusion	225
15 Summary and Conclusion	226
15.1 Summary and contributions	226
15.2 Conclusion	228
16 Appendices	229
16.1 The effect of scaling for CMOS VLSI circuits	229
16.2 Extension to fan-out and fan-in (chapter 2)	231
16.3 Coefficients for the moments of $g(r)$	232
16.4 Discussion of the validity of neglecting dielectric losses	233
16.5 Discussion of the validity of the quasi-TEM approximation	234
16.6 Signal delay for MOS VLSI circuits	235
16.7 Additional factors of 2 contributing to F	236
16.8 Calculation of $\chi(P)$ for circular and square apertures	236
16.9 Extension to fan-out and fan-in (chapter 6)	239
16.10 System size for the reflective multi-facet architecture	239
16.11 3 dimensional optical layouts with elements confined to a plane . . .	240
16.12 Average signal delay for normally conducting layouts	243
Bibliography	245

List of Tables

1.1	Problem categories	12
5.1	Normally conducting interconnection model	49
5.2	Repeatered interconnection model	51
5.3	Superconducting interconnection model	52
7.1	System size for some optical interconnection schemes	84
12.1	System linear extent for n dimensional meshes	190
12.2	Bisection-bandwidth product and delay for n dimensional meshes . .	190
14.1	Comparison of total computation time with the various methods . . .	222
16.1	Integrated circuit process parameters	230

List of Figures

1.1	Grid Model	9
1.2	Tube Model	10
1.3	Communicative and physical connection graphs	11
2.1	Layout of the connection graph	15
2.2	Binary hierarchical partitioning of the array of cells	16
2.3	Connections made by an element to other elements	19
3.1	3 dimensional heat removal	25
3.2	Heat removal from a square prism via fluid convection	26
4.1	3 parameter characterization	38
4.2	Factors determining the speed of the system	38
4.3	Ψ -surfaces versus Φ -surfaces	42
4.4	Comparison of interconnection media	43
5.1	Tube model of interconnections	48
5.2	Normally conducting interconnection model	50
5.3	Repeatered interconnection model	51
5.4	Superconducting interconnection model	52
5.5	Section of distributed RLC interconnection model	56
5.6	Physical cross section of interconnection model	57
6.1	Communication through a finite surface	70
6.2	The accessible Fourier area at a point P	71
6.3	Volume allocated for optical communication	74

6.4	A volume element in the communication volume	77
7.1	Holographic optical interconnection architectures	86
7.2	The 3 dimensional multi-facet architecture I	90
7.3	The 3 dimensional multi-facet architecture II	91
8.1	The optimal carrier wavelength	100
8.2	Absorption spectrum of sea water	101
8.3	Spectral distribution of solar radiation reaching the earth	102
10.1	Cell size for communication limited systems	109
10.2	The effect of heat removal requirements in 2 dimensions	112
10.3	Critical value of N	113
10.4	The effect of heat removal requirements in 3 dimensions	114
10.5	B versus N for normally conducting interconnections	119
10.6	Critical width and cell size	120
10.7	Comparison of optical and normally conducting interconnections . . .	122
10.8	Comparison for same values of B	123
10.9	Cell size as set by heat removal	125
10.10	Comparison in 3 dimensions	126
10.11	Comparison for same values of B in 3 dimensions	127
10.12	The effect of heat removal for normal conductors in 3 dimensions . .	128
10.13	Optimum duty ratio and resulting performance	129
10.14	S versus N for repeaters in 2 dimensions	132
10.15	S versus N for repeaters in 3 dimensions	133
11.1	Partitioning a system of N elements into N/N_1 groups of N_1 elements	142
11.2	Analysis of optimal hybrid layouts	143
11.3	N_1 versus N for (N2d,O2d) I	152
11.4	S versus N for (N2d,O2d) I	153
11.5	$\mathcal{L}$ versus N for (N2d,O2d)	154
11.6	N_1 versus N for (N2d,O2d) II.	155
11.7	S versus N for (N2d,O2d) II	156

11.8 N_1 versus N for (N2d,O2d) III	157
11.9 $\mathcal{P}$ versus N for (N2d,O2d)	157
11.10 N_1 versus N for (N2d,O2d) IV	158
11.11 N_1 versus N for (R2d,O2d)	160
11.12 N_1 versus N for (N2d,OP1) I	163
11.13 S versus N for (N2d,OP1) I	164
11.14 N_1 versus N for (N2d,OP1) II	165
11.15 S versus N for (N2d,OP1) II	166
11.16 N_1 versus N for (N2d,OP1) III	166
11.17 S versus N for (N2d,OP1) III	167
11.18 N_1 versus N for (R2d,OP1) I	168
11.19 S versus N for (R2d,OP1) I	169
11.20 N_1 versus N for (R2d,OP1) II	169
11.21 S versus N for (R2d,OP1) II	170
11.22 N_1 versus N for (N3d,O3d)	171
11.23 N_1 versus N for (N2d,O2d) V	173
11.24 N_1 versus N for (R2d,OP1) III	174
11.25 N_1 versus N for (R2d,OP1) IV	175
11.26 τ and $\mathcal{P}$ versus N_1 for (N2d,OP1)	178
13.1 The multiplexed grid architecture	201
13.2 The multiplexed global interconnection architecture	203
13.3 The multiplexed fat tree architecture	204
14.1 Decomposition of an FFT	220
14.2 Computation time for the three methods	224
16.1 Extension to fan-out and fan-in	232
16.2 Projection of a solid angle	237
16.3 Accessible Fourier area for a square aperture	238
16.4 A fan-out situation	239

List of Symbols

A	cross sectional area associated with each physical line
$\mathcal{A}$	area (base area) of 2 (3) dimensional system
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 $\equiv$ interelement spacing
d_d	linear extent of an element
d_m	center to center spacing of modules
d_{tr}	linear extent of an optical transducer
d_1	linear extent of a module of N_1 elements
e	Euclidean dimension of layout space
E	energy associated with each transmitted bit of information
$\mathcal{E}$	total energy consumption $\equiv \mathcal{PT}$
f	$\equiv W/\lambda$ for optical systems
$f^\#$	f -number of imaging system
$f(\cdot)$	functional form of connection flux distribution
$g(\cdot)$	functional form of line length distribution
h	height of dielectric
h	Planck's constant
H	bisection
$\mathcal{H}$	height of a 3 dimensional system
J_c	volume critical current density
J_{sc}	surface critical current density

k	number of graph edges (connections) per element
K	number of graph edges passing through each cell
K'	number of physical wiring tracks per cell per layer $\equiv K\chi/M$
ℓ	length of a line in real units
$\bar{\ell}$	average connection length in real units
ℓ_{max}	longest connection length in real units
L	inductance per unit length
$\mathcal{L}$	linear extent of the system
$\mathcal{L}_1$	$\equiv d_1$
m	order of moment of line length distribution
M	number of interconnection layers
n	fractal dimension of layout
N	number of elements
N_1	number of elements in each module
N_{1max}	maximum number of elements in each module for given B
$N_{1\infty}$	optimal value of N_1 for large N
$\mathcal{N}$	number of connections $\equiv kN$
p	connectivity (Rent exponent) of layout
P	number of graph edges emanating from a group of elements
$\mathcal{P}$	total power dissipation
q	modified Rent exponent $\equiv \max(p, (e-1)/e)$
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_{max}	longest connection length in grid units
$\langle r^m \rangle$	m th 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
$\mathcal{T}$	total time of computation
V	nominal voltage level
$\mathcal{V}$	volume of 3 dimensional system
w	width of conductor
W	transverse linear extent associated with each physical line
W_{min}	minimum manufacturable linewidth
Z_0	characteristic impedance
α	attenuation constant
β	$\equiv B/S$
Γ	optimization function
δ	classical skin depth
ϵ	permittivity of dielectric
ζ_m	coefficient for the m th moment of line length distribution
κ	coefficient for average connection length $\equiv \zeta_1$
λ	optical wavelength
λ	superconducting penetration depth
μ	permeability of dielectric
ν	optical frequency
ρ	resistivity of conductor
τ	worst case signal delay
τ_{ave}	average signal delay
v	velocity of propagation
χ	number of parallel physical lines used to establish each graph edge
ω	fundamental frequency component

...

*for who knows how many
ways clouds and
shadows and light can
be rearranged.*

*I would like to believe in
everything, every
possible pattern
every
possible relation,
all directions
in perfect contradiction
as they are.*

Obscure Mediterranean origin
(translated by the author)

Part I

An Interconnect Dominated Model of Computation

Chapter 1

Introduction

In this thesis we concern ourselves with *communication limits in computation*. This is not a study about distributed algorithms or graph theory, but rather about the physical limitations to information transfer among an array of points. An array of points communicating with each other is meant to suggest an abstraction of a computing system.

1.1 Background

1.1.1 The increasing importance of interconnections

In a conventional electronic circuit diagram, a solid line drawn between two device terminals represents an ideal connection. A voltage or current impressed on one end of the line is assumed to appear immediately at the other end. This abstraction is sufficient in representing real circuits as long as the physical wires are not too long and the devices are not too fast, as has been historically the case for most circuit applications [32]. With the use of faster devices and longer wires, it has been necessary to account for the capacitive, inductive and resistive effects of the wires. Thus, a computer has often been viewed as a collection of switches interconnected with (parasitic) wires, which degrade the performance expected from an ‘ideal’ system [147].

However, several general arguments and observed trends suggest that with increasing integration levels, measures of space, delay and energy tend to become dominated by the interconnections (wires), rather than the switching elements [145] [54] [5] [6] [65] [69] [63] [84]. In modern integration technology both the devices and the interconnections are characterized by the same minimum feature size, usually represented by λ . The area occupied by a device is then $\sim \lambda^2$ whereas the area occupied by an interconnection is $\sim \lambda \ell$, where ℓ is the length of the interconnection. Increase in the integration level is achieved by lowering λ or increasing system size (which results in an increase in ℓ) or a combination of both [82]. In any case the ratio of interconnection area to device area is increased.

Likewise, interconnect energy dissipation can constitute very large fractions of the system total. Eventually, interconnect area and power dissipation tend to determine how close the devices can be packed together, hence determining the communication delays between the devices and speed of operation of the overall system. In fact, even yield considerations for integrated circuits may be more concerned about defects occurring in the lines (leading to undesired shorts or disconnections) rather than those in the devices [120], as the probability of occurrence of the former increases faster than that of the latter with increasing integration levels.

It is well known that the intrinsic RC delay of an unterminated integrated circuit interconnection does not change under ideal downscaling, if its length is downscaled as well. However, scaling enables one to achieve a higher level of integration in the same area. In addition, since allowable chip areas have grown the RC delays have actually increased. The gate delays decrease with scaling so that the interconnects become the performance limiting aspect of the overall system. This discussion was given by Saraswat in 1982 [145] and the general trends predicted were confirmed in more recent work [147]. Although several quasi-ideal scaling laws [4] [54] may be applied to slow down this trend, it is not possible to circumvent it [139]. For example the aspect ratios of the connections can be made closer to unity to improve the capacitance, characteristic impedance and series resistance. However, this approach becomes useless once the fringe fields begin to dominate and crosstalk is increased to unacceptable levels.

Two important parameters of an RC limited VLSI interconnection are its switching energy and rise time. We write, following earlier authors [145] [4],

$$\begin{aligned} E_{VLSI} &= C_{line} V^2 = C \ell V^2 \\ \tau_{VLSI} &= (2R_d + R\ell) C \ell \end{aligned} \tag{1.1}$$

where R and C denote the series resistance and capacitance of the line per unit length. R_d is the impedance of the driving device and ℓ is the length of the line. V is the nominal voltage level. In writing these equations contributions of gate input and output capacitances have been neglected. A more detailed discussion is given in appendix 16.1 where more detailed design equations are shown to reduce to the above for longer line lengths and submicron scaling. Although the appendix takes as an example CMOS VLSI circuits scaled under a certain scaling rule, it is in general true that with increasing integration levels, more interconnections will tend to exceed the minimum feature size by greater ratios. Since both capacitance and series resistance increase with interconnection length the gate contributions are quickly dwarfed in comparison. In fact, for the longest interconnections or for a device technology where transistors have high current driving capability, even $2R_d$ may be neglected in the expression for τ_{VLSI} and we may write simply $\tau_{VLSI} = RC\ell^2$. For CMOS technology (as shown in appendix 16.1), the above equations are valid for lines about a few hundred times longer than the minimum feature size, when the minimum feature size is reduced to submicron dimensions.

Even if we avoid the use of resistive interconnections, the delay along any form of interconnection cannot be less than the speed of light delay which will exceed device delays with increasing system sizes.

Although the shorter links may still be device dominated, for all but very locally connected systems, the longer connections—with larger energies, delays or cross sections—will often tend to determine system performance and cost. This suggests that the ultimate limitations to achievable computational power can be discussed with little or no reference to a specific device technology.

Moreover, some computational approaches aimed at circumventing the limitations of conventional systems tend to exhibit large communication requirements, further increasing the importance of interconnections. Examples of such approaches are fine grain parallel computation, neural networks and other connectionist approaches. Knowledge representation and processing concepts where the knowledge itself is stored in the pattern of interconnections, rather than the devices themselves, have been proposed [38]. Some of these approaches are inspired by the human brain.

In accordance with these trends, in this work we do not think of wires as mere parasitics, but make them the basis of our models and abstractions.

1.1.2 Different interconnection media

As interconnections become more and more the factor limiting the performance of large scale processing systems, the use of optical or superconducting interconnections as an alternative to resistive wires has been suggested [59] [58] [128] [77] [87] [175]. Many authors have made comparative studies of optical and normally conducting interconnections [92] [94] [40] [43] [65] [121] and superconducting and normally conducting interconnections [50] [95] [100]. These studies have mostly been concerned with comparing various parameters of isolated interconnections, without exploring the simultaneous interplay of these parameters when these interconnections are embedded in a system.

1.2 Purpose

Based on the discussion of the previous section, here we outline some of our goals. First of all, we would like to develop a device independent, interconnect dominated theory of computation emphasizing the physical properties of interconnection media. We also hope that this will provide a framework in which the role and usefulness of different interconnection media can be understood. And to some limited extent, we would like to provide a basis for exploring the relationships between computational requirements, architecture and the physics of interconnection media, based on

the realization that choice of algorithm and architecture must take into account the limitations to be encountered during actual physical implementation and operation, which are strongly influenced by the physical properties of the interconnections.

More directly stated, our purpose in this work has been to set up and solve various problems relating to the physical limitations to information transfer among an array of points. Our viewpoint is more that of a physicist, rather than that of an engineer. Technological aspects of the problem will not be stressed in this thesis.

1.3 Approach

The problem we intend to attack is characterized by an intimidatingly large number of parameters. One approach would be to try to specify in full detail a system which seems can be actually constructed in the very near future, numerically account for all (or most) known higher order effects, and present the results in the form of numerical graphs. Straightforward as it is, there are two major drawbacks with such an approach.

The first is the loss of generality due to the large number of specific assumptions that must be made. In this work we will not try to fully specify a particular real life implementation. In many instances, certain different ways of doing things do not result in a numerical difference greater than a factor of the order of unity.

The second drawback is the loss of analytic transparency due to algebraic complexity. In a study of this nature where we are trying to obtain understanding of what is a very complicated problem, we must confront ourselves with the more difficult task of trying to untangle the more important issues from the less important ones, rather than try to include everything in an attempt to be 'exact'. In other words, we will try to isolate a section of reality in the form of a reasonably well defined physical problem amenable to analysis. We are aware that a certain amount of subjectivity is unavoidable in this process, and that many alternatives to our approach are possible.

We will assume practical, technological, economic, and production related problems not to be limiting factors. We will eliminate details related to mechanical construction. We will simply imagine our processor to be an array of points connected

with some kind of wire. This will mean ignoring artifacts like packages, bonding pads, boards etc. We might imagine that we are building the whole computing system on a single huge chip. In practice, systems approaching this ideal might be possible by assembling individually tested chips on a carrier substrate and connecting them with VLSI-like metallization.

The general trend in integrated circuit technology seems to be—at least partially—approaching that depicted by our idealized model. Wafer scale integration (WSI), thin film hybrid substrates and ceramic multilayer hybrids are attempts to be freed from the deleterious effects of conventional chip packages and circuit boards [120] [155] [174] [118] [117] [157] [6].

We will also make several general assumptions. We will assume that the characteristics and parameters of our processor are spatially uniform and that it is spatially isotropic, i.e. it is more or less compact in shape (square or cubic) and it does not have any preferred axis or orientation. We will also refrain from introducing so called ‘safety’ or ‘efficiency’ factors which do not have any theoretical basis and cannot be derived within the context of this study. Such factors are often close to, or of the order of unity.

So as to minimize the arbitrariness involved in deciding which physical effects should be accounted for and in choosing parameters for our numerical examples, we will try to choose what seems the best possible from a physical standpoint, anticipating technological advances. In other words, we will look into the future and attempt to analyze the basic limitations imposed on us after all technical and practical difficulties have been overcome. Determining the ultimate performance possible is a formidable task, especially if one insists on numerical accuracy. Since our aim is to develop 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 the best achievable within a factor of the order of unity. Of course, we are to a certain extent conditioned by the current trends in existing technologies and the way present computing systems are built, so we cannot exclude the possibility of breakthroughs or ingenuity not foreseen by us.

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

In conclusion, our approach may be seen to be a compromise between more detailed numerical engineering simulations which are meant to represent an actual to-be-constructed system; and approaches of a higher level of abstraction such as VLSI complexity theory. The first approach may yield more accurate results, but lacks generality and analytic transparency. The latter approach is of greater generality, but is naive from a physical standpoint and does not even distinguish between different kinds of interconnection media.




Figure 1.1: *Grid Model.*

1.4 Overview


1.4.1 Grid model

The basic abstraction of a computing system is a graph, with its nodes corresponding to the elements or switches, and its edges corresponding to the interconnections. Figure 1.1 shows our computer with its elements arrayed on a regular cartesian grid.

We would like to make our system as small as possible. This is not only a merit in itself, but will also result in smaller signal delays along the connections. However, the interelement spacing d must be large enough to

1. Accommodate the elements.
2. Allow enough space for implementing the desired pattern of connections.
3. Satisfy heat removal requirements.

These are the major physical considerations behind our analyses.

Figure 1.2: *Tube Model.*

1.4.2 Tube model

Of course, in reality, the interconnections are not 1 dimensional lines, but 3 dimensional objects. In this work an interconnection is imagined to be a flexible tube (figure 1.2) characterized by several parameters: length, cross sectional area, signal delay, minimum bit repetition interval and energy per transmitted bit. Any interconnection medium is characterized by the relationships tying these parameters together.

Books and journals dedicated to the analysis of transmission lines give detailed empirical expressions and numerical results for parameters such as the attenuation coefficient, characteristic impedance etc., which are only intermediate variables for our purpose. To carry out a system level analysis, we need a characterization in terms of the variables listed above which directly interface with our grid model. For instance, the cross sectional area defines packing density, the energy per transmitted bit is related to heat removal and so on. This characterization must be simple and transparent enough to be used in a study of a relatively high degree of abstraction, yet must still capture the essential limiting mechanisms and tradeoffs of each interconnection medium. Coming up with such models was one of the more challenging tasks in this work.

We will use these tubes (characterized by the relations tying the mentioned parameters together) to provide communication among the nodes of our system, ensuring that wireability and heat removal requirements are satisfied. This thesis is devoted to the discussion of various facets of this physical problem. In particular, we have




Figure 1.3: *Communicative (part a.) and physical (parts b. and c.) connection graphs.*

explored how the tradeoffs between the microscopic parameters listed above for a single interconnection result in tradeoffs between the macroscopic parameters of our system, such as the total number of elements, signal delay, bandwidth, system size and power dissipation.

1.4.3 Graph distinctions

Before continuing, we must distinguish between different kinds of graphs. Part a. of figure 1.3 shows a particular *communicative* connection graph among 4 nodes. An edge joining two nodes indicates that one of these nodes can send bits to the other. (The directionality of the edges makes no difference for our purpose.) The *physical* connection graph represents the actual physical wires tying the nodes together. Part b. shows a *direct* implementation of the communicative connection graph of part a. Part c. shows an example of an *indirect* implementation where it is assumed that each element has the capability of forwarding bits passed along from one of their neighbors to the other. The physical connection graph associated with an indirect implementation may involve additional dummy nodes for routing purposes, as will be exemplified in chapter 13.

In most of this work we assume that the communicative connection graph is given to us as a starting point. We are free however in choosing the physical connection

	Basic	Hybrid
Direct		
Indirect		

Table 1.1: *Problem categories.*

graph used to implement the given communicative connections. The communicative connection graph is related to the problem and algorithm, whereas the physical connection graph is related to actual physical construction.

Our analyses will not actually depend on these definitions, which are not meant to be precise and were introduced for pedagogical purposes.

1.4.4 Problem categories

We classify the problems that can be considered within our framework under 4 categories (table 1.1). The categories falling to the left involve the use of only one interconnection medium at a time for all connections, such as using normal conductors for all connections. Those to the right involve the use of more than one at the same time, such as using normal conductors for some connections and optics for other connections. The upper categories are direct implementations where the communicative connection graph is directly copied as the physical connection graph, whereas in the lower categories some form of indirect implementation is utilized. In this thesis we will mainly concentrate on direct implementations.

1.4.5 Outline

In part I we develop the general framework and models used in this work. Chapter 2 discusses wireability limitations and how we quantify the connectivity of computing systems. In chapter 3 we discuss heat removal limitations and derive a heat removal model for 3 dimensional systems. In chapter 4, we introduce several variables enabling us to tie together our grid and tube models; and also discuss how the results

to be derived in chapters 10 and 11 may be interpreted in relation to the computational requirements of particular applications [136]. Tube models for optical, normally conducting and superconducting interconnections are derived in chapter 5 [136]. In chapters 2, 3 and 5, presentation of the models are followed by their derivation and justification.

Part II of this thesis concentrates on issues specific to optical systems. In chapter 6 we derive a basic result regarding the ability of optical wave fields in providing communication among an array of points [134]. Then, in chapter 7, we compare various specific optical interconnection architectures in terms of system size, and propose two new architectures which can potentially achieve the least possible system size of any two and three dimensional architecture respectively [132] [133]. Chapters 8 and 9 present relatively fundamental results for dissipative optical systems.

In part III, we discuss general applications of the models presented in part I. First, in chapter 10, we consider the use of one interconnection medium at a time for all connections and derive tradeoff relations between signal delay, bandwidth, number of elements, system size and power dissipation for each interconnection medium [136]. Then, in chapter 11, we consider the joint use of optical and normally conducting interconnections and discuss in what combination they must be used so as to obtain a system with desired properties [135]. Chapter 12 discusses the implications of these considerations to the usefulness of an optical digital computer. In chapter 13, we briefly discuss the usefulness of indirect implementations in exploiting the high bandwidth potential of optical and superconducting interconnections [137]. Chapter 14 is an attempt to pave the way for future research regarding the comparison of locally and globally connected approaches to constructing future very large scale systems. Chapter 15 summarizes and concludes this thesis.

Chapter 2

Wireability Limitations and Connectivity Model

It was mentioned in the introductory chapter that the spacing between the elements of our system must be large enough to allow sufficient space for the interconnections to pass between them. Systems employing longer connections will require greater interelement spacing than systems employing shorter connections. We will speak of systems employing greater fractions of longer connections as being *highly connected*. In this chapter we present a model which enables us to quantify the *connectivity* of computer circuits to first order and can predict the interelement spacing required to ensure that there is enough space for implementing the desired pattern of connections.

2.1 Description of the model

For the purpose of this work, a processing system is a collection of N given similar *primitive elements* connected to each other according to a prespecified *graph* [10] [11] [14] [170]. The primitive elements may be simple switching devices or relatively complex subsystems. k will denote the number of connections (graph edges) per element¹, so that there is a total of $\mathcal{N} = kN$ connections. Within a factor of 2, we

¹For simplicity we are considering pairwise connections only, the extension to fan-out and fan-in is discussed in appendix 16.2.




Figure 2.1: *Layout of 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 interconnections and heat removal requirements.

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 accommodate their input-output ports (transducers in the case of optical interconnections).

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 2.1). In this work we do not attempt an interpretation of fractional values of e so that $e = 2$ or $e = 3$. Figure 2.2 depicts a hierarchical partitioning of our array of cells.

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




Figure 2.2: *Binary hierarchical partitioning of the array of cells.* (After [34].)

1. The average connection length $\bar{r}$ of the layout (in grid units). Higher order moments will also be useful in deriving some results.
2. The longest connection length r_{max} (in grid units).
3. The number of connections $P(N')$ emanating from a group of N' elements in the partitioning of figure 2.2.

All of the above quantities can be specified by postulating the distribution of line lengths for our system². Obviously, we cannot hope to account for all possible connection patterns. Rather, we seek a simple analytic distribution function with a few variable parameters, which we hope will be representative of the wireability requirements of typical circuits.

We will define the *line length distribution* $g(r)$ in terms of the *connection flux distribution* $f(r)$, which we will assume to be of the form

$$f(r) = kr^{e(p-1)} \left(1 - \frac{r^e}{r_{max}^e} \right) \quad 1 \leq r \leq r_{max} \quad (2.1)$$

where r denotes distances in units of grid spacing so that physical distances are given by $\ell = rd$. $r_{max} \gg 1$ denotes the longest connection length (in grid units), assumed to be of the order of the linear extent of the system. We will take $r_{max} \simeq N^{1/e}$ without

²The terms *line*, *connection* and *graph edge* are used interchangeably.

concerning ourselves with precise geometrical factors³. $Nf(r_0)$ gives the expected number of connections of length r_0 or greater. Note that $f(1) = k$ and $f(r_{max}) = 0$. The parameter $0 \leq p \leq 1$, known as the *Rent exponent*, is our measure of connectivity. The smaller p is, the quicker $f(r)$ decreases. Systems with larger Rent exponents have a larger fraction of longer connections. The factor $(1 - r^e/r_{max}^e)$ has been introduced to account for the finite extent of the system and may be ignored either when r is not close to r_{max} or when p is not close to 1. (In this work the largest value of p used in numerical examples will be 0.8, so that whether this factor is included or not will make little difference.) The line length distribution may now be defined as

$$g(r) = -\frac{df(r)}{dr} \quad 1 \leq r \leq r_{max}. \quad (2.2)$$

$Ng(r_0)\Delta r$ 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 $\int Ng(r) dr = kN$. $k^{-1}g(r)$ may be interpreted as a probability distribution defined over $[1, r_{max}]$. 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 $p = 1$.

The *fractal dimension* of information flow of our layout is defined by $n = 1/(1 - p)$ [112] [26] [27]. This parameter will be used as an equivalent measure of connectivity. A brief justification of the use of the term ‘dimension’ is given in the following section.

A useful approximation for the m th moment $\langle r^m \rangle = k^{-1} \int r^m g(r) dr$ may be derived as

$$\begin{aligned} e < mn \quad \langle r^m \rangle &= \zeta_m N^{\frac{m}{e} - \frac{1}{n}} \\ e = mn \quad \langle r^m \rangle &= \zeta'_m \ln N \\ e > mn \quad \langle r^m \rangle &= \zeta''_m \end{aligned} \quad (2.3)$$

where the coefficients are functions of m , $n = 1/(1 - p)$ and e . Because it is most often used, we will use the special symbols $\bar{r} = \langle r \rangle$ and $\kappa = \zeta_1$ for the first moment,

³There seems little to be gained by trying to specify and carrying around factors such as $\sqrt{2}$ etc.

the average connection length:

$$\begin{aligned} e < n & \quad \bar{r} = \kappa N^{\frac{1}{e} - \frac{1}{n}} = \kappa N^{p - \frac{e-1}{e}} \\ e = n & \quad \quad \bar{r} = \kappa' \ln N \\ e > n & \quad \quad \bar{r} = \kappa'' \end{aligned} \tag{2.4}$$

Notice that when $e < n$, the average connection length is proportional to the ratio of the linear extent of the system in e -space ($N^{1/e}$), to that in n -space ($N^{1/n}$).

Our line length distribution is also consistent with the following expression for the number of connections $P(N')$ emanating from a group of N' elements:

$$P(N') = kN'^p \left(\frac{N - N'}{N} \right) \tag{2.5}$$

where N is the total number of elements in the system. When $N' \leq N/2$ or so, the term in parenthesis can be ignored so that $P(N') = kN'^p$. It is also easy to see that when $p = 1$, it is equally likely for connections to be made to elements at any distance, consistent with the usual interpretation of $p = 1$. (Simply notice that in this case, among the k connections made by an element, $k(N - N')/N$ will be made to elements outside the group, and kN'/N to elements inside the group.)

Note that we are not accounting for the external connections of our system ($P(N) = 0$). We are implicitly assuming that input and output from our system is established by impressing or reading off information from the $\propto N^{(e-1)/e}$ peripheral elements.

2.2 Justification of the model

According to our model, the number of connections emanating from a group of N' elements in figure 2.2 is given by the relationship $P(N') = kN'^p$ when $N' \leq N/2$. This relationship is widely known as *Rent's rule*. Statistical variations from this formula are to be expected.

Historically, Rent's rule precedes the line length distribution. Rent's rule was originally established as an empirical relationship [103] [144] and later shown to be




Figure 2.3: *Connections made by an element to other elements.*

a consequence of the logic design process [35] [36]. Such a power law may also be justified based on a principal of *self similarity* [34] [142]. We now understand that Rent's rule is also related to the *separator* concept of VLSI complexity theory [164] [106], which provides a formal basis for the layout of given graphs, and to the theory of fractals [112]. This relationship has been used widely as a wiring model for two decades [163].

Donath [37] and Feuer [45] were the first to show that Rent's rule is consistent with a line length distribution similar to what we are assuming. Here we will give a crude derivation so that the reader can gain some insight regarding the relation between Rent's rule and the connection flux distribution. Ignoring edge effects and the distinction between cartesian and Euclidean distances, we will assume that the connection flux and line length distributions can be thought to describe the connections emanating from each element. That is, $f(r_0)$ may be interpreted to give the *expected* number of connections originating from a certain element and emanating from a spherical surface of radius r_0 centered at that element, as suggested by figure 2.3. (Likewise $g(r_0)\Delta r$ may be interpreted as the *expected* number of connections originating from an element and terminating in the interval $[r_0, r_0 + \Delta r]$.) Thus the

expected number of connections emanating from a group of N' elements is given by $P(N') \simeq N' f(N'^{1/e})$, since this group is of linear extent $\simeq N'^{1/e}$. This evaluates to $P(N') \simeq k N'^p (1 - N'/N)$, as given in the preceding section (equation 2.5).

The expressions for the moments, given in equation 2.3, are obtained easily by evaluating $\langle r^m \rangle = k^{-1} \int r^m g(r) dr$. The first and second moments were already derived by Donath [34] [37]. A discussion of the approximations leading to these equations and expressions for the coefficients are given in appendix 16.3.

We now briefly justify the use of the term ‘dimension’ for the quantity defined as $n = 1/(1 - p)$. The perimeter of a square region is proportional to the $1/2$ power of its area. The surface area of a cube is proportional to the $2/3$ power of its volume. In general, the hyperarea enclosing a hyperregion of e dimensions is proportional to the $(e - 1)/e$ power of its hypervolume. Let us now make an analogy between ‘hyperarea’ $\Leftrightarrow$ ‘number of connections emanating from a region’, and ‘hypervolume’ $\Leftrightarrow$ ‘number of elements in the region’. According to our model the number of connections emanating from the region is proportional to the p th power of the number of elements in the region. Thus, it makes sense to speak of the quantity n defined by the relation $p = (n - 1)/n$ as the *dimension* of information flow. The interested reader is referred to the work of Mandelbrot for a discussion of the relationship between inverse power law distributions and fractal forms [113] [111] [110].

Given the interrelationships between layout theory, Rent’s rule, fractal geometry and inverse power law distributions, we are convinced that our model is indeed a meaningful way of quantifying the connectivity of computer circuits to first order. Our confidence is increased by the fact that similar models have been used by many other authors in the past (for instance [86] [163] [6]). Further discussion of these interrelationships is beyond our scope and is not necessary for our purpose 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 [62] without any underlying theory. The use of such a line length distribution may also be justified empirically [54] [53].

Notice that r is actually a discrete quantity. It is possible to find graphs for which our continuous approximation leads to erroneous results. For instance, according to

equation 2.2, 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 all of our results may be cast in a form that depends only on the first few moments of $g(r)$, the length of the longest connection and the functional form $P(N')$, 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, as suggested by earlier authors [34] [37] [45] [6].

A significant quantity is the number of graph edges passing through each cell, which we denote by K . If $\chi \geq 1$ parallel physical lines are 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 accommodating the element itself [85]. A moments reflection reveals that K is given by $K = k\bar{r}$ [66] [52], since $k\bar{r}$ is the total connection length per cell in grid units. Letting W denote the transverse linear extent of a single physical line, including its share of line to line spacings, the above condition may be expressed as $d^{e-1} \geq \chi K W^{e-1}$. Combining this with the condition $d \geq d_d$, we will write $d \geq \max(d_d, (\chi k\bar{r})^{1/(e-1)} W)$. Notice that the error we incur in pretending that the interconnections and elements may cooccupy 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 interconnection 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 parts⁴ 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 $H = k\kappa N^p$. This must be multiplied by χ to obtain the number of physical lines crossing the surface in question.

We will mostly concentrate on highly connected systems, characterized by large values of p (or equivalently n). The method of analysis is easily extended to other cases.

2.3 Graph layout

In this work, we do not concern ourselves with how a given graph must be laid out on the grid of figure 2.1. We assume that the placement of the elements is determined independently and given to us as a starting point. One particular way of laying out the elements would be to place them randomly. This would result in a Rent exponent of $p = 1$, regardless of the topology of the connection graph. Although optimal graph layout is in general an NP-complete problem, several heuristic methods [73] can result in layouts with more localized and shorter connections and a smaller value of p . We are assuming that upon layout, our system exhibits a distribution of line lengths of the form $g(r)$ as defined earlier. Of course, we cannot expect arbitrarily given graphs to have this property. However, analyzing layouts which we assume result in such a line length distribution will enable us to understand the dependence of our results on the connectivity of the system.

The Rent exponent is a property of both the graph topology and the layout algorithm. However, for a given graph, there is a bound to how much the Rent exponent can be reduced with *any* layout algorithm.

The grid model and method of accounting for wireability requirements we are using is attributed to Keyes [85] [84]. We should note that several approaches to graph layout [106] [164] [12] [105] employ a somewhat different grid model attributed to Thompson [158].

⁴This surface is also referred to as the *bisection*.

2.4 Conclusion

Overall, our model is a useful paradigm enabling us to quantify the connectivity of computer circuits to first order and can predict several parameters of our layout. Although similar models have been used by many authors in the past, we must understand that it is not capable of describing everything and that it must be seen as an instrument facilitating a general analysis. Our method of analysis, however, is applicable not only to layouts with different distribution functions, but also to any layout for which the quantities $\bar{r}$, $P(N')$ etc. can be meaningfully specified.

Chapter 3

Heat Removal Model

The interelement spacing d in our grid model must be large enough so that we can successfully remove the dissipated power. Packing the elements too densely may result in unacceptable temperature rises and destruction of the system.

3.1 2 dimensional systems

For 2 dimensional systems, we assume that there is an upper limit to the amount of power we can remove per unit area, denoted by Q (W/m^2). This model has been extensively used by other authors [79] [80]. Thus the linear extent $\mathcal{L}$ of a square system uniformly dissipating a total power $\mathcal{P}$ must at least be $(\mathcal{P}/Q)^{1/2}$.

To give a feeling for the order of magnitudes involved, we note that the power density arriving from the sun is about $0.1 \text{ W}/\text{cm}^2$ [126]. The human body gives away about $10 \text{ mW}/\text{cm}^2$. Boiling water may carry away up to $10 \text{ W}/\text{cm}^2 \text{ K}$ per degree of temperature rise. For thermal radiative transfer at 100 K , we obtain $\sigma(100 \text{ K})^4 \sim 1 \text{ mW}/\text{cm}^2$ where $\sigma = 5.67 \times 10^{-8} \text{ W}/\text{m}^2 \text{ K}^4$ [71]. We will concentrate on fluid convection, which seems to be the most effective means of heat removal.

We do not embark on a quantitative discussion of 2 dimensional heat removal at this point, as our 3 dimensional model will also serve as a 2 dimensional model.




Figure 3.1: 3 dimensional heat removal.

3.2 3 dimensional systems

Whatever the modality (conduction, convection or radiation), heat transfer must take place *through* a surface. Thus for 3 dimensional systems, it is not possible to quantify our heat removal ability by specifying the amount of power we can remove per unit volume. In what follows, we will show that our ability to remove heat from a 3 dimensional system can be quantified by a quantity Q , interpreted as the maximum amount of power we can remove per unit cross section of the system, consistent with the intuitive notion of heat flow, and in analogy with the 2 dimensional case. This is most easily visualized by considering the flow of a cooling fluid through our system, as illustrated in figure 3.1. Thus the linear extent $\mathcal{L}$ of a square prism system dissipating a total power $\mathcal{P}$ must at least be $(\mathcal{P}/Q)^{1/2}$.

3.3 Derivation of the model

Let $\mathcal{P}$ denote the total power dissipated in a square prism of volume $\mathcal{L} \times \mathcal{L} \times \mathcal{H}$ with $\mathcal{H} \leq \mathcal{L}$ as illustrated in figure 3.2. $\mathcal{H}/\mathcal{L}$ will be kept fixed throughout our discussion. We assume that the total power $\mathcal{P}$ is dissipated uniformly in the solid volume between the tubes. A coolant fluid will be forced to flow through the tubes of diameter $2r_0$ and axial separation $2r_1$. Thus there are $\mathcal{L}^2/4r_1^2$ tubes each with internal surface




Figure 3.2: *Heat removal from a square prism via fluid convection.*

area $2\pi r_0 \mathcal{H}$. So as to ensure that the volume of the tubes does not exceed a certain fraction of the total volume $\mathcal{L}^2 \mathcal{H}$, we will require that $\eta = r_0/r_1 \leq \eta_{max} < 1$.

We are interested in determining the minimum value of $\mathcal{L}$ at which we can successfully remove the dissipated power from the system. Based on the conservation of energy, we might immediately claim that the weakest growth rate of $\mathcal{L}$ as a function of $\mathcal{P}$ is $\propto \mathcal{P}^{1/2}$, since the surface area of the system grows as $\propto \mathcal{L}^2$. However, it is not obvious that this minimum growth rate can be achieved under constant applied pressure difference. Not only will the increasing length of the tubes tend to hinder fluid flow, but also the solid-fluid interface will limit the amount of heat we can transfer onto the fluid.

Since our purpose is to illustrate the general principles involved in as transparent a manner as possible, rather than make engineering predictions, several textbook [71] assumptions will be employed.

We consider laminar flow of an incompressible fluid. We assume that the temperature gradients do not affect fluid flow and that the pressure is uniform at any cross section, decreasing linearly in the axial direction. We also assume that the flow is fully developed, i.e. invariant in the axial direction [71].

We ignore edge effects. Based on similar arguments as in [86], we assume that heat conduction is significant only in the transverse plane. Thus, the heat flux entering from the walls of the tubes is uniform along the axial direction, resulting in a constant temperature gradient in the axial direction.

We also assume the material parameters involved to be constant, however, one should be aware that the viscosity of most fluids is actually quite temperature dependent.

Our analysis is modeled on that of [162] [161]. The two dominating components of thermal resistance are those due to the heating of the fluid ($\Theta_{cal} = \Delta T_{cal}/\mathcal{P}$) and due to the transfer of heat at the solid-fluid interface ($\Theta_{conv} = \Delta T_{conv}/\mathcal{P}$). Θ_{cal} may be expressed as $1/\rho C_s \mathcal{F}$ where ρC_s denotes the volumetric heat capacity of the fluid and $\mathcal{F}$ is the total fluid flow rate. The flow rate through a single tube is given by $\pi r_0^4 \Delta P / 8 \mu \mathcal{H}$ where ΔP is the applied pressure difference and μ is the viscosity of the fluid [71]. Multiplying this with the number of tubes to find $\mathcal{F}$, we obtain

$$\Theta_{cal} = \frac{32}{\pi} \frac{1}{\eta^4} \frac{\mu}{\rho C_s \Delta P} \frac{\mathcal{H}}{r_1^2 \mathcal{L}^2}. \quad (3.1)$$

Θ_{conv} may be expressed as $1/hs$ where s denotes the total internal surface area of the tubes and h is the heat transfer coefficient. h is given by $Nu \kappa / D$ where κ is the conductivity of the fluid and D is the hydraulic diameter, equal to its geometrical diameter for a tube [71]. Since we are assuming a fully developed flow, we will take the Nusselt number Nu to be equal to its steady state value of $48/11$, based on similar arguments as in [162]. Thus

$$\Theta_{conv} = \frac{11}{12\pi} \frac{1}{\kappa} \frac{r_1^2}{\mathcal{H} \mathcal{L}^2}. \quad (3.2)$$

Note that it is optimum to set $\eta = \eta_{max}$. Then the value of r_1 minimizing $\Theta = \Theta_{cal} + \Theta_{conv}$ and the resulting thermal resistance is found to be

$$r_0^4 = (\eta_{max} r_1)^4 = 34.9 \frac{\mu \kappa}{\rho C_s \Delta P} \mathcal{H}^2 \quad (3.3)$$

$$\Theta = 2\Theta_{cal} = \frac{3.45}{\eta_{max}^2} \left(\frac{\mu}{\rho C_s \kappa \Delta P} \right)^{\frac{1}{2}} \frac{1}{\mathcal{L}^2}. \quad (3.4)$$

It is interesting to note that the optimum value of r_1 depends only on $\mathcal{H}$ whereas the minimum value of Θ depends only on $\mathcal{L}$. The resulting temperature rise is expressed as $\Delta T = \Delta T_{cal} + \Delta T_{conv} = 2\Delta T_{cal} = \Theta \mathcal{P}$. We can now define the more intrinsic quantity Q by the relation $Q\mathcal{L}^2 = \mathcal{P}$, so that $Q = \Delta T/\Theta\mathcal{L}^2$. If the maximum allowed temperature rise is specified, we may express Q as

$$Q = 0.29 \eta_{max}^2 \left(\frac{\rho C_s \kappa \Delta P}{\mu} \right)^{\frac{1}{2}} \Delta T. \quad (3.5)$$

Once $\mathcal{L} = (\mathcal{P}/Q)^{1/2}$ is calculated, we can go back to equation 3.3 to calculate the optimum value of r_0 for given $\mathcal{P}$.

We finally express Q in terms of ν_{max} , the maximum fractional volume occupied by the tubes, rather than η_{max} . Using $(\pi/4)\eta_{max}^2 = \nu_{max}$, we obtain

$$Q = 0.37 \nu_{max} \left(\frac{\rho C_s \kappa \Delta P}{\mu} \right)^{\frac{1}{2}} \Delta T. \quad (3.6)$$

Choosing $r_0^2 \propto \mathcal{H}$ keeps the hydraulic resistance of the tubes and the mean flow velocity v *through the tubes* constant. It can be easily shown that $v = r_0^2 \Delta P / 8\mu\mathcal{H} = 0.74(\kappa\Delta P/\rho C_s \mu)^{1/2}$, independent of $\mathcal{L}$ and $\mathcal{P}$. Likewise the mean flow velocity *through the system* $v_s = (\pi r_0^2/4r_1^2)v$ is given by $0.58 \eta_{max}^2 (\kappa\Delta P/\rho C_s \mu)^{1/2}$, in terms of which we may write the intuitively appealing $Q = \rho C_s v_s \Delta T_{cal} = \rho C_s v_s \Delta T/2$, consistent with equation 3.5.

Let us also calculate the viscous power dissipation $\mathcal{P}_p = \Delta P \mathcal{F}$ associated with the fluid flow. Using the previously derived expression for $\mathcal{F}$ and expressing everything in terms of $\mathcal{P}$, we obtain $\mathcal{P}_p = (2\Delta P/\rho C_s \Delta T)\mathcal{P}$, an exceedingly simple result. $\mathcal{P}_p$ is proportional to the power dissipated by the devices. We will later show that for typical values, $\mathcal{P}_p \ll \mathcal{P}$.

In our analysis we have ignored the effects of conduction in the solid medium in which the circuits are embedded. Consider the cell of dimensions $2r_1 \times 2r_1 \times \mathcal{H}$ enclosing each tube. Clearly, there will be no heat transfer through the boundaries of this cell (except near the edges of our system). For the purpose of calculation, let us pretend that this cell is distorted into a cylinder of diameter $2r_1$. Now, it is possible to solve Laplace's equation $\kappa_s \nabla^2 T(r) = -q_v$ inside the solid, subject to the boundary

condition that the radial derivative of the temperature (which is proportional to the heat flux) vanishes at $r = r_1$. $T(r)$ denotes the radial dependence of the temperature for $r_0 \leq r \leq r_1$, κ_s is the conductivity of the solid and q_v is the volume density of generated power and is given by $\simeq \mathcal{P}/(1 - \eta^2)\mathcal{L}^2\mathcal{H}$. Once we solve for $T(r)$, the additional temperature rise ΔT_{cond} and resulting thermal resistance $\Theta_{cond} = \Delta T_{cond}/\mathcal{P}$ due to conduction in the solid is easily found as

$$\Theta_{cond} = \lambda(\eta) \frac{1}{\kappa_s} \frac{r_1^2}{\mathcal{H}\mathcal{L}^2} \quad (3.7)$$

where $\lambda(\eta) = (\ln(1/\eta^2) - (1 - \eta^2))/4(1 - \eta^2)$. Once again the largest possible value of η is preferred. For instance, $\lambda(0.5) = 0.21$. Noticing the similarity of the above expression to equation 3.2, it is easy to see that the effects of conduction in the solid may be accounted for by replacing

$$\frac{1}{\kappa} \rightarrow \frac{1}{\kappa} + \frac{12\pi\lambda(\eta)}{11\kappa_s} \quad (3.8)$$

in our previous analysis. We see that the effects of solid conduction may be ignored with little error if $\kappa_s > 10\kappa$ or so, as would almost always be the case.

We have also assumed uniform power dissipation throughout the solid. This would be appropriate for systems where the majority of power is dissipated on the wires. If instead the power $\mathcal{P}_d$ associated with each device is dissipated within a small radius of r_d , an additional temperature rise of $\sim 3\mathcal{P}_d/8\pi\kappa_s r_d$ would be observed. (This result is derived by an elementary application of Laplace's equation. We calculate the difference in temperature rise when the total power is dissipated within a radius of r_d and when it is dissipated within a much larger radius.) If $\mathcal{P}_d = 1 \text{ mW}$, $r_d \geq 0.1 \mu\text{m}$ and $\kappa_s \sim 150 \text{ W/m K}$ (corresponding to silicon), we find that this temperature rise does not exceed ten degrees and is usually acceptable.

3.4 Numerical example

We now present numerical results assuming $\mathcal{H}/\mathcal{L} = 1$, $\rho = 10^3 \text{ kg/m}^3$, $C_s = 5 \times 10^3 \text{ J/kg K}$, $\kappa = 0.5 \text{ W/m K}$, $\mu = 10^{-3} \text{ kg/m sec}$ (corresponding to water), $\Delta P = 10^5 \text{ kg/m sec}^2$ and $\Delta T = 100 \text{ K}$. We take $\eta_{max} = 0.5$ so that the volume occupied by

the pipes is less than 25% of the total system volume. Thus the following results may be derived:

$$r_0^2 = 1.87 \times 10^{-7} \mathcal{H} = 1.78 \times 10^{-11} \mathcal{P}^{\frac{1}{2}} \quad (3.9)$$

$$\Theta = 8.73 \times 10^{-7} \frac{1}{\mathcal{L}^2} = 96.0 \frac{1}{\mathcal{P}} \quad (3.10)$$

$$Q = 1.1 \times 10^8 \quad (3.11)$$

$$v = 2.34 \quad (3.12)$$

$$v_s = 0.46 \quad (3.13)$$

$$\mathcal{P}_p = 4 \times 10^{-4} \mathcal{P} \ll \mathcal{P} \quad (3.14)$$

where everything is in SI units. 10 KW can be removed per square centimeter. For instance, for a total power of 1 KW we find $\mathcal{L} = 3$ mm, $r_0 = 24$ μ m and $\Theta = 0.096$ K/W. On the other hand, for a total power of 1 MW, we find $\mathcal{L} = 95$ mm, $r_0 = 133$ μ m and $\Theta = 9.6 \times 10^{-5}$ K/W. A megawatt can be removed from a liter, quite larger than what was previously thought possible. We also note that the viscous power dissipation is negligible in comparison to the device power dissipation.

Let us finally check some of our major assumptions. The Reynolds number $Re = v\rho D/\mu$ should be less than 2100 in order to validate our assumption of laminar flow [71]. This leads to the condition $\mathcal{P} < 128$ MW which would allow 10^{11} circuits each dissipating 1 mW in a cube of edge length of about a meter. The fluid may be considered to be fully developed (justifying our use of the steady state Nusselt number) when the distance x from the entrance of the tube satisfies $x/(D Re Pr) \geq 0.02$, provided the Prandtl number $Pr = \mu C_s/\kappa > 5$ or so [161]. Since $Pr = 10$ in our example, upon substitution we obtain $x \geq 0.35 \mathcal{H}$, so that the velocity and temperature profiles are well developed over a greater portion of the distance along the tube. (This is not a coincidence, but a direct outcome of the optimization procedure [161]. It is easy to demonstrate that the relation $\mathcal{H}/(D Re Pr) = 1/4 Nu$ holds in general, which immediately leads to the above result.) Let us also check whether axial conduction in the conductor is small with respect to radial conduction. The power carried through each tube by the coolant is $\mathcal{P}/(\mathcal{L}^2/4r_1^2)$. The axial flow may be estimated as $\sim \pi(r_1^2 - r_0^2)\kappa_s \Delta T/2\mathcal{H}$. The ratio of the latter to the former after substituting our chosen numerical values is found to be $4 \times 10^{-5}/\mathcal{H}$. Thus if $\mathcal{H} > 1$ mm, this ratio is

less than 5%, validating our assumption. Thus our analysis is applicable for systems ranging from ~ 1 mm to ~ 1 m in size.

In 1981 Tuckerman and Pease experimentally demonstrated the removal of 790 W from a $1\text{ cm} \times 1\text{ cm}$ surface using cooling fins about 0.04 cm in height [162]. This corresponds to $790\text{ W}/(1\text{ cm} \times 0.04\text{ cm}) \simeq 20\text{ KW}/\text{cm}^2$ of power being removed per unit area along the direction of fluid flow, in reasonable agreement with our predictions. (The factor of two discrepancy is easily traced down to the different values of ΔT , ΔP and ν utilized.) In fact, if we imagine that we stack 25 such assemblages on top of each other, we obtain a system which crudely resembles figure 3.2. The major difference is that the channels are narrow slits, instead of circular tubes. Provided their total area is always the same fraction of $\mathcal{L}^2$, the use of alternate cross sectional shapes for the channels alter the general results of this paper only by geometrical factors close to unity. Whereas a general proof seems difficult, it is possible to show that for narrow slits extending along the full width of our system, the value of Q is within 10% of what has been calculated for circular tubes.

3.5 How further can Q be increased?

Let us examine the various factors in equation 3.6. Water already exhibits one of the highest $\rho C_s \kappa / \mu$ ratios found among all materials (mercury is slightly better) [161]. Even if our circuits could withstand very large operating temperatures, an upper limit to ΔT is set by the difference in freezing and boiling points of the liquid used. If water is used, $\Delta T \leq 100\text{ K}$.

Without going into a detailed analysis, it is possible to claim that $\nu \sim 0.5$ or so is optimum. Increasing ν beyond 0.5 would increase Q by at most another factor of 2, while decreasing the volume available for the circuits drastically. Reducing ν below 0.5 would increase the volume available for the circuits by at most a factor of 2, while drastically reducing our heat removal ability. This argument allows us to decouple heat removal issues from wireability issues and thus greatly simplifies the use of our model.

The only remaining parameter is ΔP . Ultimately, the upper limit to ΔP will be

set by viscous power dissipation. If viscous power dissipation becomes comparable to $\mathcal{P}$, it must be added on to the total power dissipation. Thus $\mathcal{L}$ must be large enough to satisfy

$$\mathcal{L}^2 Q \geq \mathcal{P} + \mathcal{P}_p = \mathcal{P} \left[1 + \left(\frac{2\Delta P}{\rho C_s \Delta T} \right) \right] \quad (3.15)$$

since $\mathcal{P}_p = (2\Delta P / \rho C_s \Delta T) \mathcal{P}$. By substituting for Q from equation 3.5 and differentiating with respect to ΔP , we can find that the optimal value of ΔP resulting in minimum $\mathcal{L}$ is given by $\Delta P = \rho C_s \Delta T / 2$, as we might have guessed directly: the viscous power dissipation should not exceed the power dissipated by the circuits. With this value of ΔP , we obtain a more fundamental limit to Q :

$$Q = 0.1 \eta_{max}^2 (\rho^2 C_s^2 \kappa / \mu)^{1/2} \Delta T^{3/2}. \quad (3.16)$$

Once again, water exhibits one of the highest $\rho^2 C_s^2 \kappa / \mu$ ratios among all materials [161]. With the same numerical values we find that Q may not exceed $2.8 \times 10^9 \text{ W/m}^2$.

3.6 The effect on scaling of heat removal requirements

A detailed analysis of the effects of heat removal and wireability considerations will be presented in a later part of this thesis. Here we discuss an important consequence of our heat removal model in its simplest form.

Let us assume a system with a bounded degree connection graph and constant power dissipation per element.

In 2 dimensions, wireability requirements dictate that the linear extent of a system of N elements grow as $\propto N^q$ where $1/2 \leq q = \max(p, 1/2) \leq 1$. On the other hand, heat removal requirements dictate that the linear extent of the system grow as $\propto N^{1/2}$. Thus, unless $q = 1/2$, wireability requirements will surpass heat removal requirements with increasing N . Larger values of p enable greater connectedness, but result in larger layout area and delays, so that a detailed analysis is necessary to determine the optimal value of p resulting in a system with optimal properties [6].

In 3 dimensions, wireability requirements dictate that the linear extent of the system grow as $\propto N^{q/2}$ where $2/3 \leq q = \max(p, 2/3) \leq 1$. Heat removal again dictates a growth rate $\propto N^{1/2}$. Thus, for large N , the choice of p has little if any effect on the resulting system size. Since smaller values of p will not reduce system size and delays, we might as well employ high values of p , increasing connectivity. This suggests that it will be more beneficial to employ highly connected approaches in large scale 3 dimensional computing systems.

The above arguments must be modified if the power dissipation per element cannot be assumed constant. However, it should be evident that 3 dimensional systems exhibit a tendency to be heat removal limited, so that highly connected approaches can be employed with little effect on system size and delays.

3.7 An alternate approach to 3 dimensional heat removal

The only other approach to modeling 3 dimensional heat removal at a similar level of abstraction that we are aware of relies only on steady state conduction [80]. By assuming a homogeneous conductivity κ_s throughout space and assuming that a power $\mathcal{P}$ is uniformly dissipated in a spherical system of diameter $\mathcal{L}$, it is possible to show (via solution of Laplace's equation) that the minimum value of $\mathcal{L}$ for given temperature rise ΔT is $\mathcal{L} = \mathcal{P}/4\pi\kappa_s\Delta T$. Notice that the system linear extent $\propto \mathcal{P}$ rather than the best possible $\propto (\mathcal{P}/Q)^{1/2}$, which was achieved by the conductive-convective system we analyzed.

3.8 Analogies between heat removal and wireability requirements

After reading chapter 10, the reader will notice that the limitations imposed by optical and superconducting interconnections are similar to those imposed by convective heat removal. Constant cross section is required for the flow of information or heat

(via convection) over any distance. On the other hand, conductive heat removal is analogous to normally conducting interconnections. The heat removal scenario we have analyzed makes joint use of convection and conduction. This is similar to the joint use of optical and normally conducting interconnections, which will be analyzed in detail in chapter 11.

3.9 Power distribution

In this study, heat removal and communication (wireability) requirements are emphasized as the major factors limiting packing density and system performance. The effects of power distribution are not considered. This is usually considered appropriate for existing or foreseeable technology [139]. Here we briefly discuss the limitations imposed by power distribution requirements.

From a fundamental perspective, power distribution may be seen as the opposite of heat removal. Thus, we might suspect that power distribution will impose similar growth rates on system size as heat removal. The quantity of interest is the amount of power we can feed per unit cross section.

First, consider optical power distribution (such as to feed modulating devices or optical switches). Even if we use a separate spatial channel to feed each element, the cross sectional area λ^2 needed per element will often be less than that imposed by heat removal ($(10 \mu\text{m})^2$ for $Q = 10 \text{ KW/cm}^2$ and elements dissipating 10 mW each). The distributing lines can be made to consume even less space by organizing them in the form of a 2 or 3 dimensional H-tree.

With the use of superconductors, the critical current density will determine the maximum power density. The power flow along a superconducting transmission line of characteristic impedance Z_0 is given by $I^2 Z_0$ where I is the current through the conductor. In chapter 5, we will see that $Z_0 = \sqrt{\mu/\epsilon} h/w$ when $h \geq t \geq \lambda$, where μ and ϵ are the permeability and permittivity of the dielectric, λ is the superconducting penetration depth and the dimensions h , w , t and $A = 4wh$ are defined in figure 5.6. The maximum current is given by $I = J_{sc} w$ where J_{sc} is the surface critical current density. Thus the power density is given by $I^2 Z_0 / A = J_{sc}^2 \sqrt{\mu/\epsilon} / 4$. (It is also possible

to show that the power density decreases if t is reduced below λ , so that we are not overlooking any room for improvement in this direction.) With $J_{sc} = 50 \text{ mA}/\mu\text{m}$, we find that about 10^{11} W/m^2 of power can be supplied. This is much larger than the maximum amount of power we can remove and thus imposes a weaker restriction.

Finally, we consider the use of normal conductors. The axial voltage drop ΔV on a normally conducting supply line must not exceed a certain fraction of the supply voltage. For a line of length ℓ , cross sectional area A and resistivity ρ , we have $\Delta V = \rho \ell I/A$. Thus, for given ΔV , the maximum power density is $VI/A = V\Delta V/\rho \ell \sim V^2/\rho \ell$ which evaluates to about $(4 \times 10^7 \text{ W/m})/\ell$ for $V = 1 \text{ V}$ and room temperature conductivity. For a system of the order of a meter in size, this is comparable to $Q \sim 10^8 \text{ W/m}^2$. Often, due to the use of less aggressive heat removal methods, the value of Q will be less than 10^8 W/m^2 . Thus, even in the normally conducting case, power distribution requirements need not be considered until system sizes of the order of $\sim 1\text{-}10 \text{ m}$ are reached.

The space consumed by the clock distribution system (if any) is also not considered in this work. We can assume the use of a 2 or 3 dimensional (probably optical [28]) H-tree for clock distribution. The space consumed by such an H-tree can always be absorbed in the space consumed by the elements themselves with little error.

3.10 Conclusion

We have considered the problem of heat removal from a square prism of volume $\mathcal{L} \times \mathcal{L} \times \mathcal{H}$ in which a total power $\mathcal{P}$ is uniformly dissipated. Referring to the configuration of figure 3.2, we showed that choosing the number and cross sectional area of the tubes proportional to their length was optimum. With this choice, we found that a heat removal imposed lower bound to the system linear extent $\mathcal{L}$ can be written as $\mathcal{L}^2 \geq \mathcal{P}/Q$ where Q is a function of the material parameters of the coolant fluid, the applied pressure difference ΔP and the maximum allowed temperature rise ΔT . Thus Q is conveniently interpreted as the maximum amount of power we can remove per unit cross section, in analogy with the 2 dimensional case where it is customary to specify the maximum amount of power we can remove per unit area [86]. With

$\Delta P = 1 \text{ atm}$, $\Delta T = 100 \text{ K}$ and assuming water is the coolant fluid, we have estimated $Q \sim 10 \text{ KW/cm}^2$.

Chapter 4

System Characterization

In this chapter we will first introduce several variables which will serve as an interface between our grid, connectivity and heat removal models and the tube models of interconnections to be presented in chapter 5.

In this work we assume that the performance of our systems can be characterized in terms of 3 parameters (figure 4.1): the number of elements N , the bit repetition rate B along each edge of the connection graph and the inverse signal delay S (worst case or average). In this chapter we will also discuss how this characterization can be related to the computational requirements of given applications/algorithms.

4.1 Some definitions

One way to increase the processing power of a 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 steps. The physical duration of a time step (measured in seconds) is set by one of several mechanisms, as illustrated in figure 4.2. Information transfer takes place along the edges of the connection




Figure 4.1: 3 parameter characterization.




Figure 4.2: *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.

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. 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 chapter 5. 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

$$\tau = \max(\tau_d, T, T_p) \quad (4.1)$$

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 [5]. 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* [10] 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 mostly 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 bit/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 $\chi \rightarrow \infty$. 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 of¹. 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 accommodate $\sim \chi k$ 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 .

We will be interested in the HS (bisection-inverse delay) and HB (bisection-bandwidth) products of our systems. These 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 [164] [130] [7], so that these products are direct measures of system performance.

Although it would certainly be desirable, it is not possible to arbitrarily increase S , B and N simultaneously due to physical limitations. We will quantify this by

¹The use of wavelength division multiplexing constitutes an exception and must be treated separately.

deriving bounds of the form $\Phi(S, B, N) \leq C_\Phi$ for different interconnection media in chapter 10. For our present purpose, it is sufficient to realize that there will be a surface defining the region of mutually consistent values of S , B and N . This region will satisfy 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 $\iota_i \leq 1$. The optimal operating point in this region must be determined in conjunction with the requirements of the application, as discussed next.

4.2 Towards a unification of 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 connectivity p .

It is possible to an extent to tradeoff 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 [164].

These *lower* bounds should be interpreted in conjunction with the *upper* bounds




Figure 4.3: Ψ -surfaces versus Φ -surfaces. Part a. illustrates a situation where there are no common points among the regions defined by the lower and upper bounds. In part 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 .

of the form $\Phi(S, B, N) \leq C_\Phi$ mentioned at the end of the previous section, 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 medium.

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 interconnection medium is capable of handling given problems, but also to determine the appropriate choice of S , B and N . Figure 4.3 illustrates the various possibilities, where for simplicity in illustration we assume that B is not involved in the tradeoff. Part a. illustrates a situation where the Φ -surface completely lies below the Ψ -surface. This interconnection 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 (part b.). 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. Part c. shows a situation where we have the flexibility of choosing S and N from a finite interval where the Φ -surface lies above the Ψ -surface.

We now see how we may compare various interconnection media characterized by




Figure 4.4: *Comparison of interconnection media.* Part a. illustrates the Φ -surfaces associated with two different interconnection media, one of which is superior to the other. In part 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. Part c. shows a Ψ -surface which has an intersection with only the surface belonging to the first media, whereas part d. shows one which intersects only the second.

the functional forms $\Phi(S, B, N) \leq C_\Phi$ with reference to the computational requirements of a given problem and algorithm. We illustrate this in figure 4.4 where we have again assumed that B is not involved. Part a. shows the upper bounds for two different interconnection media. The surface lying to the upper right is superior to the other regardless of problem requirements and operating point. Part b. illustrates a situation where one medium is superior to the other for N greater than a certain critical value. Neither, one, or both may be able to perform the stated task in the stated amount of time. Parts c. and d. of figure 4.4 illustrate two different Ψ -surfaces, one of which has an overlap with only the first Φ -surface, the other with the second.

Let $\mathcal{T}$ denote the total time of computation, $\mathcal{E}$ the total energy consumption over

this period of time and $\mathcal{V}$ (or $\mathcal{A}$ for a 2 dimensional system) the size of our system. As shown in part c. of figure 4.3, increasing the total time $\mathcal{T}$ allowed for computation allows us to choose between many possible values of N and S (and in general also B). Each of these particular implementations will result in a particular $\mathcal{V}$ and $\mathcal{E}$. In other words, we can map the region of mutually consistent values of S , B and N (as defined by the relationship $\Phi(S, B, N) \leq C_\Phi$) into a region of mutually consistent values of $\mathcal{T}$, $\mathcal{V}$ and $\mathcal{E}$. Many points may be mapped into one. Now, given an optimization function (i.e. a figure of merit function) involving $\mathcal{T}$, $\mathcal{V}$ and $\mathcal{E}$, we can pick the optimal implementation(s). Remember that the Rent exponent, choice of algorithm and choice of interconnection medium (including the use of more than one medium at the same time) are hidden parameters in this formalism and offer additional degrees of freedom we can optimize over. In short, this procedure (in principal) enables us to find the optimal medium (or combination of media), Rent exponent, algorithm, number of elements and bit repetition rate and resulting inverse delay of the most desirable system.

4.3 Conclusion

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. Their determination is independent of the physical construction of the system. Likewise, the determination of Φ -surfaces is problem and algorithm independent. The two considerations are tied together through the parameters S , B and N .

Hillis [70] 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 [164].

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 [147]. 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 is the philosophy of this work.

(Parallel) algorithms must be developed in conjunction with their physical implementation. As an example, we consider the work of Feldman et al. [42]. 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 efficient optical implementations due to their space invariant properties.

Choice of interconnection medium is dependent on the problem, algorithm, connectivity etc. Likewise, choice of algorithm and architecture is dependent on the physical properties of the interconnections. Thus, determination of the most desirable system requires that both aspects of the problem be treated in conjunction and jointly optimized over.

In chapter 10 we will limit ourselves to an examination of the behavior of S as a function of N for constant B , to an examination of the asymptotic properties of the bisection-inverse delay and bisection-bandwidth products and some other particular optimization functions. Based on the discussion of this chapter, 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 an area for future research.

Chapter 5

Tube Models of Interconnections

In this chapter we present the interconnection models used in this work. We will first give a general discussion of the origin of the different properties of conducting and optical interconnections. Then we will present the relationships between length, cross sectional area, signal delay, minimum bit repetition interval and energy per transmitted bit for optical, normally conducting, repeatered normally conducting and superconducting interconnections. Finally, we will present a justification of these models.

5.1 Survey of interconnection alternatives

Historically, normally conducting wires have been used to interconnect the elements of electronic circuits. The use of conductors enables good confinement of the wave fields (in the sense that the internal electric and magnetic fields can terminate at the charges and currents on the conductors) and promises deep submicron scaling. However, by nature, normal conductors are lossy. This requires an increase in the cross sectional area¹ and energy per transmitted bit of such lines with increasing line lengths.

We will exclusively concentrate on double conductor normally conducting lines

¹That is, if we are to maintain length proportionate delay. To be precise, the cross sectional area *can* be kept constant, but then the delay and bandwidth degrade sharply.

supporting quasi-TEM modes. The use of hollow conductor waveguides, resulting in TE and TM modes is not beneficial, as in this case there is a limit to how narrow these lines can be before the cut-off frequency exceeds the fundamental frequency component [29].

Optical interconnections rely only on dielectric inhomogeneities for confinement and thus suffer relatively little loss. (The use of conductors, which would be extremely lossy at these frequencies, is out of the question). For the very same reason however, the cross sectional area per independent spatial channel cannot be less than the carrier wavelength squared. (It is for this reason that we prefer to modulate a high frequency optical carrier, rather than simply send baseband signals. In the latter case, even at microwave frequencies each channel would have to be several centimeters wide.) Again for the same reason, such interconnections suffer from coupling and radiation losses. Since the dielectric constant variation in nature is much less than the conductivity variation, we cannot avoid ‘spilling photons’. (This observation also has a positive side. Optical interconnections need not suffer to the same degree from termination problems.) Thus the cross sectional area and energy per transmitted bit for optical interconnections are relatively large but approximately independent of line length.

In conclusion, because of the length dependence of cross sectional area and energy per transmitted bit with normal conductors, they are preferred over shorter distances whereas optical interconnections are preferred over longer distances.

At first, superconducting interconnections seem to offer the best of both worlds, since they offer conductor confinement without loss; however they have their own limitations which must be considered in detail.

Chemical transmission and mass transport, although widely observed in biological systems, have not found their way into artificial computing systems and are not considered.

5.2 Presentation of the models

In this work an interconnection is imagined to be a flexible tube with the following parameters:




Figure 5.1: *Tube model of interconnections.*

1. Interconnection length ℓ .
2. Cross sectional area A or transverse linear extent W where $A = W^2$. These parameters define packing density for 3 and 2 dimensional systems respectively and thus include any necessary line to line separations.
3. Signal delay τ . As discussed in the previous chapter, the signal delay is given by the greater of the propagation delay T_p and the minimum temporal pulse width T , which in turn is the greater of a line imposed component T_ℓ and a device imposed component T_d .
4. Minimum pulse repetition interval T_r , which is usually equal to T , the minimum temporal pulse width along the interconnection.
5. The energy per transmitted bit, E .

Any interconnection medium is characterized by the relationships tying these parameters together. For instance, it is immediate that for any medium, the signal delay cannot be less than ℓ/c , where c is the speed of light.

The tube model we use for optical interconnections is the simplest possible. The cross sectional area is taken to be proportional to the wavelength squared: $A = W^2 = (f\lambda)^2$ where the constant f can be as small as ~ 1 for a diffraction limited system but may be larger in practice. The signal delay is taken to be the greater of the speed of light delay and the device rise time: $\tau = \max(\ell/c, T_d)$. Since the effects of dispersion and attenuation can be made small for the length scales in consideration, the minimum

	τ	T	E	ter.
$W^2 \leq 16\rho\epsilon v\ell$	$16\rho\epsilon\frac{\ell^2}{W^2}$	$16\rho\epsilon\frac{\ell^2}{W^2}$	$2\epsilon V^2\ell$	no
$W^2 \geq 16\rho\epsilon v\ell$	$\frac{\ell}{v}$	$16\rho\epsilon\frac{\ell^2}{W^2}$	$2\epsilon V^2 vT$	yes

Table 5.1: *Normally conducting interconnection model when $T_d \leq T_\ell$.* The delay $\tau = \max(T, T_p)$, pulse width $T = \max(T_\ell, T_d)$ and energy E are given as functions of length ℓ and width W . The last column indicates whether the line is to be terminated or not in that region.

pulse repetition interval T_r and energy per transmitted bit E are assumed to be constants.

In table 5.1 and figure 5.2 we see the relationships tying the length, cross sectional area, delay and energy for normally conducting lines for the case $T_d \leq T_\ell$. The symbols ρ , ϵ , μ and v denote the resistivity of the conductor, permittivity and permeability of the dielectric and propagation velocity in the dielectric respectively. V denotes the nominal voltage level. The lower right of the slanted line corresponds to terminated transmission, where the delay is proportional to length and independent of cross sectional area. The upper left region corresponds to unterminated charge-up. What is unique to our model is that it intrinsically accounts for the proper scaling effects due to the skin effect and deals with RC lines and transmission lines in a unified manner. One immediate conclusion that may be derived from our model is that the use of a single wide line is more beneficial than many narrow ones in terms of obtaining maximum information density. Yet another conclusion is that, it is beneficial to scale down wire limited layouts until we are in the RC limited region. After this, further reduction in scale does not further improve system signal delay. This conclusion is related to a well known argument stating that the rise time of an RC line remains constant when all of its dimensions are downscaled.

The use of active repeater devices along the line changes the curves for delay versus linewidth as shown in table 5.2 and figure 5.3. We should note that we employ a highly idealized model for repeaters; for instance we do not address power distribution to the repeating devices and we ignore their discrete structure. R_0C_0 denotes the intrinsic delay of the repeating devices. The use of repeaters weakens, but does not eliminate




Figure 5.2: Normally conducting interconnection model when $T_d \leq T_\ell$.

	τ	T	E	ter.
$W \leq 4\sqrt{\frac{\rho R_0 C_0}{\mu}}$	$4\sqrt{R_0 C_0 \rho \epsilon} \frac{\ell}{W}$	$R_0 C_0$	$2\epsilon V^2 \ell$	no
$W \geq 4\sqrt{\frac{\rho R_0 C_0}{\mu}}$	$\sqrt{\mu \epsilon} \ell$	$R_0 C_0$	$8\epsilon V^2 \sqrt{\frac{\rho R_0 C_0}{\mu}} \frac{\ell}{W}$	yes

Table 5.2: *Repeated interconnection model*. The delay τ , pulse width T and energy E are given as functions of length ℓ and width W . The last column indicates whether the stages of the line are to be terminated or not in that region.




Figure 5.3: *Repeated interconnection model*.

the inverse dependence of the delay on the width of the line. The delay along a repeated line also remains constant with downscaling.

In table 5.3 and figure 5.4 we can see the delay as a function of linewidth for superconducting lines. Our models take into account the proper scaling effects associated with the superconducting penetration depth λ and the critical current density J_{sc} .

An immediate conclusion is that for wire limited layouts, it is optimal to scale down the system until we are in the intermediate region. Once we are, the delay starts exhibiting an inverse dependence on width. So, even with ideal lossless superconductors, the signal delay remains constant with downscaling. Because of this, the signal delay possible with superconductors for a wire limited layout is very similar to that possible with optics, although the superconducting system may be much smaller

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

Table 5.3: *Superconducting interconnection model when $T_d \leq T_p$ (or $T_d \leq T_\ell$ in the lumped case).* The delay $\tau = \max(T_p, T)$, pulse width T and energy E are given as functions of length ℓ and width W . The last column indicates whether the line is to be terminated or not in that region.




Figure 5.4: *Superconducting interconnection model when $T_d \leq T_p$ (or $T_d \leq T_\ell$ in the lumped case).*

in size, if deep submicron scaling is employed. We are not aware of an interconnection medium that enables continual improvement of signal delay through downscaling.

5.3 Justification and derivation of the models

5.3.1 Optical interconnections

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 [90]. 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 cross talk 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 [177]. This means that the required guide separation 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 wavelength². 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.

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

$$W = 2\lambda \tag{5.1}$$

$$A = W^2 = 4\lambda^2 \tag{5.2}$$

²It 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.

independent of all other parameters. Although it would be considerably difficult to do any better than this in practice, theoretically there is still a little more room for improvement [134]. In this work we treat optical communication links as if they are 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 to be derived in chapter 6 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 [134]. 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. In fact, we will sometimes more generally write $W = f\lambda$ where the constant f can be as small as ~ 1 for a diffraction limited system but may be larger in practice. The essential feature of our model is that the cross section need not be increased with increasing line length [134]. $\lambda \sim 1 \mu\text{m}$ 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 attenuation constant α can usually be made small enough so that the condition $\alpha \ell \ll 1$ will be satisfied for the length scales in consideration. The major sources of loss (coupling and device inefficiency) are independent of length. Expressions for the required energy per transmitted bit were given previously by many authors [92] [40] [121] [39] assuming the use of laser diodes [31], light emitting diodes and light modulators [16] as output transducers. We assume the use of light modulators so that no threshold term is involved. The following is a refinement of the calculation of [121].

To generate a voltage change of V on a detector with capacitance C_D an energy

$$E = \frac{h\nu C_D V}{\eta q} \quad (5.3)$$

is required. h is Planck's constant, q is the charge of an electron, ν is the optical frequency and η is the overall differential quantum efficiency of the system, i.e. ratio of detector current to source current. We do not expect that the value of η in a large

system can be improved beyond ~ 0.1 . We must also include the energy required by the modulator, $C_M V_M^2$. Also we take into account the dissipation associated with the photocurrent [121]. The total energy is found to be

$$E = \frac{h\nu C_D V}{\eta q} + C_M V_M^2 + \frac{V_M}{(h\nu/q)} \frac{h\nu C_D V}{\eta q} \quad (5.4)$$

which numerically becomes

$$E(\text{fJ}) \simeq 24 \frac{A_D(\mu\text{m}^2)}{h_D(\mu\text{m})} + 3 \frac{A_M(\mu\text{m}^2)}{h_M(\mu\text{m})} \quad (5.5)$$

where A_D (A_M) and h_D (h_M) are the area and thickness of the detector (modulator) respectively. $V = 1\text{ V}$, $V_M = 5\text{ V}$, $\eta = 2.5\%$ and a relative permittivity of 12 for the device material were assumed. For devices of $(5\mu\text{m})^2$ area and $1\mu\text{m}$ thickness, we find $E \sim 1\text{ pJ}$.

The reader is referred to Feldman et al. [40] for a discussion of 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




Figure 5.5: *Section of distributed RLC interconnection model*

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 .

5.3.2 Normally conducting interconnections

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 5.5). As is mostly appropriate [120], the shunt dielectric conductance is ignored. (This is further discussed in appendix 16.4.) Figure 5.6 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 [121]. 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 $W = 2w$ in 2 dimensions and $A = WH = 2w \times 2(h + t)$ in 3 dimensions, independent of length. Whereas the numerical factors involved are again somewhat arbitrary, they seem to




Figure 5.6: *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. $W = 2w$ is the two dimensional minimum packing dimension. $A = WH = 4w(h+t) \simeq 4w \max(h, t) = 4wh$ is the three dimensional minimum packing area.

be representative of the geometry to which technology is converging [100]. 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 $V = 1$ V 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

$$T_\ell = (16\rho\epsilon)\frac{\ell^2}{W^2} = (16\rho\epsilon)\frac{\ell^2}{A}. \quad (5.6)$$

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 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 bit³. 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 5.6 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 [152]. 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 $E = V^2 T / Z_0$ would exceed that possible with an unterminated line $E = V^2 C \ell$. 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 traveling along the line. Thus we are agreeing to leave a line unterminated when $T > T_p$ and to terminate 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 equation 5.6 first for unterminated lines. It is known that the skin effect need not be considered in this mode of operation [152]. The rise time of the line is given by $\simeq (R_d + R\ell)C\ell$ where R_d is the drive impedance [145]. Assume for the moment that the line is not drive limited; then the rise time and energy per bit are given by

$$T_\ell = RC\ell^2 = \frac{\rho}{wt} \frac{\epsilon w}{h} \ell^2 = \frac{\rho \epsilon \ell^2}{ht} = 4\rho \epsilon \frac{w}{h} \frac{w}{t} \frac{\ell^2}{W^2} \quad (5.7)$$

³In 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 16.6.

$$E = V^2 C \ell = \epsilon V^2 \frac{w \ell}{h} \quad (5.8)$$

where $W = 2w$ has been used. 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 equation 5.6 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 [114].

The performance of present day MOS VLSI lines may be much worse than predicted by the above, because such lines are often drive limited [147] [167], 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 16.6, where we discuss how the contribution of drive resistance to the delay may be absorbed into the parameter T_d .

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. (This is further discussed in appendix 16.5.) We will show that the fundamental frequency satisfies $\omega L > R$ so that we may ignore the correction terms [116] $1/(1 + R^2/8\omega^2 L^2)$ and $(1 - jR/2\omega L)$ associated with the propagation velocity and characteristic impedance respectively.

In this case, the minimum pulse width satisfies $T < T_p = \ell/v$. 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 $R/2Z_0$ [143], we require approximately $R\ell \leq Z_0$ so that attenuation is kept at an acceptable level⁴. 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 H 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

⁴This corresponds to degradation of the signal level by $e^{-\alpha\ell} = e^{-0.5} = 0.6$.

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

$$h\delta \geq \rho \sqrt{\frac{\epsilon}{\mu}} \ell. \quad (5.9)$$

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

$$T_\ell = \rho \epsilon \frac{\ell^2}{h^2} = 4\rho \epsilon \left(\frac{w}{h}\right)^2 \frac{\ell^2}{W^2} = (16\rho \epsilon) \frac{\ell^2}{W^2} = (16\rho \epsilon) \frac{\ell^2}{A}. \quad (5.10)$$

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 $E = V^2 T / Z_0 = (w/h) \sqrt{\epsilon/\mu} V^2 T$ which also indicates that we should choose h/w large. Thus we obtain $E = 2\sqrt{\epsilon/\mu} V^2 T$ 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 $16\rho \sqrt{\epsilon/\mu} \ell = 16\rho \epsilon v \ell \leq W^2 = A$. If this condition is not satisfied, the line is high-loss and pulse transmission is not possible. Our model equations are summarized in table 5.1 and figure 5.2 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 W and ℓ we have $T_\ell < T_d$, we can reduce the width W 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 equation 5.6, 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 W by two (i.e. A by four) decreases T_ℓ by four, corresponding to a potential 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 W 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.

5.3.3 Repeated normally conducting interconnections

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 will not be concerned with the discrete structure of such lines, treating them as if they were a continuous 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 [4] 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 are given by

$$\xi \simeq \sqrt{\frac{RC\ell^2}{R_0C_0}} \quad (5.11)$$

$$T_p \simeq \sqrt{R_0C_0RC\ell^2}. \quad (5.12)$$

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

$$\xi \simeq 4\sqrt{\frac{\rho\epsilon}{R_0C_0}} \frac{\ell}{W} \quad (5.13)$$

$$T_r = T \simeq T_\ell \simeq T_d \simeq R_0C_0 \quad (5.14)$$

$$T_p = \xi T \simeq 4\sqrt{R_0C_0\rho\epsilon} \frac{\ell}{W} \quad (5.15)$$

$$E \simeq 2\epsilon V^2 \ell. \quad (5.16)$$

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 ℓ/W ; the number of stages increases linearly with distance. The length of each stage depends only on W and may be found to be

$$\ell_{stage} = \frac{\ell}{\xi} = \frac{1}{4} \sqrt{\frac{R_0C_0}{\rho\epsilon}} W \quad (5.17)$$

which together with $W^2 \leq 16\rho\epsilon v\ell_{stage}$ dictates that approximately $W \leq 4\sqrt{\rho R_0C_0/\mu}$ is necessary for inductive effects not to be considered. This evaluates to about $W < 5\mu\text{m}$ for $R_0C_0 = 100\text{psec}$ and our usual choice of physical parameters. If W is set to $4\sqrt{\rho R_0C_0/\mu}$, we obtain $T_p \simeq \sqrt{\mu\epsilon}\ell$. If W is greater than this value, each stage will become propagation limited and the delay will still be given by this expression independent of W . 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^2T/Z_0 , the energy per stage. Thus, we find that the energy per transmitted bit is given by

$$E \simeq \min \left(2\epsilon V^2\ell, 8\epsilon V^2 \sqrt{\frac{\rho R_0C_0}{\mu}} \frac{\ell}{W} \right). \quad (5.18)$$

The aspect ratio of a single stage is given by $\ell_{stage}/W = (1/4)\sqrt{R_0C_0/\rho\epsilon}$ which evaluates to ~ 2500 for $R_0C_0 = 100\text{psec}$. If $W \sim 1\mu\text{m}$, 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 $s = \sqrt{R_0C/C_0R} = \sqrt{R_0/C_0} \sqrt{\epsilon/\rho} \sqrt{t/h} w$ times that of a minimum sized transistor [4]. Remember that we took $t = h$ and $w = W/2$. With $R_0 = 20\text{K}\Omega$ and $C_0 = 5\text{fF}$ consistent with $R_0C_0 = 100\text{psec}$, and $W \sim 1\mu\text{m}$, we find $s \sim 35$ which we compare with ℓ_{stage}/W 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 5.2 and figure 5.3.

5.3.4 Superconducting interconnections

The propagation delay and characteristic impedance of a superconducting transmission line are essentially given by [116]

$$T_p = \sqrt{\mu\epsilon} \left[\frac{\lambda}{h} \coth\left(\frac{t}{\lambda}\right) + 1 \right]^{\frac{1}{2}} \ell \quad (5.19)$$

$$Z_0 = \sqrt{\frac{\mu}{\epsilon}} \frac{h}{w} \left[\frac{\lambda}{h} \coth\left(\frac{t}{\lambda}\right) + 1 \right]^{\frac{1}{2}} \quad (5.20)$$

where all parameters are defined as previously, except λ , which denotes the superconducting penetration depth throughout this section. We again refer to figure 5.6 and invoke similar geometrical constraints. Attenuation and dispersion are small enough to be safely ignored for the length scales in consideration [100]. 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 $T_c = 92.5$ K, absolute penetration depth $\lambda_0 = 1400$ Å, normal resistivity $\rho_n = 200 \mu\Omega$ cm operated at 77 K [100]. The value of the penetration depth at $T = 77$ K may be calculated as $\lambda = \lambda_0 / \sqrt{1 - (77/92.5)^4} = 1942 \simeq 2000$ 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 $J_{sc} = 8 \text{ mA}/\mu\text{m}$ was estimated [50] 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. [100] estimated $50 \text{ mA}/\mu\text{m}$ 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 $\lambda = 2000 \text{ \AA}$, the above mentioned values for J_{sc} are consistent with often cited values for the volume critical current density ($J_c = 10^6\text{-}10^7 \text{ A/cm}^2$). 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 equations 5.19 and 5.20 lead to

$$T_p E = \epsilon V^2 T \frac{w}{h} \ell. \quad (5.21)$$

This product does not depend on $W = 2w$ or the cross sectional area $A = WH \simeq 4wh$, 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,

$$T_p = \sqrt{\mu\epsilon} \ell = \frac{\ell}{v} \quad (5.22)$$

$$Z_0 = \frac{1}{2} \sqrt{\frac{\mu}{\epsilon}} \quad (5.23)$$

$$E = 2V^2 \sqrt{\frac{\epsilon}{\mu}} T. \quad (5.24)$$

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

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 $4V/(J_{sc}\sqrt{\mu/\epsilon}) \leq W$. If the critical current is high enough and the operating voltage low enough this condition is less restrictive than $W \geq 4\lambda(h/t)$. Even for $t = h$, with presently achievable critical currents as cited above, voltage values somewhat less than 1 V 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,

$$T_p = \sqrt{\mu\epsilon} \lambda \frac{\ell}{h^{\frac{1}{2}} t^{\frac{1}{2}}} \quad (5.25)$$

$$Z_0 = \sqrt{\frac{\mu}{\epsilon}} \lambda \frac{h^{\frac{1}{2}}}{wt^{\frac{1}{2}}}. \quad (5.26)$$

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

$$R_d = \frac{V}{J_c wt} > Z_0. \quad (5.27)$$

The lumped delay and energy are then expressed as $T_\ell = R_d C \ell$ [147] and $E = V^2 C \ell$,

$$T_\ell = \frac{\epsilon V}{J_c} \frac{\ell}{ht} = \frac{16\epsilon V}{J_c} \frac{\ell}{W^2} \frac{h}{t} \quad (5.28)$$

$$E = 2\epsilon V^2 \ell \quad (5.29)$$

where $W = 2w = 4h$ was used. We immediately observe from the above that $t = h$ is the optimal choice, leading to $T_\ell = (16\epsilon V/J_c)\ell/W^2$ in the region $W < 4V/(J_{sc}\sqrt{\mu/\epsilon})$. If the critical current condition is *not* violated, then using equations 5.25 and 5.26 the delay, characteristic impedance and energy are expressed as

$$T_p = \frac{4\lambda}{v} \frac{\ell}{W} \left(\frac{h}{t}\right)^{\frac{1}{2}} \quad (5.30)$$

$$Z_0 = \frac{1}{2} \sqrt{\frac{\mu}{\epsilon}} \frac{4\lambda}{W} \left(\frac{h}{t} \right)^{\frac{1}{2}} \quad (5.31)$$

$$E = \frac{V^2 T}{Z_0} \quad (5.32)$$

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 $W(t/h)^{1/2}$. Thus to minimize W we choose $t = h$, leading to

$$T_p = \frac{4\lambda}{v} \frac{\ell}{W} \quad (5.33)$$

$$Z_0 = \frac{1}{2} \sqrt{\frac{\mu}{\epsilon}} \frac{4\lambda}{W} \quad (5.34)$$

$$E = \frac{V^2 T}{Z_0} \quad (5.35)$$

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

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 H 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 5.3 and figure 5.4 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 W until $T_d = T_p$ (or $T_d = T_\ell$), ending up with a line occupying less space with the same delay.

5.4 Extensions

Though not presented, we have considered several extensions of our models. For instance, for a guided wave optical interconnection, T_ℓ is usually proportional to ℓ or ℓ^2 (depending on the applicable dispersion model [109] [67]) which may exceed T_d and dominate $T = \max(T_d, T_\ell)$ for very large ℓ . However, in most cases it should be possible to reduce T_ℓ below T_d for $\ell < 10$ m, so that we are justified in taking $T = T_d$.

It is also possible to derive models accounting for the effects of attenuation in superconducting interconnections. Again, such effects are often negligible for the length scales in consideration.

We also note some of the other effects that have not been accounted for in our models. We did not take into account the cost and non-ideal behavior associated with terminating conducting interconnections. We have ignored the effects of electromigration. Also, in their present form, our models do not allow for the effects of fan-out to be taken into consideration.

Part II

Optical Systems

Chapter 6

Lower Bound for the Communication Volume Required for an Optically Interconnected Array of Points

In the first part of this thesis, we introduced what we termed tube models of interconnections. Whereas the motivation for such a model is immediate when solid conducting wires or optical fibers/waveguides are employed, the reader may justly question whether free space optical interconnections can be made to fit into this model.

The limitations of optical communication in terms of information density are well understood in a telecommunication context, where two points are communicating with each other (figure 6.1). However, in a computation context, where many points are communicating with each other with possible overlap among many independent wave fields, it may not at first be evident what the limitations might be. One often comes across remarks in the literature referring to how the “3 dimensional non-interacting nature” of optical communication may enable one to overcome the “intrinsic limitations” of solid wires. At least one author has expressed hopes that the “space collapsing” properties of light waves might enable linear growth of system volume




Figure 6.1: *Communication through a finite surface.*

with increasing number of elements [62]. In what follows we show that the use of optical communication is similar to the use of solid wires in the sense that the total volume that must be allocated for optical communication must grow proportional to the total interconnection length of the system.

6.1 Introduction

The capacity of an information channel can be specified as the number of bits per second that can be transmitted across the channel. Since, in principle, this capacity can be increased by adding further parallel channels, another relevant quantity is the number of bits transmitted per second per cross sectional area. The maximum number of independent channels per unit area for optical fields was discussed in an information theoretical context by Gabor [51] and later elaborated by Winthrop [171], among others. Winthrop gave the following result¹ for the *number of degrees of freedom* F associated with a quasi-monochromatic optical wave field over a given

¹It may be noticed that we are omitting certain factors of 2 which Gabor and Winthrop originally included in writing similar relations. So as not to confuse our discussion we will consistently exclude these seldom applicable factors from our expressions, with the understanding that they can be readily included whenever appropriate. A discussion of the various sources which bring in an additional factor of 2 is given in appendix 16.7.




Figure 6.2: *The accessible Fourier area at a point P .* A frame of reference has been introduced with its origin coinciding with the point P and its z axis along the direction of the unit normal $\bar{n}$ to the surface S at that point. The surface S is not shown for clarity. The cone solid angle of allowed wave vectors is determined by the image of the aperture stop of the system as observed from the point P . The projection of the allowable wave vectors on the σ_x - σ_y plane determines the accessible Fourier area, as given by equation 6.2. The area of a cell of unit degree of freedom centered at P , denoted by $A(P)$, satisfies the relation $A(P)\chi(P) = 1$, as discussed in the text. (After Winthrop [171].)

surface S :

$$F = \int_S \chi(P) dS \quad (6.1)$$

where $\chi(P)$ is the accessible Fourier area at the point P relative to the surface S passing through it and is defined as

$$\chi(P) = \frac{1}{\lambda^2} \Omega_{\perp}(P) = \frac{1}{\lambda^2} \int_{\Omega(P)} \cos \theta d\Omega \quad (6.2)$$

with θ being the angle between the element of solid angle and a unit vector perpendicular to the surface in question (figure 6.2). λ is the wavelength of light used in the medium of propagation. $\Omega(P)$ denotes the cone of allowed wave vectors as limited by the image of the aperture stop of the system as observed from the point P . Notice that $0 \leq \chi(P) \leq \pi/\lambda^2$ since $\Omega(P)$ may at most be the complete hemisphere which

has projection π/λ^2 at a radius of $1/\lambda$. If $A(P)$ denotes the area associated with a cell of unit degree of freedom ($F = 1$) centered at point P on the surface and $\chi(P)$ is slowly varying at that point, then from equation 6.1 we may write

$$A(P)\chi(P) = 1. \quad (6.3)$$

This relationship is guaranteed to be preserved with free propagation [171] and in passing through arbitrary imaging elements as a consequence of Abbe's sine condition (known in its paraxial form as the Smith-Helmholtz-Lagrange invariant [15]), and is also closely connected to radiometric and thermodynamic considerations [126]. Given the accessible Fourier area at a given point on a surface along an imaging system—as determined by the image of the aperture stop as observed from that point—the area of a cell of unit degree of freedom is determined. Since $\chi(P)$ is bounded from above, $A(P)$ is bounded from below. (Determination of $\chi(P)$ for particular apertures is illustrated in appendix 16.8.)

Hence, the spatial information carrying capacity of optical wave fields is usually stated in terms of an area density as being related to communication between two regions in space, distinctly separated by a surface, as symbolically depicted in figure 6.1. A few authors [7] [43] have adapted results from VLSI complexity theory (for instance, see [158] [164]) based on solid wires to optically communicating systems by noting the fact that no more than a finite number of degrees of freedom exists over a finite surface. However, these studies considered communication among a planar array of points or through a planar surface, inhibiting the generality of the results.

Here we consider the problem of establishing optical communication among an arbitrary array of points. These 'points' may be optical switches or input and output transducers of electronic processing elements. We show that a lower bound for the total volume that must be allocated for communication is

$$\mathcal{V} = \frac{\lambda^2}{2\pi} \ell_{total} \quad (6.4)$$

where ℓ_{total} is the total interconnection length of the system, i.e. the sum of the lengths of all the component interconnects. This result accounts for all possible noninterfering overlap between independently excited optical wave fields. This essentially means that

for the purpose of calculating the volume or critical cross sections of the system, we may assume each independent optical information channel to have a minimum cross section of $\lambda^2/2\pi$ as if it were a solid wire.

In other words, the fact that any confined optical beam has a finite angular spread means that only a finite number can share the same physical volume. This in turn means that the results of area-volume complexity theory based on solid wires as a medium of communication are also applicable to systems employing electromagnetic wave propagation as a medium of communication.

6.2 Analysis

Our discussion is based on a scalar theory of light. We will assume all sources to emit spatially coherent quasi-monochromatic radiation of given center frequency ν . We will also assume that the information modulation bandwidth is greater than the linewidth of our light sources, so that the frequency deviation from the nominal optical carrier can be mainly attributed to the former effect.

We will consider the following model as illustrated in figure 6.3. It is assumed that the space allocated for communication is, in general, a multiply connected finite volume, the unshaded region in the figure. We are concerned with the problem of forming optical connections between specified transducers located at the surfaces of the shaded 'islands'. We assume that binary intensity modulation is used to impress information on the optical carriers emitted by the output transducers. The rate at which this information is generated (i.e. the temporal information modulation bandwidth) will usually be limited by the speed of the transducers or switching devices. Although it will be convenient to think of pairs of points being connected, the extension to fan-in and fan-out will be straightforward (appendix 16.9). The signals emitted by the output transducers are to be guided to the input transducers with the use of an arbitrary imaging system located inside the communication volume. The *communication length* of each connection formed may be defined naturally by multiplying the propagation delay along that interconnection with the velocity of propagation in the medium. Obviously, the communication length is greater than or




Figure 6.3: *Volume allocated for optical communication.* The unshaded region, of volume $\mathcal{V}$, is allocated for establishing optical communication between transducers located on the islands and/or enclosing surface.

equal to the Euclidean distance between the two points being connected.

We begin with the following relation for the spectral density of modes for a given volume $\mathcal{V}$, as given in many physics texts [46]:

$$\mathcal{D}(\nu) = \frac{8\pi\mathcal{V}}{c\lambda^2}, \quad (6.5)$$

c being the speed of light in the medium of propagation. As written, the above includes a factor of 2 to account for polarization. Also, we are not allowing double sidedness along the direction of propagation, as discussed in appendix 16.7. So, we divide by 4 to maintain consistency with our discussion:

$$\mathcal{D}(\nu) = \frac{2\pi\mathcal{V}}{c\lambda^2}. \quad (6.6)$$

Let $\Delta\nu \ll \nu$ denote the temporal bandwidth of the optical channels (which will most likely be limited by the speed of the transducers or optical switches). Then, according to the above equation, $(2\pi\mathcal{V}/c\lambda^2)\Delta\nu$ degrees of freedom will be at our disposal. It is of intuitive appeal to write this in an alternate form assuming that $\Delta\nu$ is a fraction ρ of the optical carrier frequency, i.e. $\Delta\nu = \rho\nu$ with $\rho \ll 1$. $\rho 2\pi\mathcal{V}/\lambda^3$ is then the maximum number of binary pulses that can be in transit at a given time in an optical communication network of volume $\mathcal{V}$. This quantity has been referred to as the *population capacity* of a communication network [63].

In order to derive our main result let us assume that we modulate our sources at a rate $B = \Delta\nu$ (i.e. we make full use of the available temporal bandwidth). This quantity will cancel out in our final result. Let each connection be numbered with the index $i = 1, 2, \dots, \mathcal{N}$, where it is assumed that there are $\mathcal{N}$ pairwise connections. If ℓ_i denotes the communication length of the i th connection, the number of bits in transit on this connection is

$$\frac{B}{1/\tau_i} = \frac{\Delta\nu}{1/\tau_i} = \frac{\Delta\nu}{c/\ell_i} = \frac{\Delta\nu\ell_i}{c} \quad (6.7)$$

where τ_i is the propagation delay along the i th connection. The total number of bits in transit at any given time on all connections is then

$$\sum_{i=1}^{\mathcal{N}} \frac{\Delta\nu\ell_i}{c} = \frac{\Delta\nu}{c} \sum_{i=1}^{\mathcal{N}} \ell_i. \quad (6.8)$$

This cannot exceed the number of degrees of freedom $\mathcal{D}(\nu)\Delta\nu$, thus we may write

$$\frac{2\pi\mathcal{V}}{c\lambda^2}\Delta\nu \geq \frac{\Delta\nu}{c} \sum_{i=1}^{\mathcal{N}} \ell_i = \frac{\Delta\nu}{c} \ell_{total} \quad (6.9)$$

or

$$\mathcal{V} \geq \frac{\lambda^2}{2\pi} \ell_{total}. \quad (6.10)$$

Thus we have shown that, under the stated assumptions, the total volume that must be allocated for optical communication must at least be $\lambda^2 \ell_{total}/2\pi$.

When it is the case that we are technologically confined to two dimensions, as in an integrated optic guided wave network, our results can be modified to show that the corresponding lower bound for the communication area is $\lambda \ell_{total}/\pi$.

6.3 A derivation based on tubes of unit degree of freedom

Winthrop develops the concept of tubes of unit degree of freedom [171]. He argues that the so called *structural information* in optical wave fields behaves like an incompressible fluid. Elementary tubes of flow may be defined such that each tube corresponds

to one spatial degree of freedom. Here we will present an intuitive derivation of our result based on this concept.

The problem of establishing communication among an array of points can be viewed as an imaging problem. Each source must be imaged onto the appropriate detector. Modulated radiation emanating from each source might undergo expansion, focusing, reflection etc. as it finds its way to a detector. We will imagine that a tube of unit degree of freedom emanates from each source and is guided to the appropriate detector. Let us denote the position along the path belonging to the i th connection as x_i . The cross sectional area and accessible Fourier area along the path of this tube satisfy the relation $A_i(x_i)\chi_i(x_i) = 1$ (equation 6.3). Tubes of unit degree of freedom associated with different sources may overlap provided their solid angle of wave vectors *do not* overlap. Also remember that the accessible Fourier area $\chi_i(x_i)$ is related to the solid angle of wave vectors $\Omega_i(x_i)$ by the relation $\chi_i(x_i) = \Omega_{\perp i}(x_i)/\lambda^2$, where $\Omega_{\perp i}(x_i)$ is the perpendicular projection of $\Omega_i(x_i)$ (equation 6.2).

Now, consider an arbitrary point in the communication volume (figure 6.4), surrounded by a small cubic volume element of diameter δ . This quantity has no physical significance and is merely an instrument in our counting argument. Several tubes of unit degree of freedom may overlap this volume element without interacting. We will choose δ small enough so that neither $A_i(x_i)$ nor $\chi_i(x_i)$ changes appreciably along this distance δ .

The i th tube will overlap $A_i(x_i)/\delta^2$ cubes over the interval $[x_i, x_i + \delta]$ along its path. However, these cubes will be available for use by other tubes with nonoverlapping wave vectors. To avoid multiple counting, let us introduce a sharing factor $sh_i(x_i) = \Omega_i(x_i)/2\pi$. This quantity accounts for the fact that this tube is only consuming a fraction of the total available solid angle 2π ($= 4\pi/2$, since we are not allowing double sidedness along the direction of propagation). Also, let us denote the distance x_i along the i th connection in units of δ as $x_i = j_i\delta$ and write $A_i(x_i) = A_{ij}$ etc. Now, summing along the path of the i th line and then over the $\mathcal{N}$ connections in our system, we obtain the total number of cubes consumed by all tubes

$$\sum_{i=1}^{\mathcal{N}} \sum_{j=1}^{\ell_i/\delta} sh_{ij} \frac{A_{ij}}{\delta^2}. \quad (6.11)$$




Figure 6.4: *A volume element in the communication volume.* A cubic volume element of diameter δ situated at an arbitrary point inside the volume allocated for optical communication is shown. In general, several independent tubes of unit degree of freedom will overlap the volume occupied by this element.

Multiplying this by δ^3 will give a lower bound for the volume $\mathcal{V}$. Using the definition of $sh_{ij} = \Omega_{ij}/2\pi$, $A_{ij} = \lambda^2/\Omega_{\perp ij}$, noting that $1 \leq \Omega_{ij}/\Omega_{\perp ij}$ and $\sum_{i=1}^{\mathcal{N}} \sum_{j=1}^{\ell_i/\delta} \delta = \sum_{i=1}^{\mathcal{N}} \ell_i = \ell_{total}$, we algebraically obtain $\mathcal{V} \geq \lambda^2 \ell_{total}/2\pi$, as before.

6.4 Discussion

Our main conclusion (that the minimum communication volume required for an optical interconnection network with total communication length ℓ_{total} is $\lambda^2 \ell_{total}/2\pi$) is stated in a global manner; it does not correspond to saying that each light path is confined to a cross section of $\lambda^2/2\pi$. Nevertheless, for the purpose of calculating lower bounds on volumes and cross sections, our result can be stated in another equivalent form if one imagines each light path to occupy constant width tubes with no further overlap allowed, i.e. by treating them like ‘solid wires’: the minimum effective nonoverlapping cross sectional area required for each independent spatial channel in an optical communication network is $\lambda^2/2\pi$.

It is not necessary that $f^\# \sim 1$ imaging be used in guiding the light emanated from the output transducers to the input transducers in order to achieve globally

an effective cross section of $\sim \lambda^2$. Use of higher $f^\#$'s means that, with a carefully designed system, the spread of allowed wave vectors can be smaller, so that through the use of different nonoverlapping regions of wave vector space, a larger number of light paths can noninterferingly overlap and share the same volume, despite the greater cross sectional area associated with each independent channel of information (i.e. each degree of freedom). This is a direct consequence of equation 6.3. Hence, even with transducers larger than the order of a wavelength it would be in principle possible to approach the lower bounds, provided that care is exercised so that the sources emit into only a single degree of freedom. Although the source will emit into a large volume of space, since the cone of wave vectors may be kept narrow, this volume is available for further usage by other beams.

One way of approximately approaching the bounds would be to use single mode waveguides with high numerical aperture, such that the guide cross sections are $\simeq \lambda^2$. Clearly, it is in principal possible to 'wire up' an arbitrary pattern of interconnections by using such waveguides. This serves as an existence proof that the lower bounds can be approached for an arbitrary pattern of interconnections, not being restricted to space invariant connections only. It further means that for $f^\# \sim 1$ imaging the advantage of being able to overlap the light paths and share the communication volume is only a small numerical factor. For larger $f^\#$'s, however, it is necessary to be able to overlap the light paths so as to efficiently utilize the available Fourier space.

The factor of 2π appearing in our equations is simply the solid angle associated with a hemisphere. Our bounds are tight in the sense that they may be closely approached if nearly complete utilization of this 2π of solid angle available for the wave vectors is accomplished at every point. In practice, squeezing out the last factor of 2π or so may not be practical. In general, if the volume required is roughly predicted by the above stated results we will call such a system *communication volume limited*. For this to hold for a certain communication architecture with the number of connections $\mathcal{N}$ as a parameter, we require that the volume be given by the above stated results, within a constant of the order of unity and independent of $\mathcal{N}$. This will be clarified through examples below and in chapter 7.

The reader may have noticed that until now we refrained from discussing how the communication lengths are specified. A full treatment employing the connectivity model of chapter 2 is postponed until chapter 10. Here we simply illustrate the major mechanism which determines ℓ_{total} for a communication volume limited system. For concreteness, let us assume that $\mathcal{N}$ connections are to be established among a more or less cubic array of points. From equation 6.10 we may write

$$\mathcal{V} \geq \frac{\lambda^2}{2\pi} \mathcal{N} \bar{\ell} \quad (6.12)$$

where $\bar{\ell}$ denotes the average interconnection length. For the purpose of this chapter $\bar{\ell}$ will be assumed to be proportional to $\mathcal{V}^{1/3}$. Thus with $\bar{\ell} = \kappa \mathcal{V}^{1/3}$ where κ is a constant, we can solve for the total volume and average propagation delay as $\mathcal{V} \geq (\lambda^2 \kappa \mathcal{N} / 2\pi)^{3/2}$ and

$$\tau_{ave} = \frac{\bar{\ell}}{c} \geq \left(\frac{\kappa^{3/2}}{\sqrt{2\pi}} \right) \frac{\mathcal{N}^{1/2}}{\nu} \quad (6.13)$$

which we write as

$$\frac{\mathcal{N}^{1/2}}{\tau_{ave}} \leq \left(\frac{\sqrt{2\pi}}{\kappa^{3/2}} \right) \nu \sim \nu \quad (6.14)$$

where ν is the frequency of the optical sources. The last equation represents a tradeoff between the propagation delay and the number of optical connections established. Our result may be considered to be a generalization of that given by Shamir for the more conventional configuration of communication between two planes [149], to an arbitrarily overlapping pattern of interconnections among a three dimensional array of points. Shamir has already noted that this result constitutes a fundamental limit for parallel processing involving global communication.

To illustrate a case which is *not* communication volume limited, let us assume that the transducers among which connections are to be established are constrained to lie on the surface of a sphere. Assume that there are a total of $2\mathcal{N}$ transducers with $\mathcal{N}$ connections to be made between them. We wish to determine the minimum radius $\mathcal{L}/2$ and volume $\mathcal{V}$ of the sphere as a function of the number of connections $\mathcal{N}$. Let the transducer areas d_{tr}^2 be expressed as $(f\lambda)^2$. The surface area, radius and hence volume of the sphere are constrained to minimum values of $4\pi(\mathcal{L}/2)^2 = 2\mathcal{N}d_{tr}^2 = 2\mathcal{N}f^2\lambda^2$,

$(\mathcal{L}/2) = (\mathcal{N}/2\pi)^{1/2} f\lambda$ and

$$\mathcal{V} = \frac{4\pi}{3} \left(\frac{\mathcal{L}}{2}\right)^3 = \frac{1}{3} \left(\frac{2}{\pi}\right)^{\frac{1}{2}} \mathcal{N}^{\frac{3}{2}} f^3 \lambda^3 \sim \mathcal{N}^{\frac{3}{2}} f^3 \lambda^3. \quad (6.15)$$

The lower bound to the communication volume is, using $\mathcal{V}_{lb} \geq \lambda^2 \mathcal{N} \bar{\ell} / 2\pi$

$$\mathcal{V}_{lb} = \frac{\lambda^2}{2\pi} \mathcal{N} \kappa \mathcal{L} = \frac{1}{\pi} \left(\frac{1}{2\pi}\right)^{\frac{1}{2}} \kappa \mathcal{N}^{\frac{3}{2}} f \lambda^3 \sim \kappa \mathcal{N}^{\frac{3}{2}} f \lambda^3 \quad (6.16)$$

where we used $\bar{\ell} = \kappa \mathcal{L}$ since $\mathcal{L}$ is the linear extent of the system. The ratio between the two previous equations for volume is approximately

$$\frac{\mathcal{V}}{\mathcal{V}_{lb}} \sim \frac{f^2}{\kappa}. \quad (6.17)$$

It is seen that for large f (i.e. transducer areas $\gg \lambda^2$) or small κ (i.e. $\bar{\ell} \ll$ system linear extent) we are doing considerably worse than as predicted by our lower bound. The situation will be even worse when $\bar{\ell}$ does not grow in proportion to $\mathcal{L}$. We will term such cases *transducer surface limited*. However, since one can arbitrarily increase the surface area enclosing a given volume (for instance by ‘wrinkling’ the surface rather than insisting on a sphere), given the freedom of rearranging the points to be connected in a more flexible manner, it should be possible to improve on this situation.

Most optical interconnection architectures that have been suggested for intra-chip or chip-to-chip communication for VLSI circuits involve a holographic optical element situated above a planar array of devices at a distance comparable to the linear extent of the device plane (which is necessary to be able to form connections between devices located at opposite ends of the layout) [59] [92]. Architectures based on conventional space invariant imaging configurations are communication volume limited only if $d_{tr} \sim \lambda$ and $\bar{\ell} \sim \mathcal{L}$. (An example of a system satisfying these conditions is an $f^\# \sim 1$ Fourier plane processor [56].) Since such architectures cannot provide an arbitrary pattern of connections, several authors have considered the use of space variant multifacet architectures [92] [41] [19]. However, these architectures have a very inhibitive growth rate of system volume (appendix 16.10) and are even farther away from being communication volume limited. We noted that high numerical aperture single mode

waveguides of cross section $\sim \lambda^2$ serve as an existence proof that an arbitrary pattern of connections may be achieved with communication volume limited systems. A problem of great practical interest is to devise free space architectures having this property. Such an architecture may involve a three dimensional layout of the points to be connected. We will suggest one such architecture in chapter 7.

6.5 Conclusions

We saw that for an *arbitrary* array of points communicating with each other, it is possible to view optical communication density limits in a global fashion by treating optical links as if they were solid wires of cross section $\lambda^2/2\pi$. This result accounts for all possible noninterfering overlap between independently excited optical wave fields. The main point is that any number of independent wave fields are allowed to overlap in coordinate space, or in Fourier space, but not both. In deriving this result no specific assumptions regarding the configuration of the points, the shape of the surface enclosing the communication volume or the imaging system were made. Thus results of area-volume complexity theory based on solid wires are also applicable to optically communicating systems.

As an alternative formulation of our result, we showed that the maximum number of binary pulses that can be in transit in an optical communication network is bounded by $\varrho 2\pi \mathcal{V}/\lambda^3$, where ϱ is the modulation bandwidth of the transducers normalized by the carrier frequency.

The utility of our global viewpoint is that it enables one to model the basic mechanisms which limit how closely one can pack an array of optically interconnected primitive computing elements to form a larger computing system. The advantage to be gained by using free space alternatives (i.e. overlap of independently excited wave fields possible) over guided wave alternatives (no overlap allowed except possibly at crossings) were seen to be (in a fundamental sense) not more than a factor of the order of unity, assuming $f^\# \sim 1$ imaging. This is because a significant fraction of the available Fourier space is already utilized and not much further overlap is allowed. As the $f^\#$'s in question increase, however, it becomes more and more important to

be able to non-interferingly overlap independent wave fields so as to make better utilization of the Fourier space, and hence the available modal volume.

Most suggested optical interconnection architectures consist of a planar array of devices. Such systems often do not enable the least possible system volume to be achieved. The problem of minimizing the volume will increase in importance as reduced device delays result in the speed of light being the dominant factor in determining the overall speed of large computing systems. Such being the case, it would be extremely beneficial to devise communication architectures which enable efficient utilization of *both* coordinate *and* Fourier spaces. We will address this problem in chapter 7.

Another important conclusion may be drawn from equation 6.10. Observe that the minimum communication volume is only linear in the total communication length. This may be interpreted as a consequence of the fact that unit cross sectional spatial response functions for electromagnetic propagation are of the form $\text{sinc}(x, y)$ or $\text{jinc}(r)$ [56] [18], functions whose self-convolutions are identical to themselves. Thus, apart from the effects of aberrations, the diffraction limited spot size does not increase upon cascading several identical imaging systems to ‘relay’ optical information over any distance. In other words, the effective cross section required per independent channel is (at least in a fundamental sense) independent of length. In contrast, the volume required for communication with conducting transmission lines is superlinear in distance. This is because longer lines must be made larger in cross section in order to maintain acceptable attenuation levels. Thus with increasing system sizes, the communication volume required for establishing optical interconnections will grow slower than that required for establishing conductor guided interconnections.

Chapter 7

Optimal Optical Interconnection Architectures

In this chapter we compare the system size of some optical interconnection architectures and introduce the folded multi-facet holographic interconnection architecture which has the potential to approach the minimum possible system size of *any* 2 dimensional architecture in providing an arbitrary pattern of interconnections among a 2 dimensional array of points. We also discuss a 3 dimensional version of this architecture.

7.1 Introduction

We assume that $\mathcal{N} = kN$ pairwise interconnections are to be established among a collection of N elements. For simplicity, the extension to fan-out is not considered. In accordance with our model, the layout area (or volume) will be expressed as Nd^2 (or Nd^3). d must be chosen large enough so that there is enough space to establish the desired interconnections. Of course, d must also be large enough so that there is enough space for the elements and satisfy heat removal requirements; however, in this chapter we concentrate on the value of d as set by communication requirements only. As introduced earlier, $\bar{\ell}$ will denote the average interconnection length of the layout, where ℓ denotes the length of a particular interconnection and the overbar

	A	B	C	D
area	-	-	-	$(k\kappa)^2 N^{2q} f^2 \lambda^2$
volume	$(k\kappa)^{\frac{3}{2}} N^{\frac{3q}{2}} f^3 \lambda^3$	$k\kappa N^{q+\frac{1}{2}} f^3 \lambda^3$	$(kN)^3 f^3 \lambda^3$	-
linear extent	$(k\kappa)^{\frac{1}{2}} N^{\frac{q}{2}} f \lambda$	$N^{\frac{1}{2}} f \lambda$	$kN f \lambda$	$k\kappa N^q f \lambda$

Table 7.1: *System size for some optical interconnection schemes for large N .* Columns A and D give the minimum system size achievable with *any* 3 or 2 dimensional optical architecture respectively. Column B gives the minimum for *any* architecture in which the points to be connected are constrained to lie on a plane but communication paths are allowed to leave the plane. Column C is for the reflective multi-facet architecture.

denotes averaging. We remember from chapter 2 that this quantity is expressed as $\bar{\ell} = \kappa N^{q-1/2} d$ for a 2 dimensional array of points, where κ is a coefficient which is often of the order of unity and q is related to the Rent exponent as $q = \max(p, 1/2)$ so that $1/2 \leq q \leq 1$. Similarly, remember that for a 3 dimensional array of points, the average connection length may be expressed as $\kappa N^{q-2/3} d$ with $q = \max(p, 2/3)$ so that $2/3 \leq q \leq 1$.

As usual, order of magnitude accuracy is sufficient for our purpose so that factors such as 2, $\sqrt{2}$ etc. will be ignored. In this chapter we will also ignore slowly varying logarithmic factors for simplicity.

7.2 System size considerations

We now refer to table 7.1 which gives the system size for various situations. λ denotes the optical wavelength. f is a dimensionless factor which in principle can approach the order of unity, but may be quite larger in practice. The results given in columns A, B and D will be derived within the framework of a more general analysis in chapter 10. The result given in column C for the reflective multi-facet architecture is derived in appendix 16.10.

The best possible 3 dimensional growth rate (column A) may be approached by using discrete fibers of diameter $\sim f\lambda$. However, the resulting system size would nevertheless be large because of the relatively large value of f . Similar comments

apply to the case where the points to be connected are constrained to lie on a plane (column B). Alternatively, it is possible to achieve this growth rate with a small value of f by using free space interconnections in conjunction with multiplexed holograms. However, this method not only results in poor diffraction efficiency, but also constrains the pattern of connections due to the ambiguity associated with the Bragg cone.

One way of achieving an arbitrary pattern of interconnections among a planar array of points is to use the reflective multi-facet architecture [92] (illustrated in part a. of figure 7.1) or one of its variants [43] [41]. However, due to diffraction considerations, the growth rate associated with this architecture (column C) is larger than the best possible (column B), although the value of f involved can be of the order of unity. In fact, unless $q = 1$, this growth rate is even worse than that achievable in 2 dimensions (column D).

The best possible growth rate for fully 2 dimensional layouts (column D) may be achieved by using waveguides with average effective line to line spacing of $\sim f\lambda$. However, the value of f must be relatively large due to crosstalk and routing considerations [76].

7.3 The folded multi-facet architecture

We now consider the folded multi-facet architecture based on the substrate-mode holographic system [3] [146] shown in part b. of figure 7.1. Such an imaging system will be used for each connection. In this manner we will be able to realize an arbitrary pattern of connections. (In certain situations involving a regular (perhaps space invariant) pattern of connections, as in Fourier plane filtering, it is possible to send more than one data channel through the imaging system, leading to a simpler design, as in [78].) This system is composed of two identical holographic optical elements (HOEs) which were recorded on the same plate. The first one, $\mathcal{H}_s$, collimates a coherent point source into a plane wave which is trapped inside the plate by total internal reflection. The second HOE, $\mathcal{H}_r$, focuses the collimated wave onto a detector. Since the holographic plate can be located very close to the source and the detector, and the light is guided inside the plate, this system can be very compact and easy




Figure 7.1: *Holographic optical interconnection architectures.* Only one source-detector pair and its associated facet(s) are shown for clarity. Part a. depicts the reflective multi-facet architecture. Part b. depicts the folded multi-facet architecture in its most primitive form. Part c. depicts a modified version which achieves the best possible growth rate of system size in 2 dimensions.

to use. Unlike the reflective multi-facet architecture, here the path length of light is proportional to the distance ℓ between the source and the detector. In fact, since the total internal reflection condition must be satisfied, the proportionality constant is merely an obliquity factor not much greater than unity (such as $\sqrt{2}$). Thus, the path length of light may be taken approximately equal to ℓ .

In order to achieve the least possible growth rate of system size, we will show below that, due to diffraction considerations, the area of both holograms should be chosen proportional to ℓ . Unfortunately, since the distance h between the device plane and the hologram plate, and the $f^\#$ of the sources are constant, the area of $\mathcal{H}_s$ is fixed and cannot be chosen proportional to ℓ .

To calculate the minimum value of d with this version of our architecture, we equate the total area occupied by the holograms associated with the kN interconnections to the total available area Nd^2

$$kN(\overline{d_s^2 + d_r^2(\ell)}) = Nd^2 \quad (7.1)$$

where $d_r(\ell) = f\lambda\ell/d_s$. f can approach the order of unity for diffraction limited operation. The overbar denotes averaging over all connections. Using the previously derived expression for the second moment $\langle \ell^2 \rangle = \zeta_2 N^p d^2$ we obtain

$$d^2 = k \left(d_s^2 + \frac{(f\lambda)^2 \zeta_2 N^p d^2}{d_s^2} \right) \quad (7.2)$$

which we can solve for d and optimize over d_s in order to obtain $d^2 = 4(f\lambda)^2 k^2 \zeta_2 N^p$ leading to a system linear extent of

$$N^{\frac{1}{2}} d \simeq k \zeta_2^{\frac{1}{2}} N^{\frac{p+1}{2}} f \lambda. \quad (7.3)$$

This growth rate is worse than the least possible $\propto N^q = N^{\max(p, 1/2)}$.

One way of achieving the least possible growth rate of system size is to use repeater holograms at equal intervals. However, we rule out this method since it would result in quickly decreasing interconnect efficiency with increasing length.

Thus, in order to achieve the least possible growth rate of system size, we modify the architecture discussed above as follows: h , and hence d_s and d_r (the diameters of $\mathcal{H}_s$ and $\mathcal{H}_r$ respectively), are fixed for connections of all lengths and are preferably

as small as possible. Two more holograms, $\mathcal{H}_a$ and $\mathcal{H}_b$, with diameters d_a and d_b are added. The route of the light from source to detector is shown in part c. of figure 7.1. $\mathcal{H}_a$ and $\mathcal{H}_b$ can be reflection holograms, or when the holographic emulsion is coated with a cover glass, transmission holograms. Unlike the area of $\mathcal{H}_s$ in part b., the areas of $\mathcal{H}_a$ and $\mathcal{H}_b$ are not fixed by the $f^\#$ of the sources and can be chosen in a manner so as to minimize total system size, as quantitatively discussed below.

To calculate the minimum value of d , we again equate the total area occupied by the holograms associated with the kN interconnections to the total available area Nd^2

$$kN(\overline{d_s^2 + d_r^2 + d_a^2(\ell) + d_b^2(\ell)}) = Nd^2 \quad (7.4)$$

where $d_b(\ell) = f\lambda\ell/d_a(\ell)$. d may be minimized by choosing $d_a^2(\ell) = d_b^2(\ell) = f\lambda\ell$. That is, we use larger facets for longer interconnections. With $d_s = d_r$ we obtain

$$d^2 = 2kd_s^2 + 2f\lambda k\kappa N^{q-\frac{1}{2}}d \quad (7.5)$$

which approximately leads to a system linear extent of

$$N^{\frac{1}{2}}d \simeq (kN)^{\frac{1}{2}}d_s + k\kappa N^q f\lambda \quad (7.6)$$

which becomes, for $q > 1/2$ and large N ,

$$N^{\frac{1}{2}}d \simeq k\kappa N^q f\lambda \quad (7.7)$$

Since the value of f need not be much greater than unity [3], the folded multifacet architecture can approach the best possible system size achievable by any 2 dimensional system both in terms of growth rate and numerical factors.

The interconnection scheme presented is similar to the use of waveguides in that each interconnection consumes area proportional to its length. Whereas guiding (i.e. focusing) takes place in a distributed manner along the length of a waveguide, in our architecture it is concentrated at the end points.

Although passage through four holograms is necessary, the overall diffraction efficiency can still be over 80% if thick phase holograms with individual diffraction efficiencies of $\sim 95\%$ are utilized. To prevent reflection from the glass-gelatin surface, the average refractive indices of these materials must be equal.

Several design issues must be addressed during practical implementation of our architecture, some of which we briefly mention. Let θ denote the angle the optical rays bouncing inside the glass plate make with the normal. θ must exceed $\sim 43^\circ$ so as to satisfy the total internal reflection condition. Another issue is that as they undergo several bounces, the rays impinge on holograms which belong to other connections. In order to avoid crosstalk, θ should differ from the Bragg angle of the impinged upon hologram. Since the holograms have very high obliquity, they are very angle sensitive. By proper selection of very thick emulsion and comparatively low depth of modulation, we may ensure that the hologram efficiency falls to nearly zero when θ differs more than $\pm 1^\circ$ - 2° from the Bragg angle. Since there is considerable flexibility in choosing θ , with careful design the light rays will impinge only at their destination holograms at the proper Bragg angle. It does not seem that the design problem of choosing θ appropriately for each interconnection is more formidable than that of routing solid wires or waveguides.

7.4 The 3 dimensional multi-facet architecture

Now we describe a 3 dimensional version of the architecture of the previous section. To the best of our knowledge, this is the first free space architecture whose growth rate of system size is equal to the least possible for any value of p . It also serves as a good illustration of how a free space optical system leads to the same growth rate as systems employing solid wires, as discussed in chapter 6.

The system is constructed in the form of a sandwich of a large number of layers, as illustrated in figure 7.2. The layers are separated by glass slabs. There are two kinds of layers. There are $N^{1/3}$ element layers, on which $N^{1/3} \times N^{1/3}$ elements are arrayed in cartesian manner. (In the figure, we choose $N = 3^3$ for illustration. In practice N will probably be much larger.) Between the element layers are an as yet unspecified number of holographic emulsion layers, spaced at a small distance h from each other. The holograms in these layers can be fabricated directly as computer generated holograms [104], can be recorded holographically as computer originated holograms [2] or can be recorded recursively from other aspherical holograms [3].




Figure 7.2: *The 3 dimensional multi-facet architecture I.* Side and top views of a system with $N = 3^3$ elements are shown.

Ultimately, in the more distant future, it may be possible to fabricate the whole system (including electronics, transducers and holograms) monolithically with MBE-like techniques. The height of the system is equal to its lateral linear extent $\mathcal{L}$, forming a perfect cube, so that there are $\simeq \mathcal{L}/h$ emulsion layers, assuming the elements are not very thick. The average refractive index of the emulsion material is equal to that of the glass slabs.

Each connection is established in at most 6 hops, using 5 holograms. The path of light for a connection between two transducers situated on elements located at diametrically opposite corners of the cube is symbolically depicted in figure 7.3.

The diameters of holograms $\mathcal{H}_s$ and $\mathcal{H}_r$, which we denote by d_s and d_r respectively, are independent of the length of the connection. Both d_s and d_r should be chosen as small as possible, however d_s must be chosen consistent with the $f^\#$ of the source and the distance of $\mathcal{H}_s$ to the source. Most probably, the holograms $\mathcal{H}_s$ and $\mathcal{H}_r$ would be located immediately above the sources and detectors. The diameters of holograms $\mathcal{H}_a$ (d_a), $\mathcal{H}_b$ (d_b) and $\mathcal{H}_c$ (d_c) are chosen according to the length of the connection. Larger holograms will be used for longer connections so as to compensate for the effects of diffraction.

For many connections, all six hops may not be necessary. However, the above




Figure 7.3: *The 3 dimensional multi-facet architecture II.* The path of light for a connection between two transducers situated on elements located at diametrically opposite corners of the cube is shown.

scheme ensures that a connection between any two elements can be established regardless of their relative position. Several other schemes likewise employing a bounded number of holograms regardless of distance are also possible. The crucial point is that the line-of sight-problem can be overcome by breaking the connection into a bounded number of hops whose total path length does not exceed the Euclidean distance between the two elements to be connected by more than a bounded geometrical factor of the order of unity. It is important to keep the number of hops bounded with increasing number of elements and system size so that the interconnect efficiency does not decrease with increasing N . For instance, if each hologram can be fabricated with $> 95\%$ diffraction efficiency, the interconnect efficiency will be $> 75\%$ regardless of N . Also note that there is great freedom in choosing the paths and hologram locations for the various connections.

Since there are $\mathcal{L}/h$ emulsion layers, the total area available for the holograms is $\mathcal{L}^2 \times \mathcal{L}/h = \mathcal{L}^3/h$. Equating this to the area required for the 5 holograms associated with each of the kN connections yields

$$\frac{\mathcal{L}^3}{h} = kN \overline{(d_s^2 + d_a^2(\ell) + d_b^2(\ell) + d_c^2(\ell) + d_r^2)} \quad (7.8)$$

where the overbar denotes averaging over all connections. The area of any 2 consecutive holograms $\mathcal{H}_x$ and $\mathcal{H}_y$ must be chosen such that $d_x = f\lambda\ell_{xy}/d_y$ where ℓ_{xy} is

the path length between them. ℓ_{xy} will not be different than the Euclidean distance between the elements by more than a numerical factor of the order of unity. Thus to simplify our analysis let us take $d_a^2 = d_b^2 = d_c^2 \simeq f\lambda\ell$ for all holograms and also assume $d_s = d_r = \text{constant}$. Then, within a numerical factor of the order of unity

$$\frac{\mathcal{L}^3}{h} \simeq kN(d_s^2 + f\lambda\kappa N^{q-1}\mathcal{L}) \quad (7.9)$$

since $\bar{\ell} = \kappa N^{q-2/3}d = \kappa N^{q-2/3}(L/N^{1/3}) = \kappa N^{q-1}\mathcal{L}$. Thus the system linear extent is found to satisfy

$$\mathcal{L} \simeq (kd_s^2h)^{\frac{1}{3}}N^{\frac{1}{3}} + (k\kappa)^{\frac{1}{2}}N^{\frac{q}{2}}\sqrt{(f\lambda)h} \quad (7.10)$$

which becomes for $q > 2/3$ and large N ,

$$\mathcal{L} \simeq (k\kappa)^{\frac{1}{2}}N^{\frac{q}{2}}\sqrt{(f\lambda)h} \quad (7.11)$$

which exhibits the least possible growth rate of any 3 dimensional system $\propto N^{q/2}$. Numerically, we are worse off than the smallest possible system size $(k\kappa)^{1/2}N^{q/2}f\lambda$ by a factor of $(h/f\lambda)^{1/2} \geq 1$. If $f \sim 1$ and $h \sim 100 \mu\text{m}$, then we are doing about an order of magnitude worse than the best possible. As an example, the linear extent $\mathcal{L}$ of a system with $N = 10^6$, $k = 50$ and $q = 0.8$ would be of the order of ten centimeters.

Again, as in the 2 dimensional case, care must be exercised so that the wave vector of any propagating beam does not unintentionally fall on the Bragg cone of a hologram it is supposed to pass through without interaction. Given that there is a very large degree of freedom at our disposal in choosing the locations of the holograms for each connection, this should not become a problem. It is possible to envision a computer aided design procedure that would determine the optimal locations for each connection.

The multiple emulsion layers between the element layers are necessary to achieve the $\propto N^{q/2}$ growth rate. If we allowed only one emulsion layer per element layer, then the total available area for the holograms would be $\mathcal{L}^2 N^{1/3}$. Equating this to $\simeq kN(f\lambda\kappa N^{q-1}\mathcal{L})$, we find a system linear extent $\mathcal{L} \propto N^{q-1/3}$, which is worse than the optimal $\propto N^{q/2}$ since $q \geq 2/3$.

In perspective, we can view the problem of optically interconnecting an array of points as an imaging problem. Conceptually, the sources can be imaged onto

their target detectors by tailoring the refractive index distribution $n(x, y, z)$ of the medium between them. In other words, an optically interconnected computer may be viewed as a volume hologram with the computing elements embedded into it. The architecture we have proposed is a particular way of realizing such a system.

7.5 Conclusion

In conclusion, the folded multi-facet architecture (which is essentially a ‘free space’ architecture) allows near diffraction limited operation and can potentially approach the smallest possible system size of any 2 dimensional system. This is difficult to achieve with waveguides, which must usually be packed at an effective line-to-line spacing much greater than $\sim \lambda$.

The 3 dimensional multi-facet architecture can achieve the least possible growth rate of system size of any architecture with increasing N . Numerically, this architecture may approach the best possible within an order of magnitude, depending on how small the spacing between the emulsion layers can be made.

Dr. Yaakov Amitai was a collaborator in devising the folded multi-facet architecture and its 3 dimensional version [132] [133]. He also experimentally demonstrated that diffraction limited operation is possible with such substrate mode holographic systems [3].

Chapter 8

The Optimal Electromagnetic Carrier Frequency Balancing Structural and Metrical Information Densities with Respect to Heat Removal Requirements

In this and the following chapter we will divert from the main theme of this thesis in order to discuss some relatively fundamental implications of our heat removal model for dissipative systems employing electromagnetic wave propagation as a medium of communication.

8.1 Introduction

Of the four basic forces in nature, only the electromagnetic force is effective on a scale comparable to biological organisms. The human visual and nervous systems are essentially based on electromagnetic interactions. Most man made computing

systems rely on electromagnetic phenomena for both regenerative nonlinearities and communication among their elements. The nonlinear operations may be based on electronic interaction.

Electromagnetic wave propagation is a most basic means of information transmission. In this chapter we will consider the use of modulated electromagnetic carrier waves to establish communication between the elements of a computing system. We assume that the signals are guided to their destinations with dielectric media only (including lenses, holograms, waveguides etc.), the use of conductors is excluded from consideration.

The elements of our computing system may be relatively simple switching devices or relatively complex processing elements. In any event we will assume them to be very small in size. We would like our overall system to be as compact as possible. Being able to handle vast amounts of information in a small volume is not only a merit in itself, as exemplified by the human eye, but also results in smaller speed of light limited communication delays between distant elements of our system. As discussed in previous chapters, spatial information density and heat removal are two major physical considerations which will limit how densely we can pack the elements of our computing system. These considerations are intimately tied together through the carrier frequency. Increasing the carrier frequency improves spatial information density; however it also increases the amount of heat we must remove per unit cross section and time. Since our heat removal ability is limited, there exists an optimal carrier frequency resulting in smallest possible system size and delay.

We will show that the optimal carrier wavelength is insensitive to system specific parameters (such as the number of elements, number of connections per element and interconnection topology) and lies near the infrared and visible bands of the spectrum.

Owing to the general nature of this discussion, our analysis is necessarily approximate. We have also preferred to leave out certain geometrical factors of the order of unity for simplicity and generality.

8.2 Preliminaries

In this section we review the two major physical mechanisms which will limit how closely we can pack the elements of our system.

Spatial information density We will assume that information transfer takes place along each independent spatial channel in the form of binary digital pulses impressed on a sinusoidal electromagnetic carrier of wavelength λ . We know from chapter 6 that the total volume allocated for communication in a system with total interconnection length ℓ_{total} must at least be

$$\sim \lambda^2 \ell_{total}. \quad (8.1)$$

Thus we see that the spatial (or *structural* [171]) information density can be increased by reducing the carrier wavelength λ , or in other words by increasing the frequency ν .

Energy dissipation and heat removal The energy of a single photon of electromagnetic radiation is given by $h\nu = hc/\lambda$ where h is Planck's constant and c is the speed of light. With increasing frequencies, the increasing energy of a single photon will require larger energies to maintain reliable communication, leading to a decrease in the so called *metrical* [51] information density. In general, based on statistical considerations $\vartheta \geq 1$ photons will be required per transmitted bit. The human eye, under optimum conditions, can detect as little as 100 photons per second [150]. If it is assumed that the eye can still differentiate events spaced about 100 msec apart under these conditions, this corresponds to 10 photons per bit. Properly designed shot noise limited systems may require ~ 100 -1000 photons, depending on the error rate we are willing to tolerate. We will follow Smith [151] in taking $\vartheta = 10^3$. The actual communication energies involved in practical systems may be quite larger than $E = \vartheta hc/\lambda$ due to various forms of overhead, as accounted for in chapter 5. However, for the purpose of this chapter, the energy per transmitted bit will be taken equal to that of ϑ photons.

The energy associated with each switching or regeneration event can usually be decreased with decreasing temperatures [79] [83]. Thus for sufficiently low temperatures and operating voltages, the communication energies will dominate electronic switching and regeneration energies. We are assuming that these energies are irreversibly dissipated.

Two remarks are appropriate at this point. First, we note that in principal, the energy associated with each transmitted bit may be reduced down to $E \simeq hB$, where B is the bit transmission rate [141] [178]. This will be discussed in chapter 9. Second, we are assuming that the energy associated with each transmitted bit of information is dissipated upon detection. It has been argued by many authors that this is not a fundamental necessity [9] [102] [49] [24]. Thus, our assumptions do not correspond to ultimate physical limitations, but rather to an idealization of dissipative shot noise limited communication systems.

The dissipated power must be removed from the system. Our heat removal ability is quantified by the parameter Q , introduced in chapter 3. Ultimately, how much we can increase Q is limited by material parameters, which are in turn related to atomic constants. Thus, even allowing for advances in materials; we do not expect that the value of Q can be further improved beyond a few more orders of magnitude, if we are to construct our processing systems from solid state materials under ‘earthly’ conditions. Whereas such a conclusion is too imprecise to have any engineering value, it will suffice for the purpose of this chapter.

8.3 Analysis

Our model computing system is to be constructed by establishing a prespecified pattern of connections among an array of $N^{1/3} \times N^{1/3} \times N^{1/3}$ elements laid out on a regular cubic grid with as yet unspecified lattice constant d . The system is confined in a cubic box of volume Nd^3 . Let k denote the average number of connections per element and $\bar{r}$ denote the average length of the connections in units of grid spacing. Ignoring numerical factors of the order of unity, we have $1 \leq \bar{r} \leq N^{1/3}$. The total

connection length in real units is $\ell_{total} = kN\bar{r}d$, since there are a total of kN connections. Using equation 8.1, we find that a volume of at least $\sim \lambda^2 kN\bar{r}d$ is necessary to establish these connections. We must choose d so that the system volume Nd^3 exceeds this minimum required volume plus the volume of the elements. Assuming the volume occupied by the elements to be negligible, we find that d must satisfy

$$d^2 \geq \lambda^2 k\bar{r}. \quad (8.2)$$

The total amount of power dissipated is $kNEB$ where $E = \vartheta hc/\lambda$ and B is the rate at which bits of data are being emitted into each connection. Let us refer to figure 3.1 and assume that we are able to remove Q of power through unit cross sectional area. Since the cross sectional area of our system is $N^{2/3}d^2$, the total amount of power we can remove is $QN^{2/3}d^2$. Requiring that this be greater than the total power dissipated, we find

$$d^2 \geq \frac{kN^{1/3}\vartheta hcB}{\lambda Q}. \quad (8.3)$$

We would like to minimize d both for sake of maintaining a compact system and—when it is a limiting factor—minimizing speed of light limited communication delays. The delay across the extent of the system is given by

$$\tau = \frac{N^{1/3}d}{c}. \quad (8.4)$$

Thus, given N and B , one may choose λ so as to minimize d and τ . Equating the right hand sides of equations 8.2 and 8.3 and solving for λ we obtain, with $\vartheta = 10^3$ and $Q = 100 \text{ W/cm}^2$,

$$\lambda = \left[\frac{ch}{Q} \right]^{1/3} \vartheta^{1/3} B^{1/3} \{1, N^{1/9}\} \sim 0.5 B^{1/3} \{1, N^{1/9}\} \mu\text{m} \quad (8.5)$$

where B is in Gbit/sec, and the notation $\{1, N^{1/9}\}$ means that any value between the two extremes is possible. Most present day computers operate at rates of 10-1000 Mbit/sec. Switches operating at ~ 100 Gbit/sec have been built. How large can N be? The connection machine [70] has $\sim 10^5$ elements. Some computers may have 10^8 transistors or more. The human brain has 10^{11} neurons [159]. Although both B

and N may vary over a large range, the optimal value of λ is quite insensitive to this variation.

Often B will not be specified independently. One way of relating B to the other parameters is to express it as a fraction of the carrier frequency ν which it may never exceed. Thus, using $B = \rho\nu$ and equations 8.2 and 8.3 we obtain

$$\lambda = \left[\frac{c^2 h}{Q} \right]^{\frac{1}{4}} \vartheta^{\frac{1}{4}} \rho^{\frac{1}{4}} \{1, N^{\frac{1}{12}}\} \sim 15 \rho^{\frac{1}{4}} \{1, N^{\frac{1}{12}}\} \mu\text{m}. \quad (8.6)$$

Even if $\rho = 1$, we find that $\lambda \sim 50 \mu\text{m}$ or so, corresponding to a frequency $\nu \sim 5 \times 10^{12}$ Hz. Presently, we are far from being able to achieve such modulation rates. $\rho \sim 10^{-3}$ would be more realistic, leading to $\lambda \sim 3 \{1, N^{\frac{1}{12}}\} \mu\text{m}$.

In most cases, the rate B at which pulses of information are emitted into each connection will be related to the cycle time of the system, determined by the worst case signal delay τ , as given by¹ equation 8.4. This approach was taken by Keyes [79]. Thus let us write $B = \beta/\tau$ where β is a constant. Most authors take $\beta = 1$ without discussion; however there is no reason why β should not be larger (i.e. we may allow pipelining). This time, using equations 8.2, 8.3 with equality, equations 8.4, 8.5 and $B = \beta/\tau$ we might solve for the optimal value of λ as

$$\lambda = \left[\frac{c^2 h}{Q} \right]^{\frac{1}{4}} \vartheta^{\frac{1}{4}} \frac{\beta^{\frac{1}{4}}}{k^{\frac{1}{8}} \{1, N^{\frac{1}{8}}\}} \sim 15 \frac{\beta^{\frac{1}{4}}}{k^{\frac{1}{8}} \{1, N^{\frac{1}{8}}\}} \mu\text{m}. \quad (8.7)$$

The optimal wavelength is very insensitive to both how it is calculated and to the various parameters. Equation 8.7 is plotted in figure 8.1 for $\bar{r} = 1$ and $\bar{r} = N^{1/3}$ with Q as a parameter.

8.4 Discussion

In this chapter we considered the use of a full 3 dimensional layout which may not always be possible to realize. Our analysis may be repeated for a fully 2 dimensional

¹Here we are implicitly assuming that the length of the longest interconnection is of the order of the linear extent of the system. We should also note that τ may not be of direct significance in some cases, such as a nearest neighbor connected system.




Figure 8.1: *The optimal carrier wavelength.* Equation 8.7 is plotted with Q as a parameter. We take $\vartheta = 10^3$ and $\beta^{1/4}/k^{1/8} = 1$. The horizontal lines correspond to $\bar{r} = 1$ and the slanted lines correspond to $\bar{r} = N^{1/3}$.

planar layout, for which the optimal wavelength is found to be of the order of $\sim 0.1 \mu\text{m}$. The optimal wavelength is smaller in this case because a fully 2 dimensional layout is much more restrictive in terms of providing communication. The optimal wavelength is also somewhat more strongly dependent on the system parameters in 2 dimensions. For layouts lying between these two extremes, the optimal wavelength will lie between this value and that found above for full 3 dimensional layouts.

Apart from system parameters with limited effect, the optimal carrier wavelength is given essentially by $\lambda \sim [c^2 h / Q]^{1/4}$ which is equal to a few microns for $Q = 100 \text{ W/cm}^2$. Human eyes, as well as those of many other living beings, operate at wavelengths around $0.5 \mu\text{m}$. Even bees, despite severe diffraction problems, have developed compound eyes enabling them to operate at only slightly higher frequencies [47]. There exist ‘windows’ at visible frequencies in the absorption spectrum of water (figure 8.2) [172] and the spectral distribution of solar radiation arriving at the earth (figure 8.3) [1] [71]. Many electronic energy levels, in particular semiconductor bandgaps, correspond to visible and infrared frequencies, enabling the construction of efficient sources and detectors at these frequencies. Although it would be unwarranted for us to draw any conclusions from these facts, one is tempted to think that they are not unrelated. After all, Q is intimately related to atomic constants through




Figure 8.2: *Absorption spectrum of sea water.* The broken lines indicate the visible spectrum [172].

material parameters.

8.5 Relation between S , B and N

Once the optimal carrier wavelength is calculated, it is an easy matter to calculate the minimum value of d and τ . Then we can show that the inverse signal delay $S = 1/\tau$ satisfies

$$S\beta^{\frac{1}{4}}N^{\frac{1}{3}}\bar{r}^{\frac{1}{8}} = k^{\frac{-3}{8}}\vartheta^{\frac{-1}{4}} \left[\frac{c^2 Q}{h} \right]^{\frac{1}{4}}. \quad (8.8)$$

We note that the tradeoff between S , $\beta = B/S$ and N is weakly dependent on the other parameters and especially on $\bar{r}$ (which depends on the pattern of connections). We will have more to say about a similar relationship between S , β and N in chapter 9.

8.6 Conclusion

We considered the use of electromagnetic radiation for communication among the elements of a computing system. In general, increasing the frequency of radiation enables higher information densities which potentially offer smaller system size and communication delays. Even when delay is not an issue, the ability to handle vast




Figure 8.3: *Spectral distribution of solar radiation reaching the earth.* The solid curve shows the distribution of solar radiation reaching the upper limits of the atmosphere. The dotted curve shows the transmittance of the atmosphere under normal conditions. The broken curve, obtained by multiplying the previous two, shows the distribution of radiation reaching sea level at normal incidence [1]. The ozone layer screens harmful ultraviolet radiation. Most of the infrared radiation with $\lambda > 1 \mu\text{m}$ is actually absorbed by CO_2 and water vapor (not shown), resulting in a greenhouse effect which keeps the surface of the earth warm. For a more detailed graph and discussion, see [71].

amounts of information in a compact volume is desirable in itself. Increasing the carrier frequency ν also increases the energy per transmitted bit. Thus, assuming these energies are dissipated, the amount of heat that must be removed from unit volume per unit time quickly increases. The maximum removable heat pertaining to 'earthly' conditions may be estimated based on physical considerations. After a certain frequency, we can no longer benefit from the high information and packing densities offered because we will fail to meet the heat removal requirements.

We have shown that there exists an optimum carrier frequency balancing information density and heat removal imposed bounds on element spacing resulting in smallest possible system size. This frequency is found to depend on the speed of light, which simply relates the frequency to the minimum resolvable dimension λ , on Planck's constant, which relates the frequency to the minimum resolvable energy, on our heat removal ability as quantified by Q and also on other system parameters. Since the optimal value of ν is only weakly dependent on all parameters, it was possible to obtain a system independent estimate of $\nu \sim 10^{14}$ which corresponds to the infrared and visible bands.

Needless to say, care must be exercised in arriving at any practical conclusions from our results which have been derived for an idealized computing system limited by simple physical considerations only.

Chapter 9

A Fundamental Consideration for Dissipative Computing

In this chapter we explore the consequences of assuming that the energy per transmitted bit is reduced down to that dictated by Heisenberg's uncertainty limit.

9.1 A fundamental consideration for dissipative computing

In chapter 8, we assumed the energy per transmitted bit to be proportional to the energy of a photon at the carrier frequency employed. This is not a fundamental necessity, but rather a consequence of the way we operate our communication channels. In principle, each photon can be put into one of many time slots, so that more than one bit of information can be conveyed per photon. The energy flow that must accompany the transfer of information has been discussed by many authors [20] [8] [102] [141]. Other references pertaining to quantum channel capacity theory may be found in [178].

We will take the following result, valid for a single spatial channel, as our starting point [141]:

$$(\text{rate of information flow})^2 \leq (\text{rate of energy flow}) \left(\frac{\pi}{3\hbar \ln^2 2} \right). \quad (9.1)$$

If we restrict our attention to binary digital transmission, the rate of information flow is simply the bit repetition rate B . The rate of energy flow is likewise expressed as EB in terms of our parameters. Ignoring the numerical factors, we then find that the minimum energy per transmitted bit is proportional to the bit repetition rate: $E \simeq hB$. Intuitively, this may be interpreted as a consequence of the Heisenberg uncertainty relation. If a temporal spread Δt is associated with each transmitted bit, this determines the bit repetition rate $B \simeq 1/\Delta t$ and $E \simeq h/\Delta t$, consistent with the above result.

The most significant property of the above result is that the energy per transmitted bit can be reduced by using a greater number of parallel spatial channels. If the total bit rate B is concentrated through a single spatial channel, we have $E \simeq hB$. If however, χ parallel channels are used, $E \simeq h(B/\chi)$. Thus there is a tradeoff between energy per transmitted bit and information density. Starting from this relationship, one can pursue an analysis similar to that of chapter 8 and derive expressions for the optimal value of χ etc. Here we will not pursue this exercise. Instead, we will explore the limitations imposed by heat removal requirements alone, assuming that the heat removal imposed system volume will be large enough to accommodate as many spatial channels as needed. (Remember that the space occupied by each spatial channel can be reduced by increasing the carrier frequency, which has no effect on the energy according to the assumptions of this chapter.)

As in chapter 8, we are assuming the energy flowing through each independent spatial channel to be dissipated upon detection. If there are a total of $\mathcal{N} = kN$ connections, the total power dissipation is $\mathcal{N}EB$, implying a system linear extent of $\mathcal{L} \geq (\mathcal{N}EB/Q)^{1/2}$ (regardless of whether the system is 2 or 3 dimensional). The delay across the extent of the system satisfies $\tau \geq \mathcal{L}/c$. Of course, the signal delay cannot be less than the temporal uncertainty Δt . Thus

$$\tau \geq \min \left(\frac{1}{c} \left(\frac{\mathcal{N}EB}{Q} \right)^{\frac{1}{2}}, \Delta t \right). \quad (9.2)$$

Since $\Delta t \simeq h/E$, τ is bounded from below. Now, it is an easy exercise to show that

the inverse signal delay $S = 1/\tau$, $\beta = B/S$ and $\mathcal{N}$ satisfy the following relation:

$$S\beta^{\frac{1}{4}}\mathcal{N}^{\frac{1}{4}} \leq \left[\frac{c^2 Q}{h} \right]^{\frac{1}{4}}. \quad (9.3)$$

The inverse delay is found to be very weakly dependent on the number of connections and the value of Q . For this reason, it is possible to put an approximate upper bound on S despite the fact that it is difficult to put a precise upper bound on Q . Based on our previous discussion of heat removal, we may write

$$S\beta^{\frac{1}{4}}\mathcal{N}^{\frac{1}{4}} \sim 10^{15} \text{ sec}^{-1} \quad (9.4)$$

independent of all system dependent parameters. Even a several order of magnitude improvement in Q , which seems unlikely, would have little effect on this expression.

The fact that the right hand side of equation 9.3 is weakly dependent on Q has a simple dimensional interpretation. If we start with the quantities c , h and Q and form a quantity with dimension of time, we find that the proper combination involves the fourth root of Q .

We do not wish to give the impression that equation 9.3 is unsurpassable. This result relies on the assumptions that the energy flow per transmitted bit E is related to the temporal spread of each transmitted bit Δt through the relation $E\Delta t \simeq h$ and that this energy is dissipated upon detection. As mentioned in chapter 8, it has been argued by several authors that dissipation is not a fundamental necessity [33].

Part III

Physical Limits to Communication in Computation

Chapter 10

Basic Analysis

In this chapter we will employ the models presented in the first part of this thesis to investigate the tradeoffs between the quantities S (inverse signal delay), B (bit repetition rate) and N (number of elements) and the cost of power and space for each interconnection medium.

10.1 Optical interconnections

10.1.1 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 [85] [84] [52] [5]. 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 sectional area’. 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:

$$\max(Nk\chi\bar{\ell}W^{e-1}, Nd_d^e) \simeq Nk\chi\bar{\ell}W^{e-1} + Nd_d^e = Nd^e. \quad (10.1)$$

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




Figure 10.1: *Cell size for communication limited systems.* We take $k = 5$, $\chi = 1$ and $W = 2\lambda$.

Ignoring d_d we find

$$d = (k\chi\bar{r})^{\frac{1}{e-1}} W. \quad (10.2)$$

Of course, d may never actually be less than d_d . Note that the same result can be obtained directly by equating $K\chi W^{e-1}$ to the cell cross section d^{e-1} , as discussed in chapter 2. 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* [63]. Space dilation occurs when $n > e$. Equation 10.2 is plotted in figure 10.1. Given k and d_d , we can use these plots to predict beyond what value of N the system volume will be determined by communication requirements, rather than by element size. The linear extent of the system may be obtained as

$$N^{\frac{1}{e}} d = N^{\frac{1}{e}} (k\chi\bar{r})^{\frac{1}{e-1}} W. \quad (10.3)$$

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 [159]). 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

$$Sr_{max}\bar{r}^{\frac{1}{e-1}} = \left(\frac{c}{2\lambda}\right) (k\chi)^{\frac{-1}{e-1}} \quad (10.4)$$

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 $r_{max} \simeq N^{1/e}$ we obtain

$$SN^{\frac{n-1}{n(e-1)}} = \left(\frac{c}{2\lambda}\right) (k\chi\kappa)^{\frac{-1}{e-1}}. \quad (10.5)$$

When $BT_r \leq 1$, we simply set $\chi = 1$. When $BT_r > 1$, we must choose¹ $\chi = BT_r$ since a single physical channel is incapable of transmitting information at a rate of B . Then

$$SB^{\frac{1}{e-1}} r_{max} \bar{r}^{\frac{1}{e-1}} = \left(\frac{c}{2\lambda}\right) (kT_r)^{\frac{-1}{e-1}} \quad BT_r \geq 1 \quad (10.6)$$

which becomes for $n > e$,

$$SB^{\frac{1}{e-1}} N^{\frac{n-1}{n(e-1)}} = \left(\frac{c}{2\lambda}\right) (kT_r\kappa)^{\frac{-1}{e-1}} \quad BT_r \geq 1. \quad (10.7)$$

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 10.5 with $\chi = 1$ or equation 10.7). If p is large, element size is small and/or devices slow, the central region may disappear.

We remind the reader that the elements must be at least large enough to accommodate $\sim k\chi$ transducers. Also, if an m -fold reduction in T_r was made possible by wavelength division multiplexing of m distinct wavelength sources, this number must be further multiplied by m .

If feasible, the use of multiple layers can contribute to 2 dimensional system performance. The width of a cell d must now satisfy $d \geq \max(K\chi/M, 1)W$ where M denotes the number of layers [85] and since d must at least be wide enough to admit the passage of one physical channel. Of course, d will have to be large enough to satisfy several other requirements, including $d \geq d_d$. It is important to note that there is a maximum useful value of M . Assuming this maximum useful value is not exceeded, we may write $d = K'W$ where $K' = K\chi/M$ is the number of physical wiring tracks

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

per cell per layer. The right hand sides of the above equations (10.5 and 10.7) are improved by a factor of M . If the number of layers is large, the effects of vertical runs must be taken into account. This is considered in detail in appendix 16.11.

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

$$\tau_L = \tau + \frac{L}{B} = \frac{1}{S} + \frac{L}{B}. \quad (10.8)$$

Let us assume $N = 10^6$, $k = 10$, $e = 2$, $n = 3$, $T = T_r = 1 \text{ nsec}$ and $L = 20$. Using equation 10.7 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.

10.1.2 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.

2 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 equation 10.7 may be derived as

$$SB^{\frac{1}{2}}N^{\frac{1}{2}} = \frac{c}{(E/Q)^{\frac{1}{2}}}k^{\frac{-1}{2}}. \quad (10.9)$$

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

$$\frac{1}{S} = \frac{N^{\frac{1}{2}}d}{c} = \frac{N^{\frac{1}{2}}}{c} \max \left(k\chi\kappa N^{p-\frac{1}{2}}(2\lambda), \left(\frac{kEB}{Q} \right)^{\frac{1}{2}} \right). \quad (10.10)$$




Figure 10.2: *The effect of heat removal requirements in 2 dimensions.* We take $k = 5$. 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.

where $\chi = \max(1, BT_r)$. Of course, as always, S can never be greater than $c/N^{1/2}d_d$ or $1/T$. Notice that when the system is heat removal limited $1/S \propto (NB)^{1/2}$, whereas when the system is wireability limited $1/S \propto N^p \max(1, BT_r)$. Equation 10.10 is plotted in figure 10.2 with EB/Q as 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 than heat removal limited. The critical value of N beyond which heat removal is no longer a limiting factor is plotted in figure 10.3 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 disappears.

As we have discussed in chapter 7 and [134] [132], many seemingly 3 dimensional optical architectures actually impose wireability requirements similar to 2 dimensional systems. In particular, certain multi-facet holographic architectures [92] [41] [43] can be very inhibitive. In fact, when $p < 1$ and the longest interconnection is 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




Figure 10.3: *Critical value of N beyond which heat removal is not a limiting factor.* We take $k = 5$ and assume T_r is small so that $\chi = 1$. EB/Q is varied over the same range as in the previous figure.

of the system grows as $\propto N$ as opposed to $\propto N^p$ which we have found for the fully 2 dimensional case (equation 10.5). Because of their flexibility in providing an arbitrary pattern of connections, these particular 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.

3 dimensions

Just as in the 2 dimensional case, Q will be specified as the power which we can remove per unit cross sectional area. 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

$$Qd^2 \geq kEBN^{\frac{1}{3}}. \quad (10.11)$$

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




Figure 10.4: *The effect of heat removal requirements in 3 dimensions.* We take $k = 5$. d_d , T_d and T_r are assumed small enough to have no effect. The second term of equation 10.12 dominates over the whole range.

Thus, using the above constraint and equation 10.2 we can show

$$\frac{1}{S} = \frac{N^{\frac{1}{3}}}{c} \max \left((k\chi\kappa N^{p-\frac{2}{3}})^{\frac{1}{2}}(2\lambda), \left(\frac{kEBN^{\frac{1}{3}}}{Q} \right)^{\frac{1}{2}} \right) \quad (10.12)$$

which is plotted in figure 10.4. 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.

On the other hand, if B is decreased according to the relation $B = \beta S$ with constant β , it is easy to show that wireability requirements dominate heat removal requirements when $p > 2/3$. What we are trying to emphasize is that whether heat removal or wireability requirements will become more important with increasing N depends also on how B is related to N and/or S .

10.1.3 Asymptotic properties

We make several observations regarding equation 10.7. For given N , B can be arbitrarily increased² 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 equation 10.7, the tradeoff between HB and S may be written in the transparent form

$$S[HB]^{\frac{1}{e-1}} = \left(\frac{c}{2\lambda}\right) T_r^{\frac{-1}{e-1}}. \quad (10.13)$$

Unlike B , S *cannot* be arbitrarily increased by reducing B , since once B drops below $1/T_r$, equation 10.5 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

$$HS \propto N^{\frac{p(e-2)}{(e-1)}}. \quad (10.14)$$

We observe that HS 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 larger 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 equation 10.10 or equation 10.12, 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

$$HS \propto N^{p-\frac{1}{2}}. \quad (10.15)$$

The resultant growth rate of HS is thus the slower of those given by equation 10.14 and the above.

It is also possible to consider other figure of merit functions of signal delay, bandwidth, system size and power dissipation and discuss their numerical and asymptotic

²Of course, as far as the interconnection network is concerned, remember that we are assuming T_{rd} to be negligibly small.

properties and derive a variety of miscellaneous conclusions. We will not further pursue such exercises.

10.2 Normally conducting interconnections

10.2.1 Relations between S , B and N

In order to accommodate $K = k\bar{r}$ wiring tracks, the linear extent of a cell d must satisfy $d^{e-1} \geq KW^{e-1}$ [85]. We must also satisfy $d \geq d_d$ and the heat removal requirement, which will be discussed later. We are free in choosing W provided it exceeds a certain minimum manufacturable value W_{min} . If d_d is small and heat removal is not an issue, we would prefer to set W to this minimum so as to make d and the overall system as small as possible. In this case, $d^{e-1} = KW_{min}^{e-1}$. However, if element size or heat removal require that we set $d^{e-1} > KW_{min}^{e-1}$, we will agree to increase W until $d^{e-1} = KW^{e-1}$. If d and hence the lengths of the lines are already set by factors other than wiring density, we increase W so as to fill up available space. In this way, we reduce the resistance of the lines as much as possible. Despite the fact that increasing W 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 W 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 $d^{e-1} = KW^{e-1}$ and equation 5.6 with $\ell_{max} = r_{max}d$ we immediately obtain

$$T_\ell = (16\rho\epsilon)r_{max}^2(k\bar{r})^{\frac{2}{e-1}}. \quad (10.16)$$

Masaki previously derived similar relationships between wire delay and connection length [114]. When $T_\ell \geq T_d$ so that $T = T_\ell$, the maximum value of B satisfies $B = 1/T_\ell$ or

$$Br_{max}^2\bar{r}^{\frac{2}{e-1}} = (16\rho\epsilon)^{-1}k^{\frac{2}{e-1}} \quad BT_d \leq 1 \quad (10.17)$$

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

$$BN^{\frac{2(n-1)}{n(e-1)}} = (16\rho\epsilon)^{-1}(k\kappa)^{\frac{2}{e-1}} \quad BT_d \leq 1. \quad (10.18)$$

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 $\chi = BT_d > 1$ parallel 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 $T_d = T_\ell = (16\rho\epsilon)\ell_{max}^2/W^2$, $d^{e-1} = \chi KW^{e-1}$ and $\chi = BT_d$ we obtain

$$B^{\frac{2}{e-1}} r_{max}^2 \bar{r}^{\frac{2}{e-1}} = (16\rho\epsilon)^{-1} k^{\frac{-2}{e-1}} T_d^{\frac{e-3}{e-1}} \quad BT_d \geq 1 \quad (10.19)$$

which becomes for $e < n$,

$$B^{\frac{2}{e-1}} N^{\frac{2(n-1)}{n(e-1)}} = (16\rho\epsilon)^{-1} (k\kappa)^{\frac{-2}{e-1}} T_d^{\frac{e-3}{e-1}} \quad BT_d \geq 1. \quad (10.20)$$

The above assumes the use of a constant W for lines of all lengths. Actually, since we are only interested in the delay and pulse width along the longest (worst case) connection³, we may make shorter lines narrower with the objective of reducing cell size. Thus, for the case $BT_d \leq 1$, let us set

$$\frac{1}{B} = T_\ell = (16\rho\epsilon) \frac{r^2 d^2}{W(r)^2} = \text{constant} \quad (10.21)$$

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) $rdW(r)^{e-1}$, the minimum value of d must satisfy

$$d^e = \int rdW(r)^{e-1} g(r) dr. \quad (10.22)$$

Solving for $W(r)$ from equation 10.21 and performing the integration we find

$$\frac{1}{B} = T_\ell = (16\rho\epsilon)(k\langle r^e \rangle)^{\frac{2}{e-1}} \quad BT_d \leq 1 \quad (10.23)$$

leading to, for $e < n$,

$$BN^{\frac{2(n-1)}{n(e-1)}} = (16\rho\epsilon)^{-1} (k\zeta_e)^{\frac{-2}{e-1}} \quad BT_d \leq 1 \quad (10.24)$$

which represents an improvement over equation 10.18 by only a constant factor! The asymptotic dependence of B on N remains unchanged. Similarly, when $BT_d \geq 1$,

³The average signal delay τ_{ave} is discussed in appendix 16.12.

one can show that equation 10.20 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 equation 10.21.

The use of multiple layers in 2 dimensions contributes greatly to performance. The number of layers M may exceed the order of ~ 10 [13]. 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 10.18 and 10.20 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 10.5 along with the corresponding relation derived using equation 10.21. The improvement possible using nonuniform linewidths is greater when p is small and less 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 W , provided W 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 [145].

It is also possible to show that the bisection-bandwidth product HB is given by

$$HB = (k\kappa N^p)^{\frac{\epsilon-3}{\epsilon-1}} (16\rho\epsilon)^{-1} \quad BT_d \leq 1 \quad (10.25)$$

$$HB = (16\rho\epsilon)^{\frac{1-\epsilon}{2}} T_d^{\frac{\epsilon-3}{2}} \quad BT_d \geq 1. \quad (10.26)$$

We will further discuss this quantity later.

Until now, we refrained from mentioning S . By definition, S may never exceed $1/\max(T_\ell, T_d, T_p)$. The 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 $W^2 < 16\rho\epsilon v\ell$. As we scale down the system photographically, all linear dimensions are decreased in proportion. Thus, below a certain critical W , this condition is satisfied so that propagation effects (i.e. inductive effects) need not be considered. Indeed, downscaling is recognized as a




Figure 10.5: B versus N for normally conducting interconnections. The branches correspond to different values of T_d (0.1, 1 and 10 nsec). We take $k = 5$. $M = 10$ layers has been assumed for the 2 dimensional cases.




Figure 10.6: *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.

useful tactic to ensure that S is not worse than $1/\max(T_\ell, T_d) = 1/\max(1/B, T_d)$ [114]. But is this critical value of W manufacturable? Using the above condition with $\ell_{\max} = r_{\max}d$ and $d = (k\chi\bar{r})^{1/(e-1)}W$, we can show that for the longest line not to be propagation limited, the scale must be reduced down to

$$W \leq 16\rho\epsilon\nu r_{\max}(k\chi\bar{r})^{\frac{1}{e-1}} \quad (10.27)$$

which is plotted in figure 10.6 for $e < n$ along with the corresponding cell size d . These W 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 10.6, 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.

10.2.2 Heat removal

2 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

$$Qd^2 \geq k(2\epsilon V^2 \bar{\ell})B = 2\epsilon V^2 k \bar{r} dB \quad (10.28)$$

$$d \geq \frac{2\epsilon V^2}{Q} k \bar{r} B \quad (10.29)$$

where $\bar{\ell} = \bar{r}d$. Let us assume $T_d \rightarrow 0$ and that the maximum possible bit repetition rate is employed, as given by equation 10.18. Thus, we may calculate the minimum value of d as set by heat removal. For $e < n$ we find

$$d = \frac{2\epsilon V^2}{Q(16\rho\epsilon)} (k\kappa)^{-1} M^2 N^{-p-\frac{1}{2}} \quad (10.30)$$

which quickly drops below the critical cell size presented in figure 10.6 for voltages $V \sim 1$ V and the modest $Q = 1$ W/cm². Thus, with increasing N , heat removal is not a limiting factor in 2 dimensions.

Figure 10.7 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




Figure 10.7: *Comparison of optical and normally conducting interconnections.* We take $k = 5$ 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 ranges from 10^{-12} m^2 to 10^{-4} m^2 in increments of 10^{-2} m^2 , as in figure 10.2.

conducting case. The curves for the optical case are in terms of the parameter EB/Q , as in figure 10.2.

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 10.8 we set B to the largest possible value allowed by normal conductors.

In general we observe that there is a critical value of N 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: keep linewidths constant and suffer quadratic increase of delay; or: 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




Figure 10.8: Comparison of optical and normally conducting interconnections for same values of B . Similar parameter values are taken as in the previous figure. E/Q ranges from $10^{-21} \text{ m}^2 \text{ sec}$ to $10^{-13} \text{ m}^2 \text{ sec}$ in increments of $10^{-2} \text{ m}^2 \text{ sec}$.

[134]. This discussion is easily quantified. If the linewidths W are kept constant with increasing system size, we eventually have $\tau \propto H^2$. If the linewidths are increased according to $W^2 \propto \ell_{\max}$ so as to maintain length proportionate delay so that $\tau \propto \ell_{\max} \simeq HW$, we again find $\tau \propto H^2$ since $\ell_{\max} \simeq HW$ implies $W \propto H$.

For large p , certain multi-facet holographic optical architectures [43] [132], as we have discussed in an earlier section, exhibit similar behavior to 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 .

As discussed above, heat removal quickly ceases to be a problem for 2 dimensional normally conducting layouts, in the sense that it has no effect on the resulting signal delay. Nevertheless we present the following analysis which enables us to determine the scale of the system as set by heat removal considerations.

We will assume that pulses of identical temporal width are launched into all lines regardless of their length. Thus the minimum value of this pulse width is set by the longest connection. (In principle, there is nothing that stops us from launching shorter pulses into the shorter lines, resulting in some energy savings. The following analysis may be modified for this case.) The minimum pulse width for the longest

line is given by

$$T = \max(T_d, T_\ell) = \max\left(T_d, (16\rho\epsilon)\frac{\ell_{max}^2}{W^2}\right) \quad (10.31)$$

where $\ell_{max} = (k\kappa N^p/M)W$, provided the number of layers is not too large. The above expression for T is independent of the scale of the system. Pulses of this duration are emitted into lines of all lengths. According to our tube model, the shorter lines, for which $T \geq T_p$, will be left unterminated, whereas the longer lines, for which $T < T_p$, will be terminated. Let us denote the breakeven length (in grid units) as r_x . Thus, lines for which $r_x d/v \leq T$ will be left unterminated and those for which $r_x d/v > T$ will be terminated. If the length $r_{max} \simeq N^{1/2}$ of the longest line in our system satisfies $r_{max}d \simeq N^{1/2}d \leq Tv$, all lines will be unterminated.

The problem is that we do not initially know d , which depends on the total power dissipated, which in turn depends not only on d but also on what fraction of the lines are unterminated. Let us initially assume that $N^{1/2}d \leq Tv$ so that all lines are unterminated. Then, the energy per bit is given by $2\epsilon V^2 \ell$, leading to (see equation 10.29)

$$d \geq \frac{2\epsilon V^2 k\kappa N^{p-\frac{1}{2}} B}{Q}. \quad (10.32)$$

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

$$Qd^2 \geq \left[\int_1^{r_x} 2\epsilon V^2 r dg(r) dr + \int_{r_x}^{r_{max}} 2\epsilon V^2 vT g(r) dr \right] B \quad (10.33)$$

which becomes

$$Qd^2 \geq k \left[2\epsilon V^2 \kappa r_x^{2p-1} d + 2\epsilon V^2 vT r_x^{2(p-1)} z \right] B \quad (10.34)$$

where $z = (1 - r_x^2/r_{max}^2) \leq 1$. Now, using $r_x = vT/d$, it is possible to solve for d as

$$d^{2p} \geq \frac{2\epsilon V^2 (vT)^{2p-1} k[\kappa + z] B}{Q}. \quad (10.35)$$

Since $z \leq 1$, we can replace $\kappa + z \simeq \kappa$ with little error. It is interesting to note that this expression forms continuity with equation 10.32 at $N^{1/2}d = vT$. Combining the




Figure 10.9: *Cell size as set by heat removal.* B is set to its largest possible value. We take $k = 5$, $M = 10$ and assume that T_d is negligibly small.

two expressions, we may write the minimum interelement spacing d as set by heat removal in the form

$$d = \min \left(\frac{2\epsilon V^2 k \kappa N^{p-\frac{1}{2}} B}{Q}, \left(\frac{2\epsilon V^2 (vT)^{2p-1} k \kappa B}{Q} \right)^{\frac{1}{2p}} \right). \quad (10.36)$$

This expression is plotted in figure 10.9. The system linear extent is of course simply given by $\mathcal{L} = N^{1/2}d$. The total power dissipation may be likewise expressed as

$$\mathcal{P} = 2\epsilon V^2 \min(\mathcal{L}, (vT)^{2p-1} \mathcal{L}^{2(1-p)}) k \kappa N^p B. \quad (10.37)$$

The heat removal limited normally conducting case is of historical importance and has been subject to many previous studies, most notably by Keyes [80] [81] [86].

3 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. 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 c. of figure 10.6. Thus, the value of S may be quite worse than $\min(B, 1/T_d)$. Of




Figure 10.10: Comparison of optical and normally conducting interconnections in 3 dimensions. We take $k = 5$, $p = 0.8$ 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 varies from 10^{-12} m^2 to 10^{-4} m^2 in increments of 10^{-2} m^2 , as in figure 10.4.

course, for ever increasing values of N , heat removal will eventually cease to be an issue since B necessarily decreases, and wiring requirements become more stringent. Before embarking on a complete analysis, let us pretend that the system *can* be downscaled enough to eliminate propagation effects (as might be possible by reducing the temperature and V). Then, a similar comparison as for the 2 dimensional case is possible and is presented in figures 10.10 and 10.11. Observe that the advantage of optical communication is greater when p is large and/or the dimension e is low.

Now we give a complete analysis of the effects of heat removal. We initially assume $N^{1/3}d/v \leq T$ so that

$$Qd^2 \geq N^{\frac{1}{3}} 2\epsilon V^2 k \kappa N^{p-\frac{2}{3}} dB. \quad (10.38)$$

If the value of d calculated from the above expression does not satisfy the assumed condition, then, just as in the 2 dimensional case we may show

$$Qd^2 \geq N^{\frac{1}{3}} k [2\epsilon V^2 \kappa r_x^{3p-2} d + 2\epsilon V^2 v T r_x^{3(p-1)} z] B \quad (10.39)$$

where $z = (1 - r_x^3/r_{max}^3) \leq 1$. Using $r_x = vT/d$, we can solve for d as

$$d^{3p-1} \geq \frac{2\epsilon V^2 (vT)^{3p-2} k \kappa N^{\frac{1}{3}} B}{Q} \quad (10.40)$$




Figure 10.11: Comparison of optical and normally conducting interconnections for same values of B in 3 dimensions. Similar assumptions as in the previous figure are made. E/Q varies from $10^{-21} \text{ m}^2 \text{ sec}$ to $10^{-13} \text{ m}^2 \text{ sec}$ in increments of $10^{-2} \text{ m}^2 \text{ sec}$.

where we again used $\kappa + z \simeq \kappa$. This expression forms continuity with equation 10.38. Thus, the minimum value of d as set by heat removal requirements is given by

$$d = \min \left(\frac{2\epsilon V^2 k \kappa N^{p-\frac{1}{3}} B}{Q}, \left(\frac{2\epsilon V^2 (vT)^{3p-2} k \kappa N^{\frac{1}{3}} B}{Q} \right)^{\frac{1}{3p-1}} \right) \quad (10.41)$$

and the resulting power dissipation may be expressed as

$$\mathcal{P} = 2\epsilon V^2 \min(\mathcal{L}, (vT)^{3p-2} \mathcal{L}^{3(1-p)}) k \kappa N^p B. \quad (10.42)$$

Now, the inverse signal delay is given by $1/S = \max(N^{1/3}d/v, T)$. Since d is an increasing function of B , we have a tradeoff between S and B . If we set B to its largest possible value of $B = 1/T$, this will result in a particular value of S . This is illustrated in figure 10.12. These curves should be compared to those for optical interconnections in figure 10.11 to determine which medium is superior. Notice that the curve for normal conductors in that figure which ignored heat removal is quite optimistic.

However, we need not set B to its largest possible value $= 1/T$. By choosing B to be smaller (i.e. by operating at a smaller duty ratio), we can reduce power dissipation, pack the elements more densely (i.e. reduce d) and thus decrease propagation delays,




Figure 10.12: *The effect of heat removal for normal conductors in 3 dimensions.*

resulting in an increase in S . There is no purpose in reducing B beyond a certain extent however, as once the scale of the system is reduced to the extent that all lines become unterminated, further reduction in scale will not improve S . Note that there is an upper limit to both of the quantities S and B regardless of the other, however they can be traded off for each other over a certain range. Given any optimization function involving B and S , we can find the optimum operating point.

As an example, let us assume that we would like to maximize $S = B$. This optimization function might be appropriate for a synchronous system where the clock rate is set by the signal delay along the longest connection.

We will use the duty ratio $x \leq 1$ as an optimization parameter. Thus the bit repetition rate may be expressed as $B = x/T$. We will assume T_d to be small enough to have no effect. In part a. of figure 10.13 we plot the optimum duty ratio for selected parameters. We observe that rather small duty ratios are optimal for a wide range of N . However, as N becomes very large, the optimal duty ratio tends to unity, because of the increase in T . As mentioned before, the maximum value of B involuntarily drops with normally conducting systems, so that eventually less and less power is dissipated and heat removal ceases to be an issue. Part b. of the figure shows the optimal value of $S = B$, as well as the maximum possible value of $B = 1/T$, and the resulting value of S when this value of B is employed (by choosing $x = 1$). As x is decreased, smaller values of B are employed, resulting in a decrease in system




Figure 10.13: *Optimum duty ratio and resulting performance.* We assume $V^2/Q = 10^{-6} \text{ V}^2 \text{ m}^2/\text{W}$ and that T_d is small enough to be ignored. Part a. shows the optimal duty ratio x . Part b. shows the resulting optimal value of $S = B$, the maximum possible value of B (when $x = 1$) and the value of S for this value of B .

size and increase in S . Of course, no matter how much B is decreased, S cannot be increased beyond $1/T$.

This example serves as a good illustration of how our analysis can predict an interesting effect and provide us with a quantitative understanding of it.

The existence of a tradeoff between switching speed and wire lengths for 3 dimensional systems based on heat removal has been independently noted by Nakayama [127] in a very recent paper.

10.2.3 Asymptotic properties

In this subsection we assume arbitrarily fast devices ($T_d \rightarrow 0$), 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 equation 10.18.

Equation 10.18 immediately leads to an important conclusion: for given N , there is an upper limit to B . This is in contrast with the optical case where B could be arbitrarily increased by suffering a decrease 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 $T_\ell \propto \ell^2/W^2$.

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

$$HS = HB \propto N^{\frac{e(e-3)}{e-1}} \quad (10.43)$$

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.

10.3 Repeatered interconnections

10.3.1 Relations between S , B and N

First assume that element size and heat removal need not be considered. As in the optical case, $T \simeq R_0C_0$ is a constant independent of other line parameters. Thus if $BT \leq 1$ we set $\chi = 1$. Then, for $e < n$, using $d^{e-1} = k\kappa N^{(n-e)/ne} W^{e-1}$ and

$$T_p = 4\sqrt{R_0C_0\rho\epsilon} \frac{N^{\frac{1}{e}}d}{W} \quad (10.44)$$

we obtain

$$SN^{\frac{n-1}{n(e-1)}} = \left(4\sqrt{R_0C_0\rho\epsilon}\right)^{-1} (k\kappa)^{\frac{-1}{e-1}} \quad (10.45)$$

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 10.5 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 10.7).

Upon comparison with the coefficient of equation 10.5, we see that repeaters are worse than optical communication by less than a factor of 10, assuming $R_0C_0 = 100$ 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 equation 10.45 is scale independent, assuming W is small enough so that inductive effects need not be accounted for ($\leq 5 \mu\text{m}$ for $R_0C_0 = 100 \text{ psec}$). When element size is accounted for, S is given by the minimum predicted by equation 10.45 and $1/S = \sqrt{\mu\epsilon} N^{1/e} d_d$.

10.3.2 Heat removal

2 dimensions

Since equation 10.45 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 equation 5.18 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 $k\chi\bar{r}W/M = d$. Since $\ell/W = rd/W = rk\chi\bar{r}/M$, we find that we must maintain

$$Qd^2 \geq k\bar{r} \min \left(2\epsilon V^2 d, 8\epsilon V^2 \sqrt{\frac{\rho R_0 C_0}{\mu}} \frac{k\chi\bar{r}}{M} \right) B \quad (10.46)$$

so that d must be at least

$$d = \min \left(\frac{2\epsilon V^2 k\bar{r} B}{Q}, \left(\frac{8\epsilon V^2 \sqrt{\rho R_0 C_0 / \mu}}{Q} \right)^{\frac{1}{2}} \frac{\chi^{\frac{1}{2}} k\bar{r} B^{\frac{1}{2}}}{M^{\frac{1}{2}}} \right). \quad (10.47)$$

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, equation 10.45 is applicable. If the second term is less than the first, we find that d is greater than the mentioned critical value. Then

$$\frac{1}{S} = T_p = \sqrt{\mu\epsilon} N^{\frac{1}{2}} d \quad (10.48)$$

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 equations 10.45 and 10.48. The resulting dependence of S on N is presented in figure 10.14.




Figure 10.14: S versus N for $B = 1$ Gbit/sec for repeaters in 2 dimensions. We take $k = 5$, $M = 10$ and $T = R_0 C_0 = 1$ 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.

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

$$d = k\bar{r}B \min \left(\frac{2\epsilon V^2}{Q}, \left(\frac{8\epsilon V^2 \sqrt{\rho R_0 C_0 / \mu R_0 C_0}}{QM} \right)^{\frac{1}{2}} \right) \quad (10.49)$$

which we may simply write as $d = k\bar{r}B(\text{constant})$. For $Q/V^2 = 10$ W/cm²V² and our usual parameters, (constant) is given by $\simeq \min(0.7, 3.5(R_0 C_0)^{3/4}/M^{1/2})(\text{m fsec})$ where $R_0 C_0$ is in nsec. Thus, if $R_0 C_0$ is not small or M is not large we will most likely be operating in the lumped regime so that equation 10.45 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 wireability requirements with increasing N (for $p > 1/2$).




Figure 10.15: S versus N for $B = 1$ Gbit/sec for repeaters in 3 dimensions. We take $k = 5$, $p = 0.8$ and $T = R_0 C_0 = 1$ nsec so that $\chi = 1$. d_d is assumed to be small enough to have no effect.

3 dimensions

We repeat the analysis presented for 2 dimensions. Now the minimum cell size is found to be

$$d = \min \left(\frac{2\epsilon V^2 k \bar{r} N^{\frac{1}{3}} B}{Q}, \left(\frac{8\epsilon V^2 \sqrt{\rho R_0 C_0 / \mu}}{Q} \right)^{\frac{1}{2}} \chi^{\frac{1}{4}} (k \bar{r})^{\frac{3}{4}} N^{\frac{1}{6}} B^{\frac{1}{2}} \right). \quad (10.50)$$

Thus the resulting delay is the greater of $T_p = \sqrt{\mu \epsilon} N^{1/3} d$ and that given by equation 10.45 and is illustrated in figure 10.15.

10.3.3 Asymptotic properties

The asymptotic behavior of repeatered systems is similar to that of optical systems, 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 equation 10.50. 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 $e < n$)

$$HS \propto N^{\frac{p}{4}} \quad (10.51)$$

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 equation 10.50 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 *cannot* be increased with increasing N , an inhibiting situation.

10.4 Superconducting interconnections

10.4.1 Relations between S , B and N

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

We refer to figure 5.4. 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 $W \leq 4\lambda$. To see this, simply notice that as we scale the system photographically, ℓ varies in linear proportion to W . The key quantity to be calculated is ℓ_{max}/W . If $B \leq 1/T_r = 1/T = 1/T_d$ so that $\chi = 1$, then for $e < n$ we obtain $\ell_{max}/W = r_{max}(k\kappa)^{1/(e-1)} N^{(n-e)/ne(e-1)}$. Since in this region $T_p = (4\lambda/v)\ell/W$, and assuming T_d is small so that $\tau = T_p$, we find, with $r_{max} \simeq N^{1/e}$

$$SN^{\frac{n-1}{n(e-1)}} = \left(\frac{v}{4\lambda}\right) (k\kappa)^{\frac{-1}{e-1}}. \quad (10.52)$$

This relation is independent of the specific choice of W , 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 W less than $0.1 \mu\text{m}$ 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 equation 10.5 derived for optical interconnections. The relation for $BT_d \geq 1$ is likewise similar to equation 10.7. If $\epsilon_r \simeq 4$ for the superconducting interconnections and optical wavelengths ($\lambda \sim 0.5 \mu\text{m}$) 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 W , so that there is no performance advantage. The potential advantage in terms of cost of area is limited by how small we can manufacture W and how much the voltage can be reduced and/or critical current increased.

10.4.2 Heat removal

2 dimensions

Heat removal has no effect on performance unless it requires that d be large enough that $W \geq 4\lambda$. Wireability dictates $d \geq k\chi\bar{r}W/M$ 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 W to be larger than 4λ . The analysis is very similar to the optical case; in fact, figure 10.2 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 W also reduces ℓ_{max} , keeping ℓ_{max}/W and the system delay constant. This not only reduces system size, as mentioned earlier, but results in less total energy consumption (since $Z_0 \propto 1/W$ in the intermediate region), making it desirable to choose W 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 $W \geq 4\lambda$ and find that heat removal requires that the scale be large enough that $W \geq 4\lambda$, we know that we are not excluding any room for improvement in the $W \leq 4\lambda$ region. When heat removal allows $W \leq 4\lambda$, S does not depend on heat removal parameters anyway and is given by equation 10.52.

3 dimensions

This case is likewise similar to the corresponding optical case, with the above remarks in mind. We only need interpret figure 10.4 with the appropriate superconducting energy E .

10.4.3 Asymptotic properties

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

10.5 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. We must not forget that a multitude of engineering issues [163] must be faced in the actual construction of such systems.

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 conducting systems. In fact, we are unable to realize 3 dimensional optical circuits that can provide an arbitrary pattern of connections and which can approach the diffraction limited $W \simeq \lambda$ we have assumed, although the architecture presented in chapter 7 comes close. 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. The construction of efficient, reliable and small size transducers is a 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 [99].

Another issue we did not directly account for is that of fan-out. Architectures involving large fan-outs tend to favor optical communication [40].

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 [24].

10.6 Summary and conclusions

Combining our system (chapters 2, 3 and 4) and physical (chapter 5) models, we derived relations of the form $\Phi(S, B, N) \leq C_\Phi$ for each interconnection medium. These relations bound the largest simultaneously possible values of S (inverse signal delay), B (bit repetition rate) and N (number of elements). An abstract formalism enabling us to relate these bounds to the computational requirements of given applications was described in chapter 4. The analysis of this chapter has provided us with a quantitative and qualitative understanding of the limitations imposed by different interconnection media, and how they compare to each other and to device limitations.

We saw that none of the interconnection media we have considered enabled continual reduction of signal delay by downscaling. (Optical interconnections cannot be

downscaled, and all kinds of conducting interconnections exhibit an inverse dependence of delay on linewidth, below a certain linewidth.)

In discussing the limitations of conducting interconnections, we allowed arbitrarily small scaling and arbitrarily fast devices. We saw that normal conductors, whether terminated or not, did not allow B to be kept constant with increasing system size (which is possible with the other media). Both B and S were found to sharply decrease with increasing N . Making longer lines wider leads to improvement by only a constant 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 media 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 very fast devices are possible, the performance for 2 dimensional layouts may approach that possible with the other media we have considered within an order of magnitude. For large system sizes, they will still be more costly in terms of power consumption. In 3 dimensions, repeaters are inferior to optics and superconductors since they result in faster growth of signal delay and slower growth of the bisection-inverse delay product with increasing N .

Optical and superconducting interconnections lead to very similar performance for same dimensional layouts and similar communication energies. Although superconducting layouts may be much smaller than optical layouts, they do not result in smaller delay because of the inverse dependence of delay on linewidth, once conductor thickness drops below the penetration depth. Optical interconnections may 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 wireability becomes more important than heat removal as N increases. In 3 dimensions however, this will usually enable lower signal delay.

In this chapter, we compared the ability of given interconnection media in providing communication among a given array of elements. In the following chapters, we discuss how these media may be used in conjunction to achieve performance not possible with any alone.

Chapter 11

Optimal Hybrid Implementations

In this chapter we will investigate how different interconnection media can be used in conjunction to realize a system not possible with any alone. More specifically, we will determine the optimal mix of optical and normally conducting interconnections maximizing a given figure of merit function.

11.1 Introduction

In chapter 10, we considered the use of only one interconnection medium at a time to implement all connections. We found that normally conducting interconnections were preferable for smaller numbers of elements whereas optical and superconducting interconnections were preferable for larger numbers of elements. This suggests that we can do better by joint use of normal conductors (for the shorter connections) and optics or superconductors (for the longer connections). The question is beyond what point should we start employing optics?

11.2 The breakeven distance approach

The way this problem was originally addressed in the literature was by deriving a breakeven distance beyond which the use of optical communication was preferable to the use of normal conductors [92] [40] [121] [65] .

For instance, Feldman [40] and Miller [121] claimed that optical communication is energetically favorable for connections of length $\ell > 1$ mm or so. We find Feldman's analysis somewhat pessimistic about normal conductors and Miller's somewhat optimistic about optics. In chapter 5 we estimated the optical communication energy to be about 1 pJ for typical parameter values (including modulator energy and 2.5% overall differential efficiency for $(5\text{ }\mu\text{m})^2$ devices operated at $V = 1$ V). The energy per transmitted bit along an unterminated normally conducting line was estimated to be about 100 fJ/mm (operating voltage $V = 1$ V), so that the breakeven length for energy for this particular choice of parameters is about ~ 1 cm. For superconducting lines, we may similarly estimate the energy per transmitted bit to be about 1 pJ (assuming 0.1 nsec pulses and an operating voltage of $V = 1$ V), similar to optics. However, note that this figure can be made much smaller if the operating voltage can be reduced.

There have also been attempts to compare the information density of optical and normally conducting interconnections [65]. However, these comparisons were based on rather arbitrary assumptions, especially regarding normal conductors. We can derive a breakeven length for information density using our tube models. First, assume that we would like to maximize the throughput of information per unit area. For optical interconnections, the maximum information density is $1/T_r(f\lambda)^2$ (in bit/m²sec) (since each channel operates at a rate $B \leq 1/T_r$ and occupies an area $(f\lambda)^2$). For normally conducting interconnections, the maximum information density is $1/(16\rho\epsilon)\ell^2$ (since $B \leq 1/T_\ell = W^2/(16\rho\epsilon)\ell^2$), which decreases with increasing length. With room temperature conductivity, $f = 2$ and $T_r = T_d = 1$ nsec, we find that the breakeven length for information density is about 2 cm, the same order as the breakeven length for energy. For superconducting interconnections, the maximum information density is $J_{sc}^2\mu/16\epsilon V^2 T_d$ (since the minimum value of $W = 4V/(J_{sc}\sqrt{\mu/\epsilon})$), which is about the same as for optical communication, assuming $J_{sc} = 25$ mA/ μ m and $V = 1$ V.

An alternative breakeven length which emphasizes signal delay, rather than bit repetition rate can also be derived. Remember from chapter 10 that the signal delay of a wireability limited layout was proportional to the linewidth divided by the

propagation velocity. For optics, this quantity is $(f\lambda/c)$ (equation 10.5). For normal conductors, maintaining length proportionate delay at a propagation velocity v requires $W^2 \geq 16\rho\epsilon v\ell$ (table 5.1). Thus, comparing the minimum value of W/v to $(f\lambda/c)$ with our usual choice of parameters we find that optical communication is preferable for lengths greater than about 0.5 mm. For superconductors, the linewidth-propagation velocity ratio is $4\lambda/v$, which is of the same order as optics for several thousand Armstrong penetrations depths.

Whereas this approach to comparing various interconnect media can be instructive, it is unsatisfactory in many ways. First of all, it enables comparison of only one quantity at a time, without paying attention to the others. Whether information density or energy is of greater importance will depend not only on whether the system is heat removal, wireability or device limited; but also on the relative emphasis we give to various measures of performance (signal delay, bit repetition rate etc.) and cost (system size, power dissipation). Since the length scale of the system is related to the properties of the interconnections through wireability and heat removal requirements in a complicated manner, we do not know initially the physical length $\ell = rd$ of a line of length r in grid units. The comparison of isolated lines of given length has little meaning when these lines are embedded in a system. Even the comparison of an all optically connected system with an all electrically connected system (discussed in chapter 10) does not tell us which connections to implement optically in a hybrid system.

11.3 Outline of the analysis

For the reasons discussed above, we will take a more general approach to this problem. We will start with a layout of N elements (gates or switches) which we will partition into N/N_1 groups of N_1 elements each (figure 11.1). All connections internal to a group will be made electrically whereas connections between elements in different groups will be made optically. For given N and B , we will calculate the optimal value of N_1 maximizing our optimization function Γ , which in general can be a function of inverse signal delay, total system size and power consumption. We will mostly be




Figure 11.1: Partitioning a system of N elements into N/N_1 groups of N_1 elements.

interested in *high performance* systems for which system size and power dissipation are only secondary considerations. We will employ the optimization function

$$\Gamma = \frac{S}{\mathcal{P}^\epsilon} \quad \epsilon \rightarrow 0 \quad (11.1)$$


which gives full precedence to maximizing inverse signal delay S for given N and B , and secondarily tries to minimize total power dissipation $\mathcal{P}$. We will also consider optimization functions putting a greater emphasis on cost of size and power.

Now we actually outline the steps of our analysis (figure 11.2). The linear extent $\mathcal{L}_1$ (also denoted as d_1) of an electrically connected group of N_1 elements must be large enough to:

1. Accommodate N_1 elements.
2. Accommodate the (electrical) wires connecting them.
3. Accommodate kN_1^p optical transducers.
4. Satisfy heat removal requirements.

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

1. Accommodate d_1 .

Figure 11.2: *Analysis of optimal hybrid layouts.*

2. Accommodate the optical channels connecting the groups.
3. Satisfy any additional heat removal requirements.

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

We note that we have relaxed our assumption regarding a uniform array of elements. Clustering the N_1 elements forming each electrically connected group together enables considerable energy savings with respect to a fully uniform layout. However, the N_1 elements in each group and the N/N_1 groups are still uniformly laid out.

We have carried out this analysis for a variety of layout constraints, combinations of media and physical parameters. It is not possible (and perhaps not useful) for us to reproduce all of our results. Rather, we will try to present representative examples, chosen for their illustrative and instructive qualities and qualitatively discuss several general conclusions deduced from the study of a large number of cases.

11.4 Notation and nomenclature

The subscripts o , n , r and s will be used to distinguish between the parameters of optical, normally conducting, repeatered and superconducting interconnections respectively. When no subscript is used, this will mean that that particular parameter is taken to be the same for all interconnection media.

The notation (N2d) will denote a fully 2 dimensional layout employing only normal conductors. The notation (N2d,O2d) will denote a fully 2 dimensional layout where normal conductors are used to establish the shorter connections and optics is used to establish the longer connections. (N3d,O3d) denotes the fully 3 dimensional analog of this layout. The notation (OP1) denotes an optical layout where the elements are constrained to lie on a plane, but optical communication paths are allowed to leave the plane. Thus (N2d,OP1) denotes a 2 dimensional array of modules (with normally conducting internal connections), interconnected with out of plane optical communication, such as a 2 dimensional array of VLSI chips communicating via a hologram situated above them. (N2d,S2d,O3d) would denote a 3 dimensional array of optically communicating ‘boards’. Each board is a 2 dimensional array of 2 dimensional modules (with normally conducting internal connections), interconnected with superconducting interconnections.

We will often speak of ‘electrical interconnections’ without specifying whether we mean normally conducting, repeatered or superconducting interconnections. Which we are referring to will be evident from the context.

11.5 Regarding numerical examples

As discussed in chapter 1, in our numerical examples we will try to look into the future and select reasonably optimistic parameters for each interconnection media. We will also consider the effects of degrading the optical parameters from what seems their best possible values.

We will assume 10 GHz devices, i.e. $T_d = 0.1$ nsec. We will allow a maximum of

10 conducting signal layers¹ but only $M_o = 1$ signal layer for 2 dimensional optical layouts. We will use a nominal voltage level of $V = 1$ V and room temperature aluminum conductivity. We will allow a maximum power dissipation per unit area of $Q = 10$ W/cm². The minimum manufacturable value of W for conducting interconnections will be taken as $W_{min} = 0.2$ μ m. The element (gate) size d_d will be assumed to be 10 times this value. We will assume there to be $k = 5$ connections per element. The optical transducers will be assumed to be $d_{tr} = 5$ μ m in diameter. We will assume the best possible optical communication energy to be $E_o = 1$ pJ but also consider degradation of this value by a factor of one hundred. Likewise, we will consider near diffraction limited operation ($f = 2$) but also consider degradation of this by up to one hundred ($f = 200$). We will consider two different values of the Rent exponent, $p = 0.6$ and 0.8 to observe the effects of connectivity on the results.

In chapter 10, we often ignored the effects of parameters such as the minimum manufacturable linewidth, transducer size etc. since the end results could be easily modified to account for their effects. Since the analysis of this section is somewhat more complicated, we will include these parameters in our expressions. However, the values chosen above for W_{min} , d_d and d_{tr} , which seem realizable in the near future, are already small enough that in most cases, totally ignoring the effects of these parameters would have little or no effect on our results.

We have set the parameter ε in equation 11.1 to 10^{-10} so as to cause the least arbitrariness. (We observed that choosing values as large as 10^{-5} or as small as 10^{-12} makes no difference. Values larger than 10^{-5} start changing the results in favor of systems exhibiting somewhat larger signal delay but less power dissipation. Values lower than 10^{-12} start causing numerical problems on PRO-MATLAB [125].)

A note on the numerical optimization procedure: the optimum value of N_1 is determined by discretizing the optimization function Γ and picking out the value of N_1 that results in the largest Γ . This procedure is repeated for discrete values of N . There are about 2 or 3 samples in every order of magnitude of N and N_1 .

¹For smaller values of N_1 , the maximum useful number of layers may be less than this value. In what follows we assume that the value of M_n etc. is set to this maximum useful value.

11.6 Some analytical considerations regarding hierarchical analysis

In this chapter, we employ an analysis similar to that of chapter 10 in hierarchical succession. That is, we first apply this analysis to a group of N_1 elements. Then we take these groups as our ‘superelements’ and apply our analysis to a collection of N/N_1 such superelements. Our use of Rent’s rule based wireability analysis in such a hierarchical manner is similar to that of Bakoglu [6], who used such an analysis to predict the optimum parameters and performance of multi-chip modules etc.

In this section we derive some results regarding the hierarchical usage of Rent’s rule based wireability analysis. For simplicity we assume $e = 2$, $p > 1/2$, $M = 1$, $\chi = 1$ and the use of optical interconnections for both levels of the hierarchy: (O2d,O2d).

We previously showed that the linear extent of an optically connected system as imposed by wireability considerations is given by

$$\mathcal{L} = k\kappa N^p W. \quad (11.2)$$

Let us derive the same result based on hierarchical considerations. The linear extent of a group of N_1 elements is $\mathcal{L}_1 = k\kappa N_1^p W$. Obviously, the linear extent of the whole system must at least be $(N/N_1)^{1/2}$ times this. Furthermore, the whole system can be viewed as a collection of N/N_1 superelements with kN_1^p connections per superelement. This implies a system linear extent of $\mathcal{L} = \max(kN_1^p \kappa (N/N_1)^p W, (N/N_1)^{1/2} \mathcal{L}_1)$ which upon substitution of the expression for $\mathcal{L}_1$ is found to be the same as equation 11.2. Thus, the hierarchical calculation of the system linear extent is consistent with its direct calculation.

The total area $\mathcal{L}^2 = \ell_{total} W$ is proportional to the total interconnection length ℓ_{total} . Thus, the above result may be interpreted as a consequence of the invariance of the total interconnection length under hierarchical analysis. The total interconnection length is given by $\ell_{total} = kN\kappa N^{p-1/2} d$ where $d = k\kappa N^{p-1/2} W$ so that

$$\ell_{total} = (k\kappa)^2 N^{2p} W. \quad (11.3)$$

Let us now calculate the same quantity based on hierarchical considerations. Again

consider N/N_1 superelements with kN_1^p connections per superelement. First let us calculate the contribution of the longer connections (between modules) to the total connection length. It is given by $kN_1^p(N/N_1)\kappa(N/N_1)^{p-1/2}d_m$ (where $d_m = \mathcal{L}/(N/N_1)^{1/2}$) which can be shown to be equal to ℓ_{total} as given by equation 11.3. Thus, the contribution of the shorter connections (those internal to the modules), given by $(N/N_1)(k\kappa)^2N_1^{2p}W$, must be negligible in comparison, as is also easily shown. Thus, in addition to showing that the hierarchical calculation of ℓ_{total} is consistent with its direct calculation, we see that the area occupied by the shorter interconnections can be neglected in comparison to that occupied by the longer ones (provided N_1 and N/N_1 are both sufficiently large).

In conclusion, the Rent's rule based wireability analysis we employ is self-consistent under hierarchical usage.

11.7 Layouts confined to a plane: normal conductors and optics

11.7.1 Description

In this section we discuss fully 2 dimensional layouts where normally conducting interconnections are used to provide connections internal to each group and optical interconnections are used to provide connections between groups: (N2d,O2d).

We refer to figure 11.2. As before, there are k connections per element. Some of these connections are made to elements in the same module and are implemented using normal conductors. Other connections are made to elements in other modules. Such connections are established optically, by tying optical transducers at the to-be-connected terminals and guiding the light emanated from the source terminal to the target terminal with some kind of optical imaging system. It matters little whether the transducers are on a separate layer or side by side with the electrical wiring (since $\max(x, y) \simeq x + y$).

The intergroup spacing d_m may have to be larger than $d_1 = \mathcal{L}_1$, because of the space necessary to accommodate the optical channels. Again, it makes little difference

whether we assume that the optical channels are on a separate layer or compete for the same space with the modules.

One implementation may involve a 2 dimensional array of VLSI chips with optical transducers located on a topmost layer or dedicated islands. The imaging system may be a glass waveguide overlay or the folded multi-facet architecture of chapter 7.

11.7.2 Analysis

Remember that our purpose is to determine, for every N and B , the value of N_1 maximizing our figure of merit function. We can immediately set an upper bound on N_1 since we know that the maximum value of B is a decreasing function of the number of electrically connected elements. This was derived previously as (equations 10.18 and 10.20 with $e = 2$)

$$\frac{1}{B} = \max \left((16\rho\epsilon) \left(\frac{k\kappa}{M_n} \right)^2 N_1^{2p}, (16\rho\epsilon T_d)^{\frac{1}{2}} \left(\frac{k\kappa}{M_n} \right) N_1^p \right) \quad (11.4)$$

assuming that the optimum value of $\chi = \max(1, BT_d)$ is chosen. We can solve each of the component equations above for N_1 and take the minimum as the largest value of N_1 compatible with given B . We will denote this value of N_1 as N_{1max} . For given B , N_1 must be chosen to lie between 1 and N_{1max} .

The linear extent of each module d_1 must satisfy

$$d_1 \geq \left(N_1^{\frac{1}{2}} d_d, N_1^{\frac{1}{2}} K'_n W_{min}, d'_1, \left(\frac{k N_1^p E_o B}{Q} \right)^{\frac{1}{2}}, (k\chi N_1^p)^{\frac{1}{2}} d_{tr}, k\chi N_1^p (f\lambda) \right). \quad (11.5)$$

The first term is trivial. d_1 must at least be large enough to accommodate $N_1^{1/2} \times N_1^{1/2}$ elements of linear extent d_d each. The second term is simply $N_1^{1/2}$ times the necessary physical wiring tracks per cell per layer times the minimum manufacturable linewidth. The number of normally conducting wiring tracks required per cell per layer is given by $K'_n = K_n \chi / M_n = k\chi \kappa N_1^{p-1/2} / M_n$. d'_1 denotes the minimum linear extent of the module as set by heat removal considerations due to dissipation on the normally conducting wires. It is given by (equation 10.36)

$$d'_1 = N_1^{\frac{1}{2}} \min \left(\frac{2\epsilon V^2 k\kappa N_1^{p-\frac{1}{2}} B}{Q}, \left(\frac{2\epsilon V^2 (vT_n)^{2p-1} k\kappa B}{Q} \right)^{\frac{1}{2p}} \right) \quad (11.6)$$

where the minimum temporal pulse width T_n is given by

$$T_n = \max \left((16\rho\epsilon) \left(\frac{k\kappa}{M_n} \right)^2 N_1^{2p}, T_d \right). \quad (11.7)$$

The term after d'_1 accounts for optical power dissipation on each module. We are assuming that the power associated with optical interconnections is dissipated at the transducers. According to Rent's rule, there will be kN_1^p connections per module, resulting in a power dissipation of $kN_1^p E_o B$. Of course, d_1 must actually satisfy $Qd_1^2 \geq (\text{electrical dissipation} + \text{optical dissipation})$ where the electrical dissipation is a function of d_1 . However, the error associated with decoupling the value of d_1 as imposed by electrical and optical power dissipation is at most a factor of $\sqrt{2}$, since the resulting linear extent is proportional to the square root of the total power dissipated. The term involving d_{tr} accounts for the fact that the module must at least be large enough to accommodate $kN_1^p \chi$ transducers, since $\chi = \max(BT_d, 1)$ optical channels are used per graph edge. The final term accounts for the fact that the linear extent of the module must be large enough to allow for the passage of $kN_1^p \chi$ optical channels through the module, otherwise the channels emanating from the transducers would not be able to get out of the boundaries of the module. (Such a problem will not arise when we allow out of plane optical communication.)

The intermodule separation d_m must satisfy

$$d_m \geq \max(d_1, K'_o(f\lambda)). \quad (11.8)$$

Apart from being large enough to accommodate the module, the intermodule spacing must be large enough to accommodate a sufficient number of optical channels. The number of tracks per module cell per layer is given by $K'_o = K_o \chi / M_o = kN_1^p \chi \kappa (N/N_1)^{p-1/2} / M_o$. Since no power is dissipated off the modules, we do not require any additional heat removal terms.

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

$$\tau = \max \left(\left(\frac{N}{N_1} \right)^{\frac{1}{2}} \frac{d_m}{c}, (16\rho\epsilon) \left(\frac{k\kappa}{M_n} \right)^2 N_1^{2p}, \frac{d_1}{v}, T_d \right). \quad (11.9)$$

The first term is the speed of light delay along the longest optical interconnection. The second and third terms correspond to the delay along the longest normally conducting interconnection and the last term accounts for device delay. The total power consumption $\mathcal{P}$ is also easily expressed as

$$\mathcal{P} = \left(\frac{N}{N_1}\right) (kN_1^p E_o + \text{electrical energy per module})B \quad (11.10)$$

where the electrical energy per module is given by (equation 10.37)

$$\min \left(2\epsilon V^2 k\kappa N_1^p d_1, 2\epsilon V^2 (vT_n)^{2p-1} k\kappa N_1^p d_1^{2(1-p)} \right). \quad (11.11)$$

Several of the terms in the above analysis are actually redundant and have been included for completeness. We will now present a stripped down version for further transparency. We assume $M_o = 1$ and $\chi = 1$, as will be the case in all of our numerical examples. Often d_d , d_{tr} and W_{min} will be small enough not to be an issue and d_1 will be small enough that all lines are unterminated. Then

$$d_1 \geq \max \left(kN_1^p (f\lambda), \left(\frac{kN_1^p E_o B}{Q} \right)^{\frac{1}{2}}, \frac{2\epsilon V^2 k\kappa N_1^p B}{Q} \right). \quad (11.12)$$

Assuming fast devices are used the delay is then given by

$$\tau = \max \left(k\kappa N^p \left(\frac{f\lambda}{c} \right), \left(\frac{N}{N_1} \right)^{\frac{1}{2}} \left(\frac{kN_1^p E_o B}{c^2 Q} \right)^{\frac{1}{2}}, \left(\frac{N}{N_1} \right)^{\frac{1}{2}} \frac{2\epsilon V^2 k\kappa N_1^p B}{cQ}, \right. \\ \left. (16\rho\epsilon) \left(\frac{k\kappa}{M_n} \right)^2 N_1^{2p} \right). \quad (11.13)$$

The first term, which is independent of N_1 , is also the delay of an all optically connected system of N elements (equation 10.5 with $\chi = 1$). Thus, we conclude that use of a hybrid layout cannot reduce the system size and delay with respect to an all optically connected system, when the system is wireability limited. This is, of course, a consequence of the discussion of section 11.6. Since the contribution of the shortest interconnections to total area is negligible anyway, it does not help to make them out of submicron normal conductors.

Since $1 \leq N_1 \leq N_{1max}$, the first term in the above equation will eventually dominate the others with increasing N . When this is the case, the choice of N_1 will

have no effect on the delay. Thus, we will choose that value which results in minimum power dissipation:

$$\mathcal{P} = NB \max(kE_o N_1^{p-1}, 2\epsilon V^2 k\kappa N_1^{p-1} d_1). \quad (11.14)$$

When d_1 is given by the first term of equation 11.12, the value of N_1 minimizing total power dissipation is given by

$$N_1^p = \frac{E_o}{2\epsilon V^2 (f\lambda) k\kappa}. \quad (11.15)$$

For larger values of B , either of the latter two terms in equation 11.12 may dominate. Interestingly, whichever of these terms dominates, it is possible to show that the value of N_1 minimizing equation 11.14 is given by the same expression:

$$N_1^p = \frac{kE_o Q}{(2\epsilon V^2 k\kappa)^2 B}. \quad (11.16)$$

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

Of course, since we are giving full precedence to minimizing signal delay, keeping the latter terms of equation 11.13 below its first term will have priority over minimizing power. For instance, so that the last term of equation 11.13 does not exceed the first, we must maintain

$$N_1 \leq \left(\frac{f\lambda/c}{16\rho\epsilon} \frac{M_n^2}{k\kappa} \right)^{\frac{1}{2p}} N^{\frac{1}{2}}. \quad (11.17)$$

Thus, the optimal value of N_1 is determined by three considerations. For large N , it will be that value which minimizes total power consumption, as given by either of




Figure 11.3: N_1 versus N for $(N2d, O2d)$ I. $p = 0.6$, $f = 2$, $E_o = 1$ pJ.

equations 11.15 and 11.16, with the restriction that it can never exceed N_{1max} . These two considerations will determine the optimal value of N_1 for large N , which we will denote as $N_{1\infty}$. The optimal value of N_1 should also not exceed the ‘envelope’ given by equation 11.17. This last restriction will relax with increasing N . Notice that the effect of varying p on all three of these considerations is the same: larger Rent exponent’s favor the use of more optics.

11.7.3 Numerical examples

Throughout all of this chapter, we will keep all but three of the physical parameters the same for all examples, as discussed in section 11.5. Thus in the following we only specify the values of p , f and E_o for each example.

In our first example we consider a system with Rent exponent $p = 0.6$, optical communication density only $f = 2$ times worse than diffraction limited and an optical communication energy $E_o = 1$ pJ. Figure 11.3 shows the optimal value of N_1 as a function of N with B as a parameter. For the two lower values of B , it is optimal to make all connections electrical (i.e. $N_1 = N$) until about $N \simeq 2 \times 10^5$, after which the optimal value of N_1 is independent of N and is that value which minimizes total power consumption. For these relatively low values of B , the size of each electrically connected group d_1 is given by $d_1 = kN_1^p(f\lambda)$. All conducting interconnections turn




Figure 11.4: S versus N for $(N2d, O2d)$ I. $B = 100$ Mbit/sec, $p = 0.6$, $f = 2$, $E_o = 1$ pJ.

out to be untermiated in this example. Thus the total power dissipation is given by equation 11.14 with $d_1 = kN_1^p(f\lambda)$. The value of N_1 that minimizes power dissipation is given by equation 11.15 and is indeed consistent with that observed in figure 11.3 for the two lowest values of B . For the two larger values of B , the latter terms of equation 11.12 dominate so that the optimal value of N_1 is given by equation 11.16, which indeed predicts the optimum values of N_1 for the two larger values of B in figure 11.3.

Figure 11.4 illustrates the resulting dependence of S on N for $B = 100$ Mbit/sec. The solid curve corresponds to the optimal choice of N_1 . The broken curve corresponds to making all connections optical and coincides with the solid curve for larger values of N , and is given by $1/S = \tau = k\kappa N^p f\lambda/c$. The dotted curve, which initially overlaps with the solid curve corresponds to making all connections electrical, and is given by $(16\rho\epsilon)(k\kappa/M_n)^2 N^{2p}$. We cannot make all connections electrical once $N > N_{1max} = 5.7 \times 10^6$, so that the dotted curve terminates at this value of N . After a certain value of N , making all connections optical is as good as the optimal hybrid combination in terms of minimizing system size and delay, as discussed in the previous section (following equation 11.13).

However, making all connections optical results in power dissipation about an order of magnitude larger than the optimal hybrid combination for the largest value




Figure 11.5: $\mathcal{L}$ versus N for $(N2d, O2d)$. $B = 100$ Mbit/sec, $p = 0.6$, $f = 2$, $E_o = 1$ pJ.

of B and about two orders of magnitude larger for the smallest value of B . The reason the disparity is greater for smaller values of B is because the optimal value of N_1 is larger when B is smaller. In other words, the all optical system ($N_1 = 1$) is farther away from the optimum. These considerations have no effect on the resulting system size and signal delay, since 2 dimensional systems tend to be wireability, rather than heat removal limited. Another consequence of this is that the resulting value of S for different values of B is identical (since also $BT_d \leq 1$ so that $\chi = 1$).

The entrance of optical interconnections after $N \sim 2 \times 10^5$ elements is accompanied by a drastic increase in system size, as illustrated in figure 11.5. The linear extent of the all electrical system is given by $2\epsilon V^2 k \kappa N^p B / Q$. Once we start using optical communication for the longer connections they dominate the system area leading to a linear extent given by $k \kappa N^p f \lambda$. The curve for S is continuous because the fast velocity of propagation of optical interconnections compensates for the increase in system size. The jump in system size can be avoided by keeping all connections electrical; however in this case the value of S will be less than that possible with a hybrid system. This is one example of a situation where the use of optics allows performance not possible with normal conductors alone, but at a significant penalty in terms of system size.

Let us now explore the effects of degrading the optical parameters. The optimal value of N_1 when $f = 50$ is plotted in figure 11.6. We observe that it is beneficial




Figure 11.6: N_1 versus N for $(N2d, O2d)$ II. $p = 0.6$, $f = 50$, $E_o = 1$ pJ.

to stick to an all electrical system ($N_1 = N$) until $N > N_{1max}$, after which the entrance of optics is unavoidable². Once this occurs, the module size d_1 is limited by the term $kN_1^p(f\lambda)$ for all values of B . The resulting large value of d_1 makes electrical connections expensive, leading to a small optimal value of N_1 , as given by equation 11.15. Notice that although increasing f leads to an overall degradation in performance and cost, it results in greater use of optics when $N > N_{1max}$.

Figure 11.7 illustrates the resulting dependence of S on N . Notice that in this case, unlike the previous example, a sudden drop in S is observed. This is because we are forced to use optical interconnections prematurely so as to maintain the given value of B , before the value of S for an all electrical system falls below that for an all optical system (as was the case in figure 11.4).

In other words, when $f > 10$ or so, we have a region in S - B - N space where a small increase in N or B is accompanied by a large increase in system size and a large decrease in S . This behavior may have algorithmic implications, as suggested by the framework described in chapter 4. Among several algorithms designed to solve a given problem, it may be preferable to employ those requiring relatively smaller values of N and/or B , if possible, since even small increases in these parameters require a large sacrifice in terms of S . For instance, if we are trying to maximize a figure of merit

²If f is too large, the increase in $\mathcal{L}$ accompanying the entrance of optics may be unreasonably large, so that increasing N beyond N_{1max} while maintaining the given value of B may not be feasible.




Figure 11.7: S versus N for $(N2d, O2d)$ II. $B = 100$ Mbit/sec, $p = 0.6$, $f = 50$, $E_o = 1$ pJ.

function of the form $S^x B^y$, where $x, y > 0$ and x is not much smaller than y , it is likely that we will settle for an operating point not involving any optical interconnections.

In conclusion, if the use of optics is to be worthwhile for two dimensional systems, it is of paramount importance to bring f as close as possible to unity. The folded multi-facet architecture of chapter 7 was devised to meet this requirement.

Now we discuss the effects of increasing the optical communication energy hundredfold (figure 11.8). For large N , the optimal value of N_1 is seen to hit N_{1max} . For smaller N , we observe the envelope $N_1 \propto N^{1/2}$ given by equation 11.17. Despite the fact that increasing the optical energy results in a drastic shift in N_1 , it has no effect on the resulting value of S , which is still as given in figure 11.4. This is because the system is still wireability limited. (An exception occurs for very large values of B , for which the system may be heat removal limited up to a certain value of N .)

Of course, now the total power dissipation is much larger and the discrepancy between the optimal system and the all optical system in this respect is even greater than before. The entrance of optics results in a large increase in total power dissipation (figure 11.9). The all optical and all electrical system power dissipations are given by kNE_oB and $(2\epsilon V^2 k \kappa N^p B)^2 / Q$ respectively. The total power dissipation for the optimal hybrid case is given by $NkN_1^{p-1} E_o B$ with $N_1 \sim 5 \times 10^6$ for larger N .

Let us now consider that p is increased to 0.8. The downward shift in the optimal




Figure 11.8: N_1 versus N for $(N2d, O2d)$ III. $p = 0.6$, $f = 2$, $E_o = 100$ pJ.




Figure 11.9: $\mathcal{P}$ versus N for $(N2d, O2d)$. $B = 100$ Mbit/sec, $p = 0.6$, $f = 2$, $E_o = 100$ pJ.




Figure 11.10: N_1 versus N for $(N2d, O2d)$ IV. $p = 0.8$, $f = 2$, $E_o = 100$ pJ.

values of N_1 is easily explained by the changes in the values of N_{1max} . Once again the envelope given by equation 11.17 is observed.

11.8 Layouts confined to a plane: repeaters and optics

11.8.1 Description

In this section we consider the use of repeatered interconnections instead of plain normal conductors: (R2d, O2d).

11.8.2 Analysis

There is no maximum value of N_1 for given B when repeaters are used. The linear extent of each module must again satisfy equation 11.5 where d'_1 is now given by (equation 10.47)

$$d'_1 = N_1^{\frac{1}{2}} \min \left(\frac{2\epsilon V^2 k \kappa N_1^{p-\frac{1}{2}} B}{Q}, \left(\frac{8\epsilon V^2 \sqrt{\rho R_0 C_0 / \mu}}{Q} \right)^{\frac{1}{2}} \frac{\chi^{\frac{1}{2}} k \kappa N_1^{p-\frac{1}{2}} B^{\frac{1}{2}}}{M_r^{\frac{1}{2}}} \right) \quad (11.18)$$

with $\chi = \max(BT_d, 1)$ as usual. The intermodule separation d_m is calculated similarly. Finally the system signal delay is given by

$$\tau = \max \left(\left(\frac{N}{N_1} \right)^{\frac{1}{2}} \frac{d_m}{c}, 4\sqrt{\rho\epsilon R_0 C_0} \left(\frac{k\kappa}{M_r} \right) N_1^p, \frac{d_1}{v}, T_d \right). \quad (11.19)$$

The total power dissipation is again given by equation 11.10 where the electrical energy per module is now given by

$$\min \left(2\epsilon V^2 k\kappa N_1^p d_1, 8\epsilon V^2 \sqrt{\frac{\rho R_0 C_0}{\mu}} \frac{\chi(k\kappa)^2}{M_r} N_1^{2p} \right). \quad (11.20)$$

As long as the first term of the above equation dominates, equations 11.15 and 11.16 will still be applicable.

11.8.3 Numerical examples

In chapter 10 we discussed that the relations between S and N for repeatered and optical layouts were identical in form and differed only by the numerical factors $(f\lambda/c)/M_o$ versus $4\sqrt{\rho\epsilon R_0 C_0}/M_r$. Optical interconnections have larger linewidths and faster propagation velocity. Repeatered interconnections can (and must) be scaled down to smaller linewidths but have a slower propagation velocity. Let us first consider a system with optical communication density $f = 50$ times worse than diffraction limited and an optical communication energy of $E_o = 1$ pJ. The same results are obtained whether $p = 0.6$ or $p = 0.8$. (For these parameters an all repeatered layout results in signal delay over 40 times less than that possible with an all optical layout.) We find that an all electrical system ($N_1 = N$) is best for the range of N and B in consideration. What essentially happens is that with 10 GHz devices repeaters allow fast propagation. Much smaller linewidths and 10 wiring layers result in much smaller system size, more than making up for the deficiency compared to the speed of light.

It is also interesting to note that no optical communication is used even when the system size exceeds a centimeter, which is the breakeven length for energy for this choice of parameters. This illustrates the inadequacy of 'breakeven length' approaches. A submicron repeatered line cannot be replaced with an $f \sim 50$ optical line even though the latter may have a smaller communication energy. (Of course,




Figure 11.11: N_1 versus N for $(R2d, O2d)$. $p = 0.8$, $f = 2$, $E_o = 1$ pJ.

for large N , the total power dissipation for the all repeatered system will be greater than possible with a hybrid system. Remember however that our original figure of merit function gives full priority to minimizing signal delay.)

Now, let us assume the value of f is reduced down to 2, which we are assuming to be its best possible value. In this case, the discrepancy between the use of all repeaters and all optics is reduced to less than a factor of two. As long as the system is wireability limited, our optimization procedure still favors all repeaters over anything else, since it gives full priority to minimizing delay. An exception occurs for very large values of B for which the system tends to be heat removal limited. Then, minimizing delay is equivalent to minimizing power dissipation. The optimal value of N_1 is illustrated in figure 11.11. For the largest value of B , the all repeatered system is heat removal limited and results in signal delay greater than that given by the wire limited equation 10.45. In this case, hybridization ($N_1 < N$) is beneficial.

Apart from such cases, when $f \sim 2$, the signal delay of an all optical system is very close to that of an all repeatered system (i.e. equations 10.5 and 10.45 are numerically very close)³. The all repeatered system has much smaller area. When $p = 0.6$, the power dissipation of the all repeatered system lies quite below that of the all optical system for the range of N in consideration (although it grows faster

³This does not mean however that we can interchangeably use optics or repeaters for the individual connections in the same system, because of scale incompatibility.

with increasing N). Thus, unless B is very large so that the system is heat removal limited (as discussed above), an all repeatered system would be preferred. When $p = 0.8$, the power dissipation for the all repeatered case grows faster so that hybridization according to equations 11.15 and 11.16 may be beneficial. (Again, based on similar arguments as those leading to equation 11.17, we must ensure that $4\sqrt{\rho\epsilon R_0 C_0} (k\kappa/M_r) N_1^p \leq (f\lambda/c)(k\kappa/M_0) N^p$.)

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

11.9 Out of plane optical communication: normal conductors and optics

11.9.1 Description

In this section we will consider layouts where the elements and conducting interconnections are confined to the plane, but optical communication paths are allowed to leave the plane: (N2d,OPl).

We again assume that the N/N_1 modules are laid out on a 2 dimensional array and, as in appendix 16.11, that the system is confined in a box with dimensions $(N/N_1)^{1/2} d_m \times (N/N_1)^{1/2} d_m \times M_o(f\lambda)$, i.e. the height of the box is measured in units of $f\lambda$ and M_o is loosely interpreted as the number of ‘layers’.

Such a system can be realized by actually using M_o guided wave layers with interchannel spacing $f\lambda$, if such a structure can be constructed. Alternatively, it can be ‘wired up’ by using discrete fibers of diameter $f\lambda$. A free space imaging system would be most desirable, since it would allow a small value of f and would be easier to construct.

11.9.2 Analysis

The initial part of the analysis closely follows that of (N2d,O2d) layouts. The linear extent of each module must satisfy

$$d_1 \geq \left(N_1^{\frac{1}{2}} d_d, N_1^{\frac{1}{2}} K'_n W_{min}, d'_1, \left(\frac{k N_1^p E_o B}{Q} \right)^{\frac{1}{2}}, (k \chi N_1^p)^{\frac{1}{2}} d_{tr} \right) \quad (11.21)$$

which is similar to equation 11.5 apart from the absence of the last term. In the present case, the optical channels are allowed to vertically leave the modules.

The rest of the analysis follows that of appendix 16.11. The minimum intermodule separation is given by

$$d_m = \max \left(d_1, \frac{k N_1^p \chi \kappa (N/N_1)^{p-\frac{1}{2}}}{M_o} (f\lambda) \right) \quad (11.22)$$

where we are free to choose M_o . Of course, there is no utility in choosing M_o greater than $k N_1^p \chi \kappa (N/N_1)^{p-1/2} (f\lambda) / d_1$. Then the signal delay can be expressed as

$$\tau = \max \left(\left(\frac{N}{N_1} \right)^{\frac{1}{2}} \frac{d_m}{c}, \frac{M_o (f\lambda)}{c}, (16\rho\epsilon) \left(\frac{k\kappa}{M_n} \right)^2 N_1^{2p}, \frac{d_1}{v}, T_d \right) \quad (11.23)$$

where the second term accounts for the vertical contribution to the delay. Again, it makes little difference whether we actually add or take the maximum of the first two terms. The optimal value of M_o minimizing signal delay can be calculated as

$$M_o = \max \left(\min \left((k\chi\kappa N^p)^{\frac{1}{2}}, \frac{k N_1^p \chi \kappa (N/N_1)^{p-\frac{1}{2}} (f\lambda)}{d_1} \right), 1 \right) \quad (11.24)$$

which leads to

$$\tau = \max \left(\frac{(k\chi\kappa N^p)^{\frac{1}{2}} (f\lambda)}{c}, \left(\frac{N}{N_1} \right)^{\frac{1}{2}} \frac{d_1}{c}, (16\rho\epsilon) \left(\frac{k\kappa}{M_n} \right)^2 N_1^{2p}, \frac{d_1}{v}, T_d \right). \quad (11.25)$$

It is also easy to show that the height of the system $\mathcal{H} = M_o(f\lambda)$ is less than or equal to its lateral linear extent $\mathcal{L}$ with the optimal choice of M_o .

The calculation of total power dissipation is likewise similar to that for (N2d,O2d) layouts, other than the slightly different expression for d_1 . Since the constraint $d_1 \geq$




Figure 11.12: N_1 versus N for $(N2d, OPl)$ I. $p = 0.6$, $f = 2$, $E_o = 1$ pJ.

$k\chi N_1^p(f\lambda)$ is eliminated, equation 11.16 will be applicable for lower values of B as well.

Most of the time, the optimal value of M_o will bring d_m down to d_1 (equation 11.22). In such cases, minimization of signal delay is equivalent to minimization of $(N/N_1)^{1/2}d_1$. If all conducting lines are unterminated, as will often be the case, d_1 will be given by $d_1 = \max((kN_1^p E_o B/Q)^{1/2}, 2\epsilon V^2 k\kappa N_1^p B/Q)$. It is possible to show that the value of N_1 minimizing $(N/N_1)^{1/2}d_1$ is precisely that given by equation 11.16. Not so surprisingly, for a 3 dimensional layout which is heat removal limited, minimizing total power is equivalent to minimizing signal delay.

Once again N_1 must be chosen so as to satisfy a similar condition as equation 11.17.

11.9.3 Numerical examples

The optimal values of N_1 and the resulting dependence of S on N for our first example are shown in figures 11.12 and 11.13. The optimal values of N_1 for large N are given by equation 11.16, as discussed above. For smaller N , the optimal values of N_1 are limited to smaller values. This is because the electrical delay term $(16\rho\epsilon)(k\kappa/M_n)^2 N_1^{2p}$ must be kept below the system delay given in figure 11.13. For instance, in the device limited region where $\tau = T_d$, requiring that the electrical delay be less than T_d leads to $N_1 \leq 10^5$, as observed in figure 11.12. After this the delay is given by $bN^{1/2}$,




Figure 11.13: S versus N for $(N2d, OPl)$ I. $B = 100$ Mbit/sec, $p = 0.6$, $f = 2$, $E_o = 1$ pJ.

where the constant b depends on B , so that we require that the value of N_1 satisfy $N_1 \leq (b/(16\rho\epsilon)(k\kappa/M_n)^2)^{1/2p} N^{1/4p}$. This equation explains the curves for the optimal value of N_1 for the two lower values of B .

Unlike the $(N2d, O2d)$ layout where the optimal value of the system linear extent and S were equal to those for the all optical case, here we see that the value of S for the hybrid case is better than both all electrical and all optical alternatives (again as a consequence of the fact that the system is heat removal limited). This example provides a good illustration of how both media can be used in conjunction to obtain performance and cost much better than possible with any alone. An order of magnitude improvement in S and a hundredfold reduction in power is achieved over the all optical case.

Increasing f by a factor of one hundred has hardly any effect on the optimal value of N_1 and the resulting performance, and is not shown. Since the interelement spacing is set by heat removal, and not wireability requirements, increasing f within certain bounds has no effect on system size and delay.

So let us consider also increasing E_o by a factor of one hundred (figures 11.14 and 11.15). The value of S for the all optical case degrades by an order of magnitude (since delay is proportional to the square root of power dissipation). The value of S for the optimal hybrid case does not undergo a full order of magnitude degradation




Figure 11.14: N_1 versus N for $(N2d, OPl)$ II. $p = 0.6$, $f = 200$, $E_o = 100$ pJ.

however, since the system ‘adapts’ itself by employing fewer optical interconnections (larger N_1). For large N , the optimal values of N_1 hit the N_{1max} ceilings. For smaller N , the optimal values of N_1 are limited to smaller values so as to keep the electrical delay below the system delay. (Since overall larger delays are observed with respect to the case when the energy was $E_o = 1$ pJ, the conditions on N_1 based on this consideration are now weaker. However, overall smaller delays are observed with respect to the fully 2 dimensional case, so that the ‘envelope’ is tighter with respect to this case (figure 11.8).) We also note that there is a narrow region around $N \sim 10^6$ when the system is wireability limited and the slope of the curve for S is $p/2 = 0.3$.

Finally, let us look at the effects of increasing p to 0.8 (figures 11.16 and 11.17).

The optimal values of N_1 are explained in a similar manner as in the previous paragraph. Smaller values of N_1 are preferred with respect to the $p = 0.6$ case. The value of S for the all optical layout (given by equation 10.9) is the same as that when $p = 0.6$. The value of S for the optimal hybrid layout is somewhat worse than when $p = 0.6$, however, the effect of increasing p on S is relatively little compared to the 2 dimensional case. Also notice that $S \propto N^{-1/2}$ regardless of p .

When $E_o \sim 1$ pJ, the discrepancy between the all optical and optimal hybrid implementations is less for larger B , both in terms of power dissipation and delay. When E_o is increased by one hundred, the optimal hybrid system suffers less than the all optical system, since it can ‘adapt’ by employing a larger fraction of electrical




Figure 11.15: S versus N for $(N2d, OPl)$ II. $B = 100$ Mbit/sec, $p = 0.6$, $f = 200$, $E_o = 100$ pJ.




Figure 11.16: N_1 versus N for $(N2d, OPl)$ III. $p = 0.8$, $f = 200$, $E_o = 100$ pJ.




Figure 11.17: S versus N for $(N2d, OPl)$ III. $B = 100$ Mbit/sec, $p = 0.8$, $f = 200$, $E_o = 100$ pJ.

connections. The room for adaptation is less for lower values of B , since in this case we are already employing a large value of N_1 . For larger values of B , there is greater room to adapt, so that in this case the degradation in performance of the optimal hybrid system is less when E_o is increased.

11.10 Out of plane optical communication: repeaters and optics

11.10.1 Description

The system is identical to that discussed in the previous section with the exception that plain normal conductors are replaced with repeatered interconnections: $(R2d, OPl)$.

11.10.2 Analysis

The analysis is an obvious combination of the analyses of sections 11.8 and 11.9.




Figure 11.18: N_1 versus N for $(R2d, OPl)$ I. $p = 0.8$, $f = 2$, $E_o = 1$ pJ.

11.10.3 Numerical examples

First consider that $p = 0.6$, $f = 2$ and $E_o = 1$ pJ (not shown). In this case the use of repeaters does not offer significant improvement over the use of plain normal conductors. This is because the expression for energy is the same for unterminated lines, and because the optimal values of N_1 are small enough that the electrical delay term for repeaters is not significantly less than that for plain normal conductors.

Figure 11.18 and 11.19 are for $p = 0.8$. The optimal values of N_1 are explained as before. Making all connections optical results in a heat removal limited system. Making all connections repeatered results in a wireability limited system for all but the largest value of B . The signal delay for the all optical system is worse than that for the optimal hybrid system by only about a factor of two, for these large values of $p = 0.8$ and $B = 10$ Gbit/sec. We will have more to say about this in chapter 12.

As a final example let us consider figures 11.20 and 11.21 where E_o and f have been increased hundredfold. Quite larger values of N_1 are preferred with respect to the previous example, since optical communication is now more expensive and since the overall drop in S allows larger electrical delays to be tolerated.

Broadly speaking, the use of repeaters often offers limited improvement over the use of plain normal conductors, since repeaters do not start exhibiting a pronounced




Figure 11.19: S versus N for $(R2d, OPl)$ I. $B = 10$ Gbit/sec, $p = 0.8$, $f = 2$, $E_o = 1$ pJ.




Figure 11.20: N_1 versus N for $(R2d, OPl)$ II. $p = 0.8$, $f = 200$, $E_o = 100$ pJ.




Figure 11.21: S versus N for $(R2d, OPl)$ II. $B = 100$ Mbit/sec, $p = 0.8$, $f = 200$, $E_o = 100$ pJ.

advantage over plain normal conductors for the values of N_1 for which it starts being beneficial to use optics (often corresponding to connection lengths less than a centimeter).

11.11 Other systems

Other than those that have been described, we have also considered several other combinations of layout constraints and interconnection media. Here we will briefly mention some general considerations and results.

We have considered the case where fully 3 dimensional electrically interconnected modules are arrayed in a 3 dimensional optically interconnected grid: (N3d, O3d). The analysis is a straightforward 3 dimensional analog of the 2 dimensional analysis. The only difference is that heat removal requirements must be considered also when determining the intermodule separation, which must satisfy

$$d_m \geq \max \left(d_1, \left(k \chi N_1^p \left(\frac{N}{N_1} \right)^{p-\frac{2}{3}} \right)^{\frac{1}{2}} (f \lambda), \right. \\ \left. \left(\left(\frac{N}{N_1} \right)^{\frac{1}{3}} \frac{(k N_1^p E_o + \text{electrical energy per module}) B}{Q} \right)^{\frac{1}{2}} \right) \quad (11.26)$$




Figure 11.22: N_1 versus N for $(N3d, O3d)$. $p = 0.8$, $f = 2$, $E_o = 1$ pJ.

since the power dissipation from a stack of $(N/N_1)^{1/3}$ modules must be removed through the cross section d_m^2 . As a numerical example let us consider figure 11.22. The optimal values of N_1 are predicted by equation 11.16 which can be shown to hold for the 3 dimensional case as well.

We have also considered systems where a 3 dimensional array of two dimensional modules are interconnected optically: $(N2d, O3d)$. The optimal value of N_1 and the resulting value of S is hardly changed with respect to the $(N2d, OPl)$ case. Since the system is heat removal limited in both cases, it does not matter whether we confine the modules to the plane or array them in 3 dimensional space. In most cases, an (OPl) layout is as good as an $(O3d)$ layout. We have also observed that $(N2d, O3d)$ layouts are not much worse than $(N3d, O3d)$ layouts. We already saw that $(R2d, OPl)$ layouts were not much better than $(N2d, OPl)$ layouts in section 11.10. As a very crude conclusion, it is probably fair to say that if an $(O3d)$ or (OPl) organization is used for higher level interconnections, whether we use $(N2d)$, $(R2d)$, $(N3d)$ or $(R3d)$ for the lower level makes little difference. The system is heat removal limited, and the energy per transmitted bit for electrical interconnections is the same for all of these alternatives. This conclusion is more true for systems for which $E_o \sim 1$ pJ or so. For such cases, the optimum value of N_1 is such that an $(R2d)$ layout does not yet start exhibiting a strong advantage over an $(N2d)$ layout etc. For larger values of E_o however, it is desirable to choose larger values of N_1 . In such cases, it may be

beneficial to use (R2d), (R3d) or (N3d) rather than (N2d), so that the electrical delays can be kept as small as possible, enabling us to choose the value of N_1 minimizing total power dissipation and system size.

We now briefly discuss how some of the results are altered when superconductors are employed instead of optics. First consider the fully two dimensional case (N2d,S2d). Remember from chapter 10 that if the superconducting penetration depth is of the same order as $f\lambda$, and if the energy per transmitted bit E_s is similar to that for optics, the use of superconducting interconnections results in similar behavior to optical interconnections. However, also remember that it is relatively difficult to realize optical systems with $f \sim 1$ and $E_o \sim 1$ pJ. With superconductors, it should be easier to approach equation 10.52, provided satisfactory termination is possible. With $V = 1$ V and $T_d = 0.1$ nsec we find $E_s = V^2 T_d / Z_0 = 1$ pJ. E_s can be made much smaller if lower voltage values are employed. Furthermore, the use of superconductors can result in smaller system size. Thus, apart from the termination problem which we have not taken into account, superconductors can do at least as well as optics, and perhaps better, for fully 2 dimensional systems.

Since 2 dimensional optics cannot offer an advantage over superconductors, we considered the case where optical communication paths are allowed to leave the plane, but superconductors are not, and that $E_s = 1$ pJ and $E_o = 100$ pJ. As usual, for small N , an all normally conducting layout is preferred. With increasing N , it is beneficial to use superconductors for the longer connections. The use of optics is not yet beneficial, because for these relatively small values of N , 2 dimensional wireability limitations are not yet very severe. The use of optics would result in a heat removal limited system, not offering any advantage over 2 dimensional superconductors. For larger values of N however, the use of optics for the longer connections is desired, since 2 dimensional wireability limitations become exceedingly severe so that it is beneficial to employ out of plane communication. As expected, the switch to optics occurs later for larger B , because of the greater importance of heat removal.




Figure 11.23: N_1 versus N for $(N2d, O2d)$ V. $p = 0.6$, $f = 2$, $E_o = 1$ pJ.

11.12 Cost based optimization

Until now we concentrated on the optimization function given by equation 11.1 which gave full precedence to minimizing signal delay and only secondarily tried to minimize power dissipation. The cost of system size was not accounted for at all. Now we consider two other example optimization functions which account for the cost of system size and power dissipation.

First, we consider dividing equation 11.1 by the system area $\mathcal{L}^2$:

$$\Gamma = \frac{S}{\mathcal{L}^2 \mathcal{P}_e}. \quad (11.27)$$

Of course, it is always possible to employ more complicated functions if desired.

Figure 11.23 shows how the optimal values of N_1 are changed for the $(N2d, O2d)$ layout considered earlier. As discussed before, the entrance of optics is accompanied by a drastic increase in system size. Thus with our new figure of merit, it is beneficial to stick to an all electrical system until larger values of N despite the fact that the resulting signal delay will be worse than that possible with a hybrid system. However, once N tends to exceed N_{1max} , the entrance of optics is unavoidable. Once optical interconnections are entered, the steep increase in system size is observed. For $N > N_{1max}$, the optimal values of N_1 are identical to those in figure 11.3.

Now let us consider $(R2d, OPl)$ systems. This time we divide equation 11.1 by the




Figure 11.24: N_1 versus N for (R2d,OPl) III. $p = 0.8$, $f = 200$, $E_o = 100$ pJ.

system volume. So as to enable comparison of 3 dimensional all optical and hybrid layouts to the 2 dimensional all electrical layout, we must agree on a minimum height $\mathcal{H}_{min}$ for our systems. We will take $\mathcal{H}_{min} = 1$ mm. Then

$$\Gamma = \frac{S}{\mathcal{L}^2 \max(M_o(f\lambda), \mathcal{H}_{min}) \mathcal{P}^e}. \quad (11.28)$$

Figure 11.24 shows the optimal values of N_1 for relatively large values of E_o and f . We observe that all electrical systems are preferred until very large values of N . After a certain value of N however, the signal delay becomes so bad that we revert to a hybrid system. After this, the optimal value of N_1 is similar to that with our original figure of merit.

Now let us consider the figure of merit

$$\Gamma = \frac{S}{\mathcal{P}} \quad (11.29)$$

which puts greater emphasis on power than equation 11.1. Again we observe an increase in the optimal values of N_1 (figure 11.25), however less than when the cost of volume was emphasized.

In conclusion, accounting for the cost of system size and power dissipation results in all electrical systems being preferred until very large numbers of elements.




Figure 11.25: N_1 versus N for $(R2d, OPl)$ IV. $p = 0.8$, $f = 200$, $E_o = 100$ pJ.

11.13 Discussion and conclusions

In this section we present some of the more general conclusions derived from the study of a large number of examples, including those presented above.

Up to a certain value of N (which usually includes the device limited region for which $\tau = T_d$), it is preferable to make all connections electrical. In chapter 10, we saw that beyond a critical value of N , even an all optical layout was preferable to an all normally conducting layout. Thus it will be beneficial to start employing optics on or before this critical value.

Broadly speaking, we have found that the use of optics may enable performance⁴ unachievable otherwise, however possibly at a significant cost of space and/or power, depending on the information density and energy per bit achievable. The use of optical interconnections may enable smaller signal delay and larger bandwidth in comparison to an all electrical system. However, they will usually not enable a reduction in our measures of cost. The system size and power dissipation of the higher performance hybrid system will often exceed that of the all electrical system.

This behavior can be a result of a number of reasons. First of all, remember that optical lines cannot be scaled down to submicron dimensions so that they lead to

⁴Remember from chapter 4 that to increase performance means to be able to access previously inaccessible points in S - B - N space in a desired direction.

large system sizes. However, they may still enable greater performance than an all electrical system, since the signal delay along a longer optical line can be less than the rise time of a shorter normally conducting line. A similar argument applies to heat removal limited systems. A hybrid system may dissipate more power than an all electrical system, leading to larger system size, but may still have smaller delay.

Another reason why we must use optics despite the large cost of system size and power dissipation is because we may want to increase N and B simultaneously beyond what is possible with normal conductors. In such cases, we might be willing to employ optical communication regardless of how large f , E_o and the resulting cost of space and power is. We can avoid falling into this situation by using repeaters; however there are still points in S - B - N space which are not accessible without the use of optics, no matter at what cost. B may be arbitrarily increased with repeaters, but in some cases resulting in larger signal delay than possible with a hybrid system (and sometimes even larger cost of space and energy).

Figure of merit functions emphasizing the cost of system size and power dissipation tend to favor all electrical systems until very large numbers of elements. If full priority is given to minimizing cost of system size and/or power dissipation, almost always an all electrical system will be preferred. (The major exception arises when we want to increase N beyond N_{1max} .) Even if equal priority is given to minimizing signal delay and measures of cost, as in our examples (section 11.12), all normally conducting systems are usually preferred until $N > N_{1max}$. When repeaters are used, all electrical layouts are preferred until even higher values of N , until eventually a hybrid layout is preferred because the signal delay becomes much worse than possible with the hybrid system, outweighing the emphasis put on measures of cost.

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

is wireability limited), although it strongly affects the optimal value of N_1 .

If $f \sim 2$ for 2 dimensional systems, then a value of $E_o \sim 1$ pJ leads to optimal values of N_1 of the order of $\sim 10^3$ - 10^4 . Increasing E_o hundredfold results in $N_1 \sim 10^6$. If repeaters are used, optical interconnections are useful only for very large values of B (when the system is heat removal limited). If $f > 10$ or so, there hardly seems to be any room for the beneficial use of optics, if repeaters are available. If not, we may resort to optics so as to achieve large values of N and B simultaneously. However if f is too large, the resulting system size may quickly reach unrealistic proportions (> 10 m), so that such a system would not be feasible anyway.

3 dimensional systems tend to be heat removal limited, so that the value of f is relatively unimportant. A hundredfold degradation from what seems the best possible ($f \sim 2$) has little effect on the results. On the other hand, the value of E_o not only determines how much optics should be used, but also has considerable effect on the resulting performance. If $E_o \sim 1$ pJ, the optimal value of N_1 is of the order of $\sim 10^4$ whereas a hundredfold increase in E_o leads to $N_1 \sim 10^6$.

The degradation in signal delay corresponding to a degradation in E_o is less for larger values of B , since such systems have greater room to adapt to a larger value of E_o by employing a larger value of N_1 , relative to systems with smaller values of B (which already employ large values of N_1). Likewise, systems with smaller values of p are hurt less by a degradation in E_o , since such systems have a smaller fraction of longer connections and are relatively less dependent on using optics. On the other hand, an increase in E_o by one hundred may lead to nearly an order of magnitude degradation in signal delay for systems with large p , nearly as much as the all optical case.

The value of p has less effect on resulting signal delay for 3 dimensional systems. The power dissipation and system size of a heat removal limited all optical system are independent of p . For a hybrid system, smaller values of p enable some energy savings and consequent reduction in system size, since a greater fraction of connections will fall below the breakeven length for energy and can be replaced with normal conductors.

Hybridization is most important when the system is heat removal limited. For 2 dimensional systems which are often wireability limited, hybridization has little




Figure 11.26: τ and $\mathcal{P}$ versus N_1 for $(N2d, OPl)$. $N = 10^8$, $B = 100$ Mbit/sec, $p = 0.6$, $f = 2$, $E_o = 1$ pJ.

effect on resulting signal delay. An all optical system is just as good in this respect (although it may have much larger power dissipation).

When E_o is of the order of ~ 1 pJ, even the longest normally conducting lines are usually short enough to be left unterminated (of the order of a centimeter or less). Terminated lines are not observed; RC lines are followed directly by optical lines. It is fortunate that the breakeven occurs at a point where terminated lines are not yet necessary, since satisfactory termination can be a significant problem. For large values of E_o (or f in 2 dimensional layouts) however, the module size may be forced to be large enough that the longer normally conducting lines have to be terminated.

Based on the considerations of this thesis, it seems fair to conclude that if the optical communication energy can be reduced to the order of a picojoule and out of plane architectures with $f \sim 100$ or less can be realized, the use of optics has the potential to contribute significantly to the performance of large scale systems.

As a final remark, we note that there is considerable latitude in the choice of N_1 , leaving room for other technological considerations. In other words, we can deviate quite a bit from the optimal value of N_1 while deviating only little from the optimal performance and cost. This is illustrated in figure 11.26 where the signal delay and power are shown as functions of N_1 for $N = 10^8$ and $B = 10^8$ Mbit/sec. We leave it to the reader to trace down the origins of the various regions in the figure.

11.14 Related work

The problem of comparing optical and normally conducting interconnections has received steadily increasing attention since 1984. As mentioned at the beginning of this chapter, many authors have made comparative studies of isolated lines. Some authors have also compared all optically connected systems to all electrically connected systems. Feldman et al. compare a 3 dimensional optical system to a 2 dimensional electrical system [43]. In a very recent study Stirk et al. compare 3 dimensional optical and electrical permutation network implementations based on yield considerations [153].

The problem of how to use both media in conjunction has received less attention. Other than this work, we are aware of only one very recent study of a similar nature. Kiamilev et al. discuss how a perfect shuffle network should be partitioned into VLSI chips which are then interconnected optically [88].

Chapter 12

On the Usefulness of Optical Digital Computing

In chapter 11, we addressed the problem of determining the optimal mix of optical and normally conducting interconnections resulting in a system with optimal properties. In some cases, we found that making all connections electrical was optimal ($N_1 = N$). However, it was never the case that making all connections optical was optimal ($N_1 = 1$). Nevertheless, it is interesting to inquire how worse off we are if we do make all connections optical. In situations where it is found that this does not result in performance and cost much worse than the optimal hybrid system, it might be meaningful to consider the construction of an all optical computer.

In this chapter we will speculate as to when an optical digital computer might be useful, based on the considerations of chapter 11. Since optical communication is preferable for establishing connections exceeding a certain length, for large system sizes the beneficial use of normally conducting wires for the shortest connections becomes an edge effect and can be ignored. This suggests that the performance and cost of an all optical computer might not be much inferior to an optimal hybrid alternative. We argue that for applications for which high bit repetition rates are useful despite large propagation delays, it might make sense to contemplate the construction of an optical digital computer.

Several researchers have addressed this issue in the past, often with negative conclusions. We believe that the increasing understanding regarding the importance of communication in computing and the realization that the architectural-logical construction of a computing system can no longer be divorced from its physical construction justifies a reevaluation of previous arguments and a search for hitherto unexplored perspectives.

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

12.1 Introduction

The possibility of an optical digital computer has attracted considerable attention. Despite the vast literature on devices, systems, architectures and algorithms for optical digital computing [23], there has been considerable discomfort as to its usefulness as compared to other approaches. Switching energy arguments on power-delay diagrams, such as in [79] [151] have resulted in the digital optical computer being mostly viewed as an esoteric curiosity. Following Goodman et al.'s seminal paper [59], the interest in 'optical computing' shifted towards 'optical interconnections'. This led to the notion of the hybrid optoelectronic computer. Nevertheless, the assumptions underlying the negative arguments have not remained unchallenged [75] and optical computing research has not ceased. In this chapter we explore several issues relating to the potential usefulness of an optical digital computer and try to identify the conditions under which it would make sense to consider it as an alternative to existing approaches.

A digital computer can be viewed as a collection of nonlinear switching elements interconnected according to a certain graph. The nonlinear switching elements often rely on electronic interaction. Usually, this is true even of so-called 'optical switches', since in most optical switches, photons interact indirectly via electrons. On the other hand, communication among the elements is often established via photons,

even in fully electronic computers [123]. The question is whether conductors are used to confine the wave fields. Thus, we define an all optical computer to be one which does not employ conducting materials for the purpose of guiding signals among its switching elements. The all optical computer, as defined, is a special case of hybrid optoelectronic computers which employ both conducting wires and optical communication for this purpose.

We will first review some of the essential ingredients of our discussion. Then we argue that an all optical computer need not be much worse than an optimal hybrid computer under certain conditions, so that an all optical computer might be preferred for its relative simplicity compared to a hybrid computer. Then we discuss some objections to our argument and some issues pertaining to the construction and utilization of an all optical computer.

12.2 System characterization and evaluation

As discussed in chapter 4, we will characterize our processing systems with 3 parameters:

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

Notice that here the delay τ is defined somewhat differently than in chapter 4. In previous chapters we explored the relationships between S , B and N for a given graph of connectivity p , and compared various implementations in terms of performance and

cost. In this chapter we will try to compare the performance of systems with different connectivities. The above definition of S will serve as a general measure of the speed at which information ‘percolates’ among the elements of our system. Also, note that in previous chapters, we mostly considered systems where the length of the longest connection was of the order of the linear extent of the system, whereas in this chapter we will also consider examples for which this does not hold.

Remember that how much we can increase all three of these quantities simultaneously depends on the choice of interconnection medium and is ultimately limited by the laws of physics, and that the particular application/algorithm at hand will determine which we will prefer to increase at the cost of the other(s).

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

12.3 Lower bounds to system size

We already noted that minimizing global propagation delays is equivalent to minimizing system linear extent. We know from earlier chapters that two major physical considerations lead to lower bounds to system size: wireability and heat removal requirements. In chapter 3, we argued that the linear extent of a 3 dimensional system with constant power dissipation per element will become limited by heat removal considerations with increasing N . In some situations, the power dissipation per element may increase with system size, strengthening our conclusion. (Note that if B does not increase with N , the existence of optical interconnections ensures that the power dissipation per element need not increase with N either.) On the other hand, if the power dissipation per element were to decrease with increasing N , this might invalidate our conclusion. We will come back to this point later.

Given that there are upper limits to device speed, the best we can do to construct ever powerful computing systems is to make them 3 dimensional and of an ever increasing number of elements N . Thus, based on the argument of the preceding

paragraph, we conclude that heat removal will be the major factor determining how densely we can pack the elements of our system. This means that employing layouts with larger fractal dimension (or Rent exponent) will not result in greater system size and delay (still assuming constant power dissipation per element). This suggests that highly interconnected approaches, which offer greater functional flexibility, will be preferred in future large scale systems.

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

12.4 Optimal hybrid partitioning

Minimization of system size and propagation delay is possible with a hybrid layout, involving both normally conducting and optical interconnections, as discussed in detail in chapter 11. Let us consider an $e = 3$ dimensional system: (N3d,O3d); with $N = 10^6$, $k = 5$ and $p = 0.8$ ($n = 5$) operated at $B = 10$ Gbit/sec. Let the optical communication energy be $E_o = 1$ pJ, based on the discussion of chapter 5; and the electrical energy be $\gamma = 100$ fJ/mm per unit length, so that the energy per transmitted bit for a normally conducting line of length ℓ is $E_n = \gamma\ell$ [121]. We have assumed that an unterminated line is charged up to 1 V. The corresponding expression for terminated lines is more complicated, as discussed in chapter 5; however it is in general true that the energy per transmitted bit must increase with line length because of resistive loss. For this choice of parameters, optical interconnections require less energy per transmitted bit for lines longer than $\ell = 1$ cm. Finally, let us be capable of removing $Q = 10$ W/cm² of power per unit cross section of our system.

We will now show that under the above assumptions, employing an optimal hybrid combination of normal conductors and optics results in a system with linear extent

$\mathcal{L} \simeq 0.5$ m. Any other than the optimal mix of optics and normal conductors will result in a larger linear extent.

In the following simplified derivations we consider only the effects of heat removal. Wiring density, bandwidth and rise time considerations are actually all coupled to energy considerations. However, since heat removal is the dominating consideration, detailed calculations as in chapter 11 give similar results.

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

First, let's find the minimum size $\mathcal{L}_1$ of an electrically connected cube of N_1 elements. The average connection length per element is given by $k\bar{\ell} \simeq kN_1^{p-2/3}d$ for an $e = 3$ dimensional layout¹. Thus the power dissipation per element is $\gamma(kN_1^{p-2/3}d)B$. Also $d = \mathcal{L}_1/N_1^{1/3}$. Heat removal requires that the total power dissipation associated with the N_1 elements not exceed $Q\mathcal{L}_1^2$ where Q is the amount of power we can remove per unit cross section. Thus

$$Q\mathcal{L}_1^2 \geq N_1 k N_1^{p-2/3} \frac{\mathcal{L}_1}{N_1^{1/3}} \gamma B \quad (12.1)$$

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

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

$$Q\mathcal{L}^2 \geq \frac{N}{N_1} [kN_1^p E_o B + \mathcal{P}_1]. \quad (12.2)$$

Substituting for $\mathcal{P}_1$, minimizing over N_1 we find

$$\mathcal{L} \geq N^{\frac{1}{2}} \left(\frac{kB}{Q} \right)^{\frac{1}{2p}} E_o^{\frac{(p-1/2)}{p}} \gamma^{\frac{(1-p)}{p}} y^{\frac{1}{2}} \quad (12.3)$$

where $y = x^{(p-1)/p} + x^{(2p-1)/p}$ and $x = (1-p)/(2p-1)$. Notice that whereas a lower value of p would result in smaller $\mathcal{L}$, it would not change the dependence of $\mathcal{L}$ on

¹In this chapter we take $\kappa \sim 1$ for simplicity.

N which is still $\propto N^{1/2}$. With our chosen parameters we find the optimal value of $N_1 \simeq 190$, resulting in a system linear extent $\mathcal{L} \simeq 0.5$ m. This result is supported by figure 11.22.

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

In general, the ratio of the linear extent of the optimal hybrid system to that of the all optical system is given by (for $n = 1/(1 - p) > 3$)

$$\left(\frac{kB\gamma^2}{QE_o} \right)^{\frac{1}{2(n-1)}} y^{\frac{1}{2}} \quad (12.4)$$

with $kB\gamma^2/QE_o = 1/200$ for our numerical example. Interestingly, this ratio is independent of N . If p is large, say $p = 0.8$ as in our example, it is insensitive to the other parameters.

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

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

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

12.5 An argument for optical digital computing

The all optically connected system described at the end of the preceding section already qualified as an all optical digital computer according to our definition, since it does not utilize any conductive wiring for communication among its elements. However, we can do even better by replacing the discrete detector-electrical switch-source/modulator combinations by their integrated versions. We will speak of an integrated version of such a combination as an optical switch. A SEED [122] [124] is an example of such a device.

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

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

We already argued that future 3 dimensional computing systems of ever increasing numbers of elements will tend to employ large values of p . If, in addition, we *assume* that large values of B are employed, we see that the conditions stated at the end of the previous paragraph coincide with the trend in constructing ever more powerful computing systems.

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

12.6 The importance of bit repetition rate B

With increasing system size and delays, some algorithms may have a tendency to be bottlenecked by communication latencies, so that working with a high bit repetition

rate B may not be of any utility. Thus, there may be a tendency to operate at slower rates with increasing N . This would invalidate our argument in many ways:

1. First of all, as evident from our analysis, when B is not high, the discrepancy between the all optical implementation and the optimal hybrid implementation increases.
2. If B decreases with N , the power dissipation per element decreases, weakening our argument that 3 dimensional systems become heat removal limited with increasing N .
3. When the value of B is less than the large intrinsic bandwidth of the optical communication channels, it is possible to employ various strategies to exploit this bandwidth in order to reduce system size. In such cases, the use of normal conductors for the shorter connections may be useful. This is discussed in chapter 13 and [137].

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

12.7 Comparison of systems with different connectivity

In chapter 3 and section 12.3 we argued that since $e = 3$ dimensional systems are heat removal limited, the value of p has no effect on the resulting system linear extent. However, since our assumption regarding constant power dissipation per element does not hold for hybrid designs, this is no longer precisely true, as evident from equation 12.3. So as to provide a basis for comparison, let us also calculate the linear extent of a system with parameters identical to that considered above, excepting p . Let all connections in this system be to nearest or second nearest neighbors only so that the average connection length per element is $\simeq kd$. (This layout has $n = e = 3$

and $p = 1 - 1/n = 2/3$.) Since all connections are short, we consider making all of them electrical. The heat removal condition may be written as

$$Q\mathcal{L}^2 \geq N\gamma(kd)B. \quad (12.5)$$

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

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

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

Tables 12.1 and 12.2 summarize these results, which will be further discussed in a later section.

n	All optical $\mathcal{L}$	Optimal hybrid $\mathcal{L}$	All electrical $\mathcal{L}$
3	0.7 m	0.27 m	0.5 m
5	0.7 m	0.5 m	3.2 m

Table 12.1: *System linear extent for $n = 3$ and $n = 5$ dimensional meshes laid out in $e = 3$ dimensions. $N = 10^6$, $k = 5$, $B = 10$ Gbit/sec, $E_o = 1$ pJ, $\gamma = 100$ fJ/mm and $Q = 10$ W/cm².*

n	All optical HB	Optimal hybrid HB	All optical τ	Optimal hybrid τ
3	500 Tbit/sec	500 Tbit/sec	12 nsec	11 nsec
5	3200 Tbit/sec	3200 Tbit/sec	3.9 nsec	3.3 nsec

Table 12.2: *Bisection-bandwidth product and global delay for $n = 3$ and $n = 5$ dimensional meshes laid out in $e = 3$ dimensions. $N = 10^6$, $k = 5$, $B = 10$ Gbit/sec, $E_o = 1$ pJ, $\gamma = 100$ fJ/mm, $Q = 10$ W/cm² and device delay = 100 psec.*

12.8 Some credible objections

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

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

A major objection to our argument rests on the following question: what if the elements are not just gates but more complicated circuitry, still of the same size and number of pinouts. In this case, even if making all connections optical is shown to

be not worse than the optimal hybrid combination, we cannot replace these detector-element-source/modulator combinations with simple optical switches. In fact, this may be considered as a proof that a hybrid system can always have more computational resources than an all optical system. The size of an optical switch cannot be less than $\sim 1 \mu\text{m}$, whereas in a microns space, several electrical switches can be squeezed if deep submicron technology is employed. However, Rent's rule will be broken: the switches will have limited communication with the outer world. Information flow at the original fractal dimension is not possible through the boundary of the elements and some form of serialism must be employed.

Thus this objection boils down to the question of what usefulness an element of internal sophistication but limited pinouts may have. Of course, an element with sophisticated internal structure would always be desirable over a simple switch with the same number of pinouts, if it is available at no cost, since it could easily simulate the simple switch. However, the question is whether such elements will have sufficient utility to make their usage worthwhile. A simple example of an element with limited pinouts but arbitrarily large internal structure is a shift register, which can serve as a memory. However, this memory will have a large access time. We might argue that a shift register is always preferable to a single gate, since a slow memory is better than none. On the other hand, this memory might have such limited usefulness in certain situations due to its slowness that we might not bother having it.

Another example is a microprocessor which has far less pinouts than Rent's rule implies [44]. Whereas such so-called 'functionally complete' units are consistent with existing systems approaches, it is not clear that this is the most advantageous way of constructing supercomputers [84]. It may be that a massive and homogeneous collection of switches connected with a uniform Rent exponent is more beneficial. Alternatively, it may be that optical nodes with some internal sophistication and function integrated into them do have some usefulness. Whether such an integrated structure is still considered to be an optical switch or rather an electronic circuit with optical ports is a matter of definition. Indeed, Lentine et al. [108] have noted the fuzziness in the definition of an 'all optical' computer: "Because of the limited functionality achievable in 'all optical' logic gates, a growing interest is seen in 'optical'

processing elements made using optoelectronic devices with greater functionality.”. Whether or not this trend is merely a legacy of existing design approaches, it seems that the first optical computers and switching systems will employ ‘optical’ nodes with relatively sophisticated functionality.

We cannot arrive at any general conclusions. Whether limited pinout elements are useful will mainly be determined by the particular application. Nevertheless, general observation seems to indicate that the percolation of information is crucial in the working of computing and cybernetic systems [84]. It is difficult to think of where human civilization would be if we had less of our five senses, even magically granted the same brains. Living systems are characterized by their ability to keep their entropy low and their orderliness high. They achieve this by energy consumption and information exchange with their surroundings [178] [21] [148]. Perhaps, such considerations may be thought to apply to artificial systems too.

12.9 What might the optical computer look like?

The example of section 12.4 illustrates the order of numerical values for which it is meaningful to consider an all optical computer. For greater performance we might increase N as much as possible. For instance, a system with 10^8 elements would be 7 m in size.

How can we quantify the performance of our system? A figure of merit of performance which jointly emphasizes parallelism, connectivity and bandwidth is the bisection-bandwidth product, defined in chapter 4. The bisection-bandwidth product of the system we considered as an example is given by $kN^pB = 3200$ Tbit/sec. A system with $N = 10^8$ elements would have $kN^pB > 10^5$ Tbit/sec.

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

The vision of the digital optical computer emerging from our considerations is as follows. It will be large ($N \sim 10^7$ - 10^8 and $\mathcal{L} > 1$ m), highly interconnected, operated

at very large repetition rates (multi Gbit/sec) and due to its size exhibit large speed of light limited delays between its distant elements (~ 10 nsec). Its bisection-bandwidth product ($kN^p B$) might be of the order of 10^{16} bit/sec. It will be suitable for situations in which a large repetition rate is useful despite large speed of light limited propagation delays. In other words, it will be appropriate for applications where one prefers large values of B at the expense of N and S .

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

It is preferable to use free space optics rather than discrete fibers. However, conventional imaging systems allow free space interconnections at high density only for a regular pattern of connections. Since the system is heat removal limited anyway, multi-facet holographic approaches [92] [43] may be used to provide an arbitrary pattern of connections for smaller N . However, the resulting system size with these approaches is crudely $\mathcal{L} \simeq kN\lambda$ (chapter 7 and [134]) so that for $N > 10^5$, they would not be desirable.

We do not know if a system with a regular pattern of connections of the same computational power as that of an irregularly connected one is always possible. Huang seems to argue in favor [74] whereas the maximum entropy approach of Keyes would seem to indicate otherwise [84] [86].

If the 3 dimensional multi-facet architecture of chapter 7 can be built, it would solve once and for all the problem of being restricted to a regular pattern of connections.

12.10 What is it good for?

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

As mentioned before, a 5 dimensional mesh with radix $N^{1/5}$ is an example of a graph with $p = 0.8$ and $k = 5$. This graph may represent the connection pattern of a 5 dimensional nearest neighbor (in 5 dimensions) connected cellular automata. In section 12.7, we also considered a 3 dimensional mesh, which may represent the connection pattern of a 3 dimensional nearest neighbor connected cellular automata. The resulting linear extent, bisection-bandwidth product and global delay for these systems were presented in tables 12.1 and 12.2. Based on a comparison of their bisection-bandwidth products and global delays, we conjecture that there exist concrete problems (such as sorting etc.) for which the time it takes to solve these problems on our 5 dimensional examples is less than on our 3 dimensional examples, despite the fact that our 5 dimensional examples have greater linear extent and connection lengths. If one can find an example of a problem which can be solved in shorter time on our optimal hybrid 5 dimensional mesh of $\mathcal{L} = 0.5$ m than on the optimal hybrid 3 dimensional mesh of $\mathcal{L} = 0.27$ m, this would serve as an existence proof that indeed for some applications, highly interconnected approaches are preferable for large $e = 3$ dimensional systems. It would be even more interesting to provide an example problem which can be solved in shorter time on the all optical 5 dimensional mesh with $\mathcal{L} = 0.7$ m, than on the optimal hybrid 3 dimensional mesh with $\mathcal{L} = 0.27$ m.

Multidimensional cellular automata offer an alternative to computing based on functionally complete entities (assuming one finds a better way of performing useful computation other than simulating conventional logic functions in the automata). The latter approaches allow us to view functionally complete entities at a lower level as black boxes, greatly facilitating design and construction. Without the benefit of such structuring, the conception and design of large scale 'arbitrarily connected'

systems would be intimidatingly complex. One way to avoid this is to resort to systems exhibiting some form of regularity. This brings us back to the question of whether requiring such regularity takes back any of the potential advantages.

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

12.11 Beyond dissipative computing

Our discussions assume dissipative computing. There is general consensus that dissipationless computing does not contradict the laws of physics. In particular, it seems that it is possible to perform a large number of operations with only limited dissipation (if we don't care to observe intermediate results) [24].

In the limit of very large systems and extremely fast devices with very little or no energy dissipation, one envisions that a nearest neighbor (in 3 dimensions) connected cellular automata can be used to simulate any other computing system with little degradation, as well as any other finite portion of the universe.

12.12 Conclusion

To build ever powerful processing systems, we must increase the number of elements N . If the bit repetition rate B along the connections of our system is large, heat removal considerations tend to dominate wireability considerations for 3 dimensional systems, suggesting (but not proving) that highly interconnected approaches may be preferred to increase parallelism and functional flexibility. For such systems, the fraction of connections with lengths greater than the energy breakeven between normally conducting and optical interconnections will be large, so that we might as well make all connections optically. It is meaningful to consider such a system because using conductive wiring to establish the shorter interconnections will not result in considerable improvement in system size and global delays. (Nor will it necessarily make them worse, and it is a subjective issue whether it is considered more simple to keep it all optical or to keep the number of optical connections minimum by using

conducting wires as much as possible.)

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

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

Ultimately, more solid answers to the questions raised in this chapter will emerge when the theory of algorithms is merged with a physically realistic theory of computer construction to create a physical theory of computation, as alluded to in chapter 4. Then it will be possible to compare and optimize jointly over many possible constructional *and* algorithmic approaches.

Chapter 13

Indirect Implementations

Until now, we limited our attention to direct implementations where the communicative connection graph was directly copied as the physical connection graph. In this chapter we will discuss the importance of organizing information flow in computation in a manner enabling multiplexing of signal paths with distinct source and destination localities, and show how this can be achieved by indirect implementation of the communicative connection graph. Such implementations allow efficient use of high bandwidth interconnection media, leading to a decrease in system size and propagation delay for wireability limited layouts. Among the three methods we consider, the fat tree architecture is found to be near optimal.

13.1 Introduction

As we have discussed in detail in earlier chapters, a major advantage of optical and superconducting interconnections is their ability to transfer large amounts of information per unit cross section over long distances. Let the maximum information flux a given communication medium can support be denoted by $\mathcal{I}$ and be measured in bit/m²sec. For the length scales involved in a computing system (< 10 m), it is possible to reduce the effects of dispersion and attenuation to the extent that $\mathcal{I}$ may be assumed to be independent of length for optical and superconducting interconnections. On the other hand, $\mathcal{I}$ is a decreasing function of communication length for

resistive interconnections, making them disadvantageous over longer distances. However, for distances less than about the order of a centimeter, they can provide greater information flux than optical or superconducting interconnections.

As defined earlier, T_r will denote the minimum pulse repetition interval for a single physical optical communication channel (i.e. corresponding to a single spatial degree of freedom). Since we are ignoring dispersion, T_r will probably be set by the speed of the switching devices or electrooptic transducers. If wavelength division multiplexing is employed, an appropriate effective value of T_r should be used.

As usual, we would like to establish a prespecified pattern of kN pairwise connections among a collection of $N \gg 1$ elements with k connections per element. B will denote the rate at which binary digital pulses are emitted into each connection. Although we restrict ourselves to a fixed connection pattern, the extension to reconfigurable or message routing systems should be possible. We also limit ourselves to single layer 2 dimensional layouts, the extension to multilayer and 3 dimensional layouts being straightforward. As usual, our purpose is to implement the given pattern of connections in a manner that results in smallest possible system area, which we assume is dominated by the space required for establishing communication (which is often true for 2 dimensional systems).

The number of binary pulses in transit at any given time in an optical communication network occupying area $\mathcal{A}$ may not exceed $\simeq \mathcal{A}/(f\lambda cT_r)$, where c and λ denote the speed of light and wavelength of radiation respectively. (This is the 2 dimensional version of the result derived in chapter 6. As defined earlier, f is a dimensionless constant factor which in principle can approach the order of unity, but may be larger in practice.) The number of pulses in transit in our system is given by $kNB\tau_{ave}$ where $\tau_{ave} = \bar{\ell}/c$ is the average delay and $\bar{\ell}$ is the average interconnection length. Remember that $\bar{\ell}$ may be expressed in terms of the system linear extent $\mathcal{L} = \mathcal{A}^{1/2} = N^{1/2}d$ as $\bar{\ell} = \kappa N^{p-1/2}d = \kappa N^{p-1}\mathcal{L}$ when $p > 1/2$. Using these relations, we find a lower bound on the area and linear extent of our system:

$$kNB\tau_{ave} \leq \frac{\mathcal{A}}{f\lambda cT_r} \quad (13.1)$$

$$\mathcal{A} \geq kN(BT_r)c\tau_{ave}f\lambda = kN(BT_r)\bar{\ell}f\lambda \quad (13.2)$$

$$\mathcal{L} = \mathcal{A}^{\frac{1}{2}} \geq k\kappa N^p (BT_r) f\lambda. \quad (13.3)$$

The above bounds represent the intrinsic information carrying capacity of optical wave fields and apply to any architecture or implementation. Notice the tradeoff between system size and B .

One way of implementing the desired pattern of connections is simply to allocate $\lceil BT_r \rceil \simeq \max(BT_r, 1)$ parallel channels between every pair of points to be connected. When $BT_r \geq 1$, such an implementation is as efficient as any other in terms of making maximum usefulness of the available capacity of the optical channels. In this case, the above lower bounds may be approached, for instance, by use of waveguides with effective¹ line to line spacing of $f\lambda$. To see this, notice that the total area required for communication is $\mathcal{A} = kN(BT_r)\bar{\ell}f\lambda$, since $BT_r \geq 1$ physical channels, each occupying an average area of $\bar{\ell}f\lambda$ are allocated per connection. Thus $\mathcal{A} = kN(BT_r)\kappa N^{p-1}A^{1/2}f\lambda$, leading to equation 13.3. However, if B is less than $1/T_r$, the channels are underutilized and the bound of equation 13.3 cannot be approached, since no matter how small B is, a channel with capacity $1/T_r$ is allocated for every pairwise connection. Thus when $B < 1/T_r$, the layout area is not any less than when $B = 1/T_r$, so that $\mathcal{L}$ can at best approach the bound

$$\mathcal{L} \geq k\kappa N^p f\lambda. \quad (13.4)$$

Equations 13.4 and 13.3 are analogous to equations 10.5 (with $\chi = 1$) and 10.7.

In the rest of this chapter we concern ourselves with methods of restoring the broken tradeoff between system size and B when $BT_r < 1$. If B is independent of N , such methods may lead to reduction of the system linear extent by a constant factor of BT_r , compared to the direct implementation just described (equation 13.4). In some cases, B may decrease with increasing N , since the computational processes become bottlenecked by the increasing propagation delays. When this is the case, restoring the mentioned tradeoff allows one to slow down the growth rate of system size as a function of N , as evident from equation 13.3. In fact, if B decreases at the same rate as $1/\tau_{ave}$, linear growth of $\mathcal{A}$ as a function of N can be achieved regardless of the value of p , as evident from equation 13.2.

¹That is, including all inefficiency factors due to routing etc.

To achieve our objective, we would like to multiplex $1/BT_r > 1$ independent signal paths into the same physical channel, so as to saturate its capacity. However, this is not straightforward when the many signal paths have distinct source and destination localities. In the following sections we describe 3 architectures which enable information flow to be organized in a manner enabling overlap between such signal paths, allowing them to be multiplexed. The reduction in the number of physical channels thus possible results in a decrease in system size and propagation delay for communication limited layouts.

13.2 The multiplexed grid architecture

The multiplexed grid architecture is based on the family of m -ary ν -dimensional meshes (grids) of $m^\nu = N$ nodes [30]. The hypercube [169] is a special case with $m = 2$ and $\nu = \log_2 N$. For sake of illustration, we consider the case $\nu = 2$ and $m = N^{1/2}$, which corresponds to the familiar planar mesh with $N^{1/2}$ nodes on an edge. An arbitrary connection is established in several nearest neighbor (in ν -space) ‘hops’, and multiplexed together with other connections with which it overlaps, as illustrated in figure 13.1. Notice that this procedure enables us to break down independent signal paths into overlapping segments which may then be multiplexed. If at least $1/BT_r$ connections can be overlapped along each edge of the mesh, then complete utilization of the available capacity $1/T_r$ of the physical channels may be achieved. Finally, the multiplexed ν dimensional mesh is laid out in 2 dimensions, as described in [30]. Of course, this is a trivial task when $\nu = 2$.

The price that must be paid in return for efficient utilization of the high capacity optical channels is the additional area and delays associated with routing of independent signal paths. Low dimensional meshes allow a larger number of connections to be overlapped, but increase the number of hops, and hence the number of device delays a signal must go through. High dimensional meshes decrease the number of hops but do not enable as many signal paths to be overlapped and multiplexed, possibly resulting in less than complete utilization of the capacity of the channels and thus larger layout area and propagation delays. The optimal value of ν minimizing




Figure 13.1: *The multiplexed grid architecture with $\nu = 2$, $m = 4$ and $N = 16$. Part a. shows two of many to-be-established connections. Part b. shows each connection established in several hops. Part c. shows overlapping portions of these connections multiplexed into high capacity channels, reducing the total number of physical channels and thus layout area.*

overall signal delay (propagation plus device) is found to decrease with increasing N and asymptotically approaches 2 for 2 dimensional layouts.

Emulation of fixed connection graphs on multidimensional meshes in the manner described may in certain cases enable one to improve over the bounds $\Phi(S, B, N) \leq C_\Phi$ derived in chapter 10 and also alter the comparison of interconnection media.

Hartmann and Ullman [63] discuss what they call a *delay balanced* architecture, which is a special case of what we are discussing with $p = 1$ (uniform message distribution assumption) and $\nu = 2$, in the context of simulating a parallel random access machine. When $p = 1$ the average and worst case signal delays differ by only a constant factor. They show for this case that if the bit repetition rate is decreased according to $B = \beta S = \beta/\tau$ where β is a constant, then $\mathcal{A} \propto N$, i.e. the area increases only linearly as a function of the number of elements. More generally (when $p \neq 1$), this will be true if $B = \beta S_{ave} = \beta/\tau_{ave}$. This is a direct consequence of equation 13.2. Since only a finite number of bits can be in transit in a system of finite area, the only way to maintain constant area per element is by keeping the number of bits in transit per element constant. The condition $B = \beta/\tau_{ave}$ enables precisely this, since $B\tau_{ave}$ is the average number of bits in transit per connection. Also note that in this

case, all physical connections are to nearest neighbors and the interelement spacing can be kept constant so that any interconnection medium results in the same performance and cost within a constant factor. However, if B is not decreased $\propto 1/\tau_{ave}$, the interelement spacing will unavoidably increase so that optical and superconducting interconnections will be preferred over normal conductors after a certain point.

We finally note that in situations where heat removal tends to be the factor limiting system size, larger values of ν are preferred, for reasons similar to those discussed in section 3.6.

13.3 The multiplexed global interconnection architecture

We now turn our attention to another architecture, illustrated in figure 13.2. The N elements among which connections are to be established are partitioned into N/N_1 'modules' of N_1 elements each. All connections between elements in one particular module to another particular module are bundled together and multiplexed into the smallest possible number of physical channels. The relatively short connections between elements in the same module are made directly and would probably be implemented with conductive wiring, because of the greater density they offer over short distances.

The larger the value of N_1 , the larger the number of connections between each module pair, so that a greater number of independent signal paths may be bundled (overlapped) and multiplexed together, resulting in a reduction of the area consumed by global communication channels. On the other hand, increasing N_1 increases the area required by the internal connections. Thus there is an optimal value of N_1 resulting in minimum system area.

The multiplexed global interconnection architecture is not very useful for connection patterns exhibiting a great degree of locality. In such systems there will not be enough connections between distant module pairs to saturate the capacity $1/T_r$ of a single physical channel. It may be useful, for instance, for the implementation of fine




Figure 13.2: *The multiplexed global interconnection architecture with $N/N_1 = 4$. Connections internal to a module (not shown) are made directly, probably with conductive wiring. A connection to a destination in another module is first wired to a common locality with other connections destined to the same target module and multiplexed together. Demultiplexing takes place at the destination module, followed by wiring to the individual destinations. Thus 2 node delays are involved for global connections.*

grain parallel random access machines [63] or connectionist systems.

As a more specific example, this architecture can be used to implement a complete graph on N nodes with $BT_r \ll 1$. This might roughly model a multiprocessor interconnection scenario where each element wants to be able to talk to every other, but only at a relatively low long term average data rate.

13.4 The multiplexed fat tree architecture

The fat tree architecture, illustrated in figure 13.3, was first advocated by Leiserson [107] in a multiprocessor interconnection context. We define the fat tree to have $[kN'^p] \simeq kN'^p$ connections emanating from subtrees containing N' elements, consistent with Rent's rule. This definition is somewhat different than that originally given by Leiserson.

For concreteness, let us assume that waveguides of effective line to line spacing of $f\lambda$ are used. Assuming that the area required for routing functions is not the limiting factor, the linear extent $\mathcal{L}(N')$ of a subtree containing N' elements can be seen to




Figure 13.3: *The multiplexed fat tree architecture.* The elements to be connected are located at the leaves, and the internal nodes provide routing functions. Each connection is established in several hops, $2\log_2 N$ in the worst case. The number of connections emanating from the subtrees increase as we go up the tree. The overlapping portions of the connections are multiplexed into the smallest possible number of physical channels.

satisfy the recursion

$$\mathcal{L}(N') = 2\mathcal{L}(N'/4) + \max(kN'^p BT_r, 1)f\lambda \quad (13.5)$$

since the kN'^p connections emanating from a subtree of N' elements can be multiplexed into $\max(kN'^p BT_r, 1)$ physical channels. $\mathcal{L}(1)$ corresponds to the linear extent of a single element d_d . From this recursion we may show that $\mathcal{L}(N)$ approximately satisfies

$$\max(kN^p(BT_r)f\lambda, N^{\frac{1}{2}}d_d) \leq \mathcal{L}(N) \leq \max(kN^p(\log_2 N^{\frac{1}{2}})(BT_r)f\lambda, N^{\frac{1}{2}}d_d). \quad (13.6)$$

The second term $N^{1/2}d_d$ is unavoidable for any 2 dimensional layout. The first term corresponds to the communication area and is what we are interested in. Upon comparison with equation 13.3, we observe that the multiplexed fat tree allows the smallest possible system size to be approached within a logarithmic factor. (Of course, if BT_r is not small enough to satisfy $BT_r < \kappa/\log_4 N$, the use of a fat tree may not prove advantageous.) What essentially happens is that the total communication area is dominated by the longer higher level connections, in the same sense that a geometric series is dominated by its leading terms. We succeed in multiplexing these to the greatest possible extent so that we can reduce the layout area near to that predicted by equation 13.3. The fat tree architecture allows a greater number of signals to be multiplexed than the multiplexed grid architecture at a cost of fewer node delays.

Once again the price paid is the area and delays associated with routing functions. It would probably be preferable to partition the N elements into N/N_1 modules with direct internal connections (implemented with conductive wires) and then use a fat tree organization to connect the N/N_1 modules. This would enable reduction of the number of device delays incurred and the routing circuitry. Larger values of N_1 would result in fewer node delays. Smaller values of N_1 would enable multiplexing at deeper levels of the system. A detailed simulation would reveal the optimal value of N_1 .

13.5 Conclusion

We have discussed the importance of organizing information flow in a manner enabling maximum multiplexing of independent signal paths, leading to a reduction in the number of physical interconnections, which results in smaller communication area and propagation delays. Among the architectures discussed, the fat tree is near optimal in this respect.

The latter two of the presented architectures provide a natural environment for the joint use of optical and conducting interconnections so as to bring out the best in both and may prove more promising than simple replacement of individual long wires with optics. Optical interconnections are used to provide high density/bandwidth multiplexed information transfer over long distances. Submicron scaled normal conductors are used to provide communication at a density unachievable with optics over shorter distances. This is also consistent with the energetic properties of the interconnection media, as discussed in previous chapters. Optical interconnections consume less energy per transmitted bit over longer distances compared to normal conductors.

Both the multiplexed global interconnection architecture and the fat tree architecture are especially suited for high density (i.e. f close to unity) free space optical implementations because of the regular pattern of interconnections.

Needless to say, a multitude of issues must be considered in contemplating the construction of a high performance computing system. We have focused our attention on the limitations imposed by the area consuming long distance connections and discussed how these limitations can be alleviated by exploiting the high bandwidth potential offered by optical and superconducting interconnections.

Detailed quantitative analysis and simulation of these architectures are beyond the scope of this thesis.

13.6 Related work

The problem of finding the optimal dimension ν of a multidimensional mesh organization was discussed by Dally in the context of a message passing multiprocessor [30],

based on rather different models and considerations. He shows that low dimensional meshes exhibit minimum first-to-last bit delay for L bit messages. More akin to our considerations is the work of Hartmann and Ullman, which we have already discussed.

Chapter 14

Comparison of Local and Global Methods for the Simulation of Physical Phenomena

In most of this work we took the communicative connection graph as our starting point. Ideally, one would take the problem to be solved as a starting point and then seek the optimal physical construct to solve that problem as quickly as possible. Since this is a very difficult task in general, here we will concentrate on a particular problem, that of simulating diffusion phenomena, so as to illustrate certain considerations. We assume that an initial distribution of particles slowly diffuses throughout a given array of cells. Although diffusion phenomena unfold via local interactions, given enough time, the initial condition at any cell may potentially influence the state of any other. Thus, the solution of this problem will eventually involve the global transfer of information across the extent of the system. We consider 3 methods of solving this problem. The first is isomorphic simulation on a locally connected grid, where our processors mimic the physical process of diffusion via local averaging. The second is to formulate the state of the system after a certain number of time steps as a convolution and to perform this convolution on again a locally connected grid, with what is known as a systolic method. The third is to evaluate this convolution using Fourier transform methods on a globally connected array of processors.

14.1 Locality, globality, physics and computation

Once the essential features of a physical phenomenon have been captured in a successful abstraction and distilled down to a set of equations, mathematical methods may be used to predict the outcome of previously unobserved instances of that phenomenon. Since analytic (symbolic) methods have limited applicability, we must usually resort to numerical methods to solve these equations. Numerical calculations are predominantly carried out with the aid of artificial computing machines.

In general, the state of the system we observe may be characterized by several quantities which are functions of the spatial and temporal coordinates (i.e. 'fields'). The many ways we can solve the equations relating these quantities form a broad spectrum, of which we will concentrate on two extremes:

1. Methods which involve only *local* operations in temporal and/or spatial coordinate space. The relaxation method for solving thermal boundary value problems is a simple example of a local method. Such methods are often *isomorphic* to the physical process they represent, in the sense that the calculation mimics the actual physical process at a relatively primitive level. Of course, physical phenomena themselves unfold in time through local interactions, since no influence may propagate faster than the finite speed of light (there is no 'action at a distance').
2. Methods involving *global* operations. Fourier spectral methods are the most widely used among such methods, because complex exponential functions are eigensolutions of linear equations [18]. The Fourier transform operation is itself global in the sense that the value of the Fourier transform of a function at any spectral point depends on the value of the function at every coordinate point. If we wait long enough, the state of any part of a physical system may potentially influence the state of any other part. Thus although physical processes unfold via local interactions, the field quantities at a given point may in general depend on those at any other point. Notice that spatial globality always comes hand in hand with temporal globality.

Our purpose is to find the optimal physical construct to solve a certain problem. In general, our figure of merit may involve measures of time, space and energy consumption. We will concentrate on solving the problem in the shortest possible time T . Say we are given one local and one global algorithm to solve this problem. They are mathematically equivalent in the sense that they will produce the same result. The number of time steps (worst case or average over all possible initial states) required for the completion of the computation can then be determined. By their very design, it is usually the case that global methods require fewer time steps. However, this algorithmic comparison has little meaning, since the physical duration of a time step may be different when we actually build the machines that will execute these algorithms.

Just as the natural phenomena they are used to predict, computing systems must also obey the laws of physics. The need for a physical theory of computation was stressed by Hillis [69] [70]. The mathematical theory of algorithms is meaningful to the extent that the underlying model computing system can be realized. Perhaps one day we will have a theory of computation and cybernetics which is more physics-like.

Towards this end, it is necessary to characterize the information flow required for the solution of a problem on a distributed computing system. This is in general an intimidating task. There have been attempts to analyze the amount of information that must be exchanged between two parties in order to compute a certain function using combinatoric methods [176] [130]. A general extension to many parties which also takes into account issues such as the communication delays among the elements, locality etc. seems exceedingly difficult. For this reason we will consider the simulation of a simple physical process to illustrate certain principles. We will compare the time it takes to solve the diffusion equation using local and global methods. Though desirable, a more general treatment is beyond our scope.

14.2 Implications to future nanoelectronic computing systems

The direct implementation of global methods in parallel hardware requires a highly connected physical connection graph among the participating processing elements. Due to the space that must be allocated for communication, this results in large system size and propagation delays. Based on this and related considerations, Hartmann and Ullman [63] and Dally [30] have argued (in the context of a general purpose message passing parallel computer) that it is more beneficial to simulate such communication networks with a low dimensional mesh. Likewise, based on the intimidating growth of wiring complexity of highly connected systems, Frazier has argued that it would be beneficial to implement future nanoelectronic systems based on quantum-coupled grid-connected cellular automata [48].

If it is indeed the case that in the limit of ultrafast devices and very large numbers of elements, planar (or 3 dimensional) mesh architectures offer better performance than highly connected ones (even with the optimal choice of interconnection medium), then we can conclude that in this limit the choice of interconnection medium (normal conductors, superconductors, optics) will be of little importance. We do not benefit from global communication because we can no longer assume constant delay along all connections. Ultimately, the transfer of information is limited by the speed of light and does not depend on how many hops the distance of travel is broken down to. A shift register begins to resemble a transmission line. A grid connected cellular automata can simulate any portion of the universe, and in particular any given computing system. However, we are presently very far from the device speeds and system sizes needed to reach this limit.

More importantly, we should note that the above arguments do not take into account the effects of heat removal. Although there is general consensus that non-dissipative computing does not contradict the laws of physics, we believe that for quite a time most digital computing systems will make use of dissipative elements. Assuming constant power dissipation per element, heat removal implies that the linear extent of a system of N elements must grow as $\propto N^{1/2}$. This is the worst case growth

rate of wireability imposed lower bounds on system linear extent for bounded degree graphs laid out in 3 dimensions. Thus, since system size and propagation delays are set by heat removal, we might beneficially employ highly connected approaches without further increase in system size. On the other hand, for some applications, the amount of information emitted into the system (and thus the power dissipated) by each element may decrease with increasing N , since the computational processes become bottlenecked by increasing propagation delays, resulting in a heat removal imposed bound on linear extent weaker than $\propto N^{1/2}$.

To us it is not yet clear whether and when highly connected approaches (such as neural networks) or locally connected approaches (such as cellular automata) or something in between is to be preferred and how this is related to the problem we wish to solve. We will merely try to illustrate certain considerations via example.

Extensions of such considerations will help us understand how future nanoelectronic systems should be contemplated and how this is related to the application at hand. Whether locally or globally connected systems are to be preferred will also determine the role of optical communication in these future systems.

14.3 Solution of a problem of high information content

It is well known that many problems such as sorting, convolution, discrete Fourier transforms etc. have an *information content* proportional to N , where N is the problem size [7] [164]. Essentially, the information content is the amount of information that must pass through an imaginary boundary dividing the system into two roughly equal parts before the problem can be solved.

For concreteness let us concentrate on a regular e dimensional cartesian array of N very small processors, $N^{1/e}$ along a side. Each processor contains an L bit precision number. Let it be necessary for NL bits of information to pass through the imaginary boundary mentioned above and there be χH independent physical channels passing through this boundary. Thus the total NL bits must form trains of $(NL)/(\chi H)$

serial bits in passing through this boundary. If the minimum bit repetition interval along each physical channel is T_r , this will take $(NL)T_r/(\chi H)$ time. If the extent of a single physical channel is denoted by W , the linear extent of the system is then $(\chi H)^{1/(e-1)}W$. Letting c denote the propagation velocity, the total computation time $\mathcal{T}$ may be written as a sum of the propagation and serial contributions

$$\mathcal{T} = (\chi H)^{\frac{1}{e-1}} \left(\frac{W}{c} \right) + \frac{NLT_r}{\chi H}. \quad (14.1)$$

The value of χH minimizing the above is found to be

$$\chi H = \left(\frac{NLT_r c(e-1)}{W} \right)^{\frac{e-1}{e}} \propto N^{\frac{e-1}{e}}. \quad (14.2)$$

With this value of χH we find

$$\mathcal{T} \simeq (NLT_r)^{\frac{1}{e}} \left(\frac{W}{c} \right)^{\frac{e-1}{e}} \propto N^{\frac{1}{e}}. \quad (14.3)$$

The above expression for $\mathcal{T}$ is a lower bound. There may be many additional bounds that must also be satisfied. In particular, the information may need to propagate through several nodes, suffering additional delays.

One way of solving such problems is to use a highly connected graph, such as the butterfly graph [164]. Such a graph would add only a few node delays on top of the value of $\mathcal{T}$ discussed above. However, $\chi H \propto N$ for such a graph, which is inconsistent with the optimal value of $\chi H \propto N^{(e-1)/e}$.

A simple mesh architecture can be used to realize the optimal value of χH . A constant number of channels will be used to transfer information among neighboring processors. Information will have to traverse $\propto N^{1/e}$ nodes in the worst case, leading to a similar growth rate of the device contribution to total computation time as given by the above expression for $\mathcal{T}$. Thus the propagation, device and serial contributions to the delay are all balanced in this architecture. Although this does not demonstrate that the problem can be actually solved in $\propto N^{1/e}$ time in general, it serves as a general indicator of the well balanced nature of this architecture.

The above simplistic derivation, which is closely related to the arguments of Hartmann and Ullman [63], suggests that when heat removal considerations are not taken

into account, local methods may be superior to global methods. We will take a closer look at each method in the context of solving a particular physical problem which is described in the next section.

14.4 Quantum diffusion as a prototype physical problem

Consider a regular cartesian e dimensional array of $N \gg 1$ cells with $N^{1/e}$ cells along each side. We will speak of N as the problem size. Without loss of generality, we assume $N^{1/e}$ to be an integer. Initially, there is a certain number $f_0[\bar{i}]$ of bosonic particles in each cell, where $\bar{i}$ is a vector of e integral indices which range from 0 to $N^{1/e} - 1$. For simplicity we assume that the array of cells and the initial distribution of particles is replicated periodically through all space (toroidal boundary conditions). Thus indice values outside the interval $[0, N^{1/e} - 1]$ may be interpreted modulo base $N^{1/e}$.

At a certain average rate of say once every $1 \mu\text{sec}$, each particle has an equal probability of jumping into any one of its 2^e nearest diagonal neighbors. We will assume that the number of particles per cell $\mathcal{M}/N$ is very large so that we may ignore the probabilistic aspects of the problem. Thus, we formulate the problem of determining the number of particles in each cell at time t as follows:

$$f_n[\bar{i}] = \frac{1}{2^e} \sum_{\bar{j} \in A_{\bar{i}}} f_{n-1}[\bar{j}] \quad (14.4)$$

where $A_{\bar{i}}$ denotes the set of 2^e cells with *all* indices differing from cell $\bar{i}$ by unity and $n = t/(1 \mu\text{sec})$. The total number of particles is of course conserved. Thus, given $f_0[\bar{i}]$, we can calculate $f_n[\bar{i}]$ recursively. This process may be expressed as

$$f_n[\bar{i}] = h_1[\bar{i}] * f_{n-1}[\bar{i}] \quad (14.5)$$

$$f_n[\bar{i}] = h_n[\bar{i}] * f_0[\bar{i}] \quad (14.6)$$

where $*$ denotes the e dimensional discrete convolution operator. $h_n[\bar{i}]$ is the n -fold self convolution of $h_1[\bar{i}]$. The value of $h_1[\bar{i}]$ is $1/2^e$ at the 2^e nearest diagonal neighbors of the origin and zero elsewhere. For instance for $e = 1$, $h_1[i] =$

$(\dots, 0, 0.5, 0, 0.5, 0, \dots)$. Notice that whereas convolving $f_{n-1}[\bar{i}]$ with $h_1[\bar{i}]$ requires only local communication, convolving $f_0[\bar{j}]$ with $h_n[\bar{i}]$ for large n requires global communication.

Multidimensional impulse responses can be written as a product of the unidimensional impulse responses. For instance, $h_n[i, j] = h_n[i]h_n[j]$, where $h_n[i]$ and $h_n[j]$ should be interpreted as two dimensional functions while taking their product. (It is also interesting to note that these functions satisfy $h_n[i, j] = h_n[i]\delta[j] * h_n[j]\delta[i]$.)

Our purpose is to calculate the state of the system at a particular final time n_f . As a special case, we will be interested in the steady state solution. Of course, in our simple example, if the total number of particles $\mathcal{M}$ is known to us beforehand, the steady state solution is also known to be a uniform distribution of $\mathcal{M}/N$ particles over all cells. However, we must not forget that this information is not initially available on the processors, which are only aware of the number of particles they contain.

In the following sections, we will consider several methods of constructing a machine that will calculate the state of the system at time n_f .

Before continuing however, we should clarify a common misconception. The parallel implementation of finite element calculations on an array of processors is often noted as an example of a case which requires only local communication among the processors [107]. However, since the state of the system at any point can eventually influence that at any other point, there are cases when it is beneficial for a processor to 'look ahead' beyond its nearest neighbors. A boundary value problem requires global communication since the resultant field distribution at any point depends on the values of the boundary at every point. Although it may depend more weakly on the value of more distant points, this doesn't weaken our argument. Consider that the voltage value at the boundary most distant from a given point is very large compared to anywhere else. Clearly this value will lead to a much different result than when it is a small value. Thus this information must somehow be conveyed to the other end.

Of course, there are examples of applications which are truly local. For instance, the problem of independently adding a large number of pairs of numbers requires no communication at all.

14.5 Solution by isomorphic simulation on a locally connected array

We may compute the state of the system at any time step by using an array of simple processing elements, one for each cell, arrayed in identical fashion as the array of cells. Each processing element is connected to its nearest diagonal neighbors via $\chi \geq 1$ communication channels each of cross sectional area (or width) W^{e-1} . Let the minimum pulse repetition interval and propagation velocity along these channels be denoted by T_r and c respectively. Let each processor be a small cube (or square) of linear dimension at least d_d (it may have to be larger so that it can accommodate its 'pinouts'). Each processor is capable of storing an L bit precision number corresponding to the number of particles in the cell to which it corresponds and can update this value by averaging the values stored in its 2^e neighbors. Let this update take time τ_d .

The interelement spacing is given by $d = \max(d_d, \chi^{1/(e-1)}W, d_Q)$ where d_Q is the interelement spacing required by heat removal considerations. In general, each iteration will take¹ $T_{it} = \max(\tau_d, d/c, LT_r/\chi, T_r)$ time. For simplicity, let us assume that τ_d is defined inclusive of d_d/c and T_r , and that W and LT_r are small enough that the expression for T_{it} reduces to $T_{it} = \max(\tau_d, d_Q/c)$ with appropriate choice of χ .

Let us now calculate d_Q . Let E_d denote the energy dissipation associated with each update on each processor, inclusive of the energy LE involved in transmitting L bits at a cost of E each. As usual, Q will denote the amount of power we can remove per unit cross section of our system. The total power dissipation is NE_d/T_{it} so that

$$Q\mathcal{L}_Q^2 = Q(N^{\frac{1}{e}}d_Q)^2 \geq \frac{NE_d}{T_{it}} \quad (14.7)$$

where $\mathcal{L}_Q$ is the system linear extent required by heat removal considerations. Solving for d_Q and substituting in $T_{it} = \max(\tau_d, d_Q/c)$ we obtain

$$T_{it} = \max\left(\tau_d, \tau_Q N^{\frac{1}{3} - \frac{2}{3e}}\right) \quad (14.8)$$

where $\tau_Q = (E_d/Qc^2)^{1/3}$. Now, the time it takes to compute the state of the system at time step n may be expressed as $\mathcal{T} = nT_{it}$.

¹As usual, we will not make the distinction between the sum and maximum of a small number of positive numbers and likewise ignore numerical factors of the order of unity.

With increasing n , the state of the system will relax towards its steady state. Of course, in general, the system will never *exactly* reach steady state (except in special cases where the solution falls in precisely to the steady state and stays there due to the discrete nature of our model). We would not expect the system to reach steady state before $\sim N^{1/e}$ time steps, since this many steps are necessary for influences to propagate across the extent of the system. Exactly how long we must wait also depends on the error ε we are willing to tolerate. One way of determining the number of time steps necessary for given ε would be to carry out simulations for a variety of initial conditions. Instead, we will estimate the number of time steps n_ε it takes for an impulse of strength $\mathcal{M}$ to diffuse into a steady state of $\mathcal{M}/N$ particles per cell with fractional error ε .

For simplicity we consider the 1 dimensional case (i.e. $e = 1$). An impulse of unit strength diffuses in the following Pascal's triangle-like manner:

$$\begin{array}{rcl}
 n = 0 & 0.000 & 0.000 & 0.000 & 0.000 & 1.000 & 0.000 & 0.000 & 0.000 & 0.000 \\
 n = 1 & 0.000 & 0.000 & 0.000 & 0.500 & 0.000 & 0.500 & 0.000 & 0.000 & 0.000 \\
 n = 2 & 0.000 & 0.000 & 0.250 & 0.000 & 0.500 & 0.000 & 0.250 & 0.000 & 0.000 \\
 n = 3 & 0.000 & 0.125 & 0.000 & 0.375 & 0.000 & 0.375 & 0.000 & 0.125 & 0.000
 \end{array} \tag{14.9}$$

and so on. Thus for an impulse of strength $\mathcal{M}$ located at the origin ($i = 0$) at $n = 0$, the number of particles in location i after n steps is

$$f_n[i] = \frac{\mathcal{M}}{2^n} \text{combination} \left(n, \frac{i+n}{2} \right) \quad -n \leq i \leq n \tag{14.10}$$

for even $(n+i)$ and 0 for odd $(n+i)$. Since N is large, n_ε will also be large. Thus, using the DeMoivre-Laplace theorem [138] the above expression may be approximated as

$$f_n[i] = \frac{\mathcal{M}}{\sqrt{2\pi n}} \exp \left(-\frac{i^2}{2n} \right) \tag{14.11}$$

where we have included an additional factor of $1/2$ so as to ensure proper normalization over all values of i , rather than just odd or even values depending on whether n is even or odd. Now, remembering that we are employing cylindrical boundary conditions, the number of particles at the midpoint between the origins at $i = 0$ and

$i = N$ (which will be the latest to reach the steady state of $\mathcal{M}/N$ particles) is given by

$$\sum_{j=-\infty}^{\infty} f_n[N/2 + jN]. \quad (14.12)$$

Using some algebra involving Fourier transform techniques [18], this summation may be expressed in the equivalent form

$$\frac{\mathcal{M}}{N} \left[1 + 2 \sum_{j=1}^{\infty} (-1)^j \exp \left(- \left(\frac{\pi \sqrt{2n}}{N} \right)^2 j^2 \right) \right]. \quad (14.13)$$

The second term in square brackets is simply the fractional error we do not want to be greater than ε . Since we are confronted with an alternating series, we can ensure the error to be less than ε by ensuring that

$$n \geq N^2 \left(\frac{\ln \frac{2}{\varepsilon}}{2\pi^2} \right) \simeq N^2 \left(0.035 + 0.117 \log_{10} \frac{1}{\varepsilon} \right). \quad (14.14)$$

Because of the weak dependence on ε , we are justified in writing $n_\varepsilon \sim N^2$. In other words, for almost all practical values of ε , taking $n \sim N^2$ is as good as $n = \infty$. We could have guessed this result beforehand. The root mean square deviation of a random walk in any dimensional space is proportional to $n^{1/2}$. Thus we might consider that we have reached steady state when $n^{1/2}$ is comparable to the linear extent N of our system. This result easily generalizes to e dimensions for which the linear extent of the system is $N^{1/e}$. Thus in e dimensions

$$n_\varepsilon = N^{\frac{2}{e}}. \quad (14.15)$$

The total time it takes to find the state of the system at time step n_f and at steady state using isomorphic grid simulation are given in table 14.1.

Conventional VLSI complexity theory would predict the same growth rate of computation time, apart from the fact that heat removal is often not considered. These calculations would take $\propto n_f N$ time steps on a single processor Von Neumann architecture.

14.6 Solution by systolic grid convolution on a locally connected array

We now discuss a method of directly performing the convolution of equation 14.6 on a nearest neighbor connected array of processors, like the one used for isomorphic simulation. For instance, consider the two dimensional case. The convolution is written out explicitly as

$$f_n[i, j] = \sum_l \sum_k f_0[k, l] h_n[i - k, j - l] \quad (14.16)$$

$$f_n[i, j] = \sum_l h_n[j - l] \sum_k f_0[k, l] h_n[i - k] \quad (14.17)$$

since $h_n[i, j]$ is separable. This amounts to first computing $g_n[i, l] = \sum_k f_0[k, l] h_n[i - k]$ for every l systolically in the i (also k) direction and then computing $\sum_l g_n[i, l] h_n[j - l]$ for every i in the j (also l) direction. The systolic convolution can be performed by rotating the values of $f_0[k, l]$ (or $g_n[i, l]$) and accumulating the properly weighted sums [98] [96]. We are assuming that the values of our impulse response can be easily generated in each processor with little hardware, probably by evaluating its analytic expression. The 2 dimensional convolution takes $2N^{1/2}\mathcal{T}_{it}$ time where $\mathcal{T}_{it}$ is the same as that during isomorphic simulation. In e dimensions, this takes $eN^{1/e} \simeq N^{1/e}$ time steps regardless of n_f .

Once again, conventional VLSI complexity theory would yield the same predictions for the computation time, apart from heat removal considerations. A single processor machine would take $N^{1+1/e}$ time steps.

14.7 Solution by Fourier techniques on a globally connected array

Equation 14.6 may be evaluated conveniently using Fourier domain techniques, since convolution in coordinate space corresponds to multiplication in Fourier space. We assume the Fourier transform of the impulse response is pretabulated and piped in in proper synchronicity. After multiplication with the Fourier transform of the initial




Figure 14.1: *Decomposition of an $N = 16$ point FFT.*

distribution, we inverse transform to get the desired result. Thus the total time of computation is about equal to the time it takes to evaluate the transform of the input function $f_0[i, j]$, which is in general not separable. (If it were separable, the problem would reduce to a unidimensional problem.)

The relationship between N point 1 dimensional Fourier transforms and $N^{1/2} \times N^{1/2}$ point 2 dimensional Fourier transforms is well known in the context of raster scan-folded spectrum techniques [22], where the two dimensional Fourier transforming capability of an optical lens is exploited for high time-bandwidth product spectral analysis of one dimensional analog signals. The FFT decomposition shown in figure 14.1 [97] may be interpreted either as a 16 point 1 dimensional transform or a 4×4 point 2 dimensional transform. The reader is referred to [55] [129] and [22] for analytic discussions. The essential idea is that regardless of its dimensionality, the value of the transform at each spectral point depends on the value of the original function at every coordinate point, so that both problems require the same pattern of information flow. A construct for solving either problem can be used for the other with often only trivial modification (such as the inclusion of additional phase shifts along a few paths, as discussed in [55]). This discussion generalizes to $N^{1/3} \times N^{1/3} \times N^{1/3}$ point 3 dimensional transforms. Thus we may speak of the complexity of evaluating an N point Fourier transform without reference to its dimensionality.

An N point Fourier transform can be computed effectively in both 2 and 3 dimensions by using a 1 or 2 dimensional lens respectively [56]. In 3 dimensions, one may situate the elements on a $N^{1/2} \times N^{1/2}$ array and use the third dimension for communication. Such a setup implies a linear extent and propagation delay of $\simeq N^{1/2}\lambda$ and $\simeq N^{1/2}\lambda/c$ respectively, assuming an $f^\# \sim 1$ optical system. Here $W = \lambda$ is the wavelength of light. In 2 dimensions, the same setup must be squashed onto the plane. The N independent spatial channels now imply a linear extent $\simeq N\lambda$, since they have only one dimension to pass through. Those familiar with the layout of the butterfly graph [164] on which the FFT is performed will immediately recognize that these results are analogous to results stating that the butterfly lays out in $\propto N^{1/2}$ linear extent in 3 dimensions and $\propto N$ linear extent in 2 dimensions, with identical growth rate of longest wire length and propagation delay. (The VLSI implementation of the butterfly graph also results in $\log_2 N$ node delays in addition to the propagation delay. This is avoided in the optical implementation.)

In conclusion, the evaluation of an N point Fourier transform on a globally connected array of processors in the manner described requires propagation delay $\simeq N^{1/(e-1)}\tau_o$ where $\tau_o = \lambda/c$.

Notice that we are (by choice) concentrating on the evaluation of the Fourier transform on a globally connected array in a particular manner. This does not exclude the existence of other asymptotically superior methods. For instance, an N point DFT can be evaluated on a nearest neighbor connected mesh in $\propto N^{1/2}$ time [158]. This method of solution would be equivalent to the systolic method discussed in the previous section.

Heat removal requires that the system linear extent be at least $(NE_d/TQ)^{1/2}$, assuming the total energy dissipation NE_d is spread over the total time of computation $\mathcal{T}$. (It is easy to show that any other choice is suboptimal.) Thus $\mathcal{T}$ is given by

$$\mathcal{T} = \max(\tau_d, N^{\frac{1}{e-1}}\tau_o, N^{\frac{1}{3}}\tau_Q) \quad (14.18)$$

where τ_Q is as defined before. Heat removal will have less and less importance as N increases. (If the total energy dissipation cannot be spread over the total time of computation, but only over a constant interval of time, the last term will become

	I (local)	I_∞ (local)	S (local)	F (global)
A	$n_f \tau_d$	$N^{\frac{2}{e}} \tau_d$	$N^{\frac{1}{e}} \tau_d$	$(\tau_d, N^{\frac{1}{e-1}} \tau_o)$
B	$n_f(\tau_d, \tau_Q)$	$N(\tau_d, \tau_Q)$	$N^{\frac{1}{2}}(\tau_d, \tau_Q)$	$(\tau_d, N\tau_o, N^{\frac{1}{3}} \tau_Q)$
C	$n_f(\tau_d, N^{\frac{1}{9}} \tau_Q)$	$(N^{\frac{2}{3}} \tau_d, N^{\frac{7}{9}} \tau_Q)$	$(N^{\frac{1}{3}} \tau_d, N^{\frac{4}{9}} \tau_Q)$	$(\tau_d, N^{\frac{1}{2}} \tau_o, N^{\frac{1}{3}} \tau_Q)$

Table 14.1: *Comparison of total computation time $\mathcal{T}$ with the various methods for calculating the state of the system at time step n_f . I: isomorphic simulation, S: systolic convolution and F: Fourier transforming. Line A gives $\mathcal{T}$ in e dimensions when heat removal is ignored. Line B is for $e = 2$ dimensions and line C for $e = 3$ dimensions. The steady state for the isomorphic simulation (I_∞) is obtained by setting $n_f = N^{2/e}$. The other methods calculate any final time step n_f , as well as the steady state, equally quickly. $\tau_o = \lambda/c$, $\tau_Q = (E_d/Qc^2)^{1/3}$. The notation (x, y) is short for $\max(x, y)$.*

$\propto N^{1/2}$, requiring minor modification of our results when $e = 3$.)

The results are again summarized in the table.

Conventional VLSI complexity theory predicts a $\propto \log N$ growth rate of delay for parallel implementation of the FFT, since propagation delays are ignored. The single processor implementation takes about $N \log_2 N \sim N$ time steps.

14.8 Comparison

The results are summarized in table 14.1. Let us ignore heat removal for the moment ($\tau_Q = 0$). First consider the limit where device delays dominate (τ_d is large or N is small). The global Fourier method is clearly superior in this case unless n_f is very small. Which of isomorphic simulation and systolic convolution is to be preferred depends on the values of n_f and N . When n_f is small and N is large, the first method is to be preferred. When n_f is large however, systolic convolution is preferred. In particular, consider the steady state which takes $N^{2/e} \tau_d$ time with isomorphic simulation in contrast to $N^{1/e} \tau_d$ time with systolic convolution.

Now let us assume ultrafast devices and compare the growth rate of computation time as a function of N and n_f . Local systolic convolution exhibits a growth rate better than global Fourier transform methods by a factor $N^{1/(e-1)}$. This result was

anticipated in section 14.3 where we discussed the solution of problems of high information content. In practice $\tau_d \gg \tau_o$. Thus, systolic convolution will be preferred over Fourier methods when $N > (\tau_d/\tau_o)^{e(e-1)}$. For instance, with $\tau_d = 1$ psec and $\tau_o = W/c = \lambda/c = 1$ fsec, the condition is approximately $N > 10^6$ in 2 dimensions and $N > 10^{18}$ in 3 dimensions. We are not the first to speak of such a large number of very simple processors, Hillis has speculated about what might happen when we have a mole ($\sim 10^{23}$) of processors [70].

Now let us consider the effects of heat removal. This has little or no effect on our conclusions in 2 dimensions. In 3 dimensions, the leading terms will be $N^{4/9}\tau_Q$ for systolic convolution versus $N^{1/2}\tau_o$ for Fourier transforming. Taking $E_d = 10$ fJ and $Q = 10$ KW/cm² so that $\tau_Q \simeq 0.1$ psec, we find that local systolic grid convolution is preferred over the global Fourier method only when $N > 10^{36}$, an unreasonably large number by any standard.

To sum up, isomorphic simulation may be preferred when n_f is small, or when we want to track the whole evolution of the system up to n_f , rather than just its final state. (The problem of piping out this data from the system remains unsolved however.) Otherwise, especially when we want to calculate the steady state, this method will not be preferred. This is because of the inefficient way in which the transfer of information occurs. The averaging at every step results in loss of information for which communication resources have already been utilized, resulting in inefficient resource utilization (much like exchanging a developed piece in a game of chess [166]). Local systolic convolution may be preferred over global Fourier methods in 2 dimensions for very large values of N . In 3 dimensions, the asymptotic superiority of local methods over global methods is so slight as not to be of any significance so that the latter are to be preferred.

Figure 14.2 illustrates $\mathcal{T}$ as a function of N for the three methods.

14.9 Discussion and extensions

Although it may seem that the diffusion process we have considered is a very special example, we stress that diffusion phenomena underly many natural occurrences, of




Figure 14.2: Computation time for the three methods with $n_f = N^{2/e}$ for isomorphic simulation. $\tau_d = 1$ psec, $\tau_o = 1$ fsec, $\tau_Q = 0.1$ psec and $d_d = 1 \mu\text{m}$. The curves are terminated when the system linear extent exceeds 10 m.

which Brownian motion [47], spread of scientific innovation [61] and gossip are a few examples. The diffusion and wave equations are often the starting point of a course in partial differential equations. In their steady state both reduce to Laplace's equation. The wave equation is often associated with coherent, orderly transfer of information, energy or particles; whereas the diffusion equation is associated with random transfer. The diffusion process is redundant in the sense that the same information is retransmitted several times only to be destroyed by the averaging process. We can do better than the isomorphic simulation of diffusion by carrying out the calculations in an orderly and efficient manner, rather than imitating the physical process itself. Similar arguments may be possible against other numerical methods which involve randomness and/or averaging, such as Monte Carlo or relaxation methods.

Convolution is one of the basic operations in signal and image processing. Thus, we would expect extensions of our discussion to have implications in these areas. Of course, some applications require convolution with only a finite 'window' function, and would probably be implemented using local methods. For other applications however, global methods may (or may not) be preferable.

Similar analysis as we have presented can be carried out in different contexts. For instance, it is well known that the same logic function can be implemented with a

fewer number of locally connected elements but with large overall logic depth, or with a larger number of globally connected elements with less logic depth [6]. The optimal implementation will lie somewhere in between.

Such studies may also be used to evaluate the usefulness of neural networks. We would not be surprised if similar conclusions regarding the usefulness of highly connected systems are reached.

We only concentrated on extreme locality and extreme globality. Intermediate approaches are possible and may offer the optimum performance. For instance, given a family of algorithms for solving the diffusion equation on multidimensional meshes of every dimension, we may pick the optimum dimension. Also, we considered only nearest neighbor communication in the grid method. In serial computation often a bounded grid of higher order neighbors is used, leading to faster convergence. In parallel hardware, this might lead to a small amount of dilation, leading to a trade-off. In recent years, a number of (serial) numerical methods involving higher order interpolation polynomials and combinations of spectral and finite element methods have emerged [17], evidence of the fact that the optimum lies somewhere in between the two extremes of global and local methods. There seems much that is unexplored as far as the use of parallel hardware is concerned.

14.10 Conclusion

We discussed various methods of simulating a particular physical process, that of particulate diffusion. We considered isomorphic simulation of the physical process, local systolic convolution and global Fourier transform methods.

Systolic convolution is asymptotically superior to Fourier methods. However, in 3 dimensions the discrepancy in the growth rate of delay is very small because of the effects of heat removal. Thus, Fourier methods result in smaller total time of computation.

We have also emphasized the importance of orderliness in information transfer and removal of redundancy. Isomorphic simulation of the diffusion process was found to be inferior because it failed to satisfy this criteria.

Chapter 15

Summary and Conclusion

15.1 Summary and contributions

We began part I with a discussion of the motivation, philosophy and general framework of this study. We then discussed wireability limitations and presented a connectivity model and method of analysis mostly due to earlier authors.

Then we turned our attention to heat removal limitations. We derived a basic result quantifying the limits to heat removal from 3 dimensional systems, showing that ten kilowatts can be removed from a square centimeter.

We discussed how the physical considerations of this study can be related to algorithmic considerations and how the choice of interconnection medium is related to such considerations.

Finally, we derived what we termed tube models of interconnections for optical, normally conducting, repeatered and superconducting interconnections. These models are simple enough to be incorporated in system level studies of a relatively high degree of abstraction, but account for most major physical limiting mechanisms such as the skin effect, superconducting penetration depth and critical current limitations.

In part II we concentrated on optical systems. We first presented a basic result regarding the limitations of electromagnetic wave propagation in providing communication among an array of points, showing that optically communicating systems are similar to those employing solid wires in the sense that the total volume that must

be allocated for communication is proportional to the total interconnection length in both cases.

We then derived and compared the system sizes for various specific optical interconnection architectures and proposed a new architecture which can approach the best possible system size of any 2 dimensional optical system, and also discussed a 3 dimensional version of this architecture.

We discussed the role of optical frequencies in balancing information density and heat removal imposed bounds in a dissipative computing system. We also derived a relatively fundamental bound between signal delay, number of elements and bit repetition rate for dissipative systems.

In part III, we first extensively investigated the tradeoffs between number of elements, inverse delay, bandwidth and the cost of power and space for each interconnection medium. Not surprisingly, we found that optical and superconducting interconnections have the best asymptotic behavior.

We then extensively discussed hybrid implementations for a variety of layout constraints and physical parameters and determined the conditions under which optical interconnections might be useful. We found that optical interconnections have relatively little to offer if the optical paths are constrained to lie on the plane (such as in a guided wave network). However, if optical paths are allowed to leave the plane, they may enable considerable increase in performance. In any event, the prize in terms of performance is often accompanied by a penalty in terms of power and/or size.

Based on the considerations of this study, we speculated about the conditions under which an all optical computer might be useful. We found that it might be meaningful to consider the construction of an all optical computer when large bit repetition rates are useful despite large propagation delays.

We discussed the importance of organizing information flow in a manner enabling multiplexing of independent signal paths with distinct source and destination localities, and discussed some architectures with this property. We showed that these architectures provide a natural environment for the joint use of optical and normally conducting interconnections so as to bring out the best in both.

Finally, we discussed the simulation of diffusion phenomena on locally and globally connected systems with very large numbers of elements and very fast devices, concluding that globally connected systems are superior, mostly due to heat removal considerations.

15.2 Conclusion

In conclusion, we have developed a physical approach to modeling and analyzing communication limits in computation. We have stressed wireability and heat removal requirements as the major limiting mechanisms and tried to explore the interplay between these considerations and the physical properties of the interconnections, which we characterized by what we termed as tube models.

It is often noted that the academic study of the so-called *systems physics* of computing systems has been neglected [161] [163]. On the one hand, considerable effort is put into the detailed study of individual devices. On the other hand, considerable effort is put into the study of systems aspects of computation, with little reference to actual physical construction. In between, the field of *microelectronic packaging* has mostly remained in the technological domain, perhaps since its interdisciplinary nature makes it difficult to break it down into self-contained academic categories [163]. Perhaps future work in this area will result in the growth of the academic discipline of systems physics.

From one viewpoint, this work can be seen as a study of the nature of computing systems, much like the study of physical phenomena or biological organisms. We have been mostly concerned with analyzing and understanding the nature of the tradeoffs and limitations involved in the construction of computing systems, rather than designing or inventing new systems. From a practical viewpoint, our approach is of utility in identifying major roadblocks to approaching fundamental limits, and determining areas in which it makes sense to invest human effort and resources.

Chapter 16

Appendices

16.1 The effect of scaling for CMOS VLSI circuits

In this appendix we will start from explicit expressions for the switching energy and rise time for CMOS VLSI circuits and discuss the effects of scaling. This will serve as an example of how the device contributions to the switching energy and delay quickly lose importance compared to interconnect contributions.

We will use the scaling rules given by Gardner [54] and used in [40]. Following this latter reference we write

$$E_{VLSI} = C_{total}V^2 = (sC_{OA} + C_{OB} + C_{in} + C_{LA}\ell w + C_{LB}\ell)V^2 \quad (16.1)$$

$$\tau_{VLSI} = R_L C_{LA}\ell^2 + R_L C_{LB}\ell^2/w + 2C_{in}R_L\ell/w + [V/(sI_0)](C_{total}). \quad (16.2)$$

The various terms are summarized in table 16.1 along with how they are scaled. Here ℓ is the length of the interconnection, w its width and s the ratio of the current driving capability of the transistor used to that of a minimum sized one. One can now substitute these in the above equations and obtain the switching energy and rise time as functions of λ :

$$C_{total} = 0.72 s\lambda^2 + 3.1 \lambda^2 + 5.9 \lambda + 0.099 r\ell + 0.061 \ell \quad (16.3)$$

$$E_{VLSI} = C_{total}V^2 \quad (16.4)$$

$$\tau_{VLSI} = 0.020 \frac{\ell^2}{\lambda^2} + 0.012 \frac{\ell^2}{r\lambda^2} + 2.4 \frac{\ell}{r\lambda} + 3800 \frac{VC_{total}}{s} \quad (16.5)$$

Parameter	Symbol	Value	Units
Minimum linewidth	w_{min}	1.5λ	μm
Gate input capacitance	C_{in}	$170 \lambda \epsilon_{ox}$	fF
Output capacitance	C_{OA}	$0.72 \lambda^2$	fF
Contact output capacitance	C_{OB}	$3.1 \lambda^2$	fF
Line capacitance	C_{LA}	$\epsilon_{ox}/(0.35 w_{min})$	fF/ μm^2
Fringing line capacitance	C_{LB}	~ 0.061	fF/ μm
Line resistance	R_L	$\rho/(0.09 \lambda)$	Ω/sq
Power supply voltage	V	$2.9 \lambda^{1/2}$	V
Inverter saturation current	I_0	0.26	mA

Table 16.1: *Integrated circuit process parameters expressed as functions of minimum feature size λ (μm) for first level aluminum lines.* ϵ_{ox} = the permittivity of silicon dioxide = $3.9 \times 8.85 \times 10^{-3}$ fF/ μm , ρ = the resistivity of aluminum = $0.0274 \Omega \mu\text{m}$. (After [40].)

where $r = w/w_{min}$. Capacitance is measured in fF, energy in fJ, time in fsec and all lengths in μm . Inspection of the equation for C_{total} reveals that for decreasing values of λ and a sufficiently high value of ℓ the gate contributions can be neglected leaving only the last two terms,

$$C_{total} = 0.099 r \ell + 0.061 \ell. \quad (16.6)$$

Similarly, in the expression for τ_{VLSI} , the first two terms will dwarf the third with growing ℓ/λ . We can then write

$$\tau_{VLSI} = 0.020 \frac{\ell^2}{\lambda^2} + 0.012 \frac{\ell^2}{r \lambda^2} + 3800 \frac{V C_{total}}{s} \quad (16.7)$$

where C_{total} is now the simplified expression given above. Note that E_{VLSI} depends on λ only through ℓ . On the other hand the expression for τ_{VLSI} depends on ℓ/λ and is otherwise independent of λ .

The ultimate scaling limits of MOS circuits have been extensively discussed [83] [119]. The minimum voltage levels around room temperature are set by the requirement $V \gg kT/q$ and also other factors such as the distribution of threshold voltages during production [83] [152]. Usually ~ 1 V is agreed upon as a reasonable minimum value, although it is possible, especially with decreasing temperatures to further lower

the voltage. According to our scaling rules the value of λ corresponding to this voltage is $0.12 \mu\text{m}$. We will assume these figures to be roughly representative of the ultimate scaling that can be achieved with MOS technology.

With this value of λ we write numerical expressions for the switching energy and delay. If s is kept within reasonable bounds (say $s \leq 200$), one can show that for all interconnections longer than about $20 \mu\text{m}$ the equations can be simplified as follows

$$E_{VLSI} = 0.55 w \ell \left(1 + \frac{0.11}{w} \right) \quad (16.8)$$

$$\tau_{VLSI} = \left(1.4 \ell^2 + 2100 \frac{\ell w}{s} \right) \left(1 + \frac{0.11}{w} \right). \quad (16.9)$$

The rightmost parentheses contain a correction factor for the fringe term. This term is bounded by 1.6 so we can further simplify to obtain

$$E_{VLSI} \simeq 0.88 w \ell \quad (16.10)$$

$$\tau_{VLSI} \simeq 2.2 \ell^2 + 3400 \frac{\ell w}{s}. \quad (16.11)$$

These are equivalent to equation 1.1. When $\ell > 280 r/s$ is satisfied the first term will begin dominating the expression for τ_{VLSI} so that for the longer interconnections driven by strong transistors the square law term will dominate.

The specific value of λ that we have chosen does not have as much significance as the ratio ℓ/λ . If transistor sizes are bounded to reasonable values ($s < 200$) the line contribution to the capacitance exceeds the gate contributions when $\ell/\lambda > 100$. This figure is consistent with breakeven values cited elsewhere [5].

16.2 Extension to fan-out and fan-in (chapter 2)

One can model the equivalent of a fan-out or fan-in situation using only pairwise links. This is depicted in part b. of figure 16.1. The layout minimizing connection length is shown in part a. Notice that the inefficiency in using pairwise links is bounded between 1 (when two target elements are located as in d.) and $1/F$ (when the target elements are as in c.), 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




Figure 16.1: *Extension to fan-out and fan-in.* Part a. shows optimal power splitting which minimizes the total connection length. Part b. illustrates the same connections wired using pairwise links only. Parts c. and d. show the two extreme cases illustrating the bounds on the inefficiency incurred by using only pairwise links.

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 [57] [40].

16.3 Coefficients for the moments of $g(r)$

The coefficients appearing in equation 2.3 are given by

$$\begin{aligned} \zeta_m &= \frac{me}{(m - e/n)(e - e/n + m)} \\ \zeta'_m &= \frac{1}{n} \\ \zeta''_m &= \frac{1}{1 - mn/e}. \end{aligned} \tag{16.12}$$

The major approximation made in the derivation of equation 2.3 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's 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 .

$\int_{r=1}^{r_{max}} g(r) dr = k$ gives the number of connections per element. In chapters 10 and 11, we will be interested in the number of connections of length $r \geq r_x$, given by $\int_{r=r_x}^{r_{max}} g(r) dr$. Using the definition of $g(r)$, this is easily shown to evaluate to $f(r_x) = kr_x^{-e/n}(1 - r_x^e/r_{max}^e)$. With $r_x \simeq N_x^{1/e}$, this becomes $\simeq kN_x^{p-1}(1 - N_x/N)$, which is $\simeq kN_x^{p-1}$ when $N_x \leq N/2$ or so.

$\int_{r=1}^{r_{max}} rg(r) dr$ gives the total interconnection length per element (in grid units). In chapters 10 and 11, we were also interested in the sum of the lengths of interconnections with length $r \leq r_x$, given by $\int_{r=1}^{r_x} rg(r) dr$. Provided p is not very close to unity (the largest value we employ is 0.8), this integral approximately evaluates to $k\kappa N_x^{1/e-1/n}$ when $n > e$, $k\kappa' \ln N$ when $n = e$ and $k\kappa''$ when $n < e$, similar to equation 2.4. In other words, if the upper limit of integration is sufficiently large compared to unity, it does not matter whether it is less than or equal to r_{max} .

In chapter 11, we apply our method of analysis hierarchically, first to groups of N_1 elements, and then to the N/N_1 groups. Thus both of these quantities must be large in order for our approximations to hold. Since the optimal value of N_1 is never very close to unity, our plots are least accurate over the limited range when $N_1/N \sim 1$, i.e. when N_1 breaks away from N .

16.4 Discussion of the validity of neglecting dielectric losses

Throughout our discussion of conducting lines we used the approximation $G = 0$. It is shown in [120] that for materials like fused quartz or polyimide the contribution of dielectric loss to attenuation is indeed negligible for frequencies and lengths of at least

10 GHz and 10-100 cm. Conductor losses will readily dominate dielectric losses in all but very wide lines, which will probably never be beneficial to use because of their low packing density. In general the quantity of interest is the loss tangent $1/(\omega\epsilon\rho_{dielectric})$, which must be small compared to unity for our approximation to hold.

The situation is somewhat different for microstriplines on Si substrates, because the losses are considerably higher than for those materials mentioned above. Detailed analysis of striplines on the silicon system can be found in [64] [147]. In general it is found that there exist three propagation modes. The quasi-TEM mode is what we have analyzed. The substrate skin effect mode is also accounted for in our analysis, if we interpret the substrate to act as the ground plane. However, the slow wave mode can only be analyzed by taking into account the finite conductivity of the substrate. We will not enter this region if either $\omega > 1/\rho_{substrate}\epsilon_{substrate}$ or $\omega > 2\rho_{substrate}/\mu h_{substrate}^2$ is satisfied, where $h_{substrate}$ is the thickness of the substrate [147]. In other words it is only a low frequency and intermediate resistivity regime that cannot be made to fit into our analysis. For a micron wide line, $(\omega/2\pi)\rho_{substrate} > 1$ with $(\omega/2\pi)$ in GHz and $\rho_{substrate}$ in Ω cm, is sufficient to ensure we are in the quasi-TEM regime.

In the case of unterminated lines the above is subject to the condition that the line is not very long, as delayed reflections will cause low frequency components. Again from the same reference, we find that a $\rho_{substrate}/\ell$ ratio of at least 10Ω must be maintained to ensure that we are in the TEM region. For longer lines, we can ensure that we are in the substrate skin effect region by keeping the resistivity sufficiently low. The reader is referred to [147] for details.

16.5 Discussion of the validity of the quasi-TEM approximation

In our analysis of transmission lines, we employed the quasi-TEM approximation, i.e. we assumed low losses so that the longitudinal electric field could be neglected in

comparison with the transversal fields. It was shown by other workers that this approximation yields reasonable results for real microstrip lines [60]. The ratio between the longitudinal electric field value and the transversal field value is [25]

$$\frac{|E_{longitudinal}|}{|E_{transversal}|} = \sqrt{\omega\epsilon\rho} \quad (16.13)$$

from which we see that this assumption will break down at sufficiently high values of the frequency and/or resistivity. For metallic conductors, say aluminum, the quantity inside the square root is around 10^{-9} per GHz so that this approximation is valid for all frequencies of interest for good conductors.

16.6 Signal delay for MOS 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 [5] [6]. 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 [168]

$$T \simeq R_d C_d + (R_d + R\ell)C\ell = R_d(C_d + C\ell) + T_\ell \quad (16.14)$$

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 [158].

Such arguments may not always be possible. However, remember that it is sub-optimal 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 .

16.7 Additional factors of 2 contributing to F


There are several factors of 2 which can be included in equation 6.1 and our major results, such as equation 6.10. Since we have centered our discussions on scalar theory, a factor of 2 may be augmented to account for the two independent polarization states.

The other factors of 2 are best understood in relation to the Nyquist sampling theorem [129] [178] and can be interpreted as a consequence of the double sidedness of the frequency domain representations. Two of these factors corresponding to the two transverse dimensions are already inherent in our results. Two more factors are associated with the longitudinal dimension and the time coordinate, i.e. both odd and even versions of both forward and backward traveling waves can independently exist. Hence, if one utilizes the channels bidirectionally a factor of 2 is added. The other one can be added only if we employ a detection scheme sensitive to the temporal phase of the optical carrier. We included neither with the understanding that they may be easily reintroduced whenever appropriate.

16.8 Calculation of $\chi(P)$ for circular and square apertures

So as to gain some insight regarding the accessible Fourier area χ , we discuss its calculation for the special cases of axially centered circular and square apertures (figure 16.2). The accessible Fourier area at a given point is determined by the image of the aperture stop observed from that point. Assume that this image is at a distance L from the point in question and has a diameter D . Define the half cone angle $\theta = \tan^{-1}(D/2L)$ and $f^\# = L/D$.

For circular apertures, we then obtain the solid angle as a function of the half

Figure 16.2: *Projection of a solid angle.*

cone angle as

$$\Omega = 2\pi(1 - \cos \theta). \quad (16.15)$$

If $\Omega_{\perp}$ is the projection of this solid angle on a plane perpendicular to the axis of the cone at a radius of unity, the accessible Fourier area is then given by

$$\chi = \frac{\Omega_{\perp}}{\lambda^2} = \frac{\pi \sin^2 \theta}{\lambda^2} = \frac{1}{\lambda^2} \Omega \left(1 - \frac{\Omega}{4\pi}\right). \quad (16.16)$$

Similar relations exist for square apertures. If we call the direction of propagation the z direction and use x and y to denote the transverse coordinates, the surfaces bounding the pyramid of allowed wave vectors are defined for $\sigma_z \geq 0$ by the equations

$$|\sigma_x| = (\tan \theta) \sigma_z \quad (16.17)$$

$$|\sigma_y| = (\tan \theta) \sigma_z \quad (16.18)$$

where $(\sigma_x, \sigma_y, \sigma_z)$ denotes the wave vector. To find the accessible Fourier area we need to find the area of the projection of the solid angle defined by these surfaces. The curves enclosing this area can be found by substituting the above in the equation

$$\sigma_x^2 + \sigma_y^2 + \sigma_z^2 = \frac{1}{\lambda^2}. \quad (16.19)$$

We then find that the accessible Fourier area is the area of intersection of the two ellipses shown in figure 16.3. We will not require the exact relation for this area; however we make the following observations:




Figure 16.3: Accessible Fourier area for a square aperture, as given by the intersection of two ellipses. The shaded area gives χ , the accessible Fourier area. This area approaches a square for small values of θ and a circle for large values of θ .

1. For smaller values of Ω and θ the accessible Fourier area can be expressed as

$$\chi \simeq \frac{4 \sin^2 \theta}{\lambda^2}, \quad (16.20)$$

i.e. a square of diameter $2 \sin \theta / \lambda \simeq 1 / f^\# \lambda$. The minimum resolvable dimension becomes $f^\# \lambda$, consistent with what is encountered in analyzing paraxial problems in cartesian coordinates [56] and also the Nyquist sampling theorem (appendix 16.7).

2. As Ω approaches 2π and θ approaches $\pi/2$, the accessible Fourier area approaches a circle of area $\chi = \pi / \lambda^2$. This is a consequence of the spherical nature of the governing physical equation (number 16.19).
3. For all values of $0 \leq \Omega \leq 2\pi$, we have $\Omega_\perp = \lambda^2 \chi \leq \Omega$, where equality holds only for $\Omega = 0$. If we write

$$\Omega_\perp = \psi(\Omega) \quad (16.21)$$

then $\psi(0) = 0$, $\psi(2\pi) = \pi$ and $\psi(\Omega) \simeq \Omega$ for small Ω . These relations are in fact true for *any* aperture shape.




Figure 16.4: A *fan-out situation*. The signal is transmitted from the left hand node and is destined to reach the right hand nodes. The point at which power is split is chosen so as to minimize the total connection length.

16.9 Extension to fan-out and fan-in (chapter 6)

The extension of the results of chapter 6 to fan-in and fan-out is straightforward. Consider for instance the fan-out situation symbolically depicted in figure 16.4. Given the locations of the source and detectors, one arranges power division to occur at such points so that the total distance required, $\ell_1 + \ell_2 + \ell_3$, is minimum. Then the contribution of this fan-out link to the total volume required is just $\lambda^2(\ell_1 + \ell_2 + \ell_3)/2\pi$.

16.10 System size for the reflective multi-facet architecture

We refer to part a. of figure 7.1. Assume we want to form interconnections between pairs of transducers which, in the worst case, may be separated from each other by a distance comparable to the total extent $\mathcal{L}$ of the devices, which are laid out on a planar surface. We also assume that the total area of the holographic optical element is about equal to the total area occupied by these devices. Since $\mathcal{N} = kN$ connections

are being formed between $\mathcal{N}$ pairs of transducers, we write

$$2\mathcal{N}d_{tr}^2 \simeq \mathcal{N}d_f^2 \quad (16.22)$$

where d_{tr} and d_f denote the width of a transducer and a reflective facet respectively. Letting $\mathcal{H}$ denote the height of the system, we also have

$$d_{tr}^2 \simeq \frac{\mathcal{H}^2}{d_f^2} \lambda^2 \quad (16.23)$$

since $\mathcal{H}/d_f$ corresponds to the $f^\#$. There also exists a limit to how large $\mathcal{L}/\mathcal{H}$ can be. Referring the reader to [93] [91] for a detailed analysis of this issue, we will satisfy ourselves with the approximation

$$\mathcal{L} \leq 2\mathcal{H}. \quad (16.24)$$

Since $\mathcal{L}^2 \simeq 2\mathcal{N}d_{tr}^2$, combining the above we find that the minimum values of d_{tr}^2 and $\mathcal{L}$ satisfy

$$\frac{d_{tr}^2}{\lambda^2} \sim \mathcal{N} \quad (16.25)$$

$$\mathcal{L} \sim \mathcal{N}\lambda \quad (16.26)$$

within numerical factors of the order of unity. Similar relations can be shown to hold for variations of this architecture [19]. The reader is referred to [41] for detailed analysis of such configurations.

16.11 3 dimensional optical layouts with elements confined to a plane

Here we consider layouts where the elements are constrained to lie on a planar $N^{1/2} \times N^{1/2}$ grid, as in the fully 2 dimensional case, but the communication paths 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 $(2\lambda)^2$ 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,

$$d^2 \geq \frac{k\chi(2\lambda)}{M}(\bar{r}d + 2M(2\lambda)) \quad (16.27)$$

$$d \geq \frac{k\chi(2\lambda)\bar{r}}{M} + (2k\chi)^{\frac{1}{2}}(2\lambda). \quad (16.28)$$

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

$$d = \max \left[\left(\frac{kEB}{Q} \right)^{\frac{1}{2}}, \frac{k\chi\bar{r}(2\lambda)}{M}, d_d \right]. \quad (16.29)$$

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

$$\tau = \frac{1}{c} \left[N^{\frac{1}{2}} d + 2M(2\lambda) \right], \quad (16.30)$$

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 $p > 1/2$)

$$\frac{1}{S} = \tau = \frac{1}{c} \max \left[N^{\frac{1}{2}} \left(\frac{kEB}{Q} \right)^{\frac{1}{2}}, (k\chi\kappa)^{\frac{1}{2}} N^{\frac{p}{2}} (2\lambda), N^{\frac{1}{2}} d_d \right] \quad (16.31)$$

where $\chi = \max(1, BT_r)$ and $d_d^2 \geq k\chi d_{tr}^2$ where d_{tr} denotes the extent of a transducer. The reader will notice that apart from the last term, this equation is identical to equation 10.12.

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

$$\frac{1}{S} = \frac{1}{c} (kNB)^{\frac{1}{2}} \max \left[\left(\frac{E}{Q} \right)^{\frac{1}{2}}, T_r^{\frac{1}{2}} \lambda, T_r^{\frac{1}{2}} d_{tr} \right] \quad (16.32)$$

where we ignored all numerical factors and assumed that $d_d^2 = k\chi d_{tr}^2$, i.e. the element size is transducer limited. Of course, the second term is redundant since $d_{tr} \geq \lambda$. 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 . Another conclusion is that, when $p = 1$, only a constant increase in delay 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.

16.12 Average signal delay for normally conducting layouts

In the main text we concentrated on the worst case signal delay τ . It is an easy exercise to show that when all linewidths are kept equal, the wireability imposed average signal delay τ_{ave} is given by

$$\tau_{ave} = (16\rho\epsilon) \left\langle \frac{\ell^2}{W^2} \right\rangle = (16\rho\epsilon) \langle r^2 \rangle \frac{d^2}{W^2} = (16\rho\epsilon) \zeta_2(k\kappa)^{\frac{2}{\epsilon-1}} N^{\frac{2n-\epsilon-1}{n(\epsilon-1)}}. \quad (16.33)$$

What if all linewidths are not kept equal? For simplicity, we assume that T_d and W_{min} are very small, that $M = 1$ and that the system can be downscaled sufficiently so that propagation effects need not be considered. Then, assuming a linewidth distribution $W(r)$, the average delay is given by

$$\tau_{ave} = \frac{1}{k} \int \tau(r) g(r) dr \quad (16.34)$$

with $\tau(r) = (16\rho\epsilon)(rd/W(r))^2$ where d satisfies

$$d^\epsilon = \int [W(r)^{\epsilon-1} rd] g(r) dr \quad (16.35)$$

since $W(r)^{\epsilon-1} rd$ is the volume (or area) occupied by each line. Upon substitution we find that the average delay is given by

$$\tau_{ave} = (16\rho\epsilon) \frac{1}{k} \left[\int \frac{r^2}{W(r)^2} g(r) dr \right] \left[\int r W(r)^{\epsilon-1} g(r) dr \right]^{\frac{2}{\epsilon-1}}. \quad (16.36)$$

Using the standard techniques of the calculus of variations [68], we find that the optimal distribution of linewidths is

$$W(r) \propto r^{\frac{1}{\epsilon+1}} \quad (16.37)$$

where the constant of proportionality is arbitrary (scale invariance). It is also easy to show that the resulting expression for τ_{ave} has a growth rate identical to equation 16.33. Only a constant factor improvement is possible.

In practice, it would probably not be feasible to continuously contour the linewidth as a function of line length, so that a staircase approximation to equation 16.37 might

be used instead. Real computing systems are built out of a hierarchy of different technologies, each higher level employing wider lines. Thus, apart from its theoretical interest, the result presented may be useful as a guideline in contemplating the interconnection hierarchy when minimization of average signal delay is an objective.

Our scaling rules cannot be applied to the lines of present day VLSI chips because h and t are fixed for all lines in the same layer. This is not optimal. However, many-layer technologies are becoming available. Thus the upper layers which are reserved for the longer and wider wires may be constructed with larger values of h and t in accordance with our scaling rules. Then, a staircase approximation to equation 16.37 is possible. Thus, the result presented may also be used as a guideline in the construction of future many-layer ULSI, WSI or hybrid WSI circuits.

Bibliography

- [1] C.W. Allen. *Astrophysical Quantities*. The Athlone Press, University of London, London, third edition, 1976.
- [2] Y. Amitai and J.W. Goodman. Recording an efficient holographic optical element from computer generated holograms. *Optics Communications*, 80:107–109, 1990.
- [3] Y. Amitai and J.W. Goodman. Design of substrate-mode holographic interconnects with different recording and readout wavelength. *Applied Optics*, 30:2376–2381, 1991.
- [4] Halil B. Bakoglu and James D. Meindl. Optimal interconnection circuits for VLSI. *IEEE Transactions on Electron Devices*, 32:903–909, 1985.
- [5] Halil Burhan Bakoglu. *Circuit and System Performance Limits on ULSI: Interconnections and Packaging*. PhD thesis, Stanford University, Stanford, California, 1986.
- [6] Halil Burhan Bakoglu. *Circuits, Interconnections and Packaging for VLSI*. Addison-Wesley, Reading, Massachusetts, 1990.
- [7] Richard Barakat and John Reif. Lower bound on the computational efficiency of optical computing systems. *Applied Optics*, 26:1015–1018, 1987.
- [8] Jacob D. Bekenstein. Energy cost of information transfer. *Physical Review Letters*, 46:623–626, 1981.
- [9] C.H. Bennett and R. Landauer. The fundamental physical limits of computation. *Scientific American*, pages 48–56, July 1985.
- [10] C. Berge. *The Theory of Graphs*. Wiley, New York, 1962.
- [11] C. Berge. *Graphs and Hypergraphs*. North Holland, Amsterdam, 1973.
- [12] Sandeep N. Bhatt and Frank Thomson Leighton. A framework for solving VLSI layout problems. *Journal of Computer and System Sciences*, 28:300–343, 1984.
- [13] Albert J. Blodgett, Jr. Microelectronic packaging. *Scientific American*, 249:86–96, July 1983.
- [14] Bela Bollobas. *Graph Theory: An Introductory Course*. Springer-Verlag, Berlin, 1979.

- [15] Max Born and Emil Wolf. *Principles of Optics*. Pergamon Press, Oxford, sixth edition, 1980.
- [16] G. D. Boyd, D. A. B. Miller, D. S. Chemla, S. L. McCall, A. C. Gossard, and J. H. English. Multiple quantum well reflection modulator. *Applied Physics Letters*, 50:1119, 1987.
- [17] John Philip Boyd. *Chebyshev and Fourier Spectral Methods*. Springer-Verlag, Berlin, 1989.
- [18] Ronald N. Bracewell. *The Fourier Transform and Its Applications*. McGraw-Hill, New York, second edition, 1986.
- [19] Eric Bradley, Paul Kit Lai Yu, and Alan R. Johnston. System issues relating to laser diode requirements for VLSI holographic optical interconnects. *Optical Engineering*, 28:201–211, 1989.
- [20] Hans J. Bremermann. Minimum energy requirements of information transfer and computing. *International Journal of Theoretical Physics*, 21:203–217, 1982.
- [21] Leon Brillouin. *Science and Information Theory*. Academic Press, New York, 1972.
- [22] D. Casasent. Optical signal processing. In D. Casasent, editor, *Optical Data Processing*, chapter 8. Springer-Verlag, Berlin, 1978.
- [23] W. Thomas Cathey, Kelvin Wagner, and William J. Miceli. Digital computing with optics. *Proceedings of the IEEE*, 77:1558–1572, 1989.
- [24] H. John Caulfield and Joseph Shamir. Wave particle duality considerations in optical computing. *Applied Optics*, 28:2184–2186, 1989.
- [25] David K. Cheng. *Field and Wave Electromagnetics*. Addison-Wesley, Reading, Massachusetts, 1983.
- [26] P. Christie and S.B. Styer. Fractal description of computer interconnection distributions. In Stuart K. Tewksbury, editor, *Microelectronic Interconnects and Packaging: System and Process Integration*, volume 1390. SPIE, 1990.
- [27] Phillip Christie, Jeffrey E. Cotter, and Allen M. Barrett. Design and simulation of optically interconnected computer systems. In Alfred P. DeFonzo, editor, *Interconnection of High Speed and High Frequency Devices and Systems*, volume 947, pages 19–24. SPIE, 1989.
- [28] Bradley D. Clymer and Joseph W. Goodman. Optical clock distribution to silicon chips. *Optical Engineering*, 25:1103–108, 1986.
- [29] Robert E. Collin. *Foundations for Microwave Engineering*. Mc-Graw Hill, New York, 1966.
- [30] William J. Dally. *A VLSI Architecture for concurrent data structures*. Kluwer Academic Publishers, Norwell, Massachusetts, 1987.

- [31] P. L. Derry and Amnon Yariv. Ultralow-threshold graded-index separate-confinement single quantum well buried heterostructure (Al,Ga)As lasers with high reflectivity coatings. *Applied Physics Letters*, 50:1773, 1987.
- [32] Charles A. Desoer and Ernst S. Kuh. *Basic Circuit Theory*. Mc-Graw Hill, New York, 1969.
- [33] David Deutsch. Is there a fundamental bound on the rate at which information can be processed? *Physical Review Letters*, 48:286–288, 1982.
- [34] W. E. Donath. Placement and average interconnection lengths of computer logic. *IEEE Transactions on Circuits and Systems*, 26:272–277, 1979.
- [35] W.E. Donath. Stochastic model of the computer logic design process. Technical Report RC 3136, IBM Thomas J. Watson Research Center, Yorktown Heights, New York, 1970.
- [36] W.E. Donath. Equivalence of memory to ‘random logic’. *IBM Journal of Research and Development*, 18:401–407, 1974.
- [37] W.E. Donath. Wire length distribution for placements of computer logic. *IBM Journal of Research and Development*, 25:152–155, 1981.
- [38] Scott E. Fahlman. *NETL, a System for Representing and Using Real-world Knowledge*. MIT press, Cambridge, Massachusetts, 1979.
- [39] M. Farhadiroushan, D.R. Selviah, and J.E. Midwinter. Asymmetric fabry-perot multiple quantum well PIN diodes and S-SEEDs for intra-chip optical interconnections. In *Technical Digest of the 1991 OSA Topical Meeting on Photonic Switching*, 1991.
- [40] Michael R. Feldman, Sadik C. Esener, Clark C. Guest, and Sing H. Lee. Comparison between optical and electrical interconnects based on power and speed considerations. *Applied Optics*, 27:1742–1751, 1988.
- [41] Michael R. Feldman and Clark C. Guest. Interconnect density capabilities of computer generated holograms for optical interconnection of very large scale integrated circuits. *Applied Optics*, 28:3134–3137, 1989.
- [42] Michael R. Feldman and Clark C. Guest. Nested crossbar connection networks for optically interconnected processor arrays for vector-matrix multiplication. *Applied Optics*, 29:1068–1076, 1990.
- [43] Michael R. Feldman, Clark C. Guest, Timothy J. Drabik, and Sadik C. Esener. Comparison between electrical and free space optical interconnects for fine grain processor arrays based on interconnect density capabilities. *Applied Optics*, 28:3820–3829, 1989.
- [44] David K. Ferry. Interconnection lengths and VLSI. *IEEE Circuits and Devices Magazine*, pages 39–42, July 1985.

- [45] Michael Feuer. Connectivity of random logic. *IEEE Transactions on Computers*, 31:29–33, 1982.
- [46] Richard P. Feynman, Robert B. Leighton, and Matthew Sands. *The Feynman Lectures on Physics, Volume 3*. Addison-Wesley, Reading, Massachusetts, 1965.
- [47] Richard P. Feynman, Robert B. Leighton, and Matthew Sands. *The Feynman Lectures on Physics, Volume 1*. Addison-Wesley, Reading, Massachusetts, 1965.
- [48] Gary Frazier. An ideology for nanoelectronics. In Stuart K. Tewksbury, editor, *Concurrent Computations*, chapter 1. Plenum Press, New York, 1988.
- [49] Edward Fredkin and Tommaso Toffoli. Conservative logic. *International Journal of Theoretical Physics*, 21:219–253, 1982.
- [50] R.C. Frye. Analysis of the trade-offs between conventional and superconducting interconnections. *IEEE Circuits and Devices Magazine*, pages 27–32, May 1989.
- [51] Dennis Gabor. Light and information. In Emil Wolf, editor, *Progress in Optics, Volume 1*, chapter 4. North-Holland Publishing Company, Amsterdam, 1961.
- [52] Abbas El Gamal. Two-dimensional stochastic model for interconnections in master slice integrated circuits. *IEEE Transactions on Circuits and Systems*, 28:127–134, 1981.
- [53] Abbas El Gamal, Jack L. Kouloheris, Dana How, and Martin Morf. BiNMOS: A basic cell for BiCMOS sea-of-gates. In *Proceedings of the IEEE Custom Integrated Circuits Conference*, pages 8.3.1–8.3.4, 1989.
- [54] Donald S. Gardner, James D. Meindl, and Krishna C. Saraswat. Interconnection and electromigration scaling theory. *IEEE Transactions on Electron Devices*, 34:633–643, 1987.
- [55] Bernard Gold and Charles M. Rader. *Digital Processing of Signals*. McGraw-Hill, New York, 1969.
- [56] Joseph W. Goodman. *Introduction to Fourier Optics*. McGraw-Hill, New York, 1968.
- [57] Joseph W. Goodman. Fan-in and fan-out with optical interconnections. *Optica Acta*, 32:1489, 1985.
- [58] Joseph W. Goodman. Optics as an interconnect technology. In Henri H. Arsenault, Tomasz Szoplik, and Bohdan Macukow, editors, *Optical Processing and Computing*, chapter 1. Academic Press, San Diego, 1989.
- [59] Joseph W. Goodman, Frederick J. Leonberger, Sun-Yuan Kung, and Ravindra Athale. Optical interconnections for VLSI systems. *Proceedings of the IEEE*, 72:850–866, 1984.

- [60] K.C. Gupta, R. Garg, and I.J. Bahl. Microstrip lines and slotlines. Technical report, Artech, Dedham, Massachusetts, 1979.
- [61] Craig Harkins. Mathematical considerations in seeking fundamental limits to professional communication processes and effects. *Proceedings of the IEEE*, 69:167–170, 1981.
- [62] Alfred C. Hartmann. Computational metrics and fundamental limits for parallel architectures. In S.K. Tewksbury, editor, *Frontiers of Computing Systems Research, Volume 1*. Plenum Press, New York, 1990.
- [63] Alfred C. Hartmann and Jeffrey D. Ullman. Model categories for theories of parallel systems. In G.J. Lipovski and M. Malek, editors, *Parallel Computing: Theory and Experience*. John Wiley and Sons, 1986.
- [64] Hideki Hasegawa, Mieko Furukawa, and Hisayoshi Yanai. Properties of microstrip line on silicon-silicon dioxide system. *IEEE Transactions on Microwave Theory and Techniques*, 19:869–881, 1971.
- [65] P.R. Haugen, S. Rychnovsky, A. Husain, and L.D. Hutcheson. Optical interconnects for high speed computing. *Optical Engineering*, 25:1076, 1986.
- [66] 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*, 2:117–144, 1978.
- [67] Paul S. Henry. Lightwave primer. *IEEE Journal of Quantum Electronics*, 21:1862–1879, 1985.
- [68] Francis B. Hildebrand. *Methods of Applied Mathematics*. Prentice-Hall Inc., Englewood Cliffs, New Jersey, 1965.
- [69] W. Daniel Hillis. New computer architectures and their relationship to physics or why computer science is no good. *International Journal of Theoretical Physics*, 21:255–262, 1982.
- [70] W. Daniel Hillis. *The Connection Machine*. The MIT press, Cambridge, Massachusetts, 1985.
- [71] Jack P. Holman. *Heat Transfer*. Mc-Graw Hill, New York, fifth edition, 1981.
- [72] John Hopcroft, Wolfgang Paul, and Leslie Valiant. On time versus space. *Journal of the Association for Computing Machinery*, 24:332–337, 1977.
- [73] T.C. Hu and Ernest S. Kuh. *VLSI Circuit Layout: Theory and Design*. IEEE Press, New York, 1985.
- [74] Alan Huang. Optical digital computers: Devices and architecture. In Stuart K. Tewksbury, editor, *Concurrent Computations*, chapter 2. Plenum Press, New York, 1988.
- [75] Alan Huang. Optical digital computing at Bell labs. In *Conference Record of the 1990 International Topical Meeting on Optical Computing*, pages 259–263, 1990.

- [76] Robert G. Hunsperger. *Integrated Optics: Theory and Technology*. Springer-Verlag, Berlin, second edition, 1984.
- [77] Lynn D. Hutcheson and Paul Haugen. Optical interconnects replace hardwire. *IEEE Spectrum*, pages 30–35, March 1987.
- [78] J. Jahns and S.J. Walker. Imaging with planar optical systems. *Optics Communications*, 76:313–317, 1990.
- [79] Robert W. Keyes. Power dissipation in information processing. *Science*, 168:796–801, 1970.
- [80] Robert W. Keyes. Physical limits in digital electronics. *Proceedings of the IEEE*, 63:740–767, 1975.
- [81] Robert W. Keyes. A figure of merit for IC packaging. *IEEE Journal of Solid State Circuits*, 13:265–266, 1978.
- [82] Robert W. Keyes. The evolution of digital electronics towards VLSI. *IEEE Transactions on Electron Devices*, 26:271–279, 1979.
- [83] Robert W. Keyes. Fundamental limits in digital information processing. *Proceedings of the IEEE*, 69:267–278, 1981.
- [84] Robert W. Keyes. Communication in computation. *International Journal of Theoretical Physics*, 21:263–273, 1982.
- [85] Robert W. Keyes. The wire-limited logic chip. *IEEE Journal of Solid State Circuits*, 17:1232–1233, 1982.
- [86] Robert W. Keyes. *The Physics of VLSI Systems*. Addison-Wesley, Reading, Massachusetts, 1987.
- [87] F. Kiamilev, Sadik C. Esener, R. Paturi, Y. Fainman, P. Mercer, C.C. Guest, and Sing H. Lee. Programmable optoelectronic multiprocessors and their comparison with symbolic substitution for digital optical computing. *Optical Engineering*, 28:396–409, 1989.
- [88] F. Kiamilev, A. Krishnamoorthy, K.S. Urquhart, P. Marchand, S. Esener, and S.H. Lee. Grain-size considerations for programmable optoelectronic multiprocessor interconnection networks. In *Optical Society of America 1990 Annual Meeting Technical Digest*, 1990.
- [89] Donald E. Knuth. *The TeXbook*. Addison-Wesley, Reading, Massachusetts, 1984.
- [90] H. Kogelnik. Theory of optical waveguides. In Theodor Tamir, editor, *Guided-Wave Optoelectronics*, chapter 2. Springer-Verlag, Berlin, 1988.
- [91] Raymond K. Kostuk. *Multiple Grating Volume Reflection Holograms with Application to Optical Interconnects*. PhD thesis, Stanford University, Stanford, California, 1986.

- [92] Raymond K. Kostuk, Joseph W. Goodman, and Lambertus Hesselink. Optical imaging applied to microelectronic chip-to-chip interconnections. *Applied Optics*, 24:2851–2858, 1985.
- [93] Raymond K. Kostuk, Joseph W. Goodman, and Lambertus Hesselink. Design considerations for holographic optical interconnects. *Applied Optics*, 26:3947–3953, 1987.
- [94] R.K. Kostuk, J.W. Goodman, and L. Hesselink. Optical interconnects. In H.M. Gibbs, G. Khitrova, and N. Peyghambarian, editors, *Nonlinear Photonics*, chapter 3. Springer-Verlag, Berlin, 1990.
- [95] Harry Kroger, Claude Hilbert, Uttam Ghoshal, David Gibson, and Larry Smith. Applications of superconductivity to packaging. *IEEE Circuits and Devices Magazine*, pages 16–21, May 1989.
- [96] H.T. Kung. Why systolic architectures? *Computer*, pages 37–46, January 1982.
- [97] H.T. Kung. Memory requirements for balanced computer architectures. In *Proceedings of the IEEE 19th Annual International Symposium on Computer Architecture*, pages 49–54, 1986.
- [98] S.Y. Kung. *VLSI Array Processors*. Prentice-Hall, Englewood Cliffs, New Jersey, 1988.
- [99] Oh-Kyong Kwon. *Chip-to-chip Interconnections for Very High-speed System-level Integration*. PhD thesis, Stanford University, Stanford, California, 1988.
- [100] 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*, 8:582–585, 1987.
- [101] Leslie Lamport. *L^AT_EX: A Document Preparation System*. Addison-Wesley, Reading, Massachusetts, 1986.
- [102] Rolf Landauer. Uncertainty principle and minimal energy dissipation in the computer. *International Journal of Theoretical Physics*, 21:283–297, 1982.
- [103] Bernard S. Landman and Roy L. Russo. On a pin versus block relationship for partitions of logic graphs. *IEEE Transactions on Computers*, 20:1469–1479, 1971.
- [104] W.H. Lee. Computer generated holograms: Techniques and applications. In Emil Wolf, editor, *Progress in Optics, Volume 16*, pages 119–232. North-Holland Publishing Company, Amsterdam, 1978.
- [105] Frank Thomson Leighton and Arnold L. Rosenberg. Three-dimensional circuit layouts. *Journal of Computer and System Sciences*, 15:793–813, 1986.
- [106] Charles E. Leiserson. *Area-Efficient VLSI Computation*. The MIT Press, Cambridge, Massachusetts, 1983.

- [107] Charles E. Leiserson. Fat-trees: Universal networks for hardware-efficient supercomputing. *IEEE Transactions on Computers*, 34:892–901, 1985.
- [108] A.L. Lentine, L.M.F. Chirovsky, M.W. Focht, J.M. Freund, G.D. Guth, R.E. Leibenguth, G.J. Przybylek, L.E. Smith, L.A. D'Asaro, and D.A.B. Miller. Integrated self electro-optic effect device photonic switching nodes. In *Technical Digest of the 1991 OSA Topical Meeting on Photonic Switching*, 1991.
- [109] Tingye Li. Lightwave telecommunication. *Physics Today*, pages 24–31, May 1985.
- [110] Benoit B. Mandelbrot. The Pareto-Levy law and the distribution of income. *International Economic Review*, 1:79–106, 1960.
- [111] Benoit B. Mandelbrot. Information theory and psycholinguistics: A theory of word frequencies. In P.F. Lazarsfeld and N.W. Henry, editors, *Readings in Mathematical Social Science*. MIT press, Cambridge, Massachusetts, 1968.
- [112] Benoit B. Mandelbrot. *Fractals: Form, Chance and Dimension*. W.H. Freeman, San Francisco, 1977.
- [113] Benoit B. Mandelbrot. *The Fractal Geometry of Nature*. W.H. Freeman, New York, 1983.
- [114] Akira Masaki. Electrical resistance as a limiting factor for high performance computer packaging. *IEEE Circuits and Devices Magazine*, pages 22–26, May 1989.
- [115] Akira Masaki and Minoru Yamada. Equations for estimating wire length in various types of 2-D and 3-D system packaging structures. *IEEE Transactions on Components, Hybrids, and Manufacturing Technology*, 10:190–198, 1987.
- [116] Richard E. Matick. *Transmission Lines for Digital and Communication Networks*. Mc-Graw Hill, New York, 1969.
- [117] Jack F. McDonald, Hans J. Greub, Randy H. Steinworth, Brian J. Donlan, and Albert S. Bergendahl. Wafer scale interconnection for GaAs packaging—applications to RISC architecture. *Computer*, pages 21–35, April 1987.
- [118] J.F. McDonald, A.J. Steckl, C.A. Neugebauer, R.O. Carlson, and A.S. Bergendahl. Multilevel interconnections for wafer scale integration. *Journal of Vacuum Science and Technology A*, 4:3127–3138, 1986.
- [119] A. Mead and L. Conway. *Introduction to VLSI Systems*. Addison-Wesley, Reading, Massachusetts, 1980.
- [120] Timothy Lyman Michalka. *Models for Wafer Scale Integration Implementation*. PhD thesis, Stanford University, Stanford, California, 1988.

- [121] D.A.B. Miller. Optics for low-energy communication inside digital processors: Quantum detectors, sources and modulators as efficient impedance converters. *Optics Letters*, 14:146–148, 1989.
- [122] D.A.B. Miller. Quantum-well self-electro-optic effect devices. *Optical and Quantum Electronics*, 22:61–98, 1990.
- [123] D.A.B. Miller. Quantum wells for optical logic and interconnection. In Gnanalingam Arjavalingam and James Pazaris, editors, *Microelectronic Interconnects and Packaging*, volume 1389. SPIE, 1990.
- [124] D.A.B. Miller, S.E. Henry, A.C. Gossard, and J.H. English. Integrated quantum well self-electrooptic effect devices: 2x2 array of optically bistable devices. *Applied Physics Letters*, 49:821–823, 1986.
- [125] Cleve Moler, John Little, and Steve Bangert. *PRO-MATLAB for Sun Workstations*. The MathWorks Inc., 20 North Main St., Suite 250, Sherborn, MA 01770, 1987.
- [126] Karl Dieter Möller. *Optics*. University Science Books, Mill Valley, California, 1988.
- [127] Wataru Nakayama. On the accomodation of coolant flow paths in high density packaging. *IEEE Transactions on Components, Hybrids and Manufacturing Technology*, 13:1040–1049, 1990.
- [128] John A. Neff. Major initiatives for optical computing. *Optical Engineering*, 26:2–9, 1987.
- [129] Alan V. Oppenheim and Ronald W. Shafer. *Digital Signal Processing*. Prentice-Hall International, London, 1975.
- [130] A. Orlitsky and A. El Gamal. Communication complexity. In Yasser S. Abu-Mostafa, editor, *Complexity in Information Theory*. Springer-Verlag, New York, 1988.
- [131] Haldun M. Ozaktas. *A Physical Approach to Communication Limits in Computation*. PhD thesis, Stanford University, Stanford, California, 1991.
- [132] Haldun M. Ozaktas, Yaakov Amitai, and Joseph W. Goodman. Comparison of system size for some optical interconnection architectures and the folded multi-facet architecture. *Optics Communications*, 1991. Accepted for publication.
- [133] Haldun M. Ozaktas, Yaakov Amitai, and Joseph W. Goodman. A three dimensional optical interconnection architecture with minimal growth rate of system size. *Optics Communications*, 1991. Accepted for publication.
- [134] Haldun M. Ozaktas and Joseph W. Goodman. Lower bound for the communication volume required for an optically interconnected array of points. *Journal of the Optical Society of America A*, 7:2100–2106, 1990.

- [135] Haldun M. Ozaktas and Joseph W. Goodman. Optimal partitioning of very large scale optoelectronic computing systems. In *Optical Society of America 1990 Annual Meeting Technical Digest*, 1990.
- [136] Haldun M. Ozaktas and Joseph W. Goodman. The limitations of interconnections in providing communication between an array of points. In S.K. Tewksbury, editor, *Frontiers of Computing Systems Research, Volume 2*. Plenum Press, New York, 1991. To appear.
- [137] Haldun M. Ozaktas and Joseph W. Goodman. Multiplexed hybrid interconnection architectures. In *Technical Digest of the 1991 OSA Topical Meeting on Optical Computing*, 1991.
- [138] Athanasios Papoulis. *Probability, Random Variables and Stochastic Processes*. McGraw-Hill, New York, second edition, 1965.
- [139] R.F.W. Pease. Advances in packaging for VLSI systems. In *Proceedings of the International Electron Devices Meeting*, pages 480–483, 1986.
- [140] William E. Pence and J. Peter Krusius. The fundamental limits for electronic packaging and systems. *IEEE Transactions on Components, Hybrids and Manufacturing Technology*, 10:176–183, 1987.
- [141] J.B. Pendry. Quantum limits to the flow of information and entropy. *Journal of Physics: A*, 16:2161–2171, 1983.
- [142] L. Pietronero. Fractals in physics: Introductory concepts. In Stig Lundqvist, Norman H. March, and Mario P. Tosi, editors, *Order and Chaos in Nonlinear Physical Systems*. Plenum Press, New York, 1988.
- [143] Simon Ramo, John R. Whinnery, and Theodore Van Duzer. *Fields and Waves in Communication Electronics*. John Wiley and Sons Inc., New York, second edition, 1984.
- [144] Roy L. Russo. On the tradeoff between logic performance and circuit-to-pin ratio for LSI. *IEEE Transactions on Computers*, 21:147–153, 1972.
- [145] Krishna C. Saraswat and Farrokh Mohammadi. Effect of scaling of interconnections on the time delay of VLSI circuits. *IEEE Transactions on Electron Devices*, 29:645–650, 1982.
- [146] F. Sauer. Fabrication of diffractive-reflective optical interconnects for infrared operation based on total internal reflection. *Applied Optics*, 28:386–388, 1989.
- [147] Timothy Alan Schreyer. *The Effects of Interconnection Parasitics on VLSI Performance*. PhD thesis, Stanford University, Stanford, California, 1989.
- [148] Erwin Schroedinger. *What is Life?* Cambridge University Press, Cambridge, 1967.
- [149] Joseph Shamir. Fundamental speed limitations on parallel processing. *Applied Optics*, 26:1567, 1987.

- [150] Anthony E. Siegman. *Lasers*. University Science Books, Mill Valley, California, 1986.
- [151] P.W. Smith. On the physical limits of digital optical switching and logic elements. *The Bell System Technical Journal*, 61:1975–1993, 1982.
- [152] Paul M. Solomon. A comparison of semiconductor devices for high-speed logic. *Proceedings of the IEEE*, 70:489–509, 1982.
- [153] Charles W. Stirk and Demetri Psaltis. Comparison of optical and electronic 3-dimensional circuits. In Gnanalingam Arjavalingam and James Pazaris, editors, *Microelectronic Interconnects and Packaging*, volume 1389. SPIE, 1990.
- [154] Ivan E. Sutherland and Donald Oestreicher. How big should a printed circuit board be? *IEEE Transactions on Computers*, 22:537–542, 1973.
- [155] Earl Swartzlander, editor. *Proceedings, International Conference on Wafer Scale Integration*, Washington D.C., 1989. IEEE Computer Society.
- [156] H.N.V. Temperley. *Graph Theory and Applications*. Ellis Horwood Limited, Chichester, 1981.
- [157] S.K. Tewksbury and L.A. Hornak. Wafer level system integration: A review. *IEEE Circuits and Devices Magazine*, pages 22–30, September 1989.
- [158] C.D. Thompson. Area-time complexity for VLSI. In *Proceedings of the 11th Annual ACM Symposium on the Theory of Computing*, pages 81–88. Association for Computing Machinery, 1979.
- [159] Richard F. Thompson. *The Brain*. W.H. Freeman and Company, New York, 1985.
- [160] Tommaso Toffoli. Physics and computation. *International Journal of Theoretical Physics*, 21:165–175, 1982.
- [161] D.B. Tuckerman. *Heat-transfer microstructures for integrated circuits*. PhD thesis, Stanford University, Stanford, California, 1984.
- [162] D.B. Tuckerman and R.F.W. Pease. High performance heat sinking for VLSI. *IEEE Electron Device Letters*, 2:126–129, 1981.
- [163] Rao R. Tummala and Eugene J. Rymaszewski, editors. *Microelectronics Packaging Handbook*. Van Nostrand Reinhold, New York, New York, 1989.
- [164] Jeffrey D. Ullman. *Computational Aspects of VLSI*. Computer Science Press, Rockville, Maryland, 1984.
- [165] E.A. Volkov. *Numerical Methods*. Mir Publishers, Moscow, 1986.
- [166] J.N. Walker. *Chess Openings for Juniors*. Oxford University Press, London, 1975.

- [167] J.T. Watt and J.D. Plummer. Effect of interconnection delay on liquid nitrogen temperature CMOS circuit performance. In *Proceedings of the International Electron Devices Meeting*, pages 393–396, 1987.
- [168] Neil H.E. Weste and Kamran Eshraghian. *Principles of CMOS VLSI design: a systems perspective*. Addison-Wesley, Reading, Massachusetts, 1985.
- [169] Paul Wiley. A parallel architecture comes of age at last. *IEEE Spectrum*, pages 46–50, June 1987.
- [170] Robin J. Wilson and Lowell W. Beineke, editors. *Applications of Graph Theory*. Academic Press, London, 1979.
- [171] John T. Winthrop. Propagation of structural information in optical wave fields. *Journal of the Optical Society of America*, 61:15–30, 1971.
- [172] William L. Wolfe and George J. Zissis, editors. *The Infrared Handbook*. Office of Naval Research, Department of the Navy, Washington D.C., 1978.
- [173] Hans Wolter. On basic analogies and principle differences between optical and electronic information. In Emil Wolf, editor, *Progress in Optics, Volume 1*, chapter 5. North-Holland Publishing Company, Amsterdam, 1961.
- [174] Bruce A. Wooley, Mark A. Horowitz, R. Fabian Pease, and Tsen-Shau Yang. Active substrate system integration. In *IEEE International Conference on Computer Design: VLSI in Computers and Processors*, pages 468–471, 1987.
- [175] Wennie H. Wu, Larry A. Bergman, Alan R. Johnston, Clark C. Guest, Sadik C. Esener, P.K.L. Yu, Michael R. Feldman, and Sing H. Lee. Implementation of optical interconnections for VLSI. *IEEE Transactions on Electron Devices*, 34:706–714, 1987.
- [176] Andrew Chi-Chih Yao. Some complexity questions related to distributive computing. In *Proceedings of the 11th Annual ACM Symposium on the Theory of Computing*, pages 209–213. Association for Computing Machinery, 1979.
- [177] Amnon Yariv. *Introduction to Optical Electronics*. Holt, Rinehart and Winston, New York, second edition, 1976.
- [178] Francis T. S. Yu. *Optics and Information Theory*. John Wiley and Sons, New York, 1976.
- [179] H.H. Zappe. Josephson quantum interference computer devices. *IEEE Transactions on Magnetics*, 13:41–47, 1977.